
How Transformers Learn Diverse Attention Correlations in Masked Vision Pretraining

Yu Huang^{*1} Zixin Wen^{*2} Yuejie Chi² Yingbin Liang³

Abstract

Masked reconstruction, which predicts randomly masked patches from unmasked ones, has emerged as an important approach in self-supervised pretraining. However, the theoretical understanding of masked pretraining is rather limited, especially for the foundational architecture of transformers. In this paper, to the best of our knowledge, we provide the first end-to-end theoretical guarantee of learning one-layer transformers in masked reconstruction self-supervised pretraining. On the conceptual side, we posit a mechanism of how transformers trained with masked vision pretraining objectives produce empirically observed **local and diverse** attention patterns, on data distributions with spatial structures that highlight *feature-position correlations*. On the technical side, our end-to-end characterization of training dynamics in softmax-attention models simultaneously accounts for input and position embeddings, which is developed based on a careful analysis tracking the interplay between feature-wise and position-wise attention correlations.

1 Introduction

Self-supervised learning has been a leading approach to pre-train neural networks for downstream applications since the introduction of BERT (Devlin et al., 2018) and GPT (Radford et al., 2018) in natural language processing (NLP). On the other hand, in vision, self-supervised learning was initially focused more on *discriminative* methods, which include contrastive learning (He et al., 2020; Chen et al., 2020) and non-contrastive learning methods (Grill et al., 2020; Chen et al., 2020; Caron et al., 2021; Zbontar et al., 2021). Inspired by masked language models in NLP, together with

the crucial progress made by (Dosovitskiy et al., 2020) in successfully implementing vision transformers (ViTs), *generative* approaches, such as masked reconstruction-based methods, have become popular in self-supervised pretraining, especially since emergences of masked autoencoders (MAE) (He et al., 2022) and SimMIM (Xie et al., 2022).

In masked vision pretraining, neural networks are instructed to reconstruct some or all patches of an image given a masked version, aiming to learn certain abstract semantics of visual contents when trained to fill in the missing patches. In practice, this approach not only proves to be very successful with remarkable finetuning performance but also shows intriguing phenomena that differ significantly from the behaviors observed in other self-supervised learning methods. The seminal work of (He et al., 2022) showed that MAE can conduct visual reasoning when filling in masked patches even with very high masking rates, suggesting that masked image modeling learns some complex relationships between visual objects. Some critical observations from recent research (Wei et al., 2022b; Park et al., 2023; Xie et al., 2023) have suggested that the ViTs trained via masked reconstruction objectives display **diverse attention patterns**: different query patches pay attention to distinct local areas. This is in sharp contrast to the discriminative self-supervised learning approaches, whose attention maps mostly capture the most significant global pattern regardless of where the query patches are, leading to a phenomenon known as “attention collapse”, as shown in Figure 1. Given these empirical observations, it naturally prompts the question: from a *theoretical* standpoint, how do ViTs learn these varied attention patterns during masked pretraining? Despite the substantial empirical effort dedicated to investigating masked vision pretraining, our theoretical understanding of it is still nascent. Most existing theories for self-supervised pretraining focused on discriminative methods (Arora et al., 2019; Chen et al., 2021a; Robinson et al., 2021; HaoChen et al., 2021; Wen & Li, 2021; Tian et al., 2021; Wang et al., 2021; Wen & Li, 2022), such as contrastive learning. Among very few attempts towards masked image modeling, (Cao et al., 2022) studied the patch-based attention via an integral kernel perspective; (Zhang et al., 2022) analyzed MAE through an augmentation graph framework, which connects MAE to contrastive learning. (Pan et al., 2022) characterized the

^{*}Equal contribution ¹University of Pennsylvania ²Carnegie Mellon University ³The Ohio State University. Correspondence to: Yu Huang <yuh42@wharton.upenn.edu>.

Work presented at TF2M workshop at ICML 2024, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).




Figure 1: Visualization of attention maps in the last layer of ViT for query patches from two different spatial locations, trained by the generative self-supervised learning approach of masked reconstruction (MAE) and two discriminative self-supervised learning methods (DINO (Caron et al., 2021), MoCo (Chen et al., 2021b)). optimization process of MAE with shallow convolutional neural networks (CNNs). However, **transformers**, as the prevailing architecture in current deep learning practice, were not addressed in the aforementioned theoretical studies of vision pretraining and, more broadly, self-supervised pretraining, leaving a significant gap in the literature.

Building on the mind-blowing empirical advances and recognizing the lack of theoretical understanding of *masked reconstruction-based methods with transformers* in self-supervised learning, we are motivated to ask:

Our Research Questions

Can we theoretically characterize what solutions transformers converge to in masked vision pretraining? How do transformers learn **diverse local attention patterns** rather than object-focused global attention in such pretraining?

Contributions. In this paper, we take a step towards answering the above questions by studying the training process of one-layer transformers trained by gradient descent (GD) for masked reconstruction objectives. We focus on spatially structured data distributions, where each image is sampled from different clusters characterized by different patch-wise associations. Under such setting:

1. We provide the global convergence guarantee of the masked reconstruction loss and characterize how attention is distributed at convergence to demonstrate the non-collapsed diverse local attention patterns obtained by masked pretraining, which to our knowledge is the first end-to-end theory of learning transformers for masked-reconstruction type self-supervised methods.
2. We analyze the training dynamics of *attention correlations* (see Definition A.1) and prove that ViTs manage to capture desirable diverse local patterns in masked pretraining by learning the target **feature-position attention correlations** for all visual features irrespective of their significance i.e., whether they are global or local features. This marks the first result of learning the

softmax attention model that jointly considers input and position encodings.

3. We design a novel empirical metric, termed as the **attention diversity metric**, to probe vision transformers trained by different methods. Our new observations (see Figure 4) provide further evidence for the diverse local patterns revealed through masked image modeling.

2 Problem Setup

In this section, we present our problem formulations for studying the training process of transformers during masked vision pretraining. We begin with some background information, followed by a description of our dataset settings. We then detail the masked pretraining strategy and the specific transformer architecture considered in this paper.

2.1 Masked Vision Pretraining with Transformers

We follow the masked reconstruction frameworks in (He et al., 2022; Xie et al., 2022). Each data sample $X \in \mathbb{R}^{d \times P}$ has the form $X = (X_{\mathbf{p}})_{\mathbf{p} \in \mathcal{P}}$, which has $|\mathcal{P}| = P$ patches, and each patch $X_{\mathbf{p}} \in \mathbb{R}^d$.

Random masking. Given an input image X , let $M(X) \rightarrow \mathbb{R}^{d \times P}$ denote the random masking operation, which randomly selects (without replacement) a subset of patches \mathcal{M} in X with a masking ratio $\gamma = \Theta(1) \in (0, 1)$ and masks them to be $M := \mathbf{0} \in \mathbb{R}^d$, i.e.,

$$M(X)_{\mathbf{p}} = \begin{cases} [X]_{\mathbf{p}} & \mathbf{p} \in \mathcal{U} := \mathcal{P} \setminus \mathcal{M} \\ \mathbf{0} & \mathbf{p} \in \mathcal{M} \end{cases}. \quad (1)$$

Masked reconstruction objective. Let $F : X \mapsto \hat{X}$ be an architecture that outputs a reconstructed image $\hat{X} \in \mathbb{R}^{d \times P}$ for any given input $X \in \mathbb{R}^{d \times P}$. To train the model $F(\cdot)$, in masked pretraining, we minimize the mean-squared reconstruction loss over the masked patches, which can be written as

$$\mathcal{L}(F) = \frac{1}{2} \mathbb{E} \left[\sum_{\mathbf{p} \in \mathcal{M}} \left\| [X]_{\mathbf{p}} - [F(M(X))]_{\mathbf{p}} \right\|_2^2 \right]. \quad (2)$$

where the expectation is with respect to both the data distribution and the masking.

Our theoretical framework is based on F being a simplified version of vision transformers (Dosovitskiy et al., 2020) which utilizes the attention mechanism (Vaswani et al., 2017), which consists of the following components: a query matrix W^Q , a key matrix W^K , and a value matrix W^V . Given an input X , the output of a self-attention layer can be described by the mapping $G(X; W^Q, W^K, W^V) =$

$$\text{softmax} \left((W^Q X)^\top W^K X \right) \cdot (W^V X)^\top, \quad (3)$$

where the $\text{softmax}(\cdot)$ function is applied row-wisely.

To simplify the theoretical analysis, we consolidate the product of query and key matrices $(W^Q)^\top W^K$ into one weight

matrix denoted as Q . Furthermore, we set W^V to be the identity matrix and fixed during the training. These simplifications are often taken in recent theoretical works (Jelassi et al., 2022; Huang et al., 2023; Zhang et al., 2023a) in order to allow tractable analysis. Moreover, to retain the crucial positional information as in practices (Dosovitskiy et al., 2020; He et al., 2022), we add positional encodings to the input embeddings, which has the following assumptions:

Assumption 2.1 (Positional encoding). We assume fixed positional encodings, which is consistent with the implementation in MAE (He et al., 2022): $E = (e_{\mathbf{p}})_{\mathbf{p} \in \mathcal{P}} \in \mathbb{R}^{d \times P}$ where positional embedding vectors $e_{\mathbf{p}}$ are orthogonal to each other and to all the features $v_{k,j}$, and are of unit-norm.

We now define the actual network for masked pretraining.

Definition 2.2 (Transformer network for masked pretraining). We assume that our vision transformer $F(X; Q)$ consists of a single-head self-attention layer with an attention weight matrix $Q \in \mathbb{R}^{d \times d}$. For an input image $X \sim \mathcal{D}$, we add positional encoding by letting $\tilde{X} = X + E$. The attention score from patch $X_{\mathbf{p}}$ to patch $X_{\mathbf{q}}$ is denoted by

$$\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}(X; Q) := \frac{e^{\tilde{X}_{\mathbf{p}}^{\top} Q \tilde{X}_{\mathbf{q}}}}{\sum_{\mathbf{r} \in \mathcal{P}} e^{\tilde{X}_{\mathbf{p}}^{\top} Q \tilde{X}_{\mathbf{r}}}}, \quad \text{for } \mathbf{p}, \mathbf{q} \in \mathcal{P}. \quad (4)$$

Then the output of the transformer for $\mathbf{p} \in \mathcal{P}$ is given by

$$[F(X; Q)]_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{P}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}(X; Q) \cdot X_{\mathbf{q}}. \quad (5)$$

Training algorithm. The learning objective in (2) is minimized via GD with learning rate $\eta > 0$. At $t = 0$, we initialize $Q^{(0)} := \mathbf{0}_{d \times d}$ as the zero matrix. The parameter is updated as follows:

$$Q^{(t+1)} = Q^{(t)} - \eta \nabla_Q \mathcal{L}(Q^{(t)}).$$

2.2 Data Distribution

We assume the data samples $X \in \mathbb{R}^{d \times P}$ are drawn independently based on some data distribution \mathcal{D} . To capture the *feature-position (FP) correlation* in the learning problem, we consider the following setup for vision data. We assume that the data distribution consists of many different clusters, where each cluster captures a distinct spatial pattern, and hence is defined by a different partition of patches with a different set of visual features. We define the data distribution \mathcal{D} formally as follows. An intuitive illustration of data generation is given in Figure 2.

Definition 2.3 (Data distribution \mathcal{D}). The data distribution \mathcal{D} has $K = O(1)$ different clusters $\{\mathcal{D}_k\}_{k=1}^K$. For every cluster \mathcal{D}_k , $k \in [K]$, there is a corresponding partition of \mathcal{P} into N_k disjoint subsets $\mathcal{P} = \bigcup_{j=1}^{N_k} \mathcal{P}_{k,j}$ which we call **areas**. For each sample $X = (X_{\mathbf{p}})_{\mathbf{p} \in \mathcal{P}}$, its sampling process is as follows:




Figure 2: Illustration of our data distribution (see Definition 2.3). Each cluster \mathcal{D}_k is segmented into distinct areas $\mathcal{P}_{k,j}$, with squares in the same color representing the same area $\mathcal{P}_{k,j}$. The global area $\mathcal{P}_{k,1}$ (depicted in orange) contains a larger count of patches compared to any other local areas.

- We draw \mathcal{D}_k uniformly at random from all clusters and draw a sample X from \mathcal{D}_k .
- Given $k \in [K]$, for any $j \in [N_k]$, all patches $X_{\mathbf{p}}$ in the area $\mathcal{P}_{k,j}$ are given the same content $X_{\mathbf{p}} = v_{k,j} z_j(X)$, where $v_{k,j} \in \mathbb{R}^d$ is the *visual* feature and $z_j(X)$ is the latent variable. We assume $\bigcup_{k=1}^K \bigcup_{j=1}^{N_k} \{v_{k,j}\}$ are orthogonal to each other with unit norm.
- Given $k \in [K]$, for any $j \in [N_k]$, $z_j(X) \in [L, U]$, where $0 \leq L < U$ are on the order of $\Theta(1)$.¹

Global and local features in an image. Image data naturally contains two types of features: the global features and the local features. For instance, in an image of an object, global features can capture the shape and texture of the object, such as the fur color of an animal, whereas local features describe specific details of local areas, such as the texture of leaves in the background. Recent empirical studies on self-supervised pretraining with transformers (Park et al., 2023; Wei et al., 2022b) and observations in Figure 1 collectively show that masked pretraining exhibits the capacity to avoid attention collapse concentrating towards those global shapes by identifying diverse local attention patterns. Consequently, unraveling their mechanisms necessitates a thorough examination of data characteristics that embody both global and local features. In this paper, we characterize these two types of features by the following assumption.

Assumption 2.4 (Global feature vs local feature). Let \mathcal{D}_k with $k \in [K]$ be a cluster from \mathcal{D} . We let $\mathcal{P}_{k,1}$ be the **global area** of cluster \mathcal{D}_k , and all the other areas $\mathcal{P}_{k,j}$, $j \in [N_k] \setminus \{1\}$ be the **local areas**. Since each area corresponds to an assigned feature, we also call them the global and local features, respectively. Moreover, we assume:

- Global area: given $k \in [K]$, we set $C_{k,1} = |\mathcal{P}_{k,1}| = \Theta(P^{\kappa_c})$ with $\kappa_c \in [0.5005, 1]$, where $C_{k,1}$ is the number of patches in the global area $\mathcal{P}_{k,1}$.
- Local area: given $k \in [K]$, we choose $C_{k,j} = \Theta(P^{\kappa_s})$ with $\kappa_s \in [0.001, 0.5]$ for $j > 1$, where $C_{k,j}$ denotes the number of patches in the local area $\mathcal{P}_{k,j}$.

The rationale for defining the global feature in this man-

¹The distribution of $z_j(X)$ can be arbitrary within the above support set.




Figure 3: The mechanism of how the masked patch attends to other patches via attention correlations.

ner stems from observing that patches representing global features ($C_{k,1}$) typically occur more frequently than those representing local features ($C_{k,j}$, for $j > 1$), since global features capture the primary visual information of an image, offering a dominant view, while local features focus on subtler details within the image.

3 Statements of Main Results

In this section, we present the main theorems of this paper. Our results are structured into two parts: *i*). analysis of convergence, which includes the global convergence guarantee of the masked reconstruction loss and characterization of the attention pattern at the end of training to demonstrate the diverse locality; *ii*). learning dynamics of attention correlations, which shows how transformers capture target Feature-Position correlations while downplaying Position-Position correlations as discussed in Appendix A.

Notations for theorem presentations. We introduce a notion of *information gap* to quantify the difference of significance between global and local areas (cf. Assumption 2.4).

$$\Delta \triangleq (1 - \kappa_s) - 2(1 - \kappa_c). \quad (6)$$

We define the *unmasked area attention* as follows:

$$\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}(M(X); Q) \triangleq \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,m}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}(X; Q).$$

Moreover, we focus on *patch-level reconstruction loss*:

$$\mathcal{L}_{\mathbf{p}}(Q) = \frac{1}{2} \mathbb{E} \left[\mathbb{1}_{\{\mathbf{p} \in \mathcal{M}\}} \left\| [F(M(X); Q)]_{\mathbf{p}} - X_{\mathbf{p}} \right\|^2 \right]. \quad (7)$$

Attention correlations: Let $\mathbf{p} \in \mathcal{P}$, and we define two types of attention correlations as:

1. Feature-Position (FP): $\Phi_{\mathbf{p} \rightarrow v_{k,m}} \triangleq e_{\mathbf{p}}^{\top} Q v_{k,m}$, $k \in [K]$ and $m \in [N_k]$;
2. Position-Position (PP): $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}} \triangleq e_{\mathbf{p}}^{\top} Q e_{\mathbf{q}}$, $\forall \mathbf{q} \in \mathcal{P}$.

Now we present our first main result regarding the convergence of masked pretraining.

Theorem 3.1 (Training convergence). *Suppose the information gap $\Delta \in [-0.5, -\Omega(1)] \cup [\Omega(1), 1]$. Under some mild assumptions², given any $0 < \epsilon < 1$, for each patch $\mathbf{p} \in \mathcal{P}$,*

²see Appendix K

1. *Loss converges:* $\mathcal{L}_{\mathbf{p}}(Q^{(T^*)}) - \mathcal{L}_{\mathbf{p}}^* \leq \epsilon$ in $T^* = O\left(\frac{1}{\eta} \log(P) P^{\max\{2(\frac{U}{L}-1), 1\}(1-\kappa_s)} + \frac{1}{\eta\epsilon} \log\left(\frac{P}{\epsilon}\right)\right)$ iterations, where $\mathcal{L}_{\mathbf{p}}^*$ is the global minimum of (7).
2. **Area-wide pattern of attention:** given cluster $k \in [K]$, and $\mathbf{p} \in \mathcal{P}_{k,j}$ for some $j \in [N_k]$, if $X_{\mathbf{p}}$ is masked, then the one-layer transformer nearly “pays all attention” to all unmasked patches in the same area $\mathcal{P}_{k,j}$ as \mathbf{p} , i.e.,

$$\left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,j}}^{(T^*)}\right)^2 \leq O(\epsilon).$$

The location of the patch \mathbf{p} determines the above area-wide attention and can be achieved no matter if \mathbf{p} belongs to the global or local areas, which jointly highlight the **diverse local patterns** for masked vision pretraining.

Next, we detail the training phases of attention correlations in the following theorem, which explicitly confirms that the model **learns** target FP correlations while **ignoring** PP correlations to achieve the desirable area-wide attention patterns (see Appendix A and B for intuitive explanations).

Theorem 3.2 (Learning Feature-Position correlations). *Following the same assumptions in Theorem 3.1, for $\mathbf{p} \in \mathcal{P}$, given $k \in [K]$, if $\mathbf{p} \in \mathcal{P}_{k,j}$ for some $j \in [N_k]$, we have*

For positive information gap $\Delta \in [\Omega(1), 1]$:

- a. **Global areas** ($j = 1$) learn FP correlation in **one-phase**: $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ monotonically increases during training with all other attention correlations nearly unchanged.
- b. **Local areas** ($j > 1$) learn FP correlation in **two-phase**: *phase I:* FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ between local area and the global area feature quickly decreases whereas all other attention correlations stay close to zero; *phase II:* FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,j}}^{(t)}$ for the target local area starts to grow until convergence with all other attention correlations nearly unchanged.

For negative information gap $\Delta \in [-0.5, -\Omega(1)]$:

- c. **All areas** learn FP correlation through **one-phase**: $\Phi_{\mathbf{p} \rightarrow v_{k,j}}^{(t)}$ monotonically increases throughout the training, with all other attention correlations close to 0.

Typically, the target FP correlations are learned directly in a single phase. However, for a *positive* information gap

Δ , when patch \mathbf{p} is located in a *local* area, the learning process contains an additional decoupling phase, to reduce the FP correlation with the non-target global features (see Appendix B for proof sketch).

