
On Mesa-Optimization in Autoregressively Trained Transformers: Emergence and Capability

Chenyu Zheng^{1,2}, Wei Huang³, Rongzhen Wang^{1,2}, Guoqiang Wu⁴,
Jun Zhu⁵, Chongxuan Li^{1,2*}

¹ Gaoling School of Artificial Intelligence, Renmin University of China

² Beijing Key Laboratory of Big Data Management and Analysis Methods

³ RIKEN AIP ⁴ School of Software, Shandong University

⁵ Dept. of Comp. Sci. & Tech., BNRist Center, THU-Bosch ML Center, Tsinghua University
{cyzheng, wangrz, chongxuanli}@ruc.edu.cn; wei.huang.vr@riken.jp;
guoqiangwu@sdu.edu.cn; dcszj@mail.tsinghua.edu.cn

Abstract

Autoregressively trained transformers have brought a profound revolution to the world, especially with their in-context learning (ICL) ability to address downstream tasks. Recently, several studies suggest that transformers learn a mesa-optimizer during autoregressive (AR) pretraining to implement ICL. Namely, the forward pass of the trained transformer is equivalent to optimizing an inner objective function in-context. However, whether the practical non-convex training dynamics will converge to the ideal mesa-optimizer is still unclear. Towards filling this gap, we investigate the non-convex dynamics of a one-layer linear causal self-attention model autoregressively trained by gradient flow, where the sequences are generated by an AR process $x_{t+1} = \mathbf{W}x_t$. First, under a certain condition of data distribution, we prove that *an autoregressively trained transformer learns \mathbf{W} by implementing one step of gradient descent to minimize an ordinary least squares (OLS) problem in-context. It then applies the learned $\widehat{\mathbf{W}}$ for next-token prediction*, thereby verifying the mesa-optimization hypothesis. Next, under the same data conditions, we explore the capability limitations of the obtained mesa-optimizer. We show that a stronger assumption related to the moments of data is the sufficient and necessary condition that the learned mesa-optimizer recovers the distribution. Besides, we conduct exploratory analyses beyond the first data condition and prove that generally, the trained transformer will not perform vanilla gradient descent for the OLS problem. Finally, our simulation results verify the theoretical results, and the code is available at <https://github.com/ML-GSAI/MesaOpt-AR-Transformer>.

1 Introduction

Foundation models based on transformers [1] have revolutionized the AI community in lots of fields, such as language modeling [2; 3; 4; 5; 6], computer vision [7; 8; 9; 10] and multi-modal learning [11; 12; 13; 14; 15]. The crux behind these large models is a very simple yet profound strategy named *autoregressive (AR) pretraining*, which encourages transformers to predict the next token when a context is given. In terms of the trained transformers, one of their most intriguing properties is the *in-context learning (ICL)* ability [2], which allows them to adapt their computation and perform downstream tasks based on the information (e.g. examples) provided in the context without any updates to their parameters. However, the reason underlying the emergence of ICL ability is still poorly understood.

*Correspondence to Chongxuan Li.

Recently, we are aware that some preliminary studies [16; 17] have attempted to understand the ICL ability from the AR training and connected its mechanisms to a popular hypothesis named *mesa-optimization* [18], which suggests that transformers learn some algorithms during the AR pretraining. In other words, the inference process of the trained transformers is equivalent to optimizing some inner objective functions on the in-context data.

Concretely, the seminal work [16] constructs a theoretical example where a single linear causally-masked self-attention layer with manually set parameters can predict the next token using one-step gradient-descent learning for an ordinary least squares (OLS) problem over the historical context. Moreover, they conduct numerous empirical studies to establish a close connection between autoregressively trained transformers and gradient-based mesa-optimization algorithms. Built upon the setting of [16], recent work [17] precisely characterizes the pretraining loss landscape of the one-layer linear transformer trained on a simple first-order AR process with a fixed full-one initial token. As a result, they find that the optimally-trained transformer recovers the theoretical construction in [16]. However, their results rely on imposing the diagonal structure on the parameter matrices of the transformer and do not discuss whether the practical non-convex dynamics can converge to the ideal global minima. Besides, it is still unclear about the impact of data distribution on the trained transformer, which has been proven to be important in practice [19; 20; 21; 22] and theory [23].

In this paper, we take a further step toward understanding the mesa-optimization in autoregressively trained transformers. Specially, without an explicit diagonal structure assumption, we analyze the non-convex dynamics of a one-layer linear transformer trained by gradient flow on a controllable first-order AR process, and try to answer the following questions rigorously:

1. *When do mesa-optimization algorithms emerge in autoregressively trained transformers?*
2. *What is the capability limitation of the mesa-optimizer if it does emerge?*

Our first main contribution is to characterize a sufficient condition (Assumption 4.1) on the data distribution for a mesa-optimizer to emerge in the autoregressively trained transformer, in Section 4.1. We note that any initial token x_1 whose coordinates x_{1i} are i.i.d. random variables with zero mean and finite moments satisfy this condition, including normal distribution $\mathcal{N}(\mathbf{0}_d, \sigma^2 \mathbf{I}_d)$. Under this assumption, the non-convex dynamics will exactly converge to the theoretical construction in [16] without any explicit structural assumption in [17], resulting in the trained transformer implementing one step of gradient descent for the minimization of an OLS problem in-context.

Our second main contribution is to characterize the capability limitation of the obtained mesa-optimizer under Assumption 4.1 in Section 4.2. We characterize a stronger assumption (Assumption 4.2) related to the moment of data distribution as the necessary and sufficient condition that the mesa-optimizer recovers the true distribution (a.k.a. predict the next token correctly). Unfortunately, we find that the mesa-optimizer can not recover the data distribution when the initial token is sampled from the standard normal distribution, which suggests that ICL by AR pretraining [16; 17] is different from ICL by few-shot learning pretraining [24; 25; 26; 27; 28; 29; 30; 23] (see details in Section 2), a setup attracting attention from many theorists recently. We think ICL in the setting of AR pretraining needs more attention from the theoretical community.

In Section 4.3, we further study the convergence of the training dynamics when Assumption 4.1 does not hold anymore by adopting the setting in [17]. In this case, as a complement of [17], we prove that under a similar but weaker structural assumption, training dynamics will converge to the theoretical construction in [16] and the trained transformer implements exact gradient-based mesa-optimization. However, we prove that without any structural assumption, the trained transformer will not perform vanilla gradient descent for the OLS problem in general. Finally, we conduct simulations to validate our theoretical findings in Section 6.

2 Additional related work

2.1 Mesa-optimization in ICL for few-shot linear regression

In addition to AR pretraining, much more empirical [31; 25; 32; 33] and theoretical studies [24; 26; 27; 28; 29; 34; 30] have given evidence to the mesa-optimization hypothesis when transformers are trained to solve few-shot linear regression problem with the labeled training instances in the context. On the experimental side, for example, the seminal empirical works by [31; 35] considers

ICL for linear regression, where they find the ICL performance of trained transformers is close to the OLS. On the theoretical side, considering a one-layer linear transformer, [26; 24] prove that the global minima of the population training loss is equivalent to one step of preconditioned gradient descent. Notably, [24] further proves that the training dynamics do converge to the global minima, and the obtained mesa-optimizer solves the linear problem. For multi-layer attention models, recent works suggest that they can perform efficient high-order optimization algorithms such as Newton’s method [36; 37; 34]. Unfortunately, the pretraining goal of these studies is different from the AR training. Therefore, it is still unclear whether these findings can be transferred to transformers autoregressively trained on sequential data.

2.2 Other explanations for ICL

In addition to the mesa-optimization hypothesis, there are other explanations for the emergence of ICL. [38; 39; 40; 41] explain ICL as inference over an implicitly learned topic model. [42] connects ICL to multi-task learning and establishes generalization bounds using the algorithmic stability technique. [43] and [44] study the implicit bias of the next(last)-token classification loss when each token is sampled from a finite vocabulary. Specially, [44] proves that self-attention with gradient descent learns an automaton that generates the next token by hard retrieval and soft composition. [45] explains ICL as kernel regression. These results are not directly comparable to ours because we study the ICL of a one-layer linear transformer with AR pretraining on the first-order AR process.

3 Problem setup

Elementary notations. We define $[n] = \{1, 2, \dots, n\}$. We use lowercase, lowercase boldface, and uppercase boldface letters to denote scalars, vectors, and matrices, respectively. For a vector \mathbf{a} , we denote its i -th element as a_i . For a matrix \mathbf{A} , we use $\mathbf{A}_{k:}$, $\mathbf{A}_{:,k}$ and A_{ij} to denote its k -th row, k -th column and (i, j) -th element, respectively. For a vector \mathbf{a} (matrix \mathbf{A}), we use \mathbf{a}^* (\mathbf{A}^*) to denote its conjugate transpose. Similarly, we use $\bar{\mathbf{a}}$ and $\bar{\mathbf{A}}$ to denote their element-wise conjugate. We denote the n -dimensional identity matrix by \mathbf{I}_n . We denote the one vector of size n by $\mathbf{1}_n$. In addition, we denote the zero vector of size n and the zero matrix of size $m \times n$ by $\mathbf{0}_n$ and $\mathbf{0}_{m \times n}$, respectively. We use \otimes and \odot to denote the Kronecker product and the Hadamard product, respectively. Besides, we denote $\text{Vec}(\cdot)$ the vectorization operator in column-wise order.

3.1 Data distribution

We consider a first-order AR process as the underlying data distribution, similar to recent works on AR pretraining [16; 17]. Concretely, to generate a sequence $(\mathbf{x}_1, \dots, \mathbf{x}_T) \in \mathbb{C}^{d \times T}$, we first randomly sample a unitary matrix $\mathbf{W} \in \mathbb{C}^{d \times d}$ uniformly from a candidate set $\mathcal{P}_{\mathbf{W}} = \{\text{diag}(\lambda_1, \dots, \lambda_d) \mid |\lambda_i| = 1, \forall i \in [d]\}$ and the initial data point \mathbf{x}_1 from a controllable distribution $\mathcal{D}_{\mathbf{x}_1}$ to be specified later, then the subsequent elements are generated according to the rule $\mathbf{x}_{t+1} = \mathbf{W}\mathbf{x}_t$ for $t \in [T - 1]$. For convenience, we denote the vector $(\lambda_1, \dots, \lambda_d)^\top$ by $\boldsymbol{\lambda}$. We note that the structural assumption on \mathbf{W} is standard in the literature on learning problems involving matrices [46; 17]. In addition, it is a natural extension of the recent studies on ICL for linear problems [24; 26; 28; 29; 30], where they focus on $y = \mathbf{w}^\top \mathbf{x} = \mathbf{1}_d^\top \text{diag}(\mathbf{w})\mathbf{x}$. Furthermore, adopting this new data model is a necessary approach to investigate AR pretraining because the regression dataset in the previous ICL theory is not suitable since each token \mathbf{x}_i does not have a relation with the context. In this paper, we mainly investigate the impact of the initial distribution $\mathcal{D}_{\mathbf{x}_1}$ on the convergence result of the AR pretraining.

3.2 Model details

Linear causal self-attention layer. Before introducing the specified transformer module we will analyze in this paper, we first recall the definition of the standard causally-masked self-attention layer [1], whose parameters $\boldsymbol{\theta}$ includes a query matrix $\mathbf{W}^Q \in \mathbb{R}^{d_k \times d_e}$, a key matrix $\mathbf{W}^K \in \mathbb{R}^{d_k \times d_e}$, a value matrix $\mathbf{W}^V \in \mathbb{R}^{d_v \times d_e}$, a projection matrix $\mathbf{W}^P \in \mathbb{R}^{d_e \times d_v}$ and a normalization factor $\rho_t > 0$. At time step t , let $\mathbf{E}_t = (\mathbf{e}_1, \dots, \mathbf{e}_t) \in \mathbb{R}^{d_e \times t}$ be the tokens embedded from the prompt sequence

$\mathbf{P}_t = (\mathbf{x}_1, \dots, \mathbf{x}_t) \in \mathbb{C}^{d \times t}$, the causal self-attention layer will output

$$\mathbf{f}_t(\mathbf{E}_t; \boldsymbol{\theta}) = \mathbf{e}_t + \mathbf{W}^P \mathbf{W}^V \mathbf{E}_t \cdot \text{softmax} \left(\frac{(\mathbf{W}^K \mathbf{E}_t)^* \mathbf{W}^Q \mathbf{e}_t}{\rho_t} \right) \in \mathbb{C}^{d_e},$$

where the softmax operator is applied to each column of the input matrix. Similar to recent theoretical works on ICL with few-shot pretraining [24; 25; 16; 17; 26; 27; 28; 29], we consider the linear attention module in this work, which modifies the standard causal self-attention by dropping the softmax operator. Reparameterizing $\mathbf{W}^{KQ} = \mathbf{W}^{K*} \mathbf{W}^Q$ and $\mathbf{W}^{PV} = \mathbf{W}^P \mathbf{W}^V$, we have $\boldsymbol{\theta} = (\mathbf{W}^{KQ}, \mathbf{W}^{PV})$, and the output can be rewritten as:

$$\mathbf{f}_t(\mathbf{E}_t; \boldsymbol{\theta}) = \mathbf{e}_t + \mathbf{W}^{PV} \mathbf{E}_t \cdot \frac{\mathbf{E}_t^* \mathbf{W}^{KQ} \mathbf{e}_t}{\rho_t}.$$

Though it is called linear attention, we note that the output $\mathbf{f}_t(\mathbf{E}_t; \boldsymbol{\theta})$ is non-linear w.r.t. the input tokens \mathbf{E}_t due to the existence of the $\mathbf{E}_t \mathbf{E}_t^*$. In terms of the normalization factor ρ_t , like existing works [24; 26], we take it to be $t - 1$ because each element in $\mathbf{E}_t \mathbf{E}_t^*$ is a Hermitian inner product of two vectors of size t .

Embeddings. We adopt the natural embedding strategy used in recent studies on AR learning [16; 17]. Given a sequence $\mathbf{P}_t = (\mathbf{x}_1, \dots, \mathbf{x}_t)$, the i -th token is defined as $\mathbf{e}_i = (\mathbf{0}_d^\top, \mathbf{x}_i^\top, \mathbf{x}_{i-1}^\top)^\top \in \mathbb{C}^{3d}$, thus the corresponding embedding matrix \mathbf{E}_t can be formally written as:

$$\mathbf{E}_t = (\mathbf{e}_1, \dots, \mathbf{e}_t) = \begin{pmatrix} \mathbf{0}_d & \mathbf{0}_d & \cdots & \mathbf{0}_d \\ \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_t \\ \mathbf{x}_0 & \mathbf{x}_1 & \cdots & \mathbf{x}_{t-1} \end{pmatrix} \in \mathbb{C}^{3d \times t},$$

where $\mathbf{x}_0 = \mathbf{0}_d$ as a complement. This embedding strategy is a natural extension of the existing theoretical works about ICL for linear regression [30; 24; 25; 26; 28; 29]. The main difference is that the latter only focus on predicting the last query token while we need to predict each historical token. We note that practical transformers do learn similar token construction in the first softmax attention layer (e.g., see Fig. 4B in [16]).

Next-token prediction. Receiving the prompt $\mathbf{P}_t = (\mathbf{x}_1, \dots, \mathbf{x}_t)$, the network’s prediction for the next element \mathbf{x}_{t+1} will be the first d coordinates of the output $\mathbf{f}_t(\mathbf{E}_t; \boldsymbol{\theta})$, aligning with the setup adopted in [16; 17]. Namely, we have

$$\hat{\mathbf{y}}_t(\mathbf{E}_t; \boldsymbol{\theta}) = [\mathbf{f}_t(\mathbf{E}_t; \boldsymbol{\theta})]_{1:d} \in \mathbb{C}^d.$$

Henceforth, we will omit the dependence on \mathbf{E}_t and $\boldsymbol{\theta}$, and use $\hat{\mathbf{y}}_t$ if it is not ambiguous. Since only the first d rows are extracted from the output by the attention layer, the prediction $\hat{\mathbf{y}}_t$ just depends on some parts of \mathbf{W}^{PV} and \mathbf{W}^{KQ} . Concretely, we denote that

$$\mathbf{W}^{PV} = \begin{pmatrix} \mathbf{W}_{11}^{PV} & \mathbf{W}_{12}^{PV} & \mathbf{W}_{13}^{PV} \\ \mathbf{W}_{21}^{PV} & \mathbf{W}_{22}^{PV} & \mathbf{W}_{23}^{PV} \\ \mathbf{W}_{31}^{PV} & \mathbf{W}_{32}^{PV} & \mathbf{W}_{33}^{PV} \end{pmatrix}, \mathbf{W}^{KQ} = \begin{pmatrix} \mathbf{W}_{11}^{KQ} & \mathbf{W}_{12}^{KQ} & \mathbf{W}_{13}^{KQ} \\ \mathbf{W}_{21}^{KQ} & \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{31}^{KQ} & \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix},$$

where $\mathbf{W}_{ij}^{PV}, \mathbf{W}_{ij}^{KQ} \in \mathbb{R}^{d \times d}$ for all $i, j \in [3]$. Then the $\hat{\mathbf{y}}_t$ can be written as

$$\hat{\mathbf{y}}_t = (\mathbf{W}_{12}^{PV} \quad \mathbf{W}_{13}^{PV}) \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix} \mathbf{e}_t^x. \quad (1)$$

Here $\mathbf{E}_t^x = (\mathbf{e}_1^x, \dots, \mathbf{e}_t^x) \in \mathbb{C}^{2d \times t}$ denotes the last $2d$ rows of the \mathbf{E}_t , where $\mathbf{e}_i^x = (\mathbf{x}_i^\top, \mathbf{x}_{i-1}^\top)^\top$. Therefore, we only need to analyze the selected parameters in Eq. 1 during the training dynamics. The derivation of Eq. 1 can be found in Appendix A.1.

3.3 Training procedure

Loss function. To train the transformer model over the next-token prediction task, we focus on minimizing the following population loss:

$$L(\boldsymbol{\theta}) = \sum_{t=2}^{T-1} L_t(\boldsymbol{\theta}) = \sum_{t=2}^{T-1} \mathbb{E}_{\mathbf{x}_1, \mathbf{W}} \left[\frac{1}{2} \|\hat{\mathbf{y}}_t - \mathbf{x}_{t+1}\|_2^2 \right], \quad (2)$$

where the expectation is taken with respect to the start point \mathbf{x}_1 and the transition matrix \mathbf{W} . Henceforth, we will suppress the subscripts of the expectation for simplicity. The population loss is a standard objective in the optimization studies [24; 47], and this objective has been used in recent works on AR modeling [16; 17]. The summation starts from $t = 2$ because we do not have any information to predict \mathbf{x}_2 given only \mathbf{x}_1 .

Initialization strategy. We adopt the following diagonal initialization strategy, and similar settings have been used in recent works on ICL for linear problem [24; 30; 17].

Assumption 3.1 (Diagonal initialization). *At the initial time $\tau = 0$, we assume that*

$$\mathbf{W}^{KQ}(0) = \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & a_0 \mathbf{I}_d & \mathbf{0}_{d \times d} \end{pmatrix}, \mathbf{W}^{PV}(0) = \begin{pmatrix} \mathbf{0}_{d \times d} & b_0 \mathbf{I}_d & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \end{pmatrix},$$

where the red submatrices are related to the $\hat{\mathbf{y}}_t$ and changed during the training process.

The most related paper [17] considers a stronger diagonal structure than ours, and it only investigates the loss landscape. Our results deepened the understanding of AR transformers by considering practical training dynamics. We think this assumption might be inevitable for a tractable analysis and leave theory for standard (Gaussian) initialization to future work. In Section 6, we also conduct experiments under standard initialization, which further supports the rationality of Assumption 3.1.

Optimization algorithm. We utilize the gradient flow to minimize the learning objective in Eq. 2, which is equivalent to the gradient descent with infinitesimal step size and governed by the ordinary differential equation (ODE) $\frac{d}{d\tau} \boldsymbol{\theta} = -\nabla L(\boldsymbol{\theta})$.

4 Main results

In this section, we present the main theoretical results of this paper. First, in Section 4.1, we prove that when $\mathcal{D}_{\mathbf{x}_1}$ satisfies some certain condition (Assumption 4.1), the trained transformer implements one step of gradient descent for the minimization of an OLS problem, which validates the rationality of the mesa-optimization hypothesis [16]. Next, in Section 4.2, we further explore the capability limitation of the obtained mesa-optimizer under Assumption 4.1, where we characterize a stronger assumption (Assumption 4.2) as the necessary and sufficient condition that the mesa-optimizer recovers the true distribution. Finally, we go beyond Assumption 4.1, where the exploratory analysis proves that the trained transformer will generally not perform vanilla gradient descent for the OLS problem.

