
Vision Transformers provably learn spatial structure

Samy Jelassi
Princeton University
sjelassi@princeton.edu

Michael E. Sander
ENS, CNRS
michael.sander@ens.fr

Yuanzhi Li
Carnegie Mellon University
yuanzhil@andrew.cmu.edu

Abstract

Vision Transformers (ViTs) have achieved comparable or superior performance than Convolutional Neural Networks (CNNs) in computer vision. This empirical breakthrough is even more remarkable since, in contrast to CNNs, ViTs do not embed any visual inductive bias of spatial locality. Yet, recent works have shown that while minimizing their training loss, ViTs specifically learn spatially localized patterns. This raises a central question: how do ViTs learn these patterns by solely minimizing their training loss using gradient-based methods from *random initialization*? In this paper, we provide some theoretical justification of this phenomenon. We propose a spatially structured dataset and a simplified ViT model. In this model, the attention matrix solely depends on the positional encodings. We call this mechanism the positional attention mechanism. On the theoretical side, we consider a binary classification task and show that while the learning problem admits multiple solutions that generalize, our model implicitly learns the spatial structure of the dataset while generalizing: we call this phenomenon patch association. We prove that patch association helps to sample-efficiently transfer to downstream datasets that share the same structure as the pre-training one but differ in the features. Lastly, we empirically verify that a ViT with positional attention performs similarly to the original one on CIFAR-10/100, SVHN and ImageNet.

1 Introduction

Transformers are deep learning models built on self-attention [65], and in the past several years they have increasingly formed the backbone for state-of-the-art models in domains ranging from Natural Language Processing (NLP) [65, 23] to computer vision [24], reinforcement learning [13, 38], program synthesis [5] and symbolic tasks [44]. Beyond their remarkable performance, several works reported the ability of transformers to simultaneously minimize their training loss and learn inductive biases tailored to specific datasets e.g. in computer vision [55], in NLP [10, 67] or in mathematical reasoning [73]. In this paper, we focus on computer vision where convolutions are considered to be an adequate and biologically plausible inductive bias since they capture local spatial information [27] by imposing a sparse local connectivity pattern. This seems intuitively reasonable: nearby pixels encode the presence of small scale features, whose patterns in turn determine more abstract features at longer and longer length scales. Several seminal works [17, 24, 55] *empirically* show that although randomly initialized, the positional encodings in Vision transformers (ViTs) [24] actually learn this local connectivity: closer patches have more similar positional encodings, as shown in Figure 1a. A priori, learning such spatial structure is surprising. Indeed, in contrast to convolutional neural networks (CNNs), ViTs are not built with the inductive bias of local connectivity and weight sharing. They start by replacing an image by a collection of D patches $(\mathbf{X}_1, \dots, \mathbf{X}_D) \in \mathbb{R}^{d \times D}$, each of dimension d . While each \mathbf{X}_i represents (an embedding of) a spatially localized portion of the original image, the relative positions of the patches \mathbf{X}_i in the image are disregarded. Instead, relative spatial information is supplied through *image-independent* positional encodings $\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_D) \in \mathbb{R}^{d \times D}$. Unlike CNNs, each layer of a ViT then learns, via trainable self-attention, a non-local set of filters that non-linearly depend on both the values of all patches \mathbf{X}_j and their positional encodings \mathbf{p}_j .

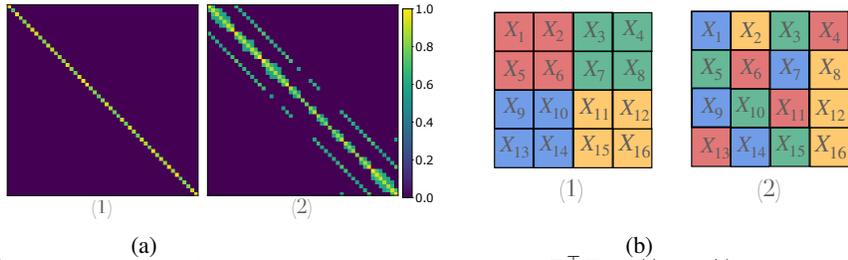


Figure 1: (a) Visualization of the positional encodings similarities $\mathbf{P}^\top \mathbf{P} = (\langle \mathbf{p}_i, \mathbf{p}_j \rangle)_{(i,j) \in [D]^2}$ at initialization (1) and after training on Imagenet (2) using a "ViT-small-patch32-224" [24]. We normalise the values $\mathbf{P}^\top \mathbf{P}$ between -1 and 1 and apply a threshold of 0.55 . In contrast with the initial arrays that are random, the final ones show local connectivity patterns: nearby patches have similar positional encodings. (b) Partition of the patches into sets \mathcal{S}_ℓ as in Definition 2.1. Squares in the same color belong to the same set \mathcal{S}_ℓ . We refer to (1) as a "spatially localized set" since all the elements in a \mathcal{S}_ℓ are spatially contiguous. This is the type of sets appearing in Figure 1a at the end of training. Definition 2.1 also covers sets with non-contiguous elements as (2).

Contributions. The empirical observation of Figure 1a sets a central question: from a *theoretical perspective*, how do ViTs manage to learn these local connectivity patterns by simply minimizing their training loss using gradient descent from random initialization? While it is known that attention can express local operations as convolution [17], it remains unclear how ViTs learn it. In this paper, we present a simple spatially-structured classification dataset for which it is sufficient (but not necessary) to learn the structure in order to generalize. We also present a simplified ViT model which we prove implicitly learns sparse spatial connectivity patterns when it minimizes its training loss via gradient descent (GD). We name this implicit bias *patch association* (defined in Definition 2.2). We prove that our ViT model leverages this bias to generalize. More precisely, we make the following contributions:

- In Section 2, we formally define the concept of performing patch association, which refer to the ability of learning spatial connectivity patterns on a dataset.
- In Section 3, we introduce a structured classification dataset and a simplified ViT model. This model is simplified in the sense that its attention matrix only depends on the positional encodings. We then present the learning problems we are interested in: empirical risk (realistic setting) and population risk (idealized setting) minimization for binary classification.
- In Section 4, we prove that a one-layer single-head ViT model trained with gradient descent on our synthetic dataset performs patch association and generalizes, in the idealized (Theorem 4.1) and realistic (Theorem 4.2) settings. We present a detailed proof, based on invariance and symmetries of coefficients in the attention matrix throughout the learning process.
- In Section 5, we show (Theorem 5.1) that after pre-training in our synthetic dataset, our model can be sample-efficiently fine-tuned to transfer to a downstream dataset that shares the same structure as the source dataset (and may have different features).
- On the experimental side, we validate in Section 6 that ViTs learn spatial structure in images from the CIFAR-100 dataset, even when the pixels of the images are permuted. This result validates that, in contrast to CNNs, ViTs learn a more general form of spatial structure that is not limited to local patterns (Figure 5). We finally show that our ViT model –where the attention matrix only depends on the positional encodings– is competitive with the vanilla ViT on the ImageNet, CIFAR-10/100 and SVHNs datasets (Figure 6 and Figure 7).

Notation. We use lower case letters for scalars, lower case bold for vectors and upper case bold for matrices. Given an integer D , we define $[D] = \{1, \dots, D\}$. Any statement made "with high probability" holds with probability at least $1 - 1/\text{poly}(d)$. Given a vector $\mathbf{a} \in \mathbb{R}^d$ and $k \leq d$, we define $\text{Top}_k\{a_j\}_{j=1}^d = \{a_{i_1}, \dots, a_{i_k}\}$ where a_{i_1}, \dots, a_{i_k} are the k -largest elements. For a function F that implicitly depend on parameters \mathbf{A} and \mathbf{v} , we often write $F_{\mathbf{A}, \mathbf{v}}$ to highlight its parameters. We use the asymptotic complexity notations when defining the different constants.

Related work

CNNs and ViTs. Many computer vision architectures can be considered as a form of hybridization between Transformers and CNNs. For example, DeTR [11] use a CNN to generate features that are fed to a Transformer. [25] show that self-attention can be initialized or regularized to behave like a convolution and [19, 30] add convolution operations to Transformers. Conversely, [8, 56, 7] introduce self-attention or attention-like operations to supplement or replace convolution in ResNet-

like models. In contrast, our paper does not consider any form of hybridization with CNN, but rather a simplification of the original ViT to explain how ViTs learn spatially structured patterns using GD.

Empirical understanding of ViTs. A long line of work consists in analyzing the properties of ViTs, such as robustness [9, 54, 51] or the effect of self-supervision [12, 14]. Closer to our work, some papers investigate why ViTs perform so well. [55] compare the representations of ViTs and CNNs and [50, 64] argue that the patch embeddings could explain the performance of ViTs. We empirically show in Section 6 that applying the attention matrices to the positional encodings – which contains the structure of the dataset – approximately recovers the baselines. Hence, our work rather suggests that the structural learning performed by the attention matrices may explain the success of ViTs.

Theory for attention models. Early theoretical works have focused on the expressivity of attention. [66, 26] addressed this question in the context of self-attention blocks and [21, 68, 34] for Transformers. On the optimization side, [76] investigate the role of adaptive methods in attention models and [59] analyze the dynamics of a single-head attention head to approximate the learning of a Seq2Seq architecture. In our work, we also consider a single-head ViT trained with gradient descent and exhibit a setting where it provably learns convolution-like patterns and generalizes.

Algorithmic regularization. The question we address concerns algorithmic regularization which characterizes the generalization of an optimization algorithm when multiple global solutions exist in over-parametrized models. This regularization arises in deep learning mainly due to the *non-convexity* of the objective function. Indeed, this latter potentially creates multiple global minima scattered in the space that vastly differ in terms of generalization. Algorithmic regularization appears in binary classification [60, 49, 16], matrix factorization [29, 3], convolutional neural networks [29, 37], generative adversarial networks [2], contrastive learning [70] and mixture of experts [15]. Algorithmic regularization is induced by and depends on many factors such as learning rate and batch size [28, 33, 41, 58, 47], initialization [1], momentum [39], adaptive step-size [42, 53, 20, 71, 78, 40], batch normalization [4, 32, 36] and dropout [61, 69]. However, all these works consider the case of feed-forward neural networks which does not apply to ViTs.

2 Defining patch association

The goal of this section is to formalize the way ViTs learn sparse spatial connectivity patterns. We thus introduce the concept of performing patch association for a spatially structured dataset.

Definition 2.1 (Data distribution with spatial structure). *Let \mathcal{D} be a distribution over $\mathbb{R}^{d \times D} \times \{-1, 1\}$ where each patch $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_D) \in \mathbb{R}^{d \times D}$ has label $y \in \{-1, 1\}$. We say that \mathcal{D} is spatially structured if*

- there exists a partition of $[D]$ into L disjoint subsets i.e. $[D] = \bigcup_{\ell=1}^L \mathcal{S}_\ell$ with $\mathcal{S}_\ell \subsetneq D$ and $|\mathcal{S}_\ell| = C$.
- there exists a labelling function f^* satisfying $\mathbb{P}[y f^*(\mathbf{X}) > 0] = 1 - d^{-\omega(1)}$ and,

$$f^*(\mathbf{X}) := \sum_{\ell \in [L]} \phi((\mathbf{X}_i)_{i \in \mathcal{S}_\ell}), \quad \text{where } \phi: \mathbb{R}^{d \times C} \rightarrow \mathbb{R} \text{ is an arbitrary function.} \quad (1)$$

Examples. A particular case for the sets \mathcal{S}_ℓ 's is the one of spatially localized sets as in Figure 1b-(1). In this case, we have $D = 16$, $C = 4$ and $\mathcal{S}_1 = \{1, 2, 5, 6\}$, $\mathcal{S}_2 = \{3, 4, 7, 8\}$, $\mathcal{S}_3 = \{9, 10, 13, 14\}$, $\mathcal{S}_4 = \{11, 12, 15, 16\}$. We emphasize that Definition 2.1 is not limited to spatially localized sets and also covers non-contiguous sets as Figure 1b-(2).

Labelling function Definition 2.1 states that there exists a labelling function that preserves the underlying structure by applying the same function ϕ to each \mathcal{S}_ℓ as in (1). For instance, when the sets \mathcal{S}_ℓ 's

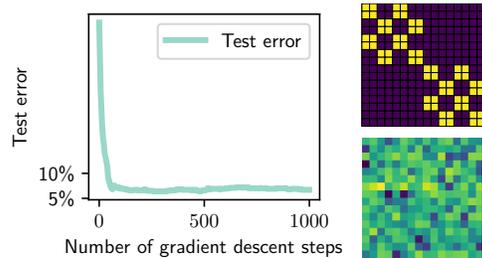


Figure 2: Left: Test error of the ViT on the convolution structured dataset. Upper Right: Grid displaying the input patches. Yellow squares represent spatially localized sets \mathcal{S}_ℓ . Those sets are taken into account when computing the convolutional function f^* . Lower Right: Learnt $\mathbf{P}^T \mathbf{P}$ looks random compared to upper one.

are spatially localized, f^* can be a one-hidden layer convolutional network. In this paper, we are interested in *patch association* which refers to the ability of an algorithm to identify the sets \mathcal{S}_ℓ 's, and is formally defined as follow.

Definition 2.2 (Patch association for ViTs). *Let \mathcal{D} be as in Definition 2.1. Let $\mathcal{M}: \mathbb{R}^{d \times D} \rightarrow \{-1, 1\}$ be a transformer and $\mathbf{P}^{(\mathcal{M})}$ its positional encodings matrix. We say that \mathcal{M} performs patch association on \mathcal{D} if for all $\ell \in [L]$ and $i \in \mathcal{S}_\ell$, we have $\text{Top}_C \{\langle \mathbf{p}_i^{(\mathcal{M})}, \mathbf{p}_j^{(\mathcal{M})} \rangle\}_{j=1}^D = \mathcal{S}_\ell$.*

Definition 2.2 states that patch association is learned when for a given $i \in \mathcal{S}_\ell$, its positional encoding mainly attends those of j such that $i, j \in \mathcal{S}_\ell$. In this way, the transformer groups the \mathbf{X}_i according to \mathcal{S}_ℓ just like the true labeling function. Definition 2.2 formally describes the empirical findings in Figure 1a-(2), where nearby patches have similar positional encodings. A natural question is then: would ViTs really learn those \mathcal{S}_ℓ after training to match the labeling function f^* ? Without further assumptions on the data distribution, we next show that the answer is no.

ViTs do not always learn patch association under Assumption 1. We give a negative answer through the following synthetic experiment. Consider the case where all the patches \mathbf{X}_j are i.i.d. standard Gaussian and f^* is a one-hidden layer CNN with cubic activation. The label y of any \mathbf{X} is then given by $y = \text{sign}(f^*(\mathbf{X}))$. As shown in Figure 2, one-layer ViT reaches small test error on the binary classification task. However, $\mathbf{P}^\top \mathbf{P}$ does not match the convolution pattern encoded in f^* . This is not surprising, since the data distribution \mathcal{D} is Gaussian, and thus lacks spatial structure. Thus, in order to prove that ViTs learn patch association, we need additional assumptions on \mathcal{D} , which we discuss in the next section.

3 Setting to learn patch association

In this section, we introduce our theoretical setting to analyze how ViTs learn patch association. We first define our binary classification dataset and finally present the ViT model we use to classify it.

Assumption 1 (Data distribution with specific spatial structure). *Let \mathcal{D} be a distribution as in Definition 2.1 and $\mathbf{w}^* \in \mathbb{R}^d$ be an underlying feature. We suppose that each data-point \mathbf{X} is defined as follow*

- Uniformly sample an index $\ell(\mathbf{X})$ from $[L]$ and for $j \in \mathcal{S}_{\ell(\mathbf{X})}$, $\mathbf{X}_j = y\mathbf{w}^* + \boldsymbol{\xi}_j$, where $y\mathbf{w}^*$ is the informative feature and $\boldsymbol{\xi}_j \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2(\mathbf{I}_D - \mathbf{w}^*\mathbf{w}^{*\top}))$ (signal set).
- For $\ell \in [L] \setminus \{\ell(\mathbf{X})\}$ and $j \in \mathcal{S}_\ell$, $\mathbf{X}_j = \delta_j \mathbf{w}^* + \boldsymbol{\xi}_j$, where $\delta_j = 1$ with probability $q/2$, -1 with same probability and 0 otherwise, and $\boldsymbol{\xi}_j \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2(\mathbf{I}_D - \mathbf{w}^*\mathbf{w}^{*\top}))$ (random sets).