References

- Ahn, K., Cheng, X., Daneshmand, H., and Sra, S. Transformers learn to implement preconditioned gradient descent for in-context learning. *arXiv preprint arXiv:2306.00297*, 2023.
- Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- Bhojanapalli, S., Chakrabarti, A., Glasner, D., Li, D., Unterthiner, T., and Veit, A. Understanding robustness of transformers for image classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10231–10241, 2021.
- Cao, S., Xu, P., and Clifton, D. A. How to understand masked autoencoders. *arXiv preprint arXiv:2202.03670*, 2022.
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., and Joulin, A. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
- Chen, S. and Li, Y. Provably learning a multi-head attention layer. *arXiv preprint arXiv:2402.04084*, 2024.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Chen, T., Luo, C., and Li, L. Intriguing properties of contrastive losses. *Advances in Neural Information Processing Systems*, 34:11834–11845, 2021a.
- Chen, X., Xie, S., and He, K. An empirical study of training self-supervised vision transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9620–9629, 2021b. URL <https://api.semanticscholar.org/CorpusID:233024948>.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Edelman, B. L., Goel, S., Kakade, S., and Zhang, C. Inductive biases and variable creation in self-attention mechanisms. In *International Conference on Machine Learning*, pp. 5793–5831. PMLR, 2022.
- Ghiasi, A., Kazemi, H., Borgnia, E., Reich, S., Shu, M., Goldblum, M., Wilson, A. G., and Goldstein, T. What do vision transformers learn? a visual exploration. *arXiv preprint arXiv:2212.06727*, 2022.
- Giannou, A., Rajput, S., Sohn, J.-y., Lee, K., Lee, J. D., and Papailiopoulos, D. Looped transformers as programmable computers. *arXiv preprint arXiv:2301.13196*, 2023.
- Greene, E. and Wellner, J. A. Exponential bounds for the hypergeometric distribution. *Bernoulli: official journal of the Bernoulli Society for Mathematical Statistics and Probability*, 23(3):1911, 2017.
- Grill, J.-B., Strub, F., Althé, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- Hahn, M. Theoretical limitations of self-attention in neural sequence models. *Transactions of the Association for Computational Linguistics*, 8:156–171, 2020.
- HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. Provable guarantees for self-supervised deep learning with spectral contrastive loss. In *Advances in Neural Information Processing Systems*, volume 34, pp. 5000–5011, 2021.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009, 2022.
- Huang, Y., Cheng, Y., and Liang, Y. In-context convergence of transformers. *arXiv preprint arXiv:2310.05249*, 2023.
- Jelassi, S., Sander, M., and Li, Y. Vision transformers provably learn spatial structure. In *Advances in Neural*

- Information Processing Systems*, volume 35, pp. 37822–37836, 2022.
- Lee, J. D., Lei, Q., Saunshi, N., and Zhuo, J. Predicting what you already know helps: Provable self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 34, pp. 309–323, 2021.
- Li, H., Wang, M., Liu, S., and Chen, P.-Y. A theoretical understanding of shallow vision transformers: Learning, generalization, and sample complexity. *arXiv preprint arXiv:2302.06015*, 2023a.
- Li, H., Wang, M., Ma, T., Liu, S., Zhang, Z., and Chen, P.-Y. What improves the generalization of graph transformer? a theoretical dive into self-attention and positional encoding. In *NeurIPS 2023 Workshop: New Frontiers in Graph Learning*, 2023b.
- Li, H., Wang, M., Lu, S., Cui, X., and Chen, P.-Y. Training nonlinear transformers for efficient in-context learning: A theoretical learning and generalization analysis, 2024.
- Liu, B., Hsu, D. J., Ravikumar, P., and Risteski, A. Masked prediction: A parameter identifiability view. In *Advances in Neural Information Processing Systems*, pp. 21241–21254, 2022.
- Melas-Kyriazi, L. Do you even need attention? a stack of feed-forward layers does surprisingly well on imagenet. *arXiv preprint arXiv:2105.02723*, 2021.
- Nichani, E., Damian, A., and Lee, J. D. How transformers learn causal structure with gradient descent. *arXiv preprint arXiv:2402.14735*, 2024.
- Pan, J., Zhou, P., and Shuicheng, Y. Towards understanding why mask reconstruction pretraining helps in downstream tasks. In *The Eleventh International Conference on Learning Representations*, 2022.
- Park, N. and Kim, S. How do vision transformers work? *arXiv preprint arXiv:2202.06709*, 2022.
- Park, N., Kim, W., Heo, B., Kim, T., and Yun, S. What do self-supervised vision transformers learn? In *The Eleventh International Conference on Learning Representations*, 2023.
- Paul, S. and Chen, P.-Y. Vision transformers are robust learners. In *Proceedings of the AAAI conference on Artificial Intelligence*, volume 36, pp. 2071–2081, 2022.
- Pokle, A., Tian, J., Li, Y., and Risteski, A. Contrasting the landscape of contrastive and non-contrastive learning. *arXiv preprint arXiv:2203.15702*, 2022.
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al. Improving language understanding by generative pre-training. 2018.
- Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C., and Dosovitskiy, A. Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems*, 34:12116–12128, 2021.
- Robinson, J., Sun, L., Yu, K., Batmanghelich, K., Jegelka, S., and Sra, S. Can contrastive learning avoid shortcut solutions? *Advances in neural information processing systems*, 34:4974–4986, 2021.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115: 211–252, 2015.
- Sanford, C., Hsu, D., and Telgarsky, M. Transformers, parallel computation, and logarithmic depth. *arXiv preprint arXiv:2402.09268*, 2024a.
- Sanford, C., Hsu, D. J., and Telgarsky, M. Representational strengths and limitations of transformers. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Tarzanagh, D. A., Li, Y., Thrampoulidis, C., and Oymak, S. Transformers as support vector machines. *arXiv preprint arXiv:2308.16898*, 2023a.
- Tarzanagh, D. A., Li, Y., Zhang, X., and Oymak, S. Max-margin token selection in attention mechanism. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b. URL <https://openreview.net/forum?id=WXc8O8ghLH>.
- Tian, Y., Chen, X., and Ganguli, S. Understanding self-supervised learning dynamics without contrastive pairs. In *International Conference on Machine Learning*, pp. 10268–10278. PMLR, 2021.
- Tian, Y., Wang, Y., Chen, B., and Du, S. Scan and snap: Understanding training dynamics and token composition in 1-layer transformer. *arXiv preprint arXiv:2305.16380*, 2023.
- Tian, Y., Wang, Y., Zhang, Z., Chen, B., and Du, S. S. JoMA: Demystifying multilayer transformers via joint dynamics of MLP and attention. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=LbJqRGNYCf>.

- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., and Jégou, H. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pp. 10347–10357. PMLR, 2021.
- Trockman, A. and Kolter, J. Z. Patches are all you need? *arXiv preprint arXiv:2201.09792*, 2022.
- Vasudeva, B., Deora, P., and Thrampoulidis, C. Implicit bias and fast convergence rates for self-attention. *arXiv preprint arXiv:2402.05738*, 2024.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Von Oswald, J., Niklasson, E., Randazzo, E., Sacramento, J., Mordvintsev, A., Zhmoginov, A., and Vladymyrov, M. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pp. 35151–35174. PMLR, 2023.
- Vuckovic, J., Baratin, A., and Combes, R. T. d. A mathematical theory of attention. *arXiv preprint arXiv:2007.02876*, 2020.
- Wang, X., Chen, X., Du, S. S., and Tian, Y. Towards demystifying representation learning with non-contrastive self-supervision. *arXiv preprint arXiv:2110.04947*, 2021.
- Wei, C., Xie, S. M., and Ma, T. Why do pretrained language models help in downstream tasks? an analysis of head and prompt tuning. *Advances in Neural Information Processing Systems*, 34:16158–16170, 2021.
- Wei, C., Chen, Y., and Ma, T. Statistically meaningful approximation: a case study on approximating turing machines with transformers. *Advances in Neural Information Processing Systems*, 35:12071–12083, 2022a.
- Wei, Y., Hu, H., Xie, Z., Zhang, Z., Cao, Y., Bao, J., Chen, D., and Guo, B. Contrastive learning rivals masked image modeling in fine-tuning via feature distillation. *arXiv preprint arXiv:2205.14141*, 2022b.
- Weiss, G., Goldberg, Y., and Yahav, E. Thinking like transformers. In *International Conference on Machine Learning*, pp. 11080–11090. PMLR, 2021.
- Wen, Z. and Li, Y. Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*, pp. 11112–11122. PMLR, 2021.
- Wen, Z. and Li, Y. The mechanism of prediction head in non-contrastive self-supervised learning. In *Advances in Neural Information Processing Systems*, pp. 24794–24809, 2022.
- Xie, Z., Zhang, Z., Cao, Y., Lin, Y., Bao, J., Yao, Z., Dai, Q., and Hu, H. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9653–9663, 2022.
- Xie, Z., Geng, Z., Hu, J., Zhang, Z., Hu, H., and Cao, Y. Revealing the dark secrets of masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14475–14485, 2023.
- Yun, C., Bhojanapalli, S., Rawat, A. S., Reddi, S. J., and Kumar, S. Are transformers universal approximators of sequence-to-sequence functions? *arXiv preprint arXiv:1912.10077*, 2019.
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pp. 12310–12320. PMLR, 2021.
- Zhang, Q., Wang, Y., and Wang, Y. How mask matters: Towards theoretical understandings of masked autoencoders. *Advances in Neural Information Processing Systems*, 35: 27127–27139, 2022.
- Zhang, R., Frei, S., and Bartlett, P. L. Trained transformers learn linear models in-context. *arXiv preprint arXiv:2306.09927*, 2023a.
- Zhang, Y., Zhang, F., Yang, Z., and Wang, Z. What and how does in-context learning learn? bayesian model averaging, parameterization, and generalization. *arXiv preprint arXiv:2305.19420*, 2023b.

A Attention Patterns and Feature-Position Correlations

To show the significance of the data distribution design and understand the nature of the masked reconstruction task, in this section, we will provide some preliminary implications of the spatial structures in Definition 2.3. Intuitively, for a given cluster \mathcal{D}_k , to reconstruct a missing patch $\mathbf{p} \in \mathcal{P}_{k,j} \cap \mathcal{M}$, the attention head should exploit all *unmasked* patches in the *target* area $\mathcal{P}_{k,j}$ to find the same visual feature $v_{k,j}$ to fill in the blank, which emphasizes the *locality* for \mathbf{p} in different areas. This approach remains effective even at high masking rates, provided that a small number of patches in $\mathcal{P}_{k,j}$ remain unmasked. We will elaborate this point by describing the *area attentions* and illustrating the intuition about how they can be learned via *attention correlations* (Definition A.1).

Area attention. We first define a new notation for a cleaner presentation. For $X \sim \mathcal{D}$ and $\mathbf{p} \in \mathcal{P}$, we write the attention of patch $X_{\mathbf{p}}$ to a subset $\mathcal{A} \subset \mathcal{P}$ of patches by

$$\widetilde{\text{Attn}}_{\mathbf{p} \rightarrow \mathcal{A}}(X; Q) := \sum_{\mathbf{q} \in \mathcal{A}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}(X; Q).$$

Let us explain why the above notion of area attention matters in understanding how attention works in masked reconstruction. Suppose we have a sample X picked from \mathcal{D}_k , and the patch $X_{\mathbf{p}}$ with $\mathbf{p} \in \mathcal{P}_{k,j}$ is masked. Then the prediction of $X_{\mathbf{p}}$ given masked input $M(X)$ can be written as

$$\begin{aligned} [F(M(X); Q)]_{\mathbf{p}} &= \sum_{\mathbf{q} \in \mathcal{P}} M(X)_{\mathbf{q}} \cdot \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}(M(X); Q) \\ &= \sum_{i \in [N_k]} z_i(X) v_{k,i} \cdot \widetilde{\text{Attn}}_{\mathbf{p} \rightarrow \mathcal{U} \cap \mathcal{P}_{k,i}}(M(X); Q) \quad (\text{since } M(X)_{\mathbf{q}} = \mathbf{0} \text{ if } \mathbf{q} \in \mathcal{M}). \end{aligned}$$

To reconstruct the original patch $X_{\mathbf{p}}$, the transformer should not only focus on the correct area $\mathcal{P}_{k,j}$, but must also prioritize attention to the *unmasked* patches within this area. This specificity is denoted by the area attention $\widetilde{\text{Attn}}_{\mathbf{p} \rightarrow \mathcal{U} \cap \mathcal{P}_{k,j}}$ over $\mathcal{U} \cap \mathcal{P}_{k,j}$, a requirement imposed by masking operations. To further explain how transformers perform such prioritization, we introduce the following quantities, which capture the major insights of our analysis that differentiate from those in (Jelassi et al., 2022).

Definition A.1. (Attention correlations) Let $\mathbf{p} \in \mathcal{P}$, and we define attention correlations as:

1. Feature-Position (FP) Correlation: $\Phi_{\mathbf{p} \rightarrow v_{k,m}} := e_{\mathbf{p}}^{\top} Q v_{k,m}$, for $k \in [K]$ and $m \in [N_k]$;
2. Position-Position (PP) Correlation: $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}} := e_{\mathbf{p}}^{\top} Q e_{\mathbf{q}}$, $\forall \mathbf{q} \in \mathcal{P}$.

Due to our (zero) initialization of $Q^{(0)}$, we have $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(0)} = \Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(0)} = 0$.

These two types of attention correlations, FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,m}}$ and PP correlation $\Phi_{\mathbf{p} \rightarrow \mathbf{q}}$, act as the exponent terms within the softmax calculations for attention scores. Given $\mathbf{p} \in \mathcal{P}_{k,j}$ is masked, the (unnormalized) attention $\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}$ directed towards an *unmasked* patch \mathbf{q} is influenced jointly by these correlations. Hence, the described attention pattern can emerge from either a substantial FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,j}}$ or a significant PP correlation $\Phi_{\mathbf{p} \rightarrow \mathbf{q}}$ for \mathbf{q} in the same area as \mathbf{p} . However, in our setting, the latter mechanism—learning via PP correlation—fails to produce desired attention patterns, as explained below on a notable gap that prior work (Jelassi et al., 2022) did not address.

Can pure positional attention explain the transformer’s ability to learn locality? (Jelassi et al., 2022) theoretically analyzed how ViTs can identify spatially localized patterns by minimizing the supervised cross-entropy loss with GD. Their analysis focused on a spatially structured dataset equivalent to our settings when $K = 1$ without distinguishing the global and local features. Hence, the association between patch \mathbf{p} and \mathbf{q} is fixed since there is only a single cluster. Based on such invariant patch-wise relations, they showed that ViTs can achieve optimal attention patterns by solely learning the “patch association”, where $e_{\mathbf{p}}^{\top} Q e_{\mathbf{q}}$ is large for \mathbf{p}, \mathbf{q} coming from the same area, corresponding to large PP correlation $\Gamma_{\mathbf{p} \rightarrow \mathbf{q}}$ in our settings. However, such an assumption of invariant patch associations is often unrealistic for vision datasets in practice, for instance, a cube-shaped building typically requires a different attention pattern from a bird inside the woods. Therefore, when multiple patterns appear in the data distribution (e.g., in our settings with $K > 1$ clusters), relying solely on PP correlations is insufficient, and is often undesirable due to variations of PP correlations across clusters. This highlights the necessity of examining FP correlations for a deeper understanding of the local representation power of transformers.

B Proof Sketch

The framework for constructing attention correlations, as outlined in Appendix A, offers insights into how transformers may achieve specific localities during masked pretraining. Within this framework, we start this section with a warm-up

discussion to further clarify why target FP correlations, rather than PP correlations, should be the focus of learning. We then offer an overview of our proof, which analyzes the dynamics of these attention correlations.

Warm-up intuition for attention correlations. Let us start by exploring the role of attention correlations in determining the area attention scores. Given a masked input $M(X)$ with $\mathbf{p} \in \mathcal{P}_{k,j} \cap \mathcal{M}$, it holds that

$$\widetilde{\text{Attn}}_{\mathbf{p} \rightarrow \mathcal{P}_{k,j}}(M(X); Q) \propto \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,j}} e^{\Phi_{\mathbf{p} \rightarrow v_{k,j}} + \Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}} + \sum_{\mathbf{q} \in \mathcal{M} \cap \mathcal{P}_{k,j}} e^{\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}}.$$

The first term on the RHS is proportional to our unmasked area attention $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,j}}$, which is balanced by the relative magnitudes between Feature-Position and Position-Position correlations. Upon initial observation, large FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,j}}$ or PP correlation $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}$ for patch $\mathbf{q} \in \mathcal{P}_{k,j}$ may both contribute to the attention towards the area $\mathcal{U} \cap \mathcal{P}_{k,j}$. However, two key issues can arise from the PP correlation, which jointly lead to the suppression of PP correlation for reconstruction: *i*). such a mechanism inadvertently directs attention to the *masked* patches, which is not desirable; *ii*). such position association could be vulnerable to the variation across different clusters, i.e., $\mathbf{p}, \mathbf{q} \in \mathcal{P}_{k,j}$ does not necessarily hold for all $k \in [K]$. Consequently, as illustrated in Figure 3, to obtain a perfect reconstruction of the masked patch, learning a relatively large FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,j}}$ from the patches to the correct features is required.

Main idea of the proof. The main idea of our analysis is to track the dynamics of those attention correlations. We first provide the following lemma of GD updates for $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ and $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$.

Lemma B.1 (Gradient of attention correlations, informal). *Given $\mathbf{p}, \mathbf{q} \in \mathcal{P}$, let $\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} = \frac{1}{\eta} (\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t+1)} - \Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)})$ with $k \in [K]$, $m \in [N_k]$, and $\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \frac{1}{\eta} (\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t+1)} - \Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)})$. We use $a_{k,\mathbf{p}}$ to indicate that the index of the area that patch \mathbf{p} is located in the cluster \mathcal{D}_k .*

- For the same area, $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)} \approx \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(t)} \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(t)}\right)^2$;
- If $a_{k,\mathbf{p}} \neq 1$, for the global area, $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \approx -\left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}\right)^2$;
- For other area $m \notin \{a_{k,\mathbf{p}}\} \cup \{1\}$, $|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)} + \mathbb{1}\{n \neq 1\} |\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{N_k}\right)$;
- $\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \sum_{k \in [K]} \beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$, where $|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{|\alpha_{k,\mathbf{p} \rightarrow a_{k,\mathbf{q}}}^{(t)}|}{C_{k,a_{k,\mathbf{q}}}}\right)$.

Here, $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ can be interpreted as the projected gradient for PP correlation $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ on the k -th cluster. Thus, Lemma B.1.d directly implies that the overall increment of PP correlation $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ should be negligible compared to the corresponding FP correlations $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{q}}}}$ on certain cluster \mathcal{D}_k since $C_{k,a_{k,\mathbf{q}}} \gg 1$ and $K = \Theta(1)$, which implies that **all PP correlations are small**.

Moreover, from Lemma B.1, it is observed that the target FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ exhibits a monotonically increasing trend, while FP-correlations for feature $v_{k,1}$ will monotonically decrease if $a_{k,\mathbf{p}} \neq 1$. Moreover, FP gradients of the patch \mathbf{p} to other unrelated local features are significantly smaller and become negligible since $N_k \gg 1$. Based on these observations, we now explain how the trend of FP-correlations determines the different learning behaviors varied across different locations $a_{k,\mathbf{p}}$ and information gap Δ .

- For local areas with $a_{k,\mathbf{p}} > 1$:** at the beginning of training, $\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(0)} = \frac{1}{P}$ for any $\mathbf{q} \in \mathcal{P}$ due to zero initialization. However, with high probability, the count of unmasked patches in the global area $\mathcal{P}_{k,1}$ significantly exceeds that in other areas.

- $\Delta = (1 - \kappa_s) - 2(1 - \kappa_c) \geq \Omega(1)$: $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(0)}| \approx \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(0)}\right)^2 = \Omega\left(\frac{1}{P^{2(1-\kappa_c)}}\right) \gg \Theta\left(\frac{1}{P^{1-\kappa_s}}\right) = \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(0)} \approx \alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(0)}$. Thus, $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ enjoys a much larger decreasing rate initially, which defines the *decoupling phase* of non-target global FP correlations. The significant decrease in $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ leads to a reduction in $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}$, which in turn makes a drop in $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|$. As the gradient decreases below $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$, $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ within the target local area begins to have a larger gradient, which triggers the *growth of target FP correlation phase*, leading to continuous growth until convergence.
- $\Delta \leq -\Omega(1)$: the above relationship reverses, and $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ starts with a fast initial increase, which moves directly into the *growth phase*, eliminating the need to decouple $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$.

- **For global areas with $a_{k,P} = 1$:** the global area is exactly the target area that FP correlations should be learned, and thus the second term in Lemma B.1 will not exist in this case. As a result, the training process also skips the initial decoupling phase and starts with the *growth* phase for $\Phi_{P \rightarrow v_{k,1}}^{(t)}$. Such an effect is independent of the value of Δ .

C Additional Related Work

Empirical studies of transformers in vision. A number of works have aimed to understand the transformers in vision from different perspectives: comparison with CNNs (Raghu et al., 2021; Ghiasi et al., 2022; Park & Kim, 2022), robustness (Bhojanapalli et al., 2021; Paul & Chen, 2022), and role of positional embeddings (Melas-Kyriazi, 2021; Trockman & Kolter, 2022). Recent studies (Xie et al., 2023; Wei et al., 2022b; Park et al., 2023) have delved into ViTs with self-supervision to uncover the mechanisms at play, particularly through visualization and analysis of metrics related to self-attention. (Xie et al., 2023) compared the masked image modeling (MIM) method with supervised models, revealing MIM’s capacity to enhance diversity and locality across all ViT layers, which significantly boosts performance on tasks with weak semantics following fine-tuning. Building on MIM’s advantages, (Wei et al., 2022b) further proposed a simple feature distillation method that incorporates locality into various self-supervised methods, leading to an overall improvement in the finetuning performance. (Park et al., 2023) conducted a detailed comparison between masked image modeling (MIM) and contrastive learning. They demonstrated that contrastive learning will make the self-attentions collapse into homogeneity for all query patches due to the nature of discriminative learning, while MIM leads to a diverse self-attention map since it focuses on local patterns.

