
Learning and Transferring Sparse Contextual Bigrams with Linear Transformers

Yunwei Ren* Zixuan Wang* Jason D. Lee
Princeton University
{yunwei.ren, wangzx, jasonlee}@princeton.edu

Abstract

Transformers have excelled in natural language modeling and one reason behind this success is their exceptional ability to combine contextual informal and global knowledge. However, the theoretical basis remains unclear. In this paper, first we introduce the Sparse Contextual Bigram (SCB), a natural extension of the classical bigram model, where the next token’s generation depends on a sparse set of earlier positions determined by the last token. We then analyze the training dynamics and sample complexity of learning SCB using a one-layer linear transformer with a gradient-based algorithm. We show that when trained from scratch, the training process can be split into an initial sample-intensive stage where the correlation is boosted from zero to a nontrivial value, followed by a more sample-efficient stage of further improvement. Additionally, we prove that, provided a nontrivial correlation between the downstream and pretraining tasks, finetuning from a pretrained model allows us to bypass the initial sample-intensive stage. We also empirically demonstrate that our algorithm can outperform SGD in this setting and discuss its relationship with the usual softmax-based transformers.

1 Introduction

Transformers have played a central role in modern deep learning, achieving significant success across various fields, including language modeling (OpenAI, 2023), computer vision (Dosovitskiy et al., 2020), and natural sciences (Jumper et al., 2021). The core of transformers is the self-attention layer (Vaswani et al., 2017), which can attend to any subset of the input sequence to output a weighted linear combination of the (transformed) tokens.

Several capabilities of the transformers contribute to their success in language modeling. First, they can extract contextual information from the input token sequences, which is essential in some arithmetic tasks (Edelman et al., 2022; Liu et al., 2022; Nanda et al., 2023; Yao et al., 2021). In addition, transformers can memorize global in-domain knowledge (Petroni et al., 2019; Zhang et al., 2023a; Haviv et al., 2022; Carlini et al., 2021). These two abilities combined enable transformers to predict the next token based on the *in-context* information as well as *global* knowledge (OpenAI, 2023) acquired during training.

To theoretically understand how transformers learn both capabilities, we propose a minimalist data-generating model, the **Sparse Contextual Bigram** (SCB). This model builds on the classical bigram model and requires learning both contextual information and the (global) transition probabilities. Here, the next token depends on the transition matrix \mathbf{P} and a sparse set of prior tokens that is determined by the last token. In particular, SCB can be represented by a one-layer linear transformer — a simplified architecture that can serve as an abstraction for studying transformer optimization (Ahn et al., 2023), which makes it suitable for theoretical analysis.

*Equal contribution.

In this paper, we investigate the training dynamics and sample complexity of training a linear transformer to learn the SCB task using a stochastic gradient-based algorithm. Our contributions are summarized as follows:

- **Data model:** We introduce the **Sparse Contextual Bigram** (SCB) model, a simple task that requires the model to learn both in-context and global information.
- **Convergence:** We prove convergence guarantees for a one-layer linear transformer trained on with the nonconvex ℓ_1 -regularized MSE loss using preconditioned projected proximal descent, given a dataset sampled from the SCB model.
- **Sample Complexity:** Under mild conditions on the data distribution, initialization, and hyperparameters, we prove that our algorithm can recover the ground-truth with polynomial dependence on the sequence length T , number of states N , and the sparsity parameter $Q \ll T$. We show that the training first goes through an initial sample-intensive stage which boosts the signal with $\text{poly}(T)$ samples, followed by a more sample-efficient stage to achieve final convergence with $\text{poly}(N, Q)$ samples. We empirically verify that our gradient-based methods converge to the ground truth with a small batch size, while unregularized stochastic gradient descent fails due to the large variance.
- **Transfer Learning:** We prove that, when there is a nontrivial correlation between the pretraining and downstream tasks, we can transfer a pre-trained model to bypass the first sample intensive stage, so that our algorithm converges to the ground truth of the downstream task with only $\text{poly}(N, Q)$ samples.

1.1 Related works

Training dynamics of transformers. Several works have studied the learnability aspects of specific transformer architectures. Jelassi et al. (2022) demonstrated that a Vision Transformer (ViT) (Dosovitskiy et al., 2020) trained through GD, augmented with positional-embedding attention matrices, can effectively capture spatial structures. Li et al. (2023) investigated the sample complexity necessary to achieve good generalization performance on a similar ViT model. Tarzanagh et al. (2023) established a connection between the optimization landscape of self-attention and the formulation of a hard-margin Support Vector Machine (SVM) problem that separates and selects specific optimal tokens and established global convergence under strong assumptions. Tian et al. (2023a,b) provided insights into the training dynamics of the self-attention and MLP layers, respectively, although they did not establish convergence guarantees.

Another line of work focuses on the training dynamics of in-context learning. Mahankali et al. (2023) was among the first to introduce linear regression as an in-context learning task, while Zhang et al. (2023b) proved global convergence of gradient flow for a single-layer linear self-attention layer on this task. Huang et al. (2023) provided a convergence guarantee for a one-layer transformer with softmax attention on a similar task where the in-context tokens are drawn from a specific data distribution. Chen et al. (2024) generalized the single-task linear regression task to a multi-task setting and proved the global convergence of multi-head attention architecture using gradient flow on the population loss with specific initialization. In contrast, our work focuses on the language modeling ability of transformers instead of their in-context learning ability.

Several recent works analyzed transformers from a Markov chain perspective. Bietti et al. (2024) studied the in-context bigram (phrased as *induction head*) from an associative memory viewpoint. Nichani et al. (2024) proved that a simplified two-layer transformer can learn the induction head and generalize it to certain latent causal graphs. Edelman et al. (2024) further investigated training process on bigram and general n -gram tasks, and observed multi-phase dynamics. Makkuva et al. (2024) studied the loss landscape of transformers trained on sequences sampled from a single Markov Chain. Our SCB model extends the classical bigram models to allow context-dependent sparse attention on previous tokens.

Several works, including Tian et al. (2023a); Zhang et al. (2023b); Huang et al. (2023); Tarzanagh et al. (2023); Nichani et al. (2024); Kim and Suzuki (2024), and ours, use a similar reparameterization, consolidating the key and query matrices into a single matrix \mathbf{W} to simplify the dynamics of the training process. Most previous studies (Tian et al., 2023a; Zhang et al., 2023b; Huang et al., 2023; Tarzanagh et al., 2023; Nichani et al., 2024; Kim and Suzuki, 2024; Wang et al., 2024; Chen et al., 2024) uses population loss to simplify the analysis. In contrast, our work goes beyond the population

loss to analyze the sample complexity of the stochastic gradient descent dynamics. Although Li et al. (2023) also investigated the sample complexity on a different task, their model requires a pre-trained initialization, while our model is trained from scratch.

Transfer Learning. Transfer learning (Devlin et al., 2018) has gained significant attention in this deep learning era. From a theoretical perspective, several works have investigated the statistical guarantees of transfer learning from the representation learning perspective (Tripuraneni et al., 2020; Du et al., 2020; Arora et al., 2019; Hanneke et al., 2023). Recent studies on transfer learning mostly focus on linear models (Li et al., 2022; Tian and Feng, 2023; Fei and Li, 2021; Zhang et al., 2022; Ju et al., 2023; Dar and Baraniuk, 2022). For dynamics of transfer learning, Lampinen and Ganguli (2018) studied the behaviors of multi-layer linear networks in a teacher-student setting, while Dhifallah and Lu (2021) analyzed single-layer perceptrons. Damian et al. (2022) showed that a two-layer neural network can efficiently learn polynomials dependent on a few directions, enabling transfer learning.

To the best of our knowledge, this is the first work studying transfer learning for transformers. Moreover, unlike previous works that assume a shared structure between the pretraining and downstream tasks, we only require them to have a non-trivial correlation, which is a much weaker assumption.

1.2 Outline of this paper

In Section 2 we formalize the problem setup, including the SCB task, the transformer architecture, and the training algorithm. Section 3 consists of our main results, and we analyze the population dynamics to provide intuitions. Section 4 contains our transfer learning results. Experimental results can be found in Section 5.

2 Setup

In this section, we describe our data-generating model, the one-layer linear transformer architecture, and the training algorithm.

Notations. We use $[T]$ to denote the set $\{1, 2, \dots, T\}$. Matrices and vectors are denoted in upper-case bold letters ($\mathbf{A}, \mathbf{V}, \mathbf{\Delta}$, etc.) and lower-case bold letters (\mathbf{a}, \mathbf{q} , etc.), respectively. For norm, $\|\cdot\|$ denotes ℓ_2 norm and $\|\cdot\|_F$ denotes the Frobenius norm. Additionally, for $\boldsymbol{\mu} \in \mathbb{R}^N$, $\|\mathbf{A}\|_{\boldsymbol{\mu}}$ denotes $\boldsymbol{\mu}$ -norm for matrix $\mathbf{A} \in \mathbb{R}^{d \times N}$ for arbitrary d , which is defined as $\|\mathbf{A}\|_{\boldsymbol{\mu}}^2 := \text{Tr}(\mathbf{A} \text{diag}(\boldsymbol{\mu}) \mathbf{A}^\top)$. We use $\mathbb{1}\{\cdot\}$ to denote the indicator function. We use $\tilde{O}(\cdot)$ to hide logarithmic factors in the asymptotic notations.

2.1 Data-generating model: Sparse Contextual Bigram

The bigram model, where the next token depends only on the current one, is arguably one of the simplest language models. To learn this model, it suffices to learn the transition probabilities $\mathbf{P} \in \mathbb{R}^{N \times N}$ where $P_{n,m} = \mathbb{P}[X_{t+1} = n \mid X_t = m]$, which is achievable through a linear model (0-layer transformer).

A natural way to extend the classical bigram model is to allow the next token to depend on a context-dependent set of previous tokens. This extension can model situations such as generating the words after the phrase “by Theorem 3.2”, which requires us to retrieve the statement of “Theorem 3.2”. Here, we propose a simple extension of this type, which we call the **Sparse Contextual Bigram** (SCB). The contextual information is encoded by a sparse probability vector \mathbf{q} determined by the last token. To generate the next token, the model retrieves the tokens referenced by \mathbf{q} and applies the transition matrix \mathbf{P} (global knowledge) to one of them according to the distribution \mathbf{q} .

Formally, our data-generating model SCB can be described as follows. Let T be the sequence length and $[N]$ the vocabulary. Let $\mathbf{P} \in \mathbb{R}^{N \times N}$ be a transition matrix, with column¹ \mathbf{P}_k being the transition probability vector of token k . Suppose that $\boldsymbol{\mu} \in \mathbb{R}_{\geq 0}^N$ is the stationary distribution of \mathbf{P} .

Each input sequence consists of $T + 1$ tokens (x_1, \dots, x_{T+1}) , i.i.d. sampled from distribution $\boldsymbol{\mu}$. The output token (label) is generated as follows. For each $k \in [N]$, there is a probability vector $\mathbf{q}^{(k)} \in \mathbb{R}_{\geq 0}^T$ that represents the tokens the model needs to attend to when the last token x_{T+1} is k . For notational

¹This differs from the convention of the usual Markov Chain literature where the rows are the transition probability vectors. We use this convention as it is more compatible with our notations.

simplicity, we will write $\mathbf{x} = (x_1, \dots, x_T)$, $\mathbf{X} = (\mathbf{e}_{x_1}, \dots, \mathbf{e}_{x_T})$ and $\mathbf{Q} = (\mathbf{q}^{(1)}, \dots, \mathbf{q}^{(N)}) \in \mathbb{R}^{T \times N}$. When $x_{T+1} = k$, the output token x_o is sampled from the distribution

$$\mathbb{P}(x_o = n \mid x_{T+1} = k, \mathbf{x}) = \sum_{t=1}^T q_t^{(k)} P_{n, x_t}, \quad \forall n \in [N]. \quad (1)$$

In words, we first sample a position $s \in [T]$ according to $\mathbf{q}^{(k)}$ and then run one step of the Markov Chain \mathbf{P} from x_s to generate x_o . Note that this model can be represented by a one-layer linear transformer (see the next subsection for details). We make the following assumptions on SCB task.

Assumption 2.1 (Sparse Contextual Bigram, SCB). *In the SCB task, we assume the following:*

- (a) (*Q-sparse*) For some $Q \ll T$ and each of $\mathbf{q}^{(k)}$, at most Q entries are nonzero.
- (b) (*Well-conditioned*) There exists some constant $C \geq 1$ such that for every $k \in [N]$ and $t \in [T]$, $q_t^{(k)} \in [1/(CQ), C/Q]$ if it is nonzero, and $\mu_k \in [1/(CN), C/N]$.
- (c) (*Nontrivial transition*) $\|\mathbf{P}\|_{\mu}^2 - \|\boldsymbol{\mu}\|^2 \geq \|\boldsymbol{\mu}\|^2$.
- (d) (*Long sequence*) $T \geq (NQ)^{10}$.

Remark on condition (c). We say the transition \mathbf{P} is trivial if the transition probability vectors are all the same, i.e., $\mathbf{P} = \boldsymbol{\mu}\mathbf{1}^{\top}$. In this case, we have $\|\mathbf{P}\|_{\mu}^2 = \langle \boldsymbol{\mu}\mathbf{1}^{\top}, \boldsymbol{\mu}\boldsymbol{\mu}^{\top} \rangle = \|\boldsymbol{\mu}\|^2$. Requiring $\|\mathbf{P}\|_{\mu}^2 - \|\boldsymbol{\mu}\|^2 \geq \|\boldsymbol{\mu}\|^2$ rules out situations where \mathbf{P} is too close to the trivial one. Also, note that for any well-conditioned $\boldsymbol{\mu}$, we have $\|\boldsymbol{\mu}\|^2 \geq \Omega(1/N)$. ♣

In this work, we focus on the case where $(\mathbf{x}, x_{T+1}, x_o)$ are given as (one data point of) the training data with (x_1, \dots, x_{T+1}) i.i.d. sampled from μ . The SCB task can be extended to a sequence-to-sequence model: we drop x_1 and append x_o to get a new input sequence $(x_2, \dots, x_{T+1}, x_o)$, and then repeat the same sampling procedure to generate another token. This generates a sequence $(x_t)_{t=1}^{\infty}$ where $(x_{T+2}, x_{T+1}, \dots)$ are not independent, and this makes our model a true language model. We leave the study of the more complicated learning-from- $(x_t)_{t=1}^{\infty}$ task to future works.

2.2 Transformer architecture

Our learner model is a one-layer single-head linear transformer (Akyürek et al., 2022; Zhang et al., 2023b; Ahn et al., 2023). A general linear transformer can be expressed as: $\mathbf{F}(\mathbf{x}, x_{T+1}; \mathbf{V}, \mathbf{A}) = \mathbf{V}\mathbf{E}(\mathbf{E}^{\top}\mathbf{A}\mathbf{E})$, where \mathbf{E} is the embedding of the input tokens and positions, and \mathbf{A}, \mathbf{V} are the parameters of the attention and output layers, respectively. In our setting, we only need a simpler model:

$$\mathbf{F}(\mathbf{x}, x_{T+1}; \mathbf{V}, \mathbf{A}) := \mathbf{V}\mathbf{X}(\mathbf{I}_T\mathbf{A}\mathbf{e}_{x_{T+1}}) =: \mathbf{V}\mathbf{X}\mathbf{a}^{(x_{T+1})}, \quad (2)$$

where $\mathbf{V} \in \mathbb{R}^{N \times N}$ and $\mathbf{A} \in \mathbb{R}^{T \times N}$ are the trainable parameters, and $\mathbf{a}^{(k)}$ denotes the k -th column of \mathbf{A} . This model uses cross-attention (replacing the last \mathbf{E} with $\mathbf{e}_{x_{T+1}}$), uses only the positional embeddings together with the last token to compute the attention weights (replacing the second \mathbf{E} with \mathbf{I}_T), and discards the positional embeddings in the output layer (replacing the first \mathbf{E} with \mathbf{X}). This is equivalent to manually set certain blocks in the weight matrices to 0, which is a common practice in the theoretical literature to simplify the analysis (Nichani et al., 2024; Huang et al., 2023; Zhang et al., 2023b).

Note that our data-generating model (1) can be represented using (2) by setting $\mathbf{A} = \mathbf{Q}$ and $\mathbf{V} = \mathbf{P}$. We will show that a modified version of SGD can approximately recover this ground-truth model.

2.3 Training algorithm

We assume that the stationary distribution $\boldsymbol{\mu}$ and certain norms of the ground-truth \mathbf{P} and \mathbf{Q} are known when choosing the initialization and learning rate. The goal here is to recover \mathbf{P} and \mathbf{Q} . Our loss function is the ℓ_1 -regularized MSE loss. The standard way to optimize an ℓ_1 -regularized loss is to use the proximal gradient descent. We adopt this algorithm with several additional pre-conditioning and a projection step to ensure some basic properties.

Formally, let the per-sample loss be defined as

$$l(\mathbf{x}, x_{T+1}, x_o; \mathbf{V}, \mathbf{A}) := \frac{1}{2} \left\| \mathbf{e}_{x_o} - \mathbf{V} \mathbf{X} \mathbf{A} \mathbf{e}_{x_{T+1}} \right\|^2. \quad (3)$$

We initialize $\mathbf{A} = \mathbf{1}_T \mathbf{1}_N^\top / T$ to have uniform attention and $\mathbf{V} = \boldsymbol{\mu} \mathbf{1}^\top$ to be the trivial transition. At each step $\tau \geq 0$, we sample B_τ fresh samples $\{\mathbf{x}^{(i)}, x_{T+1}^{(i)}, x_o^{(i)}\}_{i=1}^{B_\tau}$ to form a mini-batch. The ℓ_1 -regularized mini-batch loss is defined as

$$l_{\text{reg}}^{(B_\tau, \lambda)} \left(\{\mathbf{x}^{(i)}, x_{T+1}^{(i)}, x_o^{(i)}\}_{i=1}^{B_\tau}; \mathbf{V}, \mathbf{A} \right) := \frac{1}{B_\tau} \sum_{i=1}^{B_\tau} l(\mathbf{x}^{(i)}, x_{T+1}^{(i)}, x_o^{(i)}; \mathbf{V}, \mathbf{A}) + \lambda \sum_{k=1}^N \left\| \mathbf{a}^{(k)} \right\|_1,$$

where $\lambda > 0$ is a parameter that controls the strength of regularization. Let $\nabla_{\mathbf{V}}^{(B_\tau)} l$ and $\nabla_{\mathbf{A}}^{(B_\tau)} l$ denote the mini-batch gradients of the original l w.r.t. \mathbf{V} and \mathbf{A} , respectively. We then define the preconditioned gradients as

$$\begin{aligned} \hat{\nabla}_{\mathbf{V}}^{(B_\tau)} l &:= \left(\mathbf{I}_N - \frac{\mathbf{1}_N \mathbf{1}_N^\top}{N} \right) \left(\nabla_{\mathbf{V}}^{(B_\tau)} l \right) \text{diag}(1/\boldsymbol{\mu}) \left(\mathbf{I}_N - \frac{\boldsymbol{\mu} \boldsymbol{\mu}^\top}{\|\boldsymbol{\mu}\|^2} \right), \\ \hat{\nabla}_{\mathbf{a}^{(k)}}^{(B_\tau)} l &:= \frac{1}{\mu_k} \left(\mathbf{I}_T - \frac{\mathbf{1}_T \mathbf{1}_T^\top}{T} \right) \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l \right), \quad \forall k \in [N]. \end{aligned} \quad (4)$$

Here, the $1/\boldsymbol{\mu}$ rescaling plays a role similar to importance sampling. We multiply $\nabla_{\mathbf{V}}^{(B_\tau)} l$ with $\mathbf{I} - \mathbf{1} \mathbf{1}^\top / N$ and $\mathbf{I} - \boldsymbol{\mu} \boldsymbol{\mu}^\top / \|\boldsymbol{\mu}\|^2$ to ensure at least $\mathbf{1}_N^\top \mathbf{V} = \mathbf{1}_N^\top$, $\mathbf{1}_T^\top \mathbf{A} = \mathbf{1}_N^\top$, and $\mathbf{V} \boldsymbol{\mu} = \boldsymbol{\mu}$ always hold throughout training. Note that we project each column of \mathbf{V} to the affine space $\{\mathbf{v} \in \mathbb{R}^T : \mathbf{1}^\top \mathbf{v} = 1\}$ instead of the probability simplex. This is sufficient for our analysis and is much easier to compute than the latter. We update the output layer using

$$\mathbf{V}_{\tau+1} = \mathbf{V}_\tau - \eta_V \hat{\nabla}_{\mathbf{V}}^{(B_\tau)} l, \quad (5)$$

where $\eta_V > 0$ is the step size. Now, consider the attention layer. Due to the existence of the ℓ_1 -regularization, the update rule becomes a simple variant of the standard proximal gradient descent. Formally, for step size $\eta_A > 0$, each $k \in [N]$ and $t \in [T]$, we have

$$\begin{aligned} \mathbf{a}_{\tau+1}^{(k,')} &= \mathbf{a}_\tau^{(k)} - \frac{\eta_A}{\mu_k} \nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l, && \text{(preconditioned GD step),} \\ \mathbf{a}_{t, \tau+1}^{(k,')} &= \begin{cases} a_{t, \tau+1}^{(k,')} - \lambda, & \text{if } a_{t, \tau+1}^{(k,')} \geq \lambda, \\ 0, & \text{if } \left| a_{t, \tau+1}^{(k,')} \right| \leq \lambda, \end{cases} && \text{(proximal step),} \\ \mathbf{a}_{\tau+1}^{(k)} &= \text{Proj}_{\{\mathbf{1}^\top \mathbf{a} = 1\}} \left(\mathbf{a}_{\tau+1}^{(k,')} \right) = \mathbf{a}_{\tau+1}^{(k,')} + \left(1 - \mathbf{1}^\top \mathbf{a}_{\tau+1}^{(k,')} \right) \frac{\mathbf{1}}{T}, && \text{(projection step).} \end{aligned} \quad (6)$$

For the proximal step, we will later show that no $a_t^{(k)}$ can ever become smaller than $-\lambda$, so it suffices to consider those two cases. During the proximal step, all small $a_t^{(k)}$ are set to 0, and λ is subtracted from all large coordinates. For notational simplicity, we define $\mathbf{g}_{\lambda, \tau}^{(k)} := -\eta_A^{-1} (\mathbf{a}_{\tau+1}^{(k)} - \mathbf{a}_\tau^{(k)})$ so that we can write the update as $\mathbf{a}_{\tau+1}^{(k)} = \mathbf{a}_\tau^{(k)} - \eta_A \mathbf{g}_{\lambda, \tau}^{(k)}$. We will choose $\lambda = 0$ in certain stages of training. In this case, (6) becomes the usual projected preconditioned gradient descent and we have

$$\mathbf{a}_{\tau+1}^{(k)} = \mathbf{a}_\tau^{(k)} - \eta_A \hat{\nabla}_{\mathbf{a}^{(k)}}^{(B_\tau)} l \quad (\text{when } \lambda = 0).$$

Our algorithm consists of three stages with different hyperparameters being used in different stages and certain rounding required between stages. The pseudocode is given in Algorithm 1 and more details on the projection/normalization steps are provided in Appendix E and F. When we train the model from scratch, all three stages are used and the initialization is $\mathbf{V}_0 = \boldsymbol{\mu} \mathbf{1}^\top$ and $\mathbf{A} = \mathbf{1}_T \mathbf{1}_N^\top / T$.

Transfer learning. When doing transfer learning, the initialization will be obtained from the weights of the pre-trained model and one step of gradient update. Then, we will run Algorithm 1 from Stage 2.

Algorithm 1 Projected preconditioned ℓ_1 -proximal gradient descent

Input: Stationary distribution μ ; initialization V_0, A_0 ; learning rates $\eta_A^{(i)}, \eta_V^{(i)}, i \in [3]$; threshold λ_0 ; regularization strength $\hat{\lambda}$; times $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$

Stage 1: Run (5) and (6) with $\eta_A = \eta_A^{(1)}, \eta_V = \eta_V^{(1)}, \lambda = 0$ for \mathcal{T}_1 steps;

Thresholding-projection: $\forall k \in [n], \hat{\mathbf{a}}^{(k)} = [a_t^{(k)} \mathbb{1}\{a_t^{(k)} \geq \lambda_0\}]_t, \mathbf{a}^{(k)} \leftarrow (\mathbf{I}_T - \mathbf{1}_T \mathbf{1}_T^\top / T) \hat{\mathbf{a}}^{(k)}$

Stage 2: Run (5) and (6) with $\eta_A = \eta_A^{(2)}, \eta_V = \eta_V^{(2)}, \lambda = \hat{\lambda}$ for $\mathcal{T}_2 - \mathcal{T}_1$ steps;

Thresholding-normalization: $\forall k \in [n], \hat{\mathbf{a}}^{(k)} = [a_t^{(k)} \mathbb{1}\{a_t^{(k)} \geq \Omega(1/Q)\}]_t, \mathbf{a}^{(k)} \leftarrow \hat{\mathbf{a}}^{(k)} / \mathbf{1}^\top \hat{\mathbf{a}}^{(k)}$

Stage 3: Run (5) and (6) with $\eta_A = 0, \eta_V = \eta_V^{(3)}, \lambda = 0$ for $\mathcal{T}_3 - \mathcal{T}_2$ steps;

Output: $A_{\mathcal{T}_3}, V_{\mathcal{T}_3}$.

3 Results for training from scratch

In this section, we consider the situations where we train the model from scratch, i.e., the initialization is $V_0 = \mu \mathbf{1}^\top$ and $A_0 = \mathbf{1} \mathbf{1}^\top / T$ and discuss the ideas of the proof of the following theorem.

Theorem 3.1 (Theorem G.1). *Let $\varepsilon > 0$ be our target accuracy and $\mathcal{T}_1 = \min\{\tau \geq 0 : \max\{\alpha_{V,\tau}, \alpha_{A,\tau}\} \geq \Theta(1/(QN))\}$. We can choose the hyperparameters in Algorithm 1 such that within $\text{poly}(N, Q, 1/\varepsilon, \log T)$ steps, we have $\|A - Q\|_\mu^2 \leq \varepsilon$ and $\|V - P\|_\mu^2 \leq \varepsilon$ with probability at least $1 - \delta$ and the numbers of samples used before and after \mathcal{T}_1 are $\text{poly}(T, \delta)$ and $\text{poly}(N, Q, 1/\varepsilon, \log T, \delta)$, respectively.*

The overall strategy is analyzing the population process and then controlling the distance between the mini-batch trajectory and the population process². In Section 3.1, we discuss the key properties of the population process that simplify the analysis. After that, we describe the dynamics of the algorithm and the signal-noise-ratio (SNR) in each of the three stages of Algorithm 1 in Section 3.2~3.4.

3.1 The population process

In this subsection, we analyze the behavior of the population process and the evolution of the signal-noise ratio. More details can be found in Appendix C, where the so-called population projected process are defined and rigorously analyzed.

For ease of presentation, we assume $\lambda = 0$ and access to the population loss $\mathcal{L} := \mathbb{E} l$. In other words, we consider the projected preconditioned gradient descent. By Lemma B.8, the dynamics of the population process is controlled by

$$\begin{aligned} V_{\tau+1} &= V_\tau - \eta_V \left(\|A\|_\mu^2 (V - \mu \mathbf{1}^\top) - \langle Q, A \rangle_\mu (P - \mu \mathbf{1}^\top) \right), \\ A_{\tau+1} &= A_\tau - \eta_A \left(\left(\|V\|_\mu^2 - \|\mu\|^2 \right) \left(\mathbf{a}^{(k)} - \frac{\mathbf{1}}{T} \right) - \left(\langle V, P \rangle_\mu - \|\mu\|^2 \right) \left(\mathbf{q}^{(k)} - \frac{\mathbf{1}}{T} \right) \right). \end{aligned} \quad (7)$$

One can prove via induction on τ that V (resp. A) always stays on the straight line crossing $\mu \mathbf{1}^\top$ and P (resp. $\mathbf{1} \mathbf{1}^\top / T$ and Q). In other words, there exists some time-dependent real numbers $\alpha_{V,\tau}, \alpha_{A,\tau}, \beta_{V,\tau} := 1 - \alpha_{V,\tau}, \beta_{A,\tau} := 1 - \alpha_{A,\tau}$ such that $V_\tau = \alpha_{V,\tau} P + \beta_{V,\tau} \mu \mathbf{1}^\top$ and $A_\tau = \alpha_{A,\tau} Q + \beta_{A,\tau} \mathbf{1} \mathbf{1}^\top / T$. The same calculation yields the following equations that govern the dynamics of α_V and α_A :

$$\begin{aligned} \alpha_{V,\tau+1} &= \alpha_{V,\tau} + \eta_V K_Q (1 - \alpha_A \alpha_V) \alpha_A + \eta_V \frac{1 - \alpha_V}{T}, \\ \alpha_{A,\tau+1} &= \alpha_{A,\tau} + \eta_A K_P (1 - \alpha_V \alpha_A) \alpha_V, \end{aligned}$$

where $\alpha_{V,0} = \alpha_{A,0} = 0$, $K_P := \|P\|_\mu^2 - \|\mu\|^2 \gtrsim 1/N$ and $K_Q := \|Q\|_\mu^2 - 1/T \gtrsim 1/Q$. Choose $\eta_V = \eta/K_Q$ and $\eta_A = \eta/K_P$, and we can write the above in matrix form as

$$\begin{bmatrix} \alpha_{V,\tau+1} \\ \alpha_{A,\tau+1} \end{bmatrix} = \eta(1 - \alpha_A \alpha_V) \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha_{V,\tau} \\ \alpha_{A,\tau} \end{bmatrix} + \frac{\eta}{K_Q} \begin{bmatrix} (1 - \alpha_V)/T \\ 0 \end{bmatrix}. \quad (8)$$

²Strictly speaking, what we actually control is the distance of the mini-batch trajectory to the subspace the population process lies. This allows us to prevent the potential exponential growth of the error caused by error compounding in the analysis. For details on this technique, see Appendix C.

Hence, in order to analyze the population process, it suffices to analyze the above 2-dimensional ODE. In what follows, when we say the signal, we usually refer to these α 's or some quantities whose size is proportional to them. In particular, as one can see from (8), the size of the expected gradients is proportional to α_V and/or α_A ³.

Note that when both α_V, α_A are still small, the population dynamics of α_V, α_A are a linear system with coefficient matrix $\eta \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ and drift $\begin{bmatrix} \eta T/K_Q \\ 0 \end{bmatrix}$. The drift term will provide a small initial signal that guides the process toward the correct direction and then the linear term will amplify this signal. Since the linear term is close to 0 at initial and the initial signal provided by the drift term has order $1/T$, we should expect that poly T samples are necessary to distinguish it from noises (Stage 1). After the signal becomes reasonably large, the first term will have order $1/\text{poly}(N, Q)$, and we can then rely on it (combined with the l_1 -regularization) instead of the drift term to learn the model (Stage 2).