4.1 Trained transformer is a mesa-optimizer

In this subsection, we show that under a certain assumption of $\mathcal{D}_{\mathbf{x}_1}$, the trained one-layer linear transformer will converge to the mesa-optimizer [16; 17]. Namely, it will perform one step of gradient descent for the minimization of an OLS problem about the received prompt. The sufficient condition of the distribution $\mathcal{D}_{\mathbf{x}_1}$ can be summarized as follows.

Assumption 4.1 (Sufficient condition for the emergence of mesa-optimizer). *We assume that the distribution $\mathcal{D}_{\mathbf{x}_1}$ of the initial token $\mathbf{x}_1 \in \mathbb{R}^d$ satisfies $\mathbb{E}_{\mathbf{x}_1 \sim \mathcal{D}_{\mathbf{x}_1}} [x_{1i_1} x_{1i_2}^{r_2} \cdots x_{1i_n}^{r_n}] = 0$ for any subset $\{i_1, \dots, i_n \mid n \leq 4\}$ of $[d]$, and $r_2, \dots, r_n \in \mathbb{N}$. In addition, we assume that $\kappa_1 = \mathbb{E}[x_{1j}^4]$, $\kappa_2 = \mathbb{E}[x_{1j}^6]$ and $\kappa_3 = \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4]$ are finite constant for any $j \in [d]$.*

Finding Assumption 4.1 is non-trivial since we need to derive the training dynamics first. The key intuition of this assumption is to keep the gradient of the non-diagonal elements of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} as zero, thus they can keep diagonal structure during the training. We note that any random vectors \mathbf{x}_1 whose coordinates x_{1i} are i.i.d. random variables with zero mean and finite moments of order 2, 4, and 6 satisfy this assumption. For example, it includes the normal distribution $\mathcal{N}(\mathbf{0}_d, \sigma^2 \mathbf{I}_d)$, which is a common setting in the learning theory field [47; 48; 49; 50; 51]. Under this assumption, the final fixed point found by the gradient flow can be characterized as the following theorem.

Theorem 4.1 (Convergence of the gradient flow, proof in Section 5). *Consider the gradient flow of the one-layer linear transformer (see Eq. 1) over the population AR pretraining loss (see Eq. 2).*

Suppose the initialization satisfies Assumption 3.1, and the initial token's distribution $\mathcal{D}_{\mathbf{x}_1}$ satisfies Assumption 4.1, then the gradient flow converges to

$$\begin{pmatrix} \widetilde{\mathbf{W}}_{22}^{KQ} & \widetilde{\mathbf{W}}_{23}^{KQ} \\ \widetilde{\mathbf{W}}_{32}^{KQ} & \widetilde{\mathbf{W}}_{33}^{KQ} \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \widetilde{a} \mathbf{I}_d & \mathbf{0}_{d \times d} \end{pmatrix}, \begin{pmatrix} \widetilde{\mathbf{W}}_{12}^{PV} & \widetilde{\mathbf{W}}_{13}^{PV} \end{pmatrix} = \begin{pmatrix} \widetilde{b} \mathbf{I}_d & \mathbf{0}_{d \times d} \end{pmatrix}.$$

Though different initialization (a_0, b_0) lead to different $(\widetilde{a}, \widetilde{b})$, the solutions' product $\widetilde{a}\widetilde{b}$ satisfies

$$\widetilde{a}\widetilde{b} = \frac{\kappa_1}{\kappa_2 + \frac{\kappa_3}{T-2} \sum_{t=2}^{T-1} \frac{1}{t-1}}.$$

As far as we know, Theorem 4.1 is the first theoretical result for the training dynamics and the mesa-optimization hypothesis of autoregressive transformers. The technical challenge compared to existing ICL theory for regression [24] mainly has two parts. First, our data model breaks the independence between data at different times, which causes difficulty in decomposing and estimating the gradient terms. Second, we modify the embedding strategy (more dimensions), scale the attention model (much more parameters), and change the loss function (more terms) to perform the full AR pertaining. All these parts are not well studied in the literature and make the gradients more complicated.

Theorem 4.1 is also a non-trivial extension of recent work [17], which characterizes the global minima of the AR modeling loss when imposing the diagonal structure on all parameter matrices during the training and fixing \mathbf{x}_1 as $\mathbf{1}_d$. In comparison, Theorem 4.1 does not depend on the special structure, and further investigates when the mesa-optimizer emerges in practical non-convex optimization.

We highlight that the limiting solution found by the gradient flow shares the same structure with the careful construction in [16], though the pretraining loss is non-convex. Therefore, our result theoretically validates the rationality of the mesa-optimization hypothesis [16] in the AR pretraining setting, which can be formally presented as the following corollary.

Corollary 4.1 (Trained transformer as a mesa-optimizer, proof in Appendix A.3). *We suppose that the same precondition of Theorem 4.1 holds. When predicting the $(t+1)$ -th token, the trained transformer obtains $\widehat{\mathbf{W}}$ by implementing one step of gradient descent for the OLS problem $L_{\text{OLS},t}(\mathbf{W}) = \frac{1}{2} \sum_{i=1}^{t-1} \|\mathbf{x}_{i+1} - \mathbf{W}\mathbf{x}_i\|^2$, starting from the initialization $\mathbf{W} = \mathbf{0}_{d \times d}$ with a step size $\frac{\widetilde{a}\widetilde{b}}{t-1}$.*

4.2 Capability limitation of the mesa-optimizer

Built upon the findings in Theorem 4.1, a simple calculation (details in Appendix A.3) shows that the prediction of the obtained mesa-optimizer given a new test prompt of length T_{te} is

$$\widehat{\mathbf{y}}_{T_{te}} = \mathbf{W} \left(\frac{\widetilde{a}\widetilde{b} \sum_{i=1}^{T_{te}-1} \mathbf{x}_i \mathbf{x}_i^*}{T_{te} - 1} \right) \mathbf{x}_{T_{te}}. \quad (3)$$

It is natural to ask the question: where is the capability limitation of the obtained mesa-optimizer, and what data distribution can the trained transformer learn? Therefore, in this subsection, we study under what assumption of the initial token's distribution $\mathcal{D}_{\mathbf{x}_1}$, the one step of gradient descent performed by the trained transformer can exactly recover the underlying data distribution. First, leveraging the result from Eq. 3, we present a negative result, which proves that not all $\mathcal{D}_{\mathbf{x}_1}$ satisfies Assumption 4.1 can be recovered by the trained linear transformer.

Proposition 4.1 (AR process with normal distributed initial token can not be learned, proof in Appendix A.4). *Let $\mathcal{D}_{\mathbf{x}_1}$ be the multivariate normal distribution $\mathcal{N}(\mathbf{0}_d, \sigma^2 \mathbf{I}_d)$ with any $\sigma^2 > 0$, then the "simple" AR process can not be recovered by the trained transformer even in the ideal case with long training context. Formally, when the training sequence length T_{tr} is large enough, for any test context length T_{te} and dimension $j \in [d]$, the prediction from the trained transformer satisfies*

$$E_{\mathbf{x}_1, \mathbf{W}} \left[\frac{(\widehat{\mathbf{y}}_{T_{te}})_j}{(\mathbf{W} \mathbf{x}_{T_{te}})_j} \right] \rightarrow \frac{1}{5}.$$

Therefore, the prediction $\widehat{\mathbf{y}}_{T_{te}}$ will not converges to the true next token $\mathbf{W} \mathbf{x}_{T_{te}}$.

Proposition 4.1 suggests that ICL by AR pretraining [16; 17] is different from ICL by few-shot pretraining [24; 25; 26; 27; 28; 29; 30; 23], which attracts much more attention from the theoretical community. In the latter setting, recent works [24; 26] proves that one step of gradient descent implemented by the trained transformer can in-context learn the linear regression problem with input sampled from $\mathcal{N}(\mathbf{0}_d, \sigma^2 \mathbf{I}_d)$. However, in the AR learning setting, the trained linear transformer fails.

This negative result shows that one-step GD learned by the AR transformer can not recover the distribution, but this can be solved by more complex models. Even for more complex data (\mathbf{W} is not diagonal), [16] has empirically verified that multi-layer linear attention can perform multi-step gradient descent to learn the data distribution. Therefore, to address the issue, future works are suggested to study more complex architecture such as softmax attention [30], multi-head attention [52; 53; 54], deeper attention layers [55; 16], transformer block [56; 57; 58], and so on. Future theory considering more complex AR transformers can adopt the same data model and token embeddings in this paper, and try to use a similar proof technique to derive the training dynamics.

Proposition 4.1 implies that if we want the trained transformer to recover the data distribution by performing one step of gradient descent, a stronger condition of $\mathcal{D}_{\mathbf{x}_1}$ is needed. Under Assumption 4.1, the following sufficient and necessary condition related to the moment of $\mathcal{D}_{\mathbf{x}_1}$ is derived from Eq. 3 by letting $\hat{\mathbf{y}}_{T_{te}}$ converges to $\mathbf{W}\mathbf{x}_{T_{te}}$ when context length T_{tr} and T_{te} are large enough.

Assumption 4.2 (Condition for success of mesa-optimizer). *Based on Assumption 4.1, we further suppose that $\frac{\kappa_1}{\kappa_2} \frac{\sum_{i=1}^{T_{te}-1} \mathbf{x}_i \mathbf{x}_i^*}{T_{te}-1} \mathbf{x}_{T_{te}} \rightarrow \mathbf{x}_{T_{te}}$ for any \mathbf{x}_1 and \mathbf{W} , when T_{te} is large enough.*

Assumption 4.2 is strong and shows the poor capability of the trained one-layer linear transformer because common distribution (e.g. Gaussian distribution, Gamma distribution, Poisson distribution, etc) always fails to satisfy this condition. Besides, it is a sufficient and necessary condition for the mesa-optimizer to succeed when the distribution $\mathcal{D}_{\mathbf{x}_1}$ has satisfied Assumption 4.1, thus can not be improved in this case. We construct the following example that satisfies Assumption 4.2.

Example 4.1 (sparse vector). If the random vector $\mathbf{x}_1 \in \mathbb{R}^d$ is uniformly sampled from the candidate set of size $2d$ $\{\pm(c, 0, \dots, 0)^\top, \pm(0, c, \dots, 0)^\top, \pm(0, \dots, 0, c)^\top\}$ for any fixed $c \in \mathbb{R}$, then the distribution $\mathcal{D}_{\mathbf{x}_1}$ satisfies Assumption 4.2. The derivation can be found in Appendix A.5.

For completeness, we formally summarize the following distribution learning guarantee for the trained transformer under Assumption 3.1 and 4.1.

Theorem 4.2 (Trained transformer succeed to learn the distribution satisfies Assumption 4.2, proof in Appendix A.6). *Suppose that Assumption 3.1 and 4.1 holds, then Assumption 4.2 is the sufficient and necessary condition for the trained transformer to learn the AR process. Formally, when the training sequence length T_{tr} and test context length T_{te} are large enough, the prediction from the trained transformer satisfies*

$$\hat{\mathbf{y}}_{T_{te}} \rightarrow \mathbf{W}\mathbf{x}_{T_{te}}, \quad T_{tr}, T_{te} \rightarrow +\infty.$$

4.3 Go beyond the Assumption 4.1

The behavior of the gradient flow under Assumption 4.1 has been clearly understood in Theorem 4.1. The follow-up natural question is what solution will be found by the gradient flow when Assumption 4.1 does not hold. In this subsection, we conduct exploratory analyses by adopting the setting in [17], where the initial token \mathbf{x}_1 is fixed as $\mathbf{1}_d$.

First, sharing the similar but weaker assumption of [17], we impose \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} to stay diagonal during training by masking the non-diagonal gradients, then the trained transformer will perform one step of gradient descent, as suggested by [17]. Formally, it can be written as follows.

Theorem 4.3 (Trained transformer as mesa-optimizer with non-diagonal gradient masking, proof in Appendix A.7). *Suppose the initialization satisfies Assumption 3.1, the initial token is fixed as $\mathbf{1}_d$, and we clip non-diagonal gradients of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} during the training, then the gradient flow of the one-layer linear transformer over the population AR loss converges to the same structure as the result in Theorem 4.1, with*

$$\tilde{ab} = \frac{1}{1 + \frac{d-1}{T-2} \sum_{t=2}^{T-1} \frac{1}{t-1}}.$$

Therefore, the obtained transformer performs one step of gradient descent in this case.

Theorem 4.3 can be seen as a complement and an extension of Proposition 2 in [17] from the perspective of optimization. We note that [17] assumes all the parameter matrices to be diagonal and only analyzes the global minima without considering the practical non-convex optimization process.

Next, we adopt some exploratory analyses for the gradient flow without additional non-diagonal gradient masking. The convergence result of the gradient flow can be asserted as the following proposition. The key intuition of its proof is that when the parameters matrices share the same structure as the result in Theorem 4.1, the non-zero gradients of the non-diagonal elements of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} will occur. In addition, we note the result does not depend on Assumption 3.1.

Proposition 4.2 (Trained transformer does not perform on step of gradient descent, proof in Appendix A.8). *The limiting point found by the gradient does not share the same structure as that in Theorem 4.1, thus the trained transformer will not implement one step of vanilla gradient descent for minimizing the OLS problem $\frac{1}{2} \sum_{i=1}^{t-1} \|\mathbf{x}_{i+1} - \mathbf{W}\mathbf{x}_i\|^2$.*

To fully solve the problem and find the limiting point of the gradient flow in this case (or more generally, any case beyond Assumption 3.1 and 4.1), one can not enjoy the diagonal structure of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} anymore. When \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} are general dense matrices, computation of the gradient will be much more difficult than that in Proposition 4.2. Therefore, we leave the general rigorous result of convergence without Assumption 3.1 and 4.1 for future work.

We are aware that recent theoretical studies on ICL for linear regression have faced a similar problem. [24; 26; 23] find that when the input’s distribution does not satisfy Assumption 4.1 (e.g., $\mathcal{N}(\mathbf{0}_d, \Sigma)$), the trained transformer will implement one step of preconditioned gradient descent on for some inner objective function. We conjecture similar results will hold in the case of in-context AR learning. We will empirically verify this conjecture when \mathbf{x}_1 is a full one vector, in Section 6.

5 Proof skeleton

In this section, we outline the proof ideas of Theorem 4.1, which is one of the core findings of this paper, and also a theoretical base of the more complex proofs of Theorem 4.3 and Proposition 4.2. The full proof of this Theorem is placed in Appendix A.2.

The first key step is to observe that each coordinate of prediction $\hat{\mathbf{y}}_t$ (Eq. 1) can be written as the output of a quadratic function, which will greatly simplify the follow-up gradient operation.

Lemma 5.1 (Simplification of $\hat{\mathbf{y}}_{t,j}$, proof in Appendix A.2.1). *Each element of the network’s prediction $\hat{\mathbf{y}}_{t,j}$ ($j \in [d]$) can be expressed as the following.*

$$\hat{\mathbf{y}}_{t,j} = \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) = \text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}\top}}}{\rho_t} \cdot \mathbf{B}_j,$$

where the \mathbf{A} and \mathbf{B}_j are defined as

$$\mathbf{A} = (\mathbf{a}_1 \quad \dots \quad \mathbf{a}_{2d}) = \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix}, \quad \mathbf{B}_j = \begin{pmatrix} \mathbf{b}_{j1} \\ \mathbf{b}_{j2} \end{pmatrix} = \begin{pmatrix} \mathbf{W}_{12,j}^{PV\top} \\ \mathbf{W}_{13,j}^{PV\top} \end{pmatrix},$$

with $\mathbf{a}_i \in \mathbb{R}^{2d}$ and $\mathbf{b}_{j1}, \mathbf{b}_{j2} \in \mathbb{R}^d$.

Next, We calculate the gradient for the parameter matrices of the linear transformer and present the dynamical system result, which is the most complex part in the proof of Theorem 4.1.

Lemma 5.2 (dynamical system of gradient flow under Assumption 4.1, proof in Appendix A.2.2). *Suppose that Assumption 4.1 holds, then the dynamical process of the parameters in the diagonal of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} satisfies*

$$\begin{aligned} \frac{d}{d\tau} a &= -ab^2 \left[(T-2)\kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] + b(T-2)\kappa_1, \\ \frac{d}{d\tau} b &= -a^2b \left[(T-2)\kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] + a(T-2)\kappa_1, \end{aligned}$$

while the gradients for all other parameters were kept at zero during the training process.

Similar ODEs have occurred in existing studies, such as the deep linear networks [59] and recent ICL for linear regression [24]. Notably, these dynamics are the same as those of gradient flow on a non-convex objective function with clear global minima, which is summarized as the following.

Lemma 5.3 (Surrogate objective function, proof in Appendix A.2.3). *Suppose that Assumption 4.1 holds and denote $(T - 2)\kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1}\kappa_3$ and $(T - 2)\kappa_1$ by c_1 and c_2 , respectively. Then, the dynamics in Lemma 5.2 are the same as those of gradient flow on the following objective function:*

$$\tilde{\ell}(a, b) = \frac{1}{2c_1}(c_2 - c_1ab)^2,$$

whose global minimums satisfy $ab = c_2/c_1$.

Furthermore, We show that although the objective $\tilde{\ell}(a, b)$ is non-convex, the Polyak-Łojasiewicz (PL) inequality [60; 61] holds, which implies that gradient flow converges to the global minimum.

Lemma 5.4 (Global convergence of gradient flow, proof in Appendix A.2.4). *Suppose that Assumption 4.1 holds, then $\tilde{\ell}(a, b)$ is a non-convex function and satisfies the PL inequality as follows.*

$$\left| \frac{\partial}{\partial a} \tilde{\ell}(a, b) \right|^2 + \left| \frac{\partial}{\partial b} \tilde{\ell}(a, b) \right|^2 \geq 2c_1(a^2 + b^2) \left(\tilde{\ell}(a, b) - \min_{a,b} \tilde{\ell}(a, b) \right).$$

Therefore, the gradient flow in Lemma 5.2 converges to the global minimum of $\tilde{\ell}(a, b)$.

Finally, Theorem 4.1 can be proved by directly applying the above lemmas.

6 Simulation results

In this section, we conduct simulations to verify and generalize our theoretical results. In terms of the train set, we generate $10k$ sequences with $T_{tr} = 100$ and $d = 5$. In addition, we generate another test set with $10k$ sequences of the same shape. We train for 200 epochs with vanilla gradient descent, with different diagonal initialization of (a_0, b_0) by $(0.1, 0.1)$, $(0.5, 1.5)$, $(2, 2)$. The detailed configurations (e.g., step size) and results of different experiments can be found in Appendix B.

Initial token sampled from $\mathcal{N}(\mathbf{0}_d, \sigma^2 \mathbf{I}_d)$. We conduct simulations with $\sigma = 0.5, 1, 2$ respectively. With any initialization of (a_0, b_0) , simulations show that ab converges to $\kappa_1/\kappa_2 = 1/5\sigma^2$, and $\hat{\mathbf{y}}_{T_{te}-1}$ converges to $\mathbf{x}_{T_{te}}/5$ in expectation, which verifies Theorem 4.1 and Proposition 4.1, respectively. In the main paper, we present the convergence results with $\sigma = 0.5$ in Fig. 1a and 1b. We also verify our theory in the small-context scenarios ($T_{tr} = 5$), which is placed in Fig. 4 in Appendix B.3.

Initial token sampled from Example 4.1. We conduct simulations with scale $c = 0.5, 1, 2$ respectively. With any initialization of (a_0, b_0) , simulations show that ab converges to $\kappa_1/\kappa_2 = 1/c^2$ (see details in Appendix A.5), and $\hat{\mathbf{y}}_{T_{te}-1}$ converges to the truth $\mathbf{x}_{T_{te}}$, which verifies Theorem 4.1 and Theorem 4.2, respectively. In the main paper, we present the results with $c = 0.5$ in Fig. 1c and 1d.

Initial token fixed as $\mathbf{1}_d$. We conduct experiment with $\mathbf{x}_1 = \mathbf{1}_d$. The results Fig. 7 in Appendix B.5 show that \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} converge to dense matrices with strong diagonals, and other matrices converge to $\mathbf{0}_{d \times d}$, which means that the trained transformer performs somewhat preconditioned gradient descent. The detailed derivation is placed in Appendix B.5.

Go beyond the diagonal initialization. Finally, in order to extend our theory, we repeat experiments under Gaussian initialization with different variance ($\sigma_w = 0.001, 0.01, 0.1$). The results of Gaussian start points and sparse start points (Example 4.1) can be found in Fig. 5 of Appendix B.3 and Fig. 6 of Appendix B.4, respectively. As a result, though the convergence results of parameters are not the same as those under diagonal initialization, they keep the same diagonal structure, which can be understood as GD with adaptive learning rate in different dimensions. In addition, the test results (ratio or MSE loss) under the standard Gaussian initialization are the same as those under diagonal initialization, which further verifies the capability limitation of the trained transformers. To sum up, these experimental results demonstrate that our theoretical results have a certain representativeness, which further supports the rationality of the diagonal initialization.

