
How Do Nonlinear Transformers Learn and Generalize in In-Context Learning?

Hongkang Li¹ Meng Wang¹ Songtao Lu² Xiaodong Cui² Pin-Yu Chen²

Abstract

Transformer-based large language models have displayed impressive in-context learning capabilities, where a pre-trained model can handle new tasks without fine-tuning by simply augmenting the query with some input-output examples from that task. Despite the empirical success, the mechanics of how to train a Transformer to achieve ICL and the corresponding ICL capacity is mostly elusive due to the technical challenges of analyzing the nonconvex training problems resulting from the nonlinear self-attention and nonlinear activation in Transformers. To the best of our knowledge, this paper provides the first theoretical analysis of the training dynamics of Transformers with nonlinear self-attention and nonlinear MLP, together with the ICL generalization capability of the resulting model. Focusing on a group of binary classification tasks, we train Transformers using data from a subset of these tasks and quantify the impact of various factors on the ICL generalization performance on the remaining unseen tasks with and without data distribution shifts. We also analyze how different components in the learned Transformers contribute to the ICL performance. Furthermore, we provide the first theoretical analysis of how model pruning affects ICL performance and prove that proper magnitude-based pruning can have a minimal impact on ICL while reducing inference costs. These theoretical findings are justified through numerical experiments.

¹Department of Electrical, Computer, and System Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA ²IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA. Correspondence to: Hongkang Li <lih35@rpi.edu>, Meng Wang <wangm7@rpi.edu>, Songtao Lu <songtao@ibm.com>, Xiaodong Cui <cui@us.ibm.com>, Pin-Yu Chen <pin-yu.chen@ibm.com>.

1. Introduction

Transformers now serve as the backbone architecture for a wide range of modern, large-scale foundation models, including prominent language models like GPT-3 (Brown et al., 2020), PaLM (Chowdhery et al., 2022), LLaMa (Touvron et al., 2023), as well as versatile visual and multimodal models such as CLIP (Radford et al., 2021), DALL-E (Ramesh et al., 2021), and GPT-4 (OpenAI, 2023). One intriguing capability exhibited by certain large language models (LLMs) is known as “**in-context learning**” (ICL) (Brown et al., 2020). Given a pre-trained model $F(\Psi)$, parameterized by weights Ψ , the conventional approach fine-tunes Ψ separately for each downstream task using data from that task. In contrast, ICL allows $F(\Psi)$ to handle multiple unseen tasks simultaneously without any fine-tuning. Garg et al. (2022) is the first paper to mathematically formulate ICL. Briefly speaking, to predict $f(x_{\text{query}})$ of a query input x_{query} for a new task represented by the label function f , ICL augments x_{query} by l example input-output pairs $(x_i, f(x_i))_{i=1}^l$. The resulting so-called *prompt* is sent to the model $F(\Psi)$, and, surprisingly, the model can output a prediction close to $f(x_{\text{query}})$. Thus, ICL is an efficient alternative to the resource-consuming fine-tuning methods. ICL has shown outstanding performance in multiple tasks in practice, including question answering (Liu et al., 2022b; Wu et al., 2023b), natural language inference (Liu et al., 2022a; Wu et al., 2023b), text generation (Brown et al., 2020; Lucy & Bamman, 2021), etc.

In parallel, model pruning (Han et al., 2015; Wen et al., 2016) can reduce the inference cost by removing some weights after training. It has been extensively evaluated in various applications. Among various pruning techniques, such as gradient methods (Molchanov et al., 2016) and reconstruction error minimization (Luo et al., 2017), magnitude-based pruning (Wen et al., 2016) is the most popular approach due to its simplicity and demonstrated promising empirical results. A few recent works (Frantar & Alistarh, 2023; Ma et al., 2023; Sun et al., 2023; Liu et al., 2023) also explore the pruning of LLMs to preserve their ICL capacity while accelerating the inference.

Despite the empirical success of ICL, one fundamental and theoretical question is less investigated, which is:

How can a Transformer be trained to perform ICL and

generalize in and out of domain successfully and efficiently?

Some recent works attempt to answer this question for linear regression tasks (Li et al., 2023c; Zhang et al., 2023a). Specifically, Li et al. (2023c) investigate the generalization gap and stability of ICL. Zhang et al. (2023a) explore the training and generalization of ICL with Transformers, especially with distribution shifts during inference. Wu et al. (2023a) studies the required number of pre-training tasks for a desirable ICL property. Huang et al. (2023) characterizes the training dynamics using Transformers with softmax attention and linear MLP. However, these results are either built upon simplified Transformer models by ignoring nonlinear self-attention (Zhang et al., 2023a; Wu et al., 2023a) or nonlinear activation in the multilayer perceptron (MLP) (Huang et al., 2023; Zhang et al., 2023a; Wu et al., 2023a) or cannot characterize how to train a model to achieve the desirable ICL capability with distribution-shifted data (Huang et al., 2023; Li et al., 2023c; Wu et al., 2023a).

1.1. Major Contributions of This Work

To the best of our knowledge, our work is the first theoretical analysis of the training dynamics of Transformers with nonlinear self-attention and nonlinear MLP, together with the ICL generalization capability of the resulting model. Moreover, our paper provides the first theoretical analysis of the impact of model pruning on ICL performance. Focusing on a group of binary classification tasks, we show that training a Transformer using prompts from a subset of these tasks can return a model with the ICL capability to generalize to the rest of these tasks. We provide a quantitative analysis of the required number of training data, iterations, the length of prompts, and the resulting ICL performance. Although our analysis is centered on a simplified single-head and one-layer Transformer with softmax self-attention and ReLU MLP, our theoretical insights shed light on practical architectures. Our major contributions include:

1. A theoretical characterization of how to train Transformers to enhance their ICL capability. We consider a data model where input data include both relevant patterns that determine the labels and irrelevant patterns that do not affect the labels. We quantify how the training and the resulting ICL generalization performance are affected by various factors, such as the magnitude of relevant features and the fraction of context examples that contain the same relevant pattern as the new query. In addition to proving the ICL capability of the learned Transformer to generalize to new binary tasks based on the relevant patterns that appear in the training data, we also prove the ICL capability to generalize to tasks based on patterns that are linear combinations of the relevant patterns and are unseen in the training data.

2. Expand the theoretical understanding of the mechanism of the ICL capability of Transformers. We prove

that when sending a prompt to a properly trained Transformer, the attention weights are concentrated on contexts that share the same relevant pattern as the query. Then, the ReLU MLP layer promotes the label embedding of these examples, thus making the correct prediction for the query. Similar insights have appeared in (Huang et al., 2023). We expand the analysis to Transformers with nonlinear MLP layers and new tasks with a data distribution shift.

3. Theoretical justification of magnitude-based pruning in preserving ICL. Based on the characterization of the trained Transformer, our paper also provides the first theoretical analysis of the ICL inference performance when the trained model is pruned by removing neurons in the MLP layer. We show that pruning a set of neurons with a small magnitude has little effect on the generalization while pruning the remaining neurons leads to a large generalization error growing with the pruning rate. To the best of our knowledge, no theoretical analysis exists on how model pruning affects ICL.

1.2. Related Work

Expressive power of ICL Some existing works study the expressive power of Transformers to implement algorithms via ICL. Akyürek et al. (2023); Von Oswald et al. (2023) demonstrate that Transformers conduct gradient descent during the forward pass of Transformers with prompts as inputs. Ahn et al. (2023); Cheng et al. (2023) extend the conclusion to preconditioned and functional gradient descent via ICL. Garg et al. (2022); Bai et al. (2023); Guo et al. (2023) show the existence of Transformers that can implement a broad class of machine learning algorithms in context.

The optimization and generalization of Transformers Beyond in-context learning, there are several other works about the optimization and generalization analysis of fine-tuning or prompt tuning on Transformers. Jelassi et al. (2022); Li et al. (2023a;b); Luo et al. (2024) study the generalization of one-layer Transformer by assuming spatial association or the majority voting of tokens. Li et al. (2023d) delve into how one-layer Transformers learn semantic structure. Oymak et al. (2023) depict the trajectory of prompt tuning of attention networks. Tarzanagh et al. (2023b;a) characterize that the gradient updates of the prompt or weights converge to a max-margin SVM solution. Tian et al. (2023a;b) probe the training dynamics of Transformers for the next token prediction problem given infinitely long sequences.

Theoretical generalization analysis of pruning A few recent works consider analyzing the generalizations performance of model pruning theoretically. For example, Zhang et al. (2021) study the sample complexity of training a pruned network with a given sparse ground truth weight. Yang & Wang (2023) investigate the neural tangent kernel of the pruned model. Zhang et al. (2023c); Yang et al. (2023)

Theoretical Works	Nonlinear Attention	Nonlinear MLP	Training Analysis	Distribution-Shifted Data	Tasks
Li et al. (2023c)	✓	✓			linear regression
Zhang et al. (2023a)			✓	✓	linear regression
Huang et al. (2023)	✓		✓		linear regression
Wu et al. (2023a)			✓		linear regression
Ours	✓	✓	✓	✓	classification

Table 1. Comparison with existing works about training analysis and generalization guarantee of in-context learning

consider the generalization using magnitude pruning under a feature learning framework. However, these works are built on convolutional neural networks, and no theoretical works are for LLM or Transformer-based models.

2. Problem Formulation

This work studies the optimization and generalization of binary classification problems for in-context learning. Consider a query \mathbf{x}_{query} and its label z . Define a set of binary classification tasks \mathcal{T} , consisting of multiple task functions. The label $z \in \{+1, -1\}$ is mapped from $\mathbf{x}_{query} \in \mathbb{R}^{d_x}$ through a task f that is randomly chosen from \mathcal{T} , i.e., $z = f(\mathbf{x}_{query}) \in \{+1, -1\}$, $f \in \mathcal{T}$.

2.1. Training to Enhance ICL Capability

Following the framework of training for ICL in (Garg et al., 2022; Akyürek et al., 2023; Bai et al., 2023), we consider the problem of training such that the model has the ICL capability to generalize to new tasks using prompts. The idea is to update the model during the training process using pairs of the constructed prompt, embedded as \mathbf{P} for the query \mathbf{x}_{query} , and its label $f(\mathbf{x}_{query})$. We start by formulating \mathbf{P} and then introduce the learning model in this section.

Following (Von Oswald et al., 2023; Zhang et al., 2023a; Huang et al., 2023), the prompt embedding \mathbf{P} of query \mathbf{x}_{query} is formulated as:

$$\mathbf{P} = \begin{pmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_l & \mathbf{x}_{query} \\ \mathbf{y}_1 & \mathbf{y}_2 & \cdots & \mathbf{y}_l & \mathbf{0} \end{pmatrix} \quad (1)$$

$$:= (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_{query}) \in \mathbb{R}^{(d_x+d_y) \times (l+1)},$$

where the last column of \mathbf{P} , denoted by \mathbf{p}_{query} , includes the query \mathbf{x}_{query} with padding zeros, and the first l columns are the contexts for \mathbf{x}_{query} . We respectively call \mathbf{x}_i and \mathbf{y}_i , $i \in [l]$ context inputs and outputs, where l is also known as the prompt length. Let $\text{Emb}(\cdot)$ be the embedding function of each context output. $\mathbf{y}_i \in \mathbb{R}^{d_y}$ in (1) is defined as $\mathbf{y}_i = \text{Emb}(f(\mathbf{x}_i))$. Hence, \mathbf{P} is a function of f . The first d_x dimensions of \mathbf{p}_i are referred to as the feature embedding, while the last d_y dimensions are called the label embedding.

The **learning model** is a single-head, one-layer Transformer with one self-attention layer and one two-layer perceptron. Mathematically, it can be written as

$$F(\Psi; \mathbf{P}) = \mathbf{a}^\top \text{Relu}(\mathbf{W}_O \sum_{i=1}^l \mathbf{W}_V \mathbf{p}_i \cdot \text{attn}(\Psi; \mathbf{P}, i)), \quad (2)$$

$$\text{attn}(\Psi; \mathbf{P}, i) = \text{softmax}((\mathbf{W}_K \mathbf{p}_i)^\top \mathbf{W}_Q \mathbf{p}_{query}),$$

where $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{m_a \times (d_x+d_y)}$, $\mathbf{W}_V \in \mathbb{R}^{m_b \times (d_x+d_y)}$ are the embedding matrices for queries, keys, and values, respectively, and $\mathbf{W}_O \in \mathbb{R}^{m \times m_b}$ and $\mathbf{a} \in \mathbb{R}^m$ are parameters in the MLP layer. $\Psi := \{\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V, \mathbf{W}_O, \mathbf{a}\}$ denotes the set of all model weights. Typically, $\min(m_a, m_b) > d_x + d_y$.

The **training problem** to enhance the ICL capability solves the empirical risk minimization problem,

$$\min_{\Psi} R_N(\Psi) := \frac{1}{N} \sum_{n=1}^N \ell(\Psi; \mathbf{P}^n, z^n), \quad (3)$$

using N pairs of prompt embedding and label pairs $\{\mathbf{P}^n, z^n\}_{n=1}^N$. For the n -th pair, \mathbf{x}_{query}^n and the context input \mathbf{x}_i^n are all sampled from an unknown distribution \mathcal{D} , the task f^n is sampled from \mathcal{T} , and \mathbf{P}^n is constructed following (1). The loss function is a Hinge loss, i.e., $\ell(\Psi; \mathbf{P}^n, z^n) = \max\{0, 1 - z^n \cdot F(\Psi; \mathbf{P}^n)\}$, where $F(\Psi; \mathbf{P}^n)$ is defined in (2). Let $\mathcal{T}_{tr} = \bigcup_{n=1}^N f^n$ denote the set of tasks that appear in the training samples. Note that $\mathcal{T}_{tr} \subset \mathcal{T}$, and (3) is a *multi-task learning* problem when $|\mathcal{T}_{tr}| > 1$.

2.2. Generalization Evaluation

We define two quantities to evaluate the ICL generalization performance to new tasks as follows.

In-Domain Generalization: If the testing queries are also drawn from \mathcal{D} and all the testing tasks are drawn from \mathcal{T} , we call it *in-domain* inference, and the in-domain generalization error is defined as¹

$$\mathbb{E}_{\mathbf{x}_{query} \sim \mathcal{D}, f \in \mathcal{T}} [\ell(\Psi; \mathbf{P}, z)], \quad (4)$$

where \mathbf{P} is defined in (1). Note that the in-domain performance includes the testing performance on *unseen* tasks in

¹In terms of evaluating generalization on unseen tasks, (4) is almost equivalent to replacing $f \in \mathcal{T}$ with $f \in \mathcal{T} \setminus \mathcal{T}_{tr}$ in the subscript. This is because we later prove that all of our analysis can hold when training on a small fraction of tasks (Condition 3.2). Therefore, an $\mathcal{O}(\epsilon)$ generalization error on $f \in \mathcal{T}$ can indeed reflect an $\mathcal{O}(\epsilon)$ generalization error on $f \in \mathcal{T} \setminus \mathcal{T}_{tr}$.

$\mathcal{T} \setminus \mathcal{T}_{tr}$ that do not appear in the training samples.

Out-of-Domain Generalization: Suppose the testing queries \mathbf{x}_{query} follow the distribution \mathcal{D}' ($\mathcal{D}' \neq \mathcal{D}$), and the binary classification tasks that map the testing queries to the labels are drawn a set \mathcal{T}' ($\mathcal{T}' \neq \mathcal{T}$). Then, the *out-of-domain* generalization error can be defined as

$$\mathbb{E}_{\mathbf{x}_{query} \sim \mathcal{D}', f \in \mathcal{T}'}[\ell(\Psi; \mathbf{P}, z)]. \quad (5)$$

2.3. Training Algorithm

The model is trained using stochastic gradient descent (SGD) with step size η with batch size B , summarized in Algorithm 1 in Appendix C. \mathbf{W}_Q , \mathbf{W}_K and \mathbf{W}_V are initialized such that all diagonal entries of $\mathbf{W}_V^{(0)}$, and the first $d_{\mathcal{X}}$ diagonal entries of $\mathbf{W}_Q^{(0)}$ and $\mathbf{W}_K^{(0)}$ are set as δ with $\delta \in (0, 0.2]$, and all other entries are 0. Each entry of $\mathbf{W}_Q^{(0)}$ is generated from $\mathcal{N}(0, \xi^2)$, $\xi = 1/\sqrt{m}$ and each entry of \mathbf{a} is uniformly sampled from $\{1/m, -1/m\}$. Besides, \mathbf{a} does not update during training.

2.4. Model Pruning

We also consider the case that the learned model Ψ is pruned to reduce the inference computation. Let $\mathcal{S} \subset [m]$ denote the index set of neurons in the output layer. Pruning neurons in \mathcal{S} correspond to removing the corresponding rows in \mathbf{W}_O , resulting in the reduced matrix size of $(m - |\mathcal{S}|) \cdot m_b$.

3. Theoretical Results

We first summarize the main insights in Section 3.1. Section 3.2 formally presents our analysis model. Section 3.3 presents the formal theoretical results on the learning performance and the resulting ICL generalization. Section 3.4 provides the theoretical result that magnitude-based pruning on the out layer does not hurt ICL performance.

3.1. Main Theoretical Insights

We consider a class of binary classification tasks where the binary labels in each task are determined by two out of M_1 *in-domain-relevant patterns*. The training data include pairs of prompt embedding and labels from a small subset of these tasks. In-domain generalization evaluates the ICL capability of the learned model on tasks using all possible combinations of these M_1 patterns. Out-of-domain generalization further evaluates the binary classification tasks that are determined by pairs of *out-of-domain-relevant patterns*, which are some linear combinations of these M_1 patterns.

P1. Quantitative Learning Analysis With Guaranteed In- and Out-of-Domain Generalization. We quantitatively prove the learned model achieves desirable generalization in both in-domain and out-of-domain tasks. The required number of training data and iterations are polynomial in β^{-1}

and α^{-1} , where β represents the norm of relevant patterns, and α denotes the fraction of context inputs with the same in-domain-relevant pattern as the query. A higher α implies that the context examples offer more information about the query, consequently reducing the sample requirements and expediting the learning process.

P2. Mechanism of Transformers in Implementing ICL.

We elucidate the mechanism where the learned Transformers make predictions in- and out-of-domain in context. We quantitatively show that the self-attention layer attends to context examples with relevant patterns of the query task and promotes learning of these relevant patterns. Then, the two-layer perceptron promotes the label embeddings that correspond to these examples so as to predict the label of the query accurately.

P3. Magnitude-Based Pruning Preserves ICL. We quantify the ICL generalization if neurons with the smallest magnitude after training in the MLP layer are removed and prove that the generalization is almost unaffected even when a constant fraction of neurons are removed. In contrast, the generalization error is proved to be at least $\Omega(R)$ when R fraction of neurons with large magnitude are removed.

3.2. The Modeling of Training Data and Tasks

In-Domain Data and Tasks. Consider M_1 *in-domain-relevant (IDR)* patterns $\{\boldsymbol{\mu}_j\}_{j=1}^{M_1}$ and M_2 ($= \mathcal{O}(M_1)$) *in-domain-irrelevant (IDI)* patterns $\{\boldsymbol{\nu}_k\}_{k=1}^{M_2}$ ($M_1 + M_2 = d_{\mathcal{X}}$) in $\mathbb{R}^{d_{\mathcal{X}}}$, where these $M_1 + M_2$ patterns are pairwise orthogonal, and $\|\boldsymbol{\mu}_j\| = \|\boldsymbol{\nu}_k\| = \beta \geq 1$ (β is a constant) for $j \in [M_1], k \in [M_2]$. Each in-domain data \mathbf{x} drawn from \mathcal{D} is generated by

$$\mathbf{x} = \boldsymbol{\mu}_j + \kappa \boldsymbol{\nu}_k, \quad (6)$$

where $j \in [M_1]$ and $k \in [M_2]$ are arbitrarily selected. κ follows a uniform distribution $U(-K, K)$, $K \leq 1/2$. Denote $\text{IDR}(\mathbf{x}) := \boldsymbol{\mu}_j$ as the IDR pattern in data \mathbf{x} . Our data assumption originates from recent feature learning works on deep learning (Allen-Zhu & Li, 2023; Li et al., 2023a; Oymak et al., 2023) for language and vision data. To the best of our knowledge, only (Huang et al., 2023) theoretically analyzes the performance of ICL with softmax attention, assuming all \mathbf{x} are orthogonal to each other. Our assumption in (6) is more general than that in (Huang et al., 2023).

Each in-domain task is defined as a binary classification function that decides the label based on two IDR patterns in the query. Specifically,

Definition 3.1. (Definition of in-domain tasks) The in-domain task set \mathcal{T} includes $M_1(M_1 - 1)$ tasks such that each task $f \in \mathcal{T}$ is defined as

$$f(\mathbf{x}) = \begin{cases} +1, & \text{IDR}(\mathbf{x}) = \boldsymbol{\mu}_a, \\ -1, & \text{IDR}(\mathbf{x}) = \boldsymbol{\mu}_b, \\ \text{random from } \{+1, -1\}, & \text{otherwise,} \end{cases} \quad (7)$$




Figure 1. (A) Example of prompt embedding. $l = 3$, $\alpha = 2/3$. (B) The mechanism of a trained Transformer (2) to implement ICL.

where μ_a, μ_b are two different patterns in $\{\mu_j\}_{j=1}^{M_1}$ and are the decisive patterns for task f .

From (7), the task f outputs label +1 (or -1) if the IDR pattern is μ_a (or μ_b). If the data contains neither of these two patterns, the label is random.

Out-of-Domain Data and Tasks. Assume there are M'_1 out-of-domain-relevant (ODR) patterns $\{\mu'_j\}_{j=1}^{M'_1}$ and M'_2 out-of-domain-irrelevant (ODI) patterns $\{\nu'_k\}_{k=1}^{M'_2}$. Any data x drawn from \mathcal{D}' can be generated by

$$x = \mu'_j + \kappa' \nu'_k \quad (8)$$

where $j \in [M'_1]$ and $k \in [M'_2]$ are arbitrarily selected, and $\kappa' \sim U(K', K')$ for $K' = \mathcal{O}(1)$. We use $\text{ODR}(x) := \mu'_j$ to denote the ODR pattern of x .

The set of out-of-domain tasks \mathcal{T}' contains $M'_1(M'_1 - 1)$ binary classification problems that are defined in the same fashion as Definition 3.1, with the only difference of using $\{\mu'_j\}_{j=1}^{M'_1}$ rather than $\{\mu_j\}_{j=1}^{M_1}$ to determine labels.

Prompt Construction for Training and Testing. Let l_{tr} and l_{ts} denote the length of training and testing contexts, respectively.

Training prompt embedding: Given an input-label pair x_{query} and $f(x_{query})$ for training, the context inputs x_i in \mathbf{P} in (1) are constructed as follows. The IDR pattern is selected from $\{\mu_j\}_{j=1}^{M_1}$ following a categorical distribution parameterized by α , where $\alpha = \Theta(1) \in (0, 1]$. Specifically, each of μ_a and μ_b (the decisive patterns of task f) is selected with probability $\alpha/2$, and each of these other $M_1 - 2$ patterns elected with probability $(1 - \alpha)/(M_1 - 2)$. The context labels are determined by task f .

Testing prompt embedding: The context inputs for the testing query can be selected following a wide range of prompt selection methods (Liu et al., 2022b; Rubin et al., 2022; Wu et al., 2023b). Given an in-domain (or out-of-domain) task f that has decisive patterns μ_a and μ_b (or μ'_a and μ'_b), we only assume at least $\alpha'/2$ ($\alpha' \in (0, 1]$) fraction of context inputs contain the same IDR (or ODR) pattern as the query.

For the label embedding y_i for both training and testing, $\text{Emb}d(+1) = \mathbf{q}$, $\text{Emb}d(-1) = -\mathbf{q}$, where $\mathbf{q} \in \mathbb{R}^{d_y}$. Hence, $y_i \in \{\mathbf{q}, -\mathbf{q}\}$ for $i \in [l_{tr}]$ or $i \in [l_{ts}]$.