3.2 Stage 1: boosting the signal

At initialization, we have $\alpha_V = \alpha_A = 0$. We define Stage 1 to be the phase until at least one of them has grown from 0 to some small $\sigma_1 = 1/\text{poly}(N, Q)$. Note that in this stage, the mini-batch version of (8) is approximately equivalent to

$$\begin{bmatrix} \alpha_{V, \tau+1} \\ \alpha_{A, \tau+1} \end{bmatrix} \approx \eta \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha_{V, \tau} \\ \alpha_{A, \tau} \end{bmatrix} + \frac{\eta}{K_Q} \begin{bmatrix} 1/T \\ 0 \end{bmatrix} + \boldsymbol{\varepsilon}_{\text{noise}} + \boldsymbol{\varepsilon}_{\text{approx}}, \quad (9)$$

where $\boldsymbol{\varepsilon}_{\text{noise}}$ and $\boldsymbol{\varepsilon}_{\text{approx}}$ represent the errors introduced by the difference between the mini-batch and population gradients, and the fact that we are not exactly on the population trajectory. If we had infinite amount of samples so that both $\boldsymbol{\varepsilon}_{\text{noise}}$ and $\boldsymbol{\varepsilon}_{\text{approx}}$ were 0, then the second term on the RHS of (9) could provide a small positive signal to α and the first term would quickly boost it to σ_1 within $\mathcal{T}_1 = \log(T)/\eta$ iterations. In order for the above analysis to work, we need both $\boldsymbol{\varepsilon}_{\text{noise}}$ and $\boldsymbol{\varepsilon}_{\text{approx}}$ to be at least $O(1)/T$ small. Since, unfortunately, $\boldsymbol{\varepsilon}_{\text{noise}}$ does not scale with $1/T$, we need poly(T) samples to ensure these conditions.

We conjecture that this poly(T) dependence is unavoidable (when only a polynomial amount of computing time is available). That is because around the initialization, the only signal comes from the second term and the first term amplifies whatever the second term provides, even if it has been corrupted by the errors. It either takes poly(T) fresh samples each step to reveal the signal or poly(T) steps (whence also poly(T) samples) for the random noises to (hopefully) cancel with each other.

3.3 Stage 2: learning the model

We know that at the end of Stage 1, at least one of α_V and α_A is $\sigma_1 = 1/\text{poly}(Q, N)$ large. Hence, one may expect that the signal is $1/\text{poly}(Q, N)$ large now so that we no longer need to make the noises $1/T$ small and therefore, only poly(N, Q) samples are needed. Unfortunately, this argument will not work directly, since the variance of the mini-batch gradients scales with T .⁴ Therefore, we still need poly(T) samples to reduce the squared ℓ_2 -norm of $\boldsymbol{\varepsilon}_{\text{noise}}$ from $\Omega(T)$ to $1/\text{poly}(Q, N)$. To address this issue, we introduce the l_1 -regularizer and use a variant of proximal gradient descent.

The idea is, while the concentration in the ℓ_2 sense is difficult, controlling the ℓ_∞ -error is easy as every entry of $\nabla_{\mathbf{a}^{(k)}} l$ is bounded whence subgaussian. As a result, we can make the coordinate-wise difference between the population and mini-batch gradients $1/\text{poly}(N, Q)$ small using only poly(N, Q) $\log T$ samples by a standard concentration argument. Moreover, we have (cf. the proof of Lemma E.3)

$$-\mathbb{E} \partial_{a_t^{(k)}} l = \mu_k \alpha_V K_P \left(q_t^{(k)} - \alpha_V a_t^{(k)} \right) = \mu_k \alpha_V K_P \times \begin{cases} \Omega(1/Q), & \text{if } q_t^{(k)} \neq 0, \\ O(\alpha_V a_t^{(k)}), & \text{if } q_t^{(k)} = 0. \end{cases} \quad (10)$$

Thus, as long as $\alpha_V \geq 1/\text{poly}(N, Q)$, the ℓ_∞ -norm of the gradient noise being small is enough to create a separation between those useful entries ($q_t^{(k)} \neq 0$) and useless entries ($q_t^{(k)} = 0$) and ensure the ℓ_2 -error of those Q useful entries is small.

³We will often drop the time subscript τ and write $\alpha_V := \alpha_{V, \tau}$ for simplicity.

⁴It is possible almost explicitly to compute the covariance matrix through some tedious calculation. Intuitively, the reason it scale with T is $\nabla_{\mathbf{a}^{(k)}} l$ has T entries with most of them almost uncorrelated in a certain sense.

The above analysis suggests removing all small entries from the gradient will work. Now, we claim that ℓ_1 -regularization and proximal gradient descent naturally implement this strategy, at least approximately. We believe softmax-based attention layers also automatically implement this strategy. See Section 5 for more discussion on the relationship between our model and softmax transformers.

Note that, at the end of Stage 1 and after the thresholding-projection step — which is approximately equivalent to running one proximal step first — we know that all useful $a_t^{(k)}$ are at least $\Omega(\alpha_V/Q) = 1/\text{poly}(N, Q)$, while all useless entries are of size $O(1)/T$. By our previous discussion, we know that if λ is chosen appropriately, with $\text{poly}(N, Q, \log T)$ samples, the gradients w.r.t. those useful entries can be made approximately correct, while the gradients w.r.t. those useless entries are much smaller than λ . Thus, after one gradient step, the absolute value of each of those useless $a_t^{(k)}$ is still much smaller than λ . As a result, they will be set to 0 in the proximal step (and to $O(1)/T$ after the projection step), which is equivalent to filtering out all those entries, up to a small bias. Therefore, the proximal gradient updates stay close to the population trajectory, and the growth of the signals α_A, α_V can be analyzed using the population dynamics.

We end Stage 2 when $(\alpha_V + \alpha_A)/2 \approx 1$. Similar to Stage 1, this also only takes $\mathcal{T}_2 = \tilde{O}(1/\eta)$ steps. We also show that the difference between the mini-batch trajectory and the “population trajectory” can decrease to a small value (cf. Lemma E.10). This allows us to decouple the error introduced by Stage 1 and the target accuracy. We defer the proof details to Appendix E.

3.4 Stage 3: final rounding and convergence

The purpose of Stage 3 is to fix a minor issue regarding $|\alpha_V - \alpha_A|$. Taylor expand (8) around $(\alpha_A, \alpha_V) = (1, 1)$ and one will notice that although $(\alpha_V + \alpha_A)/2$ can converge to 1 at a linear rate (and the approximation error also decreases exponentially fast), the convergence rate of $\alpha_A - \alpha_V$ is much slower, and the process will get stuck around $(1 + \delta, 1 - \delta)$ for some small nonzero δ , instead of converging to $(1, 1)$. To accelerate this process, we directly round \mathbf{A} via normalization, which is possible only after the approximation error becomes small in Stage 2. Then we freeze \mathbf{A} and train \mathbf{V} to the desired accuracy. More details about this stage can be found in Appendix F.

4 Results for transfer learning

The transferability of neural networks and transformers and their benefits have been widely observed and studied in both practice and theory. It is often assumed that the downstream and pretraining tasks share a common structure or representations/features, and these models can learn these common structures during training, and then leverage them in fine-tuning.

In this section, we offer a different perspective: as long as there is a (potentially small) nontrivial correlation between the pretraining and downstream tasks, the pretrained model can be used to provide a nonzero initial signal, allowing us to bypass the initial sample-intensive signal-boosting stage.

Formally, we consider the following setting. Let $\hat{\mathbf{P}}$ be the transition matrix of the pretraining task and (\mathbf{P}, \mathbf{Q}) the transition matrix and \mathbf{Q} -matrix of the downstream task. We still assume Assumption 2.1. In addition, we assume $\hat{\mathbf{P}}$ and \mathbf{P} share the same stationary distribution $\boldsymbol{\mu}$, $\|\hat{\mathbf{P}}\|_{\mu}^2 = \Theta(1) \|\mathbf{P}\|_{\mu}^2$ and $\langle \hat{\mathbf{P}}, \mathbf{P} \rangle_{\mu} - \|\boldsymbol{\mu}\|^2 \geq \|\boldsymbol{\mu}\|^2$. The last condition can be viewed as the transfer learning version of condition (c) of Assumption 2.1. Note that we allow the correlation between $\hat{\mathbf{P}}$ and \mathbf{P} to be as small as $o(1)$.

Theorem 4.1 (informal version of Theorem H.3). *Consider the above setting. Initialize $\mathbf{A} = \mathbf{1}\mathbf{1}^T/T$, $\mathbf{V} = \theta\hat{\mathbf{P}} + (1 - \theta)\boldsymbol{\mu}\mathbf{1}^T$ for some small $\theta > 0$, and run one step of gradient update on \mathbf{A} . Then, running Algorithm 1 from Stage 2 allows us to recover (\mathbf{P}, \mathbf{Q}) to ε -accuracy with high probability with $\text{poly}(N, Q, 1/\varepsilon, \log T)$ samples.*

To intuitively see why using $\text{poly}(N, Q, 1/\varepsilon, \log T)$ samples is possible, recall from (9) that the reason we need $\text{poly}(T)$ samples in Stage 1 is the signal is additive and has order $O(1)/T$, so we need the size of the noise to be at most $O(1)/T$. On the other hand, when we initialize $\mathbf{V} = \theta\hat{\mathbf{P}} + (1 - \theta)\boldsymbol{\mu}\mathbf{1}^T$, we have $\alpha_V \geq \Theta(\theta/(NK_P)) \gg 1/T$. Then we can rely on α_V , instead of the $1/T$ -sized additive signal, to boost the signal of α_A to $\omega(1)/T$ in one step, which leads to a sample complexity that depends on α_V instead of the $1/T$ -sized additive signal. Then, we can reuse the analysis of Stage 2 and 3 to show that the downstream (\mathbf{P}, \mathbf{Q}) can be approximately recovered using $\text{poly}(N, Q, 1/\varepsilon, \log T)$.

Note that unlike the case of training from scratch, when performing transfer learning, the initial approximation error $\|\Delta_V\|_\mu$, i.e., the distance between V and its population projection, can be much larger than the signal α_V , and it might seem unreasonable to expect that we can leverage the small signal in the presence of a large approximation error. To handle this issue, we show that the *influence* of $\|\Delta_V\|_\mu^2$ on the dynamics scales with α_A ($\approx \alpha_V$), which is small. In addition, we also show that as long as α_A is bounded away from 0 and the batch size is large, the approximation error will not grow. This allows us to ignore the approximation errors in the signal-boosting stage until we enter the regime of the Stage 2 analysis.

5 Experiments and relationship with softmax transformers

This section contains our experimental results. We also discuss the relationship between our linear model and the usual softmax transformers.

Experiment setup We use the same shallow transformer model (2) to train on the synthetic data. The data distribution follows the SCB model (1) with a randomly sampled transition matrix P together with its stationary μ , and the ground truth attention pattern Q . We choose the number of states $N = 3$, sparsity $Q = 2$, and the sequence length $T = 5000 \gg N, Q$. We use a batch size $B = 64$ to run the online projected proximal gradient descent with $\lambda = 1e-5$ and the vanilla SGD for $\mathcal{T} = 1000$ iterations. Through the signal boosting stage $\tau \in [0, 400]$, we use $\eta_1 = 0.01$ to accelerate the process. After $\tau > 400$, we use $\eta_2 = 0.005$ for further improvement. For SGD, we add another set of experiments with $\eta'_2 = 0.001$ to prevent potential instability. For more details, see Appendix I.

5.1 Convergence

Our experiments (cf. Fig. 1) show that after switching to proximal gradient descent after Stage 1 (the signal-boosting stage), both $\|P - V\|_\mu$ and $\|A - Q\|_\mu$ decrease faster than SGD. The final distance to the ground-truth after normalization gets close to 0, and the similarity between the ground truth and parameters quickly converges close to 1. In comparison, SGD struggles to converge with the same small batch size and large learning rate, while the convergence rate is too slow when a smaller learning rate is used. This phenomenon verifies our theory that the variance of the original stochastic gradient will be too large for SGD to converge when $T \gg Q, N$, while proximal gradient descent with an ℓ_1 regularizer can resolve this issue.




Figure 1: **Convergence analysis:** We plot the distance to the ground truth $\|V - P\|_\mu, \|A - Q\|_\mu$ in different settings. After stage 1 ends at $\tau = 400$ (when $\alpha_A, \alpha_V \approx 0.1$), we use vanilla SGD and our proximal gradient method to train the transformer. Compared with SGD, the ℓ_1 regularized proximal gradient descent quickly converges, and the final solution (the star) recovers the ground truth. SGD either suffers from the large gradient variance (when η_2 is large) or a slow convergence rate (small η'_2).

5.2 Relationship between our model and softmax transformers

We claim that they have our linear model and softmax transformers have qualitatively similar behaviors: there will be a sample-intensive initial stage, and after the model and the target have a nontrivial correlation, proximal gradient descent/SGD will become much more sample efficient.

For ease of presentation, in the following, we will assume $N = 1$, write $\mathbf{a} := \mathbf{a}^{(1)}$, and assume the ground-truth $\mathbf{q} := \mathbf{q}^{(1)}$ is $\mathbf{e}_1 = (1, 0, \dots, 0)$. Most of our argument below can be generalized to the general setting at least at a heuristic level. Recall that our linear model is $f(\mathbf{X}; V, \mathbf{a}) = V\mathbf{X}\mathbf{a}$. By a softmax transformer, we mean the model $f_\sigma(\mathbf{X}; V, \mathbf{w}) = V\mathbf{X}\sigma(\mathbf{w}) =: V\mathbf{X}\mathbf{a}_\sigma$ where σ is the softmax function and $\mathbf{w} \in \mathbb{R}^T$ is the trainable first-layer weights.

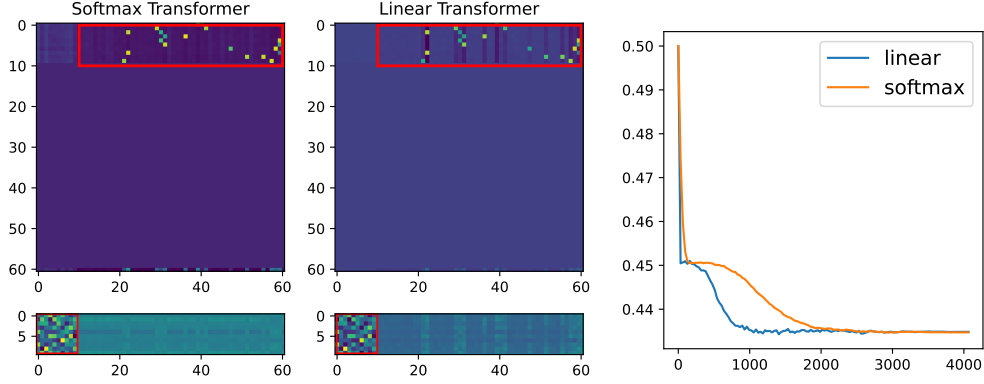


Figure 2: **Similarity between the softmax and linear attention.** We train two transformers with (1) (Left) *softmax* attention and (2) (Middle) *linear* attention layer on the SCB tasks with the same ground-truth ($T = 50, N = 10, Q = 2$). The attention pattern and the value matrix (learned transition matrix) are very similar (left two plots) and they converge to approximately the same loss (right plot).

Let l denote the (per-sample) loss. We have $\nabla_w l(f_\sigma(X)) = (\text{diag}(\mathbf{a}_\sigma) - \mathbf{a}_\sigma \mathbf{a}_\sigma^\top) (\mathbf{V}\mathbf{X})^\top \nabla l(f_\sigma(X))$. As a result, the dynamics of the attention weights are controlled by

$$\mathbf{a}(\tau + 1) \approx \mathbf{a}(\tau) - \eta \left(\mathbf{I} - \frac{\mathbf{1}\mathbf{1}^\top}{T} \right) (\mathbf{V}\mathbf{X})^\top \nabla l(f(X)), \quad \text{in our linear model,}$$

$$\mathbf{a}_\sigma(\tau + 1) \approx \mathbf{a}_\sigma(\tau) - \eta (\text{diag}(\mathbf{a}_\sigma) - \mathbf{a}_\sigma \mathbf{a}_\sigma^\top)^2 (\mathbf{V}\mathbf{X})^\top \nabla l(f_\sigma(X)), \quad \text{in softmax transformers.}$$

In other words, the main difference is that there will be a preconditioning matrix $(\text{diag}(\mathbf{a}) - \mathbf{a}\mathbf{a}^\top)^2$ in the dynamics of softmax transformers.

Near initialization, i.e., when the attention pattern is still close to the uniform attention, we have $(\text{diag}(\mathbf{a}_\sigma) - \mathbf{a}_\sigma \mathbf{a}_\sigma^\top)^2 \approx \frac{1}{T^2} \left(\mathbf{I} - \frac{\mathbf{1}\mathbf{1}^\top}{T} \right)$. In other words, our linear model and softmax transformers are approximately equivalent up to a change in learning rates.

Now, suppose that there is a nontrivial correlation between \mathbf{a}_σ and $\mathbf{q} = \mathbf{e}_1$, say, $a_{\sigma,1}$ is a small constant while all other entries are $O(1/T)$. In this case, we have $(\text{diag}(\mathbf{a}_\sigma) - \mathbf{a}_\sigma \mathbf{a}_\sigma^\top)^2 \approx a_{\sigma,1}(1 - a_{\sigma,1})\mathbf{e}_1\mathbf{e}_1^\top + O(1/T)$. Effectively, softmax transformers automatically adjust the learning rate according to $a_{\sigma,t}$ and roughly ignore those positions with a small attention weight to stabilize the gradients. Note that this is also what ℓ_1 -regularization does in our algorithm. In fact, mimicking this behavior is one of the motivations of using ℓ_1 -regularization in our linear setting. We run further experiments to highlight the resemblance between softmax attention and our linear attention model (Figure 2).

6 Conclusion and discussion

In this paper, we propose the Sparse Contextual Bigram (SCB) model, which is a natural extension of the bigram model, that requires both contextual and global information. Then, we analyze the problem of learning a SCB model using a one-layer linear transformer and a gradient-based algorithm. We prove quantitative bounds on the convergence rate and the sample complexity. In particular, we show when trained from scratch, the training process can be split into two stages, where the first stage uses a lot of samples to boost the signal from zero to a nontrivial value, while the second stage is much more sample-efficient. Then, we consider the problem in a transfer learning setting and prove that when there is a nontrivial correlation between the pretraining and downstream tasks, the first sample intensive stage can be bypassed.

Our data-generating model and results also lead to some interesting future directions. For example, can we improve the sample complexity of the first stage? What can we gain if the datapoints are sequences generated by repeatedly applying the SCB model?

Acknowledgement

JDL acknowledges support of the NSF CCF 2002272, NSF IIS 2107304, and NSF CAREER Award 2144994.

References

- OpenAI. Gpt-4 technical report, 2023.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- John M. Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Zidek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A A Kohl, Andy Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with alphafold. *Nature*, 596:583 – 589, 2021. URL <https://api.semanticscholar.org/CorpusID:235959867>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Benjamin L Edelman, Surbhi Goel, Sham Kakade, and Cyril Zhang. Inductive biases and variable creation in self-attention mechanisms. In *International Conference on Machine Learning*, pages 5793–5831. PMLR, 2022.
- Bingbin Liu, Jordan T Ash, Surbhi Goel, Akshay Krishnamurthy, and Cyril Zhang. Transformers learn shortcuts to automata. *arXiv preprint arXiv:2210.10749*, 2022.
- Neel Nanda, Lawrence Chan, Tom Liberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. *arXiv preprint arXiv:2301.05217*, 2023.
- Shunyu Yao, Binghui Peng, Christos Papadimitriou, and Karthik Narasimhan. Self-attention networks can process bounded hierarchical languages. *arXiv preprint arXiv:2105.11115*, 2021.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*, 2019.
- Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. Counterfactual memorization in neural language models. *Advances in Neural Information Processing Systems*, 36:39321–39362, 2023a.
- Adi Haviv, Ido Cohen, Jacob Gidron, Roei Schuster, Yoav Goldberg, and Mor Geva. Understanding transformer memorization recall through idioms. *arXiv preprint arXiv:2210.03588*, 2022.
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650. USENIX Association, August 2021. ISBN 978-1-939133-24-3. URL <https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting>.
- Kwangjun Ahn, Xiang Cheng, Minhak Song, Chulhee Yun, Ali Jadbabaie, and Suvrit Sra. Linear attention is (maybe) all you need (to understand transformer optimization). *arXiv preprint arXiv:2310.01082*, 2023.
- Samy Jelassi, Michael Sander, and Yuanzhi Li. Vision transformers provably learn spatial structure. *Advances in Neural Information Processing Systems*, 35:37822–37836, 2022.
- Hongkang Li, Meng Wang, Sijia Liu, and Pin-Yu Chen. A theoretical understanding of shallow vision transformers: Learning, generalization, and sample complexity. *arXiv preprint arXiv:2302.06015*, 2023.

- Davoud Ataee Tarzanagh, Yingcong Li, Christos Thrampoulidis, and Samet Oymak. Transformers as support vector machines. *arXiv preprint arXiv:2308.16898*, 2023.
- Yuandong Tian, Yiping Wang, Beidi Chen, and Simon Du. Scan and snap: Understanding training dynamics and token composition in 1-layer transformer. *arXiv preprint arXiv:2305.16380*, 2023a.
- Yuandong Tian, Yiping Wang, Zhenyu Zhang, Beidi Chen, and Simon Du. Joma: Demystifying multilayer transformers via joint dynamics of mlp and attention. *arXiv preprint arXiv:2310.00535*, 2023b.
- Arvind Mahankali, Tatsunori B Hashimoto, and Tengyu Ma. One step of gradient descent is provably the optimal in-context learner with one layer of linear self-attention. *arXiv preprint arXiv:2307.03576*, 2023.
- Ruiqi Zhang, Spencer Frei, and Peter L Bartlett. Trained transformers learn linear models in-context. *arXiv preprint arXiv:2306.09927*, 2023b.
- Yu Huang, Yuan Cheng, and Yingbin Liang. In-context convergence of transformers. *arXiv preprint arXiv:2310.05249*, 2023.
- Siyu Chen, Heejune Sheen, Tianhao Wang, and Zhuoran Yang. Training dynamics of multi-head softmax attention for in-context learning: Emergence, convergence, and optimality. *arXiv preprint arXiv:2402.19442*, 2024.
- Alberto Bietti, Vivien Cabannes, Diane Bouchacourt, Herve Jegou, and Leon Bottou. Birth of a transformer: A memory viewpoint. *Advances in Neural Information Processing Systems*, 36, 2024.
- Eshaan Nichani, Alex Damian, and Jason D Lee. How transformers learn causal structure with gradient descent. *arXiv preprint arXiv:2402.14735*, 2024.
- Benjamin L Edelman, Ezra Edelman, Surbhi Goel, Eran Malach, and Nikolaos Tsilivis. The evolution of statistical induction heads: In-context learning markov chains. *arXiv preprint arXiv:2402.11004*, 2024.
- Ashok Vardhan Makkuva, Marco Bondaschi, Adway Girish, Alliot Nagle, Martin Jaggi, Hyeji Kim, and Michael Gastpar. Attention with markov: A framework for principled analysis of transformers via markov chains. *arXiv preprint arXiv:2402.04161*, 2024.
- Juno Kim and Taiji Suzuki. Transformers learn nonlinear features in context: Nonconvex mean-field dynamics on the attention landscape. *arXiv preprint arXiv:2402.01258*, 2024.
- Zixuan Wang, Stanley Wei, Daniel Hsu, and Jason D Lee. Transformers provably learn sparse token selection while fully-connected nets cannot. *arXiv preprint arXiv:2406.06893*, 2024.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. *Advances in neural information processing systems*, 33:7852–7862, 2020.
- Simon S Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*, 2020.
- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- Steve Hanneke, Samory Kpotufe, and Yasaman Mahdaviyeh. Limits of model selection under transfer learning. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 5781–5812. PMLR, 2023.
- Sai Li, T Tony Cai, and Hongzhe Li. Transfer learning for high-dimensional linear regression: Prediction, estimation and minimax optimality. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):149–173, 2022.

- Ye Tian and Yang Feng. Transfer learning under high-dimensional generalized linear models. *Journal of the American Statistical Association*, 118(544):2684–2697, 2023.
- Zhe Fei and Yi Li. Estimation and inference for high dimensional generalized linear models: A splitting and smoothing approach. *Journal of Machine Learning Research*, 22(58):1–32, 2021.
- Xuhui Zhang, Jose Blanchet, Soumyadip Ghosh, and Mark S Squillante. A class of geometric structures in transfer learning: Minimax bounds and optimality. In *International Conference on Artificial Intelligence and Statistics*, pages 3794–3820. PMLR, 2022.
- Peizhong Ju, Sen Lin, Mark S Squillante, Yingbin Liang, and Ness B Shroff. Generalization performance of transfer learning: Overparameterized and underparameterized regimes. *arXiv preprint arXiv:2306.04901*, 2023.
- Yehuda Dar and Richard G Baraniuk. Double double descent: On generalization errors in transfer learning between linear regression tasks. *SIAM Journal on Mathematics of Data Science*, 4(4): 1447–1472, 2022.
- Andrew K Lampinen and Surya Ganguli. An analytic theory of generalization dynamics and transfer learning in deep linear networks. *arXiv preprint arXiv:1809.10374*, 2018.
- Oussama Dhifallah and Yue M Lu. Phase transitions in transfer learning for high-dimensional perceptrons. *Entropy*, 23(4):400, 2021.
- Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pages 5413–5452. PMLR, 2022.
- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. What learning algorithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*, 2022.

A Limitation

In this section, we briefly discuss the limitation of this work.

First, we consider one-layer single-head linear transformers with certain (blocks of the) weights merged or fixed. Though this simplification are widely used in theoretical works and linear and nonlinear transformers share some training behaviors (Ahn et al., 2023), this architecture is still very far away from the transformers used in practice.

We also use a non-standard training algorithm that has several manually separated stages. Some parts of the modification are made to address certain issues of linear transformers, while the other are made to simplify the analysis. It would be interesting (and more challenging) to consider more natural/practical training algorithms.

Finally, for our data-generating model, we only use it to generate one next token, instead of repeatedly apply SCB on the previous generated results to obtain a long sequence. In our setting, the contextual tokens are independent. While this simplifies the analysis, it deviates from how natural language works.

B Probabilities, expectations, and variances

We collect in this section closed-form formulas for the probabilities of certain events, and the expectations and variances of some random vectors of interest. All proofs are deferred to the end of this section.