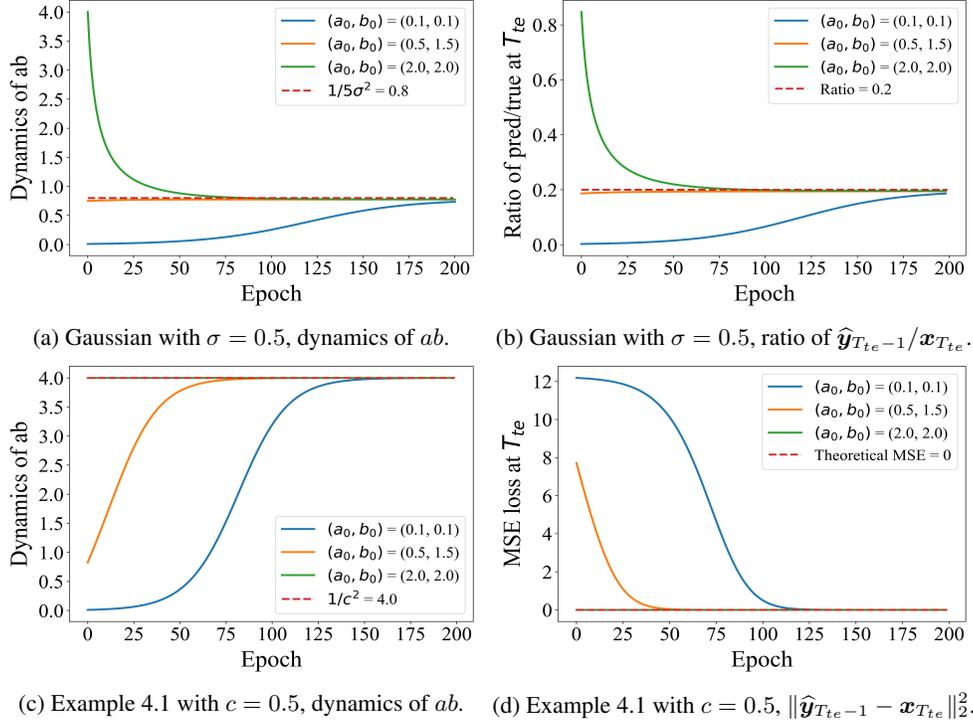


Figure 1: Simulations results on Gaussian and Example 4.1 show that the convergence of ab satisfies Theorem 4.1. In addition, the trained transformer can recover the sequence with the initial token from Example 4.1, but fails to recover the Gaussian initial token, which verifies Theorem 4.2 and Proposition 4.1, respectively.

7 Conclusion and Discussion

In this paper, we towards understanding the the mechanisms underlying the ICL by analyzing the mesa-optimization hypothesis. To achieve this goal, we investigate the non-convex dynamics of a one-layer linear transformer autoregressively trained by gradient flow on a controllable AR process. First, we find a sufficient condition (Assumption 4.1) for the emergence of mesa-optimizer. Second, we explore the capability of the mesa-optimizer, where we find a sufficient and necessary condition (Assumption 4.2) that the trained transformer recovers the true distribution. Third, we analyze the case where Assumption 4.1 does not hold, and find that the trained transformer will not perform vanilla gradient descent in general. Finally, our simulation results verify the theoretical results.

Limitations and social impact. First, our theory only focuses on the one-layer linear transformer, thus whether the results hold when more complex models are adopted is still unclear. We believe that our analysis can give insight to those cases. Second, the general case where Assumption 3.1 and 4.1 does not hold is not fully addressed in this paper due to technical difficulties. Future work can consider that setting based on our theoretical and empirical findings. Finally, this is mainly theoretical work and we do not see a direct social impact of our theory.

8 Acknowledgement

This work was supported by Beijing Natural Science Foundation (L247030); NSF of China (Nos. 62076145, 62206159); Beijing Nova Program (No. 20230484416); Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Renmin University of China; the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China (22XNKJ13); the Natural Science Foundation of Shandong Province (Nos. ZR2022QF117), the Fundamental Research Funds of Shandong University. The work was partially done at the Engineering Research Center of Next-Generation Intelligent Search and Recommendation, Ministry of Education. G. Wu was also sponsored by the TaiShan Scholars Program.

References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, pages 5998–6008, 2017.
- [2] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *NeurIPS*, 2020.
- [3] OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023.
- [4] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023.
- [5] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- [6] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [7] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, pages 8821–8831, 2021.
- [8] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *CVPR*, pages 11523–11532, 2022.
- [9] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024.
- [10] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *ICML*, volume 119, pages 1691–1703, 2020.
- [11] Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. The dawn of Imms: Preliminary explorations with gpt-4v(ision). *CoRR*, abs/2309.17421, 2023.
- [12] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, and Jifeng Dai. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. In *NeurIPS*, 2023.
- [13] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023.

- [14] Jun Zhan, Junqi Dai, Jiasheng Ye, Yunhua Zhou, Dong Zhang, Zhigeng Liu, Xin Zhang, Ruibin Yuan, Ge Zhang, Linyang Li, Hang Yan, Jie Fu, Tao Gui, Tianxiang Sun, Yugang Jiang, and Xipeng Qiu. Anygpt: Unified multimodal LLM with discrete sequence modeling. *CoRR*, abs/2402.12226, 2024.
- [15] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer, Mia Glaese, Thibault Sottiaux, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, James Molloy, Jilin Chen, Michael Isard, Paul Barham, Tom Hennigan, Ross McIlroy, Melvin Johnson, Johan Schalkwyk, Eli Collins, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Clemens Meyer, Gregory Thornton, Zhen Yang, Henryk Michalewski, Zaheer Abbas, Nathan Schucher, Ankesh Anand, Richard Ives, James Keeling, Karel Lenc, Salem Haykal, Siamak Shakeri, Pranav Shyam, Aakanksha Chowdhery, Roman Ring, Stephen Spencer, Eren Sezener, and et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *CoRR*, abs/2403.05530, 2024.
- [16] Johannes von Oswald, Eyvind Niklasson, Maximilian Schlegel, Seijin Kobayashi, Nicolas Zucchet, Nino Scherrer, Nolan Miller, Mark Sandler, Blaise Agüera y Arcas, Max Vladymyrov, Razvan Pascanu, and João Sacramento. Uncovering mesa-optimization algorithms in transformers. *CoRR*, abs/2309.05858, 2023.
- [17] Michael Eli Sander, Raja Giryes, Taiji Suzuki, Mathieu Blondel, and Gabriel Peyré. How do transformers perform in-context autoregressive learning ? In *ICML*, 2024.
- [18] Evan Hubinger, Chris van Merwijk, Vladimir Mikulik, Joar Skalse, and Scott Garrabrant. Risks from learned optimization in advanced machine learning systems. *CoRR*, abs/1906.01820, 2019.
- [19] Stephanie C. Y. Chan, Adam Santoro, Andrew K. Lampinen, Jane X. Wang, Aaditya K. Singh, Pierre H. Richemond, James L. McClelland, and Felix Hill. Data distributional properties drive emergent in-context learning in transformers. In *NeurIPS*, 2022.
- [20] Satwik Bhattamishra, Arkil Patel, Phil Blunsom, and Varun Kanade. Understanding in-context learning in transformers and llms by learning to learn discrete functions. In *ICLR*, 2024.
- [21] Kartik Ahuja and David Lopez-Paz. A closer look at in-context learning under distribution shifts. *CoRR*, abs/2305.16704, 2023.
- [22] Sewon Min, Xixi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In *EMNLP*, pages 11048–11064, 2022.
- [23] Sadegh Mahdavi, Renjie Liao, and Christos Thrampoulidis. Revisiting the equivalence of in-context learning and gradient descent: The impact of data distribution. In *ICASSP*, pages 7410–7414. IEEE, 2024.
- [24] Ruiqi Zhang, Spencer Frei, and Peter L Bartlett. Trained transformers learn linear models in-context. *Journal of Machine Learning Research*, 25(49):1–55, 2024.
- [25] Johannes von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. In *ICML*, volume 202, pages 35151–35174, 2023.
- [26] Kwangjun Ahn, Xiang Cheng, Hadi Daneshmand, and Suvrit Sra. Transformers learn to implement preconditioned gradient descent for in-context learning. In *NeurIPS*, 2023.
- [27] Kwangjun Ahn, Xiang Cheng, Minhak Song, Chulhee Yun, Ali Jadbabaie, and Suvrit Sra. Linear attention is (maybe) all you need (to understand transformer optimization). In *ICLR*, 2024.
- [28] Arvind V. Mahankali, Tatsunori Hashimoto, and Tengyu Ma. One step of gradient descent is provably the optimal in-context learner with one layer of linear self-attention. In *ICLR*, 2024.

- [29] Jingfeng Wu, Difan Zou, Zixiang Chen, Vladimir Braverman, Quanquan Gu, and Peter L. Bartlett. How many pretraining tasks are needed for in-context learning of linear regression? In *ICLR*, 2024.
- [30] Yu Huang, Yuan Cheng, and Yingbin Liang. In-context convergence of transformers. In *ICML*, 2024.
- [31] Shivam Garg, Dimitris Tsipras, Percy Liang, and Gregory Valiant. What can transformers learn in-context? A case study of simple function classes. In *NeurIPS*, 2022.
- [32] Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. Why can GPT learn in-context? language models secretly perform gradient descent as meta-optimizers. In *Findings of ACL*, pages 4005–4019. Association for Computational Linguistics, 2023.
- [33] Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. What learning algorithm is in-context learning? investigations with linear models. In *ICLR*, 2023.
- [34] Max Vladymyrov, Johannes von Oswald, Mark Sandler, and Rong Ge. Linear transformers are versatile in-context learners. *CoRR*, abs/2402.14180, 2024.
- [35] Samuel Müller, Noah Hollmann, Sebastian Pineda-Arango, Josif Grabocka, and Frank Hutter. Transformers can do bayesian inference. In *ICLR*, 2022.
- [36] Deqing Fu, Tian-Qi Chen, Robin Jia, and Vatsal Sharan. Transformers learn higher-order optimization methods for in-context learning: A study with linear models. *CoRR*, abs/2310.17086, 2023.
- [37] Angeliki Giannou, Liu Yang, Tianhao Wang, Dimitris Papailiopoulos, and Jason D. Lee. How well can transformers emulate in-context newton’s method? *CoRR*, abs/2403.03183, 2024.
- [38] Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit bayesian inference. In *ICLR*, 2022.
- [39] Xinyi Wang, Wanrong Zhu, and William Yang Wang. Large language models are implicitly topic models: Explaining and finding good demonstrations for in-context learning. *CoRR*, abs/2301.11916, 2023.
- [40] Noam Wies, Yoav Levine, and Amnon Shashua. The learnability of in-context learning. In *NeurIPS*, 2023.
- [41] Yuchen Li, Yuanzhi Li, and Andrej Risteski. How do transformers learn topic structure: Towards a mechanistic understanding. In *ICML*, volume 202, pages 19689–19729, 2023.
- [42] Yingcong Li, Muhammed Emrullah Ildiz, Dimitris Papailiopoulos, and Samet Oymak. Transformers as algorithms: Generalization and stability in in-context learning. In *ICML*, volume 202, pages 19565–19594, 2023.
- [43] Christos Thrampoulidis. Implicit bias of next-token prediction. *CoRR*, abs/2402.18551, 2024.
- [44] Yingcong Li, Yixiao Huang, Muhammed E Ildiz, Ankit Singh Rawat, and Samet Oymak. Mechanics of next token prediction with self-attention. In *International Conference on Artificial Intelligence and Statistics*, pages 685–693, 2024.
- [45] Chi Han, Ziqi Wang, Han Zhao, and Heng Ji. In-context learning of large language models explained as kernel regression. *CoRR*, abs/2305.12766, 2023.
- [46] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *NeurIPS*, pages 7411–7422, 2019.
- [47] Kulin Shah, Sitan Chen, and Adam Klivans. Learning mixtures of gaussians using the ddpm objective. *NeurIPS*, 36:19636–19649, 2023.
- [48] Haiyun He, Hanshu Yan, and Vincent YF Tan. Information-theoretic characterization of the generalization error for iterative semi-supervised learning. *Journal of Machine Learning Research*, 23(287):1–52, 2022.

- [49] Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. *NeurIPS*, 31, 2018.
- [50] Chenyu Zheng, Guoqiang Wu, and Chongxuan Li. Toward understanding generative data augmentation. *NeurIPS*, 36, 2024.
- [51] Chenyu Zheng, Guoqiang Wu, Fan Bao, Yue Cao, Chongxuan Li, and Jun Zhu. Revisiting discriminative vs. generative classifiers: Theory and implications. In *ICML*, volume 202, pages 42420–42477. PMLR, 2023.
- [52] Yingqian Cui, Jie Ren, Pengfei He, Jiliang Tang, and Yue Xing. Superiority of multi-head attention in in-context linear regression. *CoRR*, abs/2401.17426, 2024.
- [53] Siyu Chen, Heejune Sheen, Tianhao Wang, and Zhuoran Yang. Training dynamics of multi-head softmax attention for in-context learning: Emergence, convergence, and optimality (extended abstract). In Shipra Agrawal and Aaron Roth, editors, *COLT*, volume 247, page 4573, 2024.
- [54] Puneesh Deora, Rouzbeh Ghaderi, Hossein Taheri, and Christos Thrampoulidis. On the optimization and generalization of multi-head attention. *CoRR*, abs/2310.12680, 2023.
- [55] Eshaan Nichani, Alex Damian, and Jason D. Lee. How transformers learn causal structure with gradient descent. *CoRR*, abs/2402.14735, 2024.
- [56] Ruiqi Zhang, Jingfeng Wu, and Peter L. Bartlett. In-context learning of a linear transformer block: Benefits of the MLP component and one-step GD initialization. *CoRR*, abs/2402.14951, 2024.
- [57] Hongkang Li, Meng Wang, Songtao Lu, Xiaodong Cui, and Pin-Yu Chen. Training nonlinear transformers for efficient in-context learning: A theoretical learning and generalization analysis. *CoRR*, abs/2402.15607, 2024.
- [58] Juno Kim and Taiji Suzuki. Transformers learn nonlinear features in context: Nonconvex mean-field dynamics on the attention landscape. In *ICML*, 2024.
- [59] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *ICLR*, 2014.
- [60] Boris Teodorovich Polyak. Gradient methods for minimizing functionals. *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*, 3(4):643–653, 1963.
- [61] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-lojasiewicz condition. In *ECML PKDD*, volume 9851, pages 795–811, 2016.
- [62] Kaare Brandt Petersen, Michael Syskind Pedersen, et al. The matrix cookbook. *Technical University of Denmark*, 7(15):510, 2008.

Contents of Appendix

A Proofs	16
A.1 Proof of eq. (1)	16
A.2 Proof of Theorem 4.1	16
A.2.1 Proof of Lemma 5.1	16
A.2.2 Proof of Lemma 5.2	17
A.2.3 Proof of Lemma 5.3	27
A.2.4 Proof of Lemma 5.4	28
A.3 Proof of Corollary 4.1	28
A.4 Proof of Proposition 4.1	29
A.5 Derivation of Example 4.1	30
A.6 Proof of Theorem 4.2	30
A.7 Proof of Theorem 4.3	31
A.8 Proof of Proposition 4.2	36
B Experimental details and additional results	36
B.1 GPU and random seed	36
B.2 Step size in simulations	36
B.3 Additional results for Gaussian initial token	36
B.4 Additional results for Sparse initial token (Example 4.1)	36
B.5 Additional results for full-one initial token	37

Appendix A Proofs

A.1 Proof of eq. (1)

Proof. We first calculate the output by causal linear attention layer as the following:

$$\begin{aligned}
& f_t(\mathbf{E}_t; \boldsymbol{\theta}) \\
&= \mathbf{e}_t + \frac{1}{\rho_t} \mathbf{W}^{PV} \mathbf{E}_t \mathbf{E}_t^* \mathbf{W}^{KQ} \mathbf{e}_t \\
&= \mathbf{e}_t + \frac{1}{\rho_t} \mathbf{W}^{PV} \left(\sum_{i=1}^t \mathbf{e}_i \mathbf{e}_i^* \right) \mathbf{W}^{KQ} \mathbf{e}_t \\
&= \mathbf{e}_t + \frac{1}{\rho_t} \sum_{i=1}^t \left(\mathbf{W}^{PV} \mathbf{e}_i \cdot \left(\mathbf{W}^{KQ\top} \mathbf{e}_i \right)^* \right) \mathbf{e}_t \\
&= \mathbf{e}_t + \frac{1}{\rho_t} \sum_{i=1}^t \left(\begin{pmatrix} \mathbf{W}_{11}^{PV} & \mathbf{W}_{12}^{PV} & \mathbf{W}_{13}^{PV} \\ \mathbf{W}_{21}^{PV} & \mathbf{W}_{22}^{PV} & \mathbf{W}_{23}^{PV} \\ \mathbf{W}_{31}^{PV} & \mathbf{W}_{32}^{PV} & \mathbf{W}_{33}^{PV} \end{pmatrix} \mathbf{e}_i \cdot \left(\begin{pmatrix} \mathbf{W}_{11}^{KQ} & \mathbf{W}_{12}^{KQ} & \mathbf{W}_{13}^{KQ} \\ \mathbf{W}_{21}^{KQ} & \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{31}^{KQ} & \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix}^\top \mathbf{e}_i \right)^* \right) \mathbf{e}_t \\
&= \begin{pmatrix} \mathbf{0}_d \\ \mathbf{x}_t \\ \mathbf{x}_{t-1} \end{pmatrix} + \frac{1}{\rho_t} \sum_{i=1}^t \left(\begin{pmatrix} \mathbf{W}_{12}^{PV} \mathbf{x}_i + \mathbf{W}_{13}^{PV} \mathbf{x}_{i-1} \\ \times \\ \times \end{pmatrix} \cdot \begin{pmatrix} \times \\ \mathbf{W}_{22}^{KQ\top} \mathbf{x}_i + \mathbf{W}_{32}^{KQ\top} \mathbf{x}_{i-1} \\ \mathbf{W}_{23}^{KQ\top} \mathbf{x}_i + \mathbf{W}_{33}^{KQ\top} \mathbf{x}_{i-1} \end{pmatrix}^* \right) \begin{pmatrix} \mathbf{0}_d \\ \mathbf{x}_t \\ \mathbf{x}_{t-1} \end{pmatrix},
\end{aligned}$$

where \times s are the elements that will not contribute to the final $\hat{\mathbf{y}}_t$. A further simple computation shows that

$$\begin{aligned}
\hat{\mathbf{y}}_t &= \mathbf{0}_d + \frac{1}{\rho_t} \sum_{i=1}^t (\mathbf{W}_{12}^{PV} \mathbf{x}_i + \mathbf{W}_{13}^{PV} \mathbf{x}_{i-1}) (\mathbf{W}_{22}^{KQ\top} \mathbf{x}_i + \mathbf{W}_{32}^{KQ\top} \mathbf{x}_{i-1})^* \mathbf{x}_t \\
&\quad + \frac{1}{\rho_t} \sum_{i=1}^t (\mathbf{W}_{12}^{PV} \mathbf{x}_i + \mathbf{W}_{13}^{PV} \mathbf{x}_{i-1}) (\mathbf{W}_{23}^{KQ\top} \mathbf{x}_i + \mathbf{W}_{33}^{KQ\top} \mathbf{x}_{i-1})^* \mathbf{x}_{t-1} \\
&= \frac{1}{\rho_t} \sum_{i=1}^t \left((\mathbf{W}_{12}^{PV} \mathbf{x}_i + \mathbf{W}_{13}^{PV} \mathbf{x}_{i-1}) \begin{pmatrix} \mathbf{W}_{22}^{KQ\top} \mathbf{x}_i + \mathbf{W}_{32}^{KQ\top} \mathbf{x}_{i-1} \\ \mathbf{W}_{23}^{KQ\top} \mathbf{x}_i + \mathbf{W}_{33}^{KQ\top} \mathbf{x}_{i-1} \end{pmatrix}^* \right) \begin{pmatrix} \mathbf{x}_t \\ \mathbf{x}_{t-1} \end{pmatrix} \\
&= \frac{1}{\rho_t} \sum_{i=1}^t \left((\mathbf{W}_{12}^{PV} \quad \mathbf{W}_{13}^{PV}) \begin{pmatrix} \mathbf{x}_i \\ \mathbf{x}_{i-1} \end{pmatrix} \cdot \left(\begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix}^\top \begin{pmatrix} \mathbf{x}_i \\ \mathbf{x}_{i-1} \end{pmatrix} \right)^* \right) \begin{pmatrix} \mathbf{x}_t \\ \mathbf{x}_{t-1} \end{pmatrix} \\
&= \frac{1}{\rho_t} \sum_{i=1}^t \left((\mathbf{W}_{12}^{PV} \quad \mathbf{W}_{13}^{PV}) \mathbf{e}_i^{\mathbf{x}} \cdot \left(\begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix}^\top \mathbf{e}_i^{\mathbf{x}} \right)^* \right) \mathbf{e}_t^{\mathbf{x}} \\
&= \frac{1}{\rho_t} (\mathbf{W}_{12}^{PV} \quad \mathbf{W}_{13}^{PV}) \left(\sum_{i=1}^t \mathbf{e}_i^{\mathbf{x}} \mathbf{e}_i^{\mathbf{x}*} \right) \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix} \mathbf{e}_t^{\mathbf{x}} \\
&= (\mathbf{W}_{12}^{PV} \quad \mathbf{W}_{13}^{PV}) \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix} \mathbf{e}_t^{\mathbf{x}} \in \mathbb{C}^d,
\end{aligned}$$

which completes the proof. \square

A.2 Proof of Theorem 4.1

A.2.1 Proof of Lemma 5.1

For the reader's convenience, we restate the lemma as the following.