3.3. In-Domain and Out-of-Domain Generalization With Sample Complexity Analysis

In order for the learned model $F(\Psi)$ to generalize all tasks in \mathcal{T} through ICL, the training tasks in \mathcal{T}_{tr} should uniformly cover all the possibilities of IDR patterns and labels, as stated by the following condition,

Condition 3.2. For any given $j \in [M_1]$ and either label +1 or -1, the number of tasks in \mathcal{T}_{tr} that map μ_j to that label is $|\mathcal{T}_{tr}|/M_1 (\geq 1)$.

Note that Condition 3.2 is easy to meet, and $|\mathcal{T}_{tr}|$ does not have to be large. In fact, $|\mathcal{T}_{tr}|$ can be as small as M_1 . For example, let the i -th task function ($i \in [M_1 - 1]$) in \mathcal{T}_{tr} map the queries with μ_i and μ_{i+1} as IDR patterns to +1 and -1, respectively. The M_1 -th task function maps μ_{M_1} and μ_1 to +1 and -1, respectively. We can easily verify \mathcal{T}_{tr} satisfies Condition 3.2 in this case.

Following (Shi et al., 2021; Karp et al., 2021; Li et al., 2023a), we assume the training labels are balanced, i.e., $|\{n : z^n = +1\}| - |\{n : z^n = -1\}| = \mathcal{O}(\sqrt{N})$. The next theorem states the training and in-domain generalization.

Theorem 3.3. (In-Domain Generalization) Suppose Condition 3.2 holds. For any $\epsilon > 0$, when (i) the number of neurons in \mathbf{W}_O satisfies $m \geq \Omega(M_1^2 \log M_1)$, (ii) batch size $B > \Omega(\max\{\epsilon^{-2}, M_1\} \cdot \log M_1)$, (iii) the lengths of training and testing contexts are

$$l_{tr} \geq \max\{\Omega(\log M_1/\alpha), \Omega(1/(\beta^2 \alpha))\}, l_{ts} \geq \alpha'^{-1}, \quad (9)$$

(iv) and the number of iterations satisfies

$$T = \Theta(\eta^{-1} M_1 \alpha^{-\frac{2}{3}} \beta^{-2/3} \sqrt{\log M_1}), \quad (10)$$

with step size $\eta \leq 1$ and $N = BT$ samples, then with a high probability, the returned model satisfies that

$$\mathbb{E}_{x_{query} \sim \mathcal{D}, f \in \mathcal{T}} [\ell(\Psi; \mathbf{P}, z)] \leq \mathcal{O}(\epsilon). \quad (11)$$

Theorem 3.3 characterizes the sufficient condition on the model size, the required number of iterations, and the number of prompt embedding and label pairs, such that the trained model achieves an in-domain generalization error of $\mathcal{O}(\epsilon)$. Theorem 3.3 includes three major insights:

1. *In-domain generalization capability using a diminishing fraction of training tasks.* Because \mathcal{T}_{tr} can satisfy Condition

3.2 even when $|\mathcal{T}_{tr}| = M_1$, then the number of training tasks is only a fraction $(M_1 - 1)^{-1/2}$ of the total number of in-domain tasks in \mathcal{T} .

2. (*Context length*) The required length of training and testing contexts increase in the order of α^{-1} and α'^{-1} , respectively, which implies that a longer context is needed when the fraction of IDR patterns in the context is small.

3. (*Convergence and sample complexity*) The required number of iterations and the training samples is proportional to $\alpha^{-2/3}$. This indicates that a larger fraction of the IDR pattern in the context leads to more efficient convergence and generalization.

Based on the in-domain result, we can also investigate the properties of out-of-domain generalization.

Theorem 3.4. (*Out-of-Domain Generalization*)

Suppose Condition 3.2 and conditions (i)-(iv) in Theorem 3.3 hold. For any $\boldsymbol{\mu}'_1, \dots, \boldsymbol{\mu}'_{M_1}, \boldsymbol{\nu}'_1, \boldsymbol{\nu}'_{M_2}$ that are pairwise orthogonal and $\|\boldsymbol{\mu}'_j\| = \|\boldsymbol{\nu}'_k\| = \beta$, if

$$\boldsymbol{\mu}'_j \in \left\{ \sum_{i=1}^{M_1} k_{j,i} \boldsymbol{\mu}_i \mid S_j := \sum_{i=1}^{M_1} k_{j,i} \geq 1, k_{j,i} \in \mathbb{R} \right\}, \quad (12)$$

and $\boldsymbol{\nu}'_k \in \text{span}\{\boldsymbol{\nu}_1, \boldsymbol{\nu}_2, \dots, \boldsymbol{\nu}_{M_2}\}$, $j \in [M_1]$, $k \in [M_2]$, then with high probability, the learned model can achieve an out-of-domain generalization error of

$$\mathbb{E}_{\mathbf{x}_{query} \sim \mathcal{D}', f \in \mathcal{T}'} [\ell(\Psi; \mathbf{P}, z)] \leq \mathcal{O}(\epsilon). \quad (13)$$

Remark 3.5. Theorem 3.4 indicates that a one-layer Transformer can generalize well in context, even in the presence of distribution shifts between the training and testing data. The conditions for a favorable generalization encompass the following: (1) the ODR patterns are linear combinations of IDR patterns with a summation of coefficients ≥ 1 , and each ODI pattern is in the subspace spanned by IDI patterns; (2) the testing prompt is long enough, which is linear in α'^{-1} , to include context inputs involving ODR patterns.

Remark 3.6. (Comparison with existing ICL analysis) (Huang et al., 2023) analyzes the generalization performance of ICL on unseen tasks under a similar data model that includes decisive and indecisive patterns. However, (Huang et al., 2023) only analyzes in-domain unseen tasks, while our results also apply to one type of out-of-domain tasks through data shift. To the best of our knowledge, only (Zhang et al., 2023a) studies out-of-domain generalization under the setup of linear regression problems with Gaussian inputs. They conclude that, under this setup, the covariate shift, i.e., the difference between the training and testing data distributions \mathcal{D} and \mathcal{D}' , does not guarantee generalization. We consider classification problems under a data model different from (Zhang et al., 2023a). We provide the out-of-domain generalization guarantee for one type of distribution between \mathcal{D} and \mathcal{D}' .

3.4. ICL With Magnitude-Based Model Pruning

Theorem 3.7. Let \mathbf{r}_i be the i -row of $\mathbf{W}_O \mathbf{W}_V$, $i \in [m]$. Suppose Condition 3.2 and conditions (i)-(iv) in Theorem 3.3 hold, then there exists $\mathcal{L} \subset [m]$ with $|\mathcal{L}| = \Omega(m)$ s.t.,

$$\begin{aligned} \|\mathbf{r}_i^{(T)}\| &\geq \Omega(1), i \in \mathcal{L}, \\ \|\mathbf{r}_i^{(T)}\| &\leq (1/\sqrt{M_2}), i \in \mathcal{L}^c, \end{aligned} \quad (14)$$

where \mathcal{L}^c is the complementary set of \mathcal{L} . Then, for any $\epsilon > 0$ and any in- or out-of-domain $\mathbf{x}_{query} \sim \mathcal{D}$ (or \mathcal{D}') and corresponding $f \in \mathcal{T}$ (or \mathcal{T}'), pruning all neurons $i \in \mathcal{L}^c$ leads to a generalization error

$$\mathbb{E}_{\mathbf{x}_{query}, f} [\ell(\Psi_{\mathcal{L}^c}; \mathbf{P}, z)] \leq \mathcal{O}(\epsilon + M_1^{-1/2}), \quad (15)$$

where $\Psi_{\mathcal{L}^c}$ represents the model weights after removing neurons in \mathcal{L}^c in \mathbf{W}_O . In contrast, pruning $\mathcal{S} \subset \mathcal{L}$ with size $|\mathcal{S}| = Rm$, where $R \in (0, 1)$ and is a constant, and $\alpha' \geq \Omega(M_1^{-0.5})$ results in a generalization error of

$$\mathbb{E}_{\mathbf{x}_{query}, f} [\ell(\Psi_{\mathcal{S}}; \mathbf{P}, z)] \geq \Omega(R + (\alpha' M_1)^{-1}). \quad (16)$$

Remark 3.8. Theorem 3.7 proves that a constant fraction of neurons in \mathcal{L} in the trained MLP layer has large weights, while the remaining ones in \mathcal{L}^c have small weights. Pruning neurons with a smaller magnitude leads to almost the same generalization result as that of the unpruned Ψ . However, pruning neurons with a larger magnitude cause an increasing generalization error as the pruning ratio R increases. Theorem 3.7 indicates that in our setup, magnitude-based pruning on \mathbf{W}_O does not hurt the model's ICL capability.

4. The Mechanism of ICL by the Trained Transformer

Here, we provide a detailed discussion about how the generalization performance in Theorems 3.3 and 3.4 are achieved. We first introduce novel properties of the self-attention layer and the MLP layer of the learned Transformer to implement ICL in Sections 4.1 and 4.2. The high-level proof idea of Theorems 3.3 and 3.4 is presented in Appendix D.1.

4.1. Self-Attention Selects Contexts With the Same IDR/ODR Pattern as the Query

We first show the learned self-attention layer promotes context examples that share the same IDR/ODR pattern as the query. Specifically, for any vector $\mathbf{p} \in \mathbb{R}^{d_x + d_y}$ that includes input \mathbf{x} and the corresponding output embedding \mathbf{y} . We use $\text{XDR}(\mathbf{p})$ to represent the relevant pattern, which is the IDR(\mathbf{x}) for in-domain data and ODR(\mathbf{x}) for out-of-domain data. Then

Proposition 4.1. The trained model after being updated by T (characterized in (10)) iterations satisfies that, for any $(\mathbf{p}, \mathbf{W}) \in \{(\mathbf{p}_{query}, \mathbf{W}_Q^{(T)}), (\mathbf{p}_i, \mathbf{W}_K^{(T)})\}_{i=1}^l$,

$$\|[\text{XDR}(\mathbf{p})^\top, \mathbf{0}^\top] \mathbf{W} \mathbf{p}\| \geq \Omega(\sqrt{\log M_1}), \quad (17)$$




Figure 2. The properties of the trained model. (A) The average norm of $W_Q p_{query}$, $W_K p_i$, $[XDR(p_{query})^\top / \beta, \mathbf{0}^\top] \cdot W_Q p_{query}$, and $[XDR(p_i)^\top / \beta, \mathbf{0}^\top] W_K p_i$. (B) The attention weight summation on contexts with the same ODR pattern as the query and other contexts. (C) The magnitude of the feature embedding of 5 neurons in $W_O W_V$ and their angles to $\bar{\mu}$ in 400 epochs. (D) The magnitude of the label embedding of 10 neurons in $W_O W_V$ and their angles to q in 400 epochs. We choose 5 neurons for $a_i > 0$ and 5 for $a_i < 0$.

$$\|[\mathbf{a}^\top, \mathbf{0}^\top] W \mathbf{p}\| \leq \mathcal{O}(\sqrt{\log M_1}(1/M_1 + 1/M_2)), \quad (18)$$

$$\|[\mathbf{b}^\top, \mathbf{0}^\top] W \mathbf{p}\| \leq \mathcal{O}(\sqrt{\log M_1}(1/M_1 + 1/M_2)), \quad (19)$$

where \mathbf{a} is any IDR (or ODR) pattern that is different from $XDR(\mathbf{p})$ for in-domain (or out-of-domain) tasks, \mathbf{b} is any IDI (or ODI) pattern, and $\mathbf{0}$ is an all-zero vector in $\mathbb{R}^{m_a - d_x}$.

Remark 4.2. Proposition 4.1 indicates that the self-attention layer parameters $W_Q^{(T)}$ and $W_K^{(T)}$ in the returned model projects p_{query} or context embeddings p_i mainly to the directions of the corresponding IDR pattern for in-domain data or ODR pattern for out-of-domain data. This can be deduced by combining (17), (18), and (19), since components of $W \mathbf{p}$ in other directions rather than $[XDR(\mathbf{p})^\top, \mathbf{0}^\top]$ are relatively smaller. Hence, Proposition 4.1 implies that the learned $W_Q^{(T)}$ and $W_K^{(T)}$ remove the effect of IDI/ODI patterns. Meanwhile, (17) states that the $W_Q^{(T)}$ and $W_K^{(T)}$ enlarge the magnitude of the IDR or ODR patterns from $\Theta(1)$ to $\Theta(\sqrt{\log M_1})$, given that the $W_Q^{(0)}$ and $W_K^{(0)}$ are initialized with a scalar $\delta = \Theta(1)$.

Proposition 4.1 enables us to compute the attention map of the trained model. Therefore, we have the following.

Corollary 4.3. For any testing query embedding $p_{query} = [x_{query}^\top, \mathbf{0}^\top]^\top$, let $\mathcal{N}_* \in [l]$ be the set of indices of context inputs that share the same IDR (or ODR) pattern as the in-domain (or out-of-domain) x_{query} . Then, for any constant $C > 1$, by definition in (2), it holds that

$$\sum_{s \in \mathcal{N}_*} \text{attn}(\Psi; \mathbf{P}, i) \geq 1 - \Theta(1/M_1^C). \quad (20)$$

Remark 4.4. Corollary 4.3 shows that after training, the attention weights become concentrated on contexts in \mathcal{N}_* . This means that the learned self-attention layer only selects some crucial contexts that share the same IDR/ODR pattern as the query rather than all samples uniformly or randomly.

4.2. MLP Neurons Distinguish Label Embeddings Rather Than Feature Embeddings.

We next show that the trained MLP layer can distinguish the label embeddings for data from different classes.

Proposition 4.5. Let r_i introduced in Theorem 3.7 be $(r_{i_{d_x}}^\top, r_{i_{d_y}}^\top, r_i^{\prime\top})$ where $r_{i_{d_x}} \in \mathbb{R}^{d_x}$, $r_{i_{d_y}} \in \mathbb{R}^{d_y}$, and $r_i' \in \mathbb{R}^{m_b - d_x - d_y}$. Then, for any $i \in \mathcal{L}$

$$r_{i_{d_x}}^{(T)} \bar{\mu} / (\|r_{i_{d_x}}^{(T)}\| \cdot \|\bar{\mu}\|) \geq 1 - \Theta(1)/M_2, \quad (21)$$

$$r_{i_{d_y}}^{(T)} q_e / (\|r_{i_{d_y}}^{(T)}\| \cdot \|q_e\|) \geq 1 - \Theta(1)/M_1, \quad (22)$$

where $\bar{\mu} = \sum_{k=1}^{M_1} \mu_k^\top / M_1$, $q_e = q$ if $a_i > 0$ and $q_e = -q$ if $a_i < 0$, where a_i is the i -th entry of \mathbf{a} in (1).

Remark 4.6. Proposition 4.5 demonstrates that neurons with indices in \mathcal{L} have the following two properties. (P1) The first d_x entries of all the corresponding row vectors in $W_O^{(T)} W_V^{(T)}$ approximate the average of all IDR patterns μ_j , $j \in [M_1]$. (P2) The next d_y entries of the i th row of $W_O^{(T)} W_V^{(T)}$ approximates the label embedding q when $a_i > 0$ and approximates $-q$ when $a_i < 0$. (P1) indicates that the output layer focuses on all IDR patterns equally rather than any IDI pattern. (P2) indicates that the MLP layer can distinguish label embeddings for different classes.

5. Numerical Experiments

Data Generation We verify our theoretical findings using data generated as described in Section 2. Let $d_x = d_y = 30$, $\beta = 3$, $K' = 5$, $K = 0.5$. The in-context binary classification error is evaluated by $\mathbb{E}_{(\mathbf{x}, y)} [\Pr(y \cdot F(\Psi; \mathbf{P}) < 0)]$ for \mathbf{x} following either \mathcal{D} or \mathcal{D}' and \mathbf{P} constructed in (1). If not otherwise specified, we set $M_1 = 6$, $M_2 = 24$. For out-of-domain generalization, $M_1' = 3$, $\nu_i' = \nu_i$ for $i \in [M_2]$. $\mu_1' = 0.3 \cdot (\mu_1 - \mu_2) + a\mu_5 + b\mu_6$. $\mu_2' = \sqrt{2}/2 \cdot (\mu_1 + \mu_2)$. $\mu_3' = \sqrt{2}/2 \cdot (\mu_3 + \mu_4)$. For testing, we select contexts with the two decisive patterns with $\alpha'/2$ probability each and others with $(1 - \alpha')/(M_1' - 2)$ probability each to keep the context outputs balanced.

Model and Training Setup: The models we use include both the one-layer Transformer defined in (2) and the 3-layer 2-head real-world model GPT-2 (Radford et al., 2019) following (Bai et al., 2023; Wu et al., 2023a). If not otherwise specified, we set $\alpha = 0.8$, $l_{tr} = 20$ for training. The training tasks are formulated as follows to satisfy Condition

3.2. Define $\mathbf{a}_i = \mathbf{a}_{i+M_1} = \boldsymbol{\mu}_i$ for $i \in [M_1]$, and then the $((k-1) \cdot M_1 + j)$ -th task function maps the queries with \mathbf{a}_j and \mathbf{a}_{j+k} as IDR patterns to $+1$ and -1 , respectively, for $j \in [M_1]$ and $k \in [U]$. For the one-layer Transformer, we use $U = 1$ and $m_a = m_b = 60$. Hence, $|\mathcal{T}_{tr}| = 6$, and there are $|\mathcal{T} \setminus \mathcal{T}_{tr}| = 24$ in-domain unseen tasks. For GPT-2, $U = 4$. Then, $|\mathcal{T}_{tr}| = 24$, $|\mathcal{T} \setminus \mathcal{T}_{tr}| = 6$. Note that we evaluate in-domain generalization error only on unseen tasks $\mathcal{T} \setminus \mathcal{T}_{tr}$, which is generally an upper bound of that defined in (4) after sufficient training.

5.1. Experiments on the Generalization of ICL

We first verify the sufficient condition (12) for out-of-domain generalization. From the selection of $\boldsymbol{\mu}'$'s, we know that $S_1 = a + b$, $S_2 = S_3 = \sqrt{2}$. We vary a and b while satisfying $a^2 + b^2 + 2 \cdot 0.3^2 = 1$. Figure 3 (A) shows that the out-of-domain classification error archives < 0.01 when $S_1 \geq 1$ and deviates from 0 when $S_1 < 1$, which justifies the necessity of condition (12). We then investigate how the context length is affected by α' , i.e., the fraction of contexts with the same IDR/ODR pattern as the query. Figure 3 (B) indicates that a longer testing context length is needed when α' is smaller for in- or out-of-domain, which is consistent with the lower bound of l_{ts} in (9) and Theorem 3.4.




Figure 3. Out-of-domain ICL classification error on GPT-2 with (a) different S_1 on GPT-2 (b) different α' for in-domain (id) and out-of-domain (ood) generalization.

We then compare ICL with other machine learning algorithms for classification, where contexts are used as training samples for these methods. Figure 4 (A) and (B) show that when $\alpha' = 0.8$, the advance of ICL over other algorithms is not significant, while when $\alpha' = 0.6$, ICL is the most sample-efficient for a small generalization error. Thus, ICL can remove irrelevant data and is more robust to random noise in labels than other learning algorithms.

We also investigate the effect of pruning techniques on ICL. Let $\alpha = 0.6$. Figure 5 (A) shows that magnitude-based pruning does not hurt out-of-domain generalization if the pruning rate is lower than around 15%, which is the ratio of \mathbf{W}_O neurons with a small magnitude. The generalization error increases as the pruning rate increases when pruning neurons with large weights. This is consistent with Theorem 3.7 and Remark 3.8. Figure 5 (B) justifies the impact of α' in Theorem 3.7 that larger α' can improve the performance of the pruned model.

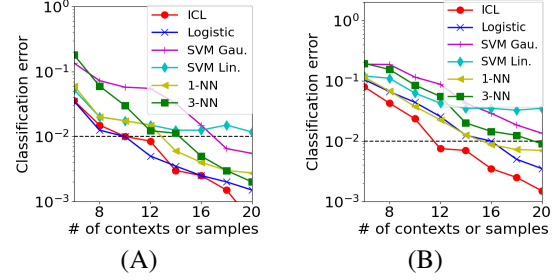


Figure 4. Binary classification performance of using ICL, logistic regression (Logistic), SVM with Gaussian kernel (SVM Gau.), SVM with linear kernel (SVM Lin.), 1-nearest neighbor (1-NN), and 3-nearest neighbor (3-NN) with one-layer Transformer when (A) $\alpha' = 0.8$ (B) $\alpha' = 0.6$.

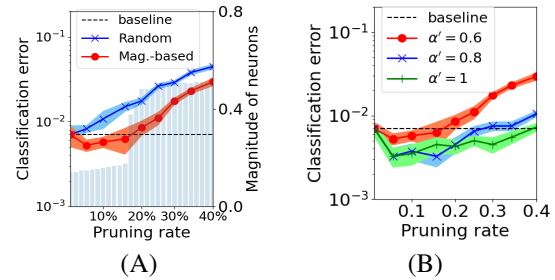


Figure 5. (A) Out-of-domain classification error (left y-axis for curves) with model pruning of the trained \mathbf{W}_O using baseline (no pruning), random pruning, and magnitude-based pruning (Mag.-based), and the magnitude of each neuron of \mathbf{W}_O (right y-axis for light blue bars) (B) Out-of-domain classification error when varying α' . These two are implemented on a one-layer Transformer.

5.2. Experiments on the Mechanism of ICL

We examine our findings regarding the mechanism of ICL in Section 4 using a one-layer Transformer formulated in (2). In Figure 2 (A) and (B), we consider out-of-domain data with $a = b = 0.64$. Figure 2 (A) shows that for any query \mathbf{p}_{query} (or context example \mathbf{p}_i for $i \in l_{ts}$), the norm of $[\mathbf{XDR}(\mathbf{p})^\top, \mathbf{0}^\top] \mathbf{W}_Q \mathbf{p}_{query}$ (or $[\mathbf{XDR}(\mathbf{p})^\top, \mathbf{0}^\top] \mathbf{W}_K \mathbf{p}_i$) is close to the norm of $\mathbf{W}_Q \mathbf{p}_{query}$ (or $\mathbf{W}_K \mathbf{p}_i$). This implies that the components of $\mathbf{W}_Q \mathbf{p}_{query}$ (or $\mathbf{W}_K \mathbf{p}_i$) in directions other than $[\mathbf{XDR}(\mathbf{p})^\top, \mathbf{0}^\top]$ are small, which is consistent with (18) and (19) in Proposition 4.1. Moreover, these norms increase from initialization during training, which justifies (17). Figure 2 (B) depicts the concentration of attention on contexts in \mathcal{N}_* after training. This verifies Corollary 4.3. Figure 2 (C) and (D) jointly verify Proposition 4.5. The color bars represent the epochs of training. We can observe that except for some neurons, $\mathbf{r}_{i_d_x}$ grows to be close to the direction of $\bar{\boldsymbol{\mu}}$ with a larger magnitude in Figure 2 (C). Moreover, Figure 2 (D) shows for $a_i > 0$ (or $a_i < 0$), $\mathbf{r}_{i_d_y}$ becomes close to \mathbf{q} (or $-\mathbf{q}$) with a large magnitude.

6. Conclusion

This paper provides theoretical analyses of the training dynamics of Transformers with nonlinear attention and nonlinear MLP, and the resulting ICL capability for new tasks with possible data shift. This paper also provides a theoretical justification for magnitude-based pruning to reduce inference costs while maintaining the ICL capability. Future directions include designing practical prompt selection algorithms and model pruning methods based on the obtained insights, as well as investigating ICL on generation tasks.