B.1 Probabilities

Lemma B.1. For any $t \in [T]$ and $k, n, m \in [N]$, we have

$$\mathbb{P}(x_o = n, x_t = m \mid x_{T+1} = k) = q_t^{(k)} P_{n,m} \mu_m + (1 - q_t^{(k)}) \mu_n \mu_m.$$

Lemma B.2. For any $s \neq t \in [T]$, $k, n, m, l \in [N]$, we have

$$\mathbb{P}(x_o = n \mid x_{T+1} = k, x_s = m, x_t = l) = q_s^{(k)} P_{n,m} + q_t^{(k)} P_{n,l} + (1 - q_s^{(k)} - q_t^{(k)}) \mu_n.$$

B.2 Gradient and Expectations

Lemma B.3. Suppose the last input token \mathbf{x}_{T+1} and $\mathbf{a} := \mathbf{A}\mathbf{x}_{T+1}$. The gradients of the objective are

$$\begin{aligned} \nabla_{\mathbf{V}} l &= (\mathbf{V}\mathbf{X}\mathbf{a} - \mathbf{e}_{x_o}) (\mathbf{X}\mathbf{a})^\top, \\ \nabla_{\mathbf{a}^{(k)}} l &= \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}\mathbf{X})^\top (\mathbf{V}\mathbf{X}\mathbf{a} - \mathbf{e}_{x_o}) \end{aligned}$$

Lemma B.4. For any $k \in [N]$ and $s, t \in [T]$, we have

$$\mathbb{E}^{(k)}[\mathbf{e}_{x_s} \mathbf{e}_{x_t}^\top] = \mathbb{E}[\mathbf{e}_{x_s} \mathbf{e}_{x_t}^\top] = \begin{cases} \text{diag}(\boldsymbol{\mu}), & s = t, \\ \boldsymbol{\mu}\boldsymbol{\mu}^\top, & s \neq t. \end{cases}$$

Lemma B.5. For any $\mathbf{V} \in \mathbb{R}^{N \times N}$ with $\mathbf{V}\boldsymbol{\mu} = \boldsymbol{\mu}$ and $s, t \in [T]$, we have

$$\sum_{n=1}^N \mathbb{E}[V_{n,x_s} V_{n,x_t}] = \begin{cases} \|\mathbf{V}\boldsymbol{\mu}\|_\mu^2, & s = t, \\ \|\boldsymbol{\mu}\|_\mu^2, & s \neq t. \end{cases}$$

Lemma B.6. For any $\mathbf{V} \in \mathbb{R}^{N \times N}$ with $\mathbf{V}\boldsymbol{\mu} = \boldsymbol{\mu}$ and $t \in [T]$, we have

$$\mathbb{E}^{(k)} V_{x_o, x_t} = q_t^{(k)} \langle \mathbf{V}, \mathbf{P} \rangle_\mu + (1 - q_t^{(k)}) \|\boldsymbol{\mu}\|_\mu^2.$$

Lemma B.7 (Expected gradients).

$$\mathbb{E} \nabla_{\mathbf{V}} l = \|\mathbf{A}\|_\mu^2 \mathbf{V} \text{diag}(\boldsymbol{\mu}) + (1 - \|\mathbf{A}\|_\mu^2) \boldsymbol{\mu}\boldsymbol{\mu}^\top$$

$$\begin{aligned}
& - \langle \mathbf{Q}, \mathbf{A} \rangle_{\mu} \mathbf{P} \text{diag}(\boldsymbol{\mu}) - \left(1 - \langle \mathbf{Q}, \mathbf{A} \rangle_{\mu}\right) \boldsymbol{\mu} \boldsymbol{\mu}^{\top}, \\
\mathbb{E} \nabla_{\mathbf{a}^{(k)}} l &= \mu_k \left(\|\mathbf{V}\|_{\mu}^2 - \|\boldsymbol{\mu}\|^2 \right) \mathbf{a}^{(k)} + \mu_k \mathbf{1} \|\boldsymbol{\mu}\|^2 \\
& - \mu_k \mathbf{q}^{(k)} \langle \mathbf{V}, \mathbf{P} \rangle_{\mu} - \mu_k \left(\mathbf{1} - \mathbf{q}^{(k)} \right) \|\boldsymbol{\mu}\|^2.
\end{aligned}$$

Lemma B.8 (Expected preconditioned gradients).

$$\begin{aligned}
\mathbb{E} \hat{\nabla}_{\mathbf{V}} l &= \|\mathbf{A}\|_{\mu}^2 (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^{\top}) - \langle \mathbf{Q}, \mathbf{A} \rangle_{\mu} (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^{\top}), \\
\mathbb{E} \hat{\nabla}_{\mathbf{a}^{(k)}} l &= \left(\|\mathbf{V}\|_{\mu}^2 - \|\boldsymbol{\mu}\|^2 \right) \left(\mathbf{a}^{(k)} - \frac{\mathbf{1}}{T} \right) - \left(\langle \mathbf{V}, \mathbf{P} \rangle_{\mu} - \|\boldsymbol{\mu}\|^2 \right) \left(\mathbf{q}^{(k)} - \frac{\mathbf{1}}{T} \right).
\end{aligned}$$

B.3 Deferred proofs of this section

B.3.1 Probabilities

Proof of Lemma B.1. For notational simplicity, define $\mathbf{x}_{-t} = (x_1, \dots, x_{t-1}, x_{t+1}, \dots, x_T)$. We compute

$$\begin{aligned}
& \mathbb{P}(x_o = n, x_t = m \mid x_{T+1} = k) \\
&= \sum_{\hat{\mathbf{m}} \in [N]^{T-1}} \mathbb{P}(x_o = n, x_t = m, \mathbf{x}_{-t} = \hat{\mathbf{m}} \mid x_{T+1} = k) \\
&= \sum_{\hat{\mathbf{m}} \in [N]^{T-1}} \mathbb{P}(x_o = n \mid x_t = m, \mathbf{x}_{-t} = \hat{\mathbf{m}}, x_{T+1} = k) \\
& \quad \times \mathbb{P}(x_t = m, \mathbf{x}_{-t} = \hat{\mathbf{m}} \mid x_{T+1} = k).
\end{aligned}$$

By the independence assumption, we have $\mathbb{P}(x_t = m, \mathbf{x}_{-t} = \hat{\mathbf{m}} \mid x_{T+1} = k) = \mathbb{P}(x_t = m, \mathbf{x}_{-t} = \hat{\mathbf{m}}) = \mathbb{P}(x_t = m) \mathbb{P}(\mathbf{x}_{-t} = \hat{\mathbf{m}})$. For the first factor, we have $\mathbb{P}(x_o = n \mid x_t = m, \mathbf{x}_{-t} = \hat{\mathbf{m}}, x_{T+1} = k) = \mathcal{Q}_t^{(k)} P_{n,m} + \sum_{s \neq t} \mathcal{Q}_s^{(k)} P_{n, \hat{m}_s}$. Therefore,

$$\begin{aligned}
& \mathbb{P}(x_o = n, x_t = m \mid x_{T+1} = k) \\
&= \sum_{\hat{\mathbf{m}} \in [T]^{N-1}} \left(\mathcal{Q}_t^{(k)} P_{n,m} + \sum_{s \neq t} \mathcal{Q}_s^{(k)} P_{n, \hat{m}_s} \right) \mathbb{P}(x_t = m) \mathbb{P}(\mathbf{x}_{-t} = \hat{\mathbf{m}}) \\
&= \left(\mathcal{Q}_t^{(k)} P_{n,m} + \sum_{s \neq t} \mathcal{Q}_s^{(k)} \sum_{\hat{\mathbf{m}} \in [T]^{N-1}} P_{n, \hat{m}_s} \mathbb{P}(\mathbf{x}_{-t} = \hat{\mathbf{m}}) \right) \mu_m.
\end{aligned}$$

Note that for any $s \neq t$

$$\sum_{\hat{\mathbf{m}} \in [T]^{N-1}} P_{n, \hat{m}_s} \mathbb{P}(\mathbf{x}_{-t} = \hat{\mathbf{m}}) = \sum_{\hat{\mathbf{m}}_{-s} \in [T]^{N-2}} \left(\sum_{\hat{m}_s=1}^N P_{n, \hat{m}_s} \mu_{\hat{m}_s} \right) \mathbb{P}(\mathbf{x}_{-s, -t} = \hat{\mathbf{m}}_{-s}) = \mu_n.$$

Thus,

$$\begin{aligned}
\mathbb{P}(x_o = n, x_t = m \mid x_{T+1} = k) &= \left(q_t^{(k)} P_{n,m} + \sum_{s \neq t} q_s^{(k)} \mu_n \right) \mu_m \\
&= q_t^{(k)} P_{n,m} \mu_m + \left(1 - q_t^{(k)} \right) \mu_n \mu_m.
\end{aligned}$$

□

Proof of Lemma B.2. For notational simplicity, let $\mathbf{x}_{-s, -t}$ denote the vector obtained by removing the s, t coordinates from \mathbf{x} . Then, we compute

$$\begin{aligned}
& \mathbb{P}(x_o = n \mid x_{T+1} = k, x_s = m, x_t = l) \\
&= \sum_{\hat{\mathbf{m}} \in [N]^{T-2}} \mathbb{P}(x_o = n, \mathbf{x}_{-s, -t} = \hat{\mathbf{m}} \mid x_{T+1} = k, x_s = m, x_t = l)
\end{aligned}$$

$$\begin{aligned}
&= \sum_{\hat{\mathbf{m}} \in [N]^{T-2}} \mathbb{P}(x_o = n \mid x_{T+1} = k, x_s = m, x_t = l, \mathbf{x}_{-s,-t} = \hat{\mathbf{m}}) \mathbb{P}(\mathbf{x}_{-s,-t} = \hat{\mathbf{m}}) \\
&= \sum_{\hat{\mathbf{m}} \in [N]^{T-2}} \left(q_s^{(k)} P_{n,m} + q_t^{(k)} P_{n,l} + \sum_{i \notin \{s,t\}} q_i^{(k)} P_{n,\hat{m}_i} \right) \mathbb{P}(\mathbf{x}_{-s,-t} = \hat{\mathbf{m}}) \\
&= q_s^{(k)} P_{n,m} + q_t^{(k)} P_{n,l} + \left(1 - q_s^{(k)} - q_t^{(k)} \right) \mu_n.
\end{aligned}$$

□

B.3.2 Expectations

Proof of Lemma B.3. For each sample \mathbf{X} , we have

$$l(\mathbf{x}, x_{T+1}, x_o; \mathbf{V}, \mathbf{A}) := \frac{1}{2} \|\mathbf{e}_{x_o} - \mathbf{V} \mathbf{X} \mathbf{A} \mathbf{e}_{x_{T+1}}\|^2.$$

and \mathbf{a} . Then we have the matrix differential:

$$d\mathbf{l} = (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o})^\top d\mathbf{V} \mathbf{X} \mathbf{a} + (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o})^\top \mathbf{V} \mathbf{X} d\mathbf{a}$$

Therefore,

$$\begin{aligned}
\nabla_{\mathbf{V}} \mathbf{l} &= (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o}) (\mathbf{X} \mathbf{a})^\top, \\
\nabla_{\mathbf{a}^{(k)}} \mathbf{l} &= \mathbb{1}\{x_{T+1} = k\} (\mathbf{V} \mathbf{X})^\top (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o})
\end{aligned}$$

□

Proof of Lemma B.4. When $s \neq t$, we have $\mathbb{E}^{(k)}[\mathbf{e}_{x_s} \mathbf{e}_{x_t}^\top] = \mathbb{E}^{(k)}[\mathbf{e}_{x_s}] \mathbb{E}^{(k)}[\mathbf{e}_{x_t}^\top] = \boldsymbol{\mu} \boldsymbol{\mu}^\top$. When $s = t$, we have $\mathbb{E}^{(k)}[\mathbf{e}_{x_t} \mathbf{e}_{x_t}^\top] = \sum_{k=1}^N \mu_k \mathbf{e}_k \mathbf{e}_k^\top = \text{diag}(\boldsymbol{\mu})$. □

Proof of Lemma B.5. When $s \neq t$, we have

$$\begin{aligned}
\sum_{n=1}^N \mathbb{E}[V_{n,x_s} V_{n,x_t}] &= \sum_{n=1}^N (\mathbb{E} V_{n,x_s})^2 = \sum_{n=1}^N \left(\sum_{k=1}^N \mu_k V_{n,k} \right)^2 \\
&= \sum_{k,l=1}^N \mu_k \mu_l \sum_{n=1}^N V_{n,k} V_{n,l} = \boldsymbol{\mu}^\top \mathbf{V}^\top \mathbf{V} \boldsymbol{\mu} = \|\boldsymbol{\mu}\|^2.
\end{aligned}$$

When $s = t$, we have $\sum_{n=1}^N \mathbb{E} V_{n,x_t}^2 = \sum_{n=1}^N \sum_{k=1}^N \mu_k V_{n,k}^2 = \|\mathbf{V}\|_{\boldsymbol{\mu}}^2$. □

Proof of Lemma B.6. Recall Lemma B.1. Then, we compute

$$\begin{aligned}
\mathbb{E}^{(k)} V_{x_o, x_t} &= \sum_{n,m=1}^N V_{n,m} \mathbb{P}(x_o = n, x_t = m \mid x_{T+1} = k) \\
&= q_t^{(k)} \sum_{n,m=1}^N V_{n,m} P_{n,m} \mu_m + \left(1 - q_t^{(k)} \right) \sum_{n,m=1}^N V_{n,m} \mu_n \mu_m \\
&= q_t^{(k)} \langle \mathbf{V}, \mathbf{P} \rangle_{\boldsymbol{\mu}} + \left(1 - q_t^{(k)} \right) \boldsymbol{\mu}^\top \mathbf{V} \boldsymbol{\mu}.
\end{aligned}$$

□

Proof of Lemma B.7. First, we consider $\nabla_{\mathbf{V}} \mathbf{l} = (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o}) (\mathbf{X} \mathbf{a})^\top$, and compute $\mathbb{E}(\mathbf{X} \mathbf{a})(\mathbf{X} \mathbf{a})^\top$ and $\mathbb{E} \mathbf{e}_{x_o} (\mathbf{X} \mathbf{a})^\top$. Write $\mathbf{X} \mathbf{a} = \sum_{t=1}^T a_t \mathbf{e}_{x_t}$. Then, we have

$$\mathbb{E}^{(k)} [(\mathbf{X} \mathbf{a})(\mathbf{X} \mathbf{a})^\top] = \sum_{s,t=1}^T a_s^{(k)} a_t^{(k)} \mathbb{E}[\mathbf{e}_{x_s} \mathbf{e}_{x_t}^\top]$$

$$\begin{aligned}
&= \sum_{t=1}^T (a_t^{(k)})^2 \mathbb{E}[\mathbf{e}_{x_t} \mathbf{e}_{x_t}^\top] + \sum_{s \neq t} a_s^{(k)} a_t^{(k)} \mathbb{E}[\mathbf{e}_{x_s} \mathbf{e}_{x_t}^\top] \\
&= \|\mathbf{a}^{(k)}\|^2 \text{diag}(\boldsymbol{\mu}) + \left(1 - \|\mathbf{a}^{(k)}\|^2\right) \boldsymbol{\mu} \boldsymbol{\mu}^\top,
\end{aligned}$$

where the last line comes from Lemma B.4. Then, we compute

$$\begin{aligned}
\mathbb{E}^{(k)}[\mathbf{e}_{x_o} (\mathbf{X} \mathbf{a})^\top] &= \sum_{t=1}^T a_t^{(k)} \mathbb{E}^{(k)}[\mathbf{e}_{x_o} \mathbf{e}_{x_t}^\top] = \sum_{t=1}^T a_t^{(k)} [\mathbb{P}(x_o = n, x_t = m \mid x_{T+1} = k)]_{n,m \in [N]} \\
&= \sum_{t=1}^T a_t^{(k)} \left(q_t^{(k)} \mathbf{P} \text{diag}(\boldsymbol{\mu}) + (1 - q_t^{(k)}) \boldsymbol{\mu} \boldsymbol{\mu}^\top \right) \\
&= \langle \mathbf{q}^{(k)}, \mathbf{a}^{(k)} \rangle \mathbf{P} \text{diag}(\boldsymbol{\mu}) + \left(1 - \langle \mathbf{q}^{(k)}, \mathbf{a}^{(k)} \rangle\right) \boldsymbol{\mu} \boldsymbol{\mu}^\top,
\end{aligned}$$

where the second line comes from Lemma B.1. Thus, for $\nabla_{\mathbf{V}} l$, we have

$$\begin{aligned}
\mathbb{E} \nabla_{\mathbf{V}} l &= \mathbf{V} \sum_{k=1}^N \mu_k \mathbb{E}^{(k)}[(\mathbf{X} \mathbf{a})(\mathbf{X} \mathbf{a})^\top] - \sum_{k=1}^N \mu_k \mathbb{E}^{(k)}[\mathbf{e}_{x_o} (\mathbf{X} \mathbf{a})^\top] \\
&= \|\mathbf{A}\|_{\mu}^2 \mathbf{V} \text{diag}(\boldsymbol{\mu}) + \left(1 - \|\mathbf{A}\|_{\mu}^2\right) \mathbf{V} \boldsymbol{\mu} \boldsymbol{\mu}^\top - \langle \mathbf{Q}, \mathbf{A} \rangle_{\mu} \mathbf{P} \text{diag}(\boldsymbol{\mu}) - \left(1 - \langle \mathbf{Q}, \mathbf{A} \rangle_{\mu}\right) \boldsymbol{\mu} \boldsymbol{\mu}^\top.
\end{aligned}$$

Now, consider $\nabla_{\mathbf{a}^{(k)}} l = \mathbb{1}_{\{x_{T+1} = k\}} (\mathbf{V} \mathbf{X})^\top (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o})$ and compute $\mathbb{E}^{(k)} (\mathbf{V} \mathbf{X})^\top (\mathbf{V} \mathbf{X})$ and $\mathbb{E}^{(k)} (\mathbf{V} \mathbf{X})^\top \mathbf{e}_{x_o}$. By Lemma B.5, for each $s, t \in [T]$, we have

$$\mathbf{e}_s^\top \mathbb{E}^{(k)}[(\mathbf{V} \mathbf{X})^\top (\mathbf{V} \mathbf{X})] \mathbf{e}_t = \mathbb{E}^{(k)}[(\mathbf{V} \mathbf{e}_{x_s})^\top \mathbf{V} \mathbf{e}_{x_t}] = \sum_{n=1}^N \mathbb{E} V_{n,x_s} V_{n,x_t} = \begin{cases} \|\mathbf{V}\|_{\mu}^2, & s = t, \\ \|\boldsymbol{\mu}\|_{\mu}^2, & s \neq t. \end{cases}$$

In matrix form, this is

$$\mathbb{E}^{(k)} (\mathbf{V} \mathbf{X})^\top \mathbf{V} \mathbf{X} = \left(\|\mathbf{V}\|_{\mu}^2 - \|\boldsymbol{\mu}\|_{\mu}^2 \right) \mathbf{I} + \mathbf{1} \mathbf{1}^\top \|\boldsymbol{\mu}\|_{\mu}^2.$$

Then, by Lemma B.6, for each $t \in [T]$, we have

$$\begin{aligned}
\mathbf{e}_t^\top \mathbb{E}[(\mathbf{V} \mathbf{X})^\top \mathbf{e}_{x_o}] &= \mathbb{E}^{(k)}[(\mathbf{V} \mathbf{e}_{x_t})^\top \mathbf{e}_{x_o}] = \mathbb{E}^{(k)} V_{x_o, x_t} \\
&= q_t^{(k)} \langle \mathbf{V}, \mathbf{P} \rangle_{\mu} + \left(1 - q_t^{(k)}\right) \|\boldsymbol{\mu}\|_{\mu}^2.
\end{aligned}$$

In matrix for, this is

$$\mathbb{E}[(\mathbf{V} \mathbf{X})^\top \mathbf{e}_{x_o}] = \mathbf{q}^{(k)} \langle \mathbf{V}, \mathbf{P} \rangle_{\mu} + \left(\mathbf{1} - \mathbf{q}^{(k)}\right) \|\boldsymbol{\mu}\|_{\mu}^2.$$

Combine these together, and we obtain

$$\begin{aligned}
\mu_k^{-1} \mathbb{E} \nabla_{\mathbf{a}^{(k)}} l &= \mathbb{E}^{(k)}[(\mathbf{V} \mathbf{X})^\top \mathbf{V} \mathbf{X}] \mathbf{a}^{(k)} - \mathbb{E}^{(k)}[(\mathbf{V} \mathbf{X})^\top \mathbf{e}_{x_o}] \\
&= \left(\|\mathbf{V}\|_{\mu}^2 - \|\boldsymbol{\mu}\|_{\mu}^2 \right) \mathbf{a}^{(k)} + \mathbf{1} \|\boldsymbol{\mu}\|_{\mu}^2 - \mathbf{q}^{(k)} \langle \mathbf{V}, \mathbf{P} \rangle_{\mu} - \left(\mathbf{1} - \mathbf{q}^{(k)}\right) \|\boldsymbol{\mu}\|_{\mu}^2.
\end{aligned}$$

□

Proof of Lemma B.8. Recall from Lemma B.7 that

$$\begin{aligned}
\mathbb{E} \nabla_{\mathbf{V}} l &= \|\mathbf{A}\|_{\mu}^2 \mathbf{V} \text{diag}(\boldsymbol{\mu}) + \left(1 - \|\mathbf{A}\|_{\mu}^2\right) \mathbf{V} \boldsymbol{\mu} \boldsymbol{\mu}^\top \\
&\quad - \langle \mathbf{Q}, \mathbf{A} \rangle_{\mu} \mathbf{P} \text{diag}(\boldsymbol{\mu}) - \left(1 - \langle \mathbf{Q}, \mathbf{A} \rangle_{\mu}\right) \boldsymbol{\mu} \boldsymbol{\mu}^\top, \\
\mathbb{E} \nabla_{\mathbf{a}^{(k)}} l &= \mu_k \left(\|\mathbf{V}\|_{\mu}^2 - \|\boldsymbol{\mu}\|_{\mu}^2 \right) \mathbf{a}^{(k)} + \mu_k \mathbf{1} \|\boldsymbol{\mu}\|_{\mu}^2 \\
&\quad - \mu_k \mathbf{q}^{(k)} \langle \mathbf{V}, \mathbf{P} \rangle_{\mu} - \mu_k \left(\mathbf{1} - \mathbf{q}^{(k)}\right) \|\boldsymbol{\mu}\|_{\mu}^2.
\end{aligned}$$

For $\nabla_{\mathbf{V}}$, we have

$$\begin{aligned}\mathbf{V}\text{diag}(\boldsymbol{\mu})\text{diag}(1/\boldsymbol{\mu})\left(\mathbf{I}-\frac{\boldsymbol{\mu}\boldsymbol{\mu}^\top}{\|\boldsymbol{\mu}\|^2}\right)&= \mathbf{V}-\frac{\boldsymbol{\mu}\boldsymbol{\mu}^\top}{\|\boldsymbol{\mu}\|^2}, \\ \mathbf{P}\text{diag}(\boldsymbol{\mu})\text{diag}(1/\boldsymbol{\mu})\left(\mathbf{I}-\frac{\boldsymbol{\mu}\boldsymbol{\mu}^\top}{\|\boldsymbol{\mu}\|^2}\right)&= \mathbf{P}-\frac{\boldsymbol{\mu}\boldsymbol{\mu}^\top}{\|\boldsymbol{\mu}\|^2}, \\ \boldsymbol{\mu}\boldsymbol{\mu}^\top\text{diag}(1/\boldsymbol{\mu})\left(\mathbf{I}-\frac{\boldsymbol{\mu}\boldsymbol{\mu}^\top}{\|\boldsymbol{\mu}\|^2}\right)&= \boldsymbol{\mu}\mathbf{1}^\top-\frac{\boldsymbol{\mu}\boldsymbol{\mu}^\top}{\|\boldsymbol{\mu}\|^2}.\end{aligned}$$

In particular, note that the $\boldsymbol{\mu}\boldsymbol{\mu}^\top/\|\boldsymbol{\mu}\|^2$ terms will cancel with each other. Thus, we have

$$\begin{aligned}\mathbb{E}\hat{\nabla}_{\mathbf{V}}l &= \|\mathbf{A}\|_\mu^2\mathbf{V}+\left(1-\|\mathbf{A}\|_\mu^2\right)\boldsymbol{\mu}\mathbf{1}^\top-\langle\mathbf{Q},\mathbf{A}\rangle_\mu\mathbf{P}-\left(1-\langle\mathbf{Q},\mathbf{A}\rangle_\mu\right)\boldsymbol{\mu}\mathbf{1}^\top \\ &= \|\mathbf{A}\|_\mu^2\left(\mathbf{V}-\boldsymbol{\mu}\mathbf{1}^\top\right)-\langle\mathbf{Q},\mathbf{A}\rangle_\mu\left(\mathbf{P}-\boldsymbol{\mu}\mathbf{1}^\top\right).\end{aligned}$$

For $\hat{\nabla}_{\mathbf{a}^{(k)}}l$, we have

$$\begin{aligned}\mathbb{E}\hat{\nabla}_{\mathbf{a}^{(k)}}l &= \left(\|\mathbf{V}\|_\mu^2-\|\boldsymbol{\mu}\|^2\right)\left(\mathbf{I}-\frac{\mathbf{1}\mathbf{1}^\top}{T}\right)\mathbf{a}^{(k)}-\left(\mathbf{I}-\frac{\mathbf{1}\mathbf{1}^\top}{T}\right)\mathbf{q}^{(k)}\langle\mathbf{V},\mathbf{P}\rangle_\mu \\ &\quad +\left(\mathbf{I}-\frac{\mathbf{1}\mathbf{1}^\top}{T}\right)\mathbf{q}^{(k)}\|\boldsymbol{\mu}\|^2 \\ &= \left(\|\mathbf{V}\|_\mu^2-\|\boldsymbol{\mu}\|^2\right)\left(\mathbf{a}^{(k)}-\frac{\mathbf{1}}{T}\right)-\left(\mathbf{q}^{(k)}-\frac{\mathbf{1}}{T}\right)\langle\mathbf{V},\mathbf{P}\rangle_\mu+\left(\mathbf{q}^{(k)}-\frac{\mathbf{1}}{T}\right)\|\boldsymbol{\mu}\|^2 \\ &= \left(\|\mathbf{V}\|_\mu^2-\|\boldsymbol{\mu}\|^2\right)\left(\mathbf{a}^{(k)}-\frac{\mathbf{1}}{T}\right)-\left(\langle\mathbf{V},\mathbf{P}\rangle_\mu-\|\boldsymbol{\mu}\|^2\right)\left(\mathbf{q}^{(k)}-\frac{\mathbf{1}}{T}\right).\end{aligned}$$

□

B.4 Concentration

In this section, we provide concentration inequalities for the gradients of the loss function. The concentration is applied on the gradient noise term

$$\begin{aligned}\mathbf{h}_{\mathbf{V},\tau} &:= \left(\hat{\nabla}_{\mathbf{V}}^{(B)}l-\mathbb{E}\hat{\nabla}_{\mathbf{V}}l\right) \\ \mathbf{h}_{\mathbf{A},\tau} &:= \left(\hat{\nabla}_{\mathbf{A}}^{(B)}l-\mathbb{E}\hat{\nabla}_{\mathbf{A}}l\right)\end{aligned}$$

where $\hat{\nabla}_{\mathbf{V}}^{(B)}l$ and $\hat{\nabla}_{\mathbf{A}}^{(B)}l$ are the preconditioned empirical gradients computed from a batch of size B . Here, we first consider the concentration of the original gradients:

$$\begin{aligned}\nabla_{\mathbf{V}}l &= (\mathbf{V}\mathbf{X}\mathbf{a}-\mathbf{e}_{x_o})(\mathbf{X}\mathbf{a})^\top, \\ \nabla_{\mathbf{a}^{(k)}}l &= \mathbb{1}\{x_{T+1}=k\}(\mathbf{V}\mathbf{X})^\top(\mathbf{V}\mathbf{X}\mathbf{a}-\mathbf{e}_{x_o})\end{aligned}$$

and then consider the concentration of the preconditioned gradients. In this paper, we focus on $\|\cdot\|_\mu$ as the mostly used metric for the gradient matrices.

First we prove a naive concentration w.r.t. any random vector \mathbf{y} with bounded second moment with any $\|\cdot\|$.

Lemma B.9. *Fix $\delta, \varepsilon > 0$. Let \mathbf{y} be a D -dimensional random vector with $\mathbb{E}\|\mathbf{y}\|^2 \leq G$. Define $\mathbf{y}^{(B)} := B^{-1}\sum_{i=1}^B\mathbf{y}_i$ where $(\mathbf{y}_i)_i$ are i.i.d. versions of \mathbf{y} . If*

$$B \geq \frac{G}{\delta\varepsilon^2},$$

then with probability at least $1-\delta$, we have $\|\mathbf{y}-\mathbb{E}\mathbf{y}\| \leq \varepsilon$.

Proof of Lemma B.9. Assume w.l.o.g. that $\mathbb{E}\mathbf{y} = \mathbf{0}$. First, note that

$$\mathbb{E}\left\|\mathbf{y}^{(B)}\right\|^2 = \frac{1}{B^2}\sum_{i,j=1}^B\mathbb{E}\langle\mathbf{y}_i,\mathbf{y}_j\rangle = \frac{1}{B^2}\sum_{i=1}^B\mathbb{E}\|\mathbf{y}_i\|^2 \leq \frac{G}{B}.$$

Hence, by the Markov inequality, we have

$$\mathbb{P}\left(\|\mathbf{y}^{(B)}\| \geq \varepsilon\right) = \mathbb{P}\left(\|\mathbf{y}^{(B)}\|^2 \geq \varepsilon^2\right) \leq \frac{\mathbb{E}\|\mathbf{y}^{(B)}\|^2}{\varepsilon^2} \leq \frac{G}{B\varepsilon^2}.$$

Thus, for fixed $\varepsilon, \delta \in (0, 1)$, if we choose $B = G/(\delta\varepsilon^2)$, then we have with probability at least $1 - \delta$, $\|\mathbf{y}^{(B)}\| \leq \varepsilon$. \square

Now we upper bound the infinity norm of the preconditioned gradients to apply concentration.

Lemma B.10. *Suppose that $\|\text{Vec } \mathbf{V}\|_\infty = O(1)$, $\mathbf{V} = \tilde{\mathbf{V}} + \Delta_{\mathbf{V}}$, $\mathbf{A} = \tilde{\mathbf{A}} + \Delta_{\mathbf{A}}$, where $\tilde{\mathbf{V}} \in \mathbb{R}^{N \times N}$ is a transition probability matrix, and in the attention matrix $\tilde{\mathbf{A}}$ each column $\tilde{\mathbf{a}}^{(k)}$ is a probability vector. Moreover, $\|\Delta_{\mathbf{A}}\|_F^2, \|\Delta_{\mathbf{V}}\|_F^2 \leq O(1/T)$. Then, we have*

$$\|\hat{\nabla}_{\mathbf{a}^{(k)}} l\|_\infty \leq O(N), \quad \|\text{Vec } \hat{\nabla}_{\mathbf{V}} l\|_\infty \leq O(N)$$

Proof. We first consider the infinity norm of the original gradient. Recall that the gradient for \mathbf{V} and \mathbf{A} are

$$\begin{aligned} \nabla_{\mathbf{a}^{(k)}} l &= \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}\mathbf{X})^\top (\mathbf{V}\mathbf{X}\mathbf{a}^{(k)} - \mathbf{e}_{x_o}) \\ \nabla_{\mathbf{V}} l &= (\mathbf{V}\mathbf{X}\mathbf{a} - \mathbf{e}_{x_o}) (\mathbf{X}\mathbf{a})^\top \end{aligned}$$

and the preconditioned gradient is:

$$\begin{aligned} \hat{\nabla}_{\mathbf{V}}^{(B_\tau)} l &:= \left(\mathbf{I}_N - \frac{\mathbf{1}_N \mathbf{1}_N^\top}{N}\right) \left(\nabla_{\mathbf{V}}^{(B_\tau)} l\right) \text{diag}(1/\mu) \left(\mathbf{I}_N - \frac{\mu\mu^\top}{\|\mu\|^2}\right), \\ \hat{\nabla}_{\mathbf{a}^{(k)}}^{(B_\tau)} l &:= \frac{1}{\mu_k} \left(\mathbf{I}_T - \frac{\mathbf{1}_T \mathbf{1}_T^\top}{T}\right) \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l\right). \end{aligned}$$

We first consider the maximum absolute value in the original gradients. For $\nabla_{\mathbf{a}^{(k)}} l$, we have

$$\begin{aligned} &\left\| \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}\mathbf{X})^\top (\mathbf{V}\mathbf{X}\mathbf{a}^{(k)} - \mathbf{e}_{x_o}) \right\|_\infty \\ &\leq \left\| \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}\mathbf{X})^\top (\mathbf{V}\mathbf{X}\mathbf{a}^{(k)}) \right\|_\infty + \left\| \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}\mathbf{X})^\top (\mathbf{e}_{x_o}) \right\|_\infty \\ &\leq \max_{s \in [T]} \left| \sum_{t=1}^T a_t (\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t} \right| + \max_{s \in [T]} |(\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{e}_{x_o}| \end{aligned}$$

The first term can be upper-bounded in the following way:

$$\begin{aligned} \max_{s \in [T]} \left| \sum_{t=1}^T a_t (\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t} \right| &= \max_{s \in [T]} \left| \sum_{t=1}^T \tilde{a}_t (\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t} + \sum_{t=1}^T \Delta_{\mathbf{a},t} (\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t} \right| \\ &\leq \left(1 + \sum_{t=1}^T \|\Delta_{\mathbf{a},t}\|_1\right) \max_{s,t} |(\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t}| \\ &\leq (1 + \sqrt{T} \|\Delta_{\mathbf{a},t}\|_2) \max_{s,t} |(\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t}| \end{aligned}$$

Since $\mathbf{V} = \tilde{\mathbf{V}} + \Delta_{\mathbf{V}}$ and $\|\Delta_{\mathbf{V}}\|_F^2, \|\Delta_{\mathbf{A}}\|_F^2 \leq O(1/T)$, we have $\max_{s,t} |(\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t}|$ upper bounded by $O(1)$. Therefore

$$\max_{s \in [T]} \left| \sum_{t=1}^T a_t (\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{V}\mathbf{e}_{x_t} \right| \leq O(1)$$

And similarly, the second term $\max_{s \in [T]} |(\mathbf{V}\mathbf{e}_{x_s})^\top \mathbf{e}_{x_o}|$ can be bounded by $O(1)$ because the infinity norm of \mathbf{V} is also upper bounded by $O(1)$. Therefore, we know $\|\nabla_{\mathbf{a}^{(k)}} l\|_\infty \leq O(1)$.