Theory of self-supervised learning. A major line of theoretical studies falls into one of the most successful self-supervised learning approaches, contrastive learning (Wen & Li, 2021; Robinson et al., 2021; Chen et al., 2021a; Arora et al., 2019), and its variant non-contrastive self-supervised learning (Wen & Li, 2022; Pokle et al., 2022; Wang et al., 2021). Some other works study the mask prediction approach (Lee et al., 2021; Wei et al., 2021; Liu et al., 2022), which is the focus of this paper. (Lee et al., 2021) provided statistical downstream guarantees for reconstructing missing patches. (Wei et al., 2021) studied the benefits of head and prompt tuning with masked pretraining under a Hidden Markov Model framework. (Liu et al., 2022) provided a parameter identifiability view to understand the benefit of masked prediction tasks, which linked the masked reconstruction tasks to the informativeness of the representation via identifiability techniques from tensor decomposition.

Theory of transformers and attention models. Prior work has studied the theoretical properties of transformers from various aspects: representational power (Yun et al., 2019; Edelman et al., 2022; Vuckovic et al., 2020; Wei et al., 2022a; Sanford et al., 2024a), internal mechanism (Tarzanagh et al., 2023a; Weiss et al., 2021), limitations (Hahn, 2020; Sanford et al., 2024b), and PAC learning (Chen & Li, 2024). Recently, there has been a growing body of research studying in-context learning with transformers due to the remarkable emergent in-context ability of large language models (Zhang et al., 2023b; Von Oswald et al., 2023; Giannou et al., 2023; Ahn et al., 2023; Zhang et al., 2023a; Huang et al., 2023; Nichani et al., 2024; Li et al., 2024). Regarding the training dynamics of attention-based models, (Li et al., 2023a) studied the training process of shallow ViTs in a classification task. Subsequent research expanded on this by exploring the graph transformer with positional encoding (Li et al., 2023b) and in-context learning performance of transformers with nonlinear self-attention and nonlinear MLP (Li et al., 2024). However, all of these analyses rely crucially on stringent assumptions on the initialization of transformers and hardly generalize to our setting. (Tian et al., 2023) mathematically described how the attention map evolves trained by SGD for one-layer transformer but did not provide any convergence guarantee, and the follow-up work (Tian et al., 2024) considered a generalized case with multiple layers. (Tarzanagh et al., 2023b; Vasudeva et al., 2024) investigated the implicit bias for self-attention models trained with GD. Furthermore, (Huang et al., 2023) proved the in-context convergence of a one-layer softmax transformer trained via GD and illustrated the attention dynamics throughout the training process. More recently, (Nichani et al., 2024) studied GD dynamics on a simplified two-layer attention-only transformer and proved that it can encode the causal structure in the first attention layer. However, none of the previous studies analyzed the training of transformers under self-supervised learning, which is the focus of this paper.

D Experiments

Previous studies on the attention mechanisms of ViT-based pre-training approaches have mainly utilized a metric known as the attention distance (Dosovitskiy et al., 2020). Such a metric quantifies the average spatial distance between the query and key tokens, weighted by their self-attention coefficients. The general interpretation is that larger attention distances indicate global understanding, and smaller values suggest a focus on local features. However, such a metric does not adequately determine if the self-attention mechanism is identifying a unique global pattern. A high attention distance could result




Figure 4: **Attention Diversity Metric:** We examined the last layer of ViTs trained by masked reconstruction (MAE), discriminative self-supervised learning (MoCo, DINO), and supervised learning (DeiT). Lower values of this metric signify focused attention on a similar area across different patches, reflecting a global pattern of focus. Conversely, higher values suggest that attention is dispersed, focusing on different, localized areas. The results show that the MAE model excels in capturing diverse local patterns (see Appendix D for more details).

from different patches focusing on varied distant areas, which does not necessarily imply that global information is being effectively synthesized. To address this limitation, we introduce a novel and revised version of average attention distance, called attention diversity metric, which is designed to assess whether various patches are concentrating on a similar region, thereby directly capturing global information.

Attention diversity metric, in distance. This metric is computed for self-attention with a single head of the specific layer. For a given image divided into $P \times P$ patches, the process unfolds as follows: for each patch, it is employed as the query patch to calculate the attention weights towards all P^2 patches, and those with the top- n attention weights are selected. Subsequently, the coordinates (e.g. (i, j) with $i, j \in [P]$) of these top- n patches are concatenated in sequence to form a $2 \times n$ -dimensional vector. The final step computes the average distance between all these $2n$ -dimensional vectors, i.e., $P^2 \times P^2$ vector pairs.

Setup. In this work, we compare the performance of ViT-B/16 encoder pre-trained on ImageNet-1K (Russakovsky et al., 2015) among the following four models: masked reconstruction model (MAE), contrastive learning model (MoCo v3 (Chen et al., 2021b)), other self-supervised model (DINO (Caron et al., 2021)), and supervised model (DeiT (Touvron et al., 2021)). We focus on 12 different attention heads in the last layer of ViT-B on different pre-trained models. The box plot visualizes the distribution of the top-10 averaged attention focus across 152 example images, as similarly done in (Dosovitskiy et al., 2020).

Implications. The experiment results based on our new metric are provided in Figure 4. Lower values of the attention diversity metric signify a focused attention on a coherent area across different patches, reflecting a global pattern of focus. On the other hand, higher values suggest that attention is dispersed, focusing on different, localized areas. It can be seen that the masked pretraining model is particularly effective in learning more diverse attention patterns, setting it apart from other models that prioritize a uniform global information with less attention diversity. This aligns with and provides further evidence for the findings in (Park et al., 2023).

E Overview of the Proof Techniques

In this section, we explain our key proof techniques in analyzing the masked pretraining of transformers. We focus on the reconstruction of a specific patch $X_{\mathbf{p}}$ for $\mathbf{p} \in \mathcal{P}$. We aim to elucidate the training phases through which the model learns FP correlations related to the area associated with \mathbf{p} across different clusters $k \in [K]$.

Our characterization of training phases differentiates between whether $X_{\mathbf{p}}$ is located in the global or local areas and further varies based on whether Δ is positive or negative. Specifically, for $\Delta \in [\Omega(1), 1]$, we observe distinct learning dynamics for FP correlations between local and global areas:

- Local area attends to FP correlation in two-phase: given $k \in [K]$, if $a_{k,\mathbf{p}} \neq 1$, then
 1. $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ first quickly decreases whereas all other $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ with $m \neq 1$ and $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ do not change much;
 2. after some point, the increase of $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ takes dominance. Such $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ will keep growing until convergence with all other FP and PP attention correlations nearly unchanged.
- Global areas learn FP correlation in one-phase: given $k \in [K]$, if $a_{k,\mathbf{p}} = 1$, the update of $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ will dominate throughout the training, whereas all other $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ with $m \neq 1$ and learned PP correlations remain close to 0.

For $\Delta \in [-0.5, -\Omega(1)]$, the behaviors of learning FP correlations are uniform for all areas. Namely, all areas learn FP correlation through one-phase: given $k \in [K]$, throughout the training, the increase of $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ dominates, whereas all other $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ with $m \neq a_{k,\mathbf{p}}$ and PP correlations $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ remain close to 0.

For clarity, this section will mainly focus on the learning of *local* feature correlations with a positive information gap $\Delta \geq \Omega(1)$ in Appendices E.2 and E.3, which exhibits a two-phase process. The other scenarios will be discussed briefly in Appendix E.4.

E.1 GD Dynamics of Attention Correlations

Based on the crucial roles that attention correlations play in determining the reconstruction loss, the main idea of our analysis is to track the dynamics of those attention correlations. We first provide the following GD updates of $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ and $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ (see Appendix F.1.1 for formal statements).

Lemma E.1 (FP correlations, informal). *Given $k \in [K]$, for $\mathbf{p} \in \mathcal{P}$, denote $n = a_{k,\mathbf{p}}$, let $\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} = \frac{1}{\eta} (\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t+1)} - \Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)})$ for $m \in [N_k]$, and suppose $X_{\mathbf{p}}$ is masked. Then*

1. for the same area, $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \approx \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right)^2$;
2. if $k \in \mathcal{B}_{\mathbf{p}}$, for the global area,

$$\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \approx -\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}\right) + \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right) \right);$$

3. for other area $m \notin \{n\} \cup \{1\}$,

$$\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} \approx \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \left(\mathbb{1}\{n \neq 1\} \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}\right)^2 - \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right).$$

From Lemma E.1, it is observed that for $\mathbf{p} \in \mathcal{P}_{k,n}$, the feature correlation $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$ exhibits a monotonically increasing trend over time because $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq 0$. Furthermore, if $n > 1$, i.e., $\mathcal{P}_{k,n}$ is the local area, $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ will monotonically decrease.

Lemma E.2 (PP attention correlations, informal). *Given $\mathbf{p}, \mathbf{q} \in \mathcal{P}$, let $\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \frac{1}{\eta} (\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t+1)} - \Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)})$, and suppose $X_{\mathbf{p}}$ is masked. Then $\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \sum_{k \in [N]} \beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$, where $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ satisfies*

1. if $a_{k,\mathbf{p}} = a_{k,\mathbf{q}} = n$, $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \approx \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right)^2$;
2. if $k \in \mathcal{B}_{\mathbf{p}} \cap \mathcal{C}_{\mathbf{q}}$, where $a_{k,\mathbf{p}} = n > 1$ and $a_{k,\mathbf{q}} = 1$:

$$\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \approx -\mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}\right) + \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right) \right);$$

3. if $a_{k,\mathbf{q}} = m \notin \{n\} \cup \{1\}$, where $a_{k,\mathbf{p}} = n$,

$$\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \approx \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(\mathbb{1}\{n \neq 1\} \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}\right)^2 - \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right).$$

Based on the above gradient update for $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$, we further introduce the following auxiliary quantity $\Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}$, which can be interpreted as the PP attention correlation ‘‘projected’’ on the k -th cluster \mathcal{D}_k , and will be useful in the later proof.

$$\Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t+1)} := \Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)} + \eta \beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}, \quad \text{with } \Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(0)} = 0. \quad (8)$$

We can directly verify that $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \sum_{k \in [K]} \Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}$.

The key observation by comparing Lemma E.1 and E.2 is that the gradient of projected PP attention $\beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ is smaller than the corresponding FP gradient $\alpha_{\mathbf{p} \rightarrow v_{k, a_{k, \mathbf{q}}}}^{(t)}$ in magnitude since $\mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \approx \frac{\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k, a_{k, \mathbf{q}}}}^{(t)}}{(1-\gamma)C_{k, a_{k, \mathbf{q}}}}$. We will show that the interplay between the increase of $\Phi_{\mathbf{p} \rightarrow v_{k, n}}^{(t)}$ and the decrease of $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)}$ determines the learning behaviors for the local patch $\mathbf{p} \in \mathcal{P}_{k, n}$ with $n > 1$, and which effect will happen first depends on the initial attention, which is also determined by the value of information gap Δ .

E.2 Phase I: Decoupling the Global FP Correlations

We now explain how the attention correlations evolve at the initial phase of the training to decouple the correlations of the non-target global features when \mathbf{p} is located in the local area for the k -th cluster. This phase can be further divided into the following two stages.

Stage 1. At the beginning of training, $\Phi_{\mathbf{p} \rightarrow v_{k, m}}^{(0)} = \Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(0)} = 0$, and hence $\mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(0)} = \frac{1}{P}$ for any $\mathbf{q} \in \mathcal{P}$, which implies that the transformer equally attends to each patch. However, with high probability, the number of unmasked global features in the global area $\mathcal{P}_{k, 1}$ is much larger than others. Hence, $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k, 1}}^{(0)} = \frac{|\mathcal{U} \cap \mathcal{P}_{k, 1}|}{P} \geq \Omega\left(\frac{1}{P^{1-\kappa_c}}\right) \gg \Theta\left(\frac{1}{P^{1-\kappa_s}}\right) = \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k, m}}^{(0)}$ for $m > 1$. Therefore, by Lemma E.1 and E.2, we immediately obtain

- $\alpha_{\mathbf{p} \rightarrow v_{k, 1}}^{(0)} = -\Theta\left(\frac{1}{P^{2(1-\kappa_c)}}\right)$, whereas $\alpha_{\mathbf{p} \rightarrow v_{k, a_{k, \mathbf{p}}}}^{(0)} = \Theta\left(\frac{1}{P^{(1-\kappa_s)}}\right)$;
- all other FP correlation gradients $\alpha_{\mathbf{p} \rightarrow v_{k, m}}^{(0)}$ with $m \neq 1, a_{k, \mathbf{p}}$ are small;
- all projected PP correlation gradients $\beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(0)}$ are small.

Since $\Delta = (1 - \kappa_s) - 2(1 - \kappa_c) \geq \Omega(1)$, it can be seen that $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)}$ enjoys a much larger decreasing rate initially. This captures the decoupling process of the feature correlations with the global feature $v_{k, 1}$ in the global area for \mathbf{p} . It can be shown that such an effect will dominate over a certain period that defines stage 1 of phase I. At the end of this stage, we will have $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)} \leq -\Omega(\log(P))$, whereas all FP attention correlation $\Phi_{\mathbf{p} \rightarrow v_{k, m}}^{(t)}$ with $m > 1$ and all projected PP correlations $\Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ stay close to 0 (see Appendix H.1).

During stage 1, the significant decrease of the global FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)}$ leads to a reduction in the attention score $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k, 1}}^{(t)}$. Meanwhile, attention scores $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k, m}}^{(t)}$ (where $m > 1$) for other patches remain consistent, reflecting a uniform distribution over unmasked patches within each area. By the end of stage 1, $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k, 1}}^{(t)}$ drops to a certain level, resulting in a decrease in $|\alpha_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)}|$ as it approaches $\alpha_{\mathbf{p} \rightarrow v_{k, n}}^{(t)}$, which indicates that stage 2 begins.

Stage 2. Soon as stage 2 begins, the dominant effect switches as $|\alpha_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)}|$ reaches the same order of magnitude as $\alpha_{\mathbf{p} \rightarrow v_{k, a_{k, \mathbf{p}}}}^{(t)}$. The following result shows that $\Phi_{\mathbf{p} \rightarrow v_{k, a_{k, \mathbf{p}}}}^{(t)}$ must update during stage 2.

Lemma E.3 (Switching of dominant effects (See Appendix H.2)). *Under the same conditions as Theorem 3.1, for $\mathbf{p} \in \mathcal{P}$, there exists \tilde{T}_1 , such that at iteration $t = \tilde{T}_1 + 1$, we have*

- $\Phi_{\mathbf{p} \rightarrow v_{k, a_{k, \mathbf{p}}}}^{(\tilde{T}_1+1)} \geq \Omega(\log(P))$, and $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(\tilde{T}_1+1)} = -\Theta(\log(P))$;
- all other FP correlations $\Phi_{\mathbf{p} \rightarrow v_{k, m}}^{(t)}$ with $m \neq 1, a_{k, \mathbf{p}}$ are small;
- all projected PP correlations $\Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ are small.

Intuition of the transition. Once $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ decreases to $-\frac{\Delta}{2L} \log(P)$, we observe that $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|$ is approximately equal to $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$. After this point, reducing $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ further is more challenging compared to the increase in $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$.

To illustrate, a minimal decrease of $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ by an amount of $\frac{0.001}{L} \log(P)$ will yield $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| \leq O(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{0.002}})$. Such a discrepancy triggers the switch of the dominant effect.

E.3 Phase II: Growth of Target Local FP Correlation

Moving beyond phase I, FP correlation $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ within the target local area \mathbf{p} already enjoys a larger gradient $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ than other $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ with $m \neq a_{k,\mathbf{p}}$ and all projected PP correlations $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$. We can show that the growth of $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ will continue to dominate until the end of training by recognizing the following two stages.

Rapid growth stage. At the beginning of phase II, $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ is mainly driven by $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(t)}$ since $1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(t)}$ remains at the constant order. Therefore, the growth of $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ naturally results in a boost in $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(t)}$, thereby promoting an increase in its own gradient $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$, which defines the rapid growth stage. On the other hand, we can prove that the following gap holds for FP and projected PP correlation gradients (see Appendix H.3):

- all other FP correlation gradients $\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ with $m \neq a_{k,\mathbf{p}}$ are small;
- all projected PP correlation gradients $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ are small.

Convergence stage. After the rapid growth stage, the desired local pattern with a high target feature-position correlation $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ is learned. In this last stage, it is demonstrated that the above conditions for non-target FP and projected PP correlations remain valid, while the growth of $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ starts to decelerate as $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ reaches $\Theta(\log(P))$, resulting in $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \approx \Omega(1)$, which leads to convergence (see Appendix H.4).

E.4 Learning Processes in Other Scenarios

In this section, we talk about the learning process in other settings, including learning FP correlations for the local area when the information gap is negative, learning FP correlations for the global area, and failure to learn PP correlations.

What is the role of positive information gap? As described in stage 1 of phase 1 in Appendix E.2, the decoupling effect happens at the beginning of the training because $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(0)} \gg \alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(0)}$ attributed to $\Delta \geq \Omega(1)$. However, in cases where $\Delta \leq -\Omega(1)$, this relationship reverses, with $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(0)}$ becoming significantly smaller than $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(0)}$. Similarly, other FP gradients $\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(0)}$ with $m \neq 1, a_{k,\mathbf{p}}$ and all the projected gradients of PP correlation $\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(0)}$ are small in magnitude. Consequently, $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ starts with a larger gradient, eliminating the need to decouple FP correlations for the global area. As a result, training skips the initial phase, and moves directly into Phase II, during which $\Phi_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(t)}$ continues to increase until it converges (see Appendix I).

Learning FP correlations for the global area. When the patch $X_{\mathbf{p}}$ is located in the global area of cluster k , i.e., $a_{k,\mathbf{p}} = 1$, the attention score $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(0)}$ directed towards the target area $\mathcal{P}_{k,1}$ is initially higher compared to other attention scores due to the presence of a significant number of unmasked patches in the global area. This leads to an initially larger gradient $\alpha_{\mathbf{p} \rightarrow v_{k,a_{k,\mathbf{p}}}}^{(0)}$. Such an effect is independent of the value of Δ . As a result, the training process skips the initial phase, which is typically necessary for the cases where $a_{k,\mathbf{p}} > 1$ with a positive information gap, and moves directly into Phase II (see Appendix J).

All PP correlations are small. Integrating the analysis from all previous discussions, we establish that for every cluster $k \in [K]$, regardless of its association with $\mathcal{C}_{\mathbf{p}}$ (global area) or $\mathcal{B}_{\mathbf{p}}$ (local area), and for any patch $X_{\mathbf{q}}$ with $\mathbf{q} \in \mathcal{P}$, the projected PP correlation $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ remains nearly zero in comparison to the significant changes observed in the FP correlation, because the gradient $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ is relatively negligible. Therefore, the overall PP correlation $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \sum_{k=1}^K \Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ also stays close to zero, given that the number of clusters $K = \Theta(1)$.

F Preliminaries

In this section, we will introduce warm-up gradient computations and probabilistic lemmas that establish essential properties of the data and the loss function, which are pivotal for the technical proofs in the upcoming sections. Throughout the appendix, we assume $N_k = N$ and $C_{k,n} = C_n$ for all $k \in [K]$ for simplicity. We will also omit the explicit dependence on X for $z_n(X)$. We use $k_X \in [K]$ to denote the cluster index that a given image X is drawn from.

F.1 Gradient Computations

We first calculate the gradient with respect to Q . We omit the superscript ‘(t)’ and write $\mathcal{L}(Q)$ as \mathcal{L} here for simplicity.