Acknowledgements

This work was supported by IBM through the IBM-Rensselaer Future of Computing Research Collaboration. We thank Dr. Hui Wan (hui.wan@gmail.com) at Google Cloud AI, USA, for the valuable discussion. We thank all anonymous reviewers for their constructive comments.

Impact Statement

This paper aims to explore the mechanisms of transformer-based neural networks in the context of in-context learning. The primary focus is on the mathematical analysis of generalization errors. To the best of our knowledge, no potential societal consequences are associated with our work.

References

- Ahn, K., Cheng, X., Daneshmand, H., and Sra, S. Transformers learn to implement preconditioned gradient descent for in-context learning. *arXiv preprint arXiv:2306.00297*, 2023.
- Akyürek, E., Schuurmans, D., Andreas, J., Ma, T., and Zhou, D. What learning algorithm is in-context learning? investigations with linear models. In *The Eleventh International Conference on Learning Representations*, 2023.
- Allen-Zhu, Z. and Li, Y. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. In *The Eleventh International Conference on Learning Representations*, 2023.
- Allen-Zhu, Z., Li, Y., and Liang, Y. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pp. 6155–6166, 2019a.
- Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pp. 242–252. PMLR, 2019b.
- Bai, Y., Chen, F., Wang, H., Xiong, C., and Mei, S. Transformers as statisticians: Provable in-context learning with in-context algorithm selection. *arXiv preprint arXiv:2306.04637*, 2023.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33: 1877–1901, 2020.
- Brutzkus, A. and Globerson, A. An optimization and generalization analysis for max-pooling networks. In *Uncertainty in Artificial Intelligence*, pp. 1650–1660. PMLR, 2021.
- Cao, Y. and Gu, Q. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 10836–10846, 2019.
- Chen, S. and Li, Y. Provably learning a multi-head attention layer. *arXiv preprint arXiv:2402.04084*, 2024.
- Chen, Z., Cao, Y., Gu, Q., and Zhang, T. A generalized neural tangent kernel analysis for two-layer neural networks. *Advances in Neural Information Processing Systems*, 33, 2020.
- Cheng, X., Chen, Y., and Sra, S. Transformers implement functional gradient descent to learn non-linear functions in context. *arXiv preprint arXiv:2312.06528*, 2023.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- Chowdhury, M. N. R., Zhang, S., Wang, M., Liu, S., and Chen, P.-Y. Patch-level routing in mixture-of-experts is provably sample-efficient for convolutional neural networks. In *International Conference on Machine Learning*, 2023.
- Chowdhury, M. N. R., Wang, M., Maghraoui, K. E., Wang, N., Chen, P.-Y., and Carothers, C. A provably effective method for pruning experts in fine-tuned sparse mixture-of-experts. *arXiv preprint arXiv:2405.16646*, 2024.
- Cui, Y., Ren, J., He, P., Tang, J., and Xing, Y. Superiority of multi-head attention in in-context linear regression. *arXiv preprint arXiv:2401.17426*, 2024.
- Daniely, A. and Malach, E. Learning parities with neural networks. *Advances in Neural Information Processing Systems*, 33:20356–20365, 2020.

- Deora, P., Ghaderi, R., Taheri, H., and Thrampoulidis, C. On the optimization and generalization of multi-head attention. *arXiv preprint arXiv:2310.12680*, 2023.
- Frantar, E. and Alistarh, D. Sparsegpt: Massive language models can be accurately pruned in one-shot. In *International Conference on Machine Learning*, pp. 10323–10337. PMLR, 2023.
- Fu, H., Chi, Y., and Liang, Y. Guaranteed recovery of one-hidden-layer neural networks via cross entropy. *IEEE Transactions on Signal Processing*, 68:3225–3235, 2020.
- Garg, S., Tsipras, D., Liang, P. S., and Valiant, G. What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35:30583–30598, 2022.
- Guo, T., Hu, W., Mei, S., Wang, H., Xiong, C., Savarese, S., and Bai, Y. How do transformers learn in-context beyond simple functions? a case study on learning with representations. *arXiv preprint arXiv:2310.10616*, 2023.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- Huang, Y., Cheng, Y., and Liang, Y. In-context convergence of transformers. In *NeurIPS 2023 Workshop on Mathematics of Modern Machine Learning*, 2023.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8571–8580, 2018.
- Jelassi, S., Sander, M., and Li, Y. Vision transformers provably learn spatial structure. *Advances in Neural Information Processing Systems*, 35:37822–37836, 2022.
- Karp, S., Winston, E., Li, Y., and Singh, A. Local signal adaptivity: Provable feature learning in neural networks beyond kernels. *Advances in Neural Information Processing Systems*, 34:24883–24897, 2021.
- Li, H., Wang, M., Liu, S., Chen, P.-Y., and Xiong, J. Generalization guarantee of training graph convolutional networks with graph topology sampling. In *International Conference on Machine Learning*, pp. 13014–13051. PMLR, 2022a.
- Li, H., Zhang, S., and Wang, M. Learning and generalization of one-hidden-layer neural networks, going beyond standard gaussian data. In *2022 56th Annual Conference on Information Sciences and Systems (CISS)*, pp. 37–42. IEEE, 2022b.
- Li, H., Wang, M., Liu, S., and Chen, P.-Y. A theoretical understanding of shallow vision transformers: Learning, generalization, and sample complexity. In *The Eleventh International Conference on Learning Representations*, 2023a. URL <https://openreview.net/forum?id=jClGv3Qjhb>.
- Li, H., Wang, M., Ma, T., Liu, S., ZHANG, Z., and Chen, P.-Y. What improves the generalization of graph transformer? a theoretical dive into self-attention and positional encoding. In *NeurIPS 2023 Workshop: New Frontiers in Graph Learning*, 2023b. URL <https://openreview.net/forum?id=BaxFC3z9R6>.
- Li, H., Zhang, S., Zhang, Y., Wang, M., Liu, S., and Chen, P.-Y. How does promoting the minority fraction affect generalization? a theoretical study of one-hidden-layer neural network on group imbalance. *IEEE Journal of Selected Topics in Signal Processing*, 2024.
- Li, S. Concise formulas for the area and volume of a hyperspherical cap. *Asian Journal of Mathematics & Statistics*, 4(1):66–70, 2010.
- Li, Y., Ildiz, M. E., Papailiopoulos, D., and Oymak, S. Transformers as algorithms: Generalization and stability in in-context learning. In *International Conference on Machine Learning*, 2023c.
- Li, Y., Li, Y., and Risteski, A. How do transformers learn topic structure: Towards a mechanistic understanding. *arXiv preprint arXiv:2303.04245*, 2023d.
- Liu, A., Swayamdipta, S., Smith, N. A., and Choi, Y. Wanli: Worker and ai collaboration for natural language inference dataset creation. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 6826–6847, 2022a.
- Liu, J., Shen, D., Zhang, Y., Dolan, W. B., Carin, L., and Chen, W. What makes good in-context examples for gpt-3? In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pp. 100–114, 2022b.
- Liu, Z., Wang, J., Dao, T., Zhou, T., Yuan, B., Song, Z., Shrivastava, A., Zhang, C., Tian, Y., Re, C., et al. Dejavu: Contextual sparsity for efficient llms at inference time. In *International Conference on Machine Learning*, pp. 22137–22176. PMLR, 2023.
- Lucy, L. and Bamman, D. Gender and representation bias in gpt-3 generated stories. In *Proceedings of the Third Workshop on Narrative Understanding*, pp. 48–55, 2021.

- Luo, J.-H., Wu, J., and Lin, W. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pp. 5058–5066, 2017.
- Luo, Y., Li, H., Shi, L., and Wu, X.-M. Enhancing graph transformers with hierarchical distance structural encoding, 2024.
- Ma, X., Fang, G., and Wang, X. LLM-pruner: On the structural pruning of large language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=J8Ajf9WfXP>.
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. *Foundations of machine learning*. MIT press, 2018.
- Molchanov, P., Tyree, S., Karras, T., Aila, T., and Kautz, J. Pruning convolutional neural networks for resource efficient inference. In *International Conference on Learning Representations*, 2016.
- OpenAI. Gpt-4 technical report. *OpenAI*, 2023.
- Oymak, S., Rawat, A. S., Soltanolkotabi, M., and Thrampoulidis, C. On the role of attention in prompt-tuning. *arXiv preprint arXiv:2306.03435*, 2023.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., and Sutskever, I. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- Rubin, O., Herzig, J., and Berant, J. Learning to retrieve prompts for in-context learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2655–2671, 2022.
- Shi, Z., Wei, J., and Liang, Y. A theoretical analysis on feature learning in neural networks: Emergence from inputs and advantage over fixed features. In *International Conference on Learning Representations*, 2021.
- Sun, J., Li, H., and Wang, M. How do skip connections affect graph convolutional networks with graph sampling? a theoretical analysis on generalization, 2024. URL <https://openreview.net/forum?id=J2pMoN2pon>.
- Sun, M., Liu, Z., Bair, A., and Kolter, J. Z. A simple and effective pruning approach for large language models. *arXiv preprint arXiv:2306.11695*, 2023.
- Tarzanagh, D. A., Li, Y., Thrampoulidis, C., and Oymak, S. Transformers as support vector machines. *arXiv preprint arXiv:2308.16898*, 2023a.
- Tarzanagh, D. A., Li, Y., Zhang, X., and Oymak, S. Max-margin token selection in attention mechanism. *CoRR*, 2023b.
- Tian, Y., Wang, Y., Chen, B., and Du, S. Scan and snap: Understanding training dynamics and token composition in 1-layer transformer. *arXiv preprint arXiv:2305.16380*, 2023a.
- Tian, Y., Wang, Y., Zhang, Z., Chen, B., and Du, S. S. Joma: Demystifying multilayer transformers via joint dynamics of mlp and attention. In *Conference on Parsimony and Learning (Recent Spotlight Track)*, 2023b.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Von Oswald, J., Niklasson, E., Randazzo, E., Sacramento, J., Mordvintsev, A., Zhmoginov, A., and Vladymyrov, M. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pp. 35151–35174. PMLR, 2023.
- Wen, W., Wu, C., Wang, Y., Chen, Y., and Li, H. Learning structured sparsity in deep neural networks. *Advances in neural information processing systems*, 29, 2016.
- Wu, J., Zou, D., Chen, Z., Braverman, V., Gu, Q., and Bartlett, P. L. How many pretraining tasks are needed for in-context learning of linear regression? *arXiv preprint arXiv:2310.08391*, 2023a.
- Wu, Z., Wang, Y., Ye, J., and Kong, L. Self-adaptive in-context learning: An information compression perspective for in-context example selection and ordering. *ACL*, 2023b.

- Yang, H. and Wang, Z. On the neural tangent kernel analysis of randomly pruned neural networks. In *International Conference on Artificial Intelligence and Statistics*, pp. 1513–1553. PMLR, 2023.
- Yang, H., Liang, Y., Guo, X., Wu, L., and Wang, Z. Theoretical characterization of how neural network pruning affects its generalization. *arXiv preprint arXiv:2301.00335*, 2023.
- Yang, K., Ackermann, J., He, Z., Feng, G., Zhang, B., Feng, Y., Ye, Q., He, D., and Wang, L. Do efficient transformers really save computation? *arXiv preprint arXiv:2402.13934*, 2024.
- Zhang, R., Frei, S., and Bartlett, P. L. Trained transformers learn linear models in-context. *arXiv preprint arXiv:2306.09927*, 2023a.
- Zhang, S., Wang, M., Liu, S., Chen, P.-Y., and Xiong, J. Why lottery ticket wins? a theoretical perspective of sample complexity on sparse neural networks. *Advances in Neural Information Processing Systems*, 34, 2021.
- Zhang, S., Li, H., Wang, M., Liu, M., Chen, P.-Y., Lu, S., Liu, S., Murugesan, K., and Chaudhury, S. On the convergence and sample complexity analysis of deep q-networks with ϵ -greedy exploration. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b.
- Zhang, S., Wang, M., Chen, P.-Y., Liu, S., Lu, S., and Liu, M. Joint edge-model sparse learning is provably efficient for graph neural networks. In *The Eleventh International Conference on Learning Representations*, 2023c.
- Zhang, Y., Li, H., Yao, Y., Chen, A., Zhang, S., Chen, P.-Y., Wang, M., and Liu, S. Visual prompting reimaged: The power of activation prompts, 2024. URL <https://openreview.net/forum?id=0b328CMwn1>.
- Zhong, K., Song, Z., Jain, P., Bartlett, P. L., and Dhillon, I. S. Recovery guarantees for one-hidden-layer neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 4140–4149, 2017. URL <https://arxiv.org/pdf/1706.03175.pdf>.
- Zou, D. and Gu, Q. An improved analysis of training over-parameterized deep neural networks. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, 2019.

A. Proof Sketch

We partially include the proof backbone in Section 4 when introducing the mechanism of the trained Transformer. We elaborate more about our proof intuition of Theorem 3.3 in the following.

We first briefly introduce our proof intuition of Theorem 3.3. We first respectively build Lemmas D.5, D.6, and D.7 to characterize gradient updates for \mathbf{W}_Q and \mathbf{W}_K , \mathbf{W}_V and \mathbf{W}_O . These Lemmas are based on an observation that a constant fraction of neurons in \mathbf{W}_O can always be activated (Lemma D.9, D.10) to avoid the non-smoothness of Relu activation. The orthogonality of patterns and Definition 3.1 enable the self-attention layer to learn in-domain-relevant (IDR) patterns rather than in-domain-irrelevant (IDI) patterns and select contexts with the same IDR pattern as the query. Then, to develop Theorem 3.3, we use Lemma D.5 to show that the attention weights converge to be close to 1 when $\eta T \geq \Omega(M_1 \sqrt{\log M_1})$. Next, we compute the network output according to the label embedding using Lemma D.6 and D.7. Finally, we derive the required number of iterations to make the generalization error $\mathcal{O}(\epsilon)$ by concentration inequalities.

We then would like to specify how we handle the Relu activation in the gradient of \mathbf{W}_K , \mathbf{W}_Q . We also want to clarify that how we can handle the training dynamics related to the softmax is different from (Huang et al., 2023) although we both derive a sparse attention distribution. Inspired by the intuition of feature-learning analyses for two-layer Relu networks (Brutzkus & Globerson, 2021; Shi et al., 2021; Zhang et al., 2023c; Allen-Zhu & Li, 2023), we initialize the model such that at least a constant fraction of the neurons of \mathbf{W}_O are activated (Lemma D.9), which are called lucky neurons as in (Zhang et al., 2023c; Li et al., 2023a). We prove that these lucky neurons are always activated (Lemma D.10) and grow with an increasing magnitude and two fixed directions of the label embedding along the training (Lemma D.7). Then, we can show that the gradient growths of \mathbf{W}_Q and \mathbf{W}_K can be lower bounded by contributions from these lucky neurons. Therefore, we are able to characterize the gradient updates of \mathbf{W}_K and \mathbf{W}_Q given a dynamic \mathbf{W}_O . This process is different from (Huang et al., 2023) since (Huang et al., 2023) does not include Relu MLP, so there is no need to study lucky neurons. Besides, we use Hinge loss, while their training loss is logistic loss, which leads to more training phases, as a difference in the training dynamics between us.

B. Addition Discussions and Extensions

B.1. The Motivation to Study NONLINEAR Transformers

The reasons we study nonlinear Transformers in this work are as follows. First, nonlinear Transformers for ICL, which are different from linear Transformers, are common in practice but less explored in theory. Nonlinear attention and nonlinear MLP are default components of standard Transformers (Vaswani et al., 2017) and are widely applied in large language models for implementing ICL in practice. Existing works show that nonlinear Transformers exhibit their empirical advantages when learning nonlinear functions (Cheng et al., 2023) or conducting dynamic programming tasks (Yang et al., 2024). However, state-of-the-art theoretical works (Zhang et al., 2023a; Huang et al., 2023; Wu et al., 2023a) ignore the nonlinearities (partially) to simplify the analysis or the presentation. Second, the analysis of nonlinear Transformers is quite different from that of Transformers without nonlinearities. For example, softmax attention has a different derivative from linear attention, which includes nonlinear exponential terms and needs a more complicated computation of the gradient updates. Relu MLP provides several non-differential points, which makes the loss landscape more challenging to analyze.

B.2. The Discussion on Single/Multi-Head Attention

There are several reasons why we only study single-head attention in the main body of the paper. First, all the previous theoretical works studying the optimization and generalization of Transformers on ICL (Zhang et al., 2023a; Huang et al., 2023; Wu et al., 2023a) only consider single-head attention in the network. Some concurrent works consider multi-head attention, but they either do not study ICL (Deora et al., 2023; Chen & Li, 2024) or do not involve convergence/generalization guarantee (Cui et al., 2024). Hence, the question of how the ICL ability on unseen tasks and out-of-domain data is obtained by training is still unexplored. Our theoretical analysis studies the convergence and generalization of ICL using Transformers with softmax attention and Relu MLP, involving generalization on unseen tasks and OOD data. Second, our empirical experiments on GPT-2 in Figure 3 are conducted with two heads to verify our theoretical findings, which means some theoretical conclusions hold in Transformers with multiple heads.

However, our analysis for single-head attention can be extended to multi-head attention to some degree. Consider a

multi-head attention layer where the layer output is a concatenation of the output of each head, i.e.,

$$\left\| \sum_{h=1}^H \mathbf{W}_{V_h} \mathbf{p}_i \cdot \text{softmax}(\mathbf{p}_i^\top \mathbf{W}_{K_h}^\top \mathbf{W}_{Q_h} \mathbf{p}_{query}) \right\|. \quad (23)$$

The overall conclusion will remain the same, given the same data formulation and initialization on each head because of the orthogonality of patterns. Specifically, we can still show that each attention head selects contexts with the same in-domain-relevant (IDR) pattern as the query. The MLP layer will still make predictions based on the label embedding, as suggested in Section 4.2. We will leave the analysis on multi-head attention with more general settings as future directions.

B.3. Extension to Multiple Patterns for One Class

We can extend our analysis to the case that several orthogonal IDR patterns correspond to the label +1, while some other orthogonal IDR patterns correspond to the label -1. Then, as long as there is always a context input that shares the same IDR/ODR pattern as the query, we can still prove that the self-attention layer selects contexts with the same IDR/ODR pattern as the query. Furthermore, we can show the MLP layer makes predictions based on the label embedding. Therefore, the mechanism remains the same as the current setting in the manuscript, where one pattern corresponds to one pattern. We will leave other cases where the data formulation is different in future works.

The reason why we use our current setting in the main body of the paper is to simplify the presentation while emphasizing our major contributions of analyzing optimization and generalization of nonlinear Transformers both in-domain and out-of-domain. As the first work on this problem, as far as we know, we believe our data formulation keeps the necessary complexity.

B.4. Additional Related Works

We introduce other existing theoretical works on learning and generalization of neural networks in this section. Some works (Zhong et al., 2017; Fu et al., 2020; Li et al., 2022b; Zhang et al., 2023b; Li et al., 2024) study the generalization performance following the model recovery framework by probing the local convexity around a ground truth parameter. The neural-tangent-kernel (NTK) analysis (Jacot et al., 2018; Allen-Zhu et al., 2019a;b; Cao & Gu, 2019; Zou & Gu, 2019; Chen et al., 2020; Li et al., 2022a; Sun et al., 2024) considers strongly overparameterized networks to linearize the neural network around the initialization. The generalization performance is independent of the feature distribution. (Daniely & Malach, 2020; Shi et al., 2021; Karp et al., 2021; Brutzkus & Globerson, 2021; Zhang et al., 2023c; Li et al., 2023a; Zhang et al., 2024; Chowdhury et al., 2023; 2024) investigate the generalization of neural networks assuming a data model consisting of discriminative patterns and background patterns. Our analysis belongs to the last line of research.

C. Additional Experiments and the Algorithm

We first present the training algorithm introduced in Section 2.3.

Algorithm 1 Training with Stochastic Gradient Descent (SGD)

- 1: **Hyperparameters:** The step size η , the number of iterations T , batch size B .
- 2: **Initialization:** Each entry of $\mathbf{W}_O^{(0)}$ and $\mathbf{a}^{(0)}$ from $\mathcal{N}(0, \xi^2)$ and $\text{Uniform}(\{+1/\sqrt{m}, -1/\sqrt{m}\})$, respectively. \mathbf{W}_Q , \mathbf{W}_K and \mathbf{W}_V are initialized such that all diagonal entries of $\mathbf{W}_V^{(0)}$, and the first $d_{\mathcal{X}}$ diagonal entries of $\mathbf{W}_Q^{(0)}$ and $\mathbf{W}_K^{(0)}$ are set as δ with $\delta \in (0, 0.2]$.
- 3: **Training by SGD:** For each iteration, we independently sample $\mathbf{x}_{query} \sim \mathcal{D}$, $f \in \mathcal{T}_{tr}$ to form a batch of training prompt and labels $\{\mathbf{P}^n, z^n\}_{n \in \mathcal{B}_t}$ as introduced in Section 3.2. Each IDR pattern is sampled equally likely in each batch. For each $t = 0, 1, \dots, T-1$ and $\mathbf{W}^{(t)} \in \Psi^{(t)}$

$$\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} - \eta \cdot \frac{1}{B} \sum_{n \in \mathcal{B}_t} \nabla_{\mathbf{W}^{(t)}} \ell(\Psi^{(t)}; \mathbf{P}^n, z^n). \quad (24)$$

- 4: **Output:** $\mathbf{W}_O^{(T)}, \mathbf{W}_V^{(T)}, \mathbf{W}_K^{(T)}, \mathbf{W}_Q^{(T)}$.
-

Then, we introduce additional experiments to verify our theory.

C.1. The impact of α

We choose $\alpha = 0.6$ and use a one-layer Transformer as in (2). Figure 6 shows that the required length of the training prompt is linear in α^{-1} , while the required number of training iterations is linear in $\alpha^{-2/3}$, which verify the theoretical findings in (9) and (10).




Figure 6. The prompt length against α , and the required number of training iterations against α .

C.2. The required number of training tasks

We choose $\alpha = 0.6$ and use a one-layer Transformer as in (2). For a given \mathcal{T} , we first generate a set of tasks that satisfies Condition 3.2 as follows. Define $\mathbf{a}_i = \mathbf{a}_{i+M_1} = \boldsymbol{\mu}_i$ for $i \in [M_1]$, and then the j -th task function map the queries with \mathbf{a}_j and \mathbf{a}_{j+1} as IDR patterns to $+1$ and -1 , respectively, for $j \in [M_1]$. Then, we get a task set \mathcal{T}_{tr0} with $|\mathcal{T}_{tr0}| = M_1$. Then, we vary the number of training tasks in the way that (1) we sample within \mathcal{T}_{tr0} to get a set \mathcal{T}_{tr} with $|\mathcal{T}_{tr}| \leq M_1$ (2) we sample within $\mathcal{T} \setminus \mathcal{T}_{tr0}$ to get a set \mathcal{T}'_{tr} , and $\mathcal{T}_{tr} = \mathcal{T}'_{tr} \cup \mathcal{T}_{tr0}$ such that $|\mathcal{T}_{tr}| \geq M_1$. Figure 7 shows that for any M_1 , the generalization error is significant as long as $|\mathcal{T}_{tr}| < M_1$, while the generalization error reaches around 0 as long as $|\mathcal{T}_{tr}| \geq M_1$ and \mathcal{T} covers all the possibilities of IDR patterns and labels. This verifies that Condition 3.2 can be met with a fraction of $(M_1 - 1)^{-1/2}$ of total number of in-domain tasks.




Figure 7. The required number of training tasks for different M_1 .

D. Proofs of the Main Theorems

We first provide several useful definitions and key lemmas for the proof of the main theorems. Table 2 shows a summary of notations used in the proof.

D.1. Proof Overview of Main Theorems

This section illustrates how Corollary 4.3 and Proposition 4.5 contribute to the final in- and out-of-domain generalization performance of ICL.