Now we consider the preconditioned gradient $\hat{\nabla}_{\mathbf{a}^{(k)}}^{(B_\tau)} l$:

$$\|\hat{\nabla}_{\mathbf{a}^{(k)}}^{(B_\tau)} l\|_\infty = \left\| \frac{1}{\mu_k} \left(\mathbf{I}_T - \frac{\mathbf{1}_T \mathbf{1}_T^\top}{T}\right) \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l\right) \right\|_\infty$$

$$\leq \left\| \frac{1}{\mu_k} \mathbf{I}_T \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l \right) \right\|_\infty + \left\| \frac{1}{\mu_k} \left(\frac{\mathbf{1}_T \mathbf{1}_T^\top}{T} \right) \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l \right) \right\|_\infty \leq O(N)$$

since $\mu_k \geq \frac{c}{N}$ for all $k \in [N]$.

We use similar technique on $\hat{\nabla}_{\mathbf{V}}^{(B_\tau)} l$. First, we prove the infinity norm upper bound on the original gradient.

$$\begin{aligned} \|\text{Vec} \nabla_{\mathbf{V}} l\|_\infty &= \|\text{Vec} (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o}) (\mathbf{X} \mathbf{a})^\top\|_\infty \\ &= \|\text{Vec} (\mathbf{V} \mathbf{X} \mathbf{a}) (\mathbf{X} \mathbf{a})^\top\|_\infty + \|\text{Vec} (\mathbf{e}_{x_o} (\mathbf{X} \mathbf{a})^\top)\|_\infty \\ &= \max_{s,t \in [T]} |(\mathbf{V} \mathbf{X} \mathbf{a})_s (\mathbf{X} \mathbf{a})_t^\top| + \left\| \sum_{t=1}^T a_t \text{Vec} (\mathbf{e}_{x_o} \mathbf{e}_{x_t}^\top) \right\|_\infty \\ &\leq \|(\mathbf{V} \mathbf{X} \mathbf{a})\|_\infty \|(\mathbf{X} \mathbf{a})^\top\|_\infty + \left\| \sum_{t=1}^T a_t \text{Vec} (\mathbf{e}_{x_o} \mathbf{e}_{x_t}^\top) \right\|_\infty \\ &= \sum_{s,t=1}^T (\tilde{a}_s + \|\Delta_{\mathbf{A},s}\|_1) (\tilde{a}_t + \|\Delta_{\mathbf{A},t}\|_1) \|\mathbf{V} \mathbf{e}_{x_s}\|_\infty \|\mathbf{e}_{x_t}\|_\infty \\ &\quad + \sum_{t=1}^T (\tilde{a}_t + \|\Delta_{\mathbf{A},t}\|_1) \|\text{Vec} (\mathbf{e}_{x_o} \mathbf{e}_{x_t}^\top)\|_\infty \leq \Theta(1). \end{aligned}$$

And therefore, the preconditioned gradient can also be bounded.

$$\begin{aligned} \|\hat{\nabla}_{\mathbf{V}}^{(B_\tau)} l\|_\infty &= \left\| \left(\mathbf{I}_N - \frac{\mathbf{1}_N \mathbf{1}_N^\top}{N} \right) \left(\nabla_{\mathbf{V}}^{(B_\tau)} l \right) \text{diag}(1/\mu) \left(\mathbf{I}_N - \frac{\mu \mu^\top}{\|\mu\|^2} \right) \right\|_\infty \\ &\leq \left\| \left(\nabla_{\mathbf{V}}^{(B_\tau)} l \right) \text{diag}(1/\mu) \left(\mathbf{I}_N - \frac{\mu \mu^\top}{\|\mu\|^2} \right) \right\|_\infty \\ &\quad + \left\| \frac{\mathbf{1}_N \mathbf{1}_N^\top}{N} \left(\nabla_{\mathbf{V}}^{(B_\tau)} l \right) \text{diag}(1/\mu) \left(\mathbf{I}_N - \frac{\mu \mu^\top}{\|\mu\|^2} \right) \right\|_\infty \\ &\leq O(N) \end{aligned}$$

since $\mu_k \geq \frac{c}{N}$ for all $k \in [N]$. Now we finished the proof. \square

With the upper bound of the infinity norm, we have the following upper bound on the second order moments of the preconditioned gradients of \mathbf{A} and \mathbf{V} .

Corollary B.11. *With the same setting in Lemma B.10 and $\|\hat{\nabla}_{\mathbf{a}^{(k)}} l\|_\infty \leq O(N)$, $\|\text{Vec} \hat{\nabla}_{\mathbf{V}} l\|_\infty \leq O(N)$. Moreover, $\|\Delta_{\mathbf{A}}\|_F^2, \|\Delta_{\mathbf{V}}\|_F^2 \leq O(1/T)$. Then, we have*

$$\mathbb{E} \|\hat{\nabla}_{\mathbf{A}} l\|_\mu^2 \leq O(TN^2), \quad \mathbb{E} \|\hat{\nabla}_{\mathbf{V}} l\|_\mu^2 \leq O(N^3)$$

Proof. We directly upper bound $\|\hat{\nabla}_{\mathbf{A}} l\|_\mu^2$ and $\|\hat{\nabla}_{\mathbf{V}} l\|_\mu^2$ using the upper bound on infinity norm. Since $\|\Delta_{\mathbf{A}}\|_F^2, \|\Delta_{\mathbf{V}}\|_F^2 \leq O(1/T)$, we have the infinity norm be upper bounded by

$$\|\hat{\nabla}_{\mathbf{a}^{(k)}} l\|_\infty \leq O(N), \quad \|\text{Vec} \hat{\nabla}_{\mathbf{V}} l\|_\infty \leq O(N)$$

Then, we can first bound the Frobenius norm $\|\hat{\nabla}_{\mathbf{A}} l\|_F^2, \|\hat{\nabla}_{\mathbf{V}} l\|_F^2$. We have $\mathbf{V} \in \mathbb{R}^{N \times N}, \mathbf{A} \in \mathbb{R}^{T \times N}$, so

$$\|\hat{\nabla}_{\mathbf{V}} l\|_F^2 \leq N^2 \|\hat{\nabla}_{\mathbf{V}} l\|_\infty^2 = O(N^4), \quad \|\hat{\nabla}_{\mathbf{A}} l\|_F^2 \leq NT \|\text{Vec} \hat{\nabla}_{\mathbf{A}} l\|_\infty^2 = O(N^3 T).$$

That leads to:

$$\begin{aligned} \|\hat{\nabla}_{\mathbf{V}} l\|_\mu^2 &= \langle \hat{\nabla}_{\mathbf{V}} l, \hat{\nabla}_{\mathbf{V}} l \text{diag}(\mu) \rangle \leq O\left(\frac{1}{N}\right) \|\hat{\nabla}_{\mathbf{V}} l\|_F^2 \leq O(N^3) \\ \|\hat{\nabla}_{\mathbf{A}} l\|_\mu^2 &= \langle \hat{\nabla}_{\mathbf{A}} l, \hat{\nabla}_{\mathbf{A}} l \text{diag}(\mu) \rangle \leq O\left(\frac{1}{N}\right) \|\hat{\nabla}_{\mathbf{A}} l\|_F^2 \leq O(N^2 T) \end{aligned}$$

where the second inequality comes from the assumption that $\mu \sim \Theta(1/N)$. \square

Now with the upper bound of the second moments of the gradients, we begin to prove the concentration of the gradients. We first consider the first-order terms that need to be bounded in the signal dynamics:

$$\frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_P K_Q}, \quad \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q}$$

Lemma B.12. Fix $\varepsilon, \delta > 0$. Under Assumption 2.1, suppose $\|\hat{\nabla}_{\mathbf{a}^{(k)}} l\|_\infty \leq \Theta(N)$, $\|\hat{\nabla}_V l\|_\infty \leq \Theta(N)$. If $B \geq \Theta(1) \max(N^4, Q^2 N^2) \frac{N^2 \log \frac{4}{\delta}}{\varepsilon^2 K_P^2 K_Q^2}$, then with probability at least $1 - \delta$, we have:

$$\left| \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_P K_Q} \right| \leq \varepsilon, \quad \left| \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q} \right| \leq \varepsilon.$$

Proof. Note $\mathbf{h}_A = \frac{1}{B} \sum_i \hat{\nabla}_{\mathbf{a}^{(k_i)}} l(i) - \mathbb{E} \hat{\nabla}_{\mathbf{a}^{(k)}} l$, thus we have the upper bound for each coordinate of the gradient error bounded by $\Theta(N)$. Similarly, we have the upper bound for each coordinate of the gradient error of \mathbf{h}_V bounded by $\Theta(N)$.

Then, we can bound the infinity norm of $\langle \hat{\nabla}_V l(i), \mathbf{P} \rangle_\mu$ and $\langle \hat{\nabla}_A l(i), \mathbf{P} \rangle_\mu$:

$$\begin{aligned} \left| \frac{\langle \hat{\nabla}_V l(i), \mathbf{P} \rangle_\mu}{K_P K_Q} \right| &= \left| \frac{\sum_{n,m=1}^N \hat{\nabla}_V l(i)_{n,m} \mathbf{P}_{n,m}}{K_P K_Q} \right| \\ &\leq \frac{N^2 \|\hat{\nabla}_V l(i)\|_\infty \|\mathbf{P}\|_\infty}{K_P K_Q} \\ &\leq \frac{\Theta(1) N^3}{K_P K_Q}. \quad (\|\mathbf{P}\|_\infty \leq 1.) \\ \left| \frac{\langle \hat{\nabla}_A l(i), \mathbf{Q} \rangle_\mu}{K_P K_Q} \right| &= \left| \frac{\sum_{n=1}^T \sum_{m=1}^N \hat{\nabla}_A l(i)_{n,m} \mathbf{Q}_{n,m}}{K_P K_Q} \right| \\ &\leq \frac{QN \|\hat{\nabla}_A l(i)\|_\infty \|\mathbf{Q}\|_\infty}{K_P K_Q} \quad (\mathbf{Q} \text{ is } Q\text{-sparse.}) \\ &\leq \frac{\Theta(1) Q N^2}{K_P K_Q}. \quad (\|\mathbf{Q}\|_\infty \leq 1.) \end{aligned}$$

Note that $\mathbb{E} [\nabla_{\mathbf{a}^{(k_i)}} l(i) - \mathbb{E} \nabla_{\mathbf{a}^{(k)}} l] = 0$, $\mathbb{E} [\nabla_V l(i) - \mathbb{E} \nabla_V l] = 0$, which means the two terms above have expectation 0. Since $\mathbf{h}_A, \mathbf{h}_V$ are both averages of B gradients of a single sample, we use Hoeffding Inequality:

$$\begin{aligned} \mathbb{P} \left(\left| \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_P K_Q} \right| \geq \varepsilon \right) &\leq 2 \exp \left(\frac{-B \varepsilon^2 K_P^2 K_Q^2}{N^6} \right) \\ \mathbb{P} \left(\left| \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q} \right| \geq \varepsilon \right) &\leq 2 \exp \left(\frac{-B \varepsilon^2 K_P^2 K_Q^2}{N^4 Q^2} \right) \end{aligned}$$

By union bound, if $B \geq \max(N^2, Q^2) \frac{N^4 \log \frac{4}{\delta}}{\varepsilon^2 K_P^2 K_Q^2}$ it has at least $1 - \delta$ probability, s.t.

$$\left| \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_P K_Q} \right| \leq \varepsilon, \quad \left| \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q} \right| \leq \varepsilon.$$

□

Then we finish this section with the concentration of the second order terms $\|\mathbf{h}_A\|^2$ and $\|\mathbf{h}_V\|^2$, which need to be bounded in the error evolution.

Lemma B.13. Fix $\varepsilon, \delta > 0$. Under Assumption 2.1, if $\mathbb{E} \|\hat{\nabla}_V l\|_\mu^2 = \text{Tmp}_3 = O(N^3)$, $\mathbb{E} \|\hat{\nabla}_A l\|_\mu^2 = \text{Tmp}_4 = O(TN^2)$, $B \geq \frac{\Theta(TN^2)}{\delta \varepsilon^2}$, then with probability at least $1 - \delta$, we have:

$$\|\mathbf{h}_A\|_\mu \leq \varepsilon, \quad \|\mathbf{h}_V\|_\mu \leq \varepsilon.$$

Proof. Similar to Lemma B.12, $\mathbb{E} \mathbf{h}_V = 0, \mathbb{E} \mathbf{h}_A = 0$. By Lemma B.9, when we pick $B \geq \frac{\Theta(TN^2)}{\delta\epsilon^2}$ we have $\|\mathbf{h}_A\|_\mu \leq \epsilon, \|\mathbf{h}_V\|_\mu \leq \epsilon$ with probability at least $1 - \delta$. \square

C The population projected process

In this section, we define the projection of the true SGD process onto the ‘‘space of population trajectories’’. Then, we derive formulas for the dynamics of projected process and the distance of the true SGD process to the space of population trajectories. All proofs — except for those short ones — are deferred to the end of this section.

C.1 Definition of the population projection

The main reason we analyze the population process first is that on the population trajectory, both layers possess special structures. Recall that

$$\begin{aligned} \mathbb{E} \hat{\mathbf{V}}_V l &= \|\mathbf{A}\|_\mu^2 (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top) - \langle \mathbf{Q}, \mathbf{A} \rangle_\mu (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top), \\ \mathbb{E} \hat{\mathbf{V}}_A l &= \left(\|\mathbf{V}\|_\mu^2 - \|\boldsymbol{\mu}\|^2 \right) \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \left(\langle \mathbf{V}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2 \right) \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right), \end{aligned}$$

and we initialize $\mathbf{V}_0 - \boldsymbol{\mu} \mathbf{1}^\top = 0$ and $\mathbf{A} - \mathbf{1} \mathbf{1}^\top / T = 0$. Note that for any $(z_\tau)_\tau$, if $z_0 = 0$ and $\dot{z} = \mathbf{A}z + Bz_*$, then $z_\tau \propto z_*$ for all $\tau \geq 0$. In other words, z_τ moves only along the direction z_* and therefore, can be characterized by a single real number. This is exactly the same case of $\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top$ and $\mathbf{A} - \mathbf{1} \mathbf{1}^\top / T$ in the population case. Hence, in the population case, \mathbf{V} stays on the line crossing $\boldsymbol{\mu} \mathbf{1}^\top$ and \mathbf{P} , and \mathbf{A} stays on the line crossing $\mathbf{1} \mathbf{1}^\top / T$ and \mathbf{Q} .

Unfortunately, mini-batch SGD does not stay exactly on the population trajectory. We can still, however, look at the projection of SGD onto the ‘‘population trajectories’’. Formally, for any \mathbf{V}, \mathbf{A} satisfying $\mathbf{1}_N^\top \mathbf{V} = \mathbf{1}_N^\top, \mathbf{1}_T^\top \mathbf{A} = \mathbf{1}_T^\top$, and $\mathbf{V} \boldsymbol{\mu} = \boldsymbol{\mu}$, we define

$$\begin{aligned} \alpha_V &:= \operatorname{argmin}_{\alpha \in \mathbb{R}} \left\| \alpha \mathbf{P} + (1 - \alpha) \boldsymbol{\mu} \mathbf{1}_N^\top - \mathbf{V} \right\|_\mu^2, \\ \alpha_A &:= \operatorname{argmin}_{\alpha \in \mathbb{R}} \left\| \alpha \mathbf{Q} + (1 - \alpha) \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} - \mathbf{A} \right\|_\mu^2. \end{aligned}$$

By setting the derivative to be 0, we can obtain the following closed-form formulas for α_V and α_A .

Note: Without specification, we drop the the time subscript τ and consider $\alpha_V := \alpha_{V, \tau}$ for simlicity.

Lemma C.1. *For any \mathbf{V}, \mathbf{A} satisfying $\mathbf{1}_N^\top \mathbf{V} = \mathbf{1}_N^\top, \mathbf{1}_T^\top \mathbf{A} = \mathbf{1}_T^\top$, and $\mathbf{V} \boldsymbol{\mu} = \boldsymbol{\mu}$, we have*

$$\alpha_V = K_{VP} / K_P \quad \text{and} \quad \alpha_A = K_{AQ} / K_Q,$$

where $K_P = \|\mathbf{P}\|_\mu^2 - \|\boldsymbol{\mu}\|^2$, $K_{VP} = \langle \mathbf{V}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2$, $K_Q = \|\mathbf{Q}\|_\mu^2 - 1/T$, $K_{AQ} = \langle \mathbf{A}, \mathbf{Q} \rangle_\mu - 1/T$.

For notational simplicity, we define $\beta_V = 1 - \alpha_V, \beta_A = 1 - \alpha_A$,

$$\tilde{\mathbf{V}} = \alpha_V \mathbf{P} + \beta_V \boldsymbol{\mu} \mathbf{1}^\top \quad \text{and} \quad \tilde{\mathbf{A}} = \alpha_A \mathbf{Q} + \beta_A \frac{\mathbf{1} \mathbf{1}^\top}{T}.$$

Then, define $\Delta_V = \mathbf{V} - \tilde{\mathbf{V}}$ and $\Delta_A = \mathbf{A} - \tilde{\mathbf{A}}$ so that we can decompose $\mathbf{V} = \tilde{\mathbf{V}} + \Delta_V$ and $\mathbf{A} = \tilde{\mathbf{A}} + \Delta_A$. By our construction, we have $\tilde{\mathbf{V}} - \boldsymbol{\mu} \mathbf{1}^\top \perp \Delta_V$ and similarly for Δ_A . We will now show that we can in fact drop $\boldsymbol{\mu} \mathbf{1}^\top$.

Lemma C.2. *For any $\mathbf{V}' = \theta \mathbf{P} + (1 - \theta) \boldsymbol{\mu} \mathbf{1}^\top$ and $\mathbf{A}' = \theta \mathbf{Q} + (1 - \theta) \mathbf{1} \mathbf{1}^\top / T$ with $\theta \in \mathbb{R}$, we have $\langle \Delta_V, \mathbf{V}' \rangle_\mu = 0$ and $\langle \Delta_A, \mathbf{A}' \rangle_\mu = 0$. In particular, we have $\langle \Delta_V, \mathbf{P} \rangle_\mu = \langle \Delta_V, \tilde{\mathbf{V}} \rangle_\mu = 0$ and $\langle \Delta_A, \mathbf{Q} \rangle_\mu = \langle \Delta_A, \tilde{\mathbf{A}} \rangle_\mu = 0$.*

Proof. Note that $\langle \boldsymbol{\mu} \mathbf{1}^\top, \Delta_V \rangle_\mu = \langle \boldsymbol{\mu}, (\mathbf{V} - \tilde{\mathbf{V}}) \boldsymbol{\mu} \rangle = 0$. Hence, $\langle \Delta_V, \mathbf{V}' \rangle = \langle \Delta_V, \mathbf{V}' - \boldsymbol{\mu} \mathbf{1}^\top \rangle = 0$. For $\langle \Delta_A, \mathbf{A}' \rangle$, it suffices to note that $\langle \Delta_A, \mathbf{1} \mathbf{1}^\top / T \rangle_\mu = \langle \mathbf{1}^\top (\mathbf{A} - \tilde{\mathbf{A}}), \mathbf{1}^\top / T \rangle_\mu = 0$. \square

The following lemma the basic definitions and results about the population projection.

Lemma C.3 (Definitions and basic results on the population projection). *Suppose that \mathbf{V}, \mathbf{A} satisfy $\mathbf{1}_N^\top \mathbf{V} = \mathbf{1}_N^\top, \mathbf{1}_T^\top \mathbf{A} = \mathbf{1}_N^\top$, and $\mathbf{V}\boldsymbol{\mu} = \boldsymbol{\mu}$. We define the following:*

$$\begin{aligned} K_P &= \|\mathbf{P}\|_\mu^2 - \|\boldsymbol{\mu}\|^2, & K_{VP} &= \langle \mathbf{V}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2, & K_V &= \|\mathbf{V}\|_\mu^2 - \|\boldsymbol{\mu}\|^2, \\ \alpha_V &= K_{VP}/K_P, & \beta_V &= 1 - \alpha_V, & \tilde{\mathbf{V}} &= \alpha_V \mathbf{P} + \beta_V \boldsymbol{\mu} \mathbf{1}^\top, \\ & & \Delta_V &= \mathbf{V} - \tilde{\mathbf{V}}, \\ K_Q &= \|\mathbf{Q}\|_\mu^2 - 1/T, & K_{AQ} &= \langle \mathbf{A}, \mathbf{Q} \rangle_\mu - 1/T, & K_A &= \|\mathbf{A}\|_\mu^2 - 1/T, \\ \alpha_A &= K_{AQ}/K_Q, & \beta_A &= 1 - \beta_A, & \tilde{\mathbf{A}} &= \alpha_A \mathbf{A} + \beta_A \mathbf{1}_T \mathbf{1}_N^\top / T, \\ & & \Delta_A &= \mathbf{A} - \tilde{\mathbf{A}}. \end{aligned}$$

Moreover, by Lemma C.2, the following hold.

$$\begin{aligned} K_{VP} &= \langle \tilde{\mathbf{V}}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2 = \alpha_V K_P, \\ K_V &= \|\tilde{\mathbf{V}}\|_\mu^2 + \|\Delta_V\|_\mu^2 - \|\boldsymbol{\mu}\|^2 = \alpha_V^2 K_P + \|\Delta_V\|_\mu^2, \\ K_{AQ} &= \langle \tilde{\mathbf{A}}, \mathbf{Q} \rangle_\mu - 1/T = \alpha_A K_Q, \\ K_A &= \|\tilde{\mathbf{A}}\|_\mu^2 + \|\Delta_A\|_\mu^2 - 1/T = \alpha_A^2 K_Q + \|\Delta_A\|_\mu^2. \end{aligned}$$

C.2 Dynamics of the population projected process and the approximation error

We write

$$\begin{aligned} \mathbf{V}_{\tau+1} &= \mathbf{V}_\tau - \eta_V \mathbb{E} \hat{\nabla}_V l - \eta_V \left(\hat{\nabla}_V^{(B)} l - \mathbb{E} \hat{\nabla}_V l \right) =: \mathbf{V}_\tau - \eta_V \mathbb{E} \hat{\nabla}_V l - \eta_V \mathbf{h}_{V,\tau}, \\ \mathbf{A}_{\tau+1} &= \mathbf{A}_\tau - \eta_A \mathbb{E} \hat{\nabla}_A l - \eta_A \left(\hat{\nabla}_A^{(B)} l - \mathbb{E} \hat{\nabla}_A l \right) =: \mathbf{A}_\tau - \eta_A \mathbb{E} \hat{\nabla}_A l - \eta_A \mathbf{h}_{A,\tau}, \end{aligned}$$

where the expectations are taken over the fresh samples at step τ .

First, we expand the expected preconditioned gradients around the population projection.

Lemma C.4 (Expanding the gradients).

$$\begin{aligned} \mathbb{E} \hat{\nabla}_A l &= K_P \alpha_V (\alpha_V \alpha_A - 1) \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) + K_P \alpha_V^2 \Delta_A + \|\Delta_V\|_\mu^2 \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right), \\ \mathbb{E} \hat{\nabla}_V l &= \alpha_A K_Q (\alpha_A \alpha_V - 1) (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top) + \frac{\alpha_V - 1}{T} (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top) \\ &\quad + \left(\alpha_A^2 K_Q + \frac{1}{T} \right) \Delta_V + \|\Delta_A\|^2 (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top). \end{aligned}$$

Then, we compute the dynamics of the projected process.

Lemma C.5 (Dynamics of the population projection).

$$\begin{aligned} \alpha_{V,\tau+1} &= \alpha_{V,\tau} + \eta_V K_Q (1 - \alpha_A \alpha_V) \alpha_A + \eta_V \frac{1 - \alpha_V}{T} - \eta_V \alpha_V \|\Delta_A\|^2 - \frac{\eta_V}{K_P} \langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu, \\ \alpha_{A,\tau+1} &= \alpha_{A,\tau} + \eta_A K_P (1 - \alpha_V \alpha_A) \alpha_V - \eta_A \alpha_A \|\Delta_V\|_\mu^2 - \frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu. \end{aligned}$$

Note that $\tilde{\mathbf{V}} = \boldsymbol{\mu} \mathbf{1}^\top + \alpha_V (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top)$ and $\tilde{\mathbf{A}} = \mathbf{1}_T \mathbf{1}_N^\top / T + \alpha_A (\mathbf{Q} - \mathbf{1}_T \mathbf{1}_N^\top / T)$. Hence, this also gives formulas for $\tilde{\mathbf{V}}_{\tau+1}$ and $\tilde{\mathbf{A}}_{\tau+1}$.

Now, we consider the dynamics of the errors.

Lemma C.6 (Dynamics of the errors).

$$\begin{aligned} \|\Delta_{A,\tau+1}\|_\mu^2 &= \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2 \right)^2 \|\Delta_A\|_\mu^2 \\ &\quad - 2\eta_A \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2 \right) \langle \Delta_A, \mathbf{h}_A \rangle_\mu \end{aligned}$$

$$\begin{aligned}
& -\frac{\eta_A^2}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu^2 + \eta_A^2 \|\mathbf{h}_A\|_\mu^2, \\
\|\Delta_{V,\tau+1}\|_\mu^2 &= \left(1 - \eta_V \left(\alpha_A^2 K_Q + \frac{1}{T}\right) - \eta_V \|\Delta_A\|^2\right)^2 \|\Delta_V\|_\mu^2 \\
&+ 2\eta_V \left(1 - \eta_V \left(\alpha_A^2 K_Q + \frac{1}{T}\right) - \eta_V \|\Delta_A\|^2\right) \langle \Delta_V, \mathbf{h}_V \rangle_\mu \\
&- \frac{\eta_V^2}{K_P} \langle \mathbf{P}, \mathbf{h}_V \rangle_\mu^2 + \eta_V^2 \|\mathbf{h}_V\|_\mu^2.
\end{aligned}$$

C.3 Omitted proofs in this section

Proof of Lemma C.1. We compute

$$\begin{aligned}
\frac{1}{2} \partial_\alpha \|\alpha \mathbf{P} + (1-\alpha) \boldsymbol{\mu} \mathbf{1}^\top - \mathbf{V}\|_\mu^2 &= \langle \alpha \mathbf{P} + (1-\alpha) \boldsymbol{\mu} \mathbf{1}^\top - \mathbf{V}, \mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top \rangle_\mu \\
&= \langle \alpha \mathbf{P} + (1-\alpha) \boldsymbol{\mu} \mathbf{1}^\top - \mathbf{V}, \mathbf{P} \rangle_\mu = \alpha K_P + \|\boldsymbol{\mu}\|^2 - \langle \mathbf{V}, \mathbf{P} \rangle_\mu.
\end{aligned}$$

Set the derivative to be 0, and we get $\alpha_V = (\langle \mathbf{V}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2) / K_P$. Similarly, we compute

$$\begin{aligned}
\frac{1}{2} \partial_\alpha \left\| \alpha \mathbf{Q} + (1-\alpha) \frac{\mathbf{1} \mathbf{1}^\top}{T} - \mathbf{A} \right\|_\mu^2 &= \left\langle \alpha \mathbf{Q} + (1-\alpha) \frac{\mathbf{1} \mathbf{1}^\top}{T} - \mathbf{A}, \mathbf{Q} - \frac{\mathbf{1} \mathbf{1}^\top}{T} \right\rangle_\mu \\
&= \left\langle \alpha \mathbf{Q} + (1-\alpha) \frac{\mathbf{1} \mathbf{1}^\top}{T} - \mathbf{A}, \mathbf{Q} \right\rangle_\mu \\
&= \alpha \left(\|\mathbf{Q}\|_\mu^2 - 1/T \right) - \left(\langle \mathbf{A}, \mathbf{Q} \rangle_\mu - 1/T \right).
\end{aligned}$$

Again, set the derivative to be 0, and we get $\alpha_A = (\langle \mathbf{A}, \mathbf{Q} \rangle_\mu - 1/T) / (\|\mathbf{Q}\|_\mu^2 - 1/T)$. \square

Proof of Lemma C.4. Recall from Lemma B.8 that

$$\begin{aligned}
\mathbb{E} \hat{\nabla}_V l &= \left(K_A + \frac{1}{T} \right) (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top) - \left(K_{AQ} + \frac{1}{T} \right) (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top), \\
\mathbb{E} \hat{\nabla}_A l &= K_V \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - K_{VP} \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}^\top}{T} \right).
\end{aligned}$$

First, consider the dynamics of \mathbf{A} . By Lemma C.3, we can further decompose it as

$$\begin{aligned}
\mathbb{E} \hat{\nabla}_A l &= \left(\alpha_V^2 K_P + \|\Delta_V\|_\mu^2 \right) \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \alpha_V K_P \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}^\top}{T} \right) \\
&= K_P \alpha_V \left(\alpha_V \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}^\top}{T} \right) \right) + \|\Delta_V\|_\mu^2 \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) \\
&= K_P \alpha_V \left(\alpha_V \left(\tilde{\mathbf{A}} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}^\top}{T} \right) \right) + K_P \alpha_V^2 \Delta_A + \|\Delta_V\|_\mu^2 \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) \\
&= K_P \alpha_V (\alpha_V \alpha_A - 1) \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) + K_P \alpha_V^2 \Delta_A + \|\Delta_V\|_\mu^2 \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right).
\end{aligned}$$

Similarly, we can rewrite the expected preconditioned gradient of \mathbf{V} as

$$\begin{aligned}
\mathbb{E} \hat{\nabla}_V l &= \left(\alpha_A^2 K_Q + \frac{1}{T} \right) (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top) - \left(\alpha_A K_Q + \frac{1}{T} \right) (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top) + \|\Delta_A\|^2 (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top) \\
&= \left(\alpha_A^2 K_Q + \frac{1}{T} \right) (\tilde{\mathbf{V}} - \boldsymbol{\mu} \mathbf{1}^\top) - \left(\alpha_A K_Q + \frac{1}{T} \right) (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top) \\
&\quad + \left(\alpha_A^2 K_Q + \frac{1}{T} \right) \Delta_V + \|\Delta_A\|^2 (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top)
\end{aligned}$$

$$\begin{aligned}
&= \alpha_A K_Q (\alpha_A \alpha_V - 1) (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top) + \frac{\alpha_V - 1}{T} (\mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top) \\
&\quad + \left(\alpha_A^2 K_Q + \frac{1}{T} \right) \Delta_V + \|\Delta_A\|^2 (\mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top).
\end{aligned}$$