Lemma A.1. Each element of the network's prediction $\hat{y}_{t,j}$ ($j \in [d]$) can be expressed as the following.

$$\hat{y}_{t,j} = \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) = \text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}_t^{\mathbf{x}}} \mathbf{E}_t^{\mathbf{x}\top}}{\rho_t} \cdot \mathbf{B}_j,$$

where the \mathbf{A} and \mathbf{B}_j are defined as

$$\mathbf{A} = (\mathbf{a}_1 \quad \dots \quad \mathbf{a}_{2d}) = \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix}, \quad \mathbf{B}_j = \begin{pmatrix} \mathbf{b}_{j1} \\ \mathbf{b}_{j2} \end{pmatrix} = \begin{pmatrix} \mathbf{W}_{12,j}^{PV\top} \\ \mathbf{W}_{13,j}^{PV\top} \end{pmatrix},$$

with $\mathbf{a}_i \in \mathbb{R}^{2d}$ and $\mathbf{b}_{j1}, \mathbf{b}_{j2} \in \mathbb{R}^d$.

Proof. Based on the result in Eq. 1, we can write

$$\begin{aligned} \hat{y}_{t,j} &= (\mathbf{W}_{12}^{PV} \quad \mathbf{W}_{13}^{PV})_j \cdot \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix} \mathbf{e}_t^{\mathbf{x}} \\ &= \mathbf{B}_j^\top \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} (\mathbf{a}_1 \quad \dots \quad \mathbf{a}_{2d}) \mathbf{e}_t^{\mathbf{x}} \\ &= \sum_{i=1}^{2d} e_{t,i}^{\mathbf{x}} \mathbf{B}_j^\top \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{a}_i \\ &= \sum_{i=1}^{2d} e_{t,i}^{\mathbf{x}} \text{tr} \left(\mathbf{B}_j^\top \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{a}_i \right) \\ &= \sum_{i=1}^{2d} e_{t,i}^{\mathbf{x}} \text{tr} \left(\mathbf{a}_i \mathbf{B}_j^\top \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \right) \\ &= \sum_{i=1}^{2d} \text{tr} \left(\mathbf{a}_i \mathbf{B}_j^\top \cdot e_{t,i}^{\mathbf{x}} \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \right) \\ &= \text{tr} \left[\begin{pmatrix} \mathbf{a}_1 \mathbf{B}_j^\top \\ \vdots \\ \mathbf{a}_{2d} \mathbf{B}_j^\top \end{pmatrix} \cdot \left(e_{t,1}^{\mathbf{x}} \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \quad \dots \quad e_{t,2d}^{\mathbf{x}} \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \right) \right] \\ &= \text{tr} \left(\text{Vec}(\mathbf{A}) \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \right) \\ &= \text{tr} \left(\mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \right) = \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \\ &= \text{tr} \left(\text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}_t^{\mathbf{x}}} \mathbf{E}_t^{\mathbf{x}\top}}{\rho_t} \cdot \mathbf{B}_j \right) = \text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}_t^{\mathbf{x}}} \mathbf{E}_t^{\mathbf{x}\top}}{\rho_t} \cdot \mathbf{B}_j, \end{aligned}$$

which finishes the proof. \square

A.2.2 Proof of Lemma 5.2

For the reader's convenience, we restate the lemma as the following.

Lemma A.2. Suppose that Assumption 4.1 holds, then the dynamical process of the parameters in the diagonal of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} satisfies

$$\begin{aligned} \frac{d}{d\tau} a &= -ab^2 \left[(T-2)\kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] + b(T-2)\kappa_1, \\ \frac{d}{d\tau} b &= -a^2b \left[(T-2)\kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] + a(T-2)\kappa_1, \end{aligned}$$

while the gradients for all other parameters were kept at zero during the training process.

Proof. The population loss $L(\boldsymbol{\theta})$ in Eq. 2 can be rewritten as

$$\begin{aligned} L(\boldsymbol{\theta}) &= \sum_{t=2}^{T-1} L_t(\boldsymbol{\theta}) = \sum_{t=2}^{T-1} \mathbb{E} \left[\frac{1}{2} \sum_{j=1}^d |\hat{y}_{t,j} - x_{t+1,j}|^2 \right] = \sum_{t=2}^{T-1} \sum_{j=1}^d \mathbb{E} \left[\frac{1}{2} |\hat{y}_{t,j} - x_{t+1,j}|^2 \right] \\ &= \sum_{t=2}^{T-1} \sum_{j=1}^d \mathbb{E} \left[\frac{1}{2} \left(\hat{y}_{t,j}^* \hat{y}_{t,j} - 2\text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) + x_{t+1,j}^* x_{t+1,j} \right) \right] \\ &= \sum_{t=2}^{T-1} \sum_{j=1}^d \mathbb{E} \left[\frac{1}{2} \hat{y}_{t,j}^* \hat{y}_{t,j} - \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) + \frac{1}{2} x_{t+1,j}^* x_{t+1,j} \right]. \end{aligned}$$

Then, we can calculate the derivatives of $L_t(\boldsymbol{\theta})$ with respect to \mathbf{B}_j and $\text{Vec}(\mathbf{A})$ as

$$\begin{aligned} \nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta}) &= \sum_{j=1}^d \mathbb{E} \left[\frac{1}{2} \nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} - \nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right] \\ &= \mathbb{E} \left[\frac{1}{2} \nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} - \nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right] \\ &= \frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right] - \mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right], \end{aligned}$$

and

$$\begin{aligned} \nabla_{\text{Vec}(\mathbf{A})} L_t(\boldsymbol{\theta}) &= \sum_{j=1}^d \mathbb{E} \left[\frac{1}{2} \nabla_{\text{Vec}(\mathbf{A})} \hat{y}_{t,j}^* \hat{y}_{t,j} - \nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right] \\ &= \frac{1}{2} \sum_{j=1}^d \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \hat{y}_{t,j}^* \hat{y}_{t,j} \right] - \sum_{j=1}^d \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]. \end{aligned}$$

Step one: calculate $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$. Based on Lemma 5.1, we have

$$\hat{y}_{t,j} = \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}).$$

Then, the $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$ can be derived as the following.

$$\begin{aligned} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right] &= \mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \right) \right] \\ &= \mathbb{E} \left[\nabla_{\mathbf{B}_j} \mathbf{B}_j^\top \cdot \text{Re} \left(x_{t+1,j}^* \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \right) \right] \\ &= \mathbb{E} \left[\text{Re} \left(x_{t+1,j}^* \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \right) \right] \\ &= \text{Re} \left(\mathbb{E} \left[x_{t+1,j}^* \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \right] \right) \\ &= \text{Re} \left(\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \text{Vec} \left(\frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{A} \mathbf{e}_t^{\mathbf{x}} \right) \right] \right) \quad (\text{use generating process}) \\ &= \text{Re} \left(\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{A} \mathbf{e}_t^{\mathbf{x}} \right) \right] \right), \end{aligned}$$

where the penultimate equality uses the property of Kronecker and Vec operator $\text{Vec}(\mathbf{A}\mathbf{X}\mathbf{B}) = (\mathbf{B}^\top \otimes \mathbf{A})\text{Vec}(\mathbf{X})$, we refer Section 10.2 in [62] for details.

For $\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t}$, we can simplify it as

$$\begin{aligned} \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} &= \frac{1}{\rho_t} \sum_{i=1}^t \mathbf{e}_i^x \mathbf{e}_i^{x*} \\ &= \frac{1}{\rho_t} \sum_{i=1}^t \begin{pmatrix} \mathbf{x}_i \mathbf{x}_i^* & \mathbf{x}_i \mathbf{x}_{i-1}^* \\ \mathbf{x}_{i-1} \mathbf{x}_i^* & \mathbf{x}_{i-1} \mathbf{x}_{i-1}^* \end{pmatrix} \\ &= \frac{1}{\rho_t} \begin{pmatrix} \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_i^* & \mathbf{W} \sum_{i=1}^{t-1} \mathbf{x}_i \mathbf{x}_i^* \\ \sum_{i=1}^{t-1} \mathbf{x}_i \mathbf{x}_i^* \mathbf{W}^* & \sum_{i=1}^{t-1} \mathbf{x}_i \mathbf{x}_i^* \end{pmatrix}. \end{aligned}$$

Based on the diagonal property of \mathbf{W} , we can simplify the $\mathbf{x}_i \mathbf{x}_i^*$ as the following.

$$\mathbf{x}_i \mathbf{x}_i^* = (\mathbf{W}^{i-1} \mathbf{x}_1)(\mathbf{W}^{i-1} \mathbf{x}_1)^* = \mathbf{M}_i \odot \widehat{\Sigma},$$

where we define $\widehat{\Sigma} = \mathbf{x}_1 \mathbf{x}_1^*$ and $\mathbf{M}_i = \lambda^{i-1} \lambda^{i-1*}$. Therefore, we have

$$\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} = \frac{1}{\rho_t} \begin{pmatrix} \sum_{i=1}^t \mathbf{M}_i \odot \widehat{\Sigma} & \mathbf{W} \sum_{i=1}^{t-1} \mathbf{M}_i \odot \widehat{\Sigma} \\ \sum_{i=1}^{t-1} \mathbf{M}_i \odot \widehat{\Sigma} \mathbf{W}^* & \sum_{i=1}^{t-1} \mathbf{M}_i \odot \widehat{\Sigma} \end{pmatrix}. \quad (4)$$

Then, leveraging the sparse property of \mathbf{A} , we can derive $\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x$ as follows.

$$\begin{aligned} \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x &= \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \begin{pmatrix} \mathbf{0}_d \\ a \mathbf{x}_t \end{pmatrix} \\ &= \frac{a}{\rho_t} \begin{pmatrix} \mathbf{W} (\sum_{i=1}^{t-1} \mathbf{M}_i \odot \widehat{\Sigma}) \mathbf{x}_t \\ (\sum_{i=1}^{t-1} \mathbf{M}_i \odot \widehat{\Sigma}) \mathbf{x}_t \end{pmatrix}. \end{aligned}$$

Therefore, for any $l \in [d]$, we have

$$\begin{aligned} \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l &= \frac{a}{\rho_t} \lambda_l \sum_{i=1}^{t-1} (\mathbf{M}_i \odot \widehat{\Sigma})_{l,:} \mathbf{x}_t \\ &= \frac{a}{\rho_t} \lambda_l \sum_{i=1}^{t-1} \left(\lambda_l^{i-1} \lambda_1^{1-i} x_{1l} x_{11} \quad \cdots \quad \lambda_l^{i-1} \lambda_d^{1-i} x_{1l} x_{1d} \right) \begin{pmatrix} \lambda_1^{t-1} x_{11} \\ \cdots \\ \lambda_d^{t-1} x_{1d} \end{pmatrix} \\ &= \frac{a}{\rho_t} \lambda_l \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_l^{i-1} \lambda_r^{t-i} x_{1l} x_{1r}^2 \\ &= \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_l^i \lambda_r^{t-i} x_{1l} x_{1r}^2. \end{aligned}$$

Similarly, for any $l \in [2d] - [d]$, we have

$$\left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l = \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_l^{i-1} \lambda_r^{t-i} x_{1l} x_{1r}^2.$$

To sum up, for any $l \in [2d]$, we have

$$\left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l = \begin{cases} \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_l^i \lambda_r^{t-i} x_{1l} x_{1r}^2, & l \in [d], \\ \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_{l-d}^{i-1} \lambda_r^{t-i} x_{1,l-d} x_{1r}^2, & l \in [2d] - [d]. \end{cases}$$

Next, we calculate the $\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right]$. For any $l \in [d]$, we have

$$\begin{aligned} \mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] &= \mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_i^i \lambda_r^{t-i} x_{1i} x_{1r}^2 \right] \\ &= \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \mathbb{E}[\lambda_j^{-t} \lambda_i^i \lambda_r^{t-i}] \mathbb{E}[x_{1j} x_{1i} x_{1r}^2]. \end{aligned}$$

We discuss them in the following categories,

1. $l \neq j$. In this case, $\mathbb{E}[x_{1j} x_{1i} x_{1r}^2] = 0$ by Assumption 4.1, thus

$$\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] = \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \mathbb{E}[\lambda_j^{-t} \lambda_i^i \lambda_r^{t-i}] 0 = 0.$$

2. $l = j, r = j$. Because λ_i s are i.i.d. and $\mathbb{E}[\lambda_i^k] = \delta(k=0)$, we have

$$\begin{aligned} \mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] &= \frac{a}{\rho_t} \sum_{r=1}^d \mathbb{E}[\lambda_j^{-t} \lambda_j^i \lambda_j^{t-i}] \mathbb{E}[x_{1j} x_{1j} x_{1j}^2] \\ &= \frac{a}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4]. \end{aligned}$$

Similarly, for any $l \in [2d] - [d]$, we have

$$\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] = 0.$$

Therefore, for any $l \in [2d]$, we have

$$\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right) \right] = \begin{cases} \frac{a}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4], & l = j, \\ 0, & l \neq j, \end{cases}$$

and the l -th element of $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$ is

$$\begin{aligned} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]_l &= \text{Re} \left(\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right) \right] \right)_l \\ &= \begin{cases} \frac{a}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4], & l = j, \\ 0, & l \neq j. \end{cases} \end{aligned}$$

Step two: calculate $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right]$. Based on Lemma 5.1, we have

$$\hat{y}_{t,j} = \text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j,$$

then we can simplify the $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right]$ as follows.

$$\begin{aligned} &\mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right] \\ &= \mathbb{E} \left[\nabla_{\mathbf{B}_j} \left(\text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right)^* \text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right] \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E} \left[\nabla_{\mathbf{B}_j} \mathbf{B}_j^\top \cdot \mathbf{e}_t^{x^*} \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \text{Vec}^\top(\mathbf{A}) \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right] \\
&= \mathbb{E} \left[\mathbf{e}_t^{x^*} \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \text{Vec}^\top(\mathbf{A}) \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right] \\
&\quad + \mathbb{E} \left[\mathbf{e}_t^{x^\top} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \text{Vec}^\top(\mathbf{A}) \overline{\mathbf{e}_t^x} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \cdot \mathbf{B}_j \right] \\
&= \mathbb{E} \left[2\text{Re} \left(\mathbf{e}_t^{x^\top} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \text{Vec}^\top(\mathbf{A}) \overline{\mathbf{e}_t^x} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \cdot \mathbf{B}_j \right) \right] \\
&= 2\text{Re} \left(\mathbb{E} \left[\mathbf{e}_t^{x^\top} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \text{Vec}^\top(\mathbf{A}) \overline{\mathbf{e}_t^x} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \cdot \mathbf{B}_j \right] \right).
\end{aligned}$$

We further derive that

$$\begin{aligned}
&\mathbb{E} \left[\mathbf{e}_t^{x^\top} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \underbrace{\text{Vec}^\top(\mathbf{A}) \overline{\mathbf{e}_t^x} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t}}_{\in \mathbb{C}} \cdot \mathbf{B}_j \right] \\
&= \mathbb{E} \left[\text{Vec}^\top(\mathbf{A}) \overline{\mathbf{e}_t^x} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \cdot \mathbf{B}_j \cdot \mathbf{e}_t^{x^\top} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \text{Vec}(\mathbf{A}) \right] \\
&= \mathbb{E} \left[\mathbf{B}_j^\top \cdot \mathbf{e}_t^{x^*} \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \mathbf{e}_t^{x^\top} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \text{Vec}(\mathbf{A}) \right] \\
&= \mathbb{E} \left[\mathbf{B}_j^\top \cdot \text{Vec} \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{A} \overline{\mathbf{e}_t^x} \right) \cdot \text{Vec} \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right) \right] \\
&= \mathbb{E} \left[\mathbf{B}_j^\top \cdot \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{A} \overline{\mathbf{e}_t^x} \cdot \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right] \\
&= \mathbb{E} \left[\sum_{k=1}^{2d} B_{jk} \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{A} \overline{\mathbf{e}_t^x} \right)_k \cdot \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right] \\
&= \mathbb{E} \left[b \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{A} \overline{\mathbf{e}_t^x} \right)_j \cdot \frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right]. \quad (\text{sparsity of } \mathbf{B})
\end{aligned}$$

For any and $l \in [2d]$, recall that

$$\left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l = \begin{cases} \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_i^i \lambda_r^{t-i} x_{1l} x_{1r}^2, & l \in [d], \\ \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_{l-d}^{i-1} \lambda_r^{t-i} x_{1,l-d} x_{1r}^2, & l \in [2d] - [d], \end{cases}$$

and for any $j \in [d]$, we have

$$\left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{A} \overline{\mathbf{e}_t^x} \right)_j = \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_j = \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_j^{-i} \lambda_r^{i-t} x_{1j} x_{1r}^2.$$

For any $l \in [d]$, with careful computing, we have

$$\begin{aligned}
&\mathbb{E} \left[b \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{A} \overline{\mathbf{e}_t^x} \right)_j \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] \\
&= \frac{a^2 b}{\rho_t^2} \sum_{i_1=1}^{t-1} \sum_{i_2=1}^{t-1} \sum_{r_1=1}^d \sum_{r_2=1}^d \mathbb{E} [\lambda_j^{-i_1} \lambda_l^{i_2} \lambda_{r_1}^{i_1-t} \lambda_{r_2}^{t-i_2}] \mathbb{E} [x_{1j} x_{i_1} x_{1r_1}^2 x_{1r_2}^2].
\end{aligned}$$

We discuss it in the following categories,

1. $l \neq j$. In this case, $\mathbb{E}[x_{1j}x_{il}x_{1r_1}^2x_{1r_2}^2] = 0$ by Assumption 4.1, thus it becomes 0.
2. $l = j, r_1 = r_2 = j$. It becomes

$$\frac{a^2b}{\rho_t^2} \sum_{i_1=1}^{t-1} \sum_{i_2=1}^{t-1} \mathbb{E}[\lambda_j^{-i_1} \lambda_j^{i_2} \lambda_j^{i_1-t} \lambda_j^{t-i_2}] \mathbb{E}[x_{1j}^6] = \frac{a^2b}{\rho_t^2} \sum_{i_1=1}^{t-1} \sum_{i_2=1}^{t-1} \mathbb{E}[x_{1j}^6] = \frac{a^2b}{\rho_t^2} (t-1)^2 \mathbb{E}[x_{1j}^6].$$

3. $l = j, r_1 = r_2 = r \neq j$. It becomes

$$\begin{aligned} \frac{a^2b}{\rho_t^2} \sum_{i_1=1}^{t-1} \sum_{i_2=1}^{t-1} \sum_{r \neq j} \mathbb{E}[\lambda_j^{-i_1} \lambda_r^{i_2} \lambda_r^{i_1-t} \lambda_r^{t-i_2}] \mathbb{E}[x_{1j}^2 x_{1r}^4] &= \frac{a^2b}{\rho_t^2} \sum_{i_1=i_2} \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] \\ &= \frac{a^2b}{\rho_t^2} (t-1) \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4]. \end{aligned}$$

4. $l = j, r_1 \neq r_2$. In this case, $\mathbb{E}[\lambda_j^{-i_1} \lambda_l^{i_2} \lambda_{r_1}^{i_1-t} \lambda_{r_2}^{t-i_2}] = \mathbb{E}[\lambda_j^{-i_1} \lambda_{r_1}^{i_1-t} \lambda_{r_2}^{t-i_2}] = 0$, thus it becomes 0.