To keep the analysis simple, the noisy patches are sampled from the orthogonal complement of \mathbf{w}^* . Note that \mathcal{D} admits the labeling function $f^*(\mathbf{X}) = \sum_{\ell \in [L]} \text{Threshold}_{0.9C}(\sum_{i \in \mathcal{S}_\ell} \langle \mathbf{w}^*, \mathbf{X}_i \rangle)$, where $\text{Threshold}_C(z) = z$ if $|z| > C$ and 0 otherwise.

We sketch a data-point of \mathcal{D} in Figure 3. Our dataset can be viewed as an extreme simplification of real-world image datasets where there is a set of adjacent patches that contain a useful feature (e.g. the nose of a dog) and many patches that have uninformative or spurious features e.g. the background of the image. We make the following assumption on the parameters of the data distribution.

Assumption 2. *We suppose that $d = \text{poly}(D)$, $C = \text{polylog}(d)$, $q = \text{poly}(C)/D$, $\|\mathbf{w}^*\|_2 = 1$ and $\sigma^2 = 1/d$. This implies $C \ll D$ and $q \ll 1$.*

Assumption 2 may be justified by considering a "ViT-base-patch16-224" model [24] on ImageNet. In this case, $d = 384$, $D = 196$. σ is set to have $\|\boldsymbol{\xi}_j\|_2 \approx \|\mathbf{w}^*\|_2$. q is chosen so that there are more spurious features than informative ones

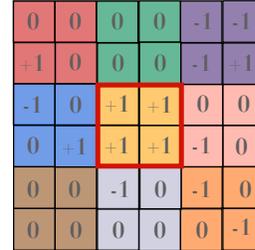


Figure 3: Visualization of a data-point \mathbf{X} in \mathcal{D} when the \mathcal{S}_ℓ 's are spatially localized. Each square depicts a patch \mathbf{X}_j and squares of the same color belong to the same set \mathcal{S}_ℓ . "0" indicates that the patch does not have a feature, "1" stands for feature $1 \cdot \mathbf{w}^*$ and "-1" for feature $-1 \cdot \mathbf{w}^*$. The large red square depicts the signal set $\ell(\mathbf{X})$. Although there are more "-1"s than "+1"s, the label of \mathbf{X} is +1 since there are only "+1"s inside the signal set.

(low signal-to-noise regime) which makes the data *non-linearly separable*. Our dataset is non-trivial to learn since generalized linear networks fail to generalize, as shown in the next theorem (see Appendix J for a proof).

Theorem 3.1. *Let \mathcal{D} be as in Assumption 1. Let $g(\mathbf{X}) = \phi\left(\sum_{j=1}^D \langle \mathbf{w}_j, \mathbf{X}_j \rangle\right)$ be a generalized linear model. Then, g does not fit the labeling function i.e. $\mathbb{P}[f^*(\mathbf{X})g(\mathbf{X}) \leq 0] \geq 1/8$.*

Intuitively, g fails to generalize because it does not have any knowledge on the underlying partition and the number of random sets is much higher than those with signal. Thus, a model must have a minimal knowledge about the \mathcal{S}_ℓ 's in order to generalize. In addition, the following Theorem 3.2 states the existence of a transformer that generalizes without learning spatial structure (see Appendix J for a proof), thus showing that the learning process has a priori no straightforward reason to lead to patch association.

Theorem 3.2. *Let \mathcal{D} be defined as in Assumption 1. There exists a (one-layer) transformer \mathcal{M} so that $\mathbb{P}[f^*(\mathbf{X})\mathcal{M}(\mathbf{X}) \leq 0] = d^{-\omega(1)}$ but for all $\ell \in [L]$, $i \in \mathcal{S}_\ell$, $\text{Top}_C \{\langle \mathbf{p}_i^{(\mathcal{M})}, \mathbf{p}_j^{(\mathcal{M})} \rangle\}_{j=1}^D \cap \mathcal{S}_\ell = \emptyset$.*

Simplified ViT model. We now define our simplified ViT model for which we show in Section 4 that it implicitly learns patch association via minimizing its training objective. We first remind the self-attention mechanism that is ubiquitously used in transformers.

Definition 3.1 (Self-attention [6, 65]). *The attention mechanism [6, 65] in the single-head case is defined as follow. Let $\mathbf{X} \in \mathbb{R}^{d \times D}$ a data point and $\mathbf{P} \in \mathbb{R}^{d \times D}$ its positional encoding. The self-attention mechanism computes*

1. the sum of patches and positional encodings i.e. $\mathbf{X} = \mathbf{X} + \mathbf{P}$.
2. the attention matrix $\mathbf{A} = \mathbf{Q}\mathbf{K}^\top$ where $\mathbf{Q} = \mathbf{X}^\top \mathbf{W}_Q$, $\mathbf{K} = \mathbf{X}^\top \mathbf{W}_K$, $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{d \times d}$.
3. the score matrix $\mathbf{S} \in \mathbb{R}^{D \times D}$ with coefficients $S_{i,j} = \exp(A_{i,j}/\sqrt{d}) / \sum_{r=1}^D \exp(A_{i,r}/\sqrt{d})$.
4. the matrix $\mathbf{V} = \mathbf{X}^\top \mathbf{W}_V$, where $\mathbf{W}_V \in \mathbb{R}^{d \times d}$.

It finally outputs $\text{SA}((\mathbf{X}; \mathbf{P})) = \mathbf{S}\mathbf{V} \in \mathbb{R}^{d \times D}$.

In this paper, our ViT model relies on a different attention mechanism –the "positional attention"– that we define as follows.

Definition 3.2 (Positional attention). *Let $\mathbf{X} \in \mathbb{R}^{d \times D}$ and $\mathbf{P} \in \mathbb{R}^{d \times D}$ the positional encoding. The positional attention mechanism takes as input the pair $(\mathbf{X}; \mathbf{P})$ and computes:*

1. the attention matrix $\mathbf{A} = \mathbf{Q}\mathbf{K}^\top$ where $\mathbf{Q} = \mathbf{P}^\top \mathbf{W}_Q$, $\mathbf{K} = \mathbf{P}^\top \mathbf{W}_K$ and $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{d \times d}$.
2. the score matrix $\mathbf{S} \in \mathbb{R}^{D \times D}$ with coefficients $S_{i,j} = \exp(A_{i,j}/\sqrt{d}) / \sum_{r=1}^D \exp(A_{i,r}/\sqrt{d})$.
3. the matrix $\mathbf{V} = \mathbf{X}^\top \mathbf{W}_V$, where $\mathbf{W}_V \in \mathbb{R}^{d \times d}$.

It outputs $\text{PA}((\mathbf{X}; \mathbf{P})) = \mathbf{S}\mathbf{V}$.

Positional attention *isolates* positional encoding \mathbf{P} from data \mathbf{X} : \mathbf{A} encodes the dynamics of \mathbf{P} and tracks whether patch association is learned. \mathbf{V} encodes the data-dependent part and monitors whether the feature is learned. Indeed, given its highly non-linear nature with respect to the input, directly analyzing self-attention is difficult. Yet, positional attention is similar to self-attention. As this latter, positional attention is also permutation-invariant and processes all tokens simultaneously. Besides, positional attention also computes a score matrix between the different tokens. This similarity matrix is also normalized in a sparse manner with the Softmax operator. The only aspect that positional attention misses from self-attention is the fact that \mathbf{S} does not depend on the input. Nevertheless, we empirically show that our positional attention model competes with self-attention in Section 6. Lastly, we make the following simplification in the parameters to ease our analysis.

Simplification 3.1. *In the positional attention mechanism, we set $d = D$, $\mathbf{W}_K = \mathbf{I}_D$ and $\mathbf{W}_Q = \mathbf{I}_D$ which implies $\mathbf{A} = \mathbf{P}^\top \mathbf{P}$. We set $\mathbf{W}_V = [\mathbf{v}, \dots, \mathbf{v}] \in \mathbb{R}^{d \times D}$ where $\mathbf{v} \in \mathbb{R}^d$. Finally, we set \mathbf{A} and \mathbf{v} as trainable parameters. Besides, without loss of generality, we train all $A_{i,j}$ for $i \neq j$ and leave the diagonals of \mathbf{A} fixed.*

In Simplification 3.1, we set \mathbf{W}_K and \mathbf{W}_Q to the identity so that $\mathbf{A} = \mathbf{P}^\top \mathbf{P}$. This Gram matrix encodes the spatial patterns learned by the ViT as shown in Figure 1a. Besides, since fitting the

labeling function requires to learn one feature \mathbf{w}^* , it is sufficient to parameterize \mathbf{W}_V with a vector \mathbf{v} . Also, although $\mathbf{A} = \mathbf{P}^\top \mathbf{P}$ and \mathbf{P} is trainable, we choose for simplicity to only optimize over \mathbf{A} . Besides, we leave the $A_{i,i}$'s fixed because Softmax is invariant under the uniform shift of the input. Under [Simplification 3.1](#), our simplified ViT model is then a two attention layer with a single head:

$$F(\mathbf{X}) = \sum_{i=1}^D \sigma \left(D \sum_{j=1}^D S_{i,j} \langle \mathbf{v}, \mathbf{X}_j \rangle \right) \quad \text{with} \quad S_{i,j} = \exp(A_{i,j}/\sqrt{d}) / \sum_{r=1}^D \exp(A_{i,r}/\sqrt{d}), \quad (\text{T})$$

where σ is an activation function. Since we aim to the simplest ViT model, we opt for a polynomial activation i.e. $\sigma(x) = x^p + \nu x$ where $p \geq 3$ is an odd integer and $\nu = 1/\text{poly}(d)$. Note that this choice of polynomial activation is common in the deep learning theory literature – see e.g. [45, 1, 72] among others. The degree p is odd to make the ViT model compatible with the labeling function and strictly larger than 1 because the data is not linearly separable ([Theorem 3.1](#)). We add a linear part in the activation function to ensure that the gradient is non-zero when \mathbf{v} has small coefficients. With these simplifications, we formally prove that F is able to learn patch association and generalize, in the two following settings.

Idealized and realistic learning problems. Given a dataset $\mathcal{Z} = \{(\mathbf{X}[i], y[i])\}_{i=1}^N$ sampled from \mathcal{D} , we solve the empirical risk minimization problem for the logistic loss defined by:

$$\min_{\mathbf{A}, \hat{\mathbf{v}}} \frac{1}{N} \sum_{i=1}^N \log(1 + e^{-y[i]F(\mathbf{X}[i])}) := \hat{\mathcal{L}}(\hat{\mathbf{A}}, \hat{\mathbf{v}}). \quad (\text{E})$$

Instead of directly analyzing [\(E\)](#), we introduce a proxy where we minimize the population risk

$$\min_{\mathbf{A}, \mathbf{v}} \mathbb{E}_{\mathcal{D}} [\log(1 + e^{-yF(\mathbf{X})})] := \mathcal{L}(\mathbf{A}, \mathbf{v}). \quad (\text{P})$$

We refer to [\(E\)](#) as the *realistic* problem while [\(P\)](#) as the *idealized* problem.

Algorithm. We solve [\(P\)](#) and [\(E\)](#) using gradient descent (GD) for T iterations. The update rule in the case of [\(P\)](#) for $t \in [T]$ and $i, j \in [D]$ is

$$A_{i,j}^{(t+1)} = A_{i,j}^{(t)} - \eta \partial_{A_{i,j}} \mathcal{L}(\mathbf{A}^{(t)}, \mathbf{v}^{(t)}), \quad \mathbf{v}^{(t+1)} = \mathbf{v}^{(t)} - \eta \nabla_{\mathbf{v}} \mathcal{L}(\mathbf{A}^{(t)}, \mathbf{v}^{(t)}), \quad (\text{GD})$$

where $\eta > 0$ is the learning rate. A similar update may be written for [\(E\)](#). We now detail how to set the parameters in [\(GD\)](#).

Parametrization 3.1. When running GD on [\(P\)](#) and [\(E\)](#), the number of iterations is any $T \geq \text{poly}(d)/\eta$. We set the learning rate as $\eta \in (0, \frac{1}{\text{poly}(d)})$. The diagonal coefficient of the attention matrix are set for $i \in [D]$ as $A_{i,i}^{(0)} = \hat{A}_{i,i}^{(0)} = \sigma_{\mathbf{A}} \mathbf{I}_D$ where $\sigma_{\mathbf{A}} = \text{polyloglog}(d)$. The off-diagonal coefficients of \mathbf{A} and the value vector are initialized as:

1. *Idealized case:* $\mathbf{v}^{(0)} = \alpha^{(0)} \mathbf{w}^*$ where $\alpha^{(0)} = \nu^{1/(p-1)}$ and $A_{i,j}^{(0)} = 0$ for $i \neq j$.
2. *Realistic case:* $\hat{\mathbf{v}}^{(0)} \sim \mathcal{N}(0, \omega^2 \mathbf{I}_d)$ and $\hat{A}_{i,j}^{(0)} \sim \mathcal{N}(0, \omega^2)$ where $i \neq j$ and $\omega = 1/\text{poly}(d)$.

We remind that in [Simplification 3.1](#), we have $\mathbf{A} = \mathbf{P}^\top \mathbf{P}$. If one initializes $\mathbf{P} \sim \mathcal{N}(0, \sigma_{\mathbf{A}} \mathbf{I}_D/D)$, then with high probability, $A_{i,i}^{(0)} = \|\mathbf{p}_i^{(0)}\|_2^2 = \Theta(\sigma_{\mathbf{A}})$ and $A_{i,j}^{(0)} = \langle \mathbf{p}_i^{(0)}, \mathbf{p}_j^{(0)} \rangle = \Theta(\sigma_{\mathbf{A}}/\sqrt{D})$ for $i \neq j$. Since $D \gg 1$, it is then reasonable to set $A_{i,j}^{(0)} = 0$. Note that, also in the idealized setting, we initialize $\mathbf{v}^{(0)}$ in $\text{span}(\mathbf{w}^*)$, even though this latter should be unknown to the algorithm. We remind that the idealized case is a proxy to ultimately characterize the realistic dynamics.

4 Learning spatial structure via matching the labeling function

As announced above, we show that our ViT [\(T\)](#) implicitly learns patch association and fits the labeling function by minimizing the training objective. We first study the dynamics in [\(P\)](#). Using the analysis in the idealized case, we then characterize the solution found in the realistic problem [\(E\)](#).

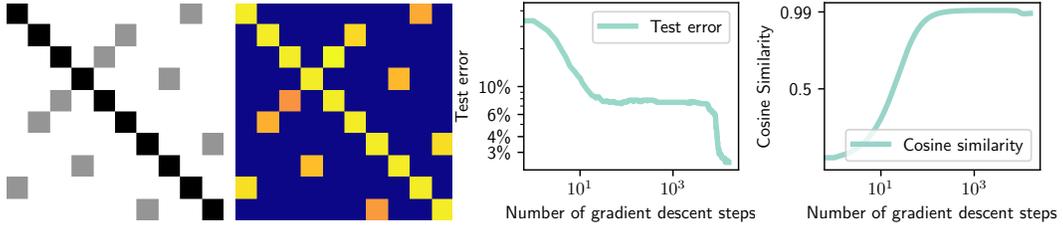


Figure 4: Illustration of [Theorem 4.2](#). We consider the exact same setting (data generation, parameter settings...) as for the realistic case. From left to right, we first display in grey the tuples (i, j) such that $(i, j) \in \mathcal{S}_\ell$. We then plot the learned matrix \mathbf{A} and see that coefficients with high value exactly correspond to their grey scale counterpart in the left plot. We also display test error and cosine similarity between \mathbf{w}^* and \mathbf{v} w.r.t the number of training steps.

4.1 Learning process in the idealized case

In this section, we analyze the dynamics of [\(P\)](#). Our main result is that after minimizing [\(P\)](#), our model [\(T\)](#) performs patch association while generalizing.

Theorem 4.1. *Assume that we run GD on [\(P\)](#) for T iterations with parameters set as in [Parametrization 3.1](#). With high probability, the ViT model [\(T\)](#)*

1. *learns patch association i.e. for all $\ell \in [L]$ and $i \in \mathcal{S}_\ell$, $\text{Top}_C \{A_{i,j}^{(T)}\}_{j=1}^D = \mathcal{S}_\ell$.*
2. *learns the labeling function f^* i.e. $\mathbb{P}_{\mathcal{D}}[f^*(\mathbf{X})F_{\mathbf{A}^{(T)}, \mathbf{v}^{(T)}}(\mathbf{X}) > 0] \geq 1 - o(1)$.*

We now sketch the main ideas to prove the theorem for which one can refer to [Appendix D](#) for a complete proof.