□

Proof of Lemma C.5. Recall that $\alpha_V = \langle \mathbf{V}, \mathbf{P} \rangle_\mu / K_P$ and $\alpha_A = \langle \mathbf{A}, \mathbf{Q} \rangle_\mu / K_Q$. First, consider the dynamics of V . By Lemma C.4 and Lemma C.3, we have

$$\begin{aligned}
\alpha_{V,\tau+1} &= \alpha_{V,\tau} - \eta_V \frac{\langle \hat{\nabla}_V \mathcal{L} + \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu}{K_P} \\
&= \alpha_{V,\tau} - \frac{\eta_V}{K_P} \alpha_A K_Q (\alpha_A \alpha_V - 1) \langle \mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top, \mathbf{P} \rangle_\mu - \frac{\eta_V}{K_P} \frac{\alpha_V - 1}{T} \langle \mathbf{P} - \boldsymbol{\mu} \mathbf{1}^\top, \mathbf{P} \rangle_\mu \\
&\quad - \frac{\eta_V}{K_P} \|\Delta_A\|^2 \langle \mathbf{V} - \boldsymbol{\mu} \mathbf{1}^\top, \mathbf{P} \rangle_\mu - \eta_V \frac{\langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu}{K_P} \\
&= \alpha_{V,\tau} + \eta_V K_Q (1 - \alpha_A \alpha_V) \alpha_A + \eta_V \frac{1 - \alpha_V}{T} - \eta_V \alpha_V \|\Delta_A\|^2 - \frac{\eta_V}{K_P} \langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu.
\end{aligned}$$

Similarly, for V , we have

$$\begin{aligned}
\alpha_{A,\tau+1} &= \alpha_{A,\tau} - \eta_A \frac{\langle \hat{\nabla}_A \mathcal{L} + \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu}{K_Q} \\
&= \alpha_{A,\tau} - \frac{\eta_A}{K_Q} K_P \alpha_V (\alpha_V \alpha_A - 1) \left\langle \mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T}, \mathbf{Q} \right\rangle_\mu \\
&\quad - \frac{\eta_A}{K_Q} \|\Delta_V\|_\mu^2 \left\langle \mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T}, \mathbf{Q} \right\rangle_\mu - \eta_A \frac{\langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu}{K_Q} \\
&= \alpha_{A,\tau} + \eta_A K_P (1 - \alpha_V \alpha_A) \alpha_V - \eta_A \alpha_A \|\Delta_V\|_\mu^2 - \frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu.
\end{aligned}$$

□

Proof of Lemma C.6. First, consider the dynamics of Δ_A , which is given by

$$\begin{aligned}
\Delta_{A,\tau+1} &= \Delta_A - \eta_A K_P \alpha_V^2 \Delta_A - \eta_A \|\Delta_V\|_\mu^2 \left(\mathbf{A} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \eta_A \mathbf{h}_A \\
&\quad + \left(\eta_A \alpha_A \|\Delta_V\|_\mu^2 + \frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu \right) \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right).
\end{aligned}$$

Decompose \mathbf{A} into $\tilde{\mathbf{A}} + \Delta_A$, rearrange terms, and we obtain

$$\begin{aligned}
\Delta_{A,\tau+1} &= \Delta_A - \eta_A K_P \alpha_V^2 \Delta_A - \eta_A \|\Delta_V\|_\mu^2 \Delta_A \\
&\quad - \eta_A \|\Delta_V\|_\mu^2 \left(\tilde{\mathbf{A}} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) + \eta_A \alpha_A \|\Delta_V\|_\mu^2 \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) \\
&\quad + \frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \eta_A \mathbf{h}_A.
\end{aligned}$$

Note that $\tilde{\mathbf{A}} - \mathbf{1} \mathbf{1}^\top / T = \alpha_A \mathbf{Q} + (1 - \alpha_A) \mathbf{1} \mathbf{1}^\top / T - \mathbf{1} \mathbf{1}^\top / T = \alpha_A (\mathbf{Q} - \mathbf{1} \mathbf{1}^\top / T)$. Hence, we have

$$\begin{aligned}
\Delta_{A,\tau+1} &= \Delta_A - \eta_A K_P \alpha_V^2 \Delta_A - \eta_A \|\Delta_V\|_\mu^2 \Delta_A \\
&\quad - \cancel{\eta_A \alpha_A \|\Delta_V\|_\mu^2 \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right)} + \eta_A \alpha_A \|\Delta_V\|_\mu^2 \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) \\
&\quad + \frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \eta_A \mathbf{h}_A
\end{aligned}$$

$$= \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2\right) \Delta_A + \frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \eta_A \mathbf{h}_A.$$

Recall that $\langle \Delta_A, \mathbf{Q} - \mathbf{1}_T^\top / T \rangle_\mu = 0$, $\|\mathbf{Q} - \mathbf{1}_T^\top / T\|_\mu^2 = K_Q$, and $\langle \mathbf{1}_T \mathbf{1}_N^\top / T, \mathbf{h}_A \rangle_\mu = 0$. Hence,

$$\begin{aligned} \|\Delta_{A,\tau+1}\|_\mu^2 &= \left\| \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2\right) \Delta_A + \frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu \left(\mathbf{Q} - \frac{\mathbf{1}_T \mathbf{1}_N^\top}{T} \right) - \eta_A \mathbf{h}_A \right\|_\mu^2 \\ &= \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2\right)^2 \|\Delta_A\|_\mu^2 + \left(\frac{\eta_A}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu\right)^2 K_Q + \eta_A^2 \|\mathbf{h}_A\|_\mu^2 \\ &\quad - 2\eta_A \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2\right) \langle \Delta_A, \mathbf{h}_A \rangle_\mu \\ &\quad - 2\frac{\eta_A^2}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu \langle \mathbf{Q}, \mathbf{h}_A \rangle_\mu \\ &= \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2\right)^2 \|\Delta_A\|_\mu^2 - 2\eta_A \left(1 - \eta_A K_P \alpha_V^2 - \eta_A \|\Delta_V\|_\mu^2\right) \langle \Delta_A, \mathbf{h}_A \rangle_\mu \\ &\quad + \eta_A^2 \|\mathbf{h}_A\|_\mu^2 - \frac{\eta_A^2}{K_Q} \langle \mathbf{h}_{A,\tau}, \mathbf{Q} \rangle_\mu^2. \end{aligned}$$

Now, consider Δ_V . Similar to the previous calculation, we have

$$\begin{aligned} \Delta_{V,\tau+1} &= \Delta_V - \eta_V \left(\alpha_A^2 K_Q + \frac{1}{T} \right) \Delta_V - \eta_V \|\Delta_A\|^2 \Delta_V - \eta_V \mathbf{h}_V \\ &\quad - \eta_V \frac{\alpha_V - 1}{T} (\mathbf{P} - \mu \mathbf{1}^\top) - \eta_V \|\Delta_A\|^2 (\tilde{\mathbf{V}} - \mu \mathbf{1}^\top) \\ &\quad - \left(\eta_V \frac{1 - \alpha_V}{T} - \eta_V \alpha_V \|\Delta_A\|^2 - \frac{\eta_V}{K_P} \langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu \right) (\mathbf{P} - \mu \mathbf{1}^\top) \\ &= \left(1 - \eta_V \left(\alpha_A^2 K_Q + \frac{1}{T} \right) - \eta_V \|\Delta_A\|^2\right) \Delta_V \\ &\quad + \frac{\eta_V}{K_P} \langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu (\mathbf{P} - \mu \mathbf{1}^\top) - \eta_V \mathbf{h}_V. \end{aligned}$$

Again, note that $\Delta_V \perp_\mu \mathbf{P} - \mu \mathbf{1}^\top$, $\|\mathbf{P} - \mu \mathbf{1}^\top\|_\mu^2 = K_P$, and $\langle \mu \mathbf{1}^\top, \mathbf{h}_V \rangle_\mu = 0$. Hence, we have

$$\begin{aligned} \|\Delta_{V,\tau+1}\|_\mu^2 &= \left(1 - \eta_V \left(\alpha_A^2 K_Q + \frac{1}{T} \right) - \eta_V \|\Delta_A\|^2\right)^2 \|\Delta_V\|_\mu^2 \\ &\quad + 2\eta_V \left(1 - \eta_V \left(\alpha_A^2 K_Q + \frac{1}{T} \right) - \eta_V \|\Delta_A\|^2\right) \langle \Delta_V, \mathbf{h}_V \rangle_\mu \\ &\quad - \frac{\eta_V^2}{K_P} \langle \mathbf{P}, \mathbf{h}_V \rangle_\mu^2 + \eta_V^2 \|\mathbf{h}_V\|_\mu^2. \end{aligned}$$

□

D Stage 1: signal boosting

In this section, we assume both α_V and α_A are close to 0. In this case, we can approximate Lemma C.5 with

$$\begin{aligned} \alpha_{V,\tau+1} &\approx \alpha_{V,\tau} + \eta \alpha_A + \eta \frac{1}{TK_Q} - \eta \alpha_V \frac{\|\Delta_A\|_\mu^2}{K_Q} - \eta \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_Q K_P}, \\ \alpha_{A,\tau+1} &\approx \alpha_{A,\tau} + \eta \alpha_V - \eta \alpha_A \frac{\|\Delta_V\|_\mu^2}{K_P} - \eta \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q}. \end{aligned}$$

We can also write this matrix form as

$$\begin{bmatrix} \alpha_{V,\tau+1} \\ \alpha_{A,\tau+1} \end{bmatrix} \approx \begin{bmatrix} \alpha_{V,\tau} \\ \alpha_{A,\tau} \end{bmatrix} + \eta \begin{bmatrix} -\|\Delta_A\|_\mu^2 / K_Q & 1 \\ 1 & -\|\Delta_V\|_\mu^2 / K_P \end{bmatrix} \begin{bmatrix} \alpha_{V,\tau} \\ \alpha_{A,\tau} \end{bmatrix}$$

$$+ \eta \begin{bmatrix} 1/(TK_Q) \\ 0 \end{bmatrix} - \eta \begin{bmatrix} \alpha_A \|\Delta_V\|_\mu^2 / (K_P K_Q) \\ \langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu / (K_P K_Q) \end{bmatrix}$$

Suppose that $\|\Delta_A\|_\mu^2 / K_Q$ and $\|\Delta_V\|_\mu^2 / K_P$ are both bounded by δ^2 . Then, we have

$$\begin{aligned} \alpha_{V,\tau+1} + \alpha_{A,\tau+1} &\geq (1 + \eta - 2\delta^2\eta) (\alpha_{V,\tau} + \alpha_{A,\tau}) + \eta \frac{1}{TK_Q} \\ &\quad - \eta \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_Q K_P} - \eta \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q}. \end{aligned} \quad (11)$$

As long as $\delta \ll 1$ and we choose a sufficiently large batch size so that the second line is bounded by $\eta/(2TK_Q)$, $\alpha_V + \alpha_A$ grows exponentially fast. Similarly, one can also bound the difference between α_V and α_A . Formally, we have the following lemma.

Lemma D.1 (Main result of Stage 1). *Define the end of Stage 1 as*

$$\mathcal{T}_1 := \min \left\{ \tau \geq 0 : \max\{\alpha_{V,\tau}, \alpha_{A,\tau}\} \geq \min \left\{ \frac{1}{2}, \frac{\Theta(1)}{QN} \right\} \right\}.$$

Suppose $\eta \leq 1/10$, $\eta_{V,\tau} = \eta/K_Q$, and $\eta_{A,\tau} = \eta/K_P$ for some $\eta \leq \min\{K_P, K_Q\}$. Let $B_\tau > 0$ be the number fresh samples we use at step τ . Suppose that $B_\tau \geq \tilde{O} \left(\frac{T^2 Q^4 N^5}{\min\{K_P^3, K_Q^3\}} \right)$ are chosen s.t. with probability $1 - \delta_\tau$

$$\max \left\{ \left| \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_Q K_P} \right|, \left| \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q} \right| \right\} \leq \frac{1}{4TK_Q}, \quad (12)$$

$$\max \left\{ \|\mathbf{h}_{V,\tau}\|_\mu, \|\mathbf{h}_{A,\tau}\|_\mu \right\} \leq \frac{\Theta(1) \min \left\{ K_Q^{1/2}, K_P^{1/2}, \frac{1}{\sqrt{QN^{1/2}}} \right\} \min\{K_Q, K_P\}}{T \log TK_Q} \quad (13)$$

Then, the following hold with probability at least $1 - \delta_P$:

(a) $\mathcal{T}_1 \leq \Theta(1) \log(TK_Q)/\eta$.

(b) Throughout Stage 1, $\|\Delta_A\|_\mu^2 \leq \Delta/T$ and $\|\Delta_V\|_\mu^2 \leq \Delta/T$, where

$$\Delta := \min \left\{ \frac{1}{4}K_Q, \frac{1}{4}K_P, \frac{\Theta(1)}{Q^4 N^3} \right\}.$$

(c) At \mathcal{T}_1 , we have $\alpha_V + \alpha_A = \Theta(\frac{1}{QN})$ and $|\alpha_V - \alpha_A| \leq \frac{\Theta(1)}{TK_Q}$.

(d) For all $\tau \leq \mathcal{T}_1$, we have $\|\mathbf{V}\|_\infty \leq \Theta(1)$, $\|\mathbf{A}\|_\infty \leq \Theta(1/Q)$.

The proof of this lemma is a large induction argument. We will first assume the bounds on $\|\Delta_A\|_\mu$ and $\|\Delta_V\|_\mu$ are true, so that the approximation (11) is valid. This will give us an upper bound on the length of Stage 1. Then, we show that within this many steps, the errors cannot exceed the given maximum values. Thus, the induction hypotheses are true and Lemma D.1 can be established.

Part I of the proof of Lemma D.1: Signal growth rate

Proof. Since $\alpha_V \alpha_A \leq \frac{1}{4}$ by definition of Stage I, we can rewrite Lemma C.5 as

$$\begin{aligned} \alpha_{V,\tau+1} + \alpha_{A,\tau+1} &\geq (1 + 3\eta/4) (\alpha_{V,\tau} + \alpha_{A,\tau}) + \eta \frac{3}{4TK_Q} \\ &\quad - \eta \alpha_V \frac{\|\Delta_A\|_\mu^2}{K_Q} - \eta \alpha_A \frac{\|\Delta_V\|_\mu^2}{K_P} - \eta \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_Q K_P} - \eta \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q} \\ &\geq (1 + 3\eta/4) (\alpha_{V,\tau} + \alpha_{A,\tau}) + \eta \frac{3}{4TK_Q} - \frac{1}{4} \eta \alpha_V - \frac{1}{4} \eta \alpha_A - \frac{1}{4TK_Q} \times 2 \end{aligned}$$

$$\geq \left(1 + \frac{\eta}{2}\right) (\alpha_{V,\tau} + \alpha_{A,\tau}) + \eta \frac{1}{4TK_Q},$$

where the second line comes from induction hypothesis (b) and (12). Recursively expand the RHS, and we obtain

$$\alpha_{V,\tau} + \alpha_{A,\tau} \geq \left(1 + \frac{\eta}{2}\right)^\tau \left(\alpha_{V,0} + \alpha_{A,0} + \frac{1}{2TK_Q}\right) - \frac{1}{2TK_Q} = \left(\left(1 + \frac{\eta}{2}\right)^\tau - 1\right) \frac{1}{2TK_Q}.$$

Since the RHS is upper bounded by 1 by definition of stage I, and we obtain

$$\mathcal{T}_1 \leq \Theta(1) \frac{\log(TK_Q)}{\eta}.$$

□

Part II of the proof of Lemma D.1: Upper bounds on $\|\Delta\|_\mu$

Proof. Recall from Lemma C.6 that

$$\begin{aligned} \|\Delta_{V,\tau+1}\|_\mu^2 &\leq \|\Delta_{V,\tau}\|_\mu^2 + 2\eta_V \|\Delta_{V,\tau}\|_\mu \|\mathbf{h}_{V,\tau}\|_\mu + \eta_V^2 \|\mathbf{h}_{V,\tau}\|_\mu^2, \\ \|\Delta_{A,\tau+1}\|_\mu^2 &\leq \|\Delta_{A,\tau}\|_\mu^2 + 2\eta_A \|\Delta_{A,\tau}\|_\mu \|\mathbf{h}_{A,\tau}\|_\mu + \eta_A^2 \|\mathbf{h}_{A,\tau}\|_\mu^2. \end{aligned}$$

By part I of the proof, Stage 1 takes at most $\Theta(1) \log(TK_Q)/\eta$ steps. Hence, it suffices to bound the increase of these $\|\Delta\|_\mu^2$ in this many steps. Recall the notation $\Delta = \min\{\frac{1}{4}K_Q, \frac{1}{4}K_P, \frac{\Theta(1)}{Q^{N^{1/2}}}\}$. By induction hypothesis (b), we have for all τ ,

$$\begin{aligned} \|\Delta_{V,\tau+1}\|_\mu^2 &\leq \|\Delta_{V,\tau}\|_\mu^2 + 2\eta_V \|\Delta_{V,\tau}\|_\mu \|\mathbf{h}_{V,\tau}\|_\mu + \eta_V^2 \|\mathbf{h}_{V,\tau}\|_\mu^2 \\ &\leq \|\Delta_{V,\tau}\|_\mu^2 + 2\eta_V \sqrt{\Delta/T} \|\mathbf{h}_{V,\tau}\|_\mu + \eta_V^2 \|\mathbf{h}_{V,\tau}\|_\mu^2 \\ &\leq \|\Delta_{V,\tau}\|_\mu^2 + \frac{\Theta(1)\eta\sqrt{\Delta/T}}{K_Q} \cdot \frac{\sqrt{\Delta/T} \min\{K_P, K_Q\}}{\log(TK_Q)} + \eta^2 \frac{\Theta(1)\Delta \cdot \min\{K_P^2, K_Q^2\}}{TK_Q^2 \log(TK_Q)^2} \\ \|\Delta_{A,\tau+1}\|_\mu^2 &\leq \|\Delta_{A,\tau}\|_\mu^2 + 2\eta_A \|\Delta_{A,\tau}\|_\mu \|\mathbf{h}_{A,\tau}\|_\mu + \eta_A^2 \|\mathbf{h}_{A,\tau}\|_\mu^2 \\ &\leq \|\Delta_{A,\tau}\|_\mu^2 + 2\eta_A \sqrt{\Delta/T} \|\mathbf{h}_{A,\tau}\|_\mu + \eta_A^2 \|\mathbf{h}_{A,\tau}\|_\mu^2 \\ &\leq \|\Delta_{A,\tau}\|_\mu^2 + \frac{\Theta(1)\eta\sqrt{\Delta/T}}{K_P} \cdot \frac{\sqrt{\Delta/T} \min\{K_P, K_Q\}}{\log(TK_Q)} + \eta^2 \frac{\Theta(1)\Delta \cdot \min\{K_P^2, K_Q^2\}}{TK_P^2 \log(TK_Q)^2} \end{aligned}$$

Since Stage I at most takes $\mathcal{T}_1 = \Theta(1) \log(TK_Q)/\eta$ steps, the increase of $\|\Delta_{A,\tau}\|_\mu^2$ and $\|\Delta_{V,\tau}\|_\mu^2$ in $\tau \leq \mathcal{T}_1$ are at most (since $\Delta_{A,0} = 0, \Delta_{V,0} = 0$):

$$\begin{aligned} \|\Delta_{V,t}\|_\mu^2 &\leq \left(\frac{\Theta(1)\eta\sqrt{\Delta/T}}{K_Q} \cdot \frac{\sqrt{\Delta/T} \min\{K_P, K_Q\}}{\log(TK_Q)} + \eta^2 \frac{\Theta(1)\Delta \cdot \min\{K_P^2, K_Q^2\}}{TK_Q^2 \log(TK_Q)^2} \right) t \\ &\leq \left(\frac{\Theta(1)\eta\sqrt{\Delta/T}}{K_Q} \cdot \frac{\sqrt{\Delta/T} \min\{K_P, K_Q\}}{\log(TK_Q)^2} + \eta^2 \frac{\Theta(1)\Delta \cdot \min\{K_P^2, K_Q^2\}}{TK_Q^2 \log(TK_Q)^2} \right) \mathcal{T}_1 \leq \Delta/T, \\ \|\Delta_{A,t}\|_\mu^2 &\leq \left(\frac{\Theta(1)\eta\sqrt{\Delta/T}}{K_P} \cdot \frac{\sqrt{\Delta/T} \min\{K_P, K_Q\}}{\log(TK_Q)} + \eta^2 \frac{\Theta(1)\Delta \cdot \min\{K_P^2, K_Q^2\}}{TK_P^2 \log(TK_Q)^2} \right) t \\ &\leq \left(\frac{\Theta(1)\eta\sqrt{\Delta/T}}{K_P} \cdot \frac{\sqrt{\Delta/T} \min\{K_P, K_Q\}}{\log(TK_Q)} + \eta^2 \frac{\Theta(1)\Delta \cdot \min\{K_P^2, K_Q^2\}}{TK_P^2 \log(TK_Q)^2} \right) \mathcal{T}_1 \leq \Delta/T. \end{aligned}$$

Therefore, we completed the induction for the error terms. □

Part III of the proof of Lemma D.1: Ending state

Proof. First, consider the distance between α_V and α_A . Similar to the part I of the proof, we rewrite Lemma C.5 as

$$\begin{aligned} \alpha_{V,\tau+1} - \alpha_{A,\tau+1} &= -(1 - \eta(1 - \alpha_V \alpha_A)) (\alpha_{V,\tau} - \alpha_{A,\tau}) + \eta \frac{1 - \alpha_V}{TK_Q} \\ &\quad - \eta \alpha_V \frac{\|\Delta_A\|_\mu^2}{K_Q} + \eta \alpha_A \frac{\|\Delta_V\|_\mu^2}{K_P} - \eta \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_Q K_P} + \eta \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q} \\ &= -(1 - \eta/2) (\alpha_{V,\tau} - \alpha_{A,\tau}) \pm O\left(\eta \frac{1}{TK_Q}\right). \end{aligned}$$

Thus, whenever $\alpha_{V,\tau} - \alpha_{A,\tau} \geq \Omega(1/(TK_Q))$, it will start to decrease. Since the amount of increase at each step is also upper bounded by $O(1/(TK_Q))$, this implies $|\alpha_{V,\tau} - \alpha_{A,\tau}| \leq O(1/(TK_Q))$. The other direction can be proved in the same way.

Finally, we bound the possible amount of overshoot. By the part I of the proof, we can also upper bound the signal term growth

$$\alpha_{V,\tau+1} + \alpha_{A,\tau+1} \leq (1 + 2\eta) (\alpha_{V,\tau} + \alpha_{A,\tau}) + O\left(\frac{\eta}{TK_Q}\right).$$

Since $\alpha_{V,\tau} + \alpha_{A,\tau} \leq 1$ for all $\tau \leq \mathcal{T}_1$, we have $\alpha_V + \alpha_A = \Theta(\frac{1}{Q_N})$ at time \mathcal{T}_1 . \square

Part IV of the proof of Lemma D.1: Upper bound on Infinity norm of V and A

Here we consider the upper bound of the weights \mathbf{V} and \mathbf{A} , which can be used in the concentration section below.

Proof. First, we upper bound the infinity norm of \mathbf{V}_τ .

$$\|\mathbf{V}_\tau\|_\infty = \|\tilde{\mathbf{V}}_\tau + \Delta_V\|_\infty \leq \|\tilde{\mathbf{V}}_\tau\|_\infty + \|\Delta_V\|_\infty \leq \|\tilde{\mathbf{V}}_\tau\|_\infty + \|\Delta_V\|_F.$$

and we can upper bound $\|\Delta_V\|_F$ by its μ -norm:

$$\|\Delta_V\|_\mu^2 = \langle \Delta_V, \Delta_V \text{diag}(\mu) \rangle \geq \frac{c}{N} \|\Delta_V\|_F^2.$$

Thus we have $\|\mathbf{V}_\tau\|_\infty \leq \|\tilde{\mathbf{V}}_\tau\|_\infty + \Theta(1)\sqrt{N} \|\Delta_V\|_\mu \leq \|\tilde{\mathbf{V}}_\tau\|_\infty + \Theta(\frac{1}{Q^2 N})$ by Induction hypothesis (b). And we can further bound $\|\tilde{\mathbf{V}}_\tau\|_\infty$:

$$\|\tilde{\mathbf{V}}_\tau\|_\infty = \|\alpha_{V,\tau} \mathbf{P} + (1 - \alpha_{V,\tau}) \mu \mathbf{1}^\top\|_\infty \leq \|\mathbf{P}\|_\infty + \|\mu \mathbf{1}^\top\|_\infty \leq \Theta(1).$$

Therefore, we have $\|\mathbf{V}_\tau\|_\infty \leq \Theta(1)$.

Similarly, for \mathbf{A} we have (since $\|\tilde{\mathbf{A}}_\tau\|_\infty = \|\mathbf{1}_T \mathbf{1}_N^\top / T + \alpha_{A,\tau} (\mathbf{Q} - \mathbf{1}_T \mathbf{1}_N^\top / T)\|_\infty \leq C/Q$.)

$$\|\mathbf{A}_\tau\|_\infty = \|\tilde{\mathbf{A}}_\tau + \Delta_A\|_\infty \leq \|\tilde{\mathbf{A}}_\tau\|_\infty + \|\Delta_A\|_\infty \leq \|\tilde{\mathbf{A}}_\tau\|_\infty + \|\Delta_A\|_F \leq \Theta(1/Q).$$

\square

Part V of the proof of Lemma D.1: Concentration

Finally, we need to ensure that with high probability, all the error terms \mathbf{h} cannot exceed the given bounds throughout $t \leq \mathcal{T}_1$. We use Lemma B.12 and Lemma B.13 to bound the concentration of the error terms.

By Lemma B.13 and union bound, we have that if $B_\tau \geq \frac{\Theta(1)\mathcal{T}_1 \cdot T^2 Q^4 N^5 \mathcal{T}_1 \log^2(TK_Q)}{\delta_\tau \min\{K_P^3, K_Q^3, 1\}}$, then with probability at least $1 - \delta_\tau/2$, the following holds for all $t \leq \mathcal{T}_1$:

$$\max \left\{ \|\mathbf{h}_{V,\tau}\|_\mu, \|\mathbf{h}_{A,\tau}\|_\mu \right\} \leq \frac{\Theta(1) \min \left\{ K_Q^{1/2}, K_P^{1/2}, \frac{1}{Q^2 N^{3/2}} \right\} \min\{K_Q, K_P\}}{\sqrt{T} \log TK_Q}$$

By Lemma B.12 and union bound, we have that if $B_\tau \geq \frac{\Theta(1)\mathcal{T}_1 \cdot T^2 \max\{N^2, Q^2\} N^4 \log\left(\frac{16 \log(TK_Q)}{\delta_\tau \eta}\right)}{K_P^2}$, then with probability at least $1 - \delta_\tau/2$, the following holds for all $t \leq \mathcal{T}_1$:

$$\left| \frac{\langle \mathbf{h}_V, \mathbf{P} \rangle_\mu}{K_P K_Q} \right| \leq \frac{1}{4TK_Q}, \quad \left| \frac{\langle \mathbf{h}_A, \mathbf{Q} \rangle_\mu}{K_P K_Q} \right| \leq \frac{1}{4TK_Q}.$$

And by union bound, we have that with probability at least $1 - \delta_\tau$, all the bounds above hold for all $t \leq \mathcal{T}_1$. Therefore, we conclude the proof of Lemma D.1.

E Stage 2: learning the model

In this Stage 2, we use a positive λ for the ℓ_1 -regularization. In this case, we can write the update rule of each $\mathbf{a}^{(k)}$ as

$$\begin{aligned} \mathbf{a}_{\tau+1}^{(k,')} &= \mathbf{a}_{\tau}^{(k)} - \frac{\eta_A}{\mu_k} \nabla_{\mathbf{a}^{(k)}}^{(B_{\tau})} l, & (\text{gradient descent step}), \\ \mathbf{a}_{\tau+1,t}^{(k,')} &= \begin{cases} a_{\tau+1,t}^{(k,')} - \lambda, & \text{if } a_{\tau+1,t}^{(k,')} \geq \lambda, \\ 0, & \text{if } |a_{\tau+1,t}^{(k,')}| \leq \lambda, \end{cases} & (\text{proximal step}), \\ \mathbf{a}_{\tau+1}^{(k)} &= \mathbf{a}_{\tau+1}^{(k,')} + \left(1 - \mathbf{1}^{\top} \mathbf{a}_{\tau+1}^{(k,')} \right) \frac{\mathbf{1}}{T}, & (\text{projection step}). \end{aligned} \quad (14)$$

For notational simplicity, we define

$$\mathbf{g}_{\tau}^{(k)} := -\eta_{\tau}^{-1} (\mathbf{a}_{\tau+1}^{(k)} - \mathbf{a}_{\tau}^{(k)}).$$

We further define $\mathbf{G}_{\lambda,\tau} = -\eta_A^{-1} (\mathbf{A}_{\tau+1} - \mathbf{A}_{\tau})$ as the full gradient of the matrix \mathbf{A} .

First, we will show that with appropriate rounding at the beginning of Stage 2, we can ensure $\mathbf{g}_{\tau}^{(k)} \approx \eta \mathbb{E} \hat{\nabla}_{\mathbf{a}^{(k)}} l$ using only $B_{\tau} \propto \log(T)$ fresh samples at each step (Section E.1).

E.1 Rounding and gradient denoising

Recall that the first step of the Stage 2 is rounding each $\mathbf{a}^{(k)}$ by setting all small coordinates to 0 and then projecting it back to the affine space the probability simplex lies. Since after Stage 1, there will be a separation between $a_t^{(k)}$ with $t \in \mathbf{q}^{(k)}$ and $t \notin \mathbf{q}^{(k)}$, this rounding step makes all $a_t^{(k)}$ with $t \notin \mathbf{q}^{(k)}$ have the same small value.

Also recall that we use ℓ_1 -regularization and proximal gradients in the Stage 2. Effectively, the ℓ_1 -regularization ensures those useless $a_t^{(k)}$ are always 0 (before projection). Then, similar to the first rounding step, projection will again make them have the same small value.

In this subsection, we formalize the above argument. We show that the rounding step can recover the support of $\mathbf{q}^{(k)}$, analyze its influence on α_A and the distance to the population subspace. Then, we analyze the effect of use of the proximal gradients and show that with $\text{poly}(N, Q) \log(T)$ fresh samples at each step, we can make sure the difference between update and the population update is small with high probability.

Lemma E.1 (Separation between noise and signal). *Assume that at the beginning of Stage 2, we have*

$$\alpha_V \geq \Theta \left(\frac{Q}{T} + Q\sqrt{N} \|\Delta_A\|_{\mu} \right).$$

Then, we can choose a threshold $\theta_{\text{tmp}} = \Theta(\alpha_V/Q)$ s.t. $a_t^{(k)} \leq \theta_{\text{tmp}}$ iff $q_t^{(k)} = 0$.