Similarly, for any $l \in [2d] - [d]$, we have

$$\mathbb{E} \left[b \left(\frac{\overline{\mathbf{E}_t^x} \mathbf{E}_t^{x\top}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_j \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] = 0.$$

To sum up, we have

$$\begin{aligned} &\mathbb{E} \left[b \left(\frac{\overline{\mathbf{E}_t^x} \mathbf{E}_t^{x\top}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_j \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] \\ &= \begin{cases} \frac{a^2b}{\rho_t^2} \left[(t-1)^2 \mathbb{E}[x_{1j}^6] + (t-1) \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] \right], & l = j, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

which implies that the l -th element of $\frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]$ is

$$\frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]_l = \begin{cases} \frac{a^2b}{\rho_t^2} \left[(t-1)^2 \mathbb{E}[x_{1j}^6] + (t-1) \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] \right], & l = j, \\ 0, & \text{otherwise.} \end{cases}$$

Step three: calculate $\nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})$ and $\nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})$. Based on steps one and two, the l -th element of $\nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})$ can be derived as follows.

$$\begin{aligned} \nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})_l &= \frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]_l - \mathbb{E} \left[\nabla_{\mathbf{B}_j} \operatorname{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right]_l \\ &= \begin{cases} \frac{a^2b}{\rho_t^2} \left[(t-1)^2 \mathbb{E}[x_{1j}^6] + (t-1) \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] \right] - \frac{a}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4], & l = j, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

Furthermore, the l -th element of $\nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})$ is

$$\begin{aligned} \nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})_l &= \sum_{t=2}^{T-1} \nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})_l \\ &= \begin{cases} \sum_{t=2}^{T-1} \left(\frac{a^2b}{\rho_t^2} \left[(t-1)^2 \mathbb{E}[x_{1j}^6] + (t-1) \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] \right] - \frac{a}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4] \right), & l = j, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

$$\begin{aligned}
&= \begin{cases} \sum_{t=2}^{T-1} \left(a^2 b \left[\mathbb{E}[x_{1j}^6] + \frac{1}{t-1} \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] \right] - a \mathbb{E}[x_{1j}^4] \right), & l = j, \\ 0, & \text{otherwise,} \end{cases} \\
&= \begin{cases} a^2 b \left[(T-2) \mathbb{E}[x_{1j}^6] + \sum_{t=2}^{T-1} \frac{1}{t-1} \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] \right] - a(T-2) \mathbb{E}[x_{1j}^4], & l = j, \\ 0, & \text{otherwise,} \end{cases} \\
&= \begin{cases} a^2 b \left[(T-2) \kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] - a(T-2) \kappa_1, & l = j, \\ 0, & \text{otherwise,} \end{cases}
\end{aligned}$$

where the last equality is from Assumption 4.1.

Step four: calculate $\mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right]$ and $\sum_{j=1}^d \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right]$.
Based on Lemma 5.1, we have

$$\widehat{y}_{t,j} = \text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j.$$

Then, the $\mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right]$ can be derived as the following.

$$\begin{aligned}
\mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right] &= \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \text{Vec}^\top(\mathbf{A}) \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right) \right] \\
&= \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Vec}^\top(\mathbf{A}) \cdot \text{Re} \left(x_{t+1,j}^* \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right) \right] \\
&= \mathbb{E} \left[\text{Re} \left(x_{t+1,j}^* \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right) \right] \\
&= \text{Re} \left(\mathbb{E} \left[x_{t+1,j}^* \cdot \mathbf{e}_t^x \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \cdot \mathbf{B}_j \right] \right) \\
&= \text{Re} \left(\mathbb{E} \left[\lambda_j^{-t} x_{1j} \cdot \text{Vec} \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x^\top} \right) \right] \right).
\end{aligned}$$

In terms of $\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x^\top}$, based on the sparsity of \mathbf{B} and Eq. 4, we can derive that

$$\begin{aligned}
\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x^\top} &= \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \right)_{:j} b \mathbf{e}_t^{x^\top} = \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x^*}}{\rho_t} \right)_{:j}^\top b \mathbf{e}_t^{x^\top} \\
&= b \left(\sum_{i=1}^t \mathbf{M}_i \odot \widehat{\Sigma} \quad \mathbf{W} \sum_{i=1}^{t-1} \mathbf{M}_i \odot \widehat{\Sigma} \right)_{:j}^\top \mathbf{e}_t^{x^\top}.
\end{aligned}$$

Then, for any $s, r \in [2d]$, we have

$$\left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x^\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x^\top} \right)_{sr} = \begin{cases} \frac{b}{\rho_t} \sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \lambda_r^{t-1} x_{1j} x_{1s} x_{1r}, & s \in [d], r \in [d], \\ \frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^i \lambda_{s-d}^{1-i} \lambda_r^{t-1} x_{1j} x_{1,s-d} x_{1r}, & s \in [2d] - [d], r \in [d], \\ \frac{b}{\rho_t} \sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \lambda_{r-d}^{t-2} x_{1j} x_{1s} x_{1,r-d}, & s \in [d], r \in [2d] - [d], \\ \frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^i \lambda_{s-d}^{1-i} \lambda_{r-d}^{t-2} x_{1j} x_{1,s-d} x_{1,r-d}, & s \in [2d] - [d], r \in [2d] - [d]. \end{cases}$$

Next, we calculate $\mathbb{E}\left[\lambda_j^{-t}x_{1j} \cdot \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x\top}\right]$. For any $s \in [2d] - [d], r \in [d]$, we have

$$\begin{aligned} \mathbb{E}\left[\lambda_j^{-t}x_{1j} \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x\top}\right)_{sr}\right] &= \mathbb{E}\left[\frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^{i-t} \lambda_{s-d}^{1-i} \lambda_r^{t-1} x_{1j}^2 x_{1,s-d} x_{1r}\right] \\ &= \frac{b}{\rho_t} \sum_{i=1}^{t-1} \mathbb{E}[\lambda_j^{i-t} \lambda_{s-d}^{1-i} \lambda_r^{t-1}] \mathbb{E}[x_{1j}^2 x_{1,s-d} x_{1r}]. \end{aligned}$$

We discuss it in the following categories,

1. $s - d \neq r$. In this case, $\mathbb{E}[x_{1j}^2 x_{1,s-d} x_{1r}] = 0$ by Assumption 4.1, thus it becomes 0.

2. $s - d = r = j$. It becomes

$$\begin{aligned} \frac{b}{\rho_t} \sum_{i=1}^{t-1} \mathbb{E}[\lambda_j^{i-t} \lambda_{s-d}^{1-i} \lambda_r^{t-1}] \mathbb{E}[x_{1j}^2 x_{1,s-d} x_{1r}] &= \frac{b}{\rho_t} \sum_{i=1}^{t-1} \mathbb{E}[\lambda_j^{i-t} \lambda_j^{1-i} \lambda_j^{t-1}] \mathbb{E}[x_{1j}^4] \\ &= \frac{b}{\rho_t} \sum_{i=1}^{t-1} \mathbb{E}[x_{1j}^4] = \frac{b}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4]. \end{aligned}$$

3. $s - d = r \neq j$. In this case, $\mathbb{E}[\lambda_j^{i-t} \lambda_{s-d}^{1-i} \lambda_r^{t-1}] = \mathbb{E}[\lambda_j^{i-t} \lambda_{s-d}^{t-i}] = 0$, thus it becomes 0.

Similarly, for any other s, r , we can calculate that

$$\mathbb{E}\left[\lambda_j^{-t} \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x\top}\right)_{sr}\right] = 0.$$

To sum up, we have

$$\mathbb{E}\left[\lambda_j^{-t} \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x\top}\right)_{sr}\right] = \begin{cases} \frac{b}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4], & s = d + j, r = j, \\ 0, & \text{otherwise,} \end{cases}$$

thus the (s, r) -th element of $\mathbb{E}\left[\nabla_{\mathbf{A}} \text{Re}\left(x_{t+1,j}^* \hat{y}_{t,j}\right)\right]$ is

$$\begin{aligned} \mathbb{E}\left[\nabla_{\mathbf{A}} \text{Re}\left(x_{t+1,j}^* \hat{y}_{t,j}\right)\right]_{sr} &= \begin{cases} \frac{b}{\rho_t} (t-1) \mathbb{E}[x_{1j}^4], & s = d + j, r = j, \\ 0, & \text{otherwise,} \end{cases} \\ &= \begin{cases} \frac{b}{\rho_t} (t-1) \kappa_1, & s = d + j, r = j, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Finally, we can calculate $\sum_{j=1}^d \mathbb{E}\left[\nabla_{\mathbf{A}} \text{Re}\left(x_{t+1,j}^* \hat{y}_{t,j}\right)\right]$ as

$$\sum_{j=1}^d \mathbb{E}\left[\nabla_{\mathbf{A}} \text{Re}\left(x_{t+1,j}^* \hat{y}_{t,j}\right)\right]_{sr} = \begin{cases} \frac{b}{\rho_t} (t-1) \kappa_1, & s - d = r, \\ 0, & \text{otherwise.} \end{cases}$$

Step five: calculate $\mathbb{E}\left[\nabla_{\text{Vec}(\mathbf{A})} \hat{y}_{t,j}^* \hat{y}_{t,j}\right]$ and $\sum_{j=1}^d \mathbb{E}\left[\nabla_{\text{Vec}(\mathbf{A})} \hat{y}_{t,j}^* \hat{y}_{t,j}\right]$. Based on Lemma 5.1, we have

$$\hat{y}_{t,j} = \mathbf{B}_j^\top \cdot \mathbf{e}_t^{x\top} \otimes \frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x*}}}{\rho_t} \cdot \text{Vec}(\mathbf{A}),$$

then we can simplify the $\mathbb{E}\left[\nabla_{\text{Vec}(\mathbf{A})}\widehat{y}_{t,j}^*\widehat{y}_{t,j}\right]$ as follows

$$\begin{aligned}
& \mathbb{E}\left[\nabla_{\text{Vec}(\mathbf{A})}\widehat{y}_{t,j}^*\widehat{y}_{t,j}\right] \\
&= \mathbb{E}\left[\nabla_{\text{Vec}(\mathbf{A})}\left(\mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A})\right)^* \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A})\right] \\
&= \mathbb{E}\left[\nabla_{\text{Vec}(\mathbf{A})}\text{Vec}(\mathbf{A})^\top \cdot \overline{\mathbf{e}}_t^{\mathbf{x}} \otimes \frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \cdot \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A})\right] \\
&= \mathbb{E}\left[\overline{\mathbf{e}}_t^{\mathbf{x}} \otimes \frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \cdot \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A})\right] \\
&\quad + \mathbb{E}\left[\mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \cdot \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}*} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \cdot \text{Vec}(\mathbf{A})\right] \\
&= \mathbb{E}\left[2\text{Re}\left(\mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \cdot \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}*} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \cdot \text{Vec}(\mathbf{A})\right)\right] \\
&= 2\text{Re}\left(\mathbb{E}\left[\mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \cdot \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}*} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \cdot \text{Vec}(\mathbf{A})\right]\right).
\end{aligned}$$

We further derive that

$$\begin{aligned}
& \mathbb{E}\left[\mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \cdot \mathbf{B}_j^\top \cdot \underbrace{\mathbf{e}_t^{\mathbf{x}*} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t}}_{\in \mathbb{C}} \cdot \text{Vec}(\mathbf{A})\right] \\
&= \mathbb{E}\left[\mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}*} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \cdot \mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j\right] \\
&= \mathbb{E}\left[\text{Vec}(\mathbf{A})^\top \cdot \overline{\mathbf{e}}_t^{\mathbf{x}} \otimes \frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \cdot \mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j\right] \\
&= \mathbb{E}\left[\text{Vec}(\mathbf{A})^\top \cdot \text{Vec}\left(\frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}*}\right) \cdot \text{Vec}\left(\frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top}\right)\right] \\
&= \mathbb{E}\left[\sum_{k,l=1}^{2d} A_{kl} \left(\frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}*}\right)_{kl} \cdot \text{Vec}\left(\frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top}\right)\right] \\
&= \mathbb{E}\left[\sum_{k=d+1}^{2d} a \left(\frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}*}\right)_{k,k-d} \cdot \text{Vec}\left(\frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top}\right)\right]. \quad (\text{sparsity of } \mathbf{A})
\end{aligned}$$

Recall that for any $s, r \in [2d]$, we have

$$\left(\frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top}\right)_{sr} = \begin{cases} \frac{b}{\rho_t} \sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \lambda_r^{t-1} x_{1j} x_{1s} x_{1r}, & s \in [d], r \in [d], \\ \frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^i \lambda_s^{1-i} \lambda_r^{t-1} x_{1j} x_{1,s-d} x_{1r}, & s \in [2d] - [d], r \in [d], \\ \frac{b}{\rho_t} \sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \lambda_r^{t-2} x_{1j} x_{1s} x_{1,r-d}, & s \in [d], r \in [2d] - [d], \\ \frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^i \lambda_s^{1-i} \lambda_r^{t-2} x_{1j} x_{1,s-d} x_{1,r-d}, & s \in [2d] - [d], r \in [2d] - [d]. \end{cases}$$

Furthermore, for any $k \in [2d] - [d]$, we can calculate that

$$\left(\frac{\mathbf{E}_t^{\mathbf{x}}\mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}*}\right)_{k,k-d} = \left(\frac{\overline{\mathbf{E}}_t^{\mathbf{x}}\overline{\mathbf{E}}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top}\right)_{k,k-d}$$

$$= \frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^{-i} \lambda_{k-d}^{i-1} \lambda_{k-d}^{1-t} x_{1j} x_{1,k-d}^2 = \frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^{-i} \lambda_{k-d}^{i-t} x_{1j} x_{1,k-d}^2,$$

With careful computing, for any $s \in [2d] - [d]$, $r \in [d]$, we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=d+1}^{2d} a \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x*} \right)_{k,k-d} \cdot \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x\top} \right)_{sr} \right] \\ &= \frac{ab^2}{\rho_t^2} \sum_{k=d+1}^{2d} \sum_{i_1=1}^{t-1} \sum_{i_2=1}^{t-1} \mathbb{E}[\lambda_j^{i_2-i_1} \lambda_{k-d}^{i_1-t} \lambda_{s-d}^{1-i_2} \lambda_r^{t-1}] \mathbb{E}[x_{1j}^2 x_{1,k-d} x_{1,s-d} x_{1r}]. \end{aligned}$$

We discuss it in the following categories,

1. $s - d \neq r$. In this case, $\mathbb{E}[x_{1j}^2 x_{1,k-d} x_{1,s-d} x_{1r}] = 0$ by Assumption 4.1, thus it becomes 0.
2. $s - d = r = k - d = j$. It becomes

$$\frac{ab^2}{\rho_t^2} \sum_{i_1=1}^{t-1} \sum_{i_2=1}^{t-1} \mathbb{E}[\lambda_j^0] \mathbb{E}[x_{1j}^6] = \frac{ab^2}{\rho_t^2} (t-1)^2 \mathbb{E}[x_{1j}^6].$$

3. $s - d = r = k - d \neq j$. It becomes

$$\begin{aligned} \frac{ab^2}{\rho_t^2} \sum_{i_1=1}^{t-1} \sum_{i_2=1}^{t-1} \mathbb{E}[\lambda_j^{i_2-i_1} \lambda_r^{i_1-i_2}] \mathbb{E}[x_{1j}^2 x_{1,r}^4] &= \frac{ab^2}{\rho_t^2} \sum_{i_1=1}^{t-1} \mathbb{E}[\lambda_j^{i_1-i_1} \lambda_r^{i_1-i_1}] \mathbb{E}[x_{1j}^2 x_{1,r}^4] \\ &= \frac{ab^2}{\rho_t^2} (t-1) \mathbb{E}[x_{1j}^2 x_{1,r}^4]. \end{aligned}$$

4. $s - d = r \neq k - d$. In these case, $\mathbb{E}[\lambda_j^{i_2-i_1} \lambda_{k-d}^{i_1-t} \lambda_{s-d}^{1-i_2} \lambda_r^{t-1}] = \mathbb{E}[\lambda_j^{i_2-i_1} \lambda_{k-d}^{i_1-t} \lambda_r^{t-i_2}] = 0$, thus it becomes 0.

Similarly, for any other s, r , we can calculate that

$$\mathbb{E} \left[\sum_{k=d+1}^{2d} a \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x*} \right)_{k,k-d} \cdot \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x\top} \right)_{sr} \right] = 0.$$

To sum up, we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{k=d+1}^{2d} a \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x*} \right)_{k,k-d} \cdot \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{x\top} \right)_{sr} \right] \\ &= \begin{cases} \frac{ab^2}{\rho_t^2} (t-1)^2 \mathbb{E}[x_{1j}^6], & s = d+r, r = j, \\ \frac{ab^2}{\rho_t^2} (t-1) \mathbb{E}[x_{1j}^2 x_{1r}^4], & s = d+r, r \neq j, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

which implies that the (s, r) -th element of $\frac{1}{2} \mathbb{E}[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j}]$ is

$$\begin{aligned} \frac{1}{2} \mathbb{E}[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j}]_{sr} &= \begin{cases} \frac{ab^2}{\rho_t^2} (t-1)^2 \mathbb{E}[x_{1j}^6], & s = d+r, r = j, \\ \frac{ab^2}{\rho_t^2} (t-1) \mathbb{E}[x_{1j}^2 x_{1r}^4], & s = d+r, r \neq j, \\ 0, & \text{otherwise,} \end{cases} \\ &= \begin{cases} \frac{ab^2}{\rho_t^2} (t-1)^2 \kappa_2, & s = d+r, r = j, \\ \frac{ab^2}{\rho_t^2} (t-1) \mathbb{E}[x_{1j}^2 x_{1r}^4], & s = d+r, r \neq j, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

Based on these results, we can derive

$$\frac{1}{2} \sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]_{sr} = \begin{cases} \frac{ab^2}{\rho_t^2} [(t-1)^2 \kappa_2 + (t-1) \kappa_3], & s-d=r, \\ 0, & \text{otherwise.} \end{cases}$$

Step six: calculate $\nabla_{\text{vec}(\mathbf{A})} L_t(\boldsymbol{\theta})$ and $\nabla_{\text{vec}(\mathbf{A})} L(\boldsymbol{\theta})$. Based on steps four and five, the (s, r) -th element of $\nabla_{\mathbf{A}} L_t(\boldsymbol{\theta})$ can be derived as follows.

$$\begin{aligned} \nabla_{\mathbf{A}} L_t(\boldsymbol{\theta})_{sr} &= \frac{1}{2} \sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right] - \sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right] \\ &= \begin{cases} \frac{ab^2}{\rho_t^2} [(t-1)^2 \kappa_2 + (t-1) \kappa_3] - \frac{b}{\rho_t} (t-1) \kappa_1, & s-d=r, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Furthermore, the (s, r) -th element of $\nabla_{\mathbf{A}} L(\boldsymbol{\theta})$ is

$$\begin{aligned} \nabla_{\mathbf{A}} L(\boldsymbol{\theta})_{sr} &= \sum_{t=2}^{T-1} \nabla_{\mathbf{A}} L_t(\boldsymbol{\theta})_{sr} \\ &= \begin{cases} \sum_{t=2}^{T-1} \left(\frac{ab^2}{\rho_t^2} [(t-1)^2 \kappa_2 + (t-1) \kappa_3] - \frac{b}{\rho_t} (t-1) \kappa_1 \right), & s-d=r, \\ 0, & \text{otherwise,} \end{cases} \\ &= \begin{cases} \sum_{t=2}^{T-1} \left(ab^2 \left[\kappa_2 + \frac{1}{t-1} \kappa_3 \right] - b \kappa_1 \right), & s-d=r, \\ 0, & \text{otherwise,} \end{cases} \\ &= \begin{cases} ab^2 \left[(T-2) \kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] - b(T-2) \kappa_1, & s-d=r, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Step seven: summarize the result by induction. From the gradient of $\nabla_{\mathbf{A}} L(\boldsymbol{\theta})$ and $\nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})$, we observe that non-zero gradients only emerge in the diagonal of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} , and they are same. Therefore, the parameter matrices keep the same structure as the initial time. Thus we can summarize the dynamic system as the following.

$$\begin{aligned} \frac{d}{d\tau} a &= -ab^2 \left[(T-2) \kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] + b(T-2) \kappa_1, \\ \frac{d}{d\tau} b &= -a^2 b \left[(T-2) \kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3 \right] + a(T-2) \kappa_1. \end{aligned}$$

which completes the proof. \square

A.2.3 Proof of Lemma 5.3

For the reader's convenience, we restate the lemma as the following.

Lemma A.3. *Suppose that Assumption 4.1 holds and denote $(T-2) \kappa_2 + \sum_{t=2}^{T-1} \frac{1}{t-1} \kappa_3$ and $(T-2) \kappa_1$ by c_1 and c_2 , respectively. Then, the dynamics in Lemma 5.2 are the same as those of gradient flow on the following objective function:*

$$\tilde{\ell}(a, b) = \frac{1}{2c_1} (c_2 - c_1 ab)^2,$$

whose global minimums satisfy $ab = c_2/c_1$.