Invariance and symmetries. In [\(P\)](#), we take the expectation over \mathcal{D} . Since [\(T\)](#) is permutation-invariant and the data distribution is symmetric, we can thus dramatically simplify the variables in [\(P\)](#). An illustration of this is the next lemma that shows that \mathbf{A} can be reduced to three variables in [\(P\)](#).

Lemma 4.1. *There exist $\beta = \sigma_{\mathbf{A}}$, $\gamma^{(t)}, \rho^{(t)} \in \mathbb{R}$ such that for all $t \geq 0$:*

1. *for all $i \in [D]$, $A_{i,i}^{(t)} = \beta$.*
2. *for all $i, j \in [D]$ such that $i, j \in \mathcal{S}_\ell$ for some $\ell \in [L]$, $A_{i,j}^{(t)} = \gamma^{(t)}$.*
3. *for all $i, j \in [D]$ such that $i \in \mathcal{S}_\ell$ and $j \in \mathcal{S}_m$ for some $\ell, m \in [L]$ with $\ell \neq m$, $A_{i,j}^{(t)} = \rho^{(t)}$.*

Besides, using the initialization in [Parametrization 3.1](#), we can show that \mathbf{v} always lies in $\text{span}(\mathbf{w}^*)$.

Lemma 4.2. *For all $t \in [T]$, there exists $\alpha^{(t)} \in \mathbb{R}$ such that $\mathbf{v}^{(t)} = \alpha^{(t)} \mathbf{w}^*$.*

In summary, [Lemma 4.1](#) and [Lemma 4.2](#) imply that instead of optimizing over \mathbf{A} and \mathbf{v} in [\(P\)](#), we can instead consider the scalar variables $\alpha^{(t)}$, $\gamma^{(t)}$ and $\rho^{(t)}$. The remaining of this section consists in analyzing the dynamics of these three quantities.

Learning patch association. We first analyze the dynamics of $\gamma^{(t)}$ and $\rho^{(t)}$. To this end, we introduce the following terms:

$$\Lambda^{(t)} = \frac{e^\beta}{e^\beta + (C-1)e^{\gamma^{(t)}} + (D-C)e^{\rho^{(t)}}}, \quad \Gamma^{(t)} = \frac{e^{\gamma^{(t)}}}{e^\beta + (C-1)e^{\gamma^{(t)}} + (D-C)e^{\rho^{(t)}}},$$

$$\Xi^{(t)} = \frac{e^{\rho^{(t)}}}{e^\beta + (C-1)e^{\gamma^{(t)}} + (D-C)e^{\rho^{(t)}}}, \quad G^{(t)} = D(\Lambda^{(t)} + (C-1)\Gamma^{(t)}).$$

Note that $\Lambda^{(t)}$, $\Gamma^{(t)}$ and $\Xi^{(t)}$ respectively correspond to the coefficients on the diagonal, those for which $i, j \in \mathcal{S}_\ell$ for some $\ell \in [L]$ and all the other coefficients of the attention matrix \mathbf{S} . Using these notations, we first derive the GD updates of $\gamma^{(t)}$ and $\rho^{(t)}$.

Lemma 4.3. *Let $t \leq T$. The attention weights $\gamma^{(t)}$ and $\rho^{(t)}$ satisfy:*

$$\gamma^{(t+1)} = \gamma^{(t)} + \eta \text{polylog}(d)(\alpha^{(t)})^p \cdot \Gamma^{(t)}(G^{(t)})^{p-1},$$

$$|\rho^{(t+1)}| \leq |\rho^{(t)}| + \eta \text{polylog}(d)(\alpha^{(t)})^p \left(\frac{1}{D} + \frac{1}{D} \Gamma^{(t)}(G^{(t)})^{p-1} \right).$$

Lemma 4.3 shows that the increment of $\gamma^{(t)}$ is larger than the one of $\rho^{(t)}$. Since $\gamma^{(0)} = \rho^{(0)} = 0$, this implies that $\gamma^{(t)} \geq \rho^{(t)}$ for all $t \geq 0$. This observation proves the first item of **Theorem 4.1**. We now explain how learning patch association leads to \mathbf{v} highly correlated with \mathbf{w}^* .

- **Event I:** At the beginning of the process, the update of $\mathbf{v}^{(t)}$ is larger than the one of $A_{i,j}^{(t)}$ which implies that only $\mathbf{v}^{(t)}$ updates during this first phase. We show that $\alpha^{(t)} = \langle \mathbf{v}^{(t)}, \mathbf{w}^* \rangle$ increases until a time $\mathcal{T}_0 > 0$ where it reaches some threshold (**Lemma D.2**). At this point, the model is nothing else than a generalized linear model that would not generalize because there are much more noisy tokens than signal ones (see **Theorem 3.1**).
- **Event II:** During this phase, the attention weights must update. Indeed, assume by contradiction that the $A_{i,j}^{(t)}$ stay around initialization and that $\mathbf{v}^{(t)}$ is optimal i.e. $\mathbf{v}^{(t)} = a^{(t)}\mathbf{w}^*$ where $a^{(t)} \gg 1$. Then, the predictor g we would have is

$$g(\mathbf{X}) = \sum_{i=1}^D \sum_{j=1}^D S_{i,j}^{(0)} \langle \mathbf{v}^{(t)}, \mathbf{X}_j \rangle \propto \sum_{i=1}^D \sum_{j=1}^D e^{A_{i,j}^{(0)}} \langle \mathbf{w}^*, \mathbf{X}_j \rangle \quad (2)$$

Such predictor g would yield high population loss because there are many more data with random labels ($qD = \text{poly}(C)$) than with the exact label. Therefore, $A_{i,j}^{(t)}$'s start to update. The gradient increment for $\gamma^{(t)}$ (which corresponds to i and j in the same set \mathcal{S}_ℓ) is much larger than the one for $\rho^{(t)}$ (**Lemma 4.3**). Thus, $\gamma^{(t)}$ increases until a time $\mathcal{T}_1 \in [\mathcal{T}_0, T]$ such that $\gamma^{(\mathcal{T}_1)} > \max_{t \in [T]} |\rho^{(t)}|$.

- **Event III:** Because we have $\gamma^{(\mathcal{T}_1)} > \max_{t \in [T]} |\rho^{(t)}|$, we again have $\alpha^{(t+1)} > \alpha^{(t)}$ as in Phase I (**Lemma D.11**). Thus, $\alpha^{(t)}$ increases again until the population risk becomes a $o(1)$.

Main insights of our analysis. Our mechanism highlights two important aspects that are proper to attention models:

- because of the initialization and the data structure, we have patch association for any time t (**Lemma 4.3**).
- our ViT model uses patch association to minimize the population loss (Event III). Without patch association, the model would only be a generalized linear model that does not minimize the loss.

4.2 From the idealized to the realistic learning process

The real learning process differs from the idealized one in that we have a finite number of samples and we initialize both $\hat{\mathbf{A}}$ and $\hat{\mathbf{v}}$ as Gaussian random variables. Using a polynomial number of samples, we show that (T) still learns patch association and generalizes.

Theorem 4.2. *Assume that we run GD on (E) for T iterations with parameters set as in [Parametrization 3.1](#). Assume that the number of samples is $N = \text{poly}(d)$. With high probability, the model*

1. *learns patch association i.e. for all $\ell \in [L]$ and $i \in \mathcal{S}_\ell$, $\text{Top}_C \{ \hat{A}_{i,j}^{(T)} \}_{j=1}^D = \mathcal{S}_\ell$.*
2. *fits the labeling function i.e. $\mathbb{P}_{\mathcal{D}}[f^*(\mathbf{X}) F_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) > 0] \geq 1 - o(1)$.*

Similarly to [46], the proof introduces a "semi-realistic" learning process that is a mid-point between the idealized and realistic processes. We show that $\hat{\mathbf{A}}^{(T)}$ and $\hat{\mathbf{v}}^{(T)}$ are close to their semi-realistic counterparts – see [Appendix E](#) for a complete proof. [Figure 4](#) numerically illustrates **Theorem 4.2**.

5 Patch association yields sample-efficient fine-tuning with ViTs

A fundamental byproduct of our theory is that after pre-training on a dataset sampled from \mathcal{D} , our model (T) sample-efficiently transfers to datasets that are structured as \mathcal{D} but differ in their features.

Downstream dataset. Let $\tilde{\mathcal{D}}$ a downstream data distribution defined as in [Assumption 1](#) such that its underlying feature is $\tilde{\mathbf{w}}^*$ with $\|\tilde{\mathbf{w}}^*\|_2 = 1$ and $\tilde{\mathbf{w}}^*$ potentially different from \mathbf{w}^* . In other words, the downstream $\tilde{\mathcal{D}}$ and source \mathcal{D} distributions share the same structure but not necessarily the same feature. We sample a downstream dataset $\tilde{\mathcal{Z}} = \{(\tilde{\mathbf{X}}[i], \tilde{\mathbf{y}}[i])\}_{i=1}^{\tilde{N}}$ from $\tilde{\mathcal{D}}$.

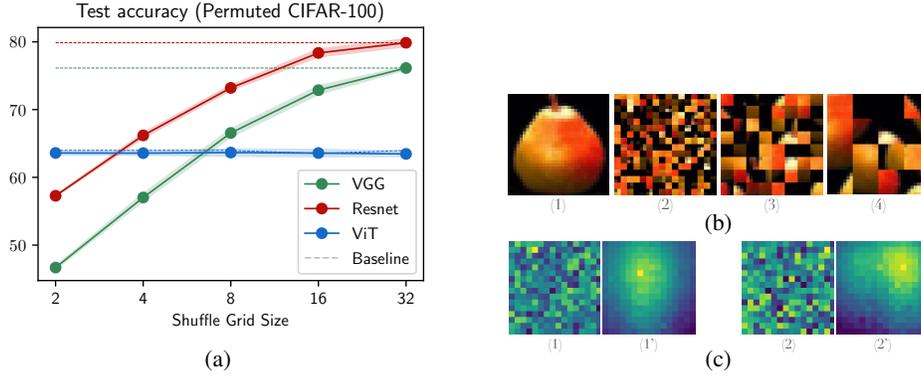


Figure 5: (a): Test accuracy obtained with ViT (patch size 2), ResNet-18 and VGG-19 on permuted (in solid lines) and on original (in dashed lines) CIFAR-100. While convolutional models are very sensitive to permutations, the ViT performs equally whether the dataset is permuted or not. (b): (2) CIFAR-100 image (1) and Permuted CIFAR-100 image when shuffle grid size is 2 (2), 4 (3) and 8 (4). (c): (1-2) Visualization of positional encoding similarities after training a ViT (patch size 2) on permuted CIFAR-100 (shuffle grid size 2). Here, we display $\mathbf{p}_i^T \mathbf{P}$ where i is some fixed index and reshape such vector into a matrix 16×16 . We observe that these similarities (1-2) do not have any spatially localized structure. However, when applying the inverse of the permutation, we recover spatially localized patterns in (1'-2').

Learning problem. We consider the model (T) pre-trained as in subsection 4.2. We assume that $\hat{\mathbf{A}}$ is kept *fixed* from the pre-trained model and we only optimize the value vector $\tilde{\mathbf{v}}$ to solve:

$$\min_{\tilde{\mathbf{v}}} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \log(1 + e^{-\tilde{y}[i]F(\tilde{\mathbf{X}}[i])}) := \tilde{\mathcal{L}}(\tilde{\mathbf{v}}). \quad (\tilde{\mathbf{E}})$$

We run GD on $(\tilde{\mathbf{E}})$ with parameters set as in Parametrization 3.1 except that the $\hat{A}_{i,j}$'s are fixed and $\tilde{\mathbf{v}}^{(0)} \sim \mathcal{N}(0, \omega^2 \mathbf{I}_d)$ with $\omega = 1/\text{poly}(d)$. Our main results states that this fine-tuning procedure requires a few samples to achieve high test accuracy in $\tilde{\mathcal{D}}$. In contrast, any algorithm without patch association needs a large number of samples to generalize.

Theorem 5.1. Let $\hat{\mathbf{A}}$ be the attention matrix obtained after pre-training as in subsection 4.2. Assume that we run GD for T iterations on $(\tilde{\mathbf{E}})$ to fine-tune the value vector. Using $\tilde{N} \leq \text{polylog}(D)$ samples, the model (T) transfers to $\tilde{\mathcal{D}}$ i.e. $\mathbb{P}_{\tilde{\mathcal{D}}}[f^*(\mathbf{X})F_{\hat{\mathbf{A}}, \tilde{\mathbf{v}}(T)}(\mathbf{X}) > 0] \geq 1 - o(1)$.

Theorem 5.2. Let $\mathcal{A}: \mathbb{R}^{d \times D} \rightarrow \{-1, 1\}$ be a binary classification algorithm without patch association knowledge. Then, it needs $D^{\Omega(1)}$ training samples to get test error $\leq o(1)$ on $\tilde{\mathcal{D}}$.

The proofs of Theorem 5.1 and Theorem 5.2 are in Appendix F. These theorems highlight that learning patch association is required for efficient transfer. We believe that they offer a new perspective on explaining why ViTs are widely used in transferring to downstream tasks. While it is possible that ViTs learn shared (with the downstream dataset) features during pretraining, our theory hints that learning the inductive bias of the labeling function is also central for transfer.

6 Numerical experiments

In this section, we first empirically verify that ViTs learn patch association while minimizing their training loss. We then numerically show that the positional attention mechanism competes with the vanilla one on small-scale datasets such as CIFAR-10/100 [43], SVHN [52] and large-scale ones such as ILSVRC-2012 ImageNet [22]. For the small datasets, we use a ViT with 7 layers, 12 heads and hidden/MLP dimension 384. For ImageNet, we train a "ViT-tiny-patch16-224" [24]. Both models are trained with standard augmentations techniques [18] and using AdamW with a cosine learning rate scheduler. We run all the experiments for 300 epochs, with batch size 1024 for Imagenet and 128 otherwise and average our results over 5 seeds. We refer to Appendix A for the training details.

ViTs learn patch association. We consider the CIFAR-100 dataset where we divide each image into grids of size $s \times s$ pixels. For a fixed $s \in \{2, 4, 8, 16, 32\}$, we permute the grids according to π_s to create the permuted CIFAR-100 dataset. We call s the grid shuffle size. Figure 5b-(1) shows a CIFAR-100 image and its corresponding shuffling in the permuted CIFAR-100 dataset Figure 5b-(2-3-4). We train a ViT and CNNs ResNet18 [31] and VGG-19 [57] on the permuted CIFAR-100 dataset.

For the ViT, we set the patch size to 2, although this is sub-optimal in terms of accuracy, because the patch size needs to stay smaller or equal to s . Indeed, intuitively, when we permute the grids in Figure 5b, we lose the local aspect of the spatial structure and create new sets \mathcal{S}_ℓ 's and a new labeling function f^* . Figure 5a reports the test accuracy of these three models for different values of s . When s is small, the image does not have a coherent structure e.g. Figure 5b-(2) and thus, CNNs struggle to generalize. As s increases e.g. Figure 5b-(4), the information inside a patch is meaningful and thus, the CNNs well-perform. Unsurprisingly, since ViTs are *permutation invariant*, their performance remains unchanged for all s – see Figure 5a. Despite this change, we verify that the ViT is able to recover the new \mathcal{S}_ℓ 's: we feed the ViT with the shuffled pear image (Figure 5b-(2)) and consider for some i the similarity matrix $\mathbf{p}_i^\top \mathbf{P}$. We see that it does not exhibit a local spatial structure in Figure 5c-(1,2). We then apply π_s^{-1} to $\mathbf{p}_i^\top \mathbf{P}$ and observe that we recover the spatially localized patterns Figure 5c-(1',2'). This experiment highlights that ViTs do not just group nearby pixels together as convolutions. They learn a more general spatial structure, in accordance to our theoretical results.

ViTs with positional attention are competitive.

We numerically verify that ViTs using positional attention compete with those with vanilla attention. In Section 3, we introduced positional attention to define our theoretical learner model. Figure 6 and Figure 7 show that ViTs using positional attention compete with vanilla ViTs on a range of datasets. These experiments strengthen our intuition that for images, having an attention matrix that only depends on the positional encodings is sufficient to have a good test accuracy.

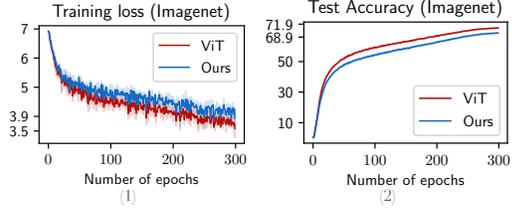


Figure 6: Training loss (1) and test accuracy (2) obtained using a ViT-tiny-patch16-224 on Imagenet. ViT using positional attention (Ours) gets 68.9% test accuracy while vanilla ViT (ViT) gets 71.9%.