Proof. Note that there exists some universal constant $c > 0$ such that $q_t^{(k)} \geq c/Q$ for all nonzero $q_t^{(k)}$ and $\mu_k \geq c/N$ for all $k \in [N]$. Recall that the population process $\tilde{\mathbf{A}} = \alpha_A \mathbf{Q} + (1 - \alpha_A) \mathbf{1}_T \mathbf{1}_N^{\top} / T$. Hence, for all $k \in [N]$ and $t \in [T]$,

$$\begin{aligned} \tilde{a}_t^{(k)} &\leq 1/T, & \text{if } t \notin \mathbf{q}^{(k)}, \\ \tilde{a}_t^{(k)} &\geq c\alpha_V/Q, & \text{if } t \in \mathbf{q}^{(k)}. \end{aligned}$$

Then, note that

$$\begin{aligned} \|\Delta_A\|_{\mu}^2 &= \sum_{t=1}^T \sum_{k=1}^N \left(a_t^{(k)} - \tilde{a}_t^{(k)} \right)^2 \mu_k \\ &\geq \frac{c}{N} \sum_{t=1}^T \sum_{k=1}^N \left(a_t^{(k)} - \tilde{a}_t^{(k)} \right)^2 \geq \frac{c}{N} \|\text{Vec}(\tilde{\mathbf{A}} - \mathbf{A})\|_2^2 \geq \frac{c}{N} \|\text{Vec}(\tilde{\mathbf{A}} - \mathbf{A})\|_{\infty}^2. \end{aligned}$$

Hence, for any $k \in [N]$ and $t \in [T]$, we have $|a_t^{(k)} - \tilde{a}_t^{(k)}| \leq \|\Delta_A\|_\mu \sqrt{N/c}$. Combine these together, and we obtain

$$\begin{aligned} a_t^{(k)} &\leq 1/T + \|\Delta_A\|_\mu \sqrt{N/c}, & \text{if } t \notin \mathbf{q}^{(k)}, \\ a_t^{(k)} &\geq c\alpha_V/Q - \|\Delta_A\|_\mu \sqrt{N/c}, & \text{if } t \in \mathbf{q}^{(k)}. \end{aligned}$$

Hence, in order to get a separation,

$$1/T + \|\Delta_A\|_\mu \sqrt{N/c} \leq \frac{1}{2} \left(c\alpha_V/Q - \|\Delta_A\|_\mu \sqrt{N/c} \right) \iff \frac{2Q}{cT} + \frac{3Q\sqrt{N}}{c^{1.5}} \|\Delta_A\|_\mu \leq \alpha_V. \quad \square$$

Lemma E.2 (Effect of rounding). *Under the conditions of Lemma E.1, choose λ as in Lemma E.1, and set*

$$\begin{aligned} \mathbf{a}^{(k)} &\leftarrow \left(\mathbf{I} - \frac{\mathbf{1}\mathbf{1}^\top}{T} \right) \mathbf{a}^{(k)} \odot \left(\mathbb{1}_{\{a_t^{(k)} \geq \lambda\}} \right)_{t=1}^T + \frac{\mathbf{1}}{T} \\ &= \mathbf{a}^{(k)} \odot \left(\mathbb{1}_{\{a_t^{(k)} \geq \lambda\}} \right)_{t=1}^T + \left(1 - \sum_{t \in \mathbf{q}^{(k)}} a_t^{(k)} \right) \frac{\mathbf{1}}{T}. \end{aligned}$$

We have $\alpha_A \leftarrow \alpha_A + O(1)/T$ and $\|\Delta_A\|_\mu^2 \leftarrow \|\Delta_A\|_\mu^2 + O(1)/T$.

Proof. For notational simplicity, put $\mathbf{b}^{(k)} = \mathbf{a}^{(k)} \odot (\mathbb{1}_{\{a_t^{(k)} \geq \lambda\}})_{t=1}^T$. By Lemma E.1, we know $\mathbf{b}^{(k)}$ is supported within $\mathbf{q}^{(k)}$. Set $b^{(k)} = \sum_{t=1}^T b_t^{(k)}$. Then, we can write

$$\mathbf{a}^{(k)} \leftarrow \mathbf{b}^{(k)} + \left(1 - b^{(k)} \right) \frac{\mathbf{1}}{T}.$$

Recall that $\alpha_A = K_{AQ}/K_Q$ where $K_{AQ} = \langle \mathbf{A}, \mathbf{Q} \rangle_\mu = \sum_{k=1}^N \mu_k \langle \mathbf{a}^{(k)}, \mathbf{q}^{(k)} \rangle$. Hence, for each $k \in [N]$, we have

$$\begin{aligned} \langle \mathbf{a}^{(k)}, \mathbf{q}^{(k)} \rangle &\leftarrow \langle \mathbf{b}^{(k)}, \mathbf{q}^{(k)} \rangle + \left(1 - b^{(k)} \right) \left\langle \frac{\mathbf{1}}{T}, \mathbf{q}^{(k)} \right\rangle \\ &\leftarrow \langle \mathbf{a}^{(k)}, \mathbf{q}^{(k)} \rangle + \left(1 - b^{(k)} \right) \frac{1}{T}. \end{aligned}$$

As a result,

$$\alpha_A \leftarrow \alpha_A + \frac{1}{T} \sum_{k=1}^N \mu_k \left(1 - b^{(k)} \right) = \alpha_A + \frac{O(1)}{T}.$$

Now, consider the effect of projection on the distance to the population subspace. We will use subscript new to indicate values after rounding and use notations such as $\mathbf{a}^{(k)}$ to denote the values before rounding. Recall that $\|\Delta_A\|_\mu^2 = \sum_{k=1}^N \mu_k \|\mathbf{a}^{(k)} - \tilde{\mathbf{a}}^{(k)}\|^2$. We have

$$\begin{aligned} \left\| \mathbf{a}_{\text{new}}^{(k)} - \tilde{\mathbf{a}}_{\text{new}}^{(k)} \right\|^2 &= \left\| \mathbf{b}^{(k)} + (1 - b^{(k)}) \frac{\mathbf{1}}{T} - \frac{\mathbf{1}}{T} - \alpha_{A,\text{new}} \left(\mathbf{q}^{(k)} - \frac{\mathbf{1}}{T} \right) \right\|^2 \\ &= \left\| \mathbf{b}^{(k)} - \alpha_{A,\text{new}} \mathbf{q}^{(k)} - \left(b^{(k)} - \alpha_{A,\text{new}} \right) \frac{\mathbf{1}}{T} \right\|^2 \\ &= \left\| \mathbf{b}^{(k)} - \alpha_{A,\text{new}} \mathbf{q}^{(k)} \right\|^2 + \left(b^{(k)} - \alpha_{A,\text{new}} \right)^2 \frac{1}{T} \\ &\quad - 2 \left(b^{(k)} - \alpha_{A,\text{new}} \right) \left\langle \mathbf{b}^{(k)} - \alpha_{A,\text{new}} \mathbf{q}^{(k)}, \frac{\mathbf{1}}{T} \right\rangle \\ &= \left\| \mathbf{b}^{(k)} - (\alpha_A + O(1)/T) \mathbf{q}^{(k)} \right\|^2 - \left(b^{(k)} - \alpha_{A,\text{new}} \right)^2 \frac{1}{T}. \end{aligned}$$

Note that

$$\left\| \mathbf{a}^{(k)} - \tilde{\mathbf{a}}^{(k)} \right\|^2 = \left\| \mathbf{a}^{(k)} - \alpha_A \mathbf{q}^{(k)} - (1 - \alpha_A) \frac{\mathbf{1}}{T} \right\|^2 = \left\| \mathbf{a}^{(k)} - \alpha_A \mathbf{q}^{(k)} \right\|^2 - (1 - \alpha_A)^2 \frac{1}{T}$$

$$\geq \left\| \mathbf{b}^{(k)} - \alpha_A \mathbf{q}^{(k)} \right\|^2 - (1 - \alpha_A)^2 \frac{1}{T}.$$

Hence,

$$\begin{aligned} & \left\| \mathbf{a}_{\text{new}}^{(k)} - \tilde{\mathbf{a}}_{\text{new}}^{(k)} \right\|^2 - \left\| \mathbf{a}^{(k)} - \tilde{\mathbf{a}}^{(k)} \right\|^2 \\ & \leq \left\| \mathbf{b}^{(k)} - \alpha_A \mathbf{q}^{(k)} - \frac{O(1)}{T} \mathbf{q}^{(k)} \right\|^2 - \left\| \mathbf{b}^{(k)} - \alpha_A \mathbf{q}^{(k)} \right\|^2 + \frac{O(1)}{T} \leq \frac{O(1)}{T}. \end{aligned}$$

Thus, $\|\Delta_A\|_\mu^2 \leftarrow \|\Delta_A\|_\mu^2 + O(1)/T$. \square

As we have seen in Stage 1, the norm of $\nabla_{\mathbf{a}^{(k)}} l$ can scale linearly with T . Hence, in order to make $\mathbf{h}_{\mathbf{a}^{(k)}, \tau}$ has size $O(1)$ in terms of $\|\cdot\|_2$, it is necessary to use $\text{poly}(T)$ samples, which is undesirable. However, note that all entries of $\nabla_{\mathbf{a}^{(k)}} l$ are bounded, and therefore are subgaussian. Hence, for each entry of $\nabla_{\mathbf{a}^{(k)}} l$, with $O(\log T)$ samples, we can make sure the relative error is small with probability at least $1 - 1/\text{poly}(T)$. By union bound, this means the $\|\cdot\|_\infty$ error can be made small using only $\log(T)$ samples. Note that there is a separation between the signal and noise parts of $\mathbb{E} \nabla_{\mathbf{a}^{(k)}} l$. This implies that we can distinguish them using $\log(T)$ samples and directly remove the noise part. Formally, we have the following lemma.

Lemma E.3 (Gradient denoising). *Given $\varepsilon \in \left(\Theta\left(\frac{Q^2 N^3}{\sqrt{T}}\right), 0.1 \right)$. Suppose that for any $k, m, n \in [N]$, $|\mathbf{a}_t^{(k)}| \leq 1/T$ if $t \notin \mathbf{q}^{(k)}$, $|V_{n,m}| \leq O(1)$, and $\|\Delta_A\|_\mu \leq c(1 - \alpha_V)/(\sqrt{N}Q)$ and $\|\Delta_V\|_\mu^2 \leq c(1 - \alpha_V)K_P$ for some small constant $c > 0$. For the target accuracy $\varepsilon_{\text{Tmp}} = \Theta\left(\frac{\varepsilon}{QN^2}\right)$, we have with probability $1 - \delta_\tau$. If we choose*

$$B_\tau \geq C \frac{N^8 Q^4}{\varepsilon^2 \alpha_V^2 K_P^2} \log\left(\frac{CNT}{\delta_\tau}\right),$$

for some large constant $C > 0$. Then, for each $k \in [N]$, with probability at least $1 - \delta_{\text{Tmp}}$, we have

$$\begin{aligned} \partial_{\mathbf{a}_t^{(k)}}^{(B_\tau)} l &= (1 \pm \varepsilon_{\text{Tmp}}) \mathbb{E} \partial_{\mathbf{a}_t^{(k)}} l \geq \Theta\left(\frac{\alpha_V K_P}{NQ}\right) & \forall t \in \mathbf{q}^{(k)}, \\ \partial_{\mathbf{a}_t^{(k)}}^{(B_\tau)} l &= O\left(\frac{1}{T} + \varepsilon_{\text{Tmp}} \frac{\alpha_V K_P}{NQ}\right) = O\left(\frac{\varepsilon K_P}{N^3 Q^2}\right) & \forall t \notin \mathbf{q}^{(k)}. \end{aligned}$$

Proof. For $s \in [T]$, recall that

$$\begin{aligned} \partial_{\mathbf{a}_s^{(k)}} l &= \mathbb{1}\{x_{T+1} = k\} (\mathbf{V} \mathbf{e}_{x_s})^\top (\mathbf{V} \mathbf{X} \mathbf{a} - \mathbf{e}_{x_o}), \\ \mathbb{E} \partial_{\mathbf{a}_s^{(k)}} l &= \mu_k K_V a_s - \mu_k K_{VP} q_s. \end{aligned}$$

First, consider the expectations. If $s \notin \mathbf{q}^{(k)}$, then by our assumption, we have $|\mathbb{E} \partial_{\mathbf{a}_s^{(k)}} l| = \mu_k K_V a_s = O(K_V/(NT))$. Meanwhile, for $s \in \mathbf{q}^{(k)}$, we have

$$\begin{aligned} -\mathbb{E} \partial_{\mathbf{a}_s^{(k)}} l &\geq \mu_k K_{VP} (q_s - a_s) + \mu_k (K_V - K_{VP}) a_s \\ &\geq \mu_k \left(\alpha_V (1 - \alpha_V) K_P q_s - \alpha_V K_P \|\Delta_A\|_\mu \sqrt{N} - \|\Delta_V\|_\mu^2 a_s \right). \end{aligned}$$

As a result, we have for any $k \in [N]$ and $s \in \mathbf{q}^{(k)}$,

$$-\mathbb{E} \partial_{\mathbf{a}_s^{(k)}} l \geq c \frac{\alpha_V K_P}{NQ} \gg O\left(\frac{K_V}{NT}\right) \iff \begin{cases} \|\Delta_A\|_\mu \leq c \frac{1 - \alpha_V}{Q \sqrt{N}}, \\ \|\Delta_V\|_\mu^2 \leq c(1 - \alpha_V) K_P, \end{cases}$$

for some small constant $c > 0$. Then, for the size of each entry, we have

$$\left| \partial_{\mathbf{a}_s^{(k)}} l \right| \leq |(\mathbf{V} \mathbf{e}_{x_s})^\top \mathbf{V} \mathbf{X} \mathbf{a}| + |(\mathbf{V} \mathbf{e}_{x_s})^\top \mathbf{e}_{x_o}| \leq \sum_{t \in \mathbf{q}^{(k)}} a_t \sum_{n=1}^N |V_{n,x_s} V_{n,x_t}| + |V_{x_o,x_s}| \leq N.$$

Therefore, $\partial_{a_s^{(k)}} l$ is $N^2 K_{\text{Tmp}}^4$ -subgaussian, whence $\partial_{a_s^{(k)}}^{(B_\tau)} l$ is $N^2 K_{\text{Tmp}}^4 / B_\tau$ -subgaussian. As a result, for each $\xi > 0$, we have

$$\mathbb{P} \left[\left| \partial_{a_s^{(k)}}^{(B_\tau)} l - \mathbb{E} \partial_{a_s^{(k)}} l \right| \geq \xi \right] \leq 2 \exp \left(-\frac{\xi^2 B_\tau}{N^2} \right).$$

Apply union bound and we get

$$\mathbb{P} [\|\nabla_A l - \mathbb{E} \nabla_A l\|_\infty \geq \xi] \leq 2TN \exp \left(-\frac{\xi^2 B_\tau}{N^2} \right) = \exp \left(\log(2NT) - \frac{\xi^2 B_\tau}{N^2} \right).$$

Recall that the separation between the expectations is $c\alpha_V K_P / (NQ)$. Hence, it suffices to choose $\xi = c\alpha_V K_P / (2NQ)$. Then, to make the failure probability at most δ_{Tmp} , we can choose B_τ as follows:

$$\exp \left(\log(2NT) - \frac{\xi^2 B_\tau}{N^2} \right) \leq \delta_\tau \iff B_\tau \geq C \frac{N^4 Q^2}{\alpha_V^2 K_P^2} \log \left(\frac{CNT}{\delta_\tau} \right),$$

for some large constant $C > 0$. Then, to boost the accuracy from $1/2$ to ε_{Tmp} , it suffices to increase the batch size to $C \frac{N^4 Q^2}{\varepsilon_{\text{Tmp}}^2 \alpha_V^2 K_P^2} \log \left(\frac{CNT}{\delta_\tau} \right)$. \square

Using this lemma, we can pick $\lambda = \Theta\left(\frac{\varepsilon K_P}{Q^2 N^3}\right)$ to sparsify our proximal gradient. Notice that our proximal gradient is a biased estimate of the true preconditioned gradient, but the separation guarantees that it is possible to make the error controllable. The following lemma calculates the error each proximal step introduces. Here we define $\hat{\mathbf{h}}_{A,\tau} = \mathbf{G}_{\lambda,\tau}^{(k)} - \mathbb{E} \hat{\nabla}_A l$ instead of \mathbf{h} because of the bias introduced by the proximal gradient.

Lemma E.4. *Under the same setting of Lemma E.3, if the batch size*

$$B_\tau \geq \max \left\{ C \frac{N^8 Q^4}{\varepsilon^2 \alpha_V^2 K_P^2} \log \left(\frac{2CNT}{\delta_\tau} \right), \frac{\Theta(N^6 Q^3)}{\delta \varepsilon^2} \right\},$$

the noise attention score $|a_t^{(k)}| \leq O(1/T)$ for all $t \notin \mathbf{q}^{(k)}$ and all $k \in [N]$, and $\lambda = \Theta\left(\frac{\varepsilon K_P}{Q^2 N^3}\right)$, then with probability $1 - \delta_\tau$, the gradient error at iteration τ

$$\|\hat{\mathbf{h}}_{A,\tau}\|_\mu := \|\mathbf{G}_{\lambda,\tau} - \mathbb{E} \hat{\nabla}_A l\|_\mu \leq O\left(\frac{\varepsilon}{QN^2}\right).$$

Proof. For notational simplicity, we drop the superscript k . The goal is to estimate the difference between $\mathbb{E} \hat{\nabla}_A l$ and $\mathbf{g}_{\lambda,\tau}$ by calculating the magnitude of the bias of the proximal gradient together with the concentration error.

We consider the population gradient first. Since we have $\{a_{t,\tau}\}_{t \notin \mathbf{q}}$ are all the same in Stage 2, we can write

$$\mathbf{a}_\tau = [\mathbf{a}_\tau]_{\mathbf{q}} + \frac{(1 - b_\tau) \mathbf{1}_{\mathbf{q}^c}}{T - Q} \quad \text{where} \quad b_\tau = \sum_{t \in \mathbf{q}} a_t,$$

and recall the projected preconditioned gradient for A .

$$\mathbb{E} \hat{\nabla}_A l = \left(\|\mathbf{V}\|_\mu^2 - \|\boldsymbol{\mu}\|^2 \right) \left(\mathbf{a}^{(k)} - \frac{\mathbf{1}}{T} \right) - \left(\langle \mathbf{V}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2 \right) \left(\mathbf{q}^{(k)} - \frac{\mathbf{1}}{T} \right).$$

Therefore, we have

$$\mathbb{E} \hat{\nabla}_A l = [\mathbb{E} \hat{\nabla}_A l]_{\mathbf{q}} + K_V \frac{(1 - b_\tau) \mathbf{1}_{\mathbf{q}^c}}{T - Q} + (\|\mathbf{V}\|_\mu^2 - \langle \mathbf{V}, \mathbf{P} \rangle_\mu) \frac{\mathbf{1}_{\mathbf{q}^c}}{T}.$$

Now, we consider \mathbf{g} by calculating the expression of $\mathbf{a}_{\tau+1}$. First, note that by Lemma E.3, we have $a''_{t,\tau+1} = (a'_{t,\tau+1} - \lambda) \mathbb{1}\{t \in \mathbf{q}\}$. This implies that $\{a_{t,\tau+1}\}_{t \notin \mathbf{q}}$ are all the same and the value is at most $1/T$. Moreover, it also implies that it suffices to focus on $[\mathbf{a}''_{\tau+1}]_{\mathbf{q}}$, for which we have

$$[\mathbf{a}''_{\tau+1}]_{\mathbf{q}} = [\mathbf{a}'_{\tau+1}]_{\mathbf{q}} - \lambda \mathbf{1}_{\mathbf{q}} = [\mathbf{a}_\tau]_{\mathbf{q}} - \frac{\eta}{\mu_k} \left[\nabla_a^{(B_\tau)} l \right]_{\mathbf{q}} - \lambda \mathbf{1}_{\mathbf{q}}.$$

Therefore, $\sum_{t \in q} a'_{t, \tau+1} = b_\tau - \frac{\eta}{\mu_k} \sum_{t \in q} \partial_{a_t}^{(B_\tau)} l - Q\lambda$ and

$$\begin{aligned} \mathbf{a}_{\tau+1} &= [\mathbf{a}_\tau]_q - \frac{\eta}{\mu_k} \left[\nabla_{\mathbf{a}}^{(B_\tau)} l \right]_q - \lambda \mathbf{1}_q + \left(1 - b_\tau + \frac{\eta}{\mu_k} \sum_{t \in q} \partial_{a_t}^{(B_\tau)} l + Q\lambda \right) \frac{\mathbf{1}}{T} \\ &= \mathbf{a}_\tau - \frac{(1 - b_\tau) \mathbf{1}_{q^c}}{T - Q} - \frac{\eta}{\mu_k} \left[\nabla_{\mathbf{a}}^{(B_\tau)} l \right]_q - \lambda \mathbf{1}_q + \left(1 - b_\tau + \frac{\eta}{\mu_k} \sum_{t \in q} \partial_{a_t}^{(B_\tau)} l + Q\lambda \right) \frac{\mathbf{1}}{T}. \end{aligned}$$

Thus, we can write an explicit update

$$\begin{aligned} \mathbf{g}_{\lambda, \tau} &= \frac{1}{\mu_k} \left[\nabla_{\mathbf{a}}^{(B_\tau)} l \right]_q + \frac{\lambda}{\eta} \mathbf{1}_q - \left(1 - b_\tau + \frac{\eta}{\mu_k} \sum_{t \in q} \partial_{a_t}^{(B_\tau)} l + Q\lambda \right) \frac{\mathbf{1}}{\eta T} + \frac{(1 - b_\tau) \mathbf{1}_{q^c}}{\eta(T - Q)} \\ &= \left[\hat{\nabla}_{\mathbf{a}}^{(B_\tau)} l \right]_q + \frac{\mathbf{1}_q \mathbf{1}^\top \nabla_{\mathbf{a}}^{(B_\tau)} l}{\mu_k T} - \left(1 - b_\tau + \frac{\eta}{\mu_k} \sum_{t \in q} \partial_{a_t}^{(B_\tau)} l + Q\lambda \right) \frac{\mathbf{1}}{\eta T} \\ &\quad + \frac{(1 - b_\tau) \mathbf{1}_{q^c}}{\eta(T - Q)} + \frac{\lambda}{\eta} \mathbf{1}_q \end{aligned}$$

Then the gradient error at step τ can be decomposed into:

$$\begin{aligned} \mathbf{g}_{\lambda, \tau} - \mathbb{E} \hat{\nabla}_{\mathbf{a}} l &= \left[\hat{\nabla}_{\mathbf{a}}^{(B_\tau)} l - \mathbb{E} \hat{\nabla}_{\mathbf{a}} l \right]_q && \text{(Concentration error)} \\ &\quad + \frac{\mathbf{1}_q \mathbf{1}^\top \nabla_{\mathbf{a}}^{(B_\tau)} l}{\mu_k T} - \left(1 - b_\tau + \frac{\eta}{\mu_k} \sum_{t \in q} \partial_{a_t}^{(B_\tau)} l + Q\lambda \right) \frac{\mathbf{1}}{\eta T} + \frac{\lambda}{\eta} \mathbf{1}_q \\ &\quad + \frac{(1 - b_\tau) \mathbf{1}_{q^c}}{\eta(T - Q)} - K_V \frac{(1 - b_\tau) \mathbf{1}_{q^c}}{T - Q} - (\|V\|_\mu^2 - \langle V, P \rangle_\mu) \frac{\mathbf{1}_{q^c}}{T} && \text{(Gradient bias error)} \end{aligned}$$

Here the gradient bias error can be further simplified to

$$\left[\frac{1}{\mu_k T} \sum_{t \in q} \partial_{a_t}^{(B_\tau)} l - \frac{1 - b_\tau + Q\lambda}{\eta T} + \frac{\lambda}{\eta} \right] \mathbf{1}_q \quad (1^*)$$

$$+ \left[\frac{Q(1 - b_\tau - Q\lambda) - Q\lambda T - \eta T K_V}{\eta T(T - Q)} - \frac{\sum_{t \in q} \partial_{a_t}^{(B_\tau)} l}{\mu_k T} - \frac{\|V\|_\mu^2 - \langle V, P \rangle_\mu}{T} \right] \mathbf{1}_{q^c} \quad (2^*)$$

First, we estimate the concentration error. Similar to Lemma B.10, we first upper bound the infinity norm of the gradient. Consider the maximum absolute value in the original gradients. For $\nabla_{\mathbf{a}^{(k)}} l$, we have

$$\begin{aligned} &\left\| \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}X)^\top (\mathbf{V}X \mathbf{a}^{(k)} - \mathbf{e}_{x_o}) \right\|_\infty \\ &\leq \left\| \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}X)^\top (\mathbf{V}X \mathbf{a}^{(k)}) \right\|_\infty + \left\| \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}X)^\top (\mathbf{e}_{x_o}) \right\|_\infty \\ &\leq \max_{s \in [T]} \left| \sum_{t=1}^T a_t (\mathbf{V}e_{x_s})^\top \mathbf{V}e_{x_t} \right| + \max_{s \in [T]} |(\mathbf{V}e_{x_s})^\top \mathbf{e}_{x_o}| \end{aligned}$$

The first term can be upper-bounded in the following way:

$$\begin{aligned} \max_{s \in [T]} \left| \sum_{t=1}^T a_t (\mathbf{V}e_{x_s})^\top \mathbf{V}e_{x_t} \right| &\leq \left(\sum_{t \in q} |a_t| + \sum_{t \notin q} |a_t| \right) \max_{s, t} |(\mathbf{V}e_{x_s})^\top \mathbf{V}e_{x_t}| \\ &\leq \Theta(1) \max_{s, t} |(\mathbf{V}e_{x_s})^\top \mathbf{V}e_{x_t}| \end{aligned}$$

The second inequality is due to $|a_t^{(k)}| \leq O(1/T)$ for $t \notin q$. Since $V_{n,m} \leq O(1)$, we have $\max_{s, t} |(\mathbf{V}e_{x_s})^\top \mathbf{V}e_{x_t}|$ upper bounded by $O(1)$. Therefore

$$\max_{s \in [T]} \left| \sum_{t=1}^T a_t (\mathbf{V}e_{x_s})^\top \mathbf{V}e_{x_t} \right| \leq O(1)$$

And similarly, the second term $\max_{s \in [T]} |(\mathbf{V} \mathbf{e}_{x_s})^\top \mathbf{e}_{x_o}|$ can be bounded by $O(1)$ because the infinity norm of \mathbf{V} is also upper bounded by $O(1)$. Therefore, we know $\|\nabla_{\mathbf{a}^{(k)}} l\|_\infty \leq O(1)$.

Now we consider the preconditioned gradient $\hat{\nabla}_{\mathbf{a}}^{(B_\tau)} l$:

$$\begin{aligned} \|\hat{\nabla}_{\mathbf{a}}^{(B_\tau)} l\|_\infty &= \left\| \frac{1}{\mu_k} \left(\mathbf{I}_T - \frac{\mathbf{1}_T \mathbf{1}_T^\top}{T} \right) \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l \right) \right\|_\infty \\ &\leq \left\| \frac{1}{\mu_k} \mathbf{I}_T \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l \right) \right\|_\infty + \left\| \frac{1}{\mu_k} \left(\frac{\mathbf{1}_T \mathbf{1}_T^\top}{T} \right) \left(\nabla_{\mathbf{a}^{(k)}}^{(B_\tau)} l \right) \right\|_\infty \leq O(N) \end{aligned}$$

since $\mu_k \geq \frac{c}{N}$ for all $k \in [N]$. Now since $\left[\hat{\nabla}_{\mathbf{a}}^{(B_\tau)} l \right]_{\mathbf{q}}$ is Q -sparse, we have $\mathbb{E} \left\| \left[\hat{\nabla}_{\mathbf{A}} l \right]_{\mathbf{q}} \right\|_\mu^2 \leq O(QN^2)$.

By Lemma B.9, when $B_\tau \geq \frac{\Theta(N^6 Q^3)}{\delta \varepsilon^2}$, with probability $1 - \frac{\delta_\tau}{2}$,

$$\left\| \left[\hat{\nabla}_{\mathbf{a}}^{(B_\tau)} l - \mathbb{E} \hat{\nabla}_{\mathbf{a}} l \right]_{\mathbf{q}} \right\| \leq O\left(\frac{\varepsilon}{QN^2} \right)$$

Then, consider the gradient bias term. With the selected $\lambda = \Theta\left(\frac{\varepsilon K_P}{Q^2 N^3}\right)$ and $\left\| \partial_{a_t}^{(B_\tau)} l \right\| \leq O\left(\frac{\varepsilon K_P}{N^3 Q^2}\right)$ for $t \neq \mathbf{q}$, with probability $1 - \delta_\tau/2$ we have the μ -norm of first term (since $K_P \leq O(N)$):

$$\begin{aligned} \|(1^*)\| &\leq \left(\left\| \frac{1}{\mu_k T} \sum_{t \neq \mathbf{q}} \partial_{a_t}^{(B_\tau)} l \right\| + \frac{Q\lambda}{T} + \lambda \right) \cdot \sqrt{Q} \\ &\leq O\left(\frac{\varepsilon}{QN^2}\right) + O\left(\frac{\varepsilon}{TN^2 \sqrt{Q}}\right) + O\left(\frac{\varepsilon}{Q^{3/2} N^2}\right) \leq O\left(\frac{\varepsilon}{QN^2}\right). \end{aligned}$$

and the second term can be upper-bounded by

$$\|(2^*)\| \leq \left(O\left(\frac{QN}{T}\right) + O\left(\frac{QN}{T}\right) + O\left(\frac{N}{T}\right) \right) \cdot \sqrt{T} \leq O\left(\frac{\varepsilon}{QN^2}\right).$$

since $\frac{1}{\sqrt{T}} \leq O\left(\frac{\varepsilon}{Q^2 N^3}\right)$, $\partial_{a_t}^{(B_\tau)} l = O(1)$, and $\|\mathbf{V}\|_\mu^2 \leq N$.

Combine all three terms and by union bound, we have with probability $1 - \delta_\tau$

$$\begin{aligned} \|\hat{\mathbf{h}}_{\mathbf{A}, \tau}\| &= \left\| \mathbf{G}_{\lambda, \tau}^{(k)} - \mathbb{E} \hat{\nabla}_{\mathbf{A}} l \right\|_\mu \\ &= \sqrt{\sum_{k=1}^N \mu_k \|\mathbf{g}_{\lambda, \tau}^{(k)} - \mathbb{E} \hat{\nabla}_{\mathbf{a}^{(k)}} l\|^2} \leq O\left(\frac{\varepsilon}{QN^2}\right) \end{aligned}$$

since $\mu_k = \Theta(1/N)$. □

E.2 Model aligning and the decrease of the errors

First, we show that the signal will continue to grow and approximation error will decrease. This decouples the error at the end of Stage 1 and the final error. In particular, we show that eventually we will have $\alpha_V + \alpha_A \approx 2$, $\|\Delta_{\mathbf{A}}\|_\mu^2 \approx 0$ and $\|\Delta_V\|_\mu^2 \approx 0^5$.