Proof. Basic calculus shows that

$$\begin{aligned}\frac{\partial}{\partial a} \tilde{\ell}(a, b) &= \frac{1}{2c_1} 2(c_2 - c_1 ab)(-c_1 b) = -b(c_2 - c_1 ab) = -\frac{d}{d\tau} a, \\ \frac{\partial}{\partial b} \tilde{\ell}(a, b) &= \frac{1}{2c_1} 2(c_2 - c_1 ab)(-c_1 a) = -a(c_2 - c_1 ab) = -\frac{d}{d\tau} b.\end{aligned}$$

Therefore, the dynamics in Lemma 5.2 are the same as those of gradient flow on $\tilde{\ell}(a, b)$, whose global minimums satisfy $ab = c_2/c_1$. \square

A.2.4 Proof of Lemma 5.4

For the reader's convenience, we restate the lemma as the following.

Lemma A.4. *Suppose that Assumption 4.1 holds, then $\tilde{\ell}(a, b)$ is a non-convex function and satisfies the PL inequality as follows.*

$$\left| \frac{\partial}{\partial a} \tilde{\ell}(a, b) \right|^2 + \left| \frac{\partial}{\partial b} \tilde{\ell}(a, b) \right|^2 \geq 2c_1(a^2 + b^2) \left(\tilde{\ell}(a, b) - \min_{a,b} \tilde{\ell}(a, b) \right).$$

Therefore, the gradient flow in Lemma 5.2 converges to the global minimum of $\tilde{\ell}(a, b)$.

Proof. First, we prove that $\tilde{\ell}(a, b)$ is non-convex. The Hessian matrix of $\tilde{\ell}(a, b)$ can be derived as follows.

$$\nabla^2 \tilde{\ell}(a, b) = \begin{pmatrix} c_1 b^2 & 2c_1 ab - c_2 \\ 2c_1 ab - c_2 & c_1 a^2 \end{pmatrix}.$$

Its determinant $c_1 a^2 b^2 - (2c_1 ab - c_2)^2 = (c_2 - c_1 ab)(3c_1 ab - c_2) < 0$ when $ab < \frac{c_2}{3c_1}$ or $ab > \frac{c_2}{c_1}$. Thus, $\tilde{\ell}(a, b)$ is non-convex.

Besides, the PL inequality holds because

$$\begin{aligned}\left| \frac{\partial}{\partial a} \tilde{\ell}(a, b) \right|^2 + \left| \frac{\partial}{\partial b} \tilde{\ell}(a, b) \right|^2 &= b^2(c_2 - c_1 ab)^2 + a^2(c_2 - c_1 ab)^2 \\ &= (a^2 + b^2)(c_2 - c_1 ab)^2 \\ &= 2c_1(a^2 + b^2) \cdot \frac{1}{2c_1}(c_2 - c_1 ab)^2 \\ &= 2c_1(a^2 + b^2) \left(\tilde{\ell}(a, b) - \min_{a,b} \tilde{\ell}(a, b) \right) \\ &\geq 2c_1(a^2 + b^2) \left(\tilde{\ell}(a, b) - \min_{a,b} \tilde{\ell}(a, b) \right).\end{aligned}$$

\square

A.3 Proof of Corollary 4.1

For the reader's convenience, we restate the corollary as the following.

Corollary A.1. *We suppose that the same precondition of Theorem 4.1 holds. When predicting the t -th token, the trained transformer implements one step of gradient descent for the minimization of the OLS problem $L_{\text{OLS}}(\mathbf{W}) = \frac{1}{2} \sum_{i=1}^{t-1} \|\mathbf{x}_{t+1} - \mathbf{W} \mathbf{x}_t\|^2$, starting from the initialization $\mathbf{W} = \mathbf{0}_{d \times d}$ with a step size $\frac{\tilde{a}\tilde{b}}{t-1}$.*

Proof. The proof is stem from the theoretical construction in [16]. First, we simplify the prediction $\hat{\mathbf{y}}_t$ as follows.

$$\hat{\mathbf{y}}_t = \begin{pmatrix} \mathbf{W}_{12}^{PV} & \mathbf{W}_{13}^{PV} \end{pmatrix} \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix} \mathbf{e}_t^x$$

$$\begin{aligned}
&= \begin{pmatrix} \tilde{b}\mathbf{I}_d & \mathbf{0}_{d \times d} \end{pmatrix} \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}^*}}{\rho_t} \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \tilde{a}\mathbf{I}_d & \mathbf{0}_{d \times d} \end{pmatrix} \mathbf{e}_t^{\mathbf{x}} \\
&= \frac{1}{\rho_t} \begin{pmatrix} \tilde{b}\mathbf{I}_d & \mathbf{0}_{d \times d} \end{pmatrix} \sum_{i=1}^t \mathbf{e}_i^{\mathbf{x}} \mathbf{e}_i^{\mathbf{x}^*} \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \tilde{a}\mathbf{I}_d & \mathbf{0}_{d \times d} \end{pmatrix} \mathbf{e}_t^{\mathbf{x}} \\
&= \frac{1}{\rho_t} \sum_{i=1}^t \tilde{b}\mathbf{x}_i (\tilde{a}\mathbf{x}_{i-1}^* \mathbf{0}_{d \times d}) \mathbf{e}_t^{\mathbf{x}} \\
&= \frac{1}{\rho_t} \sum_{i=1}^t \tilde{b}\mathbf{x}_i \tilde{a}\mathbf{x}_{i-1}^* \mathbf{x}_t = \left(\frac{\tilde{a}\tilde{b}}{t-1} \sum_{i=1}^t \mathbf{x}_i \mathbf{x}_{i-1}^* \right) \mathbf{x}_t \\
&= \left(\frac{\tilde{a}\tilde{b}}{t-1} \sum_{i=1}^{t-1} \mathbf{x}_{i+1} \mathbf{x}_i^* \right) \mathbf{x}_t.
\end{aligned}$$

Then, we connect it to the one step of gradient descent for the OLS problem $L_{\text{OLS}}(\mathbf{W}) = \frac{1}{2} \sum_{i=1}^{t-1} \|\mathbf{x}_{i+1} - \mathbf{W}\mathbf{x}_i\|^2$.

$$\begin{aligned}
&\mathbf{W} - \frac{\tilde{a}\tilde{b}}{t-1} \nabla_{\mathbf{W}} \frac{1}{2} \sum_{i=1}^{t-1} \|\mathbf{x}_{i+1} - \mathbf{W}\mathbf{x}_i\|^2 \\
&= \mathbf{W} - \frac{\tilde{a}\tilde{b}}{t-1} \sum_{i=1}^{t-1} (\mathbf{x}_{i+1} - \mathbf{W}\mathbf{x}_i)(-\mathbf{x}_i^*) \\
&= \mathbf{0} - \frac{\tilde{a}\tilde{b}}{t-1} \sum_{i=1}^{t-1} (\mathbf{x}_{i+1} - \mathbf{0}\mathbf{x}_i)(-\mathbf{x}_i^*) \\
&= \frac{\tilde{a}\tilde{b}}{t-1} \sum_{i=1}^{t-1} \mathbf{x}_{i+1} \mathbf{x}_i^*.
\end{aligned}$$

Thus, the proof is completed. \square

A.4 Proof of Proposition 4.1

For the reader's convenience, we restate the proposition as the following.

Proposition A.1. *Let $\mathcal{D}_{\mathbf{x}_1}$ be the multivariate normal distribution $\mathcal{N}(\mathbf{0}_d, \sigma^2 \mathbf{I}_d)$ with any $\sigma^2 > 0$, then the "simple" AR process can not be recovered by the trained transformer even in the ideal case with long training context. Formally, when the training sequence length T_{tr} is large enough, for any test context length T_{te} and dimension $j \in [d]$, the prediction from the trained transformer satisfies*

$$E_{\mathbf{x}_1, \mathbf{W}} \left[\frac{(\hat{\mathbf{y}}_{T_{te}})_j}{(\mathbf{W}\mathbf{x}_{T_{te}})_j} \right] \rightarrow \frac{1}{5}.$$

Therefore, the prediction $\hat{\mathbf{y}}_{T_{te}}$ will not converges to the true next token $\mathbf{W}\mathbf{x}_{T_{te}}$.

Proof. First, built upon the results in Theorem 4.1, when T_{tr} is large enough, we have

$$\tilde{a}\tilde{b} = \frac{\kappa_1}{\kappa_2 + \frac{\kappa_3}{T_{tr}-2} \sum_{t=2}^{T_{tr}-1} \frac{1}{t-1}} \rightarrow \frac{\kappa_1}{\kappa_2} = \frac{\mathbb{E}[x_{1j}^4]}{\mathbb{E}[x_{1j}^6]} = \frac{3\sigma^4}{15\sigma^6} = \frac{1}{5\sigma^2}.$$

Second, by directly calculating, we have

$$(\mathbf{W}\mathbf{x}_{T_{te}})_j = (\mathbf{W}^{T_{te}}\mathbf{x}_1)_j = \lambda_j^{T_{te}} x_{1j},$$

and

$$(\hat{\mathbf{y}}_{T_{te}})_j = \frac{\tilde{a}\tilde{b}}{T_{te}-1} \sum_{i=1}^{T_{te}-1} \sum_{k=1}^d \lambda_j^i \lambda_k^{T_{te}-i} x_{1j} x_{1k}^2.$$

Therefore, we have

$$\begin{aligned} E_{x_1, W} \left[\frac{(\widehat{y}_{T_{te}})_j}{(Wx_{T_{te}})_j} \right] &= E_{x_1, W} \left[\frac{\widetilde{ab}}{T_{te} - 1} \sum_{i=1}^{T_{te}-1} \sum_{k=1}^d \lambda_j^{i-T_{te}} \lambda_k^{T_{te}-i} x_{1k}^2 \right] \\ &= E_{x_1} \left[\frac{\widetilde{ab}}{T_{te} - 1} \sum_{i=1}^{T_{te}-1} x_{1j}^2 \right] = \widetilde{ab} \sigma^2. \end{aligned}$$

Since $\widetilde{ab} < \frac{1}{5\sigma^2}$ and converges to $\frac{1}{5\sigma^2}$ when T_{tr} is large enough, the proof is completed. \square

A.5 Derivation of Example 4.1

Proof. We first prove that the example satisfies Assumption 4.1. Because only one element of \mathbf{x}_1 sampled from Example 4.1 will be non-zero, we have $\mathbb{E}_{\mathbf{x}_1 \sim \mathcal{D}_{\mathbf{x}_1}} [x_{1i_1} x_{1i_2}^{r_2} \cdots x_{1i_n}^{r_n}] = \mathbb{E}_{\mathbf{x}_1 \sim \mathcal{D}_{\mathbf{x}_1}} [0] = 0$ for any subset $\{i_1, \dots, i_n \mid n \leq 4\}$ of $[d]$, and $r_2, \dots, r_n \in \mathbb{N}$. In addition, for any $j \in [d]$, we can derive that

$$\begin{aligned} \kappa_1 &= \mathbb{E}[x_{1j}^4] = \frac{1}{d} \cdot c^4 + \frac{d-1}{d} \cdot 0 = \frac{c^4}{d}, \\ \kappa_2 &= \mathbb{E}[x_{1j}^6] = \frac{1}{d} \cdot c^6 + \frac{d-1}{d} \cdot 0 = \frac{c^6}{d}, \\ \kappa_3 &= \sum_{r \neq j} \mathbb{E}[x_{1j}^2 x_{1r}^4] = 0. \end{aligned}$$

Second, we prove that it satisfies Assumption 4.2 as follows. Without loss of general, we assume that the first coordinate of \mathbf{x}_1 is c .

$$\begin{aligned} \frac{\kappa_1}{\kappa_2} \frac{\sum_{i=1}^{T_{te}-1} \mathbf{x}_i \mathbf{x}_i^*}{T_{te} - 1} \mathbf{x}_{T_{te}} &= \frac{1}{c^2} \text{diag}(c^2, 0, \dots, 0) (\lambda_1^{T_{te}-1} c, 0, \dots, 0)^\top \\ &= (\lambda_1^{T_{te}-1} c, 0, \dots, 0)^\top = \mathbf{x}_{T_{te}}. \end{aligned}$$

The proof is finished. \square

A.6 Proof of Theorem 4.2

For the reader's convenience, we restate the theorem as the following.

Theorem A.1. *Suppose that Assumption 4.1 holds, then Assumption 4.2 is the sufficient and necessary condition for the trained transformer to learn the AR process. Formally, when the training sequence length T_{tr} and test context length T_{te} are large enough, the prediction from the trained transformer satisfies*

$$\widehat{\mathbf{y}}_{T_{te}} \rightarrow \mathbf{W} \mathbf{x}_{T_{te}}, \quad T_{tr}, T_{te} \rightarrow +\infty.$$

Proof. First, built upon the results in Theorem 4.1, when T_{tr} is large enough, we have

$$\widetilde{ab} = \frac{\kappa_1}{\kappa_2 + \frac{\kappa_3}{T_{tr}-2} \sum_{t=2}^{T_{tr}-1} \frac{1}{t-1}} \rightarrow \frac{\kappa_1}{\kappa_2}.$$

Second, when T_{te} is large enough, by Assumption 4.2

$$\frac{\kappa_1}{\kappa_2} \frac{\sum_{i=1}^{T_{te}-1} \mathbf{x}_i \mathbf{x}_i^*}{T_{te} - 1} \mathbf{x}_{T_{te}} \rightarrow \mathbf{x}_{T_{te}}.$$

Therefore, we have

$$\widehat{\mathbf{y}}_{T_{te}} = \mathbf{W} \left(\widetilde{ab} \frac{\sum_{i=1}^{T_{te}-1} \mathbf{x}_i \mathbf{x}_i^*}{T_{te} - 1} \right) \mathbf{x}_{T_{te}} \xrightarrow{T_{tr}} \mathbf{W} \left(\frac{\kappa_1}{\kappa_2} \frac{\sum_{i=1}^{T_{te}-1} \mathbf{x}_i \mathbf{x}_i^*}{T_{te} - 1} \right) \mathbf{x}_{T_{te}} \xrightarrow{T_{te}} \mathbf{W} \mathbf{x}_{T_{te}}$$

which finishes the proof. \square

A.7 Proof of Theorem 4.3

For the reader's convenience, we restate the theorem as the following.

Theorem A.2. *Suppose the initialization satisfies Assumption 3.1, the initial token is fixed as $\mathbf{1}_d$, and we clip non-diagonal gradients of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} during the training, then the gradient flow of the one-layer linear transformer over the population AR loss converges to the same structure as the result in Theorem 4.1, with*

$$\tilde{a}\tilde{b} = \frac{1}{1 + \frac{d-1}{T-2} \sum_{t=2}^{T-1} \frac{1}{t-1}}.$$

Therefore, the obtained transformer performs one step of gradient descent in this case.

The proof is similar to the proof of Theorem 4.1 in Appendix A.2. But the calculating for the gradients is more difficult than that of Theorem 4.1. Similarly, we first present and prove the following lemmas.

Lemma A.5 (dynamical system of gradient flow). *Under the same assumption as in Theorem 4.3, the dynamical process of the parameters in the diagonal of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} satisfies*

$$\begin{aligned} \frac{d}{d\tau} a &= -ab^2 \left[(T-2) + \sum_{t=2}^{T-1} \frac{d-1}{t-1} \right] + b(T-2), \\ \frac{d}{d\tau} b &= -a^2b \left[(T-2) + \sum_{t=2}^{T-1} \frac{d-1}{t-1} \right] + a(T-2), \end{aligned}$$

while the gradients for all other parameters were kept at zero during the training process.

Proof. Recall that in Appendix A.2.2, we have already known that the population loss $L(\boldsymbol{\theta})$ in Eq. 2 can be rewritten as

$$L(\boldsymbol{\theta}) = \sum_{t=2}^{T-1} L_t(\boldsymbol{\theta}) = \sum_{t=2}^{T-1} \sum_{j=1}^d \mathbb{E} \left[\frac{1}{2} \hat{y}_{t,j}^* \hat{y}_{t,j} - \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) + \frac{1}{2} x_{t+1,j}^* x_{t+1,j} \right].$$

Besides, the derivatives of $L_t(\boldsymbol{\theta})$ with respect to $\text{Vec}(\mathbf{A})$ and \mathbf{B}_j are

$$\nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta}) = \frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right] - \mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right],$$

and

$$\nabla_{\text{Vec}(\mathbf{A})} L_t(\boldsymbol{\theta}) = \frac{1}{2} \sum_{j=1}^d \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \hat{y}_{t,j}^* \hat{y}_{t,j} \right] - \sum_{j=1}^d \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right].$$

Step one: calculate $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$. Similarly to the step one in Appendix A.2.2, the

$\mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$ can be derived as the following.

$$\begin{aligned} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right] &= \text{Re} \left(\mathbb{E} \left[x_{t+1,j}^* \cdot \mathbf{e}_t^{\mathbf{x}\top} \otimes \frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \right] \right) \\ &= \text{Re} \left(\mathbb{E} \left[\lambda_j^{-t} \cdot \text{Vec} \left(\frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{A} \mathbf{e}_t^{\mathbf{x}} \right) \right] \right) \quad (\text{use generating process}) \\ &= \text{Re} \left(\mathbb{E} \left[\lambda_j^{-t} \cdot \left(\frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{A} \mathbf{e}_t^{\mathbf{x}} \right) \right] \right). \end{aligned}$$

We note that for any $l \in [2d]$, based on the sparsity of \mathbf{A} , we have

$$\left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l = \begin{cases} \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_l^i \lambda_r^{t-i}, & l \in [d], \\ \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_{l-d}^{i-1} \lambda_r^{t-i}, & l \in [2d] - [d], \end{cases}$$

which implies that

$$\mathbb{E} \left[\lambda_j^{-t} \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right) \right] = \begin{cases} \frac{a}{\rho_t} (t-1), & l = j, \\ 0, & l \neq j, \end{cases}$$

and the l -th element of $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$ is

$$\mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]_l = \begin{cases} \frac{a}{\rho_t} (t-1), & l = j, \\ 0, & l \neq j. \end{cases}$$

Step two: calculate $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right]$. Similarly to the step two in Appendix A.2.2, we can simplify the $\mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right]$ as follows.

$$\begin{aligned} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right] &= 2 \text{Re} \left(\mathbb{E} \left[\mathbf{e}_t^{x\top} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \text{Vec}(\mathbf{A}) \cdot \text{Vec}^\top(\mathbf{A}) \mathbf{e}_t^{\bar{x}} \otimes \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \cdot \mathbf{B}_j \right] \right) \\ &= 2 \text{Re} \left(\mathbb{E} \left[b \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{A} \mathbf{e}_t^{\bar{x}} \right)_j \cdot \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right] \right). \end{aligned} \quad (\text{sparsity of } \mathbf{B})$$

For any $j \in [d]$ and $l \in [2d]$, we can calculate that

$$\left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{A} \mathbf{e}_t^{\bar{x}} \right)_j = \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_j^{-i} \lambda_r^{i-t},$$

and recall that

$$\left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l = \begin{cases} \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_l^i \lambda_r^{t-i}, & l \in [d], \\ \frac{a}{\rho_t} \sum_{i=1}^{t-1} \sum_{r=1}^d \lambda_{l-d}^{i-1} \lambda_r^{t-i}, & l \in [2d] - [d]. \end{cases}$$

With careful computing, we have

$$\mathbb{E} \left[b \left(\frac{\overline{\mathbf{E}_t^x \mathbf{E}_t^{x\top}}}{\rho_t} \mathbf{A} \mathbf{e}_t^{\bar{x}} \right)_j \cdot \left(\frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \mathbf{A} \mathbf{e}_t^x \right)_l \right] = \begin{cases} \frac{a^2 b}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)], & l = j, \\ \frac{a^2 b}{\rho_t^2} (t-1), & l \in [d] - j, \\ 0, & \text{otherwise,} \end{cases}$$

which implies that the l -th element of $\frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right]$ is

$$\frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right]_l = \begin{cases} \frac{a^2 b}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)], & l = j, \\ \frac{a^2 b}{\rho_t^2} (t-1), & l \in [d] - j, \\ 0, & \text{otherwise.} \end{cases}$$

Step three: calculate $\nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})$ and $\nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})$. Based on steps one and two, the l -th element of $\nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})$ can be derived as follows.

$$\nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})_l = \frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{B}_j} \hat{y}_{t,j}^* \hat{y}_{t,j} \right]_l - \mathbb{E} \left[\nabla_{\mathbf{B}_j} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]_l$$

$$= \begin{cases} \frac{a^2 b}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)] - \frac{a}{\rho_t} (t-1), & l = j, \\ \frac{a^2 b}{\rho_t^2} (t-1), & l \in [d] - j, \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, the l -th element of $\nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})$ is

$$\begin{aligned} \nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})_l &= \sum_{t=2}^{T-1} \nabla_{\mathbf{B}_j} L_t(\boldsymbol{\theta})_l \\ &= \begin{cases} \sum_{t=2}^{T-1} \left(\frac{a^2 b}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)] - \frac{a}{\rho_t} (t-1) \right), & l = j, \\ \sum_{t=2}^{T-1} \frac{a^2 b}{\rho_t^2} (t-1), & l \in [d] - j, \\ 0, & \text{otherwise.} \end{cases} \end{aligned} \quad (5)$$