Lemma F.1. *The gradient of the loss function with respect to Q is given by*

$$\frac{\partial \mathcal{L}}{\partial Q} = -\mathbb{E} \left[\sum_{\mathbf{p} \in \mathcal{M}} \sum_{\mathbf{q}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \mathbf{M}(X)_{\mathbf{q}}^{\top} (X_{\mathbf{p}} - [F(\mathbf{M}(X); Q)]_{\mathbf{p}}) \cdot \tilde{\mathbf{M}}(X)_{\mathbf{p}} \left(\tilde{\mathbf{M}}(X)_{\mathbf{q}} - \sum_{\mathbf{r}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}} \tilde{\mathbf{M}}(X)_{\mathbf{r}} \right)^{\top} \right].$$

Proof. We begin with the chain rule and obtain

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial Q} &= \mathbb{E} \left[\sum_{\mathbf{p} \in \mathcal{M}} \frac{\partial [F(\mathbf{M}(X); Q)]_{\mathbf{p}}}{\partial Q} ([F(\mathbf{M}(X); Q)]_{\mathbf{p}} - X_{\mathbf{p}}) \right] \\ &= \mathbb{E} \left[\sum_{\mathbf{p} \in \mathcal{M}} \sum_{\mathbf{q}} \frac{\partial \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}}{\partial Q} \mathbf{M}(X)_{\mathbf{q}}^{\top} ([F(\mathbf{M}(X); Q)]_{\mathbf{p}} - X_{\mathbf{p}}) \right]. \end{aligned} \quad (9)$$

We focus on the gradient for each attention score:

$$\begin{aligned} \frac{\partial \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}}{\partial Q} &= \sum_{\mathbf{r}} \frac{\exp \left(\tilde{\mathbf{M}}(X)_{\mathbf{p}}^{\top} Q (\tilde{\mathbf{M}}(X)_{\mathbf{r}} + \tilde{\mathbf{M}}(X)_{\mathbf{q}}) \right)}{\left(\sum_{\mathbf{r}} \exp(\tilde{\mathbf{M}}(X)_{\mathbf{p}}^{\top} Q \tilde{\mathbf{M}}(X)_{\mathbf{r}}) \right)^2} \tilde{\mathbf{M}}(X)_{\mathbf{p}} (\tilde{\mathbf{M}}(X)_{\mathbf{q}} - \tilde{\mathbf{M}}(X)_{\mathbf{r}})^{\top} \\ &= \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \sum_{\mathbf{r}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}} \tilde{\mathbf{M}}(X)_{\mathbf{p}} (\tilde{\mathbf{M}}(X)_{\mathbf{q}} - \tilde{\mathbf{M}}(X)_{\mathbf{r}})^{\top} \\ &= \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \tilde{\mathbf{M}}(X)_{\mathbf{p}} \cdot \left[\tilde{\mathbf{M}}(X)_{\mathbf{q}} - \sum_{\mathbf{r}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}} \tilde{\mathbf{M}}(X)_{\mathbf{r}} \right]^{\top}. \end{aligned}$$

Substituting the above equation into (9), we complete the proof. \square

Recall that the quantities $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ and $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ are defined in Definition A.1. These quantities are associated with the attention weights for each token, and they play a crucial role in our analysis of learning dynamics. We will restate their definitions here for clarity.

Definition F.2. (Attention correlations) Given $\mathbf{p}, \mathbf{q} \in \mathcal{P}$, for $t \geq 0$, we define two types of attention correlations as follows:

1. Feature Attention Correlation: $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} := e_{\mathbf{p}}^{\top} Q^{(t)} v_{k,m}$ for $k \in [K]$ and $m \in [N]$;
2. Positional Attention Correlation: $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} := e_{\mathbf{p}}^{\top} Q^{(t)} e_{\mathbf{q}}$.

By our initialization, we have $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(0)} = \Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(0)} = 0$.

Next, we will apply the expression in Lemma F.1 to compute the gradient dynamics of these attention correlations.

F.1.1 FORMAL STATEMENTS AND PROOF OF LEMMA E.1 AND E.2

We first introduce some notations. Given $\mathbf{r} \in \mathcal{U}$, for $\mathbf{p} \in \mathcal{P}$, $k \in [K]$ and $n \in [N]$ define the following quantities:

$$J_{\mathbf{r}}^{\mathbf{p}} := \mathbf{M}(X)_{\mathbf{r}}^{\top} (X_{\mathbf{p}} - [F(\mathbf{M}(X); Q)]_{\mathbf{p}})$$

$$I_{\mathbf{r}}^{\mathbf{p},k,n} := \left(\tilde{\mathbf{M}}(X)_{\mathbf{r}} - \sum_{\mathbf{w} \in \mathcal{P}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} \tilde{\mathbf{M}}(X)_{\mathbf{w}} \right)^{\top} v_{k,n}$$

$$K_{\mathbf{r}}^{\mathbf{p},\mathbf{q}} := \left(\tilde{\mathbf{M}}(X)_{\mathbf{r}} - \sum_{\mathbf{w} \in \mathcal{P}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} \tilde{\mathbf{M}}(X)_{\mathbf{w}} \right)^{\top} e_{\mathbf{q}}$$

Lemma F.3 (Formal statement of Lemma E.1). *Given $k \in [K]$, for $\mathbf{p} \in \mathcal{P}$, denote $n = a_{k,\mathbf{p}}$, let $\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} = \frac{1}{\eta} (\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t+1)} - \Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)})$ for $m \in [N_k]$, then*

a. for $m = n$,

$$\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} = \mathbb{E} \left[\mathbf{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{a \neq n} z_a^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right];$$

b. for $m \neq n$,

$$\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} = \mathbb{E} \left[\mathbf{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \cdot \left(\sum_{a \neq m,n} z_a^2 z_m \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 - \left(z_m z_n^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_m^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right) \right) \right].$$

Proof. From Lemma F.1, we have

$$\begin{aligned} \alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} &= e_{\mathbf{p}}^{\top} \left(-\frac{\partial \mathcal{L}}{\partial Q} \right) v_{k,m} \\ &= \mathbb{E} \left[\mathbf{1}\{\mathbf{p} \in \mathcal{M}\} \sum_{\mathbf{r} \in \mathcal{U}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}} J_{\mathbf{r}}^{\mathbf{p}} \cdot I_{\mathbf{r}}^{\mathbf{p},k,m} \right] \\ &= \mathbb{E} \left[\mathbf{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \sum_{\mathbf{r} \in \mathcal{U}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}} J_{\mathbf{r}}^{\mathbf{p}} \cdot I_{\mathbf{r}}^{\mathbf{p},k,m} \right] \end{aligned}$$

where the last equality holds since when $k_X \neq k$, $I_{\mathbf{r}}^{\mathbf{p},k,m} = 0$ due to orthogonality. Thus, in the following, we only need to consider the case $k_X = k$.

Case 1: $m = n$.

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,n}$, since $v_{k,n'} \perp v_{k,n}$ for $n' \neq n$, and $v_{k,n} \perp \{e_{\mathbf{q}}\}_{\mathbf{q} \in \mathcal{P}}$ we have

$$\begin{aligned} J_{\mathbf{r}}^{\mathbf{p}} &= z_n v_{k,n}^{\top} \left(z_n v_{k,n} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_n v_{k,n} \right) \\ &= z_n^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}} \right) \\ I_{\mathbf{r}}^{\mathbf{p},k,n} &= \left(z_n v_{k,n} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_n v_{k,n} \right)^{\top} v_{k,n} = J_{\mathbf{r}}^{\mathbf{p}} / z_n \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,n'}$ with $n' \neq n$

$$J_{\mathbf{r}}^{\mathbf{p}} = z_{n'} v_{k,n'}^{\top} \left(z_n v_{k,n} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,n'}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_{n'} v_{k,n'} \right)$$

$$\begin{aligned}
 &= -z_{n'}^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n'}} \\
 I_{\mathbf{r}}^{\mathbf{p},k,n} &= \left(z_{n'} v_{k,n'} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_n v_{k,n} \right)^\top v_{k,n} \\
 &= -z_n \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}
 \end{aligned}$$

Putting it together, then we obtain:

$$\begin{aligned}
 e_{\mathbf{p}}^\top \left(-\frac{\partial L}{\partial Q} \right) v_{k,n} &= \mathbb{E} \left[\mathbf{1} \{ \{ \mathbf{p} \in \mathcal{M}, k_X = k \} \} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \right. \\
 &\quad \left. \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{a \neq n} z_a^2 z_n \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right]
 \end{aligned}$$

Case 2: $m \neq n$. Similarly

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,n}$

$$\begin{aligned}
 J_{\mathbf{r}}^{\mathbf{p}} &= z_n v_{k,n}^\top \left(z_n v_{k,n} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_n v_{k,n} \right) \\
 &= z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}) \\
 I_{\mathbf{r}}^{\mathbf{p},k,m} &= \left(z_n v_{k,n} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,m}} \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_m v_{k,m} \right)^\top v_{k,m} \\
 &= -z_m \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}
 \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,m}$

$$\begin{aligned}
 J_{\mathbf{r}}^{\mathbf{p}} &= z_m v_{k,m}^\top \left(z_n v_{k,n} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,m}} \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_m v_{k,m} \right) \\
 &= -z_m^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}} \\
 I_{\mathbf{r}}^{\mathbf{p},k,n} &= \left(z_m v_{k,m} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,m}} \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_m v_{k,m} \right)^\top v_{k,m} \\
 &= z_n (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}})
 \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,a}, a \neq n, m$

$$\begin{aligned}
 J_{\mathbf{r}}^{\mathbf{p}} &= z_a v_{k,a}^\top \left(z_n v_{k,n} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,a}} \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_a v_{k,a} \right) \\
 &= -z_a^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}} \\
 I_{\mathbf{r}}^{\mathbf{p},k,n} &= \left(z_a v_{k,a} - \sum_{\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,m}} \mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}} z_m v_{k,m} \right)^\top v_{k,m} \\
 &= -z_m \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}
 \end{aligned}$$

Putting them together, then we complete the proof. \square

Lemma F.4 (Formal statement of Lemma E.2). *Given $\mathbf{p}, \mathbf{q} \in \mathcal{P}$, let $\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \frac{1}{\eta} (\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t+1)} - \Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)})$, then*

$$\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \sum_{k \in [K]} \beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}, \quad \text{where } \beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)} \text{ satisfies}$$

a. if $a_{k, \mathbf{p}} = a_{k, \mathbf{q}} = n$,

$$\beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(\sum_{a \neq n} z_a^2 (\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)})^2 + z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) (\mathbb{1}\{\mathbf{q} \in \mathcal{U}\} - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \right) \right];$$

b. for $a_{k, \mathbf{p}} = n \neq m = a_{k, \mathbf{q}}$,

$$\beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(\sum_{a \neq n} z_a^2 (\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)})^2 - \left(z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + \mathbb{1}\{\mathbf{q} \in \mathcal{U}\} z_m^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right) \right) \right].$$

Proof.

$$\beta_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = e_{\mathbf{p}}^\top \left(-\frac{\partial \mathcal{L}}{\partial Q} \right) e_{\mathbf{q}} = \mathbb{E}[\mathbb{1}\{\mathbf{p} \in \mathcal{M}\} \sum_{\mathbf{r} \in \mathcal{U}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}}^{(t)} J_{\mathbf{r}}^{\mathbf{P}} K_{\mathbf{r}}^{\mathbf{P}, \mathbf{q}}]$$

Then we let

$$\beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)} := \mathbb{E}[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \sum_{\mathbf{r} \in \mathcal{U}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}}^{(t)} J_{\mathbf{r}}^{\mathbf{P}} K_{\mathbf{r}}^{\mathbf{P}, \mathbf{q}}].$$

In the following, we denote $a_{k, \mathbf{p}} = n$ and $a_{k, \mathbf{q}} = m$ for simplicity.

Case 1: $m = n$. If $\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,n}$:

- For $\mathbf{r} = \mathbf{q}$

$$\begin{aligned} J_{\mathbf{r}}^{\mathbf{P}} &= z_n v_{k,n}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} z_n v_{k,n} \right) \\ &= z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}) \\ K_{\mathbf{r}}^{\mathbf{P}, \mathbf{q}} &= (e_{\mathbf{q}} - (\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} e_{\mathbf{q}} + \sum_{\mathbf{w} \neq \mathbf{q}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} e_{\mathbf{w}}))^\top e_{\mathbf{q}} \\ &= 1 - \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}. \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,n}$, and $\mathbf{r} \neq \mathbf{q}$

$$\begin{aligned} J_{\mathbf{r}}^{\mathbf{P}} &= z_n v_{k,n}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} z_n v_{k,n} \right) \\ &= z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}) \end{aligned}$$

$$\begin{aligned} K_r^{\mathbf{p},\mathbf{q}} &= (e_r - (\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}e_{\mathbf{q}} + \sum_{\mathbf{w}\neq\mathbf{q}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}}e_{\mathbf{w}}))^\top e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \end{aligned}$$

Thus

$$\begin{aligned} &\sum_{\mathbf{r}\in\mathcal{U}\cap\mathcal{P}_{k,n}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{r}} J_r^{\mathbf{p}} \cdot K_r^{\mathbf{p},\mathbf{q}} \\ &= z_n^2 \left(1 - \sum_{\mathbf{w}\in\mathcal{U}\cap\mathcal{P}_{k,n}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} \right) \\ &\quad \cdot \left(- \sum_{\mathbf{r}\in\mathcal{U}\cap\mathcal{P}_{k,n}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{r}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} + \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \right) \\ &= z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}})^2 \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}^{(t)} \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,a}$, $a \neq n$

$$\begin{aligned} J_r^{\mathbf{p}} &= z_a v_{k,a}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w}\in\mathcal{U}\cap\mathcal{P}_{k,a}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} z_a v_{k,a} \right) \\ &= -z_a^2 \sum_{\mathbf{w}\in\mathcal{U}\cap\mathcal{P}_{k,a}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} \\ K_r^{\mathbf{p},\mathbf{q}} &= (e_r - (\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}e_{\mathbf{q}} + \sum_{\mathbf{w}\neq\mathbf{q}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}}e_{\mathbf{w}}))^\top e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \end{aligned}$$

Thus

$$\sum_{\mathbf{r}\in\mathcal{U}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{r}} J_r^{\mathbf{p}} K_r^{\mathbf{p},\mathbf{q}} = \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \cdot \left(z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}})^2 + \sum_{a\neq n} z_a^2 (\text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,a}})^2 \right)$$

If $\mathbf{q} \in \mathcal{M} \cap \mathcal{P}_{k,n}$:

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,n}$,

$$\begin{aligned} J_r^{\mathbf{p}} &= z_n v_{k,n}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w}\in\mathcal{U}\cap\mathcal{P}_{k,n}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} z_n v_{k,n} \right) \\ &= z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}) \\ K_r^{\mathbf{p},\mathbf{q}} &= (e_r - (\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}e_{\mathbf{q}} + \sum_{\mathbf{w}\neq\mathbf{q}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}}e_{\mathbf{w}}))^\top e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,a}$, $a \neq n$

$$\begin{aligned} J_r^{\mathbf{p}} &= z_a v_{k,a}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w}\in\mathcal{U}\cap\mathcal{P}_{k,a}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} z_a v_{k,a} \right) \\ &= -z_a^2 \sum_{\mathbf{w}\in\mathcal{U}\cap\mathcal{P}_{k,a}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} \end{aligned}$$

$$\begin{aligned} K_r^{\mathbf{p},\mathbf{q}} &= (e_r - (\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}e_{\mathbf{q}} + \sum_{\mathbf{w}\neq\mathbf{q}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}}e_{\mathbf{w}}))^\top e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \end{aligned}$$

Thus

$$\begin{aligned} &\sum_{r\in\mathcal{U}} \text{attn}_{\mathbf{p}\rightarrow r} J_r^{\mathbf{p}} K_r^{\mathbf{p},\mathbf{q}} \\ &= \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \cdot \left(z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}})^2 - z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}) + \sum_{a\neq n} z_a^2 (\text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,a}})^2 \right) \end{aligned}$$

Putting it together,

$$\begin{aligned} \beta_{k,\mathbf{p}\rightarrow\mathbf{q}}^{(t)} &= \mathbb{E} [\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \cdot \\ &\quad \left(-z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}) \mathbb{1}\{\mathbf{q} \in \mathcal{M}\} + z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}})^2 + \sum_{m\neq n} z_m^2 (\text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}})^2 \right)] \end{aligned}$$

Case 2: $m \neq n$. Similarly, if $\mathbf{q} \in \mathcal{U} \cap \mathcal{P}_{k,m}$:

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,n}$,

$$\begin{aligned} J_r^{\mathbf{p}} &= z_n v_{k,n}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} z_n v_{k,n} \right) \\ &= z_n^2 (1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}) \\ K_r^{\mathbf{p},\mathbf{q}} &= (e_r - \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}e_{\mathbf{q}} - \sum_{\mathbf{w}\neq\mathbf{q}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}}e_{\mathbf{w}})^\top e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \end{aligned}$$

- For $\mathbf{r} = \mathbf{q}$

$$\begin{aligned} J_r^{\mathbf{p}} &= z_m v_{k,m}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w} \in \mathcal{U} \cap \mathcal{P}_{k,m}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} z_m v_{k,m} \right) \\ &= -z_m^2 \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}} \\ K_r^{\mathbf{p},\mathbf{q}} &= (e_{\mathbf{q}} - \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}e_{\mathbf{q}} - \sum_{\mathbf{w}\neq\mathbf{q}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}}e_{\mathbf{w}})^\top e_{\mathbf{q}} \\ &= 1 - \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,a}$, $a \neq n$, and $\mathbf{r} \neq \mathbf{q}$

$$\begin{aligned} J_r^{\mathbf{p}} &= z_a v_{k,a}^\top \left(z_n v_{k,n} - \sum_{\mathbf{w} \in \mathcal{U} \cap \mathcal{P}_{k,a}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}} z_a v_{k,a} \right) \\ &= -z_a^2 \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,a}} \\ K_r^{\mathbf{p},\mathbf{q}} &= (e_r - \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}e_{\mathbf{q}} - \sum_{\mathbf{w}\neq\mathbf{q}} \text{attn}_{\mathbf{p}\rightarrow\mathbf{w}}e_{\mathbf{w}})^\top e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p}\rightarrow\mathbf{q}} \end{aligned}$$

Thus

$$\begin{aligned} & \sum_{\mathbf{r} \in \mathcal{U}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}} J_{\mathbf{r}}^{\mathbf{p}} K_{\mathbf{r}}^{\mathbf{p}, \mathbf{q}} \\ &= \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \cdot \left(-z_n^2 (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}} - z_m^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}} + \sum_{a \neq n} z_a^2 (\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}})^2 \right) \end{aligned}$$

If $\mathbf{q} \in \mathcal{M} \cap \mathcal{P}_{k,m}$:

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,n}$,

$$\begin{aligned} J_{\mathbf{r}}^{\mathbf{p}} &= z_n v_{k,n}^{\top} \left(z_n v_{k,n} - \sum_{\mathbf{w} \in \mathcal{U} \cap \mathcal{P}_{k,n}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} z_n v_{k,n} \right) \\ &= z_n^2 (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}) \\ K_{\mathbf{r}}^{\mathbf{p}, \mathbf{q}} &= (e_{\mathbf{r}} - \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} e_{\mathbf{q}} - \sum_{\mathbf{w} \neq \mathbf{q}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} e_{\mathbf{w}})^{\top} e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \end{aligned}$$

- For $\mathbf{r} \in \mathcal{U} \cap \mathcal{P}_{k,a}$, $a \neq n$

$$\begin{aligned} J_{\mathbf{r}}^{\mathbf{p}} &= z_a v_{k,a}^{\top} \left(z_n v_{k,n} - \sum_{\mathbf{w} \in \mathcal{U} \cap \mathcal{P}_{k,a}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} z_a v_{k,a} \right) \\ &= -z_a^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}} \\ K_{\mathbf{r}}^{\mathbf{p}, \mathbf{q}} &= (e_{\mathbf{r}} - \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} e_{\mathbf{q}} - \sum_{\mathbf{w} \neq \mathbf{q}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{w}} e_{\mathbf{w}})^{\top} e_{\mathbf{q}} \\ &= -\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \end{aligned}$$

Thus

$$\begin{aligned} & \sum_{\mathbf{r} \in \mathcal{U}} \text{attn}_{\mathbf{p} \rightarrow \mathbf{r}} J_{\mathbf{r}}^{\mathbf{p}} K_{\mathbf{r}}^{\mathbf{p}, \mathbf{q}} \\ &= \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \cdot \left(-z_n^2 (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}} + \sum_{a \neq n} z_a^2 (\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}})^2 \right). \end{aligned}$$

Therefore

$$\begin{aligned} \beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)} &= \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}} \cdot \right. \\ &\quad \left. \left(-z_n^2 (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}} - \mathbb{1}\{\mathbf{q} \in \mathcal{U}\} z_m^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}} \right. \right. \\ &\quad \left. \left. + \sum_{a \neq n} z_a^2 (\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}})^2 \right) \right]. \end{aligned}$$

□

Based on the above gradient update for $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$, we further introduce the following auxiliary quantity, which will be useful in the later proof.

$$\Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t+1)} := \Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)} + \eta \beta_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}, \quad \text{with } \Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(0)} = 0 \quad (10)$$

It is easy to verify that $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \sum_{k \in [K]} \Upsilon_{k, \mathbf{p} \rightarrow \mathbf{q}}^{(t)}$.