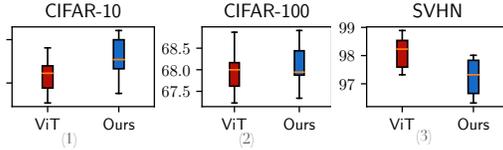


Figure 7: Test accuracy obtained with a ViT using vanilla attention (ViT) and positional attention (Ours) on CIFAR-10 (1), CIFAR-100 (2) and SVHN (3). Our model competes with the vanilla ViT. Patch size 4 and average over 10 seeds for this experiment.

Conclusion, limitations and future works

Our work is a first step towards understanding how Transformers learn tailored inductive biases when trained with gradient descent. Our analysis heavily relies on the positional attention mechanism that disentangles patches and positional encodings. In practice, self-attention mixes these two quantities. An interesting direction is to understand the impact of patch embeddings on the inductive bias learned by ViTs. Moreover, our experiment on the Gaussian data shows that ViTs do not always learn the correct inductive bias under Definition 2.1: characterizing the distributions under which ViTs recover the structure of the function is an important question. Lastly, this work also paves the way to many extensions beyond convolution. For example, can ViTs learn other inductive biases? What are the inductive biases learnt by Transformers in NLP? Answering those questions is central to better understand the underlying mechanism of attention.

Acknowledgments and Disclosure of Funding

The authors would like to thank Boris Hanin for helpful discussions and feedback on this work.

References

[1] Zeyuan Allen-Zhu and Yuanzhi Li. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv preprint arXiv:2012.09816*, 2020.

[2] Zeyuan Allen-Zhu and Yuanzhi Li. Forward super-resolution: How can gans learn hierarchical generative models for real-world distributions. *arXiv preprint arXiv:2106.02619*, 2021.

- [3] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. *arXiv preprint arXiv:1905.13655*, 2019.
- [4] Sanjeev Arora, Zhiyuan Li, and Kaifeng Lyu. Theoretical analysis of auto rate-tuning by batch normalization. *arXiv preprint arXiv:1812.03981*, 2018.
- [5] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- [6] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [7] Irwan Bello. Lambdanetworks: Modeling long-range interactions without attention. *arXiv preprint arXiv:2102.08602*, 2021.
- [8] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3286–3295, 2019.
- [9] Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. Understanding robustness of transformers for image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10231–10241, 2021.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [11] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [12] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021.
- [13] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34, 2021.
- [14] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021.
- [15] Zixiang Chen, Yihe Deng, Yue Wu, Quanquan Gu, and Yuanzhi Li. Towards understanding mixture of experts in deep learning. *arXiv preprint arXiv:2208.02813*, 2022.
- [16] Lenaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338. PMLR, 2020.
- [17] Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. On the relationship between self-attention and convolutional layers. *arXiv preprint arXiv:1911.03584*, 2019.
- [18] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018.
- [19] Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. *Advances in Neural Information Processing Systems*, 34:3965–3977, 2021.
- [20] Amit Daniely. Sgd learns the conjugate kernel class of the network. In *Advances in Neural Information Processing Systems*, pages 2422–2430, 2017.

- [21] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018.
- [22] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [23] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [24] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [25] Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *International Conference on Machine Learning*, pages 2286–2296. PMLR, 2021.
- [26] Benjamin L Edelman, Surbhi Goel, Sham Kakade, and Cyril Zhang. Inductive biases and variable creation in self-attention mechanisms. *arXiv preprint arXiv:2110.10090*, 2021.
- [27] Kunihiko Fukushima. Neocognitron for handwritten digit recognition. *Neurocomputing*, 51:161–180, 2003.
- [28] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- [29] Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. *Advances in Neural Information Processing Systems*, 31, 2018.
- [30] Jianyuan Guo, Kai Han, Han Wu, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. Cmt: Convolutional neural networks meet vision transformers. *arXiv preprint arXiv:2107.06263*, 2021.
- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [32] Elad Hoffer, Ron Banner, Itay Golan, and Daniel Soudry. Norm matters: efficient and accurate normalization schemes in deep networks, 2019.
- [33] Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *arXiv preprint arXiv:1705.08741*, 2017.
- [34] Jiri Hron, Yasaman Bahri, Jascha Sohl-Dickstein, and Roman Novak. Infinite attention: Nngp and ntk for deep attention networks. In *International Conference on Machine Learning*, pages 4376–4386. PMLR, 2020.
- [35] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer, 2016.
- [36] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [37] Meena Jagadeesan, Ilya Razenshteyn, and Suriya Gunasekar. Inductive bias of multi-channel linear convolutional networks with bounded weight norm. In *Conference on Learning Theory*, pages 2276–2325. PMLR, 2022.
- [38] Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34, 2021.

- [39] Samy Jelassi and Yuanzhi Li. Towards understanding how momentum improves generalization in deep learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 9965–10040. PMLR, 17–23 Jul 2022.
- [40] Samy Jelassi, Arthur Mensch, Gauthier Gidel, and Yuanzhi Li. Adam is no better than normalized SGD: Dissecting how adaptivity improves GAN performance, 2022.
- [41] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- [42] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [43] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [44] Guillaume Lample and François Charton. Deep learning for symbolic mathematics. *arXiv preprint arXiv:1912.01412*, 2019.
- [45] Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pages 2–47. PMLR, 2018.
- [46] Yuanzhi Li, Tengyu Ma, and Hongyang R Zhang. Learning over-parametrized two-layer relu neural networks beyond ntk. *arXiv preprint arXiv:2007.04596*, 2020.
- [47] Yuanzhi Li, Colin Wei, and Tengyu Ma. Towards explaining the regularization effect of initial large learning rate in training neural networks. *arXiv preprint arXiv:1907.04595*, 2019.
- [48] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [49] Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. *arXiv preprint arXiv:1906.05890*, 2019.
- [50] Luke Melas-Kyriazi. Do you even need attention? a stack of feed-forward layers does surprisingly well on imagenet. *arXiv preprint arXiv:2105.02723*, 2021.
- [51] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *Advances in Neural Information Processing Systems*, 34, 2021.
- [52] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [53] Behnam Neyshabur, Ruslan Salakhutdinov, and Nathan Srebro. Path-sgd: Path-normalized optimization in deep neural networks. *arXiv preprint arXiv:1506.02617*, 2015.
- [54] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. *arXiv preprint arXiv:2105.07581*, 2(3), 2021.
- [55] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems*, 34, 2021.
- [56] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. *Advances in Neural Information Processing Systems*, 32, 2019.
- [57] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

- [58] Samuel L. Smith, Pieter-Jan Kindermans, Chris Ying, and Quoc V. Le. Don't decay the learning rate, increase the batch size, 2018.
- [59] Charlie Snell, Ruiqi Zhong, Dan Klein, and Jacob Steinhardt. Approximating how single head attention learns. *arXiv preprint arXiv:2103.07601*, 2021.
- [60] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- [61] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [62] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [63] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.
- [64] Asher Trockman and J Zico Kolter. Patches are all you need? *arXiv preprint arXiv:2201.09792*, 2022.
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [66] James Vuckovic, Aristide Baratin, and Remi Tachet des Combes. A mathematical theory of attention. *arXiv preprint arXiv:2007.02876*, 2020.
- [67] Alex Warstadt and Samuel R Bowman. Can neural networks acquire a structural bias from raw linguistic data? *arXiv preprint arXiv:2007.06761*, 2020.
- [68] Colin Wei, Yining Chen, and Tengyu Ma. Statistically meaningful approximation: a case study on approximating turing machines with transformers. *arXiv preprint arXiv:2107.13163*, 2021.
- [69] Colin Wei, Sham Kakade, and Tengyu Ma. The implicit and explicit regularization effects of dropout, 2020.
- [70] Zixin Wen and Yuanzhi Li. Toward understanding the feature learning process of self-supervised contrastive learning. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11112–11122. PMLR, 18–24 Jul 2021.
- [71] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nathan Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. *arXiv preprint arXiv:1705.08292*, 2017.
- [72] Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, pages 3635–3673. PMLR, 2020.
- [73] Yuhuai Wu, Markus N Rabe, Wenda Li, Jimmy Ba, Roger B Grosse, and Christian Szegedy. Lime: Learning inductive bias for primitives of mathematical reasoning. In *International Conference on Machine Learning*, pages 11251–11262. PMLR, 2021.
- [74] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.
- [75] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.

- [76] Jingzhao Zhang, Sai Praneeth Karimireddy, Andreas Veit, Seungyeon Kim, Sashank Reddi, Sanjiv Kumar, and Suvrit Sra. Why are adaptive methods good for attention models? *Advances in Neural Information Processing Systems*, 33:15383–15393, 2020.
- [77] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13001–13008, 2020.
- [78] Difan Zou, Yuan Cao, Yuanzhi Li, and Quanquan Gu. Understanding the generalization of adam in learning neural networks with proper regularization. *arXiv preprint arXiv:2108.11371*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]** See [Section 4](#) and [Section 5](#).
 - (b) Did you describe the limitations of your work? **[Yes]** See [Section 6](#).
 - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]** This is a theory paper.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** See [Section 3](#).
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See Appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** See supplementary material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See Appendix.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]** See [Section 6](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** See Appendix.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]**
 - (b) Did you mention the license of the assets? **[N/A]**
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

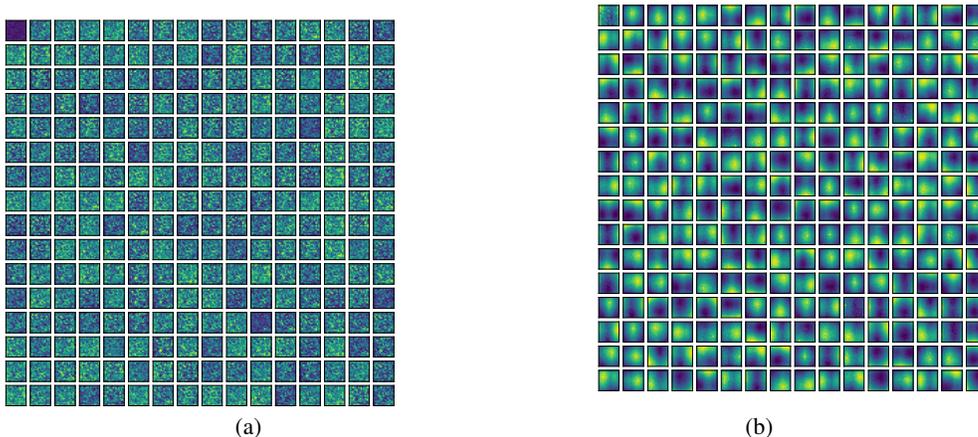


Figure 8: (a) Visualization of the positional encodings similarities when feeding the ViT with Permuted CIFAR-100 data. Each array represents $\langle \mathbf{p}_i, \mathbf{p}_j \rangle_{j=1}^D$ for a fixed i . (b) displays the positional encodings similarities obtained after inverting the permutation. We see that the ViT is able to recover the convolution-like structure in all the cases.

A Additional experimental details

In this section, we provide additional details on our experiments and additional plots.

A.1 Details on the implementation

We used Pytorch and Nvidia Tesla V100 GPUs. We conduct experiments on small-scale (CIFAR-10/100 and SVHN) and large-scale datasets (ImageNet). The choice of architecture and training parameters depend on the size of the dataset as we detail below.

Small-scale experiments. We use the code available at <https://github.com/omihub777/ViT-CIFAR>. The model is made of 7 layers, 12 heads, hidden and MLP dimension 384, dropout 0. We use "mean-pooling" and not the CLS pooling. We set the patch size to 2 in the experiment Figure 5 and to 4 in the experiment Figure 7. Indeed, we empirically found that setting patch size 4 was the optimal choice. We apply label smoothing [62] with coefficient 0.1 and do not apply any cutmix [75] nor mixup [74]. We use Adam [42] as optimizer and set the learning rate to 10^{-3} , minimum learning rate to 10^{-5} , β_1 to 0.9, β_2 to 0.999, batch size to 128, weight decay to $5 \cdot 10^{-5}$, number of warmup epochs to 5 and number of total epochs to 200. The scheduler is a cosine learning rate. We used the AutoAugment procedure [18] as in the repository to generate data augmentations. The model has been trained over a single GPU.

Regarding the convolutional models in the experiment Figure 5, we trained a ResNet-18 and a VGG-19 with batch normalization. We trained the two architectures using the same training procedure and hyperparameters as for the ViT.

Large-scale experiments. We use the code available at <https://github.com/facebookresearch/deit>. Due to limited computational resources, we train a ViT-tiny-patch16-224 [24] where "CLS-pooling" is applied. A detailed table with the hyperparameters used for the ImageNet experiment may be found in Table 9 (column "DeiT-B") in [63]. We set no dropout but set stochastic depth [35] 0.1. We used label smoothing 0.1. Regarding the augmentations, we set RandAugment [18] 9/0.5, mixup 0.8, cutmix 1, erasing probability [77] 0.25. Lastly, we trained the model using AdamW [48] and set the batch size to 1024, learning rate to $5 \cdot 10^{-4} \cdot \frac{\text{batchsize}}{512}$ as in [28], weight decay to 0.05, warmup epochs 5 and number of total epochs to 300. The total number of epochs is 300. The model has been trained over 16 GPUs (8 nodes and 2 GPUs per node) and batch size for each device is 64.

A.2 Additional plots

In Figure 5, we plot the positional encoding similarities for a few patches. Figure 8 provides these plots for all the patches. One should think of Figure 5 as a Figure displaying just two of the arrays present in Figure 8. We consistently verify that the ViT is always able to recover the convolution-like patterns which shows that it is able to learn the right patch association.

B Induction hypothesis

In this section, we present the induction hypothesis that we use in the analysis of the idealized case. This hypothesis is ultimately proved in [subsection D.7](#).

Induction hypothesis B.1. *During the idealized learning process, the following holds for $t \leq T$.*

- the softmax denominator is large i.e. $\Lambda^{(t)} + (C - 1)\Gamma^{(t)} + (D - C)\Xi^{(t)} = \Theta(D)$.
- $\Xi^{(t)}$ is not too small i.e. $\Xi^{(t)} = \Theta(1/D)$.
- $\Gamma^{(t)}$ and $\Lambda^{(t)}$ are in a good range i.e.

$$\Lambda^{(t)} = \frac{e^{\text{polyloglog}(d)}}{D}, \quad (C - 1)\Gamma^{(t)} \in \left[\frac{\Omega(C)}{D}, \frac{\lambda_0}{D} \right],$$

where $\lambda_0 = \Theta(D^{0.01})$.

C Notations

In this section, we introduce the different notations used in the proofs.

General purpose. We first define notations that are used everywhere in the appendix.

- Sigmoid function: Given $z \in \mathbb{R}$, $\mathfrak{S}(z) = 1/(1 + e^{-z})$.
- Softmax function: Given $\mathbf{z} = (z_1, \dots, z_D) \in \mathbb{R}^D$, $(\text{softmax}(z_1, \dots, z_D))_m = \frac{e^{z_m}}{\sum_{j=1}^D e^{z_j}}$.
- Loss for a data-point (\mathbf{X}, y) : $L(\mathbf{X}) = \log(1 + e^{-yF(\mathbf{X})})$.

Analysis in idealized case. We now provide notations used in the analysis of the idealized case.

- $\kappa(\mathbf{X}) = \sum_{\ell \neq \ell(\mathbf{X})} \sum_{j \in \mathcal{S}_\ell} \delta_j + C$.
- for $i \in [D]$, $\mathbf{O}_i^{(t)} = \sum_{j=1}^D S_{i,j}^{(t)} \mathbf{X}_j$.

Analysis in realistic case. We now provide notations used in the analysis of the realistic case.