If we clip the non-diagonal gradient of \mathbf{W}_{12}^{PV} , we have

$$\begin{aligned} \nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})_l &= \begin{cases} \sum_{t=2}^{T-1} \left(\frac{a^2 b}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)] - \frac{a}{\rho_t} (t-1) \right), & l = j, \\ 0, & \text{otherwise,} \end{cases} \\ &= \begin{cases} a^2 b \left[(T-2) + \sum_{t=2}^{T-1} \frac{d-1}{t-1} \right] - a(T-2), & l = j, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Step four: calculate $\mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$ **and** $\sum_{j=1}^d \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$. Similarly to the step four in Appendix A.2.2, the $\mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$ can be derived as the following.

$$\begin{aligned} \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right] &= \text{Re} \left(\mathbb{E} \left[x_{t+1,j}^* \cdot \mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}_t^{\mathbf{x}}} \mathbf{E}_t^{\mathbf{x}\top}}{\rho_t} \cdot \mathbf{B}_j \right] \right) \\ &= \text{Re} \left(\mathbb{E} \left[\lambda_j^{-t} \cdot \text{Vec} \left(\frac{\overline{\mathbf{E}_t^{\mathbf{x}}} \mathbf{E}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top} \right) \right] \right). \end{aligned}$$

For any $s, r \in [2d]$, we have

$$\left(\frac{\overline{\mathbf{E}_t^{\mathbf{x}}} \mathbf{E}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top} \right)_{sr} = \begin{cases} \frac{b}{\rho_t} \left(\sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \right) \lambda_r^{t-1}, & s \in [d], r \in [d], \\ \frac{b}{\rho_t} \left(\sum_{i=1}^{t-1} \lambda_j^i \lambda_{s-d}^{1-i} \right) \lambda_r^{t-1}, & s \in [2d] - [d], r \in [d], \\ \frac{b}{\rho_t} \left(\sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \right) \lambda_{r-d}^{t-2}, & s \in [d], r \in [2d] - [d], \\ \frac{b}{\rho_t} \left(\sum_{i=1}^{t-1} \lambda_j^i \lambda_{s-d}^{1-i} \right) \lambda_{r-d}^{t-2}, & s \in [2d] - [d], r \in [2d] - [d], \end{cases}$$

which implies that

$$\mathbb{E} \left[\lambda_j^{-t} \left(\frac{\overline{\mathbf{E}_t^{\mathbf{x}}} \mathbf{E}_t^{\mathbf{x}\top}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top} \right)_{sr} \right] = \begin{cases} \frac{b}{\rho_t} (t-1), & s = d+j, r = j, \\ \frac{b}{\rho_t}, & s \neq d+j, r = j, \\ 0, & \text{otherwise.} \end{cases}$$

and the (s, r) -th element of $\mathbb{E} \left[\nabla_{\mathbf{A}} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]$ is

$$\mathbb{E} \left[\nabla_{\mathbf{A}} \text{Re} \left(x_{t+1,j}^* \hat{y}_{t,j} \right) \right]_{sr} = \begin{cases} \frac{b}{\rho_t} (t-1), & s = d+j, r = j, \\ \frac{b}{\rho_t}, & s \neq d+j, r = j, \\ 0, & \text{otherwise.} \end{cases}$$

Finally, we can calculate $\sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right]$ as

$$\sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right]_{sr} = \begin{cases} \frac{b}{\rho_t} (t-1), & s-d=r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ \frac{b}{\rho_t}, & s-d \neq r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise.} \end{cases}$$

Step five: calculate $\mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]$ **and** $\sum_{j=1}^d \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]$. Similarly to the step five in Appendix A.2.2, the $\mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]$ is simplified as follows.

$$\begin{aligned} \mathbb{E} \left[\nabla_{\text{Vec}(\mathbf{A})} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right] &= 2\text{Re} \left(\mathbb{E} \left[\mathbf{e}_t^{\mathbf{x}} \otimes \frac{\overline{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}\top}}}{\rho_t} \mathbf{B}_j \cdot \mathbf{B}_j^\top \cdot \mathbf{e}_t^{\mathbf{x}*} \otimes \frac{\overline{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}\top}}}{\rho_t} \cdot \text{Vec}(\mathbf{A}) \right] \right) \\ &= 2\text{Re} \left(\mathbb{E} \left[\sum_{k=d+1}^{2d} a \left(\frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}*} \right)_{k,k-d} \cdot \text{Vec} \left(\frac{\overline{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top} \right) \right] \right). \end{aligned}$$

For any $k \in [2d] - [d]$, we can calculate that

$$\left(\frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}*} \right)_{k,k-d} = \frac{b}{\rho_t} \left(\sum_{i=1}^{t-1} \lambda_j^{-i} \lambda_{k-d}^{i-1} \right) \lambda_{k-d}^{1-t} = \frac{b}{\rho_t} \sum_{i=1}^{t-1} \lambda_j^{-i} \lambda_{k-d}^{i-t},$$

and recall that for any $s, r \in [2d]$, we have

$$\left(\frac{\overline{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top} \right)_{sr} = \begin{cases} \frac{b}{\rho_t} \left(\sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \right) \lambda_r^{t-1}, & s \in [d], r \in [d], \\ \frac{b}{\rho_t} \left(\sum_{i=1}^{t-1} \lambda_j^i \lambda_{s-d}^{1-i} \right) \lambda_r^{t-1}, & s \in [2d] - [d], r \in [d], \\ \frac{b}{\rho_t} \left(\sum_{i=1}^t \lambda_j^{i-1} \lambda_s^{1-i} \right) \lambda_{r-d}^{t-2}, & s \in [d], r \in [2d] - [d], \\ \frac{b}{\rho_t} \left(\sum_{i=1}^{t-1} \lambda_j^i \lambda_{s-d}^{1-i} \right) \lambda_{r-d}^{t-2}, & s \in [2d] - [d], r \in [2d] - [d]. \end{cases}$$

With careful computing, we have

$$\begin{aligned} &\mathbb{E} \left[\sum_{k=d+1}^{2d} a \left(\frac{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}*}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}*} \right)_{k,k-d} \cdot \left(\frac{\overline{\mathbf{E}_t^{\mathbf{x}} \mathbf{E}_t^{\mathbf{x}\top}}}{\rho_t} \mathbf{B}_j \mathbf{e}_t^{\mathbf{x}\top} \right)_{sr} \right] \\ &= \begin{cases} \frac{ab^2}{\rho_t^2} (t-1)^2, & s = d+j, r = j, \\ \frac{ab^2}{\rho_t^2} (t-1), & s \neq d+j, r = j, \\ \frac{ab^2}{\rho_t^2} (t-1), & s = d+r, r \neq j, \\ \frac{ab^2}{\rho_t^2} (t-1), & s = d+j, r \neq j, \\ \frac{ab^2}{\rho_t^2}, & \text{remains in } \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise,} \end{cases} \end{aligned}$$

which implies that the l -th element of $\frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right]$ is

$$\frac{1}{2} \mathbb{E} \left[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right] = \begin{cases} \frac{ab^2}{\rho_t^2} (t-1)^2, & s = d+j, r = j, \\ \frac{ab^2}{\rho_t^2} (t-1), & s \neq d+j, r = j, \\ \frac{ab^2}{\rho_t^2} (t-1), & s = d+r, r \neq j, \\ \frac{ab^2}{\rho_t^2} (t-1), & s = d+j, r \neq j, \\ \frac{ab^2}{\rho_t^2}, & \text{remains in } \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise.} \end{cases}$$

Based on these results, we can derive

$$\frac{1}{2} \sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right] = \begin{cases} \frac{ab^2}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)], & s-d=r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ \frac{ab^2}{\rho_t^2} [2(t-1) + d-2], & s-d \neq r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise.} \end{cases}$$

Step six: calculate $\nabla_{\text{Vec}(\mathbf{A})} L_t(\boldsymbol{\theta})$ and $\nabla_{\text{Vec}(\mathbf{A})} L(\boldsymbol{\theta})$. Based on steps four and five, the (s, r) -th element of $\nabla_{\mathbf{A}} L_t(\boldsymbol{\theta})$ can be derived as follows.

$$\begin{aligned} \nabla_{\mathbf{A}} L_t(\boldsymbol{\theta})_{sr} &= \frac{1}{2} \sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \widehat{y}_{t,j}^* \widehat{y}_{t,j} \right] - \sum_{j=1}^d \mathbb{E} \left[\nabla_{\mathbf{A}} \text{Re} \left(x_{t+1,j}^* \widehat{y}_{t,j} \right) \right] \\ &= \begin{cases} \frac{ab^2}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)] - \frac{b}{\rho_t} (t-1), & s-d=r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ \frac{ab^2}{\rho_t^2} [2(t-1) + d-2] - \frac{b}{\rho_t}, & s-d \neq r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Furthermore, the (s, r) -th element of $\nabla_{\mathbf{A}} L(\boldsymbol{\theta})$ is

$$\begin{aligned} \nabla_{\mathbf{A}} L(\boldsymbol{\theta})_{sr} &= \sum_{t=2}^{T-1} \nabla_{\mathbf{A}} L_t(\boldsymbol{\theta})_{sr} \\ &= \begin{cases} \sum_{t=2}^{T-1} \left(\frac{ab^2}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)] - \frac{b}{\rho_t} (t-1) \right), & s-d=r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ \sum_{t=2}^{T-1} \left[\frac{ab^2}{\rho_t^2} [2(t-1) + d-2] - \frac{b}{\rho_t} \right], & s-d \neq r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise.} \end{cases} \quad (6) \end{aligned}$$

If we clip the non-diagonal gradient of \mathbf{W}_{32}^{KQ} , we have

$$\begin{aligned} \nabla_{\mathbf{A}} L(\boldsymbol{\theta})_{sr} &= \sum_{t=2}^{T-1} \nabla_{\mathbf{A}} L_t(\boldsymbol{\theta})_{sr} \\ &= \begin{cases} \sum_{t=2}^{T-1} \left(\frac{ab^2}{\rho_t^2} [(t-1)^2 + (d-1)(t-1)] - \frac{b}{\rho_t} (t-1) \right), & s-d=r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise,} \end{cases} \\ &= \begin{cases} ab^2 \left[(T-2) + \sum_{t=2}^{T-1} \frac{d-1}{t-1} \right] - b(T-2), & s-d=r, A_{sr} \in \mathbf{W}_{32}^{KQ}, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Step seven: summarize the result by induction. From the gradient of $\nabla_{\mathbf{A}} L(\boldsymbol{\theta})$ and $\nabla_{\mathbf{B}_j} L(\boldsymbol{\theta})$, we observe that non-zero gradients only emerge in the diagonal of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} , and they are same. Therefore, the parameter matrices keep the same structure as the initial time. Thus we can summarize the dynamic system as the following.

$$\begin{aligned} \frac{d}{d\tau} a &= -ab^2 \left[(T-2) + \sum_{t=2}^{T-1} \frac{d-1}{t-1} \right] + b(T-2), \\ \frac{d}{d\tau} b &= -a^2b \left[(T-2) + \sum_{t=2}^{T-1} \frac{d-1}{t-1} \right] + a(T-2), \end{aligned}$$

which completes the proof. \square

Lemma A.6. *Suppose that the preconditions of Theorem 4.3 hold, and denote $(T-2) + \sum_{t=2}^{T-1} \frac{d-1}{t-1}$ and $T-2$ by c_1 and c_2 , respectively. Then, the dynamics are the same as those of gradient flow on the following objective function:*

$$\tilde{\ell}(a, b) = \frac{1}{2c_1} (c_2 - c_1 ab)^2,$$

whose global minimums satisfy $ab = c_2/c_1$.

Table 1: Step size in different simulations.

\mathbf{x}_1	σ/c	step size
Gaussian	0.5	0.001
	1	0.0001
	2	0.000002
Example 4.1	0.5	0.03
	1	0.001
	2	0.0001
$\mathbf{1}_d$	-	0.0005

Proof. The proof is the same as that of Lemma 5.3 in Appendix A.2.3. □

Using the results from the above lemmas and Lemma 5.4, we can conclude Theorem 4.3.

A.8 Proof of Proposition 4.2

For the reader’s convenience, we restate the proposition as the following.

Proposition A.2. *The limiting point found by the gradient does not share the same structure as that in Theorem 4.1, thus the trained transformer will not implement one step of gradient descent for minimizing $\frac{1}{2} \sum_{i=1}^{t-1} \|\mathbf{x}_{i+1} - W \mathbf{x}_i\|^2$.*

Proof. From Eq. 6 and Eq. 5, we know that when the parameters matrices share the same structure as the result in Theorem 4.1, the non-zero gradients of the non-diagonal elements of \mathbf{W}_{32}^{KQ} and \mathbf{W}_{12}^{PV} will occur, which implies the result. □

Appendix B Experimental details and additional results

B.1 GPU and random seed

The random seed in the experiments is fixed as 1. All experiments are done on a single GeForce RTX 3090 GPU in one hour.

B.2 Step size in simulations

The step size of the gradient descent in different simulations is summarized in Table 1.

B.3 Additional results for Gaussian initial token

In practice, we estimate the ratio of two vectors by calculating the mean of the element-wise divide between two vectors. The results for $\sigma = 1, 2$ and $T = 100$ with diagonal initialization are presented in Figure 2. The results for $\sigma = 1$ and $T = 5$ (small-context scenarios) with diagonal initialization are presented in Figure 4. The results for $\sigma = 1$ and $T = 100$ with Gaussian initialization ($\sigma_w = 0.001, 0.01, 0.1$) are presented in Figure 5.

B.4 Additional results for Sparse initial token (Example 4.1)

The results for $c = 1, 2$ and $T = 100$ with diagonal initialization are presented in Figure 3. The results for $c = 1, 2$ and $T = 100$ with Gaussian initialization ($\sigma_w = 0.001, 0.01, 0.1$) are presented in Figure 6.

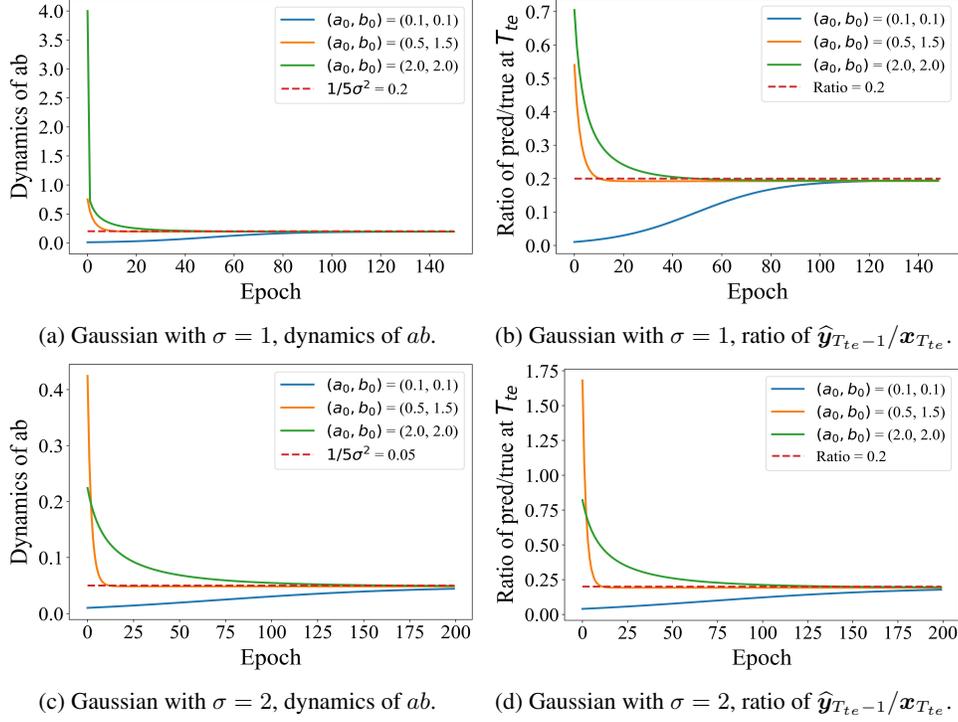


Figure 2: Simulations results in Gaussian initial token with $\sigma = 1, 2$. The results show that the convergence of ab satisfies Theorem 4.1. In addition, the trained transformer can not recover the Gaussian initial token, which verifies Proposition 4.1.

B.5 Additional results for full-one initial token

\mathbf{W}^{KQ} and \mathbf{W}^{PV} with different diagonal initializations are presented in Figure 7. From the results, we know that the gradient descent converges to

$$\begin{pmatrix} \widetilde{\mathbf{W}}_{22}^{KQ} & \widetilde{\mathbf{W}}_{23}^{KQ} \\ \widetilde{\mathbf{W}}_{32}^{KQ} & \widetilde{\mathbf{W}}_{33}^{KQ} \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{W}_1 & \mathbf{0}_{d \times d} \end{pmatrix}, \quad \begin{pmatrix} \widetilde{\mathbf{W}}_{12}^{PV} & \widetilde{\mathbf{W}}_{13}^{PV} \end{pmatrix} = (\mathbf{W}_2 \quad \mathbf{0}_{d \times d}),$$

where \mathbf{W}_1 and \mathbf{W}_2 are some dense matrices. Similarly to the proof of Corollary 4.1 in Appendix 4.1, we have

$$\begin{aligned} \hat{\mathbf{y}}_t &= (\mathbf{W}_{12}^{PV} \quad \mathbf{W}_{13}^{PV}) \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \begin{pmatrix} \mathbf{W}_{22}^{KQ} & \mathbf{W}_{23}^{KQ} \\ \mathbf{W}_{32}^{KQ} & \mathbf{W}_{33}^{KQ} \end{pmatrix} \mathbf{e}_t^x \\ &= (\mathbf{W}_2 \quad \mathbf{0}_{d \times d}) \frac{\mathbf{E}_t^x \mathbf{E}_t^{x*}}{\rho_t} \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{W}_1 & \mathbf{0}_{d \times d} \end{pmatrix} \mathbf{e}_t^x \\ &= \frac{1}{\rho_t} (\mathbf{W}_2 \quad \mathbf{0}_{d \times d}) \sum_{i=1}^t \mathbf{e}_i^x \mathbf{e}_i^{x*} \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{W}_1 & \mathbf{0}_{d \times d} \end{pmatrix} \mathbf{e}_t^x \\ &= \frac{1}{\rho_t} \sum_{i=1}^t \mathbf{W}_2 \mathbf{x}_i (\mathbf{x}_{i-1}^* \mathbf{W}_1 \quad \mathbf{0}_{d \times d}) \mathbf{e}_t^x \\ &= \frac{1}{\rho_t} \sum_{i=1}^t \mathbf{W}_2 \mathbf{x}_i \mathbf{x}_{i-1}^* \mathbf{W}_1 \mathbf{x}_t. \end{aligned}$$

which can be seen as a somewhat preconditioned gradient descent on the OLS problem.

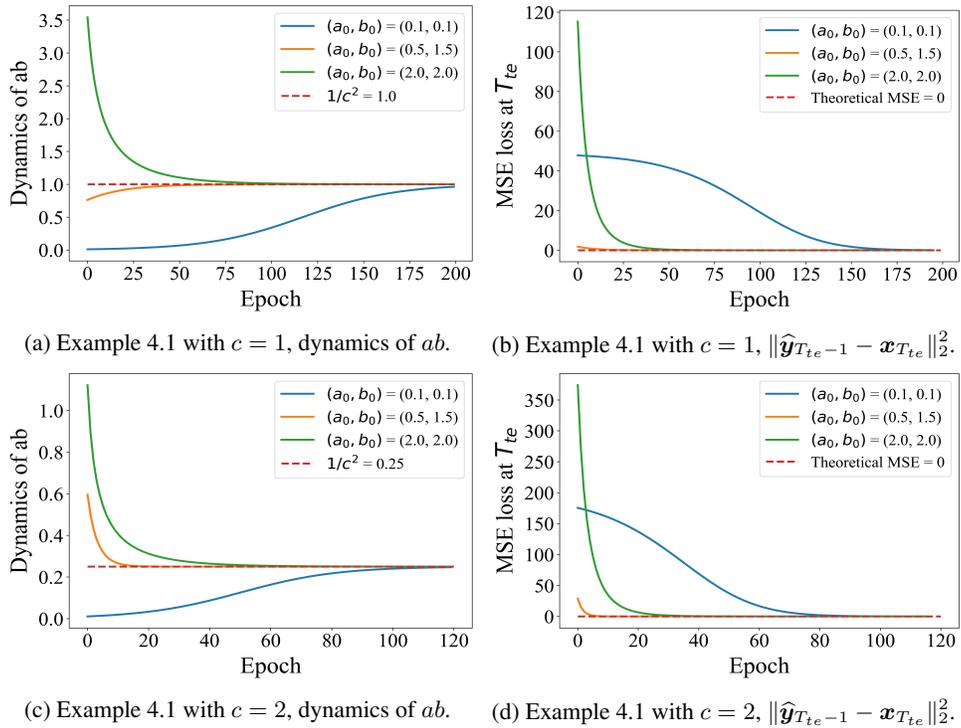


Figure 3: Simulations results on Example 4.1. Results show that the convergence of ab satisfies Theorem 4.1. In addition, the trained transformer can recover the sequence with the initial token from Example 4.1, which verifies Theorem 4.2.