F.2 High-probability Event

We first introduce the following exponential bounds for the hypergeometric distribution $\text{Hyper}(m, D, M)$. $\text{Hyper}(m, D, M)$ describes the probability of certain successes (random draws for which the object drawn has a specified feature) in m draws, without replacement, from a finite population of size M that contains exactly D objects with that feature, wherein each draw is either a success or a failure.

Proposition F.5 ((Greene & Wellner, 2017)). *Suppose $S \sim \text{Hyper}(m, D, M)$ with $1 \leq m, D \leq M$. Define $\mu_M := D/M$. Then for all $t > 0$*

$$P(|S - m\mu_M| > t) \leq 2 \exp\left(-\frac{t^2}{4m\mu_M + 2t}\right).$$

We then utilize this property to prove the high-probability set introduced in Appendix E.1.

Lemma F.6. *For $k \in [K]$ $n \in [N]$, define*

$$\mathcal{E}_{k,n}(\gamma, P) := \{\mathbf{M} : |\mathcal{P}_{k,n} \cap \mathcal{U}| = \Theta(C_n)\}, \quad (11)$$

we have

$$\mathbb{P}(\mathbf{M} \in \mathcal{E}_{k,n}) \geq 1 - 2 \exp(-c_{n,1}C_n) \quad (12)$$

where $c_{n,0} > 0$ is some constant.

Proof. Under the random masking strategy, given $k \in [K]$ and $n \in [N]$, $Y_{k,n} = |\mathcal{U} \cap \mathcal{P}_{k,n}|$ follows the hypergeometric distribution, i.e. $Y_{k,n} \sim \text{Hyper}((1-\gamma)P, C_n, P)$. Then by tail bounds, for $t > 0$, we have:

$$\mathbb{P}[|Y_{k,n} - (1-\gamma)C_n| > t] \leq 2 \exp\left(-\frac{t^2}{4(1-\gamma)C_n + 2t}\right)$$

Letting $t = \Theta(C_n)$, we have

$$\mathbb{P}[Y_{k,n} = \Theta(C_n)] \geq 1 - 2e^{-c_{n,1}C_n}.$$

□

We further have the following fact, which will be useful for proving the property of loss objective in the next subsection.

Lemma F.7. *For $k \in [K]$ and $n \in [N]$, we have*

$$\mathbb{P}(|\mathcal{U} \cap \mathcal{P}_{k,n}| = 0) \leq \exp(-c_{n,0}C_n). \quad (13)$$

where $c_{n,0} > 0$ is some constant.

Proof. By the form of probability density for $\text{Hyper}((1-\gamma)P, C_n, P)$, we have

$$\begin{aligned} \mathbb{P}(|\mathcal{U} \cap \mathcal{P}_{k,n}| = 0) &= \frac{\binom{C_n}{0} \binom{P-C_n}{(1-\gamma)P}}{\binom{P}{(1-\gamma)P}} \\ &\leq \gamma^{C_n} = \exp(-c_{n,0}C_n). \end{aligned}$$

□

F.3 Properties of Loss Function

Recall the training and regional reconstruction loss we consider are given by:

$$\mathcal{L}(Q) := \frac{1}{2} \mathbb{E} \left[\sum_{\mathbf{p} \in \mathcal{P}} \mathbb{1}\{\mathbf{p} \in \mathcal{M}\} \|[F(\mathbf{M}(X); Q, E)]_{\mathbf{p}} - X_{\mathbf{p}}\|^2 \right] \quad (14)$$

$$\mathcal{L}_{\mathbf{p}}(Q) = \frac{1}{2} \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}\} \|[F(\mathbf{M}(X), E)]_{\mathbf{p}} - X_{\mathbf{p}}\|^2 \right] \quad (15)$$

In this part, we will present several important lemmas for such a training objective. We first single out the following lemma, which connects the loss form with the attention score.

Lemma F.8 (Loss Calculation). *The population loss $L(Q)$ can be decomposed into the following form:*

$$\begin{aligned} \mathcal{L}(Q) &= \sum_{\mathbf{p} \in \mathcal{P}} \mathcal{L}_{\mathbf{p}}(Q), \text{ where} \\ \mathcal{L}_{\mathbf{p}}(Q) &= \frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \cdot \right. \\ &\quad \left. \left(z_{a_{k,\mathbf{p}}}^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(t)} \right)^2 + \sum_{a \neq a_{k,\mathbf{p}}} z_a^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right] \end{aligned}$$

Proof.

$$\begin{aligned} &\mathcal{L}_{\mathbf{p}}(Q) \\ &= \frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \|[F(\mathbf{M}(X), E)]_{\mathbf{p}} - X_{\mathbf{p}}\|^2 \right] \\ &= \frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \left\| \sum_{m \in [N]} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}} z_m v_{k,m} - z_{a_{k,\mathbf{p}}} v_{k,a_{k,\mathbf{p}}} \right\|^2 \right] \\ &\stackrel{(i)}{=} \frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \left(z_{a_{k,\mathbf{p}}}^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}} \right)^2 + \sum_{m \neq a_{k,\mathbf{p}}} z_m^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}} \right)^2 \right) \right] \end{aligned}$$

where (i) follows since the features are orthogonal. □

We then introduce some additional crucial notations for the loss objectives.

$$\mathcal{L}_{\mathbf{p}}^* = \min_{Q \in \mathbb{R}^{d \times d}} \mathcal{L}_{\mathbf{p}}(Q), \quad (16a)$$

$$\mathcal{L}_{\mathbf{p}}^{\text{low}} = \frac{1}{2} \left(\sigma_z^2 + \frac{L^2}{N-1} \right) \sum_{k \in [K]} \mathbb{P} \left(|\mathcal{U} \cap \mathcal{P}_{k,z_{a_{k,\mathbf{p}}}}| = 0 \right) \quad (16b)$$

$$\tilde{\mathcal{L}}_{\mathbf{p}}(Q) = \sum_{k=1}^K \tilde{\mathcal{L}}_{k,\mathbf{p}}(Q), \quad \text{where}$$

$$\begin{aligned} \tilde{\mathcal{L}}_{k,\mathbf{p}}(Q) &= \frac{1}{2} \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k, \mathbf{M} \in \mathcal{E}_{k,z_{a_{k,\mathbf{p}}}}\} \cdot \right. \\ &\quad \left. \left(z_{a_{k,\mathbf{p}}}^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_{k,\mathbf{p}}}}^{(t)} \right)^2 + \sum_{a \neq a_{k,\mathbf{p}}} z_a^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right] \quad (16c) \end{aligned}$$

Here $\sigma_z^2 = \mathbb{E}[Z_n(X)^2]$. \mathcal{L}_p^* denotes the minimum value of the population loss in (15), and $\mathcal{L}_p^{\text{low}}$ represents the unavoidable errors for $\mathbf{p} \in \mathcal{P}$, given that all the patches in $\mathcal{P}_{k,a_k,\mathbf{p}}$ are masked. We will show that $\mathcal{L}_p^{\text{low}}$ serves as a lower bound for \mathcal{L}_p^* , and demonstrate that the network trained with GD will attain nearly zero error compared to $\mathcal{L}_p^{\text{low}}$. Our convergence will be established by the sub-optimality gap with respect to $\mathcal{L}_p^{\text{low}}$, which necessarily implies the convergence to \mathcal{L}_p^* . (It also implies $\mathcal{L}_p^* - \mathcal{L}_p^{\text{low}}$ is small.)

Lemma F.9. For \mathcal{L}_p^* and $\mathcal{L}_p^{\text{low}}$ defined in (16a) and (16b), respectively, we have $\mathcal{L}_p^{\text{low}} \leq \mathcal{L}_p^*$ and they are both at the order of $\Theta\left(\exp\left(-\left(c_1 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k,\mathbf{p}}\}\} c_2 P^{\kappa_s}\right)\right)\right)$ where $c_1, c_2 > 0$ are some constants.

Proof. We first prove $\mathcal{L}_p^{\text{low}} \leq \mathcal{L}_p^*$:

$$\begin{aligned} \mathcal{L}_p^* &= \min_{Q \in \mathbb{R}^{d \times d}} \frac{1}{2} \sum_{k=1}^K \mathbb{E} [\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \cdot \\ &\quad \left(z_{a_{k,\mathbf{p}}}^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_k,\mathbf{p}}}^{(t)}\right)^2 + \sum_{a \neq a_{k,\mathbf{p}}} z_a^2 z_{a_{k,\mathbf{p}}} \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)}\right)^2 \right)] \\ &\geq \min_{Q \in \mathbb{R}^{d \times d}} \frac{1}{2} \sum_{k=1}^K \mathbb{E} [\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k\} \mathbb{1}\{|\mathcal{U} \cap \mathcal{P}_{k,a_k,\mathbf{p}}| = 0\} \cdot \\ &\quad \left(z_{a_{k,\mathbf{p}}}^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_k,\mathbf{p}}}^{(t)}\right)^2 + \sum_{a \neq a_{k,\mathbf{p}}} z_a^2 z_{a_{k,\mathbf{p}}} \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)}\right)^2 \right)] \end{aligned}$$

Notice that when all patches in $\mathcal{P}_{k,a_k,\mathbf{p}}$ are masked, $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_k,\mathbf{p}}}^{(t)} = 0$. Moreover,

$$\sum_{m \neq a_{k,\mathbf{p}}} z_m^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \geq \frac{L^2}{N-1}$$

by Cauchy–Schwarz inequality. Thus

$$\mathcal{L}_p^* \geq \frac{1}{2} \sum_{k=1}^K \left(\sigma_z^2 + \frac{L^2}{N-1} \right) \mathbb{P}(|\mathcal{U} \cap \mathcal{P}_{k,a_k,\mathbf{p}}| = 0) = \mathcal{L}_p^{\text{low}}.$$

$\mathcal{L}_p^{\text{low}} = \Theta\left(\exp\left(-\left(c_1 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k,\mathbf{p}}\}\} c_2 P^{\kappa_s}\right)\right)\right)$ immediately comes from Lemma F.7. Furthermore, we only need to show $\mathcal{L}_p^* = O\left(\exp\left(-\left(c_1 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k,\mathbf{p}}\}\} c_2 P^{\kappa_s}\right)\right)\right)$. This can be directly obtained by choosing $Q = \sigma I_d$ for some sufficiently large σ and hence omitted here. \square

Lemma F.10. Given $\mathbf{p} \in \mathcal{P}$, for any Q , we have

$$\tilde{\mathcal{L}}_p(Q) \leq L_p(Q) - \mathcal{L}_p^{\text{low}} \leq \tilde{\mathcal{L}}_p(Q) + O\left(\exp\left(-\left(c_3 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k,\mathbf{p}}\}\} c_4 P^{\kappa_s}\right)\right)\right).$$

where $c_3, c_4 > 0$ are some constants.

Proof. The lower bound is directly obtained by the definition and thus we only prove the upper bound.

$$\begin{aligned} &L_p(Q) - \tilde{\mathcal{L}}_p(Q) \\ &= \frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\mathbb{1}\{\mathbf{p} \in \mathcal{M}, k_X = k, \mathbf{M} \in \mathcal{E}_{k,z_{a_{k,\mathbf{p}}}}^c\} \cdot \left(z_{a_{k,\mathbf{p}}}^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_k,\mathbf{p}}}^{(t)}\right)^2 + \sum_{a \neq a_{k,\mathbf{p}}} z_a^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)}\right)^2 \right) \right] \end{aligned}$$

$$\begin{aligned} &\leq \sum_{k=1}^K U^2 \mathbb{P}(M \in \mathcal{E}_{k, z_{a_k, \mathbf{p}}}^c) \\ &\leq O\left(\exp\left(-\left(c_3 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k, \mathbf{p}}\}\} c_4 P^{\kappa_s}\right)\right)\right). \end{aligned}$$

where the last inequality follows from Lemma F.6. □

G Overall Induction Hypotheses and Proof Plan

Our main proof utilizes the induction hypotheses. In this section, we introduce the main induction hypotheses for the positive and negative information gaps, which will later be proven to be valid throughout the entire learning process.

G.1 Positive Information Gap

We first state our induction hypothesis for the case that the information gap Δ is positive.

Induction Hypothesis G.1. For $t \leq T$, given $\mathbf{p}, \mathbf{q} \in \mathcal{P}$, for $k \in [K]$, the following holds

- a. $\Phi_{\mathbf{p} \rightarrow v_{k, a_k, \mathbf{p}}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k, a_k, \mathbf{p}}}^{(t)} \in [0, \tilde{O}(1)]$;
- b. if $a_{k, \mathbf{p}} \neq 1$, then $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)}$ is monotonically decreasing and $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)} \in [-\tilde{O}(1), 0]$;
- c. $|\Phi_{\mathbf{p} \rightarrow v_{k, m}}^{(t)}| = \tilde{O}\left(\frac{1}{P^{1-\kappa_s}}\right)$ for $m \notin \{1\} \cup \{a_{k, \mathbf{p}}\}$;
- d. for $\mathbf{q} \neq \mathbf{p}$, $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \tilde{O}\left(\frac{1}{P^{\kappa_s}}\right)$;
- e. $\Upsilon_{\mathbf{p} \rightarrow \mathbf{p}}^{(t)} = \tilde{O}\left(\frac{1}{P}\right)$.

G.2 Negative Information Gap

Now we turn to the case that $\Delta \leq -\Omega(1)$.

Induction Hypothesis G.2. For $t \leq T$, given $\mathbf{p}, \mathbf{q} \in \mathcal{P}$, for $k \in [K]$, the following holds

- a. $\Phi_{\mathbf{p} \rightarrow v_{k, a_k, \mathbf{p}}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k, a_k, \mathbf{p}}}^{(t)} \in [0, \tilde{O}(1)]$;
- b. if $a_{k, \mathbf{p}} \neq 1$, then $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)}$ is monotonically decreasing and $\Phi_{\mathbf{p} \rightarrow v_{k, 1}}^{(t)} \in [-\tilde{O}\left(\frac{1}{P^{-\Delta}}\right), 0]$;
- c. $|\Phi_{\mathbf{p} \rightarrow v_{k, m}}^{(t)}| = \tilde{O}\left(\frac{1}{P^{1-\kappa_s}}\right)$ for $m \notin \{1\} \cup \{a_{k, \mathbf{p}}\}$;
- d. for $\mathbf{q} \neq \mathbf{p}$, $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \tilde{O}\left(\frac{1}{P^{\kappa_s}}\right)$;
- e. $\Upsilon_{\mathbf{p} \rightarrow \mathbf{p}}^{(t)} = \tilde{O}\left(\frac{1}{P}\right)$.

G.3 Proof Outline

In both settings, we can classify the process through which transformers learn the feature attention correlation $\Phi_{\mathbf{p} \rightarrow v_{k, a_k, \mathbf{p}}}^{(t)}$ into two distinct scenarios. These scenarios hinge on the spatial relation of the area \mathbf{p} within the context of the k -th partition \mathcal{D}_k , specifically, whether \mathbf{p} is located in the global area of the k -th cluster, i.e. whether $a_{k, \mathbf{p}} = 1$. The learning dynamics exhibit different behaviors of learning the local FP correlation in the local area with different Δ , while the behaviors for features located in the global area are very similar, unaffected by the value of Δ . Therefore, through Appendices H to J, we delve into the learning phases and provide technical proofs for the local area with $\Delta \geq \Omega(1)$, local area with $\Delta \leq -\Omega(1)$ and the global area respectively. Finally, we will put this analysis together to prove that the Induction Hypothesis G.1 (resp. Induction Hypothesis G.2) holds during the entire training process, thereby validating the main theorems in Appendix K.

H Analysis for the Local Area with Positive Information Gap

In this section, we focus on a specific patch $\mathbf{p} \in \mathcal{P}$ with the k -th cluster for $k \in [K]$, and present the analysis for the case that $X_{\mathbf{p}}$ is located in the local area for the k -th cluster, i.e. $a_{k,\mathbf{p}} > 1$. We will analyze the case that $\Delta \geq \Omega(1)$. Throughout this section, we denote $a_{k,\mathbf{p}} = n$ for simplicity. We will analyze the convergence of the training process via two phases of dynamics. At the beginning of each phase, we will establish an induction hypothesis, which we expect to remain valid throughout that phase. Subsequently, we will analyze the dynamics under such a hypothesis within the phase, aiming to provide proof of the hypothesis by the end of the phase.

H.1 Phase I, Stage 1

In this section, we shall discuss the initial stage of phase I. Firstly, we present the induction hypothesis in this stage.

We define the stage 1 of phase I as all iterations $t \leq T_1$, where

$$T_1 \triangleq \max \left\{ t : \Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq -\frac{1}{U} \left(\frac{\Delta}{2} - 0.01 \right) \log(P) \right\}.$$

We state the following induction hypotheses, which will hold throughout this period:

Induction Hypothesis H.1. For each $0 \leq t \leq T_1$, $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, the following holds:

- $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \in [0, O\left(\frac{(\frac{\Delta}{2} - 0.01) \log(P)}{P^{0.02}}\right)]$;
- $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically decreasing and $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in [-\frac{1}{U} (\frac{\Delta}{2} - 0.01) \log(P), 0]$;
- $|\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right)$ for $m \neq 1, n$;
- $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ for $a_{k,\mathbf{q}} = n$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$;
- $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{|\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right) + O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} = 1$;
- $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} \neq 1, n$.

H.1.1 PROPERTY OF ATTENTION SCORES

We first introduce several properties of the attention score if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold.

Lemma H.1. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold at iteration $t \leq T_1$, then the following holds

- $1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \geq \Omega(1)$;
- If $\mathbf{M} \in \mathcal{E}_{k,n}$, $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} = \Theta\left(\frac{1}{P^{1-\kappa_s}}\right)$;
- Moreover, if $\mathbf{M} \in \mathcal{E}_{k,1}$, we have $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} = \Omega\left(\frac{1}{P^{\frac{1-\kappa_s}{2} - 0.01}}\right)$;
- For $\mathbf{q} \in \mathcal{M} \cap (\mathcal{P}_{k,n} \cup \mathcal{P}_{k,1})$, $\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{P}\right)$.

Lemma H.2. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold at iteration $t \leq T_1$, then for $m \neq n, 1$, the following holds:

- For any $\mathbf{q} \in \mathcal{P}_{k,m}$, $\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \leq O\left(\frac{1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{P}\right)$.

2. Moreover, $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \leq O\left(\frac{1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{N}\right)$.

The above properties can be easily verified through direct calculations by using the definition in (4) and conditions in Induction Hypothesis H.1, which are omitted here for brevity.