The establishment of generalization

1. (*Self-Attention*) We can deduce from Corollary 4.3 that, for a query with IDR pattern $\boldsymbol{\mu}_j$ ($j \in [M_1]$) and label $+1$, the weighted summation of contexts and the query by the attention score, i.e., $\sum_{i=1}^l \mathbf{p}_i \text{attn}(\Psi; \mathbf{P}, i)$, is close to $[\boldsymbol{\mu}_j^\top, \mathbf{q}^\top]^\top$. This is because as long as the training/testing prompt length satisfies (9), large attention weights are assigned on \mathbf{p}_i of which

Table 2. Summary of Notations

Notations	Annotation
$\mathbf{x}_s^n, \mathbf{y}_s^n$	\mathbf{x}_s^n is the data for classification. \mathbf{y}_s^n is the embedding of the label for \mathbf{x}_s^n .
\mathbf{P}^n, z^n	\mathbf{P}^n is a prompt that consists of l pairs of \mathbf{x}_s^n and \mathbf{y}_s^n , $s \in [l]$. The last column of \mathbf{P}^n contains \mathbf{p}_{query}^n , which is the query of \mathbf{P}^n . $z^n \in \{+1, -1\}$ is the binary label of \mathbf{p}_{query}^n , which is also the label of \mathbf{P}^n when we formulate the problem as a supervised learning problem.
$F(\Psi; \mathbf{P}^n), \ell(\Psi; \mathbf{P}^n, z^n)$	$F(\Psi; \mathbf{P}^n)$ is the Transformer output for \mathbf{P}^n with Ψ as the parameter. $\ell(\Psi; \mathbf{P}^n, z^n)$ is the loss function value given \mathbf{P}^n and the corresponding label z^n .
$\mathbf{p}_s^n, \boldsymbol{\mu}_j, \boldsymbol{\nu}_k$	\mathbf{p}_s^n is the s -th example with the corresponding label in \mathbf{P}^n . If $s = query$, \mathbf{p}_s^n is the query. $\boldsymbol{\mu}_j$ and $\boldsymbol{\nu}_k$ are the IDR and IDI patterns in the feature embedding of \mathbf{p}_s^n as the corresponding coefficients, respectively.
\mathbf{q}	\mathbf{q} is the label space embedding.
M_1, M_2, M	M_1 is the number of IDR patterns. M_2 is the number of IDI patterns. $M = M_1 + M_2$.
α, a	α is the probability of selecting examples that contain either of the two decisive IDR patterns in each \mathbf{P}^n . $a = 1/ a_i $ where a_i is the entry of each neuron in \mathbf{W}_O . $a = m$.
$\kappa, \kappa', K, K', \beta$	κ and κ' are the coefficients of the IDI pattern and the ODI pattern in the input \mathbf{x} , respectively. κ and κ' follow uniform distribution $U(-K, K)$ and $U(-K', K')$ with $K \leq 1/2$ and $K' \leq \mathcal{O}(1)$, respectively. β is the norm of in-/out-of-domain-(ir)relevant (IDR/ODR/IDI/ODI) patterns.
$\mathcal{W}_n, \mathcal{U}_n$	The sets of lucky neurons. \mathcal{W}_n is the set of neurons of \mathbf{W}_O that can activate the terms inside $\text{Relu}(\cdot)$ in $F(\Psi; \mathbf{P}^n)$ for $z^n = +1$ at initialization. \mathcal{U}_n is the set of neurons of \mathbf{W}_O that can activate the Relu part of $F(\Psi; \mathbf{P}^n)$ for $z^n = -1$ at initialization.
\mathcal{W}, \mathcal{U}	$\mathcal{W} = \cup_{n \in [N]} \mathcal{W}_n$. $\mathcal{U} = \cup_{n \in [N]} \mathcal{U}_n$
\mathcal{N}_j^n	The set of examples in \mathbf{P}^n that contains $\boldsymbol{\mu}_j$ as the IDR pattern.
γ_t	γ_t is the summation of attention weight on examples that have different IDR patterns from the query.
ζ_t	ζ_t is smallest positive value inside the $\text{Relu}(\cdot)$ in $F(\Psi; \mathbf{P}^n)$ for all the \mathbf{W}_O neuron and all $n \in [N]$.
\mathcal{B}_b	\mathcal{B}_b is the SGD batch at the b -th iteration.
l_{tr}	l_{tr} is the prompt length of the training data.
l_{ts}	l_{ts} is the prompt length of the testing data.
$\mathcal{O}(), \Omega(), \Theta()$	We follow the convention that $f(x) = \mathcal{O}(g(x))$ (or $\Omega(g(x)), \Theta(g(x))$) means that $f(x)$ increases at most, at least, or in the order of $g(x)$, respectively.
\gtrsim, \lesssim	$f(x) \gtrsim g(x)$ (or $f(x) \lesssim g(x)$) means that $f(x) \geq \Omega(g(x))$ (or $f(x) \lesssim \mathcal{O}(g(x))$).

the IDR pattern is $\boldsymbol{\mu}_j$, and the label embedding is \mathbf{q} by (20). Similarly, if its label is -1 , the weighted summation of contexts and the query outputs $[\boldsymbol{\mu}_j^\top, -\mathbf{q}^\top]^\top$.

2. (*MLP*) By Proposition 4.5, we know that a large enough proportion of positive (or negative) neurons $i \in [m]$ have the label embedding of $\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}$ close to $\pm \mathbf{q}$ (22). They can thus map the weighted summation of contexts and the query by attention with $+\mathbf{q}$ (or $-\mathbf{q}$) to positive (or negative) values. This leads to a correct prediction in-domain (Theorem 3.3).

3. (*Out-of-Domain Generalization*) Since Corollary 4.3 also applies to ODR patterns, then for a query with an ODR pattern $\boldsymbol{\mu}'_j, j \in [M'_1]$, the resulting weighted summation of contexts and the query is close to $[\boldsymbol{\mu}'_j^\top, \mathbf{q}^\top]^\top$ or $[\boldsymbol{\mu}'_j^\top, -\mathbf{q}^\top]^\top$. Then, by combining (21), (22) and the condition on ODR pattern characterized in (12), we can ensure that the MLP layer produces a desired prediction out of the domain (Theorem 3.4).

D.2. Preliminaries

Lemma D.1. (*Multiplicative Chernoff bounds, Theorem D.4 of (Mohri et al., 2018)*) Let X_1, \dots, X_m be independent random variables drawn according to some distribution \mathcal{D} with mean p and support included in $[0, 1]$. Then, for any

$\gamma \in [0, \frac{1}{p} - 1]$, the following inequality holds for $\hat{p} = \frac{1}{m} \sum_{i=1}^m X_i$:

$$\Pr(\hat{p} \geq (1 + \gamma)p) \leq e^{-\frac{m\gamma^2}{3}}, \quad (25)$$

$$\Pr(\hat{p} \leq (1 - \gamma)p) \leq e^{-\frac{m\gamma^2}{2}}. \quad (26)$$

Definition D.2. (Vershynin, 2010) We say X is a sub-Gaussian random variable with sub-Gaussian norm $K > 0$, if $(\mathbb{E}|X|^p)^{\frac{1}{p}} \leq K\sqrt{p}$ for all $p \geq 1$. In addition, the sub-Gaussian norm of X , denoted $\|X\|_{\psi_2}$, is defined as $\|X\|_{\psi_2} = \sup_{p \geq 1} p^{-\frac{1}{2}} (\mathbb{E}|X|^p)^{\frac{1}{p}}$.

Lemma D.3. ((Vershynin, 2010) Proposition 5.1, Hoeffding's inequality) Let X_1, X_2, \dots, X_N be independent centered sub-gaussian random variables, and let $K = \max_i \|\mathbf{X}_i\|_{\psi_2}$. Then for every $\mathbf{a} = (a_1, \dots, a_N) \in \mathbb{R}^N$ and every $t \geq 0$, we have

$$\Pr\left(\left|\sum_{i=1}^N a_i X_i\right| \geq t\right) \leq e \cdot \exp\left(-\frac{ct^2}{K^2 \|\mathbf{a}\|^2}\right), \quad (27)$$

where $c > 0$ is an absolute constant.

Definition D.4. For any data index n and iteration t , we can find i such that $\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)} \top \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) > 0$ by the initialization with high probability. Define

1. $\zeta_{i,n,t} := \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)} \top \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n)$.
2. $\zeta_{i,t} = \min_n \{\zeta_{i,n,t}\}$.
3. $\zeta_t = \min_i \{\zeta_{i,t}\}$.
4. $\gamma_{t,n} = 1 - \sum_{s \in \mathcal{N}_*^n} \text{softmax}((\mathbf{W}_K^{(t)} \mathbf{p}_s^n) \top (\mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n))$.
5. $\gamma_t = \max_{n \in [N]} \{\gamma_{t,n}\}$.

Lemma D.5. (gradient updates of \mathbf{W}_Q and \mathbf{W}_K) By the SGD training method described in Section 2.3, we have the following equations. Given the definition of in-/out-of-domain data as in (1) and the in-/out-of-domain data distribution \mathcal{D} in (6) and \mathcal{D}' in (8), we study the gradient updates in the directions of queries or contexts. Note that we require $m \gtrsim M_1^2$, $B \gtrsim M_1 \log M_1$, $l = l_{tr} \gtrsim 1$, $\beta \in [1, O(1)]$.

We first consider the case when the feature embeddings of the query \mathbf{x}_{query} and the example $\mathbf{x}_q, q \in [l]$ are $\boldsymbol{\mu}_j$. The label embedding is $\mathbf{0}$ for the query and $\pm \mathbf{q}$ for non-query examples. Then, for any $l, a \in [M_1], k \in [M_2], t_0 \geq 1$, where $\boldsymbol{\mu}_l$ forms a task in \mathcal{T}_{tr} with $\boldsymbol{\mu}_j$ and $\boldsymbol{\mu}_a$ does not,

$$(\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \sum_{b=0}^{t_0} \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \gtrsim \eta \frac{1}{M_1} \sum_{b=0}^{t_0} \zeta_b \delta \gamma_b \beta^4, \quad (28)$$

$$\begin{aligned} & \left| (\boldsymbol{\mu}_l^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{t_0-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \lesssim e^{-\Theta\left(\frac{\eta t_0}{M_1}\right)^2} \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \sum_{b=0}^{t_0-1} \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|, \end{aligned} \quad (29)$$

$$\begin{aligned} & \left| (\boldsymbol{\mu}_a^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{t_0-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \lesssim \frac{1}{M_1} \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \sum_{b=0}^{t_0-1} \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|, \end{aligned} \quad (30)$$

$$\left| (\boldsymbol{\nu}_k^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \lesssim \frac{1}{M_2} \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|, \quad (31)$$

$$(\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=t_0+1} \mathbf{p}_q \gtrsim \eta \frac{1}{M_1} \sum_{b=0}^{t_0} \zeta_b \delta \gamma_b \beta^4, \quad (32)$$

$$\begin{aligned} & \left| (\boldsymbol{\mu}_l^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=t_0} \mathbf{p}_q \right| \\ & \lesssim e^{-\Theta\left(\frac{\eta t_0}{M_1}\right)^2} \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=t_0} \mathbf{p}_q \right|, \end{aligned} \quad (33)$$

$$\left| (\boldsymbol{\mu}_a^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=t_0} \mathbf{p}_q \right| \lesssim \frac{1}{M_1} \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=t_0} \mathbf{p}_q \right|, \quad (34)$$

$$\left| (\boldsymbol{\nu}_k^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=t_0} \mathbf{p}_q \right| \lesssim \frac{1}{M_2} \left| (\boldsymbol{\mu}_j^\top, \mathbf{q}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=t_0} \mathbf{p}_q \right|. \quad (35)$$

In the above, equations (28), (29), (30), and (31) characterize the directions of gradient updates of \mathbf{W}_Q when projected with $(\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top$. Similarly, equations (32), (33), (34), and (35) characterize the directions of gradient updates of \mathbf{W}_K when projected with \mathbf{p}_q , $q \in [l]$.

Lemma D.6. (gradient updates of \mathbf{W}_V) For \mathbf{p}_j^n defined in (1) and $t_0 \geq 1$, if $l = l_{tr} \gtrsim \max\{1, \frac{1}{\alpha \beta^2}\}$ and $BT \gtrsim \Theta(M_1^2)$, $B \gtrsim M_1$, we have that for \mathbf{p}_j of which the corresponding label embedding is \mathbf{q} ,

$$\begin{aligned} & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \sum_{b=0}^{t_0} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V^{(b)}} \mathbf{p}_j \\ & = \eta \sum_{b=0}^{t_0} \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right), \end{aligned} \quad (36)$$

where

$$-V_i(b) \gtrsim \beta^2(1 - \gamma_t)/a, \quad i \in \mathcal{W}_n, \quad (37)$$

$$-V_i(b) \leq \frac{1}{\beta^2 + 1} V_j(b), \quad i \in \mathcal{U}_n, j \in \mathcal{W}_n, \quad (38)$$

$$|V_i(b)| \lesssim \sqrt{\frac{\log B}{B}} \cdot \frac{1}{a}, \quad i \notin \mathcal{W}_n \cup \mathcal{U}_n. \quad (39)$$

If the corresponding label embedding is $-\mathbf{q}$, we have that (36) holds with

$$-V_i(b) \gtrsim \beta^2(1 - \gamma_t)/a, \quad i \in \mathcal{U}_n, \quad (40)$$

$$-V_i(b) \leq \frac{1}{\beta^2 + 1} V_j(b), \quad i \in \mathcal{W}_n, j \in \mathcal{U}_n, \quad (41)$$

$$|V_i(b)| \lesssim \sqrt{\frac{\log B}{B}} \cdot \frac{1}{a}, \quad i \notin \mathcal{W}_n \cup \mathcal{U}_n. \quad (42)$$

We can also derive

$$\begin{aligned} & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V^{(t)}} (\boldsymbol{\nu}_k^\top, \mathbf{0}^\top)^\top \\ & = \eta \sum_{b=0}^{t_0} \left(\sum_{i \in \mathcal{W}_n} V_i'(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i'(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i'(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right), \end{aligned} \quad (43)$$

where

$$|V'_i(b)| \leq |V_i(b)| \cdot \frac{1}{M_2}. \quad (44)$$

Lemma D.7. (gradient updates of \mathbf{W}_O) We are given $\Theta(1) \geq \beta \geq 1$ and $m \gtrsim M_1^2$, $BT \gtrsim M_1 \log M_1$, $B \gtrsim M_1$, $t = t_0 \geq \Theta(1)$. Denote the set of examples that share the same IDR pattern as \mathbf{p}_{query}^n as \mathcal{B}_b^n in the b -th iteration. For $i \in \mathcal{W}$, $b \neq a$, and \mathbf{p}_{query}^n corresponding to \mathbf{q} and $\boldsymbol{\mu}_a$,

$$\eta \frac{1}{|\mathcal{B}_b^n|} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(t_0)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top = \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2}\right)^{t_0}, \quad (45)$$

$$\eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(t_0)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta t_0}{2a}. \quad (46)$$

For $i \in \mathcal{U}$, $b \neq a$, and \mathbf{p}_{query}^n corresponding to \mathbf{q} and $\boldsymbol{\mu}_a$,

$$\eta \frac{1}{|\mathcal{B}_b^n|} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(t_0)}} (\boldsymbol{\mu}_a^\top, -\mathbf{q}^\top)^\top = \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2}\right)^{t_0}, \quad (47)$$

$$\begin{aligned} & \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(b)}} (\boldsymbol{\mu}_a^\top, -\mathbf{q}^\top)^\top \\ & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta t_0}{2a}. \end{aligned} \quad (48)$$

For $i \in \mathcal{W} \cup \mathcal{U}$ and $c \in [M_2]$,

$$\|\mathbf{W}_{O(i, \cdot)}^{(t_0)}\| \gtrsim \sqrt{M_1} \delta(\beta^2 + 1)^{\frac{1}{2}} \frac{\alpha \eta t_0}{2a}, \quad (49)$$

$$\eta \frac{1}{B} \sum_{b=0}^{t_0} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(b)}} (\boldsymbol{\nu}_c^\top, \pm \mathbf{q}^\top)^\top \leq \frac{1}{M_2} \eta \frac{1}{B} \sum_{b=0}^{t_0} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(b)}} (\boldsymbol{\mu}_b^\top, \mathbf{q}^\top)^\top. \quad (50)$$

For $i \notin \mathcal{W} \cup \mathcal{U}$, we have

$$\eta \frac{1}{B} \sum_{b=0}^{t_0} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(b)}} (\boldsymbol{\mu}_a^\top, \pm \mathbf{q}^\top)^\top \leq \eta t \sqrt{\frac{\log B}{B}} \frac{1}{a}. \quad (51)$$

Definition D.8. Define

$$\mathbf{V}^n(t) := \sum_{s=1}^{l+1} \mathbf{W}_V^{(t)} \mathbf{p}_s^n \text{softmax}(\mathbf{p}_s^n^\top \mathbf{W}_K^{(t)})^\top \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n, \quad (52)$$

for \mathbf{P}^n . Let $\mathbf{W}_{O(i, \cdot)} = (\mathbf{O}_{i,1}, \mathbf{O}_{i,2}, \mathbf{0}^\top)$ where $\mathbf{O}_{i,1} \in \mathbb{R}^{d_x}$, $\mathbf{O}_{i,2} \in \mathbb{R}^{d_y}$. Let $\mathbf{V}^n(t) = (\mathbf{V}_{n,1}(t)^\top, \mathbf{V}_{n,2}(t)^\top, \mathbf{0}^\top)^\top$ where $\mathbf{V}_{i,1}(t) \in \mathbb{R}^{d_x}$, $\mathbf{V}_{i,2}(t) \in \mathbb{R}^{d_y}$. Define $\mathcal{W}_n, \mathcal{U}_n$ as the sets of lucky neurons such that

$$\mathcal{W}_n = \{i : \mathbf{O}_{i,1}^{(0)} \mathbf{V}_{n,1}(0) > 0, \mathbf{O}_{i,2}^{(0)} \mathbf{V}_{n,2}(0) > 0, a_i > 0\}, \quad (53)$$

$$\mathcal{U}_n = \{i : \mathbf{O}_{i,1}^{(0)} \mathbf{V}_{n,1}(0) > 0, \mathbf{O}_{i,2}^{(0)} \mathbf{V}_{n,2}(0) > 0, a_i < 0\}. \quad (54)$$

Define

$$\mathcal{N}_j^{n,i} = \{i : i \in [l+1], \mathbf{x}_i^n = \boldsymbol{\mu}_j + \kappa_i^n \boldsymbol{\nu}_k + \mathbf{n}_i^n, k \in [M_2]\}, \quad (55)$$

$$\mathcal{M}_k^{n,i} = \{i : i \in [l+1], \mathbf{x}_i^n = \boldsymbol{\mu}_j + \kappa_i^n \boldsymbol{\nu}_k + \mathbf{n}_i^n, j \in [M_1]\}, \quad (56)$$

as the set of example inputs with $\boldsymbol{\mu}_j$ as the IDR patterns or with $\boldsymbol{\nu}_k$ as the IDI patterns, respectively.

$$\mathcal{W} = \bigcup_{n=1}^N \mathcal{W}_n, \quad \mathcal{U} = \bigcup_{n=1}^N \mathcal{U}_n. \quad (57)$$

Lemma D.9. By the definition of lucky neurons in (53) and (54), and the initialization described in Section 2.3, the number of lucky neurons $|\mathcal{W}_n|, |\mathcal{U}_n|$ satisfies

$$|\mathcal{W}_n|, |\mathcal{U}_n| \geq \Omega(m). \quad (58)$$

Hence,

$$|\mathcal{W}|, |\mathcal{U}| \geq \Omega(m). \quad (59)$$

Lemma D.10. Under the condition that $m \gtrsim M_1^2 \log M_1$, we have the following results.

1. When $t \geq 0$, for $V^n(t)$ where \mathbf{p}_{query}^n corresponds to the label +1,

$$\mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{V}^n(t)] = 1, i \in \mathcal{W}_n, \quad (60)$$

for $V^n(t)$ where \mathbf{p}_{query}^n corresponds to the label -1,

$$\mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{V}^n(t)] = 1, i \in \mathcal{U}_n. \quad (61)$$

2. When $t \geq \Theta(1)$, for $i \in \mathcal{W}$, we have that for $V^n(t)$ where \mathbf{p}_{query}^n corresponds to the label +1,

$$\mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{V}^n(t)] = 1. \quad (62)$$

For $i \in \mathcal{U}$, we have that for $V^n(t)$ where \mathbf{p}_{query}^n corresponds to the label -1,

$$\mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{V}^n(t)] = 1. \quad (63)$$

Lemma D.11. With in-domain tasks defined in Definition 3.1 and Condition 3.2, the number of training tasks should satisfy $|\mathcal{T}_{tr}| \geq M_1$ to make Condition 3.2 hold.

D.3. Proof of Theorem 3.3

Proof. We first look at the required length of the prompt. Define m_i as the corresponding IDR pattern in the i -th demonstration. Consider the categorical distribution where the probabilities of selecting μ_a and μ_b are $\alpha/2$ respectively. By the Chernoff bound of Bernoulli distribution in Lemma D.1, we can obtain

$$\Pr \left(\frac{1}{l_{tr}} \sum_{i=1}^{l_{tr}} \mathbb{1}[m_i = \mu_a] \leq (1-c) \frac{\alpha}{2} \right) \leq e^{-l_{tr} c^2 \frac{\alpha}{2}} = M_1^{-C}, \quad (64)$$

for some $c \in (0, 1)$ and $C > 0$. Hence, with a high probability, combining the condition $l_{tr} \geq (\alpha\beta^2)^{-1}$ in Lemma D.6,

$$l_{tr} \gtrsim \max \left\{ \Omega \left(\frac{2 \log M_1}{\alpha} \right), \Omega \left(\frac{1}{\alpha\beta^2} \right) \right\}. \quad (65)$$

By the condition in Lemma D.5, we have that

$$B \geq \Omega(M_1 \log M_1). \quad (66)$$

We know that there exists gradient noise caused by imbalanced IDR patterns in each batch. Therefore, by Hoeffding's inequality (27), for any $\mathbf{W} \in \Psi$,

$$\Pr \left(\left\| \frac{1}{|\mathcal{B}_b|} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\Psi; \mathbf{P}^n, z^n)}{\partial \mathbf{W}} - \mathbb{E} \left[\frac{\partial \ell(\Psi; \mathbf{P}^n, z^n)}{\partial \mathbf{W}} \right] \right\| \geq \left| \mathbb{E} \left[\frac{\partial \ell(\Psi; \mathbf{P}^n, z^n)}{\partial \mathbf{W}} \right] \right| \epsilon \right) \leq e^{-B\epsilon^2} \leq M_1^{-C}, \quad (67)$$

if $B \gtrsim \epsilon^{-2} \log M_1$. Therefore, we require

$$B \gtrsim \max\{\epsilon^{-2}, M_1\} \log M_1. \quad (68)$$

(a) We have that for i such that $a_i > 0$ but $i \notin \mathcal{W}$ by the definition of the Relu function,

$$a_i \text{Relu}(\mathbf{W}_{O(i,\cdot)}^{(T)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(T)} \mathbf{p}_s^n) \text{softmax}((\mathbf{W}_K^{(T)} \mathbf{p}_s^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n))) \geq 0. \quad (69)$$

(b) Furthermore, we have that for $i \in \mathcal{W}$, and for \mathbf{p}_s^n that shares the same IDR pattern as \mathbf{p}_{query}^n , with a high probability of $1 - M_1^{-C}$,

$$\begin{aligned} & \eta \sum_{b=0}^{T-1} \mathbf{W}_{O(j,\cdot)}^{(T)} \sum_{j \in \mathcal{W}_n} \mathbf{W}_{O(j,\cdot)}^{(b)\top} \\ & \geq \eta \sum_{b=0}^{T-1} M_1 \delta (\beta^2 + 1)^{\frac{1}{2}} \frac{\alpha \eta T}{2a} \cdot \delta (\beta^2 + 1)^{\frac{1}{2}} \frac{\alpha \eta b}{2a} \\ & \gtrsim \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2}, \end{aligned} \quad (70)$$

where the first step comes from (49) in Lemma D.7, and the second step is by $\sum_{b=0}^{T-1} b = \Theta(T^2)$. Then, we can obtain

$$\begin{aligned} & \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} \mathbf{p}_s^n \\ & = \mathbf{W}_{O(i,\cdot)}^{(T)} (\delta \mathbf{p}_s^n + \sum_{b=0}^{T-1} \eta (\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)})^\top) \\ & \gtrsim \delta^2 (\beta^2 + 1) \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2}, \end{aligned} \quad (71)$$

where the first step is by (36) in Lemma D.6, and the last step comes from Lemma D.7.