- Score matrix: $\mathbf{S} \in \mathbb{R}^{D \times D}$ with coefficients $\hat{S}_{i,j} = \exp(\hat{A}_{i,j}/\sqrt{d}) / \sum_{r=1}^D \exp(\hat{A}_{i,r}/\sqrt{D})$.
- Given a data-point $(\mathbf{X}[i], y[i])$ and $j \in [D]$, $\mathbf{O}_j^{(t)}[i] = \sum_{k=1}^D S_{j,k}^{(t)} \mathbf{X}_k[i]$

D Learning process in the idealized setting

D.1 Roadmap of the proof

From [Lemma 4.2](#), we know that $\mathbf{v}^{(t)} \in \text{span}(\mathbf{w}^*)$ for all $t \in [T]$. The main idea of the proof consists in analyzing the GD dynamics of $\alpha^{(t)} = \langle \mathbf{v}^{(t)}, \mathbf{w}^* \rangle$ that satisfy

$$\begin{aligned} \alpha^{(t+1)} &= \alpha^{(t)} + \eta \mathbb{E} \left[\underbrace{y \mathfrak{S}(-yF_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X})) \sum_{i \in \mathcal{S}_\ell(\mathbf{X})} \sigma' \left(\sum_{j=1}^D S_{i,j}^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_j \rangle \right) \sum_{k=1}^D S_{i,k}^{(t)} \langle \mathbf{w}^*, \mathbf{X}_k \rangle}_{\mathcal{S}^{(t)}} \right] \\ &+ \eta \mathbb{E} \left[\underbrace{y \mathfrak{S}(-yF_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X})) \sum_{i \notin \mathcal{S}_\ell(\mathbf{X})} \sigma' \left(\sum_{j=1}^D S_{i,j}^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_j \rangle \right) \sum_{m=1}^D S_{i,m}^{(t)} \langle \mathbf{w}^*, \mathbf{X}_m \rangle}_{\mathcal{N}^{(t)}} \right]. \quad (\text{GD-}\alpha) \end{aligned}$$

We divide the idealized learning process as follows.

- **Event I** ($t \in [0, \mathcal{T}_0]$, [subsection D.2](#)): at initialization, $\alpha^{(0)}$ is small. Therefore, the sigmoid $\mathfrak{S}(-yF(\mathbf{X}))$ is large. Besides, around $\alpha^{(0)}$, it stays constant i.e. $\mathfrak{S}(-yF_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X})) \approx \mathfrak{S}(-yF_{\mathbf{A}^{(0)}, \mathbf{v}^{(0)}}(\mathbf{X}))$ in (GD- α). This implies that $\mathfrak{N}^{(t)} = 0$ which yields $\alpha^{(t)}$ to increase until reaching a specific value where the sigmoid is not constant anymore.
- **Event II** ($t \in [\mathcal{T}_0, \mathcal{T}_1]$, [subsection D.3](#)): at time \mathcal{T}_0 , $\alpha^{(t)}$ is large. This fact along with [Lemma 4.3](#) imply that $\gamma^{(t)}$ increases. Eventually, $\Gamma^{(\mathcal{T}_1)}$ becomes large enough so that $\mathfrak{S}^{(t)} \geq \max_{\tau \leq T} |\mathfrak{N}^{(\tau)}|$.
- **Event III** ($t \in [\mathcal{T}_1, T]$, [subsection D.4](#)): Since $\mathfrak{S}^{(t)} \geq \max_{\tau \leq T} |\mathfrak{N}^{(\tau)}|$, $\alpha^{(t)}$ increases again. It increases until the population risk is at most $o(1)$.

After T iterations, $\alpha^{(t)}$ is large and the population risk thus converges ([subsection D.5](#)). Since the logistic loss is a surrogate for the 0-1 loss, we prove that the learner model fits the labeling function ([subsection D.6](#)) which implies the first statement of [Theorem 4.1](#).

Remark : Since we initialize $\alpha^{(0)} \geq \nu^{1/(p-1)}$, [Lemma D.7](#) implies that we can overlook the linear part of the activation in this section. Therefore, we only consider $\sigma(x) = x^p$ in the idealized process.

D.2 Event I: $\alpha^{(t)}$ initially increases

A first question that arises is: starting from $\alpha^{(0)}$, what is the value of $\alpha^{(t)}$ that makes the sigmoid non-constant? The following lemma addresses this question.

Lemma D.1. *The value $\alpha^{(t)}$ at which the sigmoid $\mathfrak{S}(-yF(\mathbf{X}))$ becomes non-constant is:*

$$\tilde{\alpha} = \frac{\Theta(1)}{C^2 \lambda_0}.$$

Proof of Lemma D.1. The update of $\alpha^{(t)}$ is

$$\begin{aligned} \alpha^{(t+1)} &= \alpha^{(t)} + \eta \mathbb{E}[y \mathfrak{S}(-yF(\mathbf{X})) \mathcal{G}^{(t)}] \\ &= \alpha^{(t)} + \eta \mathbb{E}[y \mathfrak{S}(-\alpha^{(0)}) \mathcal{G}^{(t)}] + \eta \mathbb{E}[y (\mathfrak{S}(-yF(\mathbf{X})) - \mathfrak{S}(-\alpha^{(0)})) \mathcal{G}^{(t)}], \end{aligned} \quad (3)$$

where $\mathcal{G}^{(t)} := \sum_{j=1}^D \langle \mathbf{O}_j^{(t)}, \mathbf{w}^* \rangle^p$. Since $x \mapsto \mathfrak{S}(-x)$ is 1/4-Lipschitz, we rewrite (3) as:

$$\begin{aligned} |\alpha^{(t+1)} - \alpha^{(t)} - \eta \mathbb{E}[y \mathfrak{S}(-\alpha^{(0)}) \mathcal{G}^{(t)}]| &\leq \frac{\eta}{4} \mathbb{E}[|F(\mathbf{X}) - y \alpha^{(0)}| \cdot |y \mathcal{G}^{(t)}|] \\ &\leq \Theta(C\eta) (\alpha^{(t)})^p \mathbb{E}[(C^{1+1/p} \lambda_0)^p |y \mathcal{G}^{(t)}|] \\ &\leq \Theta(\eta) (\alpha^{(t)})^p \mathbb{E}[(C^2 \lambda_0)^p |y \mathcal{G}^{(t)}|], \end{aligned} \quad (4)$$

where we applied [Lemma D.3](#), [Induction Hypothesis B.1](#) and the fact that $\alpha^{(0)}$ is small in the penultimate inequality. Using [Lemma D.4](#) and [Lemma D.5](#), we have $\mathbb{E}[|y \mathcal{G}^{(t)}|] \leq \Theta(1) \mathbb{E}[y \mathcal{G}^{(t)}]$ which yields

$$|\alpha^{(t+1)} - \alpha^{(t)} - \eta \mathbb{E}[y \mathfrak{S}(-\alpha^{(0)}) \mathcal{G}^{(t)}]| \leq \Theta(\eta) (\alpha^{(t)})^p \mathbb{E}[(C^2 \lambda_0)^p y \mathcal{G}^{(t)}]. \quad (5)$$

(4) shows that when $\alpha^{(t)}$ is small, we have $\alpha^{(t+1)} \approx \alpha^{(t)} + \eta \mathbb{E}[y \mathfrak{S}(-\alpha^{(0)}) \mathcal{G}^{(t)}]$. Besides, we have $\mathbb{E}[y \mathcal{G}^{(t)}] \geq 0$ so $\alpha^{(t)}$ increases. However, during this increase, the right-hand side of (4) increases and this approximation does not hold anymore. Therefore, the sigmoid is approximately constant when $\alpha^{(t)}$ satisfies:

$$\Theta(1) (\alpha^{(t)})^p \mathbb{E}[(C^2 \lambda_0)^p y \mathcal{G}^{(t)}] \leq \mathbb{E}[\mathfrak{S}(-\alpha^{(0)}) y \mathcal{G}^{(t)}] \iff \alpha^{(t)} \leq \frac{\Theta(1)}{C^2 \lambda_0}. \quad (6)$$

□

Lemma D.2. *Let $\mathcal{T}_0 = \Theta\left(\frac{1}{\eta C (\alpha^{(0)})^{p-1}}\right)$. For all $t \in [0, \mathcal{T}_0]$, we have $\mathfrak{N}^{(t)} = 0$. Therefore, $\alpha^{(t)}$ is updated as*

$$\alpha^{(t+1)} = \alpha^{(t)} + \Theta(\eta C) (G^{(t)})^p (\alpha^{(t)})^{p-1}.$$

Consequently, $\alpha^{(t)}$ is non-decreasing and after \mathcal{T}_0 iterations, we have $\alpha^{(t)} \geq \frac{\Omega(1)}{C^2 \lambda_0}$ for $t \geq \mathcal{T}_0$.

Proof of Lemma D.2. For $t \in [0, \mathcal{T}_0]$, we know that the sigmoid $\mathfrak{S}(-yF(\mathbf{X}))$ is constant. We apply Lemma D.4 and Lemma D.6 to respectively bound $\mathcal{S}^{(t)}$ and $\mathcal{N}^{(t)}$ in the update of $\alpha^{(t)}$.

$$\alpha^{(t+1)} = \alpha^{(t)} + \Theta(\eta C)(\alpha^{(t)})^{p-1}(G^{(t)})^p. \quad (7)$$

We apply Induction Hypothesis B.1 and $\Omega(C) \geq 0$ and $e^{\text{polyloglog}(d)} \leq \lambda_0$ in (7) and obtain:

$$\begin{cases} \alpha^{(t+1)} \geq \alpha^{(t)} + \Theta(\eta C)e^{\text{polyloglog}(d)}(\alpha^{(t)})^{p-1} \\ \alpha^{(t+1)} \leq \alpha^{(t)} + \Theta(\eta C)\lambda_0^p(\alpha^{(t)})^{p-1} \end{cases}. \quad (8)$$

(8) indicates that $\alpha^{(t)}$ is a non-decreasing sequence. Therefore, there exists a time \mathcal{T}_0 such that $\alpha^{(\mathcal{T}_0)} = \frac{\Theta(1)}{C^2\lambda_0}$. Using Lemma K.1, the time \mathcal{T}_0 is equal to:

$$\mathcal{T}_0 = \frac{1}{\eta C(\alpha^{(0)})^{p-1}e^{\text{polyloglog}(d)}} + (\lambda_0)^p e^{-\text{polyloglog}(d)} \left[-\log(C^2\lambda_0\alpha^{(0)}) \right]. \quad (9)$$

□

D.2.1 Auxiliary lemmas

In this section, we present the auxiliary lemmas needed to prove the main results of subsection D.2. We first present a lemma that bounds the learner model.

Lemma D.3. *Let $t \in [0, T]$. The learner model F is bounded for all $(\mathbf{X}, y) \sim \mathcal{D}$ as:*

$$yF_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X}) \leq \Theta(1)(C^{1+1/p}\alpha^{(t)}\lambda_0)^p.$$

Proof of Lemma D.3. By definition, the learner model is:

$$yF_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X}) = y \sum_{i=1}^D \langle \mathbf{v}^{(t)}, \mathbf{O}_i^{(t)} \rangle^p = \sum_{i \in \mathcal{S}_\ell(\mathbf{X})} \langle \mathbf{v}^{(t)}, \mathbf{O}_i^{(t)} \rangle^p + \sum_{i \notin \mathcal{S}_\ell(\mathbf{X})} \langle \mathbf{v}^{(t)}, \mathbf{O}_i^{(t)} \rangle^p. \quad (10)$$

We successively apply Lemma D.4, Induction Hypothesis B.1 and Lemma D.5 to bound (10).

$$yF_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X}) \leq (\alpha^{(t)})^p \left(C\Theta(DCT^{(t)})^p + o(\lambda_0)^p \right). \quad (11)$$

Finally, we apply Induction Hypothesis B.1 in (11) to obtain the desired result. □

We now present lemmas that bound $\mathcal{S}^{(t)}$ and $\mathcal{N}^{(t)}$.

Lemma D.4. *Let $t \in [T]$ and $i \in \mathcal{S}_\ell(\mathbf{X})$. We have $y\langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p = \Theta(G^{(t)})^p$. As long as the population risk is not $o(1)$, we have $\mathcal{S}^{(t)} = C\Theta(G^{(t)})^p$ for all $t \leq T$.*

Proof of Lemma D.4. We apply Lemma G.3 and obtain:

$$\begin{aligned} y\langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p &= D^p \left((\Lambda^{(t)} + (C-1)\Gamma^{(t)}) + y\Xi^{(t)} \sum_{h \in [L] \setminus \{\ell(\mathbf{X})\}} \sum_{r \in \mathcal{S}_h} \delta_r \right)^p \\ &= \Theta \left(D(\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^p \\ &= \Theta(G^{(t)})^p. \end{aligned}$$

□

Lemma D.5. *Let $t \in [T]$. We have $\sum_{i \notin \mathcal{S}_\ell(\mathbf{X})} y\langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p \leq o(\lambda_0)^p$. In particular, this implies $\mathcal{N}^{(t)} \leq o(\lambda_0)^p$ for all $t \leq T$.*

Proof of Lemma D.5. We have:

$$\begin{aligned} & \sum_{k \in [L] \setminus \{\ell(\mathbf{X})\}} \sum_{i \in \mathcal{S}_k} y \langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p \\ &= D^p \sum_{k \in [L] \setminus \{\ell(\mathbf{X})\}} \sum_{m \in \mathcal{S}_k} y \left(\Lambda^{(t)} \delta_m + \Gamma^{(t)} \sum_{s \in \mathcal{S}_k \setminus \{m\}} \delta_s + \Xi^{(t)} \left[C y + \sum_{h \in [L] \setminus \{\ell(\mathbf{X}), k\}} \sum_{r \in \mathcal{S}_h} \delta_r \right] \right)^p. \end{aligned} \quad (12)$$

We distinguish two cases.

– $\delta_r = 0$ for all $r \in \mathcal{S}_\ell$: we apply [Lemma K.3](#) and obtain:

$$\begin{aligned} \sum_{k \neq \ell(\mathbf{X})} \sum_{i \in \mathcal{S}_k} y \langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p &= \sum_{k \neq \ell(\mathbf{X})} \sum_{m \in \mathcal{S}_k} y \left(D \Xi^{(t)} \left[C y + \sum_{h \neq \{\ell(\mathbf{X}), k\}} \sum_{r \in \mathcal{S}_h} \delta_r \right] \right)^p \\ &\leq D \cdot \Theta(D \Xi^{(t)} q D \log(d))^p. \end{aligned} \quad (13)$$

– $\exists r \in \mathcal{S}_\ell$ such that $\delta_r \neq 0$: let $i \in \mathcal{S}_k$. We apply [Lemma K.4](#) and obtain:

$$y \langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p \leq \Theta(D(\Lambda^{(t)} + O(1)\Gamma^{(t)}))^p \cdot \mathbf{1}_{\exists r \in \mathcal{S}_k, \delta_r \neq 0}. \quad (14)$$

We now sum (14) and apply [Lemma K.3](#) to obtain:

$$\begin{aligned} \sum_{k \neq \ell(\mathbf{X})} \sum_{i \in \mathcal{S}_k} y \langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p &\leq \Theta(D(\Lambda^{(t)} + O(1)\Gamma^{(t)}))^p \sum_{k \neq \ell(\mathbf{X})} \sum_{i \in \mathcal{S}_k} |\delta_r| \\ &\leq \Theta(D(\Lambda^{(t)} + O(1)\Gamma^{(t)}))^p q D \log(d). \end{aligned} \quad (15)$$

We finally apply [Induction Hypothesis B.1](#) to have $\Gamma^{(t)} \leq \lambda_0/C$ in (15) and get

$$\sum_{k \neq \ell(\mathbf{X})} \sum_{i \in \mathcal{S}_k} y \langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p \leq \Theta(\lambda_0/C)^p q D \log(d). \quad (16)$$

We finally plug (13) and (16) in (12) and obtain:

$$\sum_{k \neq \ell(\mathbf{X})} \sum_{i \in \mathcal{S}_k} y \langle \mathbf{w}^*, \mathbf{O}_i^{(t)} \rangle^p \leq D \cdot \Theta(D \Xi^{(t)} q D \log(d))^p + \Theta(\lambda_0/C)^p q D \log(d) \leq o(\lambda_0)^p.$$

□

Lemma D.6. Let $t \in [0, \mathcal{T}_0]$. We have $\mathcal{N}^{(t)} = 0$.