H.1.2 BOUNDING THE GRADIENT UPDATES FOR FP CORRELATIONS

Lemma H.3. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold at iteration $0 \leq t \leq T_1$, then $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq 0$ and satisfies:

$$\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} = \Theta\left(\frac{C_n}{P}\right) = \Theta\left(\frac{1}{P^{1-\kappa_s}}\right).$$

Proof. By Lemma E.2, we have

$$\begin{aligned} & \alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}\right)^2 \right) \right] \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}\right)^2 \right) \right] \\ &+ \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n}^c \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}\right)^2 \right) \right] \\ &\leq \mathbb{P}(\mathbf{M} \in \mathcal{E}_{k,n}) \\ &\quad \cdot \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}\right)^2 \right) \middle| \mathcal{E}_{k,n} \right] \\ &\quad + O(1) \cdot \mathbb{P}(\mathbf{M} \in \mathcal{E}_{k,n}^c) \\ &\leq O\left(\frac{C_n}{P}\right) + O(\exp(-c_{n,1} C_n)) \\ &\leq O\left(\frac{C_n}{P}\right), \end{aligned}$$

where the second inequality invokes Lemma H.1 and Lemma F.6, and the last inequality is due to $\exp(-c_{n,1} C_n) \ll \frac{C_n}{P}$. Similarly, we can show that $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq \Omega\left(\frac{C_n}{P}\right)$. \square

Lemma H.4. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold at iteration $0 \leq t \leq T_1$, then $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} < 0$ and satisfies

$$|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| \geq \Omega\left(\frac{1}{P^{2\left(\frac{1-\kappa_s}{2} - 0.01\right)}}\right) = \Omega\left(\frac{1}{P^{0.98-\kappa_s}}\right).$$

Proof. We first single out the following fact:

$$\begin{aligned} & -z_1 z_n^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - z_1^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}\right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} + \sum_{a \neq 1,n} z_a^2 z_1 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)}\right)^2 \\ &\leq z_1 \left(\max_{a \neq 1,n} z_a^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} - z_n^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - z_1^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \right) \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}\right) \end{aligned}$$

$$= -z_1(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \left(z_n^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_1^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \max_{a \neq 1,n} z_a^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right). \quad (17)$$

Therefore, by Lemma E.1, we have

$$\begin{aligned} \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \right. \\ &\quad \left. \left(-z_1(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \left(z_n^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_1^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \max_{a \neq 1,n} z_a^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right) \right) \right] \\ &\quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1}^c \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \sum_{a \neq 1,n} z_1^2 z_a \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right] \\ &\leq \mathbb{P}(M \in \mathcal{E}_{k,1}) \cdot \left(-\left(\Omega(1) \cdot \Omega\left(\frac{1}{P^{2 \times (\frac{1-\kappa_s}{2} - 0.01)}} \right) \right) \right) + O(1) \cdot \mathbb{P}(M \in \mathcal{E}_{k,1}^c) \\ &\leq -\Omega\left(\frac{1}{P^{2 \times (\frac{1-\kappa_s}{2} - 0.01)}} \right) = -\Omega\left(\frac{1}{P^{0.98-\kappa_s}} \right) \end{aligned}$$

where the second inequality invokes Lemma H.1 and the last inequality comes from Lemma F.6. \square

Lemma H.5. *At each iteration $t \leq T_1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold, then for any $m > 1$ with $m \neq n$, the following holds*

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{N} \right) = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}} \right).$$

Proof. By Lemma E.1, for $m \neq n$, we have

$$\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} \leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \cdot \left(\sum_{a \neq m,n} z_a^2 z_m \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right] \quad (18)$$

$$\begin{aligned} -\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \cdot \left(z_m z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right. \right. \\ &\quad \left. \left. + z_m^3 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right) \right] \quad (19) \end{aligned}$$

For (18), we have

$$\begin{aligned} &\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} \\ &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \cdot \left(\sum_{a \neq m,n} z_a^2 z_m \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right] \\ &\quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, (\mathcal{E}_{k,1} \cap \mathcal{E}_{k,n})^c \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \cdot \left(\sum_{a \neq m,n} z_a^2 z_m \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right] \\ &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{N} \right) \right. \\ &\quad \left. \cdot \left(z_1^2 z_m \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \right)^2 + O\left(\frac{1}{N} \right) \right) \right] + O(1) \cdot \mathbb{P}(M \in (\mathcal{E}_{k,1} \cap \mathcal{E}_{k,n})^c) \\ &\leq O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{N} \right) + O(1) \cdot \mathbb{P}(M \in (\mathcal{E}_{k,1} \cap \mathcal{E}_{k,n})^c) \end{aligned}$$

$$\leq O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{P^{1-\kappa_s}}\right)$$

where the second inequality is due to Lemma H.2, the last inequality follows from Lemma H.4 and Lemma F.6.

On the other hand, for (19), we can use the similar argument by invoking Lemma H.2 and Lemma H.3, and thus obtain

$$-\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} \leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{1-\kappa_s}}\right).$$

Putting them together, we have

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right).$$

□

H.1.3 BOUNDING THE GRADIENT UPDATES FOR POSITIONAL CORRELATIONS

Lemma H.6. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold at iteration $0 \leq t \leq T_1$, then for $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$ and $a_{k,\mathbf{q}} = n$, we have $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \geq 0$ and satisfies:

$$\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right).$$

Furthermore, we have $|\beta_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$.

Proof. By Lemma E.2, for $\mathbf{q} \in \mathcal{P}_{k,n}$ with $\mathbf{q} \neq \mathbf{p}$, we have

$$\begin{aligned} \beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} &= \\ &\underbrace{\mathbb{E} \left[\mathbf{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{U}\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(z_n^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right]}_{H_1} \\ &+ \underbrace{\mathbb{E} \left[\mathbf{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{M}\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(-z_n^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right) \right) \right]}_{H_2} \\ &+ \underbrace{\mathbb{E} \left[\mathbf{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{M}\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(\sum_{m \neq n} z_m^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right]}_{H_3}. \end{aligned}$$

Firstly, for H_1 , notice that

$$\begin{aligned} (C_n - 1)H_1 &= \mathbb{E} \left[\mathbf{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &= \Theta(\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}). \end{aligned}$$

For H_2 , since $\mathbf{p}, \mathbf{q} \in \mathcal{M}$, by Lemma H.1, we can upper bound $\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ by $O\left(\frac{1}{P}\right)$, thus

$$-H_2 \leq \mathbb{E} \left[\mathbf{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} O\left(\frac{1}{P}\right) \cdot \left(z_n^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right) \right) \right] \leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P}\right).$$

Further notice that H_3 can be upper bounded by $O(H_1)$, putting it together, we have

$$\beta_{k,\mathbf{p}\rightarrow\mathbf{q}}^{(t)} = \Theta\left(\frac{\alpha_{\mathbf{p}\rightarrow v_{k,n}}^{(t)}}{C_n}\right).$$

Turn to $\beta_{k,\mathbf{p}\rightarrow\mathbf{p}}^{(t)}$, when $\mathbf{q} = \mathbf{p}$,

$$\begin{aligned} \beta_n^{(t)} &= \mathbb{E} \left[\underbrace{\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \text{attn}_{\mathbf{p}\rightarrow\mathbf{p}}^{(t)} \cdot \left(-z_n^2 \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \left(1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right) \right)}_{J_2} \right] \\ &+ \mathbb{E} \left[\underbrace{\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \text{attn}_{\mathbf{p}\rightarrow\mathbf{p}}^{(t)} \cdot \left(\sum_{m \neq n} z_m^2 \left(\text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}}^{(t)} \right)^2 \right)}_{J_3} \right]. \end{aligned}$$

We can bound J_2 in a similar way as H_2 . Thus, we only focus on further bounding J_3 :

$$\begin{aligned} J_3 &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} O\left(\frac{1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,1}}^{(t)} - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)}}{P}\right) \cdot \left(\sum_{m \neq n} z_m^2 \left(\text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &\leq O\left(\frac{|\alpha_{\mathbf{p}\rightarrow v_{k,1}}^{(t)}|}{P}\right). \end{aligned}$$

where the first inequality holds by invoking Lemma H.1 and the last inequality follows similar arguments as analysis for (18). \square

Lemma H.7. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold at iteration $0 \leq t \leq T_1$, then for $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$ and $a_{k,\mathbf{q}} = 1$, we have $\beta_{k,\mathbf{p}\rightarrow\mathbf{q}}^{(t)}$ satisfies:

$$|\beta_{k,\mathbf{p}\rightarrow\mathbf{q}}^{(t)}| = O\left(\frac{|\alpha_{\mathbf{p}\rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p}\rightarrow v_{k,1}}^{(t)}|}{P}\right) + O\left(\frac{|\alpha_{\mathbf{p}\rightarrow v_{k,1}}^{(t)}|}{C_1}\right).$$

Proof. By Lemma E.2, for $\mathbf{q} \in \mathcal{P}_{k,1}$, we have

$$\begin{aligned} \beta_{k,\mathbf{p}\rightarrow\mathbf{q}}^{(t)} &= \\ &- \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{U}\} \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}^{(t)} \cdot \right. \\ &\quad \left. \left(z_1^2 \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,1}}^{(t)} \left(1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,1}}^{(t)} \right) + z_n^2 \left(1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right) \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} - \sum_{a \neq 1,n} z_a^2 \left(\text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right] \quad (20) \\ &- \mathbb{E} \left[\underbrace{\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{M}\} \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}^{(t)} \cdot \left(z_n^2 \left(1 - \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right) \text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right)}_{G_2} \right] \\ &+ \mathbb{E} \left[\underbrace{\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{M}\} \text{attn}_{\mathbf{p}\rightarrow\mathbf{q}}^{(t)} \cdot \left(\sum_{a \neq n} z_a^2 \left(\text{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,a}}^{(t)} \right)^2 \right)}_{G_3} \right]. \end{aligned}$$

For (20) denoted as G_1 , following the direct calculations, we have

$$-(C_1 - 1)G_1 = \Theta(\alpha_{\mathbf{p}\rightarrow v_{k,1}}^{(t)})$$

We can further bound G_2 and G_3 in a similar way as H_2 and H_3 in Lemma H.6 and thus obtain

$$\begin{aligned} -G_2 &\leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P}\right), \\ G_3 &\leq O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{P}\right). \end{aligned}$$

which completes the proof. \square

Lemma H.8. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold at iteration $0 \leq t \leq T_1$, then for $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$ and $n \neq a_{k,\mathbf{q}}$, $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ satisfies:

$$|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right).$$

Proof. By Lemma E.2, for $\mathbf{q} \in \mathcal{P}_{k,m}$, we have

$$\begin{aligned} \beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} &= \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{U}\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \right. \\ &\quad \left. \left(z_m^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}) + z_n^2 (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \sum_{a \neq n,m} z_a^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right] \quad (21) \\ &= \underbrace{\mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{M}\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(z_n^2 (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right) \right]}_{I_2} \\ &\quad + \underbrace{\mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}, \mathbf{q} \in \mathcal{M}\} \text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \cdot \left(\sum_{a \neq n} z_a^2 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right) \right]}_{I_3}. \end{aligned}$$

(21) can be upper bounded by $O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}|}{C_m}\right) = O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{NC_m}\right) = O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{P}\right)$, where the first equality holds by invoking Lemma H.5. I_2 and I_3 can be bounded similarly as G_2 and G_3 , which is omitted here. \square

H.1.4 AT THE END OF PHASE I, STAGE 1

Lemma H.9. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.1 hold for all $0 \leq t \leq T_1 = O\left(\frac{\log(P)P^{0.98-\kappa_s}}{\eta}\right)$, At iteration $t = T_1 + 1$, we have

- $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(T_1+1)} \leq -\frac{1}{U} \left(\frac{\Delta}{2} - 0.01 \right) \log(P)$;
- $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(T_1+1)} = O\left(\frac{1}{P^{(1-\kappa_c) + \frac{\Delta}{U} (\frac{\Delta}{2} - 0.01)}}\right)$.

Proof. By comparing Lemma H.3 and Lemma H.4, we have $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| \gg \alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$. Then the existence of $T_{1,k} = O\left(\frac{\log(P)P^{0.98-\kappa_s}}{\eta}\right)$ directly follows from Lemma H.4. \square

H.2 Phase I, Stage 2

During stage 1, $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ significantly decreases to decouple the FP correlations with the global feature, resulting in a decrease in $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}$, while other $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}$ with $m > 1$ remain approximately at the order of $O\left(\frac{1}{P^{1-\kappa_s}}\right) (\Theta\left(\frac{1}{P^{1-\kappa_s}}\right))$. By

the end of phase I, $(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)})^2$ decreases to $O(\frac{1}{P^{1.96-2\kappa_s}})$, leading to a decrease in $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|$ as it approaches towards $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$. At this point, stage 2 begins. Shortly after entering this phase, the prior dominant role of the decrease of $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ in learning dynamics diminishes as $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|$ reaches the same order of magnitude as $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$.

We define stage 2 of phase I as all iterations $T_1 < t \leq \tilde{T}_1$, where

$$\tilde{T}_1 \triangleq \max \left\{ t > T_1 : \Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \leq \left(\frac{\Delta}{2L} + \frac{0.01}{L} + \frac{c_1^*(1-\kappa_s)}{U} \right) \log(P) \right\}.$$

for some small constant $c_1^* > 0$.

For computational convenience, we make the following assumptions for κ_c and κ_s , which can be easily relaxed with the cost of additional calculations.

$$\frac{\Delta}{2} \left(\frac{1}{L} - \frac{1}{U} \right) + \frac{0.01}{L} + \frac{0.01}{U} \leq \frac{c_0^*(1-\kappa_s)}{U} \quad (22a)$$

$$\left(1 - \frac{c_1^*L}{U} \right) (1 - \kappa_s) \leq (1 - \kappa_c) + \frac{U}{L} \left(\frac{\Delta}{2} + 0.01 \right) \quad (22b)$$

Here c_0^* is some small. We state the following induction hypotheses, which will hold throughout this period:

Induction Hypothesis H.2. For each $T_1 < t \leq \tilde{T}_1$, $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, the following holds:

- $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \in [0, \frac{c_0^* + c_1^*}{U} \log(P)]$;
- $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically decreasing and $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in [-\frac{1}{L} (\frac{\Delta}{2} + 0.01) \log(P), -\frac{1}{U} (\frac{\Delta}{2} - 0.01) \log(P)]$;
- $|\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right)$ for $m \neq 1, n$;
- $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ for $a_{k,\mathbf{q}} = n$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$;
- $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{|\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right) + O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} = 1$;
- $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} \neq 1, n$.

H.2.1 PROPERTY OF ATTENTION SCORES

We first single out several properties of attention scores that will be used for the proof of Induction Hypothesis H.2.

Lemma H.10. *if Induction Hypothesis G.1 and Induction Hypothesis H.2 hold at iteration $T_1 + 1 \leq t \leq \tilde{T}_1$, then the following holds*

- $1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \geq \Omega(1)$;
- if $M \in \mathcal{E}_{k,n}$, $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \in \left[\Omega\left(\frac{1}{P^{1-\kappa_s}}\right), O\left(\frac{1}{P^{(1-c_1^*-c_0^*)(1-\kappa_s)}}\right) \right]$;
- Moreover, $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} = O\left(\frac{1}{P^{(1-\kappa_c) + \frac{U}{L}(\frac{\Delta}{2} - 0.01)}}\right)$; if $M \in \mathcal{E}_{k,1}$, we have $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} = \Omega\left(\frac{1}{P^{(1-\kappa_c) + \frac{U}{L}(\frac{\Delta}{2} + 0.01)}}\right)$;
- for $\mathbf{q} \in \mathcal{M} \cap (\mathcal{P}_{k,n} \cup \mathcal{P}_{k,1})$, $\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}}{P}\right)$.

Lemma H.11. *if Induction Hypothesis G.1 and Induction Hypothesis H.2 hold at iteration $T_1 + 1 \leq t \leq \tilde{T}_1$, then for $m \neq n$, the following holds:*

- for any $\mathbf{q} \in \mathcal{P}_{k,m}$, $\text{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \leq O\left(\frac{1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}}{P}\right)$;
- Moreover, $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \leq O\left(\frac{1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{N}\right)$.

H.2.2 BOUNDING THE GRADIENT UPDATES OF FP CORRELATIONS

Lemma H.12. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.2 hold at iteration $T_1 + 1 \leq t \leq \tilde{T}_1$, then $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq 0$ and satisfies:

$$\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} = \Omega\left(\frac{1}{P^{1-\kappa_s}}\right).$$

Proof. By Lemma E.2, we have

$$\begin{aligned} & \alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &\quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n}^c \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &\geq \mathbb{P}(\mathbf{M} \in \mathcal{E}_{k,n}) \\ &\quad \cdot \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \middle| \mathcal{E}_{k,n} \right] \\ &\geq \Omega\left(\frac{C_n}{P}\right) \end{aligned}$$

where the last inequality invokes Lemma H.10. □

Lemma H.13. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.2 hold at iteration $T_1 + 1 \leq t \leq \tilde{T}_1$, then $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} < 0$ and satisfies

$$|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| \geq \Omega\left(\frac{1}{P^{2(1-\kappa_c)} + \frac{L}{L}(\Delta + 0.02)}\right).$$

Proof. Following (17), we have

$$\begin{aligned} & -z_1 z_n^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - z_1^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} + \sum_{a \neq 1,n} z_a^2 z_1 \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \\ & \leq -z_1 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \right) \left(z_n^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_1^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \max_{a \neq 1,n} z_a^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right) \end{aligned}$$

Therefore, by Lemma E.1, we obtain

$$\begin{aligned} \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \right. \\ &\quad \left(-z_1 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \right) \right. \\ &\quad \left. \cdot \left(z_n^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_1^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \max_{a \neq 1,n} z_a^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right) \right) \left. \right] \\ &\quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1}^c \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \sum_{a \neq 1,n} z_1^2 z_a \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right] \end{aligned}$$

$$\begin{aligned} &\leq \mathbb{P}(M \in \mathcal{E}_{k,1}) \cdot \left(-\Omega(1) \cdot \Omega\left(\frac{1}{P^{2(1-\kappa_c) + \frac{2U}{L}(\frac{\Delta}{2} + 0.01)}}\right) \right) + O(1) \cdot \mathbb{P}(M \in \mathcal{E}_{k,1}^c) \\ &\leq -\Omega\left(\frac{1}{P^{2(1-\kappa_c) + \frac{U}{L}(\Delta + 0.02)}}\right) \end{aligned}$$

where the second inequality invokes Lemma H.10 and the last inequality comes from Lemma F.6. The upper bound can be obtained by using similar arguments and invoking the upper bound for $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}$ in Lemma H.10. \square

Lemma H.14. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.2 hold at iteration $T_1 + 1 \leq t \leq \tilde{T}_1$, then for any $m > 1$ with $m \neq n$, the following holds

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right).$$

The proof is similar to Lemma H.5, and thus omitted here.

H.2.3 BOUNDING THE GRADIENT UPDATES OF POSITIONAL CORRELATIONS

We then summarize the properties for gradient updates of positional correlations, which utilize the identical calculations as in Section H.1.3.

Lemma H.15. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.2 hold at iteration $T_1 + 1 \leq t \leq \tilde{T}_1$, then

- if $a_{k,\mathbf{q}} = n$ and $\mathbf{q} \neq \mathbf{p}$, $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \geq 0$; $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ and $|\beta_n^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$.
- if $a_{k,\mathbf{q}} = 1$, $|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right) + O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right)$.
- if $a_{k,\mathbf{q}} = m$ and $m \neq 1, n$, $|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$.

H.2.4 END OF PHASE I, STAGE 2

Lemma H.16. Induction Hypothesis H.2 holds for all iteration $T_1 + 1 \leq t \leq \tilde{T}_1 = T_1 + O\left(\frac{\log(P)P^{1-\kappa_s}}{\eta}\right)$, and at iteration $t = \tilde{T}_1 + 1$, we have

- $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(\tilde{T}_1+1)} \geq \frac{c_1^*(1-\kappa_s)\log(P)}{U}$;
- $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(\tilde{T}_1+1)} \geq -\left(\frac{\Delta}{2L} + \frac{0.01}{L}\right)\log(P)$.

Proof. The existence of $\tilde{T}_1 = T_1 + O\left(\frac{\log(P)P^{1-\kappa_s}}{\eta}\right)$ directly follows from Lemma H.12 and Lemma H.13. Moreover, since $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} < 0$, then

$$\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(\tilde{T}_1+1)} \leq \left(\frac{\Delta}{2L} + \frac{0.01}{L} + \frac{c_1^*(1-\kappa_s)}{U}\right)\log(P) - \frac{1}{U}\left(\frac{\Delta}{2} - 0.01\right) \leq \frac{(c_0^* + c_1^*)(1-\kappa_s)}{U}\log(P)$$

where the last inequality invokes (22a). Now suppose $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(\tilde{T}_1+1)} < \frac{c_1^*(1-\kappa_s)\log(P)}{U}$, then $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(\tilde{T}_1+1)} < -\left(\frac{\Delta}{2L} + \frac{0.01}{L}\right)\log(P)$. Denote the first time that $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ reaches $-\left(\frac{\Delta}{2L} + \frac{0.001}{L}\right)\log(P)$ as \tilde{T} . Note that $\tilde{T} < \tilde{T}_1$ since $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$, the change of $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$, satisfies $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| \ll \log(P)$. Then for $t \geq \tilde{T}$, the following holds:

- $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \geq \Omega\left(\frac{1}{P^{1-\kappa_s}}\right)$;

$$2. \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \leq O\left(\frac{1}{P^{\frac{1-\kappa_s}{2}+0.001}}\right).$$

Therefore,

$$\begin{aligned} |\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \right. \\ &\quad \left. z_1 \left(z_n^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) + z_1^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} (1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \right) \right] \\ &\quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1}^c \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \sum_{a \neq 1,n} z_1^2 z_a \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right)^2 \right] \\ &\leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{\frac{1-\kappa_s}{2}+0.001}}\right) + \mathbb{P}(\mathbf{M} \in \mathcal{E}_{k,1}) \cdot \left(O(1) \cdot O\left(\frac{1}{P^{\frac{1-\kappa_s}{2}+0.001}}\right) \right) + O(1) \cdot \mathbb{P}(\mathbf{M} \in \mathcal{E}_{k,1}^c) \\ &\leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{\frac{1-\kappa_s}{2}+0.001}}\right) + O\left(\frac{1}{P^{(1-\kappa_s)+0.002}}\right). \end{aligned}$$

Lemma H.12 still holds, and thus

$$|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| \leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{0.002}}\right).$$

Since $|\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(\tilde{T}_1+1)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(\tilde{T})}| \geq \Omega(\log(P))$, we have

$$\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(\tilde{T}_1+1)} \geq |\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(\tilde{T}_1+1)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(\tilde{T})}| \cdot \Omega(P^{0.002}) + \Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(\tilde{T})} \gg \Omega(P^{0.002} \log(P)),$$

which contradicts the assumption that $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(\tilde{T}_1+1)} < \frac{c_1^*(1-\kappa_s) \log(P)}{U}$. □

H.3 Phase II, Stage 1

For $n > 1$, we define stage 1 of phase II as all iterations $\tilde{T}_1 + 1 \leq t \leq T_2$, where

$$T_2 \triangleq \max \left\{ t : \Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \leq \frac{(1-\kappa_s)}{L} \log(P) \right\}.$$

We state the following induction hypotheses, which will hold throughout this stage:

Induction Hypothesis H.3. For each $\tilde{T}_1 + 1 \leq t \leq T_2$, $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, the following holds:

- a. $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \in \left[\frac{c_1^*(1-\kappa_s)}{U} \log(P), \frac{(1-\kappa_s)}{L} \log(P) \right]$;
- b. $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically decreasing and

$$\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in \left[-\frac{1}{L} \left(\frac{\Delta}{2} + 0.01 \right) \log(P) - o(1), -\frac{1}{U} \left(\frac{\Delta}{2} - 0.01 \right) \log(P) \right];$$

- c. $|\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right)$ for $m \neq 1, n$;
- d. $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ for $a_{k,\mathbf{q}} = n$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$;
- e. $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{|\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right) + O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} = 1$;
- f. $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} \neq 1, n$.