For notational simplicity, define $\delta_A^2 = \|\Delta_{\mathbf{A}}\|_\mu^2 / K_Q$ and $\delta_V^2 = \|\Delta_V\|_\mu^2 / K_P$. Recall Lemma C.5 and Lemma C.6 and that we choose $\eta_V = \eta / K_Q$, $\eta_A = \eta / K_P$. The dynamics of the signals and the errors can be described using⁶

$$\alpha_{V, \tau+1} = \alpha_{V, \tau} + \eta (1 - \alpha_A \alpha_V) \alpha_A + \frac{\eta}{K_Q} \frac{1 - \alpha_V}{T} - \eta \alpha_V \delta_A^2 - \eta \frac{\langle \mathbf{h}_{V, \tau}, \mathbf{P} \rangle_\mu}{K_P K_Q},$$

⁵However, we cannot ensure $\alpha_A \approx \alpha_V$ since $\alpha_A - \alpha_V$ is not contractive toward the end of training due to the magnitude of gradient noise. This issue can be fixed by a final rounding stage (See Appendix F).

⁶Here $\mathbf{h}_{\mathbf{A}, \tau} \rightarrow \hat{\mathbf{h}}_{\mathbf{A}, \tau}$ because each gradient step is changed to the ℓ_1 -regularized gradient. It does not change the main parts in the population process.

$$\alpha_{A,\tau+1} = \alpha_{A,\tau} + \eta(1 - \alpha_V \alpha_A) \alpha_V - \eta \alpha_A \delta_V^2 - \eta \frac{\langle \hat{\mathbf{h}}_{A,\tau}, \mathbf{Q} \rangle_\mu}{K_P K_Q},$$

and

$$\begin{aligned} \delta_{A,\tau+1}^2 &= \left(1 - \eta \alpha_V^2 - \eta \delta_A^2\right)^2 \delta_A^2 - 2\eta \left(1 - \eta \alpha_V^2 - \eta \delta_V^2\right) \frac{\langle \Delta_A, \hat{\mathbf{h}}_{A,\tau} \rangle_\mu}{K_P K_Q} \\ &\quad - \eta^2 \frac{\langle \hat{\mathbf{h}}_{A,\tau}, \mathbf{Q} \rangle_\mu^2}{K_P^2 K_Q^2} + \eta^2 \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu^2}{K_P^2 K_Q}, \\ \delta_{V,\tau+1}^2 &= \left(1 - \eta \left(\alpha_A^2 + \frac{1}{K_Q T} + \delta_A^2\right)\right)^2 \delta_V^2 + 2\eta \left(1 - \eta \left(\alpha_A^2 + \frac{1}{K_Q T} + \delta_A^2\right)\right) \frac{\langle \Delta_V, \mathbf{h}_V \rangle_\mu}{K_P K_Q} \\ &\quad - \eta^2 \frac{\langle \mathbf{P}, \mathbf{h}_V \rangle_\mu^2}{K_Q^2 K_P^2} + \eta^2 \frac{\|\mathbf{h}_V\|_\mu^2}{K_P K_Q^2}. \end{aligned}$$

E.2.1 Lemmas for the dynamics

Before we come to the final convergence analysis, we first simplify the dynamics with some basic lemmas.

Lemma E.5 (Dynamics of the errors). *For the errors, we have*

$$\begin{aligned} \delta_{A,\tau+1}^2 &\leq \exp\left(-2\eta \alpha_V^2\right) \delta_{A,\tau}^2 + 3\eta \left(\delta_A + \eta \frac{\|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}}\right) \frac{\|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}}, \\ \delta_{V,\tau+1}^2 &\leq \exp\left(-2\eta \alpha_A^2\right) \delta_{V,\tau}^2 + 3\eta \left(\delta_V + \eta \frac{\|\mathbf{h}_V\|_\mu}{\sqrt{K_P} K_Q}\right) \frac{\|\mathbf{h}_V\|_\mu}{\sqrt{K_P} K_Q}. \end{aligned}$$

Proof. First, we write

$$\begin{aligned} \delta_{A,\tau+1}^2 &= \left(1 - \eta \alpha_V^2 - \eta \delta_A^2\right)^2 \delta_A^2 - 2\eta \left(1 - \eta \alpha_V^2 - \eta \delta_V^2\right) \frac{\langle \Delta_A, \hat{\mathbf{h}}_A \rangle_\mu}{K_P K_Q} \\ &\quad - \eta^2 \frac{\langle \hat{\mathbf{h}}_{A,\tau}, \mathbf{Q} \rangle_\mu^2}{K_P^2 K_Q^2} + \eta^2 \frac{\|\hat{\mathbf{h}}_A\|_\mu^2}{K_P^2 K_Q} \\ &\leq \left(1 - \eta \alpha_V^2 - \eta \delta_A^2\right)^2 \delta_A^2 + 2\eta \frac{\delta_A \|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}} + 3\eta^2 \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu^2}{K_P^2 K_Q} \\ &\leq \left(1 - \eta \alpha_V^2 - \eta \delta_A^2\right)^2 \delta_A^2 + 3\eta \left(\delta_A + \frac{\eta}{K_P} \frac{\|\hat{\mathbf{h}}_A\|_\mu}{\sqrt{K_Q}}\right) \frac{\|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}}. \end{aligned}$$

For the first term, we have

$$\left(1 - \eta \alpha_V^2 - \eta \delta_A^2\right)^2 \leq \exp\left(-\eta \alpha_V^2 - \eta \delta_A^2\right) \leq \exp\left(-2\eta \alpha_V^2\right).$$

Thus, for δ_A , we have

$$\delta_{A,\tau+1}^2 \leq \exp\left(-2\eta \alpha_V^2\right) \delta_{A,\tau}^2 + 3\eta \left(\delta_A + \eta \frac{\|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}}\right) \frac{\|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}}.$$

Similarly, for δ_V , we have

$$\begin{aligned}\delta_{V,\tau+1}^2 &\leq \left(1 - \eta \left(\alpha_A^2 + \frac{1}{K_Q T} + \delta_A^2\right)\right)^2 \delta_V^2 + 3\eta \frac{\delta_V \|\mathbf{h}_V\|_\mu}{\sqrt{K_P K_Q}} + 3\eta^2 \frac{\|\mathbf{h}_V\|_\mu^2}{K_P K_Q^2} \\ &\leq \exp(-2\eta\alpha_A^2) \delta_V^2 + 3\eta \left(\delta_V + \eta \frac{\|\mathbf{h}_V\|_\mu}{\sqrt{K_P K_Q}}\right) \frac{\|\mathbf{h}_V\|_\mu}{\sqrt{K_P K_Q}}\end{aligned}$$

□

Lemma E.6 (Dynamics of $\alpha_A - \alpha_V$). *The difference between the signals evolves as follows*

$$\begin{aligned}(\alpha_{V,\tau+1} - \alpha_{A,\tau+1})^2 &\leq \exp(-2\eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A)^2 \\ &\quad + 8\eta \left(\frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right).\end{aligned}$$

Proof. First, we write

$$\begin{aligned}\alpha_{V,\tau+1} - \alpha_{A,\tau+1} &= \alpha_{V,\tau} - \alpha_{A,\tau} - \eta(1 - \alpha_A \alpha_V) (\alpha_V - \alpha_A) \\ &\quad + \frac{\eta}{K_Q} \frac{1 - \alpha_V}{T} - \eta \alpha_V \delta_A^2 - \eta \frac{\langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu}{K_P K_Q} + \eta \alpha_A \delta_V^2 + \eta \frac{\langle \hat{\mathbf{h}}_{A,\tau}, \mathbf{Q} \rangle_\mu}{K_P K_Q} \\ &=: (1 - \eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A) + \text{Tmp}.\end{aligned}$$

Therefore, we have

$$\begin{aligned}(\alpha_{V,\tau+1} - \alpha_{A,\tau+1})^2 &= ((1 - \eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A) + \text{Tmp})^2 \\ &= (1 - \eta(1 - \alpha_A \alpha_V))^2 (\alpha_V - \alpha_A)^2 \\ &\quad + 2(1 - \eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A) \text{Tmp} + \text{Tmp}^2 \\ &\leq \exp(-2\eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A)^2 \\ &\quad + 3|\alpha_V - \alpha_A| |\text{Tmp}| + \text{Tmp}^2.\end{aligned}$$

Then, for Tmp , we compute

$$0.5\eta^{-1} |\text{Tmp}| \leq \frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}}.$$

In particular, this implies $|\text{Tmp}| \leq 1$. Thus, we have

$$\begin{aligned}(\alpha_{V,\tau+1} - \alpha_{A,\tau+1})^2 &\leq \exp(-2\eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A)^2 + 4|\text{Tmp}| \\ &\leq \exp(-2\eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A)^2 \\ &\quad + 8\eta \left(\frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right).\end{aligned}$$

□

Lemma E.7. *Suppose that both α_V, α_A are at most 1. Then, we have*

$$\begin{aligned}1 - \frac{\alpha_{V,\tau+1} + \alpha_{A,\tau+1}}{2} &\leq \left(1 - \eta \frac{\alpha_V + \alpha_A}{2}\right) \left(1 - \frac{\alpha_V + \alpha_A}{2}\right) + \eta \frac{\alpha_V + \alpha_A}{2} (\delta_A^2 + \delta_V^2) \\ &\quad + \frac{\eta}{2} \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\eta}{2} \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}}.\end{aligned}$$

Proof. First, we write

$$\begin{aligned}\alpha_{V,\tau+1} + \alpha_{A,\tau+1} &= \alpha_{V,\tau} + \alpha_{A,\tau} + \eta(1 - \alpha_A \alpha_V)(\alpha_V + \alpha_A) \\ &\quad + \frac{\eta}{K_Q} \frac{1 - \alpha_V}{T} - \eta \alpha_V \delta_A^2 - \eta \alpha_A \delta_V^2 - \eta \frac{\langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu}{K_P K_Q} - \eta \frac{\langle \hat{\mathbf{h}}_{A,\tau}, \mathbf{Q} \rangle_\mu}{K_P K_Q} \\ &=: \alpha_{V,\tau} + \alpha_{A,\tau} + \eta \text{Tmp}_{\text{sig}} + \eta \text{Tmp}_{\text{err}}.\end{aligned}$$

For the signal growth, note that $\alpha_A \alpha_V \leq (\alpha_V^2 + \alpha_A^2)/2 \leq (\alpha_V + \alpha_A)/2$. Hence,

$$(1 - \alpha_A \alpha_V)(\alpha_V + \alpha_A) \geq \frac{1}{2}(\alpha_V + \alpha_A)(2 - \alpha_V - \alpha_A).$$

For the error terms, we have

$$\left| \frac{\langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu}{K_P K_Q} + \frac{\langle \hat{\mathbf{h}}_{A,\tau}, \mathbf{Q} \rangle_\mu}{K_P K_Q} \right| \leq \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P} K_Q} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}}.$$

Thus, we have

$$\begin{aligned}\alpha_{V,\tau+1} + \alpha_{A,\tau+1} &\geq \alpha_{V,\tau} + \alpha_{A,\tau} + \eta(\alpha_V + \alpha_A) \left(1 - \frac{\alpha_V + \alpha_A}{2} - \delta_A^2 - \delta_V^2\right) \\ &\quad - \eta \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P} K_Q} - \eta \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}},\end{aligned}$$

and, therefore,

$$\begin{aligned}1 - \frac{\alpha_{V,\tau+1} + \alpha_{A,\tau+1}}{2} &\leq \left(1 - \eta \frac{\alpha_V + \alpha_A}{2}\right) \left(1 - \frac{\alpha_V + \alpha_A}{2}\right) + \eta \frac{\alpha_V + \alpha_A}{2} (\delta_A^2 + \delta_V^2) \\ &\quad + \frac{\eta}{2} \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P} K_Q} + \frac{\eta}{2} \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}}.\end{aligned}$$

□

Lemma E.8. Put $\varepsilon_A = 1 - \alpha_A$ and $\varepsilon_V = 1 - \alpha_V$. When $|\varepsilon_A|, |\varepsilon_V| \leq 1/2$, we have

$$\begin{aligned}(\varepsilon_{A,\tau+1} + \varepsilon_{V,\tau+1})^2 &\leq \exp(-2\eta) (\varepsilon_A + \varepsilon_V)^2 + 8\eta |\varepsilon_A + \varepsilon_V| \varepsilon_A \varepsilon_V + 16\eta^2 \varepsilon_A^2 \varepsilon_V^2 \\ &\quad + 8\eta |\varepsilon_A + \varepsilon_V| \left(\frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P} K_Q} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right).\end{aligned}$$

Proof. Similar to the proof of the previous lemma, we write

$$\begin{aligned}\alpha_{V,\tau+1} + \alpha_{A,\tau+1} &= \alpha_{V,\tau} + \alpha_{A,\tau} + \eta(1 - \alpha_A \alpha_V)(\alpha_V + \alpha_A) \\ &\quad + \frac{\eta}{K_Q} \frac{1 - \alpha_V}{T} - \eta \alpha_V \delta_A^2 - \eta \alpha_A \delta_V^2 - \eta \frac{\langle \mathbf{h}_{V,\tau}, \mathbf{P} \rangle_\mu}{K_P K_Q} - \eta \frac{\langle \hat{\mathbf{h}}_{A,\tau}, \mathbf{Q} \rangle_\mu}{K_P K_Q} \\ &=: \alpha_{V,\tau} + \alpha_{A,\tau} + \eta \text{Tmp}_{\text{sig}} + \eta \text{Tmp}_{\text{err}}.\end{aligned}$$

For the signal term, we have

$$\begin{aligned}\text{Tmp}_{\text{sig}} &= (1 - (1 - \varepsilon_A)(1 - \varepsilon_V))(2 - \varepsilon_A - \varepsilon_V) \\ &= (2 - \varepsilon_A - \varepsilon_V)(\varepsilon_A + \varepsilon_V) - \varepsilon_A \varepsilon_V (2 - \varepsilon_A - \varepsilon_V).\end{aligned}$$

For the error term, we have

$$|\text{Tmp}_{\text{err}}| \leq \frac{1}{K_Q T} + 2\delta_A^2 + 2\delta_V^2 + \frac{2\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P} K_Q} + \frac{2\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}}.$$

Combine these together, and we obtain

$$\begin{aligned} \varepsilon_{A,\tau+1} + \varepsilon_{V,\tau+1} &= \varepsilon_{A,\tau} + \varepsilon_{V,\tau} - \eta(2 - \varepsilon_A - \varepsilon_V)(\varepsilon_A + \varepsilon_V) \\ &\quad + \eta\varepsilon_A\varepsilon_V(2 - \varepsilon_A - \varepsilon_V) \pm 2\eta \left(\frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right) \\ &=: (1 - \eta(2 - \varepsilon_A - \varepsilon_V))(\varepsilon_A + \varepsilon_V) + \text{Tmp}. \end{aligned}$$

Thus,

$$\begin{aligned} (\varepsilon_{A,\tau+1} + \varepsilon_{V,\tau+1})^2 &= ((1 - \eta(2 - \varepsilon_A - \varepsilon_V))(\varepsilon_A + \varepsilon_V) + \text{Tmp})^2 \\ &\leq \exp(-2\eta(2 - \varepsilon_A - \varepsilon_V))(\varepsilon_A + \varepsilon_V)^2 + 2|\varepsilon_A + \varepsilon_V||\text{Tmp}| + \text{Tmp}^2. \end{aligned}$$

Note that

$$|\text{Tmp}| \leq 4\eta \left(\varepsilon_A \varepsilon_V + \frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right).$$

Recall $|\varepsilon_A|, |\varepsilon_V| \leq 1/2$. Thus,

$$\begin{aligned} (\varepsilon_{A,\tau+1} + \varepsilon_{V,\tau+1})^2 &\leq \exp(-2\eta)(\varepsilon_A + \varepsilon_V)^2 \\ &\quad + 8\eta|\varepsilon_A + \varepsilon_V| \left(\varepsilon_A \varepsilon_V + \frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\mathbf{h}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right) \\ &\quad + 16\eta^2 \left(\varepsilon_A^2 \varepsilon_V^2 + \left(\frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\mathbf{h}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right)^2 \right) \\ &\leq \exp(-2\eta)(\varepsilon_A + \varepsilon_V)^2 + 8\eta|\varepsilon_A + \varepsilon_V|\varepsilon_A \varepsilon_V + 16\eta^2 \varepsilon_A^2 \varepsilon_V^2 \\ &\quad + 8\eta|\varepsilon_A + \varepsilon_V| \left(\frac{1}{K_Q T} + \delta_A^2 + \delta_V^2 + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\mathbf{h}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right). \end{aligned}$$

□

Lemma E.9. *Suppose that $(X_\tau)_\tau$ satisfies $X_{\tau+1} \leq e^{-A}X_\tau + B$ for some $A \in (0, 1]$, $B \geq 0$. If $X_0 \leq 2B/A$, then we have $X_\tau \leq 3B/A$ for all $\tau \geq 0$. If $X_0 \geq 2B/A$, then we have $X_\tau \leq 3B/A$ for all $\tau \geq \frac{6}{A} \log\left(\frac{X_0 A}{3B}\right)$.*

Proof. Since $e^{-A} \leq 1 - A/2$ for $A \in (0, 1]$, we have $X_{\tau+1} \leq X_\tau - AX_\tau/2 + B$. Hence, whenever $X_\tau \geq 2B/A$, we will have $X_{\tau+1} \leq X_\tau$. Moreover, if $X_\tau < 2B/A$, we have $X_{\tau+1} \leq 2B/A + B \leq 3B/A$. This proves the first part of the lemma.

Now, suppose that $X_0 \geq 2B/A$. When $X_\tau \geq 3B/A$, we have

$$X_{\tau+1} \leq X_\tau - AX_\tau/2 + B \leq X_\tau - AX_\tau/6 \leq e^{-A/6}X_\tau.$$

Thus, it takes at most $\frac{6}{A} \log\left(\frac{X_0 A}{3B}\right)$ steps to reduce X_τ from X_0 to $3B/A$. After that, the previous analysis applies. □

E.2.2 Main lemma of Stage 2

We split the analysis of Stage 2 into two substage. Let $c_\alpha \in (0, 0.05)$ be a small constant. Define

$$\mathcal{T}_{2.1} := \inf \{ \tau \geq \mathcal{T}_1 : (\alpha_{V,\tau} + \alpha_{A,\tau})/2 \geq 1 - c_\alpha \}.$$

We call $\{ \tau : \tau \leq \mathcal{T}_{2.1} \}$ stage 2.1 and $\{ \tau : \tau \geq \mathcal{T}_{2.1} \}$ stage 2.2. For notational simplicity, we define

$$\text{Err}_{\text{small}} = \max_\tau \left\{ \frac{1}{K_Q T} + \frac{\|\mathbf{h}_{V,\tau}\|_\mu}{\sqrt{K_P K_Q}} + \frac{\|\hat{\mathbf{h}}_{A,\tau}\|_\mu}{K_P \sqrt{K_Q}} \right\}.$$

Note that this can be made essentially arbitrarily small by choosing a large enough batch size.

Lemma E.10. *Suppose that the following hold at the beginning of Stage 2 (after thresholding and projection):*

- (a) $\alpha_V, \alpha_A \geq \alpha^{(2)}$
- (b) $\delta_A^2 + \delta_V^2 \leq (\delta^{(2)})^2$.

Let δ_* be our target value for δ_A and δ_V . Choose λ as in Lemma E.4. Suppose that

$$\begin{aligned} \max \left\{ \delta^{(2)}, \text{Err}_{\text{small}} \right\} \text{Err}_{\text{small}} &\leq \delta_*^2, \quad \text{Err}_{\text{small}} \leq O(\alpha^{(2)}), \\ \max \left\{ (\alpha_V, \tau_1 - \alpha_A, \tau_1)^2, \delta^{(2)}, \text{Err}_{\text{small}} \right\} &\leq O\left(\frac{1}{\log(1/\delta_*)}\right). \end{aligned}$$

Then, within $O(1/(\eta\alpha^{(2)}) + \log(1/\delta_*)/\eta)$ steps, we will have $\delta_A^2, \delta_V^2 \leq \delta_*^2$ and $\alpha_V, \alpha_A \in (0.9, 1.1)$.

Remark. Note that our conditions on α_V, α_A and $\delta^{(2)}$ are much weaker than what one can obtain from Stage 1. This allows us to apply the analysis here to transfer learning. ♣

The following proof should be treated as a large induction argument though we do not explicitly write down the induction as in the proof of Stage 1. In particular, we will show (by induction) that the approximation errors δ_A and δ_V are small, so that most of the naïve bounds on the entries of \tilde{A} and \tilde{V} can be transferred to A and V . In particular, $|V_{n,m}| = O(1)$ for all $n, m \in [N]$ so that our bounds in Section E.1 are valid, and $|a_t^{(k)}| = O(1)$ for all $k \in [N], t \in [N]$, which implies that after the projection step, $a_t^{(k)} = O(1/T)$ for all $t \notin q^{(k)}$.

Proof of the Lemma E.10

Common results for Stage 2.1 and 2.2 First, we prove some basic results that hold for both Stage 2.1 and 2.2. First, we show (by induction) that $\alpha_V, \alpha_A \geq \alpha^{(2)}$ and $\delta_V^2 + \delta_A^2 \leq \delta^{(2)}$ hold throughout Stage 2. Recall from Lemma E.5 that

$$\begin{aligned} \delta_{A,\tau+1}^2 &\leq \exp\left(-2\eta\alpha_V^2\right) \delta_{A,\tau}^2 + 3\eta \left(\delta_A + \eta \frac{\|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}} \right) \frac{\|\hat{\mathbf{h}}_A\|_\mu}{K_P \sqrt{K_Q}} \\ &\leq \exp\left(-2\eta(\alpha^{(2)})^2\right) \delta_{A,\tau}^2 + \eta \text{Err}_{\text{small}}. \end{aligned}$$

Hence, as long as $\text{Err}_{\text{small}} \leq O(1)\delta^{(2)}/(\alpha^{(2)})^2$, we can ensure $\delta_{A,\tau+1}^2 + \delta_{V,\tau+1}^2 \leq (\delta^{(2)})^2$ always hold.

Stage 2.1: signal growth By Lemma E.7, we have

$$1 - \frac{\alpha_{V,\tau+1} + \alpha_{A,\tau+1}}{2} \leq \left(1 - \eta \frac{\alpha_V + \alpha_A}{2}\right) \left(1 - \frac{\alpha_V + \alpha_A}{2}\right) + \eta \frac{\alpha_V + \alpha_A}{2} (\delta_A^2 + \delta_V^2) + \eta \text{Err}_{\text{small}}.$$

When $(\alpha_V + \alpha_A)/2 \leq 1 - c_\alpha$ and $\alpha_V, \alpha_A \geq \alpha^{(2)}$, we have

$$\eta \frac{\alpha_V + \alpha_A}{2} \left(1 - \frac{\alpha_V + \alpha_A}{2}\right) \geq \eta \frac{\alpha_V + \alpha_A}{2} c_\alpha \geq \eta \alpha^{(2)} c_\alpha.$$

Hence, as long as $\text{Err}_{\text{small}} \leq \alpha^{(2)} c_\alpha / 2$ and $\delta^{(2)} \leq c_\alpha / 2$, we have

$$1 - \frac{\alpha_{V,\tau+1} + \alpha_{A,\tau+1}}{2} \leq \left(1 - \frac{\eta \alpha_V + \alpha_A}{2}\right) \left(1 - \frac{\alpha_V + \alpha_A}{2}\right) \leq \exp\left(-\frac{\eta \alpha^{(2)}}{2}\right) \left(1 - \frac{\alpha_V + \alpha_A}{2}\right).$$

Thus, stage 2.1 takes at most $O(1/(\eta\alpha^{(2)}))$ steps.

Stage 2.1: difference between the α 's By Lemma E.6, we have

$$(\alpha_{V,\tau+1} - \alpha_{A,\tau+1})^2 \leq \exp(-2\eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A)^2 + 8\eta \left((\delta^{(2)})^2 + \text{Err}_{\text{small}} \right).$$

By the AM-GM inequality, we have $\alpha_A \alpha_V \leq ((\alpha_A + \alpha_V)/2)^2 \leq (1 - c_\alpha)^2$. Therefore, $1 - \alpha_A \alpha_V \geq c_\alpha$ and the above inequality can be further rewritten as

$$(\alpha_{V,\tau+1} - \alpha_{A,\tau+1})^2 \leq \exp(-2c_\alpha \eta) (\alpha_V - \alpha_A)^2 + 8\eta \left((\delta^{(2)})^2 + \text{Err}_{\text{small}} \right).$$

Thus, by (the proof of) Lemma E.9, we have

$$(\alpha_{V,\tau} - \alpha_{A,\tau})^2 \leq \max \left\{ 2(\alpha_{V,\tau_1} - \alpha_{A,\tau_2})^2, \frac{4 \left((\delta^{(2)})^2 + \text{Err}_{\text{small}} \right)}{c_\alpha} \right\}.$$

Stage 2.2: error decrease By Lemma E.5, we have

$$\begin{aligned} \delta_{A,\tau+1}^2 &\leq \exp(-2\eta\alpha_V^2) \delta_{A,\tau}^2 + 3\eta \max \left\{ \delta^{(2)}, \text{Err}_{\text{small}} \right\} \text{Err}_{\text{small}} \\ &\leq \exp(-\eta) \delta_{A,\tau}^2 + 3\eta \max \left\{ \delta^{(2)}, \text{Err}_{\text{small}} \right\} \text{Err}_{\text{small}}. \end{aligned}$$

Thus, by Lemma E.9, we have

$$\delta_{A,\tau}^2 \leq O(1) \max \left\{ \delta^{(2)}, \text{Err}_{\text{small}} \right\} \text{Err}_{\text{small}} \leq \delta_*^2, \quad \forall \tau \geq O \left(\frac{\log(1/\delta_*)}{\eta} \right).$$

In other words, Stage 2.2 takes at most $O \left(\frac{\log(1/\delta_*)}{\eta} \right)$ steps.

Stage 2.2: stability of $\alpha_V \pm \alpha_A$ Recall from Lemma E.8 and Lemma E.6 that

$$\begin{aligned} (\varepsilon_{A,\tau+1} + \varepsilon_{V,\tau+1})^2 &\leq \exp(-2\eta) (\varepsilon_A + \varepsilon_V)^2 + 8\eta |\varepsilon_A + \varepsilon_V| \varepsilon_A \varepsilon_V + 16\eta^2 \varepsilon_A^2 \varepsilon_V^2 \\ &\quad + 8\eta |\varepsilon_A + \varepsilon_V| \left(\delta^{(2)} + \text{Err}_{\text{small}} \right), \\ (\alpha_{V,\tau+1} - \alpha_{A,\tau+1})^2 &\leq \exp(-2\eta(1 - \alpha_A \alpha_V)) (\alpha_V - \alpha_A)^2 + 8\eta \left(\delta^{(2)} + \text{Err}_{\text{small}} \right). \end{aligned}$$

We wish to maintain the induction hypotheses $|\alpha_V - \alpha_A| = |\varepsilon_V - \varepsilon_A| = \theta_-$ for some $\theta_- = o(1)$ and $(\alpha_A + \alpha_V)/2 \leq 1 + c/\log(1/\delta_*) =: 1 + \theta_+$.

First, assume these conditions are true. Then, we have

$$(\alpha_{V,\tau+1} - \alpha_{A,\tau+1})^2 \leq \exp(O(1)\eta/\log(1/\delta_*)) (\alpha_V - \alpha_A)^2 + 8\eta \left(\delta^{(2)} + \text{Err}_{\text{small}} \right).$$

Since Stage 2.2 takes at most $O(\log(1/\delta_*)/\eta)$ steps, when the constant c is small, we have

$$\begin{aligned} (\alpha_{V,\tau} - \alpha_{A,\tau})^2 &\leq 2 \left((\alpha_{V,\tau_{2,1}} - \alpha_{A,\tau_{2,1}})^2 + 8 \left(\delta^{(2)} + \text{Err}_{\text{small}} \right) \right) \\ &\leq O(1) \max \left\{ (\alpha_{V,\tau_{2,1}} - \alpha_{A,\tau_{2,1}})^2, \delta^{(2)}, \text{Err}_{\text{small}} \right\} =: \theta_-. \end{aligned}$$

We can choose the parameters appropriately so that $\theta_- < o(\theta_+)$. Then, when $\varepsilon_A + \varepsilon_V \in (-2\theta_+, -\theta_+)$, we have, for some large universal constant $C > 0$,

$$\begin{aligned} (\varepsilon_{A,\tau+1} + \varepsilon_{V,\tau+1})^2 &\leq \exp(-2\eta) \theta_+^2 + C\eta \theta_+^3 + C\eta^2 \theta_+^4 + C\eta \theta_+ \left(\delta^{(2)} + \text{Err}_{\text{small}} \right) \\ &\leq \exp(-2\eta) \theta_+^2 + C\eta \theta_+^3 + C\eta \theta_+ \theta_- \\ &\leq \theta_+^2. \end{aligned}$$

This establishes the induction hypotheses on $\alpha_V \pm \alpha_A$.

F Stage 3: Final Convergence

As mentioned last section, we cannot ensure $\alpha_A \approx \alpha_V$ since $\alpha_A - \alpha_V$ is not contractive toward the end of training. However, we can add a final rounding step and then continue train the model to recover the ground-truth with ε -error.

First, we formally define our rounding procedure. Let $c > 0$ be a small constant (cf. the proof of Lemma F.1). For each $k \in [N]$, define

$$\hat{\mathbf{a}}^{(k)} := \left[a_t^{(k)} \mathbb{1} \left\{ a_t^{(k)} \geq c/Q \right\} \right]_{t \in [T]} \quad \text{and} \quad \mathbf{a}^{(*,k)} := \frac{\hat{\mathbf{a}}^{(k)}}{\mathbf{1}^\top \hat{\mathbf{a}}^{(k)}}.$$

This $\mathbf{A}^* := [\mathbf{a}^{(*,k)}]_{k \in [N]}$ is our rounded version of \mathbf{A} . For the error between \mathbf{A}^* and \mathbf{Q} , we have the following lemma.

Lemma F.1 (Rounding \mathbf{A}). *Let $\varepsilon \in (\Theta(Q^2/T^2), 0.1)$ be our target accuracy. Suppose that $\alpha_A = 1 - \varepsilon_A$ for some $|\varepsilon_A| \leq 0.1$ and $\|\Delta_A\|_\mu^2 \leq \delta_A^2$ for some $0 < \delta_A \ll 1/(Q\sqrt{N})$. Then, after rounding, we have*

$$\|\mathbf{A}^* - \mathbf{Q}\|_\mu^2 \leq O\left(\frac{Q^2}{T^2} + \delta_A^2 NQ\right).$$

In particular, to achieve ε accuracy in terms of $\|\cdot\|_\mu$, we only need $|\varepsilon_A| \leq 0.1$ and $\delta_A^2 \leq O(\varepsilon/\sqrt{NQ})$.