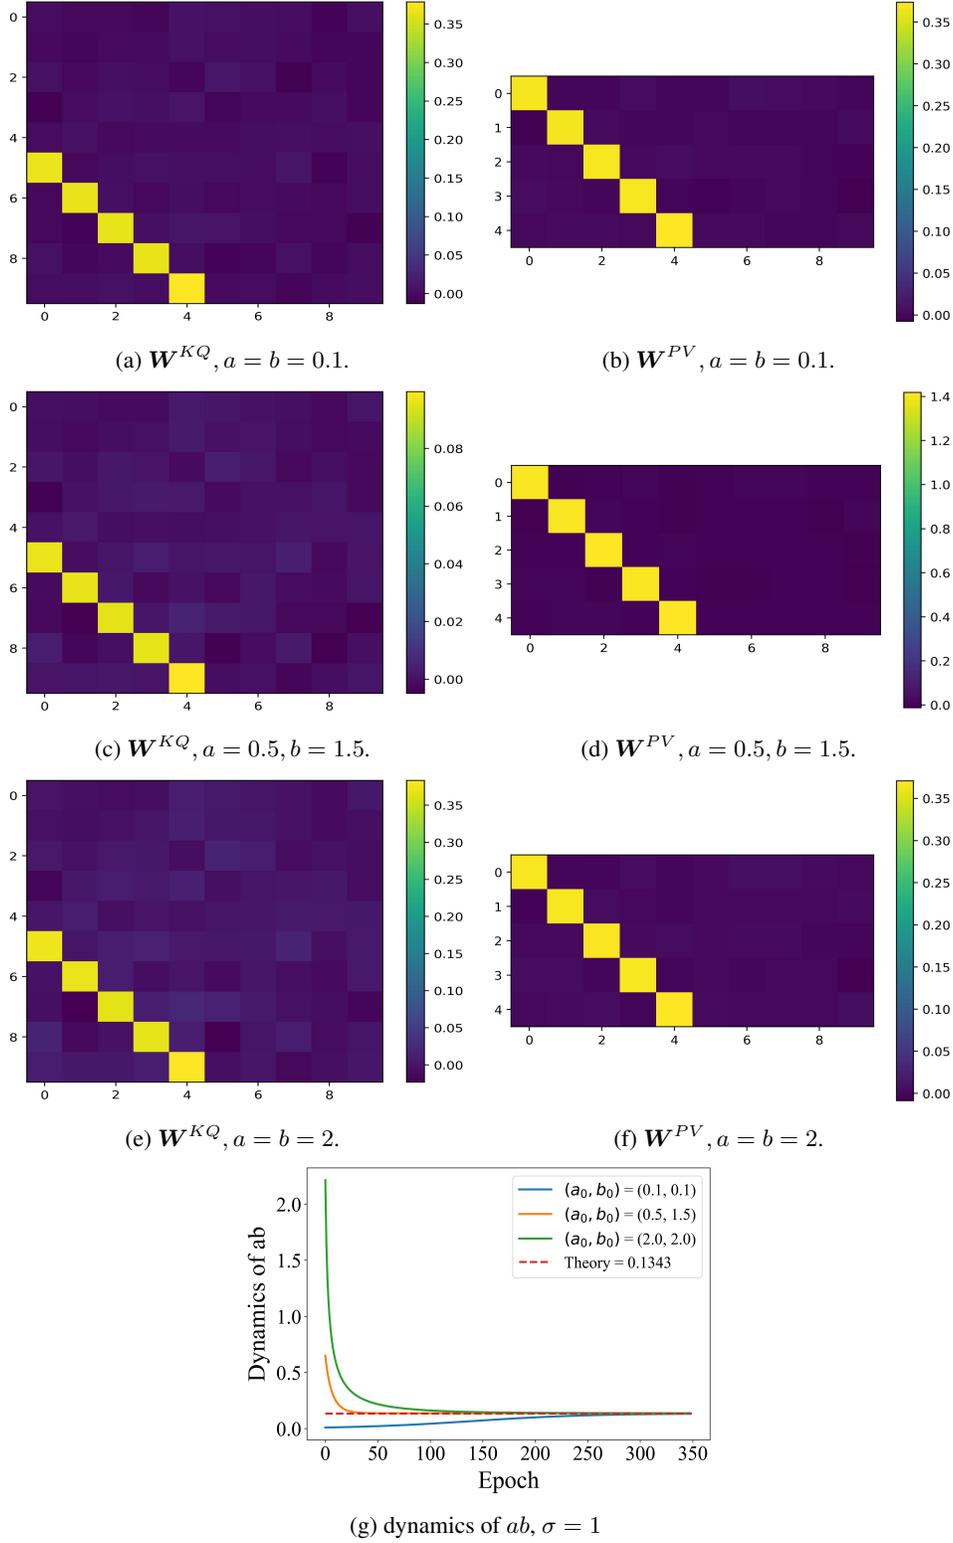


Figure 4: Results of small-context scenarios ($d = T = 5$) with Gaussian start point ($\sigma = 1$) and diagonal initialization, and the theoretical limit is $1/(5\sigma^2 + \frac{d-1}{T-2}\sigma^2 \sum_{t=2}^{T-1} \frac{1}{t-1})$ by Theorem 4.1. The **read blocks** in Assumption 3.1 are presented, which are related to the final prediction. The product ab does converge to the results in Theorem 4.1.

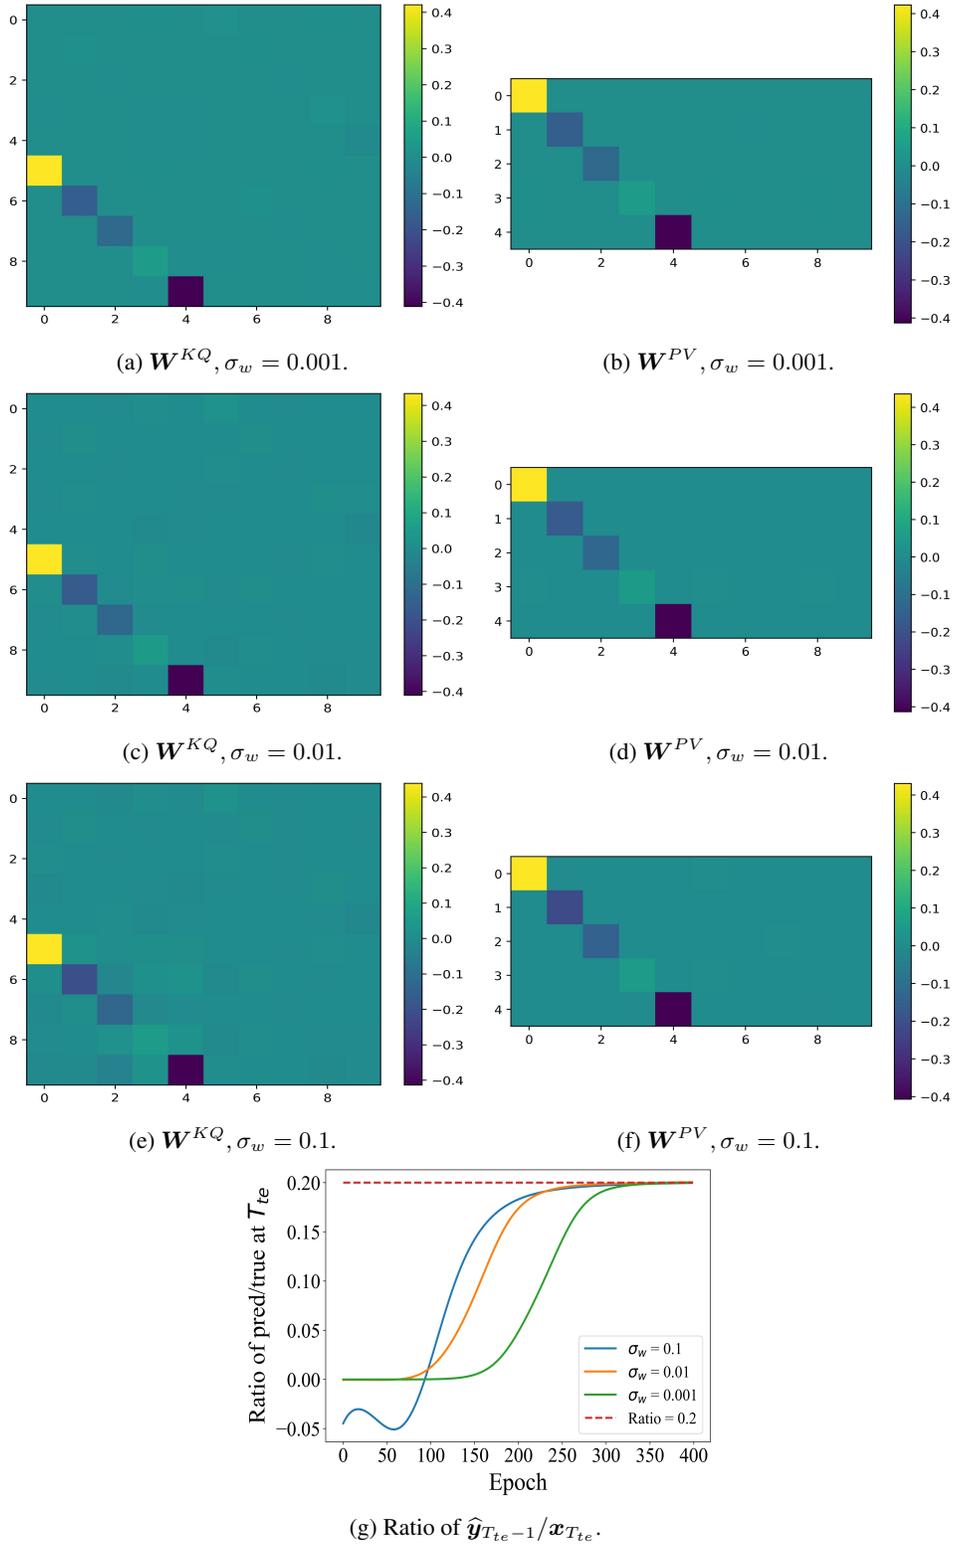


Figure 5: Results of Gaussian start point ($\sigma = 1$) and standard Gaussian initialization with different variance σ_w . The read blocks in Assumption 3.1 are presented, which are related to the final prediction. The parameter matrices retain the same strong diagonal structure and test performance as those of the diagonal initialization.

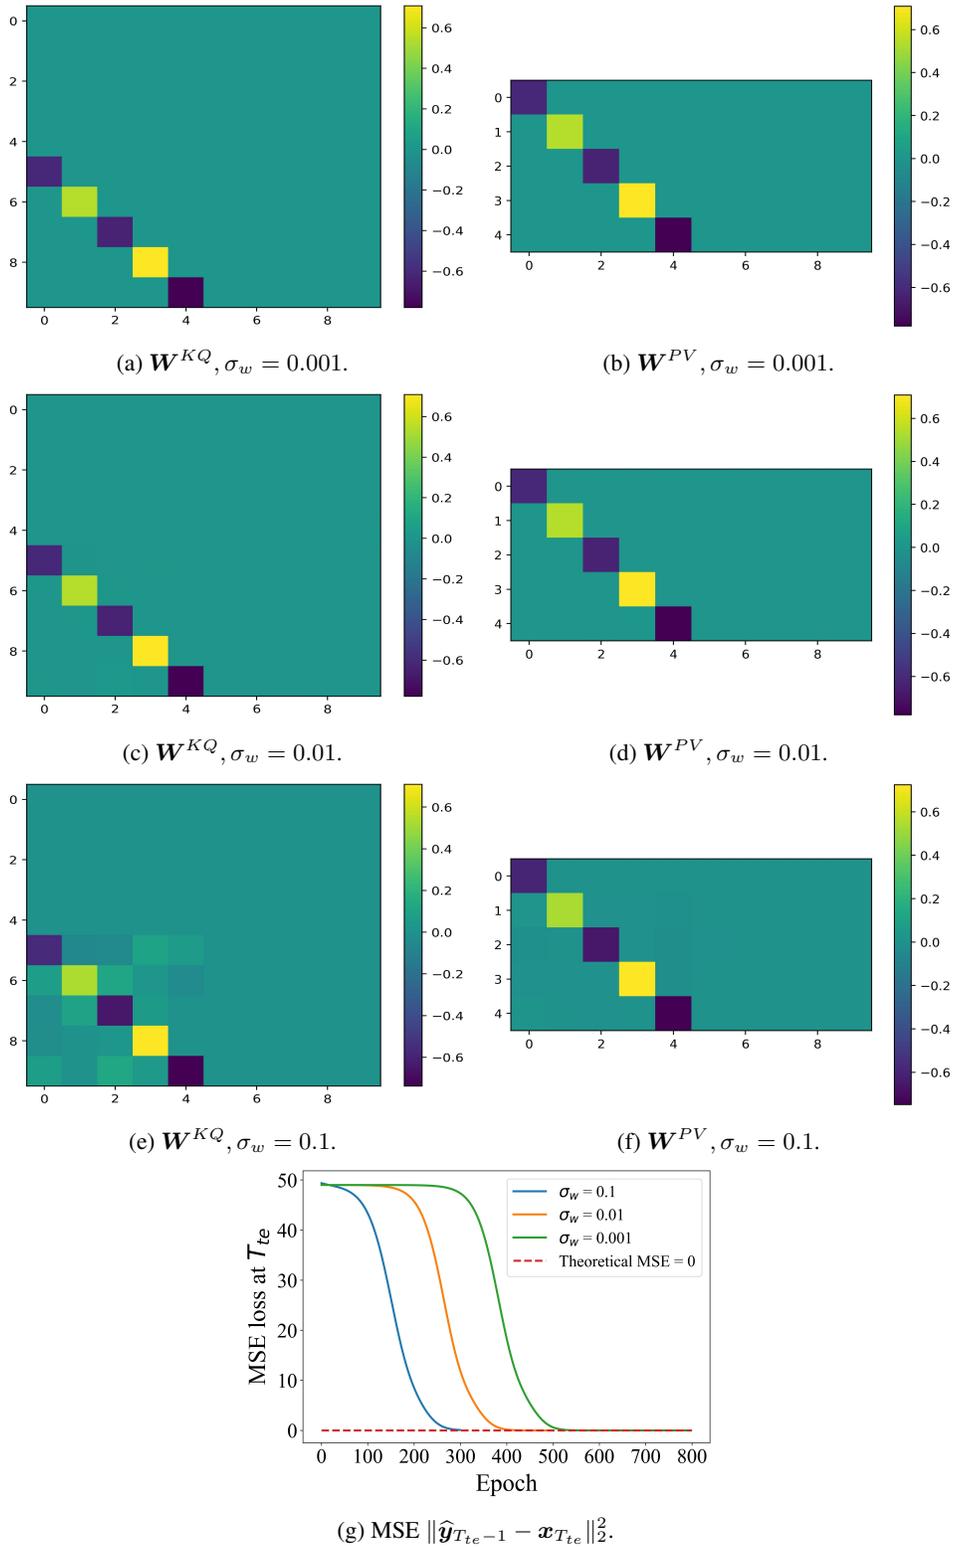


Figure 6: Results of Example 4.1 ($c = 1$) and standard Gaussian initialization with different variance σ_w . The **read blocks** in Assumption 3.1 are presented, which are related to the final prediction. The parameter matrices retain the same strong diagonal structure and test performance as those of the diagonal initialization.

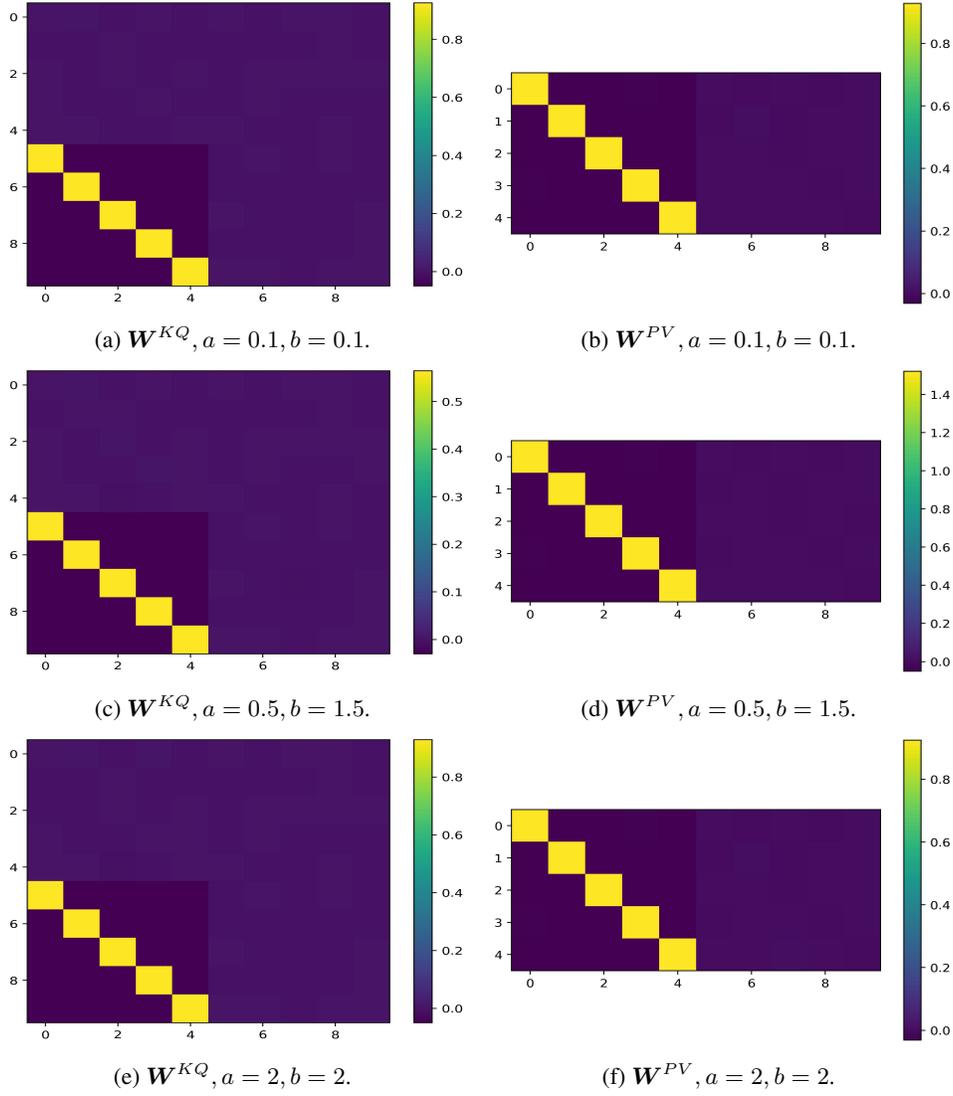


Figure 7: \mathbf{W}^{KQ} and \mathbf{W}^{PV} of full-one start points with different diagonal initialization. The **read blocks** in Assumption 3.1 are presented, which are related to the final prediction.

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: Our claims in the abstract and introduction accurately match our theoretical results and reflect the paper's contribution and scope.

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: Please see Section 7 for details.

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: All assumptions are presented clearly in the main paper (Assumption 3.1, 4.1 and 4.2). Complete proof can be found in Section 5 and Appendix A. The correctness of our proofs can be verified by simulations in Section 6 and Appendix B.

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: We provide complete configurations in Section 6 and Appendix B. We also upload the code at <https://github.com/ML-GSAI/MesaOpt-AR-Transformer> for reproduction as well.

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 upload the complete code and sufficient instructions in the supplemental material. It is available at <https://github.com/ML-GSAI/MesaOpt-AR-Transformer>.

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, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Please see details in Section 6, Appendix B and attached code.

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: We do not present an error bar due to insufficient computing resources. However, we have run all experiments with three different initializations to verify the correctness of our theory, see details in Section 6 and Appendix B.

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: Please see details in Appendix B.

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: We conform research conducted in the paper satisfies the NeurIPS Code of Ethics.

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: [Yes]

Justification: Please see details in Section 7.

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 poses no 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 simulations are simple and we do 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.