H.3.1 PROPERTY OF ATTENTION SCORES

We first single out several properties of attention scores that will be used for the proof of Induction Hypothesis H.3.

Lemma H.17. *if Induction Hypothesis G.1 and Induction Hypothesis H.3 hold at iteration $\tilde{T}_1 + 1 \leq t \leq T_2$, then the following holds*

1. if $M \in \mathcal{E}_{k,n}$, $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \geq \Omega\left(\frac{1}{P(1-\frac{c_X^* L}{U})(1-\kappa_s)}\right)$. Moreover, if $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}$ does not reach the constant level, $1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} = \Omega(1)$; otherwise, $1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} = \Omega\left(\frac{1}{P(\frac{U}{L}-1)(1-\kappa_s)}\right)$.
2. $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} = O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{P(1-\kappa_c) + \frac{U}{L}(\frac{\Delta}{2} - 0.01)}\right)$; if $M \in \mathcal{E}_{k,1}$, we have $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} = \Omega\left(\frac{1}{P(1-\kappa_c) + \frac{U}{L}(\frac{\Delta}{2} + 0.01)}\right)$;
3. for $\mathbf{q} \in \mathcal{M} \cap (\mathcal{P}_{k,n} \cup \mathcal{P}_{k,1})$, $\mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{P}\right)$

Lemma H.18. *if Induction Hypothesis G.1 and Induction Hypothesis H.3 hold at iteration $\tilde{T}_1 + 1 \leq t \leq T_2$, then for $m \neq n$, the following holds:*

1. for any $\mathbf{q} \in \mathcal{P}_{k,m}$, $\mathbf{attn}_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \leq O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{P}\right)$.
2. Moreover, $\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \leq O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}}{N}\right)$.

H.3.2 BOUNDING THE GRADIENT UPDATES OF FP CORRELATIONS

Lemma H.19. *if Induction Hypothesis G.1 and Induction Hypothesis H.3 hold at iteration $\tilde{T}_1 + 1 \leq t \leq T_2$, then $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq 0$ and satisfies:*

$$\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq \min \left\{ \Omega\left(\frac{1}{P(1-\frac{c_X^* L}{U})(1-\kappa_s)}\right), \Omega\left(\frac{1}{P^2(\frac{U}{L}-1)(1-\kappa_s)}\right) \right\}.$$

Proof. By Lemma E.2, we have

$$\begin{aligned} & \alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ & \quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n}^c \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ & \gtrsim \mathbb{P}(M \in \mathcal{E}_{k,n}) \cdot \\ & \quad \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \middle| \mathcal{E}_{k,n} \right] \\ & \quad + O(1) \cdot \mathbb{P}(M \in \mathcal{E}_{k,n}^c) \\ & \gtrsim \min \left\{ \Omega\left(\frac{1}{P(1-\frac{c_X^* L}{U})(1-\kappa_s)}\right), \Omega\left(\frac{1}{P^2(\frac{U}{L}-1)(1-\kappa_s)}\right) \right\} \end{aligned}$$

where the last inequality invokes Lemma H.17 by observing that for $M \in \mathcal{E}_{k,n}$,

$$\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)})^2 \geq \min \left\{ \Omega \left(\frac{1}{P(1 - \frac{c_s^* L}{U})(1 - \kappa_s)} \right), \Omega(1), \Omega(1) \cdot \Omega \left(\frac{1}{P2 \times (\frac{U}{L} - 1)(1 - \kappa_s)} \right) \right\}.$$

□

Lemma H.20. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.3 hold at iteration $\tilde{T}_1 + 1 \leq t \leq T_2$, then $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} < 0$ and satisfies

$$\begin{aligned} |\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| &\geq \min \left\{ \Omega \left(\frac{1}{P(1 - \frac{c_s^* L}{U})(1 - \kappa_s)} \right), \Omega \left(\frac{1}{P(\frac{U}{L} - 1)(1 - \kappa_s)} \right) \right\} \cdot \Omega \left(\frac{1}{P(1 - \kappa_c) + \frac{L}{U}(\frac{\Delta}{2} - 0.01)} \right), \\ |\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| &\leq \max \left\{ O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P(1 - \kappa_c) + \frac{L}{U}(\Delta/2 - 0.01)} \right), O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P2(1 - \kappa_c) + \frac{L}{U}(\Delta - 0.02) - (1 - \frac{c_s^* L}{U})(1 - \kappa_s)} \right) \right\}. \end{aligned}$$

Proof. Following (17), we have

$$\begin{aligned} &-z_1 z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - z_1^3 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} + \sum_{a \neq 1, n} z_a^2 z_1 (\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)})^2 \\ &\leq -z_1 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \left(z_n^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_1^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \max_{a \neq 1, n} z_a^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right). \end{aligned}$$

Therefore, by Lemma E.1, we obtain

$$\begin{aligned} \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} &\leq \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \right. \\ &\quad \left. \left(-z_1 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \left(z_n^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_1^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} - \max_{a \neq 1, n} z_a^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} \right) \right) \right] \\ &\quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1}^c \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \sum_{a \neq 1, n} z_a^2 z_1 (\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)})^2 \right] \\ &\leq -\min \left\{ \Omega \left(\frac{1}{P(1 - \frac{c_s^* L}{U})(1 - \kappa_s)} \right), \Omega \left(\frac{1}{P(\frac{U}{L} - 1)(1 - \kappa_s)} \right) \right\} \cdot \Omega \left(\frac{1}{P(1 - \kappa_c) + \frac{L}{U}(\frac{\Delta}{2} - 0.01)} \right) \end{aligned}$$

where the second inequality invokes Lemma H.17 and (22b). Moreover,

$$\begin{aligned} |\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| &\lesssim \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot \right. \\ &\quad \left. \left(z_1 z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} + z_1^3 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \right) \right] \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} z_1 z_n^2 \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} \cdot (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} \right] \\ &\quad + \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,1} \cap \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} z_1^3 (\mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)})^2 \cdot (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)}) \right] \\ &\leq \max \left\{ O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P(1 - \kappa_c) + \frac{L}{U}(\frac{\Delta}{2} - 0.01)} \right), O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P2(1 - \kappa_c) + \frac{2L}{U}(\frac{\Delta}{2} - 0.01) - (1 - \frac{c_s^* L}{U})(1 - \kappa_s)} \right) \right\} \end{aligned}$$

where the second inequality invokes Lemma H.17. □

Lemma H.21. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.3 hold at iteration $\tilde{T}_1 + 1 \leq t \leq T_2$ for any $m > 1$ with $m \neq n$, the following holds

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1 - \kappa_s}} \right).$$

The proof is similar to Lemma H.5, and thus omitted here.

H.3.3 BOUNDING THE GRADIENT UPDATES OF POSITIONAL CORRELATIONS

We then summarize the properties for gradient updates of positional correlations, which utilizes the identical calculations as in Section H.1.3.

Lemma H.22. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.3 hold at iteration $\tilde{T}_1 + 1 \leq t \leq T_2$, then

- if $a_{k,\mathbf{q}} = n$ and $\mathbf{q} \neq \mathbf{p}$, $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \geq 0$; $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ and $|\beta_n^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$.
- if $a_{k,\mathbf{q}} = 1$, $|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right) + O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right)$.
- if $a_{k,\mathbf{q}} = m$ and $m \neq 1, n$, $|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$.

H.3.4 END OF PHASE II, STAGE 1

Lemma H.23. Induction Hypothesis H.3 holds for all $\tilde{T}_1 + 1 \leq t \leq T_2$, and at iteration $t = T_2 + 1$, we have

- $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} > \frac{(1-\kappa_s)}{L} \log(P)$;
- $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} = \Omega(1)$ if $\mathbf{M} \in \mathcal{E}_{k,n}$.

Proof. By comparing Lemma H.19 and Lemma H.20-H.23, we have $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \gg |\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}|, |\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}|$. Then the existence of $T_2 = \tilde{T}_1 + O\left(\frac{\log(P)P^\Lambda}{\eta}\right)$ directly follows from Lemma H.19, where

$$\Lambda = \max \left\{ \left(1 - \frac{c_1^* L}{U}\right), 2\left(\frac{U}{L} - 1\right) \right\} \cdot (1 - \kappa_s).$$

The second statement can be directly verified by noticing that $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} > \frac{(1-\kappa_s)}{L} \log(P)$ while all other attention correlations are sufficiently small. \square

H.4 Phase II, Stage 2

In this final stage, we establish that these structures indeed represent the solutions toward which the algorithm converges.

Given any $0 < \epsilon < 1$, for $n > 1$, define

$$T_2^\epsilon \triangleq \max \left\{ t > T_2 : \Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \leq \log \left(c_5 \left(\left(\frac{3}{\epsilon} \right)^{\frac{1}{2}} - 1 \right) N \right) \right\}. \quad (23)$$

where c_5 is some largely enough constant.

We state the following induction hypotheses, which will hold throughout this stage:

Induction Hypothesis H.4. For $n > 1$, suppose $\text{polylog}(P) \gg \log(\frac{1}{\epsilon})$, for each $T_2 + 1 \leq t \leq T_2^\epsilon$, $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, the following holds:

- $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \in [\frac{(1-\kappa_s)}{L} \log(P), O(\log(P/\epsilon))]$;
- $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically decreasing and $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in \left[-\frac{1}{L} \left(\frac{\Delta}{2} + 0.01 \right) \log(P) - o(1), -\frac{1}{U} \left(\frac{\Delta}{2} - 0.01 \right) \log(P) \right]$;
- $|\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right)$ for $m \neq 1, n$;
- $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ for $a_{k,\mathbf{q}} = n$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$;

- e. $|\Upsilon_{k,\mathbf{p}\rightarrow\mathbf{q}}^{(t)}| = O\left(\frac{|\Phi_{\mathbf{p}\rightarrow v_{k,1}}^{(t)}|}{C_1}\right) + O\left(\frac{\Phi_{\mathbf{p}\rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p}\rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} = 1$;
- f. $|\Upsilon_{k,\mathbf{p}\rightarrow\mathbf{q}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p}\rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p}\rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} \neq 1, n$.

H.4.1 PROPERTY OF ATTENTION SCORES

We first single out several properties of attention scores that will be used for the proof of Induction Hypothesis H.4.

Lemma H.24. *if Induction Hypothesis G.1 and Induction Hypothesis H.4 hold at iteration $T_{n,2} < t \leq T_{n,2}^\epsilon$, then the following holds*

1. if $M \in \mathcal{E}_{k,n}$, $\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} = \Omega(1)$ and $(1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)})^2 \geq O(\epsilon)$.
2. Moreover, $\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,1}}^{(t)} = O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)}}{P^{(1-\kappa_c) + \frac{t}{T}(\frac{\epsilon}{2} - 0.01)}}\right)$; if $M \in \mathcal{E}_{k,1}$, we have $\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,1}}^{(t)} = \Omega\left(\frac{1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)}}{P^{(1-\kappa_c) + \frac{t}{T}(\frac{\epsilon}{2} + 0.01)}}\right)$;
3. for $\mathbf{q} \in \mathcal{M} \cap (\mathcal{P}_{k,n} \cup \mathcal{P}_{k,1})$, $\mathbf{attn}_{\mathbf{p}\rightarrow\mathbf{q}}^{(t)} = O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)}}{P}\right)$.

Lemma H.25. *if Induction Hypothesis G.1 and Induction Hypothesis H.4 hold at iteration $T_{n,2} < t \leq T_{n,2}^\epsilon$, then for $m \neq n$, the following holds:*

1. for any $\mathbf{q} \in \mathcal{P}_{k,m}$, $\mathbf{attn}_{\mathbf{p}\rightarrow\mathbf{q}}^{(t)} \leq O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)}}{P}\right)$.
2. $\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \leq O\left(\frac{1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)}}{N}\right)$, and if $M \in \mathcal{E}_{k,m}$, $\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} = \Theta\left(\frac{1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)}}{N}\right)$.

H.4.2 BOUNDING THE GRADIENT UPDATES OF FP CORRELATIONS

Lemma H.26. *For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.4 hold at iteration $T_2 + 1 \leq t \leq T_2^\epsilon$, then $\alpha_{\mathbf{p}\rightarrow v_{k,n}}^{(t)} \geq 0$ and satisfies:*

$$\alpha_{\mathbf{p}\rightarrow v_{k,n}}^{(t)} \geq \Omega(\epsilon).$$

Proof. By Lemma E.2, we have

$$\begin{aligned} & \alpha_{\mathbf{p}\rightarrow v_{k,n}}^{(t)} \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &= \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &+ \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n}^c \cap \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \right] \\ &\succeq \mathbb{P}(M \in \mathcal{E}_{k,n}) \cdot \\ & \mathbb{E} \left[\mathbb{1}\{k_X = k, \mathbf{p} \in \mathcal{M}\} \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \cdot \left(z_n^3 \left(1 - \mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,n}}^{(t)} \right)^2 + \sum_{m \neq n} z_m^2 z_n \left(\mathbf{Attn}_{\mathbf{p}\rightarrow\mathcal{P}_{k,m}}^{(t)} \right)^2 \right) \middle| \mathcal{E}_{k,n} \right] \\ &+ O(1) \cdot \mathbb{P}(M \in \mathcal{E}_{k,n}^c) \end{aligned}$$

$$\gtrsim \Omega(\epsilon)$$

where the last inequality invokes Lemma H.24, Lemma F.6 and the fact that

$$\epsilon \geq \exp(-\text{polylog}(K)) \gg \exp(-c_{n,1}C_n).$$

□

Lemma H.27. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.4 hold at iteration $T_{n,3} < t \leq T_{n,4}^\epsilon$, then $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} < 0$ and satisfies

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq \max \left\{ O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{(1-\kappa_c) + \frac{L}{U}(\Delta/2 - 0.01)}}\right), O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{2(1-\kappa_c) + \frac{L}{U}(\Delta - 0.02) - (1 - \frac{c_1^* L}{U})(1-\kappa_s)}}\right) \right\}$$

The proof follows the similar arguments Lemma H.20 by noticing that $\epsilon \gg \mathbb{P}(M \in \mathcal{E}_{k,m}^c)$ for any $m \neq n$.

Lemma H.28. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.4 hold at iteration $T_2 < t \leq T_2^\epsilon$, then for any $m > 1$ with $m \neq n$, the following holds

$$-O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{1-\kappa_s}}\right) \leq \alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} \leq 0$$

Proof. We first note that

$$\begin{aligned} & -z_1 z_n^2 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}\right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - z_m^3 \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}\right) \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(t)} + \sum_{a \neq 1, n} z_a^2 z_m \left(\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)}\right)^2 \\ & \leq z_m \left(\max_{a \neq m, n} z_a^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)} - z_n^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - z_m^2 \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}\right) \left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)}\right) \\ & \lesssim -\Omega(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \end{aligned}$$

since when $M \in \mathcal{E}_{k,n}$, we have $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)} = \Omega(1) \gg \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a}}^{(t)}$. Thus, we have

$$\begin{aligned} 0 \geq \alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} & \gtrsim -\mathbb{E} \left[\mathbb{1}\{k_X = k, \mathcal{E}_{k,n} \cap \mathbf{p} \in \mathcal{M}\} \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(t)} \cdot \Omega(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(t)}) \right] \\ & \geq -O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{1-\kappa_s}}\right). \end{aligned}$$

□

H.4.3 BOUNDING THE GRADIENT UPDATES OF POSITIONAL CORRELATIONS

We then summarize the properties for gradient updates of positional correlations, which utilizes the identical calculations as in Section H.1.3.

Lemma H.29. For $n > 1$, if Induction Hypothesis G.1 and Induction Hypothesis H.4 hold at iteration $T_2 + 1 \leq t \leq T_2^\epsilon$, then

- if $a_{k,\mathbf{q}} = n$ and $\mathbf{q} \neq \mathbf{p}$, $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \geq 0$; $\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ and $|\beta_n^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$.
- if $a_{k,\mathbf{q}} = 1$, $|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right) + O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right)$.
- if $a_{k,\mathbf{q}} = m$ and $m \neq 1, n$, $|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$.

H.4.4 END OF PHASE II, STAGE 2

Lemma H.30. For $n > 1$, and $0 < \epsilon < 1$, suppose $\text{polylog}(P) \gg \log(\frac{1}{\epsilon})$. Then Induction Hypothesis H.4 holds for all $T_2 < t \leq T_2^\epsilon = T_2 + O\left(\frac{\log(P\epsilon^{-1})}{\eta\epsilon}\right)$, and at iteration $t = T_2^\epsilon + 1$, we have

1. $\tilde{\mathcal{L}}_{k,\mathbf{p}}(Q^{T_2^\epsilon+1}) < \frac{\epsilon}{2K}$;
2. If $M \in \mathcal{E}_{k,n}$, we have $(1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(T_2^\epsilon+1)})^2 \leq O(\epsilon)$.

Proof. The existence of $T_{2,k}^\epsilon = T_{2,k} + O\left(\frac{\log(P\epsilon^{-1})}{\eta\epsilon}\right)$ directly follows from Lemma H.26. We further derive

$$\begin{aligned} \tilde{\mathcal{L}}_{k,\mathbf{p}}(Q^{T_2^\epsilon+1}) &= \\ & \frac{1}{2} \mathbb{E} \left[\mathbf{1}\{k_X = k, \mathbf{p} \in \mathcal{M} \cap M \in \mathcal{E}_{k,n}\} \left(z_n^2 (1 - \mathbf{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}})^2 + \sum_{m \neq n} z_m^2 (\mathbf{Attn}_{n,m})^2 \right) \right] \\ & \leq \frac{1}{2K} \cdot \gamma \cdot U^2 \cdot (1 + o(1)) \cdot O(\epsilon) \\ & \leq \frac{\epsilon}{2K} \end{aligned}$$

where the first inequality is due to direct calculations by the definition of T_2^ϵ , and the second inequality can be obtained by setting $c_{n,2}$ in (23) sufficiently large. \square

I Analysis for Local Areas with Negative Information Gap

In this section, we focus on a specific patch $\mathbf{p} \in \mathcal{P}$ with the k -th cluster for $k \in [K]$, and present the analysis for the case that $X_{\mathbf{p}}$ is located in the local area for the k -th cluster, i.e. $a_{k,\mathbf{p}} > 1$. Throughout this section, we denote $a_{k,\mathbf{p}} = n$ for simplicity. When $\Delta \leq -\Omega(1)$, we can show that the gap of attention correlation changing rate for the positive case does not exist anymore, and conversely $\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \gg \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ from the beginning. We can reuse most of the gradient calculations in the previous section and only sketch them in this section.

Stage 1: we define stage 1 as all iterations $0 \leq t \leq T_{\text{neg},1}$, where

$$T_{\text{neg},1} \triangleq \max \left\{ t : \Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \leq \frac{(1 - \kappa_s)}{L} \log(P) \right\}.$$

We state the following induction hypothesis, which will hold throughout this stage:

Induction Hypothesis I.1. For each $0 \leq t \leq T_{\text{neg},1}$, $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, the following holds:

- a. $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \in \left[0, \frac{(1 - \kappa_s)}{L} \log(P)\right]$;
- b. $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically decreasing and $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in \left[-O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P^{-\Delta}}\right), 0\right]$;
- c. $|\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right)$ for $m \neq 1, n$;
- d. $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right)$ for $a_{k,\mathbf{q}} = n$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$;
- e. $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{|\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right) + O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} = 1$;
- f. $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} \neq 1, n$.