Therefore, by combining Lemma D.9 and Lemma D.10, we have that

$$\begin{aligned} & \sum_{i \in \mathcal{W}} a_i \text{Relu}(\mathbf{W}_{O(i,\cdot)}^{(T)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(T)} \mathbf{p}_s^n) \text{softmax}((\mathbf{W}_K^{(T)} \mathbf{p}_s^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n))) \\ & \gtrsim (1 - \gamma_T) \cdot \left(\delta^2 (\beta^2 + 1) \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2} \right), \end{aligned} \quad (72)$$

when γ_T is order-wise smaller than 1. We next give a bound for γ_T , which is give by Definition D.4,

$$\gamma_T \geq 1 - \sum_{s \in \mathcal{N}_s^n} \text{softmax}((\mathbf{W}_K^{(T)} \mathbf{p}_s^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n)), \quad (73)$$

from Defition D.4 for μ_j as the IDR pattern in \mathbf{p}_{query}^n . We can tell from (72) and Definition D.4,

$$\zeta_b \gtrsim \delta^2 (\beta^2 + 1) \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2}. \quad (74)$$

Then, if $\zeta_T \gtrsim 1$ and $T \gtrsim M_1$, by Lemma D.5, with high probability,

$$\begin{aligned} & (\mathbf{W}_K^{(T)} \mathbf{p}_s^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n) \\ & \gtrsim (\eta \frac{1}{M_1} \sum_{b=0}^{T-1} \zeta_b \delta \gamma_b \beta^2 + \delta)^2 \\ & \gtrsim (\eta \sum_{b=0}^{T-1} \gamma_b \beta^2 \cdot (\delta^2 \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2}) + \delta)^2 \\ & := (\eta \sum_{b=0}^{T-1} \gamma_b \cdot \Delta_T + \delta)^2, \end{aligned} \quad (75)$$

where the first step comes from the fact that the gradient update projections of \mathbf{W}_Q and \mathbf{W}_K onto queries or examples are close to the corresponding IDR pattern the most by Lemma D.5. In the last inequality of (75), we only consider the term related to T and γ_b . For any \mathbf{p}_l^n that shares a different IDR pattern as \mathbf{p}_{query}^n , we have

$$(\mathbf{W}_K^{(T)} \mathbf{p}_l^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n) \lesssim \frac{1}{M_1} (\mathbf{W}_K^{(T)} \mathbf{p}_s^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n), \quad (76)$$

by Lemma D.5. Then, given the definition of softmax,

$$\begin{aligned} & \sum_{s \in \mathcal{N}_j^n} \text{softmax}(\mathbf{p}_s^n^\top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ & \geq \frac{\sum_{s \in \mathcal{N}_j^n} e^{\Theta(\delta^2) + (\eta \sum_{b=0}^{T-1} \gamma_b \cdot \Delta_T)^2}}{\sum_{s \in \mathcal{N}_j^n} e^{\Theta(\delta^2) + (\eta \sum_{b=0}^{T-1} \gamma_b \cdot \Delta_T)^2} + \sum_{s \in [l] - \mathcal{N}_j^n} e^{\frac{1}{M_1} (\eta \sum_{b=0}^{T-1} \gamma_b \cdot \Delta_T)^2}} \\ & \geq 1 - \frac{2 - \alpha}{\alpha} e^{-(\eta \sum_{b=0}^{T-1} \gamma_b \cdot \Delta_T)^2}, \end{aligned} \quad (77)$$

where the first step is by (75) and (76). Combining with (73), we can derive

$$\begin{aligned} \gamma_T & \leq \frac{2 - \alpha}{\alpha} e^{-(\eta \sum_{b=0}^{T-1} \gamma_b \cdot \Delta_T)^2} = \frac{2 - \alpha}{\alpha} e^{-(\eta \sum_{b=0}^{T-2} \gamma_b \cdot \Delta_T)^2} \cdot e^{-\eta^2 \Delta_T^2 (2\gamma_{T-1} \sum_{b=0}^{T-2} \gamma_b + \gamma_{T-1}^2)} \\ & = \gamma_{T-1} \cdot e^{-\eta^2 \Delta_T^2 (2\gamma_{T-1} \sum_{b=0}^{T-2} \gamma_b + \gamma_{T-1}^2)}. \end{aligned} \quad (78)$$

When T is large, γ_T is approaching zero. Hence, the equality of (78) is close to being achieved, in which case,

$$\gamma_T \approx \gamma_{T-1} \cdot e^{-\eta^2 \Delta_T^2 (2\gamma_{T-1} \sum_{b=0}^{T-2} \gamma_b + \gamma_{T-1}^2)}. \quad (79)$$

We can observe that when $\sum_{b=0}^{t_0-1} \eta \gamma_b \Delta_T \geq \sqrt{\log M_1}$, γ_{t_0} reaches $\Theta(1/M_1 \cdot \frac{2-\alpha}{\alpha})$. Similarly, when $\sum_{b=0}^{t'_0-1} \eta \gamma_b \Delta_T \leq \sqrt{\log C}$ for some $C > 1$, $\gamma_{t'_0}$ is still $\Theta(1)$, which indicates $t'_0 \lesssim \eta^{-1} M_1 \sqrt{\log C}$ if we only care about η and M_1 as variables. Therefore, we require that the final T satisfies $T \gtrsim \eta^{-1} M_1 \sqrt{\log M_1}$.

(c) We next look at i where $a_i < 0$. If $i \in \mathcal{U}$, we have that for s such that the y -embedding of \mathbf{p}_s^n is \mathbf{q} , the summation of corresponding softmax value is $1 - \gamma_T$. Furthermore,

$$\begin{aligned} & \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} \mathbf{p}_s^n \\ & \lesssim -\delta^2 \frac{\alpha \eta T}{2a} - \delta^2 \alpha^2 \frac{(\eta T)^3 M_1}{a^2}, \end{aligned} \quad (80)$$

if $\beta \leq \Theta(1)$. Hence,

$$\text{Relu}(\mathbf{W}_{O(i,\cdot)}^{(T)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(T)} \mathbf{p}_s^n) \text{softmax}((\mathbf{W}_K^{(T)} \mathbf{p}_s^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n))) = 0. \quad (81)$$

(d) If $i \notin \mathcal{W} \cup \mathcal{U}$ and $s \in \mathcal{W}$, we have,

$$\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} \mathbf{p}_s^n \lesssim \frac{1}{\sqrt{M_1}} \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} \mathbf{p}_s^n, \quad (82)$$

by Lemma D.6 and $B \gtrsim M_1$. The final lower bound of $F(\Psi; \mathbf{P}^n)$ is based on the lower bound introduced by $i \in \mathcal{W}$. Then, combining (69), (72), (81), and (82), we can derive

$$\begin{aligned} & F(\Psi; \mathbf{P}^n) \\ & \gtrsim (1 - \gamma_T) \cdot \left(\delta^2 (\beta^2 + 1) \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2} \right). \end{aligned} \quad (83)$$

Therefore, as long as

$$T = \Theta(\eta^{-1} M_1 \delta^{-2/3} \beta^{-2/3} \alpha^{-2/3} \sqrt{\log M_1}), \quad (84)$$

for some large $C > 1$, we can obtain

$$F(\Psi; \mathbf{P}^n) \geq 1. \quad (85)$$

Similarly, we can derive that for $z^n = -1$,

$$F(\Psi; \mathbf{P}^n) \leq -1. \quad (86)$$

Then, we study in-domain generalization. By (67), for any given testing prompt embedding \mathbf{P} with $z = +1$, we have

$$F(\Psi; \mathbf{P}) \geq 1 - \epsilon, \quad (87)$$

and if $z = -1$,

$$F(\Psi; \mathbf{P}) \leq -1 + \epsilon. \quad (88)$$

Therefore,

$$\mathbb{E}_{\mathbf{x}_{query} \sim \mathcal{D}, f \in \mathcal{T}}[\ell(\Psi; \mathbf{P}, y)] \leq \epsilon. \quad (89)$$

□

D.4. Proof of Theorem 3.4

Proof. Note that we require that the fraction of contexts with the same ODR pattern as the query is at least α' . Since we need that there exists at least one context that contains the same ODR pattern as the query, we have

$$l_{ts} \geq \frac{1}{\alpha'}. \quad (90)$$

Consider $\mathbf{p}_{query}^{n'}$ such that the label is $+1$. Let $\boldsymbol{\mu}'_j = \sum_{j=1}^{M_1} c_j \boldsymbol{\mu}_j$ where $\sum_{j=1}^{M_1} c_j^2 = 1$ and $\boldsymbol{\nu}'_k = \sum_{j=1}^{M_2} g_j \boldsymbol{\nu}_j$ where $\sum_{j=1}^{M_2} g_j^2 = 1$. Following the derivation of (74) and (75), we have that for $s \in \mathcal{N}^n$,

$$\begin{aligned} & (\mathbf{W}_K^{(T)} \mathbf{p}_s^{n'})^\top \mathbf{W}_Q^{(T)} \mathbf{p}_{query}^{n'} \\ & \gtrsim \sum_{j=1}^{M_1} c_j^2 \cdot \left(\eta \sum_{b=0}^{T-1} \gamma_b \beta^2 \cdot \left(\delta^2 \frac{\alpha \eta T}{2a} + \delta^2 (\beta + 1)^2 \alpha^2 \frac{(\eta T)^3 M_1}{a^2} \right) \right)^2 \\ & \gtrsim \log M_1. \end{aligned} \quad (91)$$

For ODR patterns, by Proposition 4.1, we have for \mathbf{p}_l^n that has a different ODR pattern than \mathbf{p}_{query}^n ,

$$(\mathbf{W}_K^{(T)} \mathbf{p}_l^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n) \lesssim \left(\frac{1}{M_1} + \frac{1}{M_2} \right) (\mathbf{W}_K^{(T)} \mathbf{p}_s^n)^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^n). \quad (92)$$

Therefore, by similarly defining $\mathcal{N}_j^n = \{\mathbf{p}_s^{n'} : \text{The testing-relevant pattern of } \mathbf{P}^n \text{ is } \boldsymbol{\mu}'_j\}$, we can derive

$$\sum_{s \in \mathcal{N}_j^n} \text{softmax}((\mathbf{W}_K^{(T)} \mathbf{p}_s^{n'})^\top (\mathbf{W}_Q^{(T)} \mathbf{p}_{query}^{n'})) \geq 1 - \frac{2 - \alpha'}{\alpha'} \Theta\left(\frac{1}{M_1}\right) \geq 1 - \frac{2}{\alpha'} \Theta\left(\frac{1}{M_1}\right). \quad (93)$$

Note that for $\mathbf{p}_s^{n'} = \sum_{i=1}^{M_1} c_i \boldsymbol{\mu}_i + \sum_{j=1}^{M_2} \kappa_s^{n'} g_j \boldsymbol{\nu}_j$, when $M_1 \geq M_2$, we can find a set of $\boldsymbol{\mu}_j + \kappa_s^{n'} \boldsymbol{\nu}_j$ from $j = 1$ to $j = M_2$ with g_j as the coefficients. When $M_1 < M_2 = \Theta(M_1)$, we can find a set of $\boldsymbol{\mu}_t + \kappa_s^{n'} \boldsymbol{\nu}_j$ from $j = 1$ to $j = M_2$ with $t \in [M_1]$, g_j as the coefficients likewise. The remaining $\boldsymbol{\mu}_i$ has coefficients of which the summation is smaller than 1.

Therefore, we have that for a certain $i \in \mathcal{W}$ and $\mathbf{p}_s^{n'}$ where the corresponding label-space embedding is \mathbf{q} ,

$$\begin{aligned}
 & \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} \mathbf{p}_s^{n'} \\
 = & \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} \left(\sum_{i=1}^{M_1} c_i \boldsymbol{\mu}_i^\top + \kappa_s^{n'} \sum_{j=1}^{M_2} g_j \boldsymbol{\nu}_j^\top + \boldsymbol{o}_s^{n'} \top, \mathbf{q}^\top \right)^\top \\
 \geq & \sum_{i=1}^{M_1} c_i \left(\delta^2 \beta^2 \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 \alpha^2 \frac{(\eta T)^3 M_1}{a^2}) \right) \\
 & + \left(\delta^2 \frac{\alpha \eta T}{2a} + \delta^2 \alpha^2 \frac{(\eta T)^3 M_1}{a^2} \right) (1 - \epsilon) \\
 \geq & (\delta^2 (\beta^2 + 1) \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2}) \cdot (1 - \epsilon),
 \end{aligned} \tag{94}$$

where the first equality comes from the definition of $\mathbf{p}_s^{n'}$. The first inequality of (94) is derived from (67). The last inequality is by the condition $\sum_{i=1}^{M_1} c_i \geq 1$. Therefore, we can derive that

$$\begin{aligned}
 F(\Psi; \mathbf{P}^{n'}) & \geq (1 - \gamma_T) \left(\delta^2 (\beta^2 + 1) \frac{\alpha \eta T}{2a} + \delta^2 (\beta^2 + 1) \alpha^2 \frac{(\eta T)^3 M_1}{a^2} \right) \cdot (1 - \epsilon) \\
 & \geq 1 - \epsilon,
 \end{aligned} \tag{95}$$

where the first step is by following (85), and the remaining steps are from basic mathematical computation. Likewise, for \mathbf{p}_{query}^n such that the label is -1 , we can obtain

$$F(\Psi; \mathbf{P}^{n'}) < -(1 - \epsilon). \tag{96}$$

Therefore, we have

$$\mathbb{E}_{\mathbf{x}_{query} \sim \mathcal{D}', f \in \mathcal{T}'} [\ell(\Psi; \mathbf{P}, y)] \leq \epsilon. \tag{97}$$

□

D.5. Proof of Theorem 3.7

Proof. We cover the proof in the proof of Proposition 4.5. Please see Section E.4 for more details. □

E. Proofs of Key Lemmas and propositions

E.1. Proof of Lemma D.11

Proof. We first show that if $|\mathcal{T}_{tr}| < M_1$, Condition 3.2 cannot hold. Then, We show that there exists \mathcal{T}_{tr} with $|\mathcal{T}_{tr}| \geq M_1$ such that Condition 3.2 holds.

(1) If $|\mathcal{T}_{tr}| < M_1$, then $|\mathcal{T}_{tr}|/M_1 < 1$, which is contradict to $|\mathcal{T}_{tr}|/M_1 \geq 1$ in Condition 3.2.

(2) The following example satisfies $|\mathcal{T}_{tr}| \geq M_1$. In this example, the i -th task function ($i \in [M_1]$) in \mathcal{T}_{tr} maps the query with $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_{i+1}$ as IDR patterns to $+1$ and -1 , respectively, where we denote $\boldsymbol{\mu}_{M_1+1} := \boldsymbol{\mu}_1$. Hence, the numbers of tasks that map $\boldsymbol{\mu}_j$ to $+1$ and -1 are both 1 for any $j \in [M_1]$. In this case, $|\mathcal{T}_{tr}| = M_1$.

□

E.2. Proof of Proposition 4.1

Proof. We first show the results for in-domain patterns.

(1) We investigate the results about \mathbf{W}_Q and then \mathbf{W}_K . For \mathbf{p}_{query} with $\boldsymbol{\mu}_j$ as the IDR pattern and $\mathbf{a} \in$

$\{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\} \setminus \{\boldsymbol{\mu}_j\}$, by (30), we have

$$\begin{aligned}
 & (\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \\
 &= (\mathbf{a}^\top, \mathbf{0}^\top) (\mathbf{W}_Q^{(0)} + \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b}) \mathbf{p}_{query} \\
 &\lesssim \frac{1}{M_1} \cdot (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b} \mathbf{p}_{query},
 \end{aligned} \tag{98}$$

if \mathbf{a} does not form a task in \mathcal{T}_{tr} with $\boldsymbol{\mu}_j$. If \mathbf{a} forms a task in \mathcal{T}_{tr} with $\boldsymbol{\mu}_j$, and $\eta T = \Theta(M_1 \sqrt{\log M_1})$, by (29)

$$\begin{aligned}
 & (\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \\
 &\lesssim \frac{1}{M_1} \cdot (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b} \mathbf{p}_{query}.
 \end{aligned} \tag{99}$$

For $\mathbf{a} \perp \boldsymbol{\mu}_j$ but $\mathbf{a} \notin \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\}$, by (31), we have

$$\begin{aligned}
 & (\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top)^\top \\
 &= (\mathbf{a}^\top, \mathbf{0}^\top) (\mathbf{W}_Q^{(0)} + \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b}) \mathbf{p}_{query} \\
 &\lesssim \frac{1}{M_2} \cdot (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b} \mathbf{p}_{query}.
 \end{aligned} \tag{100}$$

By (28) and the initialization, we have

$$(\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \geq (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b} \mathbf{p}_{query} \gtrsim \sqrt{\log M_1} + \delta \gtrsim \sqrt{\log M_1}, \tag{101}$$

where the $\sqrt{\log M_1}$ in the second step comes from that $\eta T \geq \Theta(M_1) \sqrt{\log M_1}$. Hence, by combining (98), (99), and (101), we can derive that

$$\begin{aligned}
 \|\mathbf{W}_Q^{(T)} \mathbf{p}_{query}\| &\lesssim \sqrt{1 + \frac{1}{M_1^2} \cdot M_1 + \frac{1}{M_1^2} + \frac{1}{M_2^2} \cdot M_2 (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b} \mathbf{p}_{query}} \\
 &= \sqrt{1 + \frac{1}{M_1} + \frac{1}{M_2} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q^{(t)}} \Big|_{t=b} \mathbf{p}_{query}},
 \end{aligned} \tag{102}$$

where in the first step, the first $1/M_1^2 \cdot M_1$ comes from (98) with $M_1 - 1$ choices of \mathbf{a} . The second $1/M_1^2$ comes from (99), i.e., $(1/M_1)^2 \cdot \Theta(1)$ since there are only a constant number of such cases. The third $1/M_2^2 \cdot M_2$ is from (100) with M_2 choices of \mathbf{a} . Therefore, by (101) and (102), for $\mathbf{a} \in \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\} \setminus \{\boldsymbol{\mu}_j\}$, we have

$$(\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \lesssim \frac{\sqrt{\log M_1}}{M_1}. \tag{103}$$

For $\mathbf{a} \perp \boldsymbol{\mu}_j$ but $\mathbf{a} \notin \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\}$,

$$(\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \lesssim \frac{\sqrt{\log M_1}}{M_2}. \tag{104}$$

For \mathbf{W}_K , we can make derivations following the above steps. For \mathbf{p}_q with $\boldsymbol{\mu}_j$ as the IDR pattern and $\mathbf{a} \in \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\} \setminus \{\boldsymbol{\mu}_j\}$, by (34), we have

$$\begin{aligned} & (\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_K^{(T)} \mathbf{p}_q \\ & \lesssim \frac{1}{M_1} \cdot (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=b} \mathbf{p}_q, \end{aligned} \quad (105)$$

if \mathbf{a} does not form a task in \mathcal{T}_{tr} with $\boldsymbol{\mu}_j$. If \mathbf{a} forms a task in \mathcal{T}_{tr} with $\boldsymbol{\mu}_j$, and $\eta T = \Theta(M_1 \sqrt{\log M_1})$, by (33),

$$\begin{aligned} & (\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_K^{(T)} \mathbf{p}_q \\ & \lesssim \frac{1}{M_1} \cdot (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=b} \mathbf{p}_q. \end{aligned} \quad (106)$$

For $\mathbf{a} \perp \boldsymbol{\mu}_j$ but $\mathbf{a} \notin \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\}$, by (35), we have

$$\begin{aligned} & (\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_K^{(T)} \mathbf{p}_q \\ & \lesssim \frac{1}{M_2} \cdot (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=b} \mathbf{p}_q. \end{aligned} \quad (107)$$

By (32) and the initialization, we have

$$(\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_K^{(T)} \mathbf{p}_q \geq \sqrt{\log M_1} + \delta \geq \sqrt{\log M_1}. \quad (108)$$

Hence, by combining (105), (106), and (107), we can derive that

$$\|\mathbf{W}_K^{(T)} \mathbf{p}_i\| \lesssim \sqrt{1 + \frac{1}{M_1} + \frac{1}{M_1^2} + \frac{1}{M_2} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{b=0}^{T-1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K^{(t)}} \Big|_{t=b} (\boldsymbol{\mu}_j^\top, \pm \mathbf{q})^\top}. \quad (109)$$

Therefore, by (108) and (109), for $\mathbf{a} \in \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\} \setminus \{\boldsymbol{\mu}_j\}$, we have

$$(\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_K^{(T)} \mathbf{p}_q \lesssim \frac{\sqrt{\log M_1}}{M_1}. \quad (110)$$

For $\mathbf{a} \perp \boldsymbol{\mu}_j$ but $\mathbf{a} \notin \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{M_1}\}$,

$$(\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_K^{(T)} \mathbf{p}_q \lesssim \frac{\sqrt{\log M_1}}{M_2}. \quad (111)$$

(2) For out-of-domain patterns, we have the following derivation. Let $\boldsymbol{\mu}'_j = \sum_{i=1}^{M_1} k_{ji} \boldsymbol{\mu}_i$ where $\sum_{i=1}^{M_1} k_{ji} \geq 1$ and $\sum_{i=1}^{M_1} k_{ji}^2 = 1$. Then, for a query \mathbf{p}'_{query} , of which the corresponding ODR pattern is $\boldsymbol{\mu}'_j$, we have that by (28), (29), (30), and (31),

$$|(\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}'_{query}| \geq |k_j| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \left(1 - \frac{\Theta(1)}{M_1} - \frac{\Theta(1)}{M_2}\right), \quad (112)$$

$$|(\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(t)} \mathbf{p}'_{query}| \leq |k_j| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(t)} \mathbf{p}_{query} \left(1 + \frac{\Theta(1)}{M_1} + \frac{\Theta(1)}{M_2}\right), \quad (113)$$

for any \mathbf{p}_{query} with $\boldsymbol{\mu}_j$ as the IDR pattern. Meanwhile,

$$|(\boldsymbol{\nu}_k^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(t)} \mathbf{p}'_{query}| \leq \frac{1}{M_2} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(t)} \mathbf{p}_{query}. \quad (114)$$

Therefore,

$$\|\mathbf{W}_Q^{(t)} \mathbf{p}'_{query}\| \geq \sqrt{\log M_1} \left(1 - \frac{\Theta(1)}{M_1} - \frac{\Theta(1)}{M_2}\right) \gtrsim \sqrt{\log M_1}, \quad (115)$$

$$\|\mathbf{W}_Q^{(t)} \mathbf{p}'_{query}\| \leq \sqrt{\log M_1} \left(1 + \frac{\Theta(1)}{M_1} + \frac{\Theta(1)}{M_2}\right) \lesssim \sqrt{\log M_1}, \quad (116)$$

$$|(\boldsymbol{\mu}'_j{}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}'_{query}| \gtrsim \sum_{i=1}^{M_1} |k_{ji}| \sqrt{\log M_1} \geq \sqrt{\log M_1}. \quad (117)$$

For $\boldsymbol{\mu}_a \in \{\boldsymbol{\mu}'_1, \dots, \boldsymbol{\mu}'_{M'_1}\} \setminus \{\boldsymbol{\mu}'_j\}$, let $\boldsymbol{\mu}_a = \sum_{i=1}^{M_1} k_{ai} \boldsymbol{\mu}_i$, we have

$$(\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \lesssim \left(\frac{1}{M_1} + \frac{1}{M_2}\right) \sum_{i=1}^{M_1} |k_{ai} k_{ji}| \leq \sqrt{\log M_1} \left(\frac{1}{M_1} + \frac{1}{M_2}\right), \quad (118)$$

where the first step is by (112) and (113), and the second step is by Cauchy-Schwarz inequality given that $\sum_{i=1}^{M_1} k_{ji}^2 = \sum_{i=1}^{M_1} k_{ai}^2 = 1$. For $\mathbf{a} \perp \boldsymbol{\mu}'_j$ but $\mathbf{a} \notin \{\boldsymbol{\mu}'_1, \dots, \boldsymbol{\mu}'_{M'_1}\}$,

$$(\mathbf{a}^\top, \mathbf{0}^\top) \mathbf{W}_Q^{(T)} \mathbf{p}_{query} \lesssim \sqrt{\log M_1} \left(\frac{1}{M_1} + \frac{1}{M_2}\right). \quad (119)$$

Likewise, we can derive the conclusion for the testing context with $\mathbf{W}_K^{(T)}$.