Proof of Lemma D.6. By definition of $\mathcal{N}^{(t)}$, we have:

$$\begin{aligned} & \mathcal{N}^{(t)} \\ &= \frac{1}{2} \left(\mathbb{E} \left[\mathfrak{S}(-F(\mathbf{X})) \sum_{j \notin \mathcal{S}_\ell(\mathbf{X})} \langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p \middle| y = 1 \right] - \mathbb{E} \left[\mathfrak{S}(F(\mathbf{X})) \sum_{j \notin \mathcal{S}_\ell(\mathbf{X})} \langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p \middle| y = -1 \right] \right) \\ &= \Theta(1) \left(\mathbb{E} \left[\sum_{j \notin \mathcal{S}_\ell(\mathbf{X})} \langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p \middle| y = 1 \right] - \mathbb{E} \left[\sum_{j \notin \mathcal{S}_\ell(\mathbf{X})} \langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p \middle| y = -1 \right] \right), \end{aligned} \quad (17)$$

where we use $\mathfrak{S}(-F(\mathbf{X})) \approx \mathfrak{S}(-F_{\mathbf{A}^{(0)}, \mathbf{v}^{(0)}}(\mathbf{X}))$ for $t \in [0, \mathcal{T}_0]$ in the last equality of (17). We now show that each of the summands in (17) is zero. Without loss of generality, let's focus on the first summand. The same reasoning holds for the second one. In particular, for $j \in \mathcal{S}_k$ with $k \neq \ell$, we now compute $\mathbb{E}[\langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p | y = 1]$. Using the binomial theorem and the independence of the δ_r 's, we have:

$$\begin{aligned} & \mathbb{E}[\langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p | y = 1] \\ &= \sum_{a=0}^p \binom{p}{a} (\Lambda^{(t)})^{p-a} \mathbb{E}[(\delta_j)^{p-a}] \mathbb{E} \left[\left(\Gamma^{(t)} \sum_{s \in \mathcal{S}_k \setminus \{j\}} \delta_s + \Xi^{(t)} \left[C + \sum_{h \neq \{\ell(\mathbf{X}), k\}} \sum_{r \in \mathcal{S}_h} \delta_r \right] \right)^a \right]. \end{aligned} \quad (18)$$

For a even, we have $p - a$ odd which implies $\mathbb{E}[(\delta_u)^{p-a}] = 0$. Therefore, the summands with even a are zero. We now focus on the case a odd. We again apply the binomial theorem and the independence of the δ_r 's to get:

$$\begin{aligned} & \mathbb{E} \left[\left(\Gamma^{(t)} \sum_{s \in \mathcal{S}_k \setminus \{j\}} \delta_s + \Xi^{(t)} \left[C + \sum_{h \neq \{\ell(\mathbf{X}), k\}} \sum_{r \in \mathcal{S}_h} \delta_r \right] \right)^a \right] \\ &= \sum_{b=0}^a \binom{a}{b} (\Gamma^{(t)})^{a-b} \mathbb{E} \left[\left(\sum_{s \in \mathcal{S}_k \setminus \{j\}} \delta_s \right)^{a-b} \right] (\Xi^{(t)})^b \mathbb{E} \left[\left(Cy + \sum_{h \neq \{\ell(\mathbf{X}), k\}} \sum_{r \in \mathcal{S}_h} \delta_r \right)^b \right]. \end{aligned} \quad (19)$$

For b even, we have $a - b$ odd. This implies that $\mathbb{E}[(\sum_{s \in \mathcal{S}_k \setminus \{j\}} \delta_s)^{a-b}] = 0$. In the case b odd, we exactly use the same argument and obtain:

$$\mathbb{E} \left[\left(Cy + \sum_{h \neq \{\ell(\mathbf{X}), k\}} \sum_{r \in \mathcal{S}_h} \delta_r \right)^b \right] = 0. \quad (20)$$

(20) implies (19) is zero and which lastly implies (18) is zero. We conclude that $\mathbb{E}[\langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p | y = 1] = \mathbb{E}[\langle \mathbf{w}^*, \mathbf{O}_j^{(t)} \rangle^p | y = -1] = 0$ and thus $\mathfrak{N}^{(t)} = 0$. □

Lemma D.7. Let $(\mathbf{X}, \cdot) \sim \mathcal{D}$ and $j \in [D]$. Assume that $\alpha^{(t)} \geq \nu^{1/(p-1)}$. Then, we have:

$$\sigma' \left(D \sum_{k=1}^D \mathbf{S}_{j,k}^{(t)} \langle \mathbf{v}, \mathbf{X}_k \rangle \right) = \Theta(1) \left(D \sum_{k=1}^D \mathbf{S}_{j,k}^{(t)} \langle \mathbf{v}, \mathbf{X}_k \rangle \right)^{p-1}.$$

Proof of Lemma D.7. We remind that the derivative of the activation function $\sigma'(x) = px^{p-1} + \nu$. We first remark that for all x , $\sigma'(x) \geq px^{p-1}$. Besides, we have:

$$px^{p-1} + \nu \leq \frac{3px^{p-1}}{2} \iff x \geq \left(\frac{2\nu}{p} \right)^{1/(p-1)}. \quad (21)$$

In our case, we have $x = D\alpha^{(t)} \sum_{k=1}^D \mathbf{S}_{j,k}^{(t)} \langle \mathbf{w}^*, \mathbf{X}_k \rangle$. Using [Induction Hypothesis B.1](#), we have $x \leq C\lambda_0\alpha^{(t)}$. Therefore, a sufficient condition for (21) to hold is $\alpha^{(t)} \geq \nu^{1/(p-1)}/(C\lambda_0)$. Since $\lambda_0, C \ll \text{poly}(d)$, we can simplify this condition as $\alpha^{(t)} \geq \nu^{1/(p-1)}$. □

D.3 Event II: $\Gamma^{(t)}$ significantly increases

In this section, we show the increase of $\alpha^{(t)}$ for $t \in [0, \mathcal{T}_0]$ leads to the increase of $\Gamma^{(t)}$. At time $\mathcal{T}_1 > \mathcal{T}_0$, $\Gamma^{(t)}$ is significantly large.

Lemma D.8. Let $\mathcal{T}_1 = \mathcal{T}_0 + \Theta\left(\frac{\lambda_0^p}{\eta e^{\text{poly} \log \log(d)}}\right)$. For all $t \in [\mathcal{T}_1, T]$, we have $\Gamma^{(t)} \geq \frac{\Omega(\lambda_0)}{D}$. This implies $\mathfrak{S}^{(t)} \geq \max_{\tau \leq T} |\mathfrak{N}^{(\tau)}|$.

Proof of Lemma D.8. Let $t \in [\mathcal{T}_0, T]$ and $\tau \in [\mathcal{T}_0, t]$. Using [Corollary G.1](#) and [Induction Hypothesis B.1](#), $\gamma^{(\tau)}$ satisfies:

$$\gamma^{(\tau+1)} \geq \gamma^{(\tau)} + \Omega(\eta C) (\alpha^{(\tau)})^p e^{\text{poly} \log \log(d)}. \quad (22)$$

Summing (22) for $\tau = \mathcal{T}_0, \dots, t-1$ yields

$$\gamma^{(t)} \geq \gamma^{(\mathcal{T}_0)} + e^{\text{poly} \log \log(d)} \Omega(\eta C) \sum_{\tau=\mathcal{T}_0}^{t-1} (\alpha^{(\tau)})^p. \quad (23)$$

We successively apply [Lemma D.9](#) and $(a - b)^p \geq a^p - pb$ for $a \ll b$ to lower bound [\(23\)](#) to obtain:

$$\gamma^{(t)} \geq \gamma^{(\mathcal{T}_0)} + e^{\text{polyloglog}(d)} \Omega(\eta C) \sum_{\tau=\mathcal{T}_0}^{t-1} \left(\frac{\Omega(1)}{C^{2p} \lambda_0^p} - o(\eta) \right). \quad (24)$$

We apply [Induction Hypothesis B.1](#) in [\(24\)](#) to obtain a bound on $\Gamma^{(t)}$.

$$\begin{aligned} C\Gamma^{(t)} &\geq C\Gamma^{(\mathcal{T}_0)} \exp \left(e^{\text{polyloglog}(d)} \Omega(\eta C) \sum_{\tau=\mathcal{T}_0}^{t-1} \left(\frac{\Omega(1)}{C^{2p} \lambda_0^p} - o(\eta) \right) \right) \\ &\geq \frac{\Omega(C^2)}{D} \exp \left(e^{\text{polyloglog}(d)} \Omega(\eta C) \sum_{\tau=\mathcal{T}_0}^{t-1} \left(\frac{\Omega(1)}{C^{2p} \lambda_0^p} - o(\eta) \right) \right). \end{aligned} \quad (25)$$

[\(25\)](#) shows that $\Gamma^{(t)}$ is an non-decreasing sequence. We thus deduce the time \mathcal{T}_1 such that $C\Gamma^{(t)} \geq \Omega(\lambda_0)/D$.

$$\begin{aligned} \frac{C^2}{D} \exp \left(e^{\text{polyloglog}(d)} \Omega(\eta C) \left(\frac{\Omega(1)}{C^{2p} \lambda_0^p} - o(\eta) \right) (\mathcal{T}_1 - \mathcal{T}_0) \right) &= \frac{\lambda_0}{D} \\ \implies \mathcal{T}_1 = \mathcal{T}_0 + \Theta \left(\frac{C^{2p} \lambda_0^p}{\eta e^{\text{polyloglog}(d)}} \log \left(\frac{\lambda_0}{C^2} \right) \right). \end{aligned} \quad (26)$$

We now prove the second part of the lemma. We respectively apply [Lemma D.4](#) and [Lemma D.5](#) to bound $\mathcal{S}^{(t)}$ and $\max_{\tau \leq T} |\mathcal{N}^{(\tau)}|$.

$$\mathcal{S}^{(t)} \geq \Omega(\lambda_0)^p \quad \text{and} \quad \max_{\tau \leq [T]} |\mathcal{N}^{(\tau)}| \leq o(\lambda_0)^p. \quad (27)$$

[\(27\)](#) implies for all $t \in [\mathcal{T}_1, T]$, $\mathcal{S}^{(t)} \geq \max_{\tau \leq [T]} |\mathcal{N}^{(\tau)}|$. □

D.3.1 Auxiliary lemmas

In this section, we present the auxiliary lemmas needed to prove the main results in [subsection D.3](#).

Lemma D.9. *Let $t \geq \mathcal{T}_0$. Then, we always have $\alpha^{(t)} \geq \frac{\Omega(1)}{C^2 \lambda_0} - o(\eta)$.*

Proof of [Lemma D.9](#). For $t \in [0, \mathcal{T}_0]$, $\alpha^{(t)}$ increases and eventually satisfies $\alpha^{(t)} \geq \frac{\Omega(1)}{C^2 \lambda_0}$ ([Lemma D.2](#)). However, for $t \geq \mathcal{T}_0$, $\alpha^{(t)}$ may be non-increasing. Here, we want to quantify the maximum amount of decrease for $t > \mathcal{T}_0$. The worst-case scenario is when $\alpha^{(t)} = \frac{\Omega(1)}{C^2 \lambda_0}$. We bound $\alpha^{(t+1)}$ by using [Lemma D.4](#), [Lemma D.10](#) and [Lemma D.5](#).

$$\alpha^{(t+1)} \geq \alpha^{(t)} + \Theta(C\eta)(\alpha^{(t)} G^{(t)})^p - \eta(\alpha^{(t)})^p o(\lambda_0)^p. \quad (28)$$

We now apply [Induction Hypothesis B.1](#) in [\(28\)](#) and get:

$$\begin{aligned} \alpha^{(t+1)} &\geq \frac{\Omega(1)}{C^2 \lambda_0} + \Theta(\eta) \frac{e^{\text{polyloglog}(d)}}{C^{2p-1} \lambda_0^p} - \frac{o(\eta)}{C^{2p}} \\ &\geq \frac{\Omega(1)}{C^2 \lambda_0} - \frac{o(\eta)}{C^{2p}}. \end{aligned} \quad (29)$$

At time $t+1$, we potentially have $\alpha^{(t+1)} < \frac{\Omega(1)}{C^2 \lambda_0}$. In this case, $\alpha^{(t+1)}$ starts to increase again because it is in the range of α 's that satisfies Event I (and therefore the update rule in [Lemma D.2](#) holds). Thus, for all $t \geq \mathcal{T}_0$, we have $\alpha^{(t)} \geq \frac{\Omega(1)}{C^2 \lambda_0} - o(\eta)$. □

Lemma D.10. *When the population risk is $\Omega(1)$, we have $\mathbb{E}[\mathfrak{S}(-yF(\mathbf{X}))] \geq \Omega(1)$.*

Proof of [Lemma D.10](#). Let (\mathbf{X}, y) be a data-point. We distinguish two cases:

- $yF(\mathbf{X}) > 0$: we apply [Lemma K.6](#) which implies $\mathfrak{S}(-yF(\mathbf{X})) \geq \log(1 + e^{yF(\mathbf{X})})$. Since the population loss is $\Omega(1)$, this implies the aimed result.
- $yF(\mathbf{X}) \leq 0$: we have necessarily $\mathfrak{S}(-yF(\mathbf{X})) \geq \Omega(1)$ since the sigmoid function is large for non-positive values.

Therefore, we have $\mathbb{E}[\mathfrak{S}(-yF(\mathbf{X}))] \geq \Omega(1)$. \square

D.4 Event III: $\alpha^{(t)}$ keeps increases again

For $t \in [\mathcal{T}_0, \mathcal{T}_1]$, $\Gamma^{(t)}$ increases until reaching $C\Gamma^{(t)} \geq \Omega(\lambda_0)/D$. In this section, we show that this implies that $\alpha^{(t)}$ increases again.

Lemma D.11. *Let $\mathcal{T}_1 = \mathcal{T}_0 + \Theta\left(\frac{\lambda_0^p}{\eta e^{\text{polyloglog}(d)}}\right)$ and $t \in [\mathcal{T}_1, T]$. Since $\mathcal{S}^{(t)} \geq \max_{\tau \leq T} |\mathfrak{N}^{(\tau)}|$, $\alpha^{(t)}$ updates as*

$$\alpha^{(t+1)} = \alpha^{(t)} + \Theta(\eta C)(G^{(t)})^p (\alpha^{(t)})^{p-1}.$$

Consequently, $\alpha^{(t)}$ is non-decreasing until the population risk satisfies $\mathcal{L}(\mathbf{A}^{(t)}, \mathbf{v}^{(t)}) \leq o(1)$. Eventually, $\alpha^{(T)}$ is as large as $\alpha^{(T)} = \frac{\text{polylog}(d)}{C^2 \lambda_0}$.

Proof of [Lemma D.11](#). Since $\mathcal{S}^{(t)} \geq \max_{\tau \leq T} |\mathfrak{N}^{(\tau)}|$ ([Lemma D.8](#)), the update of $\alpha^{(t)}$ is:

$$\alpha^{(t+1)} - \alpha^{(t)} = \Theta(\eta C)(\alpha^{(t)})^{p-1} (G^{(t)})^p \mathbb{E}[\mathfrak{S}(-yF(\mathbf{X}))], \quad (30)$$

Since the population loss is at least $\Omega(1)$ for $t \in [\mathcal{T}_1, T]$, [Lemma D.10](#) implies that $\mathbb{E}[\mathfrak{S}(-yF(\mathbf{X}))] \geq \Omega(1)$. Besides, we apply [Induction Hypothesis B.1](#) and $\Omega(C) \geq 0$ and $e^{\text{polyloglog}(d)} \leq \lambda_0$ in (30) and obtain:

$$\begin{cases} \alpha^{(t+1)} \geq \alpha^{(t)} + \Theta(\eta C) e^{\text{polyloglog}(d)} (\alpha^{(t)})^{p-1} \\ \alpha^{(t+1)} \leq \alpha^{(t)} + \Theta(\eta C) \lambda_0^p (\alpha^{(t)})^{p-1} \end{cases}. \quad (31)$$

(31) and [Lemma D.12](#) show that $\alpha^{(t)}$ increases until reaching $\alpha^{(t)} \leq \frac{\text{polylog}(d)}{C^2 \lambda_0}$. \square

D.4.1 Auxiliary results

Lemma D.12. *The values of α such that $\mathbb{E}[\mathfrak{S}(-yF(\mathbf{X}))] \geq \Omega(1)$ is*

$$\alpha \leq \frac{\text{polylog}(d)}{C^2 \lambda_0}.$$

Proof of [Lemma D.12](#). We say that the sigmoid term is small for a constant κ that satisfies

$$\sum_{\tau=0}^T \frac{1}{1 + \exp(\kappa)} \leq \text{polylog}(d) \implies \kappa \geq \log(T) \iff \kappa \geq \text{polylog}(d). \quad (32)$$

Intuitively, (32) means that the sum of the sigmoid terms for all time steps is bounded (up to a logarithmic dependence). In our case, by using (32) and [Lemma D.3](#), the sigmoid $\mathfrak{S}(-yF(\mathbf{X}))$ is small when

$$(C^2 \alpha \lambda_0)^p \geq \text{polylog}(d) \implies \alpha \geq \frac{\text{polylog}(d)}{C^2 \lambda_0}. \quad (33)$$