Through similar calculations for phase II, stage 1 in Appendix H.3, we obtain the following lemmas to control the gradient updates for attention correlations.

Lemma I.1. *If Induction Hypothesis G.2 and Induction Hypothesis I.1 hold for $0 \leq t \leq T_{\text{neg},1}$, then we have*

$$\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq \min \left\{ \Omega \left(\frac{1}{P^{1-\kappa_s}} \right), \Omega \left(\frac{1}{P^{2(\frac{U}{L}-1)(1-\kappa_s)}} \right) \right\}, \quad (24a)$$

$$0 \geq \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \geq -O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P-\Delta} \right), \quad (24b)$$

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}} \right) \text{ for all } m \neq n, 1 \quad (24c)$$

$$\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n} \right) \text{ for } a_{k,\mathbf{q}} = n, \mathbf{q} \neq \mathbf{p} \quad (24d)$$

$$|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P} \right) + O \left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1} \right) \text{ for } a_{k,\mathbf{q}} = 1, \quad (24e)$$

$$|\beta_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}|, |\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right) \text{ for all } a_{k,\mathbf{p}} \neq n, 1. \quad (24f)$$

Here $\Delta < 0$ implies $|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}| \ll \alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$. Induction Hypothesis I.1 can be directly proved by Lemma I.1 and we have

$$T_{\text{neg},1} = O \left(\frac{P^{\max\{1, 2(\frac{U}{L}-1)\} \cdot (1-\kappa_s)} \log(P)}{\eta} \right). \quad (25)$$

Stage 2: Given any $0 < \epsilon < 1$, define

$$T_{\text{neg},1}^\epsilon \triangleq \max \left\{ t > T_1 : \Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \leq \log \left(c_6 \left(\left(\frac{3}{\epsilon} \right)^{\frac{1}{2}} - 1 \right) P^{1-\kappa_s} \right) \right\}. \quad (26)$$

where c_6 is some largely enough constant. We then state the following induction hypotheses, which will hold throughout this stage:

Induction Hypothesis I.2. For $n > 1$, suppose $\text{polylog}(P) \gg \log(\frac{1}{\epsilon})$, for $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, and each $T_{\text{neg},1} < t \leq T_{\text{neg},1}^\epsilon$, the following holds:

- $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \in \left[\frac{(1-\kappa_s)}{L} \log(P), O(\log(P/\epsilon)) \right]$;
- $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically decreasing and $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in \left[-O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P-\Delta} \right), 0 \right]$;
- $|\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| = O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}} \right)$ for $m \neq 1, n$;
- $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n} \right)$ for $a_{k,\mathbf{q}} = n$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right)$;
- $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O \left(\frac{|\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1} \right) + O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right)$ for $a_{k,\mathbf{q}} = 1$;
- $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right)$ for $a_{k,\mathbf{q}} \neq 1, n$.

Lemma I.2. *If Induction Hypothesis G.2 and Induction Hypothesis I.2 hold for $T_{\text{neg},1} < t \leq T_{\text{neg},1}^\epsilon$, then we have*

$$\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} \geq \Omega(\epsilon), \quad (27a)$$

$$0 \geq \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \geq -O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P-\Delta}\right), \quad (27b)$$

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}}\right) \text{ for all } m \neq n, 1 \quad (27c)$$

$$\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{C_n}\right) \text{ for } a_{k,\mathbf{q}} = n, \mathbf{q} \neq \mathbf{p} \quad (27d)$$

$$|\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}}{P}\right) + O\left(\frac{|\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|}{C_1}\right) \text{ for } a_{k,\mathbf{q}} = 1, \quad (27e)$$

$$|\beta_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}|, |\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,n}}^{(t)} - \alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right) \text{ for all } a_{k,\mathbf{p}} \neq n, 1. \quad (27f)$$

Induction Hypothesis I.2 can be directly proved by Lemma I.2. Furthermore, at the end of this stage, we will have:

Lemma I.3. *Suppose $\text{polylog}(P) \gg \log(\frac{1}{\epsilon})$, then Induction Hypothesis I.2 holds for all $T_{\text{neg},1} < t \leq T_{\text{neg},1}^\epsilon = T_{\text{neg},1} + O\left(\frac{\log(P\epsilon^{-1})}{\eta\epsilon}\right)$, and at iteration $t = T_{\text{neg},1}^\epsilon + 1$, we have*

1. $\tilde{\mathcal{L}}_{k,\mathbf{p}}(Q^{T_{\text{neg},1}^\epsilon+1}) < \frac{\epsilon}{2K}$;
2. If $M \in \mathcal{E}_{k,n}$, we have $\left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,n}}^{(T_{\text{neg},1}^\epsilon+1)}\right)^2 \leq O(\epsilon)$.

J Analysis for the Global area

When $a_{\mathbf{p},k} = 1$, i.e. the patch lies in the global area, the analysis is much simpler and does not depend on the value of Δ . We can reuse most of the gradient calculations in Appendix H and only sketch them in this section.

For $X_{\mathbf{p}}$ in the global region $\mathcal{P}_{k,1}$, since the overall attention $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(0)}$ to the target feature already reaches $\Omega\left(\frac{C_1}{P}\right) = \Omega\left(\frac{1}{P^{1-\kappa_c}}\right)$ due to the large number of unmasked patches featuring $v_{k,1}$ when $M \in \mathcal{E}_{k,1}$, which is significantly larger than $\text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,m}}^{(0)} = \Theta\left(\frac{1}{P^{1-\kappa_s}}\right)$ for all other $m > 1$. This results in large $\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ initially, and thus the training directly enters phase II.

Stage 1: we define stage 1 as all iterations $0 \leq t \leq T_{c,1}$, where

$$T_{c,1} \triangleq \max \left\{ t : \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \leq \frac{(1-\kappa_c)}{L} \log(P) \right\}.$$

We state the following induction hypotheses, which will hold throughout this stage:

Induction Hypothesis J.1. For each $0 \leq t \leq T_{c,1}$, $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, the following holds:

- a. $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in \left[0, \frac{(1-\kappa_c)}{L} \log(P)\right]$;
- b. $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ is monotonically decreasing for $m > 1$ and $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} \in \left[-O\left(\frac{\log(P)}{N}\right), 0\right]$;
- c. $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{C_1}\right)$ for $a_{k,\mathbf{q}} = 1$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$;
- d. $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O\left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P}\right)$ for $a_{k,\mathbf{q}} \neq 1$.

Through similar calculations for phase II, stage 1 in Appendix H.3, we obtain the following lemmas to control the gradient updates for attention correlations.

Lemma J.1. *If Induction Hypothesis G.1 (or Induction Hypothesis G.2) and Induction Hypothesis J.1 hold for $0 \leq t \leq T_{c,1}$, then we have*

$$\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \geq \min \left\{ \Omega \left(\frac{1}{P^{(1-\kappa_c)}} \right), \Omega \left(\frac{1}{P^{2(\frac{U}{L}-1)(1-\kappa_c)}} \right) \right\}, \quad (28a)$$

$$|\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}} \right) \quad \text{for all } m \neq 1, \quad (28b)$$

$$\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{C_1} \right), \text{ for } a_{k,\mathbf{q}} = 1, \mathbf{q} \neq \mathbf{p}, \quad (28c)$$

$$|\beta_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}|, |\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right) \quad \text{for all } a_{k,\mathbf{q}} > 1. \quad (28d)$$

Induction Hypothesis J.1 can be directly proved by Lemma J.1 and we have

$$T_{c,1} = O \left(\frac{P^{\max\{1, 2(\frac{U}{L}-1)\} \cdot (1-\kappa_c)} \log(P)}{\eta} \right). \quad (29)$$

Stage 2: Given any $0 < \epsilon < 1$, define

$$T_{c,1}^\epsilon \triangleq \max \left\{ t > T_{c,1} : \Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \leq \log \left(c_7 \left(\left(\frac{3}{\epsilon} \right)^{\frac{1}{2}} - 1 \right) P^{1-\kappa_c} \right) \right\}. \quad (30)$$

where c_7 is some largely enough constant. We then state the following induction hypotheses, which will hold throughout this stage:

Induction Hypothesis J.2. For $n > 1$, suppose $\text{polylog}(P) \gg \log(\frac{1}{\epsilon})$, $\mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$, for each $T_{c,1} + 1 \leq t \leq T_{c,1}^\epsilon$, the following holds:

- $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ is monotonically increasing, and $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \in \left[\frac{(1-\kappa_c)}{L} \log(P), O(\log(P/\epsilon)) \right]$;
- $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$ is monotonically decreasing for $n > 1$ and $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)} \in \left[-O(\frac{\log(P)}{N}), 0 \right]$;
- $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{C_1} \right)$ for $a_{k,\mathbf{q}} = 1$, $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}| = O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right)$;
- $|\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O \left(\frac{\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right)$ for $a_{k,\mathbf{q}} \neq 1$.

We also have the following lemmas to control the gradient updates for attention correlations.

Lemma J.2. *If Induction Hypothesis G.1 (or Induction Hypothesis G.2) and Induction Hypothesis J.1 hold for $T_{c,1} + 1 \leq t \leq T_{c,1}^\epsilon$, then we have*

$$\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)} \geq \Omega(\epsilon), |\alpha_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}| \leq O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P^{1-\kappa_s}} \right) \quad \text{for all } m \neq 1 \quad (31a)$$

$$\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \Theta \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{C_1} \right), \text{ for } a_{k,\mathbf{q}} = 1, \mathbf{q} \neq \mathbf{p} \quad (31b)$$

$$|\beta_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)}|, |\beta_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}| = O \left(\frac{\alpha_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}}{P} \right) \quad \text{for all } a_{k,\mathbf{q}} > 1. \quad (31c)$$

Induction Hypothesis J.2 can be directly proved by Lemma J.2. Furthermore, at the end of this stage, we will have:

Lemma J.3. *Suppose $\text{polylog}(P) \gg \log(\frac{1}{\epsilon})$, then Induction Hypothesis J.2 holds for all $T_{c,1} < t \leq T_{c,1}^\epsilon = T_{c,1} + 1$, we have $O \left(\frac{\log(P\epsilon^{-1})}{\eta\epsilon} \right)$, and at iteration $t = T_{c,1}^\epsilon + 1$, we have*

1. $\tilde{\mathcal{L}}_{k,\mathbf{p}}(Q^{T_{c,1}^\epsilon+1}) < \frac{\epsilon}{2K}$;
2. If $M \in \mathcal{E}_{k,1}$, we have $\left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,1}}^{(T_{c,1}^\epsilon+1)}\right)^2 \leq O(\epsilon)$.

K Proof of Main Theorems

K.1 Proof of Induction Hypotheses

We are now ready to show Induction Hypothesis [G.1](#) (resp. Induction Hypothesis [G.2](#)) holds through the learning process.

Theorem K.1 (Positive Information Gap). *For sufficiently large $P > 0$, $\eta \ll \log(P)$, $\Omega(1) \leq \Delta < 1$, Induction Hypothesis [G.1](#) holds for all iterations $t = 0, 1, \dots, T = O\left(\frac{e^{\text{polylog}(P)}}{\eta}\right)$.*

Theorem K.2 (Negative Information Gap). *For sufficiently large $P > 0$, $\eta \ll \log(P)$, $-0.5 < \Delta \leq -\Omega(1)$, Induction Hypothesis [G.2](#) holds for all iterations $t = 0, 1, \dots, T = O\left(\frac{e^{\text{polylog}(P)}}{\eta}\right)$.*

Proof of Theorem K.1. It is easy to verify Induction Hypothesis [G.1](#) holds at iteration $t = 0$ due to the initialization $Q^{(0)} = \mathbf{0}_{d \times d}$. At iteration $t > 0$:

- Induction Hypothesis [G.1a](#). can be proven by Induction Hypothesis [H.1-H.4 a](#) and Induction Hypothesis [J.1-J.2 a](#), combining with the fact that $\log(1/\epsilon) \ll \text{polylog}(P)$.
- Induction Hypothesis [G.1b](#). can be obtained by invoking Induction Hypothesis [H.1-H.4 b](#).
- Induction Hypothesis [G.1c](#). can be obtained by invoking Induction Hypothesis [H.1-H.4 c](#) and Induction Hypothesis [J.1-J.2 b](#).
- To prove Induction Hypothesis [G.1d](#)., for $\mathbf{q} \neq \mathbf{p}$, $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \sum_{k=1}^K \Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$. By item d-f in Induction Hypothesis [H.1-H.4](#) and item c-d in Induction Hypothesis [J.1-J.2](#), we can conclude that no matter the relative areas \mathbf{q} and \mathbf{p} belong to for a specific cluster, for all $k \in [K]$, throughout the entire learning process, the following upper bound always holds:

$$\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{q}}^{(t)} \leq \max_{t \in [T]} (|\Phi_{\mathbf{p} \rightarrow v_{k,n}}^{(t)}| + |\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}|) \max \left\{ O\left(\frac{1}{C_1}\right), O\left(\frac{1}{C_n}\right), O\left(\frac{1}{P}\right) \right\} \leq \tilde{O}\left(\frac{1}{C_n}\right).$$

Moreover, since $K = \Theta(1)$, we then have $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)} = \tilde{O}\left(\frac{1}{C_n}\right)$, which completes the proof.

- The proof for Induction Hypothesis [G.1d](#). is similar as before, by noticing that $\Upsilon_{k,\mathbf{p} \rightarrow \mathbf{p}}^{(t)} = \tilde{O}\left(\frac{1}{P}\right)$ for each $k \in [K]$, which is due to Induction Hypothesis [H.1-H.4 d](#) and Induction Hypothesis [J.1-J.2 c](#).

The proof of Theorem [K.2](#) mirrors that of Theorem [K.1](#), with the only difference being the substitution of relevant sections with Induction Hypothesis [G.2](#). For the sake of brevity, this part of the proof is not reiterated here.

K.2 Proof of Theorem [3.1](#) and Theorem [3.2](#) with Positive Information Gap

Theorem K.3. *Suppose $\Omega(1) \leq \Delta \leq 1$. For any $0 < \epsilon < 1$, suppose $\text{polylog}(P) \gg \log\left(\frac{1}{\epsilon}\right)$. We apply GD to train the loss function given in [\(2\)](#) with $\eta \ll \text{poly}(P)$. Then for each $\mathbf{p} \in \mathcal{P}$, we have*

1. *The loss converges: after $T^* = O\left(\frac{\log(P)P^{\max\{2(\frac{L}{L}-1), 1\}}(1-\kappa_s)}{\eta} + \frac{\log(P\epsilon^{-1})}{\eta\epsilon}\right)$ iterations, $\mathcal{L}_{\mathbf{p}}(Q^{(T^*)}) - \mathcal{L}_{\mathbf{p}}^* \leq \epsilon$, where $\mathcal{L}_{\mathbf{p}}^*$ is the global minimum of patch-level construction loss in [\(7\)](#).*
2. *Attention score concentrates: given cluster $k \in [K]$, if $X_{\mathbf{p}}$ is masked, then the one-layer transformer nearly ‘‘pays all attention’’ to all unmasked patches in the same area $\mathcal{P}_{k,a_k,\mathbf{p}}$, i.e., $\left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k,a_k,\mathbf{p}}}^{(T^*)}\right)^2 \leq O(\epsilon)$.*
3. **Local area learning feature attention correlation through two-phase:** given $k \in [K]$, if $a_{k,\mathbf{p}} > 1$, then we have
 - (a) $\Phi_{\mathbf{p} \rightarrow v_{k,1}}^{(t)}$ first quickly decrease with all other $\Phi_{\mathbf{p} \rightarrow v_{k,m}}^{(t)}$, $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ not changing much;

(b) after some point, the increase of $\Phi_{\mathbf{p} \rightarrow v_k, a_k, \mathbf{p}}^{(t)}$ takes dominance. Such $\Phi_{\mathbf{p} \rightarrow v_k, a_k, \mathbf{p}}^{(t)}$ will keep growing until convergence with all other feature and positional attention correlations nearly unchanged.

4. **Global area learning feature attention correlation through one-phase:** given $k \in [K]$, if $a_{k, \mathbf{p}} = 1$, throughout the training, the increase of $\Phi_{\mathbf{p} \rightarrow v_k, 1}^{(t)}$ dominates, whereas all $A_{1, m}^{(t)}$ with $m \neq 1$ and position attention correlations remain close to 0.

Proof. The first statement is obtained by letting $T^* = \max\{T_2^\epsilon, T_{c,1}^\epsilon\} + 1$ in Lemma H.30 and Lemma J.3, combining with Lemma F.9 and Lemma F.10, which lead to

$$\begin{aligned} \mathcal{L}_{\mathbf{p}}(Q^{(T^*)}) - \mathcal{L}_{\mathbf{p}}^* &\leq \mathcal{L}_{\mathbf{p}}(Q^{(T^*)}) - \mathcal{L}_{\mathbf{p}}^{\text{low}} \\ &\leq \tilde{\mathcal{L}}_{\mathbf{p}}(Q^{T^*}) + O\left(\exp\left(-\left(c_3 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k, \mathbf{p}}\}\} c_4 P^{\kappa_s}\right)\right)\right) \\ &\leq K \cdot \frac{\epsilon}{2K} + O\left(\exp\left(-\left(c_3 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k, \mathbf{p}}\}\} c_4 P^{\kappa_s}\right)\right)\right) \\ &< \epsilon. \end{aligned}$$

The second statement follows from Lemma H.30 and Lemma J.3. The third and fourth statements directly follow from the learning process described in Appendix H and Appendix J when Induction Hypothesis G.1 holds. \square

K.3 Proof of Theorem 3.1 and Theorem 3.2 with Negative Information Gap

Theorem K.4. Suppose $-0.5 \leq \Delta \leq \Omega(1)$. For any $0 < \epsilon < 1$, suppose $\text{polylog}(P) \gg \log(\frac{1}{\epsilon})$. We apply GD to train the loss function given in (2) with $\eta \ll \text{poly}(P)$. Then for each $\mathbf{p} \in \mathcal{P}$, we have

1. The loss converges: after $T^* = O\left(\frac{\log(P) P^{\max\{2(\frac{L}{L}-1), 1\}} (1-\kappa_s)}}{\eta} + \frac{\log(P\epsilon^{-1})}{\eta\epsilon}\right)$ iterations, $\mathcal{L}_{\mathbf{p}}(Q^{(T^*)}) - \mathcal{L}_{\mathbf{p}}^* \leq \epsilon$, where $\mathcal{L}_{\mathbf{p}}^*$ is the global minimum of patch-level construction loss in (7).
2. Attention score concentrates: given cluster $k \in [K]$, if $X_{\mathbf{p}}$ is masked, then the one-layer transformer nearly ‘‘pays all attention’’ to all unmasked patches in the same area $\mathcal{P}_{k, a_k, \mathbf{p}}$, i.e., $\left(1 - \text{Attn}_{\mathbf{p} \rightarrow \mathcal{P}_{k, a_k, \mathbf{p}}}^{(T^*)}\right)^2 \leq O(\epsilon)$.
3. All areas learning feature attention correlation through one-phase: given $k \in [K]$, throughout the training, the increase of $\Phi_{\mathbf{p} \rightarrow v_k, a_k, \mathbf{p}}^{(t)}$ dominates, whereas all $\Phi_{\mathbf{p} \rightarrow v_k, m}^{(t)}$ with $m \neq 1$ and position attention correlations $\Upsilon_{\mathbf{p} \rightarrow \mathbf{q}}^{(t)}$ remain close to 0.

Proof. The first statement is obtained by letting $T^* = \max\{T_{\text{neg},1}^\epsilon, T_{c,1}^\epsilon\} + 1$ in Lemma I.3 and Lemma J.3, combining with Lemma F.9 and Lemma F.10, which lead to

$$\begin{aligned} \mathcal{L}_{\mathbf{p}}(Q^{(T^*)}) - \mathcal{L}_{\mathbf{p}}^* &\leq \mathcal{L}_{\mathbf{p}}(Q^{(T^*)}) - \mathcal{L}_{\mathbf{p}}^{\text{low}} \\ &\leq \tilde{\mathcal{L}}_{\mathbf{p}}(Q^{T^*}) + O\left(\exp\left(-\left(c_3 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k, \mathbf{p}}\}\} c_4 P^{\kappa_s}\right)\right)\right) \\ &\leq K \cdot \frac{\epsilon}{2K} + O\left(\exp\left(-\left(c_3 P^{\kappa_c} + \mathbb{1}\{1 \notin \cup_{k \in [K]} \{a_{k, \mathbf{p}}\}\} c_4 P^{\kappa_s}\right)\right)\right) \\ &< \epsilon. \end{aligned}$$

The second statement follows from Lemma I.3 and Lemma J.3. The third and fourth statements directly follow from the learning process described in Appendix I and Appendix J when Induction Hypothesis G.2 holds. \square