Remark. In particular, this lemma implies that as long as ε_A is not too large, after rounding, the error depends solely on $\|\Delta_A\|_\mu^2$. \clubsuit

Proof. For notational simplicity, we omit the superscript k for now. Write

$$\mathbf{a} = \tilde{\mathbf{a}} + \Delta_{\mathbf{a}} = \alpha_V \mathbf{q} + (1 - \alpha_V) \frac{\mathbf{1}}{T} + \Delta_{\mathbf{a}} = (1 - \varepsilon_A) \mathbf{q} + \varepsilon_A \frac{\mathbf{1}}{T} + \Delta_{\mathbf{a}}.$$

Note that $\|\Delta_{\mathbf{a}}\|_\infty \leq \|\Delta_{\mathbf{a}}\|_2 \leq O(1)\sqrt{N} \|\Delta_{\mathbf{a}}\|_\mu \leq O(\sqrt{N}\delta_A)$. Since all nonzero q_s are lower bounded by $\Omega(1/Q)$, we have, for any $s \in \mathbf{q}$ and $t \notin \mathbf{q}$,

$$a_s \geq \frac{\Omega(1)}{Q} - \frac{|\varepsilon_A|}{T} - O(\sqrt{N}\delta_A) = \frac{\Omega(1)}{Q},$$

$$|a_t| \leq \frac{|\varepsilon_A|}{T} + O(\sqrt{N}\delta_A) \ll \frac{1}{Q}.$$

Hence, we can choose a small constant $c > 0$, so that

$$\hat{\mathbf{a}} := [a_t \mathbb{1} \{ |a_t| \geq c/Q \}]_{t \in [T]} = (1 - \varepsilon_A) \mathbf{q} + \varepsilon_A \frac{\mathbf{1}_q}{T} + [\Delta_{\mathbf{a}}]_q,$$

where for $\mathbf{v} \in \mathbb{R}^T$, $\mathbf{v}_q \in \mathbb{R}^T$ is defined as $[v_t \mathbb{1} \{ t \in \mathbf{q} \}]_{t \in [T]}$ here. Now, consider the difference between \mathbf{q} and $\hat{\mathbf{a}}/\mathbf{1}^\top \hat{\mathbf{a}}$. We have

$$\mathbf{1}^\top \hat{\mathbf{a}} = 1 - \varepsilon_A + \frac{\varepsilon_A Q}{T} \pm O(Q\sqrt{N}\delta_A),$$

and therefore,

$$\begin{aligned} \mathbf{a}_* &:= \frac{\hat{\mathbf{a}}}{\mathbf{1}^\top \hat{\mathbf{a}}} = \frac{(1 - \varepsilon_A) \mathbf{q} + \varepsilon_A \mathbf{1}_q/T + [\Delta_{\mathbf{a}}]_q}{(1 - \varepsilon_A) + \varepsilon_A Q/T \pm O(Q\sqrt{N}\delta_A)} \\ &= \frac{(1 - \varepsilon_A) \mathbf{q} + \varepsilon_A \mathbf{1}_q/T + [\Delta_{\mathbf{a}}]_q}{(1 - \varepsilon_A)} \left(1 \pm O(\varepsilon_A Q/T + Q\sqrt{N}\delta_A) \right) \\ &= (\mathbf{q} \pm O(\varepsilon_A \mathbf{1}_q/T + [\Delta_{\mathbf{a}}]_q)) \left(1 \pm O(\varepsilon_A Q/T + Q\sqrt{N}\delta_A) \right) \\ &= \mathbf{q} \pm O_2(\varepsilon_A Q/T + \sqrt{QN}\delta_A). \end{aligned}$$

Thus,

$$\|\mathbf{A}^* - \mathbf{Q}\|_\mu^2 = \sum_{k=1}^N \mu_k \left\| \mathbf{a}_*^{(k)} - \mathbf{q}^{(k)} \right\|^2 = O\left(\frac{\varepsilon_A^2 Q^2}{T^2} + \delta_A^2 NQ\right).$$

□

Lemma F.2. Let $\varepsilon \in (\Theta(1)/(K_Q T), 0.1)$ be our target accuracy. Suppose that $\|A - Q\|_\mu \leq \delta_{A,*} \leq 0.01\varepsilon$ and $\|\Delta_V\|_\mu \leq \delta_V \leq 0.01$ at the beginning of Stage 3, and $\|h_V\|_\mu \leq c\varepsilon\sqrt{K_P K_Q}$ for all τ and a sufficiently small constant c . Then, we have $\|V - P\|_\mu^2 \leq \varepsilon$ for all $\tau \geq \mathcal{T}_2 + \Theta(\log(1/\varepsilon)/\eta)$.

Proof. Under the condition $\|A - Q\|_\mu \leq \delta_{A,*}$, we have $\|\Delta_A\|_\mu \leq \|A - Q\|_\mu \leq \delta_{A,*}$ and

$$\alpha_A = \frac{1}{K_Q} \left(\langle A, Q \rangle_\mu - \frac{1}{T} \right) = \frac{1}{K_Q} \left(K_Q + \langle \Delta_A, Q \rangle_\mu \right) = 1 \pm O\left(\sqrt{Q}\delta_{A,*}\right).$$

Recall that we only train V in Stage 3. Note that by Lemma C.3,

$$\|V - P\|_\mu^2 = \|\tilde{V} - P\|_\mu^2 + \|\Delta_V\|_\mu^2 = (1 - \alpha_V)^2 K_P + \|\Delta_V\|_\mu^2 \leq (1 - \alpha_V)^2 + \|\Delta_V\|_\mu^2.$$

Hence, to get ε accuracy, it suffices to have $(1 - \alpha_V)^2 \leq \varepsilon/2$ and $\|\Delta_V\|_\mu^2 \leq \varepsilon/2$.

First, for α_V , by Lemma C.5 and $\eta_V = \eta/K_Q$, we have

$$\begin{aligned} \alpha_{V,\tau+1} &= \alpha_{V,\tau} + \eta_V K_Q (1 - \alpha_V \alpha_A) \alpha_V + \eta_V \frac{1 - \alpha_V}{T} - \eta_V \alpha_V \|\Delta_A\|_\mu^2 - \frac{\eta_V}{K_P} \langle h_{V,\tau}, P \rangle_\mu \\ &= \alpha_{V,\tau} + \eta (1 - \alpha_V) \alpha_V \pm \eta O\left(\frac{1}{K_Q T} + \delta_{A,*} \alpha_V\right) \pm \eta O\left(\frac{\langle h_{V,\tau}, P \rangle_\mu}{K_P K_Q}\right). \end{aligned}$$

Hence,

$$\begin{aligned} (1 - \alpha_{V,\tau+1})^2 &= \left((1 - \eta \alpha_V) (1 - \alpha_{V,\tau}) \pm \eta O\left(\frac{1}{K_Q T} + \delta_{A,*} \alpha_V\right) \pm \eta O\left(\frac{\langle h_{V,\tau}, P \rangle_\mu}{K_P K_Q}\right) \right)^2 \\ &\leq \exp(-2\eta \alpha_V) (1 - \alpha_{V,\tau})^2 + \eta O\left(\frac{1}{K_Q T} + \delta_{A,*} \alpha_V\right) + \eta O\left(\frac{\langle h_{V,\tau}, P \rangle_\mu}{K_P K_Q}\right) \\ &\leq \exp(-\eta) (1 - \alpha_{V,\tau})^2 + 0.1\eta\varepsilon. \end{aligned}$$

For $\|\Delta_A\|_\mu^2$, by (the proof of) Lemma E.5, we have

$$\begin{aligned} \|\Delta_{V,\tau+1}\|_\mu^2 &\leq \exp(-2\eta \alpha_A^2) \|\Delta_{V,\tau}\|_\mu^2 + 3\eta \left(\|\Delta_{V,\tau}\|_\mu + \eta \frac{\|h_V\|_\mu}{\sqrt{K_P K_Q}} \right) \frac{\|h_V\|_\mu}{\sqrt{K_P K_Q}} \\ &\leq \exp(-\eta) \|\Delta_{V,\tau}\|_\mu^2 + 0.1\eta\varepsilon. \end{aligned}$$

Thus, by Lemma E.9, we have $(1 - \alpha_V)^2 \leq \varepsilon/2$ and $\|\Delta_V\|_\mu^2 \leq \varepsilon/2$ for all $\tau \geq \mathcal{T}_2 + \Theta(\log(1/\varepsilon)/\eta)$. \square

Corollary F.3. Let $\varepsilon \in (\Theta(1)/(K_Q T), 0.1)$ be our target accuracy. Suppose that $\alpha_A \in (0.9, 1.1)$, $\|\Delta_A\|_\mu^2 \leq O(\varepsilon/(Q\sqrt{N}))$, and $\|\Delta_V\|_\mu^2 \leq 0.01$. Then with $\text{poly}(N, Q, 1/\varepsilon)$ samples, we have with high probability that $\|A - Q\|_\mu^2 \leq \varepsilon$ and $\|V - P\|_\mu^2 \leq \varepsilon$ after Stage 3, which takes $O(\log(1/\varepsilon)/\eta)$ steps.

Proof. It suffices to combine the previous two lemmas, the concentration results in Section B, and apply union bound. \square

G Proof of the main theorem

In this section, we combine the results from the last three sections and prove the following formal version of Theorem 3.1.

Theorem G.1. Let $\varepsilon > 0$ be our target accuracy and $\mathcal{T}_1 = \min\{\tau \geq 0 : \max\{\alpha_{V,\tau}, \alpha_{A,\tau}\} \geq \Theta(1/(QN))\}$. We can choose the hyperparameters in Algorithm 1 such that within $O(\log(T)/\eta_1 + 1/(\eta\alpha^{(2)}) + \log(1/\varepsilon)/\eta)$ steps, we have $\|A - Q\|_\mu^2 \leq \varepsilon$ and $\|V - P\|_\mu^2 \leq \varepsilon$ with probability at least $1 - \delta$ and the number of samples used before and after \mathcal{T}_1 are $\text{poly}(T, \delta)$ and $\text{poly}(N, Q, 1/\varepsilon, \log T, \delta)$, respectively.

Proof. The results for Stage 1 follow directly from Lemma D.1. Now, consider the results for Stage 2 and 3. First, by Corollary F.3, it suffices to make sure at time \mathcal{T}_2 , we have $\alpha_A \in (0.9, 1.1)$, $\|\Delta_A\|_\mu^2 \leq O(\varepsilon/(Q\sqrt{N}))$ and $\|\Delta_V\|_\mu^2 \leq 0.01$ (with high probability). By Lemma D.1, we know the following hold at time \mathcal{T}_1 w.h.p:

$$\alpha_A, \alpha_V = \Theta(1/(QN)) \quad \text{and} \quad \|\Delta_A\|_\mu^2, \|\Delta_V\|_\mu^2 \leq O(1/T).$$

Therefore, by Lemma E.1, if we choose the threshold to be $\Theta(1/(Q^2N))$, then after thresholding and projection, we have $a_t^{(k)} = O(1/T)$ for all $t \notin \mathbf{q}^{(k)}$. Thus, by Lemma E.3 and Lemma E.4, with $\text{poly}(Q, N, 1/\varepsilon, \log(T))$ samples, we have with high probability that $a_t^{(k)} = O(1/T)$ for all $t \notin \mathbf{q}^{(k)}$ holds and $\|\hat{\mathbf{h}}_A\|_\mu$ satisfies the requirements in Lemma E.10 throughout Stage 2. Thus, by Lemma E.10 with $\delta_*^2 = O(\varepsilon/Q\sqrt{N})$, at the end of Stage 2, we have with high probability that $\alpha_A \in (0.9, 1.1)$, $\|\Delta_A\|_\mu^2 \leq O(\varepsilon/(Q\sqrt{N}))$ and $\|\Delta_V\|_\mu^2 \leq 0.01$. When combined with Corollary F.3, this completes the proof. \square

H Transfer learning

Lemma H.1 (Initialization). *Suppose that we have learned $\hat{\mathbf{P}}$ and $\langle \hat{\mathbf{P}}, \mathbf{P} \rangle_\mu \geq 2\|\boldsymbol{\mu}\|^2$, $\|\hat{\mathbf{P}}\|_\mu^2 = \Theta(1)\|\mathbf{P}\|_\mu^2$. Let $\mathbf{V} = \theta\hat{\mathbf{P}} + (1-\theta)\boldsymbol{\mu}\mathbf{1}^\top$ for some $\theta \in (0, 1)$. We have*

$$\alpha_V \in \left[\Theta\left(\frac{\theta}{NK_P}\right), \Theta(\theta) \right] \quad \text{and} \quad \|\Delta_V\|_\mu^2 \leq \Theta(1)\theta^2 K_P.$$

Proof. First, consider Δ_V . Since $\|\Delta_V\|_\mu$ is the distance to a projection, we have

$$\|\Delta_V\|_\mu^2 \leq \|\mathbf{V} - (\theta\mathbf{P} + (1-\theta)\boldsymbol{\mu}\mathbf{1}^\top)\|_\mu^2 = \theta^2 \|\hat{\mathbf{P}} - \mathbf{P}\|_\mu^2 \leq \Theta(1)\theta^2 \|\mathbf{P}\|_\mu^2 = \Theta(1)\theta^2 K_P.$$

For α_V , we compute

$$\langle \mathbf{V}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2 = \theta \left(\langle \hat{\mathbf{P}}, \mathbf{P} \rangle_\mu - \|\boldsymbol{\mu}\|^2 \right) \geq \theta \|\boldsymbol{\mu}\|^2 \geq \Theta(1)\theta/N.$$

In particular, this implies $\alpha_V \geq \Theta\left(\frac{\theta}{NK_P}\right)$. \square

Lemma H.2 (First gradient step). *Let $\mathbf{A} = \mathbf{1}\mathbf{1}^\top/T$ and \mathbf{V} given by Lemma H.1. Let $\varepsilon_{\text{tmp}} \leq O(1/(NK_P))$. Run one gradient step with $\text{poly}(N, Q, \log T, 1/\varepsilon_{\text{tmp}})$ samples and $\eta = 1$, remove all entries of $\mathbf{a}^{(k)}$ with $|a_t^{(k)}| \leq \Theta(\theta/(NK_P))$ and the replace $\mathbf{a}^{(k)}$ with the projection $(\mathbf{I} - \mathbf{1}\mathbf{1}^\top/T)\mathbf{a}^{(k)}$.*

With high probability, we have $\alpha_A = (1 \pm O(\varepsilon_{\text{tmp}}))\alpha_V$ and $\|\Delta_A\|_\mu^2 \leq O\left(\frac{\varepsilon_{\text{tmp}}^2 \alpha_V^2}{Q} + \frac{1}{T}\right)$.

Proof. The proof idea is essentially the same as Lemma E.3, though the reinitialization of the first layer allows better estimations in several places. Recall that for each $s \in [T]$, we have

$$\begin{aligned} \partial_{a_s^{(k)}} l &= \mathbb{1}\{x_{T+1} = k\} (\mathbf{V}\mathbf{e}_{x_s})^\top (\mathbf{V}\mathbf{X}\mathbf{1}/T - \mathbf{e}_{x_o}), \\ \mathbb{E} \partial_{a_s^{(k)}} l &= \mu_k K_V/T - \mu_k K_{VP} q_s. \end{aligned}$$

Also recall from Lemma C.3 that $K_{VP} = \alpha_V K_P$ and $K_V = \alpha_V^2 K_P + \|\Delta_V\|_\mu^2$. For any $s \in \mathbf{q}^{(k)}$, we have

$$-\mathbb{E} \partial_{a_s^{(k)}} l = \mu_k \left(\alpha_V K_P q_s - \frac{\alpha_V^2 K_P + \|\Delta_V\|_\mu^2}{T} \right) \geq \frac{\mu_k \alpha_V K_P q_s}{2}$$

where the inequality comes from $\mu_k \alpha_V K_P q_s \geq \Omega(\alpha_V/N^2/Q) \gg 1/T$. Meanwhile, for any $s \notin [T]$, we have $\mathbb{E} \partial_{a_s^{(k)}} l = O(1/T)$.

Meanwhile, for any $s \in [T]$ we have

$$\left| \partial_{a_s^{(k)}} l \right| \leq |(\mathbf{V} \mathbf{e}_{x_s})^\top \mathbf{V} \mathbf{X} \mathbf{1} / T| + |(\mathbf{V} \mathbf{e}_{x_s})^\top \mathbf{e}_{x_o}| \leq \frac{1}{T} \sum_{t=1}^T |\langle \mathbf{V}_{:,x_s}, \mathbf{V}_{:,x_t} \rangle| + |V_{x_o, x_s}| \leq O(N).$$

Thus, by some standard concentration argument similar to the one in Lemma E.3, we can show that with $\text{poly}(N, Q, 1/\alpha_V, 1/\varepsilon_{\text{Tmp}}, \log T)$ samples, we can make sure with high probability,

$$\begin{aligned} \partial_{a_t^{(k)}}^{(B_\tau)} l &= (1 \pm \varepsilon_{\text{Tmp}}) \mathbb{E} \partial_{a_t^{(k)}} l = (1 \pm \varepsilon_{\text{Tmp}}) \mu_k \alpha_V K_P q_s & \forall t \in \mathbf{q}^{(k)}, \\ \partial_{a_t^{(k)}}^{(B_\tau)} l &= O\left(\frac{1}{T} + \varepsilon_{\text{Tmp}} \frac{\alpha_V K_P}{NQ}\right) = O\left(\varepsilon_{\text{Tmp}} \frac{\alpha_V K_P}{NQ}\right) & \forall t \notin \mathbf{q}^{(k)}. \end{aligned}$$

Thus, after one gradient step with $\eta = 1$, we have

$$\begin{aligned} a_t^{(k)} &= (1 \pm \varepsilon_{\text{Tmp}}) \alpha_V q_t^{(k)} & \forall t \in \mathbf{q}^{(k)}, \\ a_t^{(k)} &= O\left(\varepsilon_{\text{Tmp}} \frac{\alpha_V}{Q}\right) & \forall t \notin \mathbf{q}^{(k)}. \end{aligned}$$

Recall from Lemma H.1 that $\alpha_V \in [\Theta(\theta/(NK_P)), \Theta(\theta)]$. Hence, we can choose the threshold λ_0 to be $\Theta(\theta/(NK_P))$ and $\varepsilon_{\text{Tmp}} \leq 1/(NK_P)$ so that

$$\hat{\mathbf{a}}^{(k)} = \left[a_t^{(k)} \mathbb{1}\{a_t^{(k)} \geq \lambda_0\} \right]_{t \in [T]} = \left[(1 \pm \varepsilon_{\text{Tmp}}) \alpha_V q_t^{(k)} \mathbb{1}\{t \in \mathbf{q}^{(k)}\} \right]_{t \in [T]}.$$

Now, set

$$\mathbf{a}^{(k)} \leftarrow (\mathbf{I} - \mathbf{1}\mathbf{1}^\top / T) \hat{\mathbf{a}}^{(k)} = (1 \pm O_\infty(\varepsilon_{\text{Tmp}})) \alpha_V \mathbf{q}^{(k)} + O_\infty(1) \frac{\mathbf{1}}{T}.$$

Note that this implies that after the first step, we have

$$\alpha_A = \frac{\sum_k \mu_k \langle \mathbf{a}^{(k)}, \mathbf{q}^{(k)} \rangle - 1/T}{K_Q} = \frac{\alpha_V (1 \pm O(\varepsilon_{\text{Tmp}})) \|\mathbf{Q}\|_\mu^2 + O(Q/T)}{K_Q} = (1 \pm O(\varepsilon_{\text{Tmp}})) \alpha_V,$$

and

$$\|\Delta_A\|_\mu^2 \leq \sum_{k=1}^N \mu_k \left\| \mathbf{a}^{(k)} - \left(\alpha_V \mathbf{q}^{(k)} + (1 - \alpha_V) \mathbf{1} / T \right) \right\|_\mu^2 \leq O\left(\frac{\varepsilon_{\text{Tmp}}^2 \alpha_V^2}{Q} + \frac{1}{T}\right).$$

□

Theorem H.3 (Main theorem for transfer learning). *Let $\varepsilon > 0$ be our target accuracy. Consider the same setting of Lemma H.1 and Lemma H.2. Choose $\theta = \sqrt{O(1/\log(1/\varepsilon))}$ and $\varepsilon_{\text{Tmp}} = O\left(\frac{1}{\log(1/\varepsilon)}\right)$. Then, after one step of update on \mathbf{A} as in Lemma H.2, \mathbf{A} and \mathbf{V} satisfies the conditions of Lemma E.10, and therefore we can learn (\mathbf{P}, \mathbf{Q}) to ε -accuracy using $\text{poly}(N, Q, 1/\varepsilon, 1/\delta)$ samples with probability at least $1 - \delta$ within $\text{poly}(N, Q, 1/\varepsilon, 1/\delta)$ steps.*

I Additional Experiment

For all our experiments, we use Numpy and run on a normal laptop which takes about 20 minutes.

Setup. In all our experiments, we choose $T = 5000, Q = 2, N = 3$. The architecture is

$$\mathbf{F}(\mathbf{x}, x_{T+1}; \mathbf{V}, \mathbf{A}) := \mathbf{V} \mathbf{X} (\mathbf{I}_T \mathbf{A} \mathbf{e}_{x_{T+1}}) =: \mathbf{V} \mathbf{X} \mathbf{a}^{(x_{T+1})}, \quad (15)$$

and the data model is the SCB (1) data-generating model. The batch size is $B = 64$ and the regularization hyperparameter is $\lambda = 1e-5$. The total time is $\mathcal{T} = 1000$ iterations where stage 1 takes $\tau \in [0, 400]$ with learning rate $\eta_1 = 0.01$. After $\tau > 400$, we use $\eta_2 = 0.005$ for further improvement (stages 2 and 3).

Hyperparameter selection. Due to the limitation of computational resources, we do experiments with $N \in [3, 20]$ for real-world batched gradient experiments, and $N \in \{100, 500\}, T = 100000$ experiments by using Gaussian noise SGD simulations based on the dynamics of Lemma C.5 and C.6. As T needs to scale with N polynomially, it would be beyond our computation capability to

experiment with larger N . As for other hyperparameters, λ is chosen based on our theoretical results (Theorem 3.1 and G.1): $\lambda \sim \Theta(\epsilon K_P / Q^2 N^3)$ in Lemma E.4. The batch size can be chosen from standard $\{64, 128, 256\}$, while smaller batch size will lead to divergence for both SGD and regularized GD. η is chosen as the largest learning rate without divergence.

Besides the original parameters, we consider the approximation error after the normalization step (stage 3) in real-time. That is thresholding and normalizing the attention block

$$\forall k \in [n], \hat{\mathbf{a}}^{(k)} = [a_t^{(k)} \mathbb{1}\{a_t^{(k)} \geq \Omega(1/Q)\}]_t. \mathbf{a}^{(k)} \leftarrow \hat{\mathbf{a}}^{(k)} / \mathbf{1}^\top \hat{\mathbf{a}}^{(k)}$$

and we do further gradient descent to recover V . In this case, we will directly use the linear regression solution on population loss for V .

Here we report in addition: (1) original signal projection on the population process trajectory α_A, α_V and the distance to the trajectory Δ_A, Δ_V . (2) The approximation error/similarity before and after normalization for both SGD and proximal gradient descent. We conclude that in all metrics proximal gradient descent performs better than the vanilla gradient descent with a small batch size (when the noise is large).




Figure 3: **Signals α_A, α_V and the distance to population process Δ_A, Δ_V .** For the SGD, the distance to the population process of the attention matrix A keeps growing and dominates the signal term. That explains the failure to learn the correct attention pattern, which leads to saturation of the signal. In comparison, our proximal methods dramatically help reduce the gradient noise and keep close to the population process. Though $\|\Delta_A\|$ eventually grows up due to the bias of the gradient estimate (the original signal growth is also slowed down), after normalization it can still approximately learn the correct pattern. Both $\|\Delta_V\|$ stay small empirically.

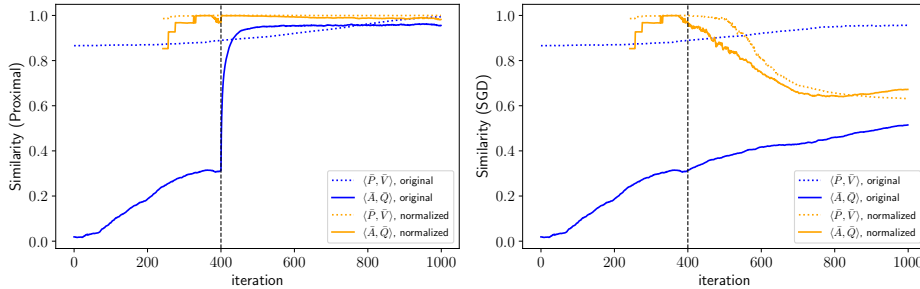


Figure 4: **Similarity with the ground-truth.** The figure shows after Stage 1, normalization helps further improve the solution of the proximal method. Meanwhile, with or without normalization, our proximal method always outperforms the vanilla SGD, which fails to recover the ground-truth.

We tried different orders of state number N and show that ℓ_1 regularization is necessary and outperforms SGD when batch size is small (gradient noise is large). Due to computation limitation, we experiment with real batched gradient on $N \leq 20, T \leq 5000$, and do SGD simulation by combining our population gradient + Gaussian noise (to mimic the batch gradient noise) for $N \leq 500, T \leq 100000$.

We also corroborate our previous experiments with the new test loss plot to show the convergence of training. Note that since there are multiple global minima for the linear attention, ℓ_1 regularization dynamics eventually will make A and Q deviate from the ground-truth while representing the same function. That is why the loss converges but the distance to the ground-truth increases after some

point, making the final normalization step essential to recover the ground-truth. Another point is that according to our theory, the regularization will eventually distort the learned pattern when trained for too many iterations. Empirically, the loss also increases a little after it converges. Therefore, we must stop early and normalize before the distortion happens.

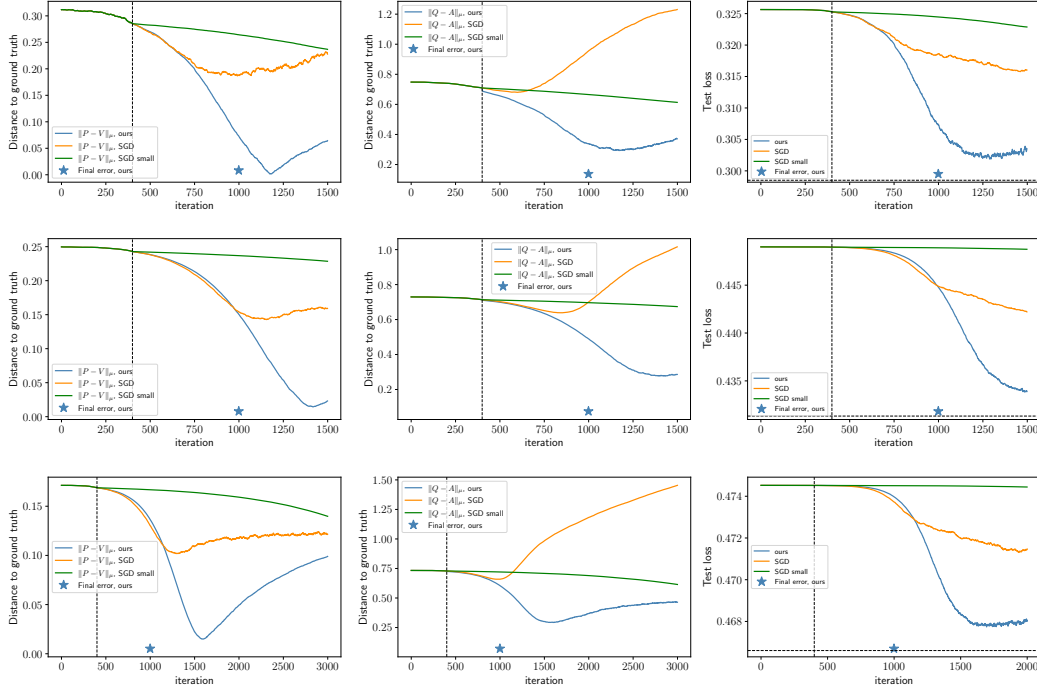


Figure 5: **Convergence analysis.** We plot the distance to the ground-truth and the test loss for $N = 3, 10, 20$ (from top to bottom). It shows that when gradient noise is large, ℓ_1 regularized algorithm with normalization and early-stopping can almost perfectly recover the ground-truth (the star), while SGD struggles to learn the target function. Figure 3 (Appendix I) also shows when the gradient noise is large, SGD never learns ground-truth Q even with normalization.

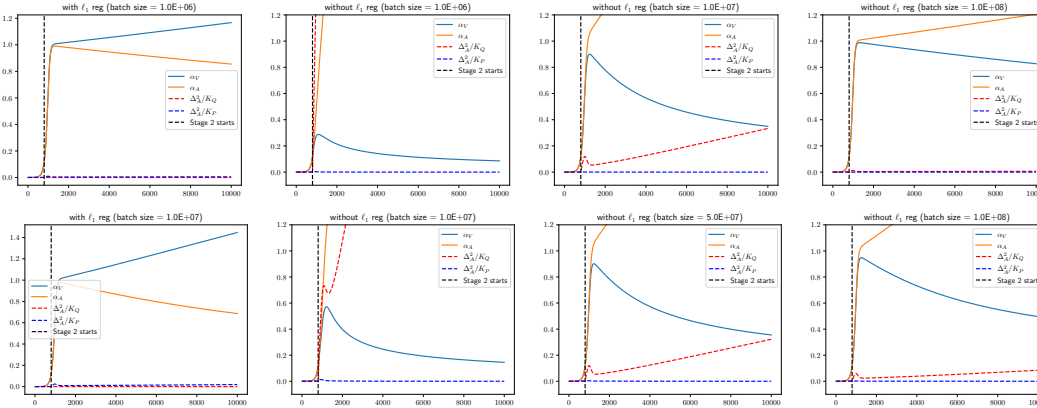


Figure 6: **Simulation with larger N and T .** We simulate the SGD/ ℓ_1 regularized dynamics by replacing the batched noise with Gaussian noise in the dynamics formula in Lemma C.5 and C.6. The gaussian noise variance scales with the inverse of batch size. The experiments show that the conclusions drawn from the small N cases still hold in those simulations: when $T = 100000$, $N = 100/500$, our ℓ_1 regularized algorithm can recover the ground-truth since the distance to the population trajectory (Δ_A, Δ_V) stays very small, while the error along SGD trajectories quickly increases with the same batch size.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We accurately summarized our claims and contributions in the abstract as well as the introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: It is included in Appendix A.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The assumptions and conditions are included in the Setup section 2.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All information is included in Section 5 and Appendix I.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will upload the codes in the supplementary materials.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All information is included in Section 5 and Appendix I.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Our experiments are for dynamics simulations. We do not include performance/accuracy results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: All information is included in Section 5 and Appendix I.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: This work is a pure theory paper without potential harmful effect.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This work is a theoretical paper without direct societal impact.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper is a theoretical work without such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.