\square

D.5 Convergence rate of the population loss

Lemma D.13. *Let $t \in [\mathcal{T}_1, T]$. Then, the population loss linearly converges to zero i.e.*

$$\mathcal{L}(\mathbf{A}^{(t)}, \mathbf{v}^{(t)}) \leq \frac{\text{polylog}(d)}{\eta \lambda_0^{2p} (t - \mathcal{T}_1 + 1)} \quad (34)$$

Proof of Lemma D.13. To ease the explanation in this proof, we use the $\tilde{\Omega}, \tilde{\Theta}, \tilde{O}$ notations to hide the logarithmic dependence. We hide for instance the constant C under this notation. From Lemma D.11, we know that $\alpha^{(t)}$ is lower bounded as:

$$\alpha^{(t+1)} \geq \alpha^{(t)} + \Theta(\eta C) (\alpha^{(t)})^{p-1} (G^{(t)})^p \mathbb{E}[\mathfrak{S}(-yF(\mathbf{X}))]. \quad (35)$$

Using Lemma D.3, we have $\mathbb{E}[\mathfrak{S}(-yF(\mathbf{X}))] \geq \mathfrak{S}(-(\alpha^{(t)} G^{(t)})^p)$. Plugging this in (35) yields:

$$\alpha^{(t+1)} \geq \alpha^{(t)} + \Theta(\eta) G^{(t)} \frac{(\alpha^{(t)})^{p-1} (G^{(t)})^{p-1}}{1 + \exp((\alpha^{(t)} G^{(t)})^p)}. \quad (36)$$

Since $0 < \alpha^{(t)} G^{(t)} \leq \tilde{O}(\lambda_0^{p-1})$, we apply Lemma K.5 and get:

$$\alpha^{(t+1)} \geq \alpha^{(t)} + \frac{\tilde{\Omega}(\eta)}{\lambda_0^p} G^{(t)} \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}). \quad (37)$$

Lemma D.8 implies that $G^{(t)} \geq \tilde{\Omega}(\lambda_0)$. Therefore, we have:

$$\alpha^{(t+1)} \geq \alpha^{(t)} + \frac{\tilde{\Omega}(\eta)}{\lambda_0^p} \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}). \quad (38)$$

Let's now assume by contradiction that for $t \in [\mathcal{T}_1, T]$, we have:

$$\log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) > \frac{\tilde{\Omega}(1)}{\eta \lambda_0^{2p} (t - \mathcal{T}_1 + 1)}. \quad (39)$$

For $t \in [\mathcal{T}_1, T]$, we know that $\alpha^{(t)} G^{(t)}$ is non-decreasing which implies that $(\alpha^{(t)} G^{(t)})^p$ is also non-decreasing. Since $x \mapsto \log(1 + \exp(-x))$ is non-increasing, this implies for $s \leq t$ that

$$\frac{\tilde{\Omega}(1)}{\eta \lambda_0^{2p} (t - \mathcal{T}_1)} < \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) \leq \log(1 + e^{-(\alpha^{(s)} G^{(s)})^p}). \quad (40)$$

Plugging (40) in the update (36) yields for $s \in [\mathcal{T}_1, t]$:

$$\alpha^{(s+1)} > \alpha^{(s)} + \frac{\tilde{\Omega}(1)}{\lambda_0^p (t - \mathcal{T}_1 + 1)}. \quad (41)$$

Let $t \in [\mathcal{T}_1, T]$. We now sum (41) for $s = \mathcal{T}_1, \dots, t$ and obtain:

$$\alpha^{(t+1)} > \alpha^{(\mathcal{T}_1)} + \frac{\tilde{\Omega}(1)(t - \mathcal{T}_1 + 1)}{\lambda_0^p (t - \mathcal{T}_1 + 1)} > \frac{\tilde{\Omega}(1)}{\lambda_0^p}, \quad (42)$$

where we used the fact that $\alpha^{(\mathcal{T}_1)} \geq \alpha^{(\mathcal{T}_0)} \geq \tilde{\Omega}(1)/(C\lambda_0) > 0$ (Lemma D.2) in the last inequality. Therefore, we have for $t \in [\mathcal{T}_1, T]$, $\alpha^{(t)} \geq \tilde{\Omega}(1/\lambda_0^p) > 0$. Let's now show that (42) implies a contradiction. Indeed, we have:

$$\begin{aligned} \eta \lambda_0^{2p} (t - \mathcal{T}_1 + 1) \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) &\leq \eta \lambda_0^{2p} T \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) \\ &\leq \eta \lambda_0^{2p} T \log(1 + e^{-\tilde{\Omega}(1)}), \end{aligned} \quad (43)$$

where we used $G^{(t)} \geq \Omega(\lambda_0)$ (Lemma D.8) and (42) in the last inequality. We now apply Lemma K.6 and obtain:

$$\eta \lambda_0^{2p} (t - \mathcal{T}_1 + 1) \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) \leq \frac{\eta \lambda_0^{2p} T}{1 + \exp(\tilde{\Omega}(1))}. \quad (44)$$

Given the values of T, η, λ_0 , we finally have:

$$\eta \lambda_0^{2p} (t - \mathcal{T}_1 + 1) \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) < \tilde{O}(1), \quad (45)$$

which contradicts (39). Therefore, we obtain the convergence rate:

$$\log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) \leq \frac{\tilde{O}(1)}{\eta \lambda_0^{2p} (t - \mathcal{T}_1 + 1)} \quad (46)$$

We apply Lemma D.14 to bound the left-hand side of (46) and get the aimed result. \square

D.5.1 Auxiliary lemmas

Lemma D.14. *Let $t \in [\mathcal{T}_1, T]$. We have:*

$$\Theta(1) \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}) \geq \mathcal{L}(\mathbf{A}^{(t)}, \mathbf{v}^{(t)}).$$

Proof of Lemma D.14. The proof is similar to the one of Lemma D.3. We apply Lemma D.4, Lemma D.8 and Lemma D.5 and get:

$$y F_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X}) \geq (\alpha^{(t)})^p \left(C(G^{(t)})^p - o(\lambda_0)^p \right) \geq \Theta(1)(\alpha^{(t)} G^{(t)})^p. \quad (47)$$

Using (47), we deduce:

$$\mathbb{E}[\log(1 + e^{-y F_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X})})] \leq \log(1 + e^{-\Theta(1)(\alpha^{(t)} G^{(t)})^p}) \leq \Theta(1) \log(1 + e^{-(\alpha^{(t)} G^{(t)})^p}), \quad (48)$$

where we applied Lemma K.7 in the last inequality. \square

D.6 Fitting the labeling function

We now show that the learner model fits the labeling function.

Lemma D.15. *After T iterations, the population risk converges i.e. $\mathcal{L}(\mathbf{A}^{(T)}, \mathbf{v}^{(T)}) \leq O(1/\text{poly}(d))$. Therefore, $\mathbb{P}_{\mathcal{D}}[f^*(\mathbf{X}) F_{\mathbf{A}^{(T)}, \mathbf{v}^{(T)}}(\mathbf{X}) > 0] \geq 1 - o(1)$.*

Proof of Lemma D.15. Since the logistic loss is a surrogate for the 0-1 loss, we have:

$$\mathbb{E}_{\mathcal{D}}[\mathbf{1}_{y F_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X}) < 0}] \leq \mathbb{E}_{\mathcal{D}}[\log(1 + e^{-y F_{\mathbf{A}^{(t)}, \mathbf{v}^{(t)}}(\mathbf{X})})]. \quad (49)$$

We now apply Lemma D.13 to bound the right-hand side of (49). Given the value of T , we have:

$$\mathbb{P}_{\mathcal{D}}[y F_{\mathbf{A}^{(T)}, \mathbf{v}^{(T)}}(\mathbf{X})] = \mathbb{E}_{\mathcal{D}}[\mathbf{1}_{y F_{\mathbf{A}^{(T)}, \mathbf{v}^{(T)}}(\mathbf{X}) < 0}] \leq \frac{\text{polylog}(d)}{\eta \lambda_0^{2p} T} \leq \frac{1}{\text{poly}(d)}. \quad (50)$$

We now use (50) and Definition 2.1 to show that the learner model fits the labeling function. Indeed, we rewrite $\mathbb{P}_{\mathcal{D}}[f^*(\mathbf{X}) F_{\mathbf{A}^{(T)}, \mathbf{v}^{(T)}}(\mathbf{X})]$ as

$$\begin{aligned} \mathbb{P}_{\mathcal{D}}[f^*(\mathbf{X}) F_{\mathbf{A}^{(T)}, \mathbf{v}^{(T)}}(\mathbf{X})] &\geq \mathbb{P}_{\mathcal{D}}[y f^*(\mathbf{X}) > 0] \mathbb{P}_{\mathcal{D}}[y F_{\mathbf{A}^{(T)}, \mathbf{v}^{(T)}}(\mathbf{X}) > 0] \\ &\geq (1 - d^{-\omega(1)}) \left(1 - \frac{1}{\text{poly}(d)} \right) \\ &= 1 - o(1). \end{aligned} \quad (51)$$

\square

D.7 Proof of the induction hypothesis

In this section, we prove Induction Hypothesis B.1.

Proof of Induction Hypothesis B.1. We start by proving that $|\rho^{(t)}| = \Theta(1)$ for all $t \in [T]$.

Let $t \in [\mathcal{T}_1, T]$ and $\tau \in [t]$. Using Corollary G.2 and Induction Hypothesis B.1, we upper bound $|\rho^{(\tau)}|$ as:

$$|\rho^{(\tau+1)}| \leq |\rho^{(\tau)}| + \frac{\eta \Theta(C)}{D} (\alpha^{(\tau)})^p \left(1 + \frac{\lambda_0^{p+1}}{D} \right). \quad (52)$$

Summing (52) for $\tau = 0, \dots, t-1$ and using $\rho^{(0)} = 0$ lead to

$$|\rho^{(t)}| \leq \frac{\eta \Theta(C)}{D} \left(1 + \frac{\lambda_0^{p+1}}{D} \right) \sum_{\tau=0}^{t-1} (\alpha^{(\tau)})^p \quad (53)$$

We now apply Lemma D.16 to bound the sum of $\alpha^{(t)}$'s in (53).

$$|\rho^{(t)}| \leq \Theta \left(\frac{1}{D(\alpha^{(0)})^{p-1} (\lambda_0)^p e^{\text{polyloglog}(d)}} + \frac{1}{e^{\text{polyloglog}(d)}} \right). \quad (54)$$

Given the values of the different parameters, (54) implies that $|\rho^{(t)}| \leq \Theta(1)$.

We now prove $e^{\gamma^{(t)}} \in [\Omega(1), \lambda_0]$. Since $e^{\gamma^{(t)}}$ is non-decreasing (Corollary G.1), we have $e^{\gamma^{(t)}} \geq e^{\gamma^{(0)}} \geq \Omega(1)$ for all $t \geq 0$. We now prove the upper bound on $e^{\gamma^{(t)}}$. We assume that for all $\tau \leq t$, $e^{\gamma^{(\tau)}} \leq \lambda_0$. Let's show this inequality for $t + 1$. Using Corollary G.1, we have:

$$\begin{aligned} e^{\gamma^{(t+1)}} &\leq e^{\gamma^{(t)}} \exp\left(\Theta(C\eta)(\alpha^{(t)})^p \Gamma^{(t)}(G^{(t)})^{p-1}\right) \\ &= \exp\left(\Theta(C\eta) \sum_{\tau=0}^t (\alpha^{(\tau)})^p \Gamma^{(\tau)}(G^{(\tau)})^{p-1}\right). \end{aligned} \quad (55)$$

We now apply the induction hypothesis in (55) and get:

$$\begin{aligned} e^{\gamma^{(t+1)}} &\leq e^{\gamma^{(t)}} \exp\left(\Theta(C\eta)(\alpha^{(t)})^p \Gamma^{(t)}(G^{(t)})^{p-1}\right) \\ &= \exp\left(\frac{\Theta(\lambda_0^p C\eta)}{D^p} \sum_{\tau=0}^t (\alpha^{(\tau)})^p\right) \\ &\leq \exp\left(\frac{\Theta(\lambda_0^p C\eta)}{D^p} \sum_{\tau=0}^T (\alpha^{(\tau)})^p\right). \end{aligned} \quad (56)$$

We apply Lemma D.16 in (56) and obtain:

$$\begin{aligned} e^{\gamma^{(t+1)}} &\leq \exp\left(\Theta\left(\frac{1}{D^p(\alpha^{(0)})^{p-1}e^{\text{polyloglog}(d)}} + \frac{\lambda_0^p}{D^{p-1}e^{\text{polyloglog}(d)}}\right)\right) \\ &= \Theta(1) \exp\left(\frac{\lambda_0^p}{D^{p-1}e^{\text{polyloglog}(d)}}\right) \\ &\leq \Theta\left(1 + \frac{\lambda_0^p}{D^{p-1}e^{\text{polyloglog}(d)}} + \frac{\lambda_0^{2p}}{D^{2(p-1)}e^{\text{polyloglog}(d)}}\right), \end{aligned} \quad (57)$$

where we used the inequality $e^x \leq 1 + x + x^2$ for $x \leq 1$ in (57). Given the values of the different parameters, we deduce that $e^{\gamma^{(t+1)}} \leq \lambda_0$.

We now prove $e^{\beta^{(t)}} = e^{\text{polyloglog}(d)}$ for $t \in [0, T]$. Since $\beta^{(t)}$ is not updated i.e. $\beta^{(t)} = \beta^{(0)}$ and $\beta^{(0)} = \sigma_M = \text{polyloglog}(d)$, we therefore have the aimed result.

Lastly, we prove that $e^{\beta^{(t)}} + (C-1)e^{\gamma^{(t)}} + (D-C)e^{\rho^{(t)}} = \Theta(D)$ for $t \in [T]$. Since $e^{\gamma^{(t)}} \geq \Theta(1)$, $e^{\rho^{(t)}} = \Theta(1)$ and $e^{\beta^{(t)}} \geq \Theta(1)$, we have:

$$e^{\beta^{(t)}} + (C-1)e^{\gamma^{(t)}} + (D-C)e^{\rho^{(t)}} \geq \Theta(D). \quad (58)$$

On the other hand, we have $e^{\gamma^{(t)}} \leq \lambda_0$, $e^{\beta^{(t)}} = e^{\text{polyloglog}(d)}$ and $e^{\rho^{(t)}} = \Theta(1)$ which imply:

$$\begin{aligned} e^{\beta^{(t)}} + (C-1)e^{\gamma^{(t)}} + (D-C)e^{\rho^{(t)}} &\leq e^{\text{polyloglog}(d)} + (C-1)\lambda_0 + (D-C)\Theta(1) \\ &\leq C\lambda_0 + (D-C)\Theta(1) \\ &\leq \Theta(D). \end{aligned} \quad (59)$$

□

D.7.1 Auxiliary lemmas

Lemma D.16. *The sum of the $\alpha^{(t)}$'s is bounded as:*

$$\sum_{\tau=0}^T (\alpha^{(\tau)})^p = \Theta\left(\frac{1}{\eta C(\alpha^{(0)})^{p-1}(\lambda_0)^p e^{\text{polyloglog}(d)}} + \frac{D}{C\eta e^{\text{polyloglog}(d)}}\right).$$

Proof of Lemma D.16. We first decompose the sum of $\alpha^{(t)}$'s.

$$\sum_{\tau=0}^T (\alpha^{(\tau)})^p = \sum_{\tau=0}^{\mathcal{T}_0-1} (\alpha^{(\tau)})^p + \sum_{\tau=\mathcal{T}_0}^{\mathcal{T}_1-1} (\alpha^{(\tau)})^p + \sum_{\tau=\mathcal{T}_1}^T (\alpha^{(\tau)})^p. \quad (60)$$

We apply Lemma D.2 and Lemma D.8 to rewrite (60).

$$\sum_{\tau=0}^T (\alpha^{(\tau)})^p = \frac{\Theta(1)}{\eta C(\alpha^{(0)})^{p-1} (\lambda_0)^p e^{\text{polyloglog}(d)}} + \frac{\log(\lambda_0)}{\eta e^{\text{polyloglog}(d)}} + \sum_{\tau=\mathcal{T}_1}^T (\alpha^{(\tau)})^p. \quad (61)$$

Now, we aim to obtain the value of the last summand in (61). Using Corollary G.1, we have

$$C\Gamma^{(T)} \geq C\Gamma^{(\mathcal{T}_1)} \exp\left(\frac{C\eta}{D} e^{\text{polyloglog}(d)} \sum_{\tau=\mathcal{T}_1}^T (\alpha^{(\tau)})^p\right). \quad (62)$$