□

E.3. Proof of Corollary 4.3

Proof. From (75) to (79), we can derive the conclusion for IDR patterns. For ODR patterns, from (93), we can obtain the conclusion. Note that $\frac{2-\alpha}{\alpha} = \Theta(1)$ since $\alpha = \Theta(1)$. □

E.4. Proof of Proposition 4.5

Proof. For any i , we denote $\mathbf{W}_{O(i,\cdot)}^{(b)} = (\mathbf{O}_{i,1}^{(b)}, \mathbf{O}_{i,2}^{(b)}, \mathbf{0}^\top)$ where $\mathbf{O}_{i,1}^{(b)} \in \mathbb{R}^{d_x}$ and $\mathbf{O}_{i,2}^{(b)} \in \mathbb{R}^{d_y}$. Following the derivation of (71), we can obtain that for $s \in [l]$ or s is the query,

$$\begin{aligned} & \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} (\mathbf{p}_s^n, \mathbf{0}^\top)^\top \\ &= \mathbf{W}_{O(i,\cdot)}^{(T)} (\delta(\mathbf{p}_s^n, \mathbf{0}^\top)^\top + \sum_{b=0}^{T-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{O}_{i,1}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{O}_{i,1}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{O}_{i,1}^{(b)} \right)^\top) \\ & \gtrsim \delta^2 \beta^2 \frac{\alpha \eta T}{2a} + \delta^2 \beta^2 \alpha^2 \frac{(\eta T)^3 M_1}{a^2}, \end{aligned} \quad (120)$$

for any $j \in [M_1]$, and

$$\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} (\kappa_i^n \boldsymbol{\nu}_k^\top, \mathbf{0}^\top)^\top \leq \frac{1}{M_2} \cdot \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top)^\top, \quad (121)$$

for $k \in [M_2]$ by (44) and (50) in Lemma D.6 and D.7, respectively. Then, we have

$$\begin{aligned} & \frac{(\frac{1}{M_1} \sum_{j=1}^{M_1} \boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}}{\|(\frac{1}{M_1} \sum_{j=1}^{M_1} \boldsymbol{\mu}_j^\top, \mathbf{0}^\top)\| \| \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} \|} \geq \frac{1}{\sqrt{1 + \frac{1}{M_2^2} \cdot M_2}} \\ & \geq 1 - \frac{\Theta(1)}{M_2}, \end{aligned} \quad (122)$$

because $BT \geq \Theta(M_1^2)$. For any $i \in \mathcal{W}$,

$$\begin{aligned} & \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} (\mathbf{0}^\top, \mathbf{q}^\top)^\top \\ &= \mathbf{W}_{O(i,\cdot)}^{(T)} (\delta(\mathbf{0}^\top, \mathbf{q}^\top)^\top + \sum_{b=0}^{T-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{O}_{i,2}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{O}_{i,2}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{O}_{i,2}^{(b)} \right)^\top) \\ & \gtrsim \delta^2 \frac{\alpha \eta T}{2a} + \delta^2 \alpha^2 \frac{(\eta T)^3 M_1}{a^2}. \end{aligned} \quad (123)$$

Note that by gradient updates of \mathbf{W}_O and \mathbf{W}_V , there are no gradient components perpendicular to \mathbf{p} except some Gaussian noise. Hence,

$$\begin{aligned} \frac{(\mathbf{0}^\top, \mathbf{q}^\top) \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}}{\|(\mathbf{0}^\top, \mathbf{q}^\top)\| \|\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}\|} &\geq \frac{1}{\sqrt{1+\xi}} \\ &\geq 1 - \frac{\Theta(1)}{M_1}. \end{aligned} \quad (124)$$

For any $i \in \mathcal{U}$,

$$\begin{aligned} &\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)} (\mathbf{0}^\top, -\mathbf{q}^\top)^\top \\ &= \mathbf{W}_{O(i,\cdot)}^{(T)} (\delta(\mathbf{0}^\top, -\mathbf{q}^\top)^\top + \sum_{b=0}^{T-1} \eta (\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{O}_{i,2}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{O}_{i,2}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{O}_{i,2}^{(b)})^\top) \\ &\gtrsim \delta^2 \frac{\alpha \eta T}{2a} + \delta^2 \alpha^2 \frac{(\eta T)^3 M_1}{a^2}. \end{aligned} \quad (125)$$

Similarly to (124), we have

$$\frac{(\mathbf{0}^\top, -\mathbf{q}^\top) \mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}}{\|(\mathbf{0}^\top, -\mathbf{q}^\top)\| \|\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}\|} \geq 1 - \frac{\Theta(1)}{M_1}. \quad (126)$$

Hence, for $i \in \mathcal{W} \cup \mathcal{U}$,

$$\|\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}\| \gtrsim \beta^{-1} = \Omega(1). \quad (127)$$

By (82), we have that for $i \notin \mathcal{W} \cup \mathcal{U}$,

$$\|\mathbf{W}_{O(i,\cdot)}^{(T)} \mathbf{W}_V^{(T)}\| \lesssim \sqrt{\frac{1}{M_1} + \frac{1}{M_2^2} \cdot M_2} = \frac{1}{\sqrt{M_2}}, \quad (128)$$

where $1/M_1$ is the square of (82). $1/M_2^2$ is the square of the scaling in RHS of (121), and M_2 is the number of IDI patterns.

If we prune all neurons $i \notin \mathcal{W} \cup \mathcal{U}$, we have that

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_{query}, f} [\ell(\Psi; \mathbf{P}, y)] &\leq 1 - (1 - \frac{2}{\alpha' M_1}) \frac{1 - \epsilon}{(1 - \frac{2}{\alpha M_1})} (1 - \frac{1}{\sqrt{M_1}}) \\ &\leq 1 - (1 - \frac{2}{\alpha' M_1}) (1 - \epsilon) (1 - \frac{1}{\sqrt{M_1}}) \\ &\leq 1 - (1 - \frac{2}{\alpha' M_1} - \epsilon - \frac{1}{\sqrt{M_1}}) \\ &\leq \mathcal{O}(\epsilon + \frac{1}{\sqrt{M_1}} + \frac{1}{\alpha' M_1}) \\ &\leq \mathcal{O}(\epsilon + \frac{1}{\sqrt{M_1}}), \end{aligned} \quad (129)$$

where the first step combines (83), (82), and $2/(\alpha M_1)$ and $2/(\alpha' M_1)$ comes from (77) and (93). The last step comes from $\alpha' \geq M_1^{-1/2}$. Meanwhile, if we prune R fraction of neurons in $\mathcal{W} \cup \mathcal{U}$, given (67), we have for the trained model Ψ ,

$$F(\Psi; \mathbf{P}^n) \leq (1 + \epsilon)(1 - R) \cdot \frac{(1 - \frac{2}{\alpha' M_1})}{(1 - \frac{2}{\alpha M_1})}. \quad (130)$$

Then,

$$\begin{aligned}
 \mathbb{E}_{\mathbf{x}_{query}, f} [\ell(\Psi; \mathbf{P}, y)] &\geq 1 - \left(1 - \frac{2}{\alpha' M_1}\right) \frac{1 + \epsilon}{\left(1 - \frac{2}{\alpha M_1}\right)} (1 - R) \\
 &\geq 1 - \left(1 - \frac{2}{\alpha' M_1}\right) (1 + \epsilon) \left(1 + \frac{4}{\alpha M_1}\right) (1 - R) \\
 &= 1 - \left(1 - R - \frac{2}{\alpha' M_1} + \frac{2R}{\alpha' M_1}\right) \left(1 + \epsilon + \frac{4}{\alpha M_1} + \frac{4\epsilon}{\alpha M_1}\right) \\
 &\geq R + \frac{2}{\alpha' M_1} - \frac{2R}{\alpha' M_1} - \left(1 - R - \frac{2}{\alpha' M_1}\right) \left(\epsilon + \frac{4 + 4\epsilon}{\alpha M_1}\right) \\
 &\geq \Omega\left(R + \frac{1}{\alpha' M_1}\right),
 \end{aligned} \tag{131}$$

where the second step is by $(1 - x)^{-1} \leq 1 + 2x$ for small $x > 0$, and the last step is by $R = \Theta(1)$. \square

E.5. Proof of Lemma D.5

Proof. We first study the gradient of $\mathbf{W}_Q^{(t+1)}$ in part (a) and the gradient of $\mathbf{W}_K^{(t+1)}$ in part (b). The proof is derived with a framework of induction combined with Lemma D.6 and D.7.

(a) From the training loss function, by basic mathematical computation, we can obtain

$$\begin{aligned}
 &\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial F(\mathbf{p}_{query}^n)} \frac{\partial F(\mathbf{p}_{query}^n)}{\partial \mathbf{W}_Q} \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i, \cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{\top} \mathbf{W}_K^{\top} \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 &\quad \cdot \left(\mathbf{W}_{O(i, \cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{\top} \mathbf{W}_K^{\top} \mathbf{W}_Q \mathbf{p}_{query}^n) \right. \\
 &\quad \cdot \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{\top} \mathbf{W}_K^{\top} \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{W}_K (\mathbf{p}_s^n - \mathbf{p}_r^n) \mathbf{p}_{query}^{\top} \left. \right) \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i, \cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{\top} \mathbf{W}_K^{\top} \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 &\quad \cdot \left(\mathbf{W}_{O(i, \cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{\top} \mathbf{W}_K^{\top} \mathbf{W}_Q \mathbf{p}_{query}^n) \right. \\
 &\quad \cdot \left. \left(\mathbf{W}_K \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{\top} \mathbf{W}_K^{\top} \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{W}_K \mathbf{p}_r^n \right) \mathbf{p}_{query}^{\top} \right).
 \end{aligned} \tag{132}$$

If $t = 0$, we have that

$$(\mathbf{W}_K^{(t)} \mathbf{p}_s^n)^{\top} (\mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) = \mathbf{p}_s^{\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n. \tag{133}$$

When $z^n = +1$, let \mathbf{x}_{query}^n be a noisy version of $\boldsymbol{\mu}_j + \kappa_{query}^n \boldsymbol{\nu}_k$ where $j \in \{1, 2, \dots, M_1\}$ and $k \in \{1, 2, \dots, M_2\}$. Define m_i as the corresponding IDR pattern in the i -th demonstration. Consider the categorical distribution where the probability of selecting $\boldsymbol{\mu}_q$ is $\alpha/2$. We know there exists a $\boldsymbol{\mu}_j$ such that the probability of selecting $\boldsymbol{\mu}_j$ is also $\alpha/2$. Selecting other $\boldsymbol{\mu}_t$ for $t \neq p, j$ has a probability of $(1 - \alpha)/(M_1 - 2)$. Selecting any IDI pattern $\boldsymbol{\nu}_k$ has a probability of $1/M_2$. By the Chernoff bound of Bernoulli distribution in Lemma D.1, given $l \geq \Theta(\max\{M_1, M_2\} \log M_1)$, we can obtain

$$\Pr \left(\sum_{i=1}^l \mathbb{1}[m_i = \boldsymbol{\mu}_j] \leq l \cdot \frac{\alpha}{2} \right) \leq e^{-C \log M_1} = M_1^{-C}, \tag{134}$$

$$\Pr \left(\sum_{i=1}^l \mathbb{1}[m_i = \boldsymbol{\mu}_s] \leq l \cdot \frac{\alpha}{2} \right) \leq e^{-C \log M_1} = M_1^{-C}, \quad (135)$$

$$\Pr \left(\sum_{i=1}^l \mathbb{1}[m_i = \boldsymbol{\mu}_t] \geq l \cdot \frac{1}{M_1} \right) \leq e^{-C \log M_1 \cdot M_1 \cdot \frac{1}{M_1}} = M_1^{-C}, \quad (136)$$

$$\Pr \left(\sum_{i=1}^l \mathbb{1}[m_i = \boldsymbol{\nu}_k] \geq l \cdot \frac{1}{M_2} \right) \leq e^{-C \log M_1 \cdot M_2 \cdot \frac{1}{M_2}} = M_1^{-C}, \quad (137)$$

for some $C > 0$. Therefore, since that $\frac{1}{\sqrt{M_2}} \cdot e^{\delta^2} \lesssim \frac{\alpha}{2} = \Theta(1)$,

$$\begin{aligned} \sum_{s \in \mathcal{N}_j^{n,i} \cap \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta + 1)} &\leq l \cdot \frac{1}{\sqrt{M_2}} e^{\delta^2} \cdot e^{\delta^2(\beta \cdot \beta)} \\ &\lesssim l \cdot \frac{\alpha}{2} \cdot e^{\delta^2(\beta \cdot \beta)} \\ &\lesssim \sum_{s \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta)}. \end{aligned} \quad (138)$$

Similarly,

$$\begin{aligned} \sum_{s \in \mathcal{M}_k^{n,i} - \mathcal{N}_j^{n,i}} e^{\delta^2} &\leq l \cdot \frac{1}{\sqrt{M_2}} e^{\delta^2} \cdot e^{\delta^2(\beta \cdot \beta)} \\ &\lesssim l \cdot \frac{\alpha}{2} \cdot e^{\delta^2(\beta \cdot \beta)} \\ &\lesssim \sum_{s \in [l] - \mathcal{N}_j^{n,i} - \mathcal{N}_k^{n,i}} e^{\delta^2}, \end{aligned} \quad (139)$$

where the last step is by the fact that there exists $\boldsymbol{\mu}_s$ for $p \in \{1, 2, \dots, M_2\} \setminus \{j\}$ such that selecting $\boldsymbol{\mu}_s$ has a probability of $\alpha/2$. Let $i \in \mathcal{W}$, $s \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}$, then

$$\begin{aligned} &\text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{\text{query}}^n) \\ &\geq e^{\delta^2(\beta \cdot \beta)} \cdot \left(\sum_{s \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta)} + \sum_{s \in \mathcal{N}_j^{n,i} \cap \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta + 1)} \right. \\ &\quad \left. + \sum_{s \in [l] - \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2} + \sum_{s \in \mathcal{M}_k^{n,i} - \mathcal{N}_j^{n,i}} e^{\delta^2} \right)^{-1} \\ &\gtrsim \frac{e^{\delta^2(\beta \cdot \beta)}}{\sum_{s \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta)} + \sum_{s \in [l] - \mathcal{N}_j^{n,i} - \mathcal{N}_k^{n,i}} e^{\delta^2}}, \end{aligned} \quad (140)$$

where the second step is by (138) and (139). Similarly, for $s \in \mathcal{N}_j^{n,i} \cap \mathcal{M}_k^{n,i}$,

$$\begin{aligned} &\text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{\text{query}}^n) \\ &\gtrsim \frac{e^{\delta^2(\beta \cdot \beta)}}{\sum_{s \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta)} + \sum_{s \in [l] - \mathcal{N}_j^{n,i} - \mathcal{N}_k^{n,i}} e^{\delta^2}}. \end{aligned} \quad (141)$$

For $s \in \mathcal{M}_k^{n,i} - \mathcal{N}_j^{n,i}$,

$$\begin{aligned} &\text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{\text{query}}^n) \\ &\lesssim \frac{e^{\delta^2}}{\sum_{s \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta)} + \sum_{s \in [l] - \mathcal{N}_j^{n,i} - \mathcal{N}_k^{n,i}} e^{\delta^2}}. \end{aligned} \quad (142)$$

For $s \in [l] - \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}$,

$$\begin{aligned} & \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ & \lesssim \frac{e^{\delta^2}}{\sum_{s \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2(\beta \cdot \beta)} + \sum_{s \in [l] - \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} e^{\delta^2}}. \end{aligned} \quad (143)$$

By (53) and (54) in Definition D.8, we have that for $i \in \mathcal{W}_n$,

$$\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) > 0. \quad (144)$$

Then we derive

$$\begin{aligned} & \mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n \\ & = \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \\ & = \left(\sum_{r \in \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} + \sum_{r \in \mathcal{N}_j^{n,i} \cap \mathcal{M}_k^{n,i}} + \sum_{r \in \mathcal{M}_k^{n,i} - \mathcal{N}_j^{n,i}} + \sum_{r \in [l] - \mathcal{N}_j^{n,i} - \mathcal{M}_k^{n,i}} \right) \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ & \quad \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \mathbf{W}_K^{(t)} \mathbf{p}_r^n). \end{aligned} \quad (145)$$

One can observe that

$$\begin{aligned} & \sum_{s \in \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \\ & = \sum_{s \in \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - (\sum_{r \in \mathcal{N}_j^{n,i}} + \sum_{r \notin \mathcal{N}_j^{n,i}}) \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \\ & = \sum_{r \notin \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \cdot \sum_{s \in \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_s^n \\ & \quad - \sum_{s \in \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \cdot \sum_{r \notin \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n + \mathbf{n} \\ & = \sum_{s \in \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \cdot \sum_{r \notin \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} (\mathbf{p}_s^n - \mathbf{p}_r^n) + \mathbf{n}. \end{aligned} \quad (146)$$

Hence, by Definition D.4,

$$1 > \sum_{r \in [l] - \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \geq \gamma_t > 0. \quad (147)$$

Since that the feature space embedding of $(\mathbf{p}_r^n \top, \mathbf{0}^\top)^\top$ are orthogonal to $\mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n$ for $r \in [l] - \mathcal{N}_j^{n,i}$, we have that with high probability, for $s \in \mathcal{N}_j^{n,i}$,

$$\begin{aligned} & (\mathbf{x}_s^n \top, \mathbf{0}^\top) \sum_{r \in [l] - \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_r^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \\ & \geq \gamma_t \beta^2 \delta, \end{aligned} \quad (148)$$

where γ_t comes from the definition. β is from the definition of the data. Meanwhile, for r such that μ_r is the IDR pattern with the probability of $\alpha/2$ to be selected,

$$\begin{aligned} & \left| (\mathbf{x}_r^{n\top}, \mathbf{0}^\top) \sum_r \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \right| \\ & \leq (\mathbf{x}_s^{n\top}, \mathbf{0}^\top) \sum_{r \in [l] - \mathcal{N}_j^{n,i}} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \cdot \frac{1 - \alpha}{1 - \alpha/2}, \end{aligned} \quad (149)$$

where $(1 - \alpha)/(1 - \alpha/2)$ comes from the fraction of attention weights on μ_r in $[l] - \mathcal{N}_j^{n,i}$. If μ_r is the pattern that does not decide the label of the current \mathbf{P}^n , we have

$$\begin{aligned} & \left| (\mathbf{x}_r^{n\top}, \mathbf{0}^\top) \sum_r \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \right| \\ & \leq \frac{\gamma_t}{l}, \end{aligned} \quad (150)$$

where l in the denominator comes from that with high probability, at most 1 μ_r appears in one data for a certain r . Therefore, for $i \in \mathcal{W}_n$, we denote that $\zeta'_{i,n} = \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s \in \mathcal{N}_j^{n,i}} (\mathbf{W}_V^{(t)} \mathbf{p}_s^{n(t)}) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n)$. Then, if $\zeta'_{i,n} > 0$, we have that for μ_q that has the same IDR pattern as \mathbf{x}_{query} ,

$$\begin{aligned} & (\mu_q^\top, \mathbf{0}^\top) \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s \in \mathcal{N}_j^{n,i}} (\mathbf{W}_V^{(t)} \mathbf{p}_s^{n(t)}) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ & \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\ & \geq \zeta'_{i,n} \delta \beta^4 \gamma_t, \end{aligned} \quad (151)$$

where γ_t comes from that $\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{p}_s^n$ is much larger in average, $i \in \mathcal{W}_n$, if than other $\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{p}_t^n$ if $s \in \mathcal{N}_*^n$ while $s \notin \mathcal{N}_*^n$. For j such that μ_j is the IDR pattern with the probability of $\alpha/2$ to be selected but different from q ,

$$\begin{aligned} & (\mu_j^\top, \mathbf{0}^\top) \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s \in \mathcal{N}_j^{n,i}} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ & \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\ & \leq \frac{1 - \alpha}{1 - \alpha/2} (\mathbf{x}_{query}^\top, \mathbf{0}^\top) \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s \in \mathcal{N}_j^{n,i}} (\mathbf{W}_V^{(t)} \mathbf{p}_s^{n(t)}) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n). \end{aligned} \quad (152)$$

For μ_j that does not decide the label of the current \mathbf{P}^n , we have

$$\begin{aligned} & (\mu_j^\top, \mathbf{0}^\top) \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s \in \mathcal{N}_j^{n,i}} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ & \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\ & \leq \frac{\zeta'_{i,n} \delta \gamma_t \beta^4}{l}. \end{aligned} \quad (153)$$

To deal with $s \in [l] - \mathcal{N}_j^{n,i}$, we cover this part when summing up all the neurons. Since that each entry of $\mathbf{W}_{O(i,\cdot)}$ follows $\mathcal{N}(0, \xi^2)$, we have

$$\Pr(\|\mathbf{W}_{O(i,1:d_{\mathcal{X}})} \mathbf{x}_{query}^n\| \leq \beta \xi) \leq \beta \xi, \quad (154)$$

by the standard property of Gaussian distribution. Meanwhile, by Hoeffding's inequality (27),

$$\Pr(\|\mathbf{W}_{O_{(i,1:d_X)}} \mathbf{x}_{query}^n\| \geq \beta\xi \log M_1) \leq M_1^{-C}, \quad (155)$$

for some $C > 1$. Hence, with a high probability, by Hoeffding's inequality (27),

$$\begin{aligned} \left| \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O_{(i,1:d_X)}}^{(t)} \mathbf{x}_{query}^n \right| &\lesssim \frac{|\mathcal{W}_n|}{m} \Phi(0) \beta\xi + \beta\xi \cdot \frac{\log M_1}{\sqrt{m}} \\ &\lesssim \beta\xi, \end{aligned} \quad (156)$$

$$\begin{aligned} \left| \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O_{(i,1:d_X)}}^{(t)} \mathbf{x}_{query}^n \right| &\gtrsim \frac{|\mathcal{W}_n|}{m} \Phi(0) \beta\xi - \beta\xi \cdot \frac{\log M_1}{\sqrt{m}} \\ &\gtrsim \beta\xi. \end{aligned} \quad (157)$$

For p such that the probability of selecting $\boldsymbol{\mu}_p$ is $\alpha/2$, we have

$$\left| \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O_{(i,1:d_X)}}^{(t)} \boldsymbol{\mu}_p \right| \lesssim \left| \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O_{(i,1:d_X)}}^{(t)} \boldsymbol{\mu}_j \right| \cdot e^{-\delta\beta^2}. \quad (158)$$

We have that for $z^n = 1$, we can then derive that for $s \in \mathcal{N}_j^i$, by Definition D.4, for $\boldsymbol{\mu}_q$ which is the IDR pattern of \mathbf{x}_{query} ,

$$\begin{aligned} &(\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O_{(i,\cdot)}}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ &\cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^n \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^n \top (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\ &\geq \zeta_{i,t} \delta \gamma_t (1 - e^{-\delta^2 \beta^2}) \beta^4, \end{aligned} \quad (159)$$

where $\zeta_{i,t}$ is used as a lower bound after taking an average of $i \in \mathcal{W}_n$. Similarly, for j such that $\boldsymbol{\mu}_j$ has a probability of $\alpha/2$ to be selected, but different from q we have

$$\begin{aligned} &(\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O_{(i,\cdot)}}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ &\cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^n \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^n \top (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\ &\lesssim e^{-\delta^2 \beta^2} \cdot (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O_{(i,\cdot)}}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\ &\cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^n \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^n \top (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top. \end{aligned} \quad (160)$$

For $j \in [l] - \mathcal{N}_q^{n,i}$ with probability of $(1 - \alpha)/(M_1 - 2)$ to be selected, with high probability, at most 1 example has μ_j in each data. Then,

$$\begin{aligned}
 & (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\
 & \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\
 & \leq \frac{1}{l} \cdot (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\
 & \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top.
 \end{aligned} \tag{161}$$

If $i \in \mathcal{U}_n$, since that $z^n = 1$, the indicator by the Relu activation returns zero. Hence, we do not need to compute this case. If $i \notin \mathcal{W}_n \cup \mathcal{U}_n$, by the uniform distribution of a_i , we have that, for $\boldsymbol{\mu}_q$ which is the IDR pattern of \mathbf{x}_{query} ,

$$\begin{aligned}
 & (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_i^n) \geq 0] \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right) \\
 & \leq \sqrt{\frac{\log m}{m}} \cdot (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\
 & \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top,
 \end{aligned} \tag{162}$$

where $\sqrt{\log m/m}$ is because a_i can be either +1 and -1 following a uniform distribution in this case. For $\boldsymbol{\mu}_j$ that has a probability of $\alpha/2$ to be selected but different from $\boldsymbol{\mu}_q$, we have

$$\begin{aligned}
 & (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \sum_{i \notin \mathcal{W}_i \cup \mathcal{U}} a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_i^n) \geq 0] \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right) \\
 & \leq e^{-\delta\beta^2} \sqrt{\frac{\log m}{m}} \cdot (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\
 & \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top.
 \end{aligned} \tag{163}$$

For \mathbf{x}_j^n with μ_j that has a probability of $(1 - \alpha)/(M_1 - 2)$ to be selected, we have

$$\begin{aligned}
 & (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \sum_{i \notin \mathcal{W}_i \cup \mathcal{U}} a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0] \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} \right) (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\
 & \leq \frac{1}{l} \cdot \sqrt{\frac{\log m}{m}} \cdot (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \frac{1}{m} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \\
 & \cdot \left(\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top.
 \end{aligned} \tag{164}$$

Therefore, by (159), (162), (163), and (164), we have that for one \mathbf{x}_{query} ,

$$\begin{aligned}
 & \left| \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) (\mathbf{x}_{query}^\top, \mathbf{0}^\top) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \cdot \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0] \right. \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. \left. (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^{n\top} \right) (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\
 & \geq \eta \frac{1}{BM_1} \sum_{n \in \mathcal{B}_b} \frac{1}{m} \sum_{i \in \mathcal{W}_n} \zeta_{i,t} \delta \gamma_t (1 - e^{-\delta^2 \beta^2} \sqrt{\frac{\log m}{m}} - \frac{1}{l} \sqrt{\frac{\log m}{m}}) \beta^4 \\
 & \gtrsim \eta \frac{1}{M_1} \zeta_t \delta \gamma_t \beta^4,
 \end{aligned} \tag{165}$$

as long as

$$m \gtrsim 1, \tag{166}$$

and

$$B \gtrsim M_1 \log M_1, \tag{167}$$

to ensure that

$$\Pr \left(\sum_{n=1}^B \mathbb{1}[m_n = \mu_j] \leq B(1 - c) \cdot \frac{1}{M_1} \right) \leq e^{-c^2 B \cdot \frac{1}{M_1}} = e^{-c \log M_1} = M_1^{-C}, \tag{168}$$

for some $c \in (0, 1)$ and $C > 1$, where m_i denotes the IDR pattern in the query of the n -th data. Meanwhile, for $j \in [l] - \mathcal{N}_q^{n,i}$ that has a IDR pattern which forms a task in \mathcal{T}_{tr} with the IDR pattern of \mathbf{x}_{query} , more indicators of $i \in \mathcal{U}_n$ is activated.

$$\begin{aligned}
 & - \left| \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)}] \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \cdot \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0 \right| \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. \left(\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n \right) \mathbf{p}_{query}^{n\top} \right) (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \Big| \\
 & \geq - \frac{1}{2} e^{-\delta^2 \beta^2} \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)}] \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \cdot \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0 \Big| \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. \left(\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n \right) \mathbf{p}_{query}^{n\top} \right) (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top.
 \end{aligned} \tag{169}$$

For other $j \in [l] - \mathcal{N}_q^{n,i}$,

$$\begin{aligned}
 & \left| \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)}] \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \cdot \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0 \right| \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. \left(\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n \right) \mathbf{p}_{query}^{n\top} \right) (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \Big| \\
 & \leq \frac{1}{M_1} \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) (\boldsymbol{\mu}_q^{n\top}, \mathbf{0}^\top) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)}] \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \cdot \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0 \Big| \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\
 & \cdot \left. \left(\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n \right) \mathbf{p}_{query}^{n\top} \right) (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top,
 \end{aligned} \tag{170}$$

where M_1 comes from the fact that the softmax value between \mathbf{p}_{query}^n and \mathbf{p}_r^n with $\boldsymbol{\mu}_j$ as the IDR pattern of \mathbf{p}_r^n is $\Theta(1 - \gamma_t)/M_1$ in average of $B \gtrsim M_1 \log M_1$ samples. Then, by combining (165), (169), and (170), we have

$$\begin{aligned}
 & (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\
 & \gtrsim \eta \frac{1}{M_1} \zeta_t \delta \gamma_t \beta^4.
 \end{aligned} \tag{171}$$

By (169) and (170), we have that for $\boldsymbol{\mu}_j$ which forms a task in \mathcal{T}_{tr} with the $\boldsymbol{\mu}_q$,

$$\begin{aligned}
 & - \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\
 & \geq - \frac{1}{2} e^{-\delta^2 \beta^2} \left| (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|.
 \end{aligned} \tag{172}$$

For μ_j which does not form a task in \mathcal{T}_{tr} with the μ_q ,

$$\begin{aligned} & \left| (\mu_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \leq \frac{1}{M_1} \left| (\mu_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|. \end{aligned} \quad (173)$$

Similarly, for $k \in [M_2]$,

$$\begin{aligned} & \left| (\nu_k^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \lesssim \frac{1}{M_2} \left| (\mu_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|, \end{aligned} \quad (174)$$

where M_2 comes from that for ν_k that is added to μ_j , the contribution of gradient is $1/M_2$ times of replacing ν_k with μ_j . Hence, $1/M_2 \cdot (1 + 1/M_1 \cdot M_1) = 2/M_2 = \Theta(1/M_2)$. For the label embedding, we have

$$\begin{aligned} & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=0} [; d_{\mathcal{X}} + 1 : d_{\mathcal{X}} + d_{\mathcal{Y}}] \\ & = \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^\top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0] \\ & \quad \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^\top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\ & \quad \left. \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^\top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \mathbf{p}_{query}^\top \right) [; d_{\mathcal{X}} + 1 : d_{\mathcal{X}} + d_{\mathcal{Y}}]. \\ & = 0. \end{aligned} \quad (175)$$

We then have

$$\begin{aligned} & \left| \mathbf{q}^\top \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^\top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_l^n) \geq 0] \right. \\ & \quad \cdot \left(\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^\top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \right. \\ & \quad \left. \cdot (\mathbf{W}_K^{(t)} \mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^\top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \mathbf{W}_K^{(t)} \mathbf{p}_r^n) \right) [d_{\mathcal{X}} + 1 : d_{\mathcal{X}} + d_{\mathcal{Y}}] \Big| \\ & = 0. \end{aligned} \quad (176)$$

Hence, the conclusion holds when $t = 1$. Suppose that the statement also holds when $t = t_0$. When $t = t_0 + 1$, the gradient update is the same as in (165) and (169). Note that the indicator of \mathcal{W}_n will not change along the training. The only difference is the changes in ζ_t and γ_t . Thus, we can obtain that for μ_q with the same IDR pattern as \mathbf{x}_{query} ,

$$\begin{aligned} & (\mu_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \Big|_{t=t_0+1} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \\ & \gtrsim \eta \frac{1}{M_1} \sum_{b=0}^{t_0} \zeta_b \delta \gamma_b \beta^4 \\ & \gtrsim \eta \frac{1}{M_1} \sum_{b=0}^{t_0} \zeta_b \delta \gamma_b \beta^4, \end{aligned} \quad (177)$$

as long as (167) holds. We also have

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0+1} [:, d_{\mathcal{X}} + 1 : d_{\mathcal{X}} + d_{\mathcal{Y}}] = \mathbf{0}. \quad (178)$$

Similarly, for $j \neq q$ and $j \in [M_1]$ where μ_j does not form a task in \mathcal{T}_{tr} with the μ_q ,

$$\begin{aligned} & \left| (\mu_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \lesssim \frac{1}{M_1} \left| (\mu_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|. \end{aligned} \quad (179)$$

For $j \neq q$ and $j \in [M_1]$ where μ_j forms a task in \mathcal{T}_{tr} with μ_q ,

$$\begin{aligned} & - \left| (\mu_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \gtrsim - e^{-\delta^2 \beta^2 - (\eta \frac{1}{M_1} \sum_{b=0}^{t_0-1} \zeta_b \delta \gamma_b \beta^2)^2} \left| (\mu_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \gtrsim - e^{-\Theta(\frac{\eta t_0}{M_1})^2} \left| (\mu_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|, \end{aligned} \quad (180)$$

where the first step comes from the fact that a negative gradient update makes the softmax value of μ_i much smaller. The last step is obtained in the order related to η, t, M_1 as variables. Meanwhile, for $k \in [M_2]$,

$$\begin{aligned} & \left| (\nu_k^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right| \\ & \lesssim \frac{1}{M_2} \left| (\mu_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \left. \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_Q} \right|_{t=t_0} (\mathbf{x}_{query}^\top, \mathbf{0}^\top)^\top \right|. \end{aligned} \quad (181)$$

(b) Then we study the updates of \mathbf{W}_K . We can compute the gradient as

$$\begin{aligned} & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n, \Psi)}{\partial \mathbf{W}_K} \\ & = \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \cdot \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\ & \quad \cdot \left(\mathbf{W}_{O(i,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{W}_Q^\top \mathbf{p}_{query}^n \right. \\ & \quad \left. \cdot \left(\mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_r^n \right)^\top \right). \end{aligned} \quad (182)$$

If we investigate $\mathbf{W}_K^{(t)} \mathbf{p}_s^n$, we can tell that the output is a weighed summation of multiple $\mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n$. Similarly, the output of $\mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n$ is a weighed summation of multiple $\mathbf{W}_K^{(t)} \mathbf{p}_s$. Given the initialization $\mathbf{W}_Q^{(0)}$ and $\mathbf{W}_K^{(0)}$, the update of $\mathbf{W}_K^{(t)} \mathbf{p}_s^n$ and $\mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n$ only contains the contribution from the feature space embeddings at the initialization. One

difference is that since that \mathbf{q} appears with $1/2$ probability in all \mathbf{p}_s^n ,

$$\begin{aligned}
 & \left| \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \cdot \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \right. \\
 & \cdot \left(\mathbf{W}_{O(i,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{W}_Q^\top \mathbf{p}_{query}^n \right. \\
 & \left. \cdot \left(\mathbf{p}_s^n - \sum_{r=1}^{l+1} \text{softmax}(\mathbf{p}_r^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_r^n \right)^\top \right) [d_{\mathcal{X}} + 1 : d_{\mathcal{X}} + d_{\mathcal{Y}}] \mathbf{q} \left| \right. \\
 & \leq \sqrt{\frac{\log B}{B}} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0+1} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top)^\top.
 \end{aligned} \tag{183}$$

Following the steps in Part (a), we can obtain that for $\boldsymbol{\mu}_q$ as the IDR pattern of \mathbf{x}_q , $e \in [l]$,

$$\begin{aligned}
 & (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0+1} (\mathbf{x}_e^\top, \mathbf{0}^\top)^\top \\
 & \gtrsim \eta \frac{1}{M_1} \sum_{b=0}^{t_0} \zeta_b \delta \gamma_b \beta^2,
 \end{aligned} \tag{184}$$

and combining (183),

$$\begin{aligned}
 & (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0+1} \mathbf{p}_q \\
 & \gtrsim \eta \frac{1}{M_1} \sum_{b=0}^{t_0} \zeta_b \delta \gamma_b \beta^2 (1 - \sqrt{\frac{\log B}{B}}) \\
 & \gtrsim \eta \frac{1}{M_1} \sum_{b=0}^{t_0} \zeta_b \delta \gamma_b \beta^2,
 \end{aligned} \tag{185}$$

where the last step holds as long as (167). For $\boldsymbol{\mu}_j$ which forms a task in \mathcal{T}_{tr} with the $\boldsymbol{\mu}_q$,

$$\begin{aligned}
 & - \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0} \mathbf{p}_q \right| \\
 & \gtrsim -e^{-\Theta(\frac{\eta t_0}{M_1})^2} \left| (\boldsymbol{\mu}_q^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0} \mathbf{p}_q \right|.
 \end{aligned} \tag{186}$$

For $\boldsymbol{\mu}_j$ which does not form a task in \mathcal{T}_{tr} with the $\boldsymbol{\mu}_q$,

$$\begin{aligned}
 & \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0} \mathbf{p}_q \right| \\
 & \leq \frac{1}{M_1} \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0} \mathbf{p}_q \right|.
 \end{aligned} \tag{187}$$

Meanwhile, for $k \in [M_2]$, similar to (181),

$$\begin{aligned}
 & \left| (\boldsymbol{\nu}_k^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0} \mathbf{p}_q \right| \\
 & \leq \frac{1}{M_2} \left| (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_K} \Big|_{t=t_0} \mathbf{p}_q \right|.
 \end{aligned} \tag{188}$$

□

E.6. Proof of Lemma D.6

Proof.

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V} \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial F(\mathbf{p}_{query}^n)} \frac{\partial F(\mathbf{p}_{query}^n)}{\partial \mathbf{W}_V} \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i=1}^m a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 & \quad \cdot \mathbf{W}_{O(i,\cdot)}^\top \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n\top}.
 \end{aligned} \tag{189}$$

Let \mathbf{x}_i^n and \mathbf{x}_j^n correspond to IDR patterns $\boldsymbol{\mu}_a$ and $\boldsymbol{\mu}_b$, respectively. For \mathbf{p}_{query}^n which corresponds to the IDR feature $\boldsymbol{\mu}_a$,

$$\begin{aligned}
 \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n\top} (\mathbf{x}_i^{n\top}, \mathbf{q}^\top)^\top &\gtrsim \beta^2 (1 - \gamma_t) \cdot 1 - \frac{1}{\frac{\alpha}{2} l} \\
 &\gtrsim \beta^2 (1 - \gamma_t) \cdot 1,
 \end{aligned} \tag{190}$$

where the first step holds since that by (137), with high probability, no other \mathbf{x}_k^n where $k \neq l+1$ shares the same IDI pattern as \mathbf{x}_{query}^n . The last step holds if

$$l_{tr} \gtrsim \frac{1}{\alpha \beta^2}. \tag{191}$$

Meanwhile, by a different IDR pattern of \mathbf{x}_j^n ,

$$\sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n\top} (\mathbf{x}_j^{n\top}, \mathbf{q}^\top)^\top \lesssim \beta^2 \gamma_t. \tag{192}$$

When $t = 0$, for all $i \in \mathcal{W}_n$, we have that by Lemma D.10, for \mathbf{p}_{query}^n that corresponds to $\boldsymbol{\mu}_a$,

$$\mathbf{W}_{O(i,\cdot)}^{(t)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) > 0. \tag{193}$$

Therefore, for any $\mathbf{p}_j^n = (\mathbf{x}_j^{n\top}, \mathbf{y}_j^{n\top})^\top$ where $f^{(n)}(\tilde{\mathbf{x}}_j^n) = +1$, and

$$\mathbf{x}_j^n = \boldsymbol{\mu}_a + \kappa_j^n \boldsymbol{\nu}_b, \tag{194}$$

we have

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V^{(t)}} \mathbf{p}_j^n \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{k=1}^m a_k \mathbb{1}[\mathbf{W}_{O(k,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 & \quad \cdot \mathbf{W}_{O(k,\cdot)}^\top \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n\top} \mathbf{p}_j^n.
 \end{aligned} \tag{195}$$

We then have that by combining (190) and (192),

$$\begin{aligned}
 & -\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i \in \mathcal{W}_n} a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 & \quad \cdot \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n\top} \mathbf{p}_j^n \\
 & \gtrsim \eta \beta^2 (1 - \gamma_t).
 \end{aligned} \tag{196}$$

Since that for \mathbf{p}_s^n and \mathbf{p}_j^n with different label embeddings, their inner product is smaller than $-1 + \beta$ if they share the same IDR pattern, or smaller than -1 if they share different IDR patterns. On average, in a batch, this product is close to -1 . Hence

$$\begin{aligned}
 & -\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i \in \mathcal{U}_n} a_i \mathbb{1}[\mathbf{W}_{O_{(k, \cdot)}}] \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0 \\
 & \cdot \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n \top} \mathbf{p}_j^n \\
 & \leq \frac{1}{\beta^2 + 1} \cdot (-\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i \in \mathcal{W}_n} a_i \mathbb{1}[\mathbf{W}_{O_{(i, \cdot)}}] \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0) \\
 & \cdot \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n \top} \mathbf{p}_j^n.
 \end{aligned} \tag{197}$$

Meanwhile, since that

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} a_i \mathbb{1}[\mathbf{W}_{O_{(i, \cdot)}}] \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0 \\
 & \cdot \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n \top} \mathbf{p}_j^n \\
 & \lesssim \sqrt{\frac{\log B}{B}} \cdot \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{j \in \mathcal{W}_n} a_j \mathbb{1}[\mathbf{W}_{O_{(j, \cdot)}}] \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0 \\
 & \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n \top} \mathbf{p}_j^n,
 \end{aligned} \tag{198}$$

where $\sqrt{\log B/B}$ is because that z^n is selected from $\{+1, -1\}$ with equal probability. Hence, we can denote and derive that when $t = t_0 + 1$,

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \sum_{b=0}^{t_0} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V^{(b)}} \mathbf{p}_j^n \\
 & = \eta \sum_{b=0}^{t_0} \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O_{(i, \cdot)}}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O_{(i, \cdot)}}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O_{(i, \cdot)}}^{(b)} \right),
 \end{aligned} \tag{199}$$

where

$$-V_i(b) \gtrsim \beta^2(1 - \gamma_t)1/a, \quad i \in \mathcal{W}_n, \tag{200}$$

$$-V_i(b) \leq \frac{1}{\beta^2 + 1} V_j(b), \quad i \in \mathcal{U}_n, j \in \mathcal{W}_n, \tag{201}$$

$$|V_i(b)| \lesssim \sqrt{\frac{\log B}{B}} \cdot \frac{1}{a}, \quad i \notin \mathcal{W}_n \cup \mathcal{U}_n. \tag{202}$$

Similarly, for any $\mathbf{p}_j^n = (\mathbf{x}_j^{n \top}, \mathbf{y}_j^{n \top})^\top$ where $f^{(n)}(\tilde{\mathbf{x}}_j^n) = -1$,

$$\mathbf{x}_j^n = \boldsymbol{\mu}_a + \kappa_j^n \boldsymbol{\nu}_b, \tag{203}$$

we have

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V^{(t)}} \mathbf{p}_j^n \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{k=1}^m a_k \mathbb{1}[\mathbf{W}_{O(k,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 & \quad \cdot \mathbf{W}_{O(k,\cdot)}^\top \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n\top} \mathbf{p}_j^n,
 \end{aligned} \tag{204}$$

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \sum_{b=0}^{t_0} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V^{(b)}} \mathbf{p}_j^n \\
 &= \eta \sum_{b=0}^{t_0} \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W} \cup \mathcal{U}} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \right),
 \end{aligned} \tag{205}$$

where

$$-V_i(b) \gtrsim \beta^2 (1 - \gamma_t) 1/a, \quad i \in \mathcal{U}_n, \tag{206}$$

$$-V_i(b) \leq \frac{1}{\beta^2 + 1} V_j(b), \quad i \in \mathcal{W}_n, j \in \mathcal{U}_n, \tag{207}$$

$$|V_i(b)| \lesssim \sqrt{\frac{\log B}{B}} \cdot \frac{1}{a}, \quad i \notin \mathcal{W}_n \cup \mathcal{U}_n. \tag{208}$$

We can also derive

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{b=0}^{t_0} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_V^{(t)}} (\mathbf{v}_k^\top, \mathbf{0}^\top)^\top \\
 &= \eta \frac{1}{B} \sum_{b=0}^{t_0} \sum_{n \in \mathcal{B}_b} (-z^n) \sum_{k=1}^m a_k \mathbb{1}[\mathbf{W}_{O(k,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 & \quad \cdot \mathbf{W}_{O(k,\cdot)}^\top \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \mathbf{p}_s^{n\top} (\mathbf{v}_k^\top, \mathbf{0}^\top)^\top \\
 &:= \eta \sum_{b=0}^{t_0} \left(\sum_{i \in \mathcal{W}_n} V'_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V'_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W} \cup \mathcal{U}} V'_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \right),
 \end{aligned} \tag{209}$$

where

$$|V'_i(b)| \leq |V_i(b)| \cdot \frac{1}{M_2}, \tag{210}$$

since that $1/M_2$ fraction of \mathbf{p}_s^n has \mathbf{v}_K as the IDI pattern in average.

□

E.7. Proof of Lemma D.7

Proof.