We finally apply Lemma D.8 and Induction Hypothesis B.1 in (62) to get:

$$\sum_{\tau=\mathcal{T}_1}^T (\alpha^{(\tau)})^p \leq \frac{\Theta(D)}{C\eta e^{\text{polyloglog}(d)}}. \quad (63)$$

To obtain the aimed result, we plug (63) in (61) and use $\frac{\log(\lambda_0)}{\eta e^{\text{polyloglog}(d)}} \leq \frac{\Theta(D)}{C\eta e^{\text{polyloglog}(d)}}$. \square

E From idealized to real learning process

In Appendix D, we analyzed the ideal learning process. We now aim to bridge the gap between the idealized and realistic cases. Given our initialization, $\hat{\mathbf{v}}^{(t)}$ has a component in $\text{span}(\mathbf{w}^*)^\perp$ i.e.

$$\hat{\mathbf{v}}^{(t)} = \hat{\alpha}^{(t)} \mathbf{w}^* + \varepsilon_{\mathbf{v}}^{(t)} \mathbf{u}^{(t)}$$

where $\mathbf{u}^{(t)} \in \mathbb{R}^d$ such that $\mathbf{u}^{(t)} \perp \mathbf{w}^*$ and $\|\mathbf{u}^{(t)}\|_2 = 1$. Thus, one main difference between the two cases is that we initialize $\alpha^{(0)} \geq \nu^{1/(p-1)}$ in the idealized case while $\hat{\alpha}^{(0)} \leq \omega$. Thus, the proof strategy consists in i) $t \in [0, \mathcal{T}]$, $\hat{\alpha}^{(t)}$ increases until having $\hat{\alpha}^{(t)} \geq \nu^{1/(p-1)}$ (subsubsection E.1.1) while $\varepsilon_{\mathbf{v}}^{(t)}$ (subsubsection E.1.2) and $\hat{A}_{i,j}^{(t)}$ (subsubsection E.1.3) stay tiny. ii) $t \in [\mathcal{T}, T]$, compare the realistic and idealized iterates. We remind the GD update of $\hat{\alpha}^{(t)}$.

$$\begin{aligned} & \hat{\alpha}^{(t+1)} - \hat{\alpha}^{(t)} \\ &= \eta \underbrace{\frac{D}{N} \sum_{i=1}^N y[i] \mathfrak{S}(-y[i]F(\mathbf{X}[i])) \sum_{j \in \mathcal{S}_\ell(\mathbf{X}[i])} \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_k[i] \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{w}^*, \mathbf{X}_r[i] \rangle}_{\hat{\mathcal{S}}^{(t)}}}_{\hat{\mathcal{N}}^{(t)}} \\ &+ \eta \underbrace{\frac{D}{N} \sum_{i=1}^N y[i] \mathfrak{S}(-y[i]F(\mathbf{X}[i])) \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_k[i] \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{w}^*, \mathbf{X}_r[i] \rangle}_{\hat{\mathcal{N}}^{(t)}}}. \end{aligned} \quad (\text{GD-}\hat{\alpha})$$

E.1 Bound on the iterates during the initial steps ($t \in [0, \mathcal{T}]$)

Since we randomly initialize $\hat{\mathbf{v}}^{(0)}$ with tiny variance, we need to take into account the linear part of the activation function. Lemma E.8 shows that we can overlook the power part of the activation and consider $\sigma(x) = \nu x$ as long as $\hat{\alpha}^{(t)} \geq \nu^{1/(p-1)}$.

E.1.1 $\hat{\alpha}^{(t)}$ initially increases

Lemma E.1. Let $\mathcal{T} = \Theta\left(\frac{1}{\eta \nu^{(p-2)/(p-1)} e^{\text{polyloglog}(d)}}\right)$. For all $t \in [0, \mathcal{T}]$, $\hat{\alpha}^{(t)}$ is updated as

$$\hat{\alpha}^{(t+1)} = \hat{\alpha}^{(t)} + \Theta(\eta \nu) e^\beta.$$

Consequently, $\hat{\alpha}^{(t)}$ is non-decreasing and after \mathcal{T} iterations, we have $\hat{\alpha}^{(t)} \geq \nu^{1/(p-1)}$ for $t \geq \mathcal{T}$.

Proof of Lemma E.1. Let $t \geq 0$. We apply Lemma E.5 and Lemma E.6 to respectively bound $\widehat{\mathcal{S}}^{(t)}$ and $\widehat{\mathcal{N}}^{(t)}$ in the update of $\widehat{\alpha}^{(t)}$.

$$\widehat{\alpha}^{(t+1)} = \widehat{\alpha}^{(t)} + \Theta(C\eta\nu)e^\beta. \quad (64)$$

(64) indicates that $\widehat{\alpha}^{(t)}$ is a non-decreasing sequence. Therefore, there exists a time \mathcal{T} such that $\widehat{\alpha}^{(\mathcal{T})} = \nu^{1/(p-1)}$. Summing (64) for $t = 0, \dots, \mathcal{T} - 1$ yields $\mathcal{T} = \Theta\left(\frac{1}{\eta\nu^{(p-2)/(p-1)}e^\beta}\right)$. \square

E.1.2 Bound on ε_v

We now show that for $t \in [0, \mathcal{T}]$, the orthogonal component $\varepsilon_v^{(t)}$ stays small.

Lemma E.2. Assume that we run GD on the empirical risk (E) for T iterations with parameters set as in Parametrization 3.1. For $t \in [0, \mathcal{T}]$, the orthogonal component ε_v satisfies

$$\varepsilon_v^{(t+1)} \leq \left(1 + \eta\nu \frac{\text{poly}(D)}{\sqrt{d}}\right) \varepsilon_v^{(t)} + \eta\zeta,$$

where $\zeta = \frac{\text{poly}(D)}{\sqrt{N}}$.

Proof of Lemma E.2. Let $\mathbf{P} = (\mathbf{I} - \mathbf{w}^* \mathbf{w}^{*\top})$ and $\hat{v}^{(t)} = \widehat{\alpha}^{(t)} \mathbf{w}^*$. The projected update of \hat{v} satisfies:

$$\begin{aligned} & \|\mathbf{P}\hat{v}^{(t+1)} - \mathbf{P}\hat{v}^{(t)}\|_2 \\ & \leq D\eta\nu \left\| \frac{1}{N} \sum_{i=1}^N y[i] \mathfrak{S}(-y[i]F_{\hat{v}}(\mathbf{X}[i])) \sum_{m=1}^D \sum_{b=1}^D \widehat{\mathcal{S}}_{m,b}^{(t)} \mathbf{P}\mathbf{X}_b[i] \right\|_2 \end{aligned} \quad (65)$$

$$- \mathbb{E} \left[y \mathfrak{S}(-yF_{\hat{v}}(\mathbf{X})) \sum_{m=1}^D \sum_{b=1}^D \widehat{\mathcal{S}}_{m,b}^{(t)} \mathbf{P}\mathbf{X}_b \right] \Big\|_2 \quad (66)$$

$$+ D\eta\nu \left\| \mathbb{E} \left[y \mathfrak{S}(-yF_{\hat{v}}(\mathbf{X})) \sum_{m=1}^D \sum_{b=1}^D \widehat{\mathcal{S}}_{m,b}^{(t)} \mathbf{P}\mathbf{X}_b \right] \right\|_2 \quad (67)$$

$$- \mathbb{E} \left[y \mathfrak{S}(-yF_{\hat{v}}(\mathbf{X})) \sum_{m=1}^D \sum_{b=1}^D \widehat{\mathcal{S}}_{m,b}^{(t)} \mathbf{P}\mathbf{X}_b \right] \Big\|_2. \quad (68)$$

Summand 1: $\|(65) - (66)\|_2$. Using the matrix Hoeffding inequality, we have with high probability,

$\|(81) - (82)\|_2 \leq 16\sqrt{\log(d) \sum_{i=1}^N M_i^2}$, where

$\left\| \frac{\eta\nu D y[i] \mathfrak{S}(-y[i]F_{\hat{v}}(\mathbf{X}^{(i)})) \sum_{m=1}^D \sum_{b=1}^D \widehat{\mathcal{S}}_{m,b}^{(t)} \mathbf{P}\mathbf{X}_b^{(i)} \right\|_2^2 \leq M_i^2$. Induction Hypothesis B.1 and $\|\mathbf{P}\mathbf{X}_u^{(t)}\|_2 \leq \sigma^2 d \log(d)$ imply $M_i^2 \leq \frac{\eta^2 \nu^2 \text{poly}(D)}{N^2}$. We deduce that $\|(65) - (66)\|_2 \leq \eta\nu \frac{\text{poly}(D)}{\sqrt{N}}$.

Summand 2: $\|(67) - (68)\|_2$. We use the 1-Lipschitzness of the sigmoid function and get:

$$\begin{aligned} & \|(67) - (68)\|_2 \\ & \leq \eta\nu \left\| \mathbb{E} \left[Dy \sum_{m=1}^D \sum_{b=1}^D \widehat{\mathcal{S}}_{m,b}^{(t)} \mathbf{P}\mathbf{X}_u \right. \right. \\ & \quad \left. \left[\mathfrak{S}\left(-y\nu D \sum_{a=1}^D \sum_{a'=1}^D \widehat{\mathcal{S}}_{a,a'}^{(t)} \langle \hat{v}^{(t)}, \mathbf{X}_{a'} \rangle\right) - \mathfrak{S}\left(-y\nu D \sum_{a=1}^D \sum_{a'=1}^D \widehat{\mathcal{S}}_{a,a'}^{(t)} \langle \hat{v}^{(t)}, \mathbf{X}_{a'} \rangle\right) \right] \right\|_2 \\ & \leq \Theta(\eta\nu D) \mathbb{E} \left[\sum_{m=1}^D \sum_{r=1}^D \widehat{\mathcal{S}}_{m,r}^{(t)} \sum_{a=1}^D \left| D \sum_{a'=1}^D \widehat{\mathcal{S}}_{a,a'}^{(t)} \langle \hat{v}^{(t)} - \hat{v}^{(t)}, \mathbf{X}_{a'} \rangle \right| \sum_{b=1}^D \widehat{\mathcal{S}}_{m,u}^{(t)} \|\mathbf{P}\mathbf{X}_b\|_2 \right] \\ & \leq \Theta(\eta\nu D) \mathbb{E} \left[\sum_{m=1}^D \lambda_0 \varepsilon_v^{(t)} \cdot \left| \langle \mathbf{u}^{(t)}, \sum_{a'=1}^D \boldsymbol{\xi}_{a'} \rangle \right| \cdot \sum_{b=1}^D \lambda_0 \|\boldsymbol{\xi}_b\|_2 \right]. \end{aligned} \quad (69)$$

where we applied [Induction Hypothesis B.1](#) in (69). Since with high probability, $\|\xi_b\|_2 \leq \sigma\sqrt{d\log(d)}$, $\langle \mathbf{u}^{(t)}, \sum_{r=1}^D \xi_r \rangle \leq \sqrt{D\log(d)}\sigma$, we finally have:

$$\|(67) - (68)\|_2 \leq \eta\nu \cdot \text{poly}(D)\varepsilon_v^{(t)}\sigma = \eta\nu \frac{\text{poly}(D)}{\sqrt{d}}\varepsilon_v^{(t)}.$$

Combining the bounds on Summands 1 and 2 yields the aimed result. \square

We now use [Lemma E.2](#) to show that ε_v stays small.

Lemma E.3. *For all $t \leq \mathcal{T}$, $\varepsilon_v^{(t)} \leq \omega\sqrt{d\log(d)} + \nu^{1/(p-1)}\frac{\text{poly}(D)}{\sqrt{N}}$. By setting $N = \text{poly}(d)$, we have: $\varepsilon_v^{(t)} \leq 1/\text{poly}(d)$.*

Proof of Lemma E.3. Unraveling [Lemma E.2](#) for $t = 0, \dots, \mathcal{T}$ and using $\varepsilon_v^{(0)} \leq \omega\sqrt{d\log(d)}$ (with high probability) leads to:

$$\varepsilon_v^{(\mathcal{T})} \leq \omega\sqrt{d\log(d)} + \eta\nu\zeta \frac{\left(1 + \eta\nu\frac{\text{poly}(D)}{\sqrt{d}}\right)^{\mathcal{T}} - 1}{\eta\nu\frac{\text{poly}(D)}{\sqrt{d}}} \leq \omega\sqrt{d\log(d)} + 2\mathcal{T}\eta\nu\zeta, \quad (70)$$

where we used $(1+x)^y \leq 1+2yx$ for $x \ll 1$ and $y \geq 0$. Plugging the value of \mathcal{T} in (70) yields the aimed result. \square

E.1.3 $\hat{A}_{a,b}^{(t)}$ stays small

We finally show that $\hat{A}_{a,b}^{(t)}$ remains tiny for $t \in [0, \mathcal{T}]$.

Lemma E.4. *Let $a, b \in [D]$. We have $|\hat{A}_{a,b}^{(t)}| \leq \omega\sqrt{d\log(d)} + \frac{\Theta(\nu^{2/(p-1)})}{De^\beta}$.*

Proof of Lemma E.4. We remind that the GD update of $\hat{A}_{a,b}^{(t)}$ is

$$\hat{A}_{a,b}^{(t+1)} = \hat{A}_{a,b}^{(t)} + \frac{\nu\eta}{N} \sum_{i=1}^N \mathfrak{S}(-y[i]F(\mathbf{X}[i]))\hat{S}_{a,b}^{(t)} \sum_{m \neq b} \hat{S}_{a,m}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_m[i] - \mathbf{X}_b[i] \rangle. \quad (71)$$

The proof is by induction. We assume that $\hat{A}_{a,b}^{(t)} \leq \omega\sqrt{d\log(d)} + \frac{\Theta(\nu^{2/(p-1)})}{De^\beta}$. We first apply Cauchy-Schwarz on (71) and get:

$$|\hat{A}_{a,b}^{(t+1)}| \leq |\hat{A}_{a,b}^{(t)}| + \nu\eta\hat{S}_{a,b}^{(t)} \sum_{m \neq b} \hat{S}_{a,m}^{(t)} \|\hat{\mathbf{v}}^{(t)}\|_2. \quad (72)$$

Using the induction hypothesis, we have $\hat{S}_{a,b}^{(t)} \leq \Theta(1)/D$. Thus, we have

$$|\hat{A}_{a,b}^{(t+1)}| \leq |\hat{A}_{a,b}^{(t)}| + \frac{\Theta(\nu\eta)}{D} \|\hat{\mathbf{v}}^{(t)}\|_2 \leq |\hat{A}_{a,b}^{(t)}| + \frac{\Theta(\nu\eta)}{D} \sqrt{(\hat{\alpha}^{(t)})^2 + (\varepsilon_v^{(t)})^2}. \quad (73)$$

We sum (73) and get:

$$|\hat{A}_{a,b}^{(\mathcal{T})}| \leq |\hat{A}_{a,b}^{(0)}| + \frac{\Theta(\nu\eta)}{D} \sum_{t=0}^{\mathcal{T}-1} \sqrt{(\hat{\alpha}^{(t)})^2 + (\varepsilon_v^{(t)})^2}. \quad (74)$$

We now use [Lemma E.1](#) and [Lemma E.3](#) in (74) and get:

$$\begin{aligned} |\hat{A}_{a,b}^{(\mathcal{T})}| &\leq \omega\sqrt{d\log(d)} + \frac{\Theta(\nu^{1/(p-1)})}{De^\beta} \sqrt{\nu^{2/(p-1)} + \left(\omega\sqrt{d\log(d)} + \nu^{1/(p-1)}\frac{\text{poly}(D)}{\sqrt{N}}\right)^2} \\ &\leq \omega\sqrt{d\log(d)} + \frac{\Theta(\nu^{2/(p-1)})}{De^\beta}. \end{aligned} \quad (75)$$

\square

E.1.4 Auxiliary lemmas

Lemma E.5. Let \mathcal{T} be the time where $\hat{\alpha}^{(t)} \geq \nu^{1/(p-1)}$. Let $t \in [0, \mathcal{T}]$, $i \in [N]$ and $j \in \mathcal{S}_{\ell(\mathbf{X}[i])}$. We have:

$$y[i] \langle \mathbf{w}^*, \mathbf{O}_j^{(t)}[i] \rangle = e^\beta.$$

This implies $\hat{\mathcal{S}}^{(t)} = Ce^\beta$ for all $t \in [0, \mathcal{T}]$.

Proof of Lemma E.5. We successively apply Lemma E.4 and Lemma K.3 to get:

$$\begin{aligned} y[i] \langle \mathbf{w}^*, \mathbf{