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial F(\mathbf{p}_{query}^n)} \frac{\partial F(\mathbf{p}_{query}^n)}{\partial \mathbf{W}_{O(i, \cdot)}} \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) a_i \mathbb{1}[\mathbf{W}_{O(i, \cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\
 & \quad \cdot \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n).
 \end{aligned} \tag{211}$$

We have that

$$\begin{aligned}
 & \mathbf{W}_V^{(t)} \mathbf{p}_s^n \\
 &= \delta (\mathbf{p}_s^{n \top}, \mathbf{0}^\top)^\top + \sum_{b=0}^{t-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right)^\top.
 \end{aligned} \tag{212}$$

Consider a certain $\mathbf{p}_s^n = (\mathbf{x}_s^{n \top}, \mathbf{y}_s^{n \top}, \mathbf{0}^\top)^\top$ where $f^{(n)}(\tilde{\mathbf{x}}_s^n) = +1$, and

$$\mathbf{x}_s^n = \boldsymbol{\mu}_a + \kappa_s^n \boldsymbol{\nu}_b. \tag{213}$$

When $t = 0$, we can obtain that for $i \in \mathcal{W}_n$ and $\boldsymbol{\mu}_a$ as the IDR pattern of \mathbf{p}_{query}^n ,

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) (\delta \mathbf{p}_s^{n \top} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top + \sum_{b=0}^{t-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right. \\
 & \quad \left. + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right)^\top (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top) \\
 & \geq \frac{\alpha \eta}{2a} \delta (\beta^2 + 1).
 \end{aligned} \tag{214}$$

Then, we have the following results by Lemma D.6, and the magnitude of $\boldsymbol{\mu}_a$, $\boldsymbol{\mu}_b$, and \mathbf{q} ,

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, -\mathbf{q}^\top)^\top \leq \frac{\beta^2 - 1}{\beta^2 + 1} \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top, \tag{215}$$

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, \mathbf{q}^\top)^\top \leq \frac{1}{\beta^2 + 1} \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top, \tag{216}$$

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, -\mathbf{q}^\top)^\top \leq -\frac{1}{\beta^2 + 1} \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top, \tag{217}$$

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\nu}_c^\top, \mathbf{0}^\top)^\top \leq \frac{1}{M_2} \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top. \tag{218}$$

Denote the set of data that share one same IDR pattern as \mathbf{p}_{query}^n as \mathcal{B}_b^n in the b -th iteration. Therefore, when $t = 1$, we have

$$\begin{aligned}
 & \eta \frac{1}{|\mathcal{B}_b^n|} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n \top} \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) (\delta \mathbf{p}_s^{n \top} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top) + \sum_{b=0}^{t-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right. \\
 & \quad \left. + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right)^\top (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 & \geq \frac{\alpha}{2} \left(\frac{\eta}{a} (\delta(\beta^2 + 1)) - \frac{\eta}{a} \cdot \eta \cdot \sqrt{\frac{\log B}{B}} \frac{m}{a} \cdot \xi \log M_1 \right. \\
 & \quad \left. + \frac{\eta}{a} \cdot \eta \frac{m}{a} (\beta^2(1 - \gamma_t)) \left(\frac{\eta}{a} \delta(\beta^2 + 1) - \xi \right) \right) \\
 & \gtrsim \frac{\alpha}{2} \left(\frac{\eta}{a} (\delta(\beta^2 + 1)) + \frac{\eta}{a} \cdot \eta \frac{m}{a} (\beta^2(1 - \gamma_t)) \frac{\eta}{a} \delta(\beta^2 + 1) \right) \\
 & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2} \right),
 \end{aligned} \tag{219}$$

where the first inequality comes from that the update in the previous step makes the output of $\mathbf{W}_{O(i, \cdot)}^{(b)}$ for $i \in \mathcal{W}_n$ be positive. The second step holds when $B \gtrsim M_1$. We also have

$$\begin{aligned}
 & \eta \frac{1}{|\mathcal{B}_b - \mathcal{B}_b^n|} \sum_{n \in \mathcal{B}_b - \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2} \right).
 \end{aligned} \tag{220}$$

For $i \in \mathcal{U}_n$, we also have

$$\begin{aligned}
 & \eta \frac{1}{|\mathcal{B}_b^n|} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, -\mathbf{q}^\top)^\top \\
 & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2} \right),
 \end{aligned} \tag{221}$$

$$\begin{aligned}
 & \eta \frac{1}{|\mathcal{B}_b - \mathcal{B}_b^n|} \sum_{n \in \mathcal{B}_b - \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, -\mathbf{q}^\top)^\top \\
 & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2} \right),
 \end{aligned} \tag{222}$$

if \mathbf{p}_j^n corresponds to label -1 in this task. For $i \notin \mathcal{W}_n(t) \cup \mathcal{U}_n(t)$, we have

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(t)}} (\mathbf{p}_j^{n \top}, \mathbf{0})^\top \leq \eta \sqrt{\frac{\log B}{B}} \frac{1}{a}. \tag{223}$$

Suppose that the conclusion holds when $t \leq t_0$. Then when $t = t_0 + 1$, we have that for $i \in \mathcal{W}_n$, $b \neq a$, and \mathbf{p}_{query}^n

corresponding to \mathbf{q} and $\boldsymbol{\mu}_a$,

$$\begin{aligned}
 & \eta \frac{1}{|\mathcal{B}_b^n|} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(t)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 &= \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^\top \mathbf{W}_K^\top \mathbf{W}_Q \mathbf{p}_{query}^n) (\delta \mathbf{p}_s^\top (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top) + \sum_{b=0}^{t-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right. \\
 & \quad \left. + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i, \cdot)}^{(b)} \right)^\top (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} + \frac{\alpha \eta}{2a} \cdot \eta \sum_{b=0}^{t_0} \delta(\beta^2 + 1) \frac{\eta m}{a^2} \left(1 + \frac{\eta^2 m}{a^2}\right)^b \\
 &= \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2} \cdot \frac{(1 + \frac{\eta^2 m}{a^2})^{t_0+1} - 1}{\frac{\eta^2 m}{a^2}}\right) \\
 &= \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2}\right)^{t_0+1},
 \end{aligned} \tag{224}$$

where the first inequality is by plugging the condition in the induction. The last two steps come from basic mathematical computation. Then,

$$\begin{aligned}
 & \eta \frac{1}{|\mathcal{B}_b^n|} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}^{(t)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 & \gtrsim \delta(\beta^2 + 1) \sum_{b=0}^{t_0+1} \frac{\alpha \eta}{2a} \left(1 + \frac{\eta^2 m}{a^2}\right)^b \\
 & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} (t_0 + 1),
 \end{aligned} \tag{225}$$

where lower bound in the last step is also a tight estimation of the second to last step if $\eta^2 T m / a^2 \ll 1$. Then, we have

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, \mathbf{q}^\top)^\top \\
 & \gtrsim \frac{1}{M_1} \cdot \delta \beta^2 (\beta^2 + 1) \frac{\alpha \eta}{2a} (t_0 + 1).
 \end{aligned} \tag{226}$$

By Lemma D.10, when $t \geq \Theta(1)$, we have

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b - \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, \mathbf{q}^\top)^\top \\
 & \gtrsim \frac{M_1 - 1}{M_1} \cdot \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} (t_0 + 1).
 \end{aligned} \tag{227}$$

Hence,

$$\begin{aligned}
 & \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, \mathbf{q}^\top)^\top \\
 & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} (t_0 + 1),
 \end{aligned} \tag{228}$$

which holds for $i \in \cup_{n \in [N]} \mathcal{W}_n = \mathcal{W}$. Meanwhile,

$$\eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\nu}_c^\top, \mathbf{0}^\top)^\top \leq \frac{1}{BT} \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, \mathbf{q}^\top)^\top. \tag{229}$$

For $i \in \mathcal{U}_n$ and \mathbf{p}_{query}^n corresponding to $-\mathbf{q}$ and $\boldsymbol{\mu}_a$, similarly to (226), (227), and (228), we have

$$\begin{aligned} & \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, -\mathbf{q}^\top)^\top \\ & \gtrsim \frac{1}{M_1} \cdot \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} (t_0 + 1), \end{aligned} \quad (230)$$

and when $t \geq \Theta(1)$,

$$\begin{aligned} & \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b - \mathcal{B}_b^n} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, -\mathbf{q}^\top)^\top \\ & \gtrsim \frac{M_1 - 1}{M_1} \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} (t_0 + 1), \end{aligned} \quad (231)$$

$$\begin{aligned} & \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, -\mathbf{q}^\top)^\top \\ & \gtrsim \delta(\beta^2 + 1) \frac{\alpha \eta}{2a} (t_0 + 1), \end{aligned} \quad (232)$$

which also holds for $i \in \cup_{n \in [N]} \mathcal{U}_n = \mathcal{U}$. Meanwhile,

$$\eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\nu}_c^\top, \pm \mathbf{q}^\top)^\top \leq \frac{1}{M_2} \eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_b^\top, \mathbf{q}^\top)^\top. \quad (233)$$

Then, for $i \in \mathcal{W}_n \cup \mathcal{U}_n$,

$$\|\mathbf{W}_{O(i, \cdot)}^{(t_0+1)}\| \gtrsim \sqrt{M_1} \delta(\beta^2 + 1)^{\frac{1}{2}} \frac{\alpha \eta}{2a} (t_0 + 1). \quad (234)$$

For $i \notin \mathcal{W} \cup \mathcal{U}$, we have

$$\eta \frac{1}{B} \sum_{b=0}^{t_0+1} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i, \cdot)}} (\boldsymbol{\mu}_a^\top, -\mathbf{q}^\top)^\top \leq \eta \sqrt{\frac{\log B(t_0 + 1)}{B(t_0 + 1)}} \frac{1}{a}. \quad (235)$$

□

E.8. Proof of Lemma D.9

Proof. We know that the Gaussian initialization of $\mathbf{W}_{O(i, \cdot)}^{(0)}$ generates a uniform distribution on the $d_{\mathcal{X}} - 1$ -sphere for the first $d_{\mathcal{X}}$ dimensions. Therefore,

$$\Pr(i \in \mathcal{W}_n) = A_{d_{\mathcal{X}}}^{cap}(\phi) / A_{d_{\mathcal{X}}}, \quad (236)$$

where $A_{d_{\mathcal{X}}}$ is the surface area of an $d_{\mathcal{X}} - 1$ -sphere. $A_{d_{\mathcal{X}}}^{cap}(\phi)$ is the surface area of a $d_{\mathcal{X}} - 1$ -spherical cap with ϕ as the colatitude angle. By Equation 1 in (Li, 2010), we have

$$\Pr(i \in \mathcal{W}_n) = \frac{1}{2} I_{\sin^2 \phi} \left(\frac{d_{\mathcal{X}} - 1}{2}, \frac{1}{2} \right) = \frac{\int_0^{\sin^2 \phi} t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt}{2 \int_0^1 t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt}, \quad (237)$$

where $I(\cdot, \cdot)$ is the regularized incomplete beta function. Since that

$$\phi \leq \pi/2 - \Theta(1/M_1), \quad (238)$$

to avoid concentration error of $\Theta(\sqrt{1/m})$ if $m \gtrsim M_1^2$, we have that when $d_{\mathcal{X}} = M_1 + M_2 = M = \Theta(M)$,

$$\begin{aligned}
 & \frac{\int_0^{\sin^2 \phi} t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt}{\int_0^1 t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt} \\
 & \geq \frac{\int_0^{\cos^2 1/M} t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt}{\int_0^1 t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt} \\
 & \geq 1 - \frac{\int_{1-1/M^2}^1 t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt}{\int_0^1 t^{\frac{d_{\mathcal{X}}-3}{2}} (1-t)^{-\frac{1}{2}} dt} \\
 & \geq 1 - \frac{\int_{1-1/M^2}^1 (1-t)^{-\frac{1}{2}} dt}{\int_{1-1/M}^1 \Theta(1) \cdot (1-t)^{-\frac{1}{2}} dt} \\
 & = 1 - \frac{\frac{2}{M}}{\Theta(1) \cdot (\frac{2}{\sqrt{M}} - \frac{2}{M})} \\
 & \geq \Theta(1),
 \end{aligned} \tag{239}$$

where the second inequality comes from that $\cos^2(1/M) = (1 + \cos(2/M))/2 \geq 1 - 1/M^2 \geq 1 - 1/M$, and the third to last step is by $(1 - 1/M)^{\frac{d_{\mathcal{X}}-3}{2}} \geq \Theta(1)$, and the last step is by $M \geq \Theta(1)$. For the second $d_{\mathcal{Y}}$ dimensions of $\mathbf{W}_{O(i,\cdot)}^{(0)}$, we can derive a similar result by replacing $d_{\mathcal{X}}$ with $d_{\mathcal{Y}}$ in (239). This implies that

$$|\mathcal{W}_n| \geq \Omega(1) \cdot \Omega(1) \cdot m \geq \Omega(m). \tag{240}$$

Likewise, the conclusion holds for \mathcal{U}_n . Since that $\mathcal{W} = \cup_{n \in [N]} \mathcal{W}_n$, $\mathcal{U} = \cup_{n \in [N]} \mathcal{U}_n$, we have

$$|\mathcal{W}|, |\mathcal{U}| \geq \Omega(m). \tag{241}$$

□

E.9. Proof of Lemma D.10

Proof. We prove this lemma in two steps. In the first step, we prove the conclusion by replacing \mathcal{W} with \mathcal{W}_n , and replacing \mathcal{U} with \mathcal{U}_n . We will also cover the proof of $\mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{W}_{O(i,\cdot)}^{(t)\top} > 0$ and $\sum_{b=0}^{t-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \mathbf{W}_{O(i,\cdot)}^{(t)\top} > 0$ in the induction as a support. In the second step, we prove the results for \mathcal{W} and \mathcal{U} .

(1) When $t = 0$. For any $i \in \mathcal{W}_n$, we have that by definition of \mathcal{W}_n

$$\mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{V}^n(0) > 0, \tag{242}$$

$$\mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{W}_{O(i,\cdot)}^{(t)\top} > 0. \tag{243}$$

Hence, the conclusion holds. When $t = 1$, we have

$$\begin{aligned}
 & \mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{V}^n(t) \\
 & = \mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{W}_{O(i,\cdot)}^{(t)\top} + \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) (\delta \mathbf{p}_s^n \top \mathbf{W}_{O(i,\cdot)}^{(t)\top} + \sum_{b=0}^{t-1} \eta (\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \\
 & \quad + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)}) \top \mathbf{W}_{O(i,\cdot)}^{(t)\top}).
 \end{aligned} \tag{244}$$

By (211) and definition of \mathbf{W}_n , we have

$$\begin{aligned} & \mathbf{W}_{O(i,\cdot)}^{(0)} \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (+z^n) a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)} \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K \top \mathbf{W}_Q \mathbf{p}_{query}^n) \geq 0] \\ & \cdot \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t-1)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t-1)} \top \mathbf{W}_Q^{(t-1)} \mathbf{p}_{query}^n) \\ & > 0. \end{aligned} \quad (245)$$

Hence,

$$\mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{W}_{O(i,\cdot)}^{(t)} \top = \sum_{b=0}^{t-1} \mathbf{W}_{O(i,\cdot)}^{(b)} \mathbf{W}_{O(i,\cdot)}^{(t)} \top > 0, \quad (246)$$

and

$$\sum_{b=0}^{t-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \mathbf{W}_{O(i,\cdot)}^{(t)} \top > 0. \quad (247)$$

By the gradient update when $t = 0$, we know that the largest component in the feature embedding is the IDR pattern for \mathbf{p}_{query}^n , and the label embedding is close to being in the direction of the label embedding of \mathbf{p}_{query}^n . Hence,

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)} \top \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \delta \mathbf{p}_s^n \top \mathbf{W}_{O(i,\cdot)}^{(t)} \top > 0. \quad (248)$$

Denote θ_n^i as the angle between the feature embeddings of $\mathbf{V}^n(0)$ and $\mathbf{W}_{O(i,\cdot)}^{(0)}$. Since that the feature embedding of $\mathbf{W}_{O(i,\cdot)}^{(0)}$ is initialized uniformed on the $d_{\mathcal{X}} - 1$ -sphere, we have $\mathbb{E}[\theta_n^i] = 0$. By Hoeffding's inequality (27), we have

$$\left\| \frac{1}{|\mathcal{W}_n|} \sum_{i \in \mathcal{W}_n} \theta_n^i - \mathbb{E}[\theta_n^i] \right\| = \left\| \frac{1}{|\mathcal{W}_n|} \sum_{i \in \mathcal{W}_n} \theta_n^i \right\| \leq \sqrt{\frac{\log M_1}{m}}, \quad (249)$$

with probability of at least $1 - M_1^{-10}$. When $m \gtrsim M^2 \log M_1$, we can obtain that

$$\left\| \frac{1}{|\mathcal{W}_n|} \sum_{i \in \mathcal{W}_n} \theta_n^i - \mathbb{E}[\theta_n^i] \right\| \leq \Theta\left(\frac{1}{M_1}\right). \quad (250)$$

Therefore, for $i \in \mathcal{W}_n$, as long as $m \gtrsim M_1^2 \log M_1$, we have

$$\mathbf{W}_{O(i,\cdot)}^{(0)} \sum_{b=0}^{t-1} \sum_{i \in \mathcal{W}_n} \mathbf{W}_{O(i,\cdot)}^{(b)} \top > 0. \quad (251)$$

Given $B \gtrsim M_1 \log M_1$, by Lemma D.6, and combining (251), we have that

$$\begin{aligned} & \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)} \top \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \sum_{b=0}^{t-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \right. \\ & \left. + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \right) \top \mathbf{W}_{O(i,\cdot)}^{(t)} \top \\ & > 0. \end{aligned} \quad (252)$$

Therefore, the conclusion holds when $t = 1$.

Suppose that the conclusion holds when $t \leq t_0$. When $t = t_0 + 1$, by (244), we can check that

$$\begin{aligned}
 & \mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{W}_{O(i,\cdot)}^{(t)\top} \\
 = & \mathbf{W}_{O(i,\cdot)}^{(0)} (\mathbf{W}_{O(i,\cdot)}^{(0)\top} + \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (+z^n) a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}] \sum_{s=1}^{l+1} (\mathbf{W}_V \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t-1)\top} \mathbf{W}_Q^{(t-1)} \mathbf{p}_{query}^n) \geq 0] \\
 & \cdot \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t-1)} \mathbf{p}_s^n) \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t-1)\top} \mathbf{W}_Q^{(t-1)} \mathbf{p}_{query}^n) \\
 & > 0 + 0 = 0,
 \end{aligned} \tag{253}$$

where the second 0 comes from (248) and the conditions that such conclusion in (253) holds when $t \leq t_0$. Combining (228) and the fact that the weighted summation of \mathbf{p}_s^n is close to be in the direction of $\boldsymbol{\mu}_j$ and \mathbf{q} in the feature label embeddings, respectively, where $\boldsymbol{\mu}_j$ is the IDR pattern of the \mathbf{p}_{query} , we have

$$\eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \delta \mathbf{p}_s^{n\top} \mathbf{W}_{O(i,\cdot)}^{(t)\top} > 0, \tag{254}$$

$$\begin{aligned}
 & \sum_{b=0}^{t-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \mathbf{W}_{O(i,\cdot)}^{(t)\top} \\
 = & (\sum_{b=0}^{t_0-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \eta \sum_{i \in \mathcal{W}_n} V_i(t_0) \mathbf{W}_{O(i,\cdot)}^{(t_0)}) (\mathbf{W}_{O(i,\cdot)}^{(t_0)} + \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i,\cdot)}^{(t_0)}} \Big|_{t=t_0})^\top \\
 > & 0 + (\sum_{b=0}^{t_0-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \eta \sum_{i \in \mathcal{W}_n} V_i(t_0) \mathbf{W}_{O(i,\cdot)}^{(t_0)}) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{\partial \ell(\tilde{\mathbf{P}}^n, z^n; \Psi)}{\partial \mathbf{W}_{O(i,\cdot)}^{(t_0)}} \Big|_{t=t_0}^\top \\
 = & (\sum_{b=0}^{t_0-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(t_0)} + \eta \sum_{i \in \mathcal{W}_n} V_i(t_0) \mathbf{W}_{O(i,\cdot)}^{(t_0)}) \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} (-z^n) a_i \mathbb{1}[\mathbf{W}_{O(i,\cdot)}^{(t_0)}] \sum_{s=1}^{l+1} (\mathbf{W}_V^{(t_0)} \mathbf{p}_s^n) \\
 & \cdot \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t_0)\top} \mathbf{W}_Q^{(t_0)} \mathbf{p}_{query}^n) \geq 0] \cdot \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^{n\top} \mathbf{W}_K^{(t_0)\top} \mathbf{W}_Q^{(t_0)} \mathbf{p}_{query}^n) (\delta(\mathbf{p}_s^{n\top}, \mathbf{0}^\top)^\top \\
 & + \sum_{b=0}^{t_0-1} \eta (\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)})^\top), \\
 & > 0,
 \end{aligned} \tag{255}$$

where the first step is by the formula of the gradient descent, and the second step is by $(\sum_{b=0}^{t_0-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \eta \sum_{i \in \mathcal{W}_n} V_i(t_0) \mathbf{W}_{O(i,\cdot)}^{(t_0)}) \mathbf{W}_{O(i,\cdot)}^{(t_0)} > 0$ from the induction steps. The last step comes from the fact that $\|\sum_{b=0}^{t_0-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)}\|^2 > 0$ and $\sum_{b=0}^{t_0-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} \cdot \eta \sum_{i \in \mathcal{W}_n} V_i(t_0) \mathbf{W}_{O(i,\cdot)}^{(t_0)\top} > 0$ by the induction, and that $V_i(t)$ for $i \notin \mathcal{W}_n$ is much smaller than that in \mathcal{W}_n given $B \gtrsim M_1$. Then,

$$\sum_{b=0}^{t-1} \eta (\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)})^\top \mathbf{W}_{O(i,\cdot)}^{(t)\top} > 0, \tag{256}$$

since the norm $\mathbf{W}_{O(i,\cdot)}^{(t)}$ for $i \notin \mathcal{W}_n$ is no larger than that for $i \in \mathcal{W}_n$. Combining (253), (254), and (256), we have

$$\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{V}^n(t) > 0. \tag{257}$$

Hence, we finish the induction.

(2) When $t \gtrsim \Theta(\beta)$, for $i \in \mathcal{W}$, by checking (244), we can deduce that

$$\mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{W}_{O(i,\cdot)}^{(t)\top} + \eta \frac{1}{B} \sum_{n \in \mathcal{B}_b} \frac{1}{a} \sum_{s=1}^{l+1} \text{softmax}(\mathbf{p}_s^n \top \mathbf{W}_K^{(t)\top} \mathbf{W}_Q^{(t)} \mathbf{p}_{query}^n) \delta \mathbf{p}_s^n \top \mathbf{W}_{O(i,\cdot)}^{(t)\top} > 0, \quad (258)$$

since that the accumulated label embedding term of $\mathbf{W}_{O(i,\cdot)}^{(t)}$ contributed positively to $\mathbf{p}_s^n \top \mathbf{W}_{O(i,\cdot)}^{(t)\top}$ and is larger than that of the feature embedding contribution by (154) and (155) (the gradient updates is close in the direction of the IDR pattern of \mathbf{p}_{query}^n when $m \gtrsim M_1^2$). Since that $\|\mathbf{W}_{O(i,\cdot)}^{(0)} (\boldsymbol{\mu}_j^\top, \mathbf{0}^\top)^\top\| \leq \beta \xi$ for any $j \in [M_1]$, the effect of $\mathbf{W}_{O(i,\cdot)}^{(0)} \mathbf{W}_{O(i,\cdot)}^{(t)\top}$ to the sign is much smaller than the remaining terms in (258). Hence, we show (258).

Then, since that the label embedding of $\mathbf{W}_{O(i,\cdot)}$, $\mathbf{W}_{O(j,\cdot)}$ are both close to \mathbf{q} for $i, j \in \mathcal{W}$, and that the feature embedding of $\mathbf{W}_{O(i,\cdot)}$, $i \in \mathcal{W}_n$ is close to the IDR pattern of \mathbf{p}_{query}^n , which is not the negative direction of the feature embedding of $\mathbf{W}_{O(j,\cdot)}$, $j \in \mathcal{W}_{n'}$, we have for $j \in \mathcal{W} \setminus \mathcal{W}_n$,

$$\sum_{b=0}^{t-1} \eta \sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)\top} \mathbf{W}_{O(j,\cdot)}^{(t)\top} > 0. \quad (259)$$

Given that $V_i(t)$ for $i \notin \mathcal{W}_n$ is much smaller than that in \mathcal{W}_n given $B \gtrsim M_1$ and the norm $\mathbf{W}_{O(i,\cdot)}^{(t)}$ for $i \notin \mathcal{W}_n$ is no larger than that for $i \in \mathcal{W}_n$, we have

$$\sum_{b=0}^{t-1} \eta \left(\sum_{i \in \mathcal{W}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \in \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)} + \sum_{i \notin \mathcal{W}_n \cup \mathcal{U}_n} V_i(b) \mathbf{W}_{O(i,\cdot)}^{(b)\top} \mathbf{W}_{O(i,\cdot)}^{(t)\top} \right) > 0. \quad (260)$$

Therefore, we can derive that for $i \in \mathcal{W}$,

$$\mathbf{W}_{O(i,\cdot)}^{(t)} \mathbf{V}^n(t) > 0. \quad (261)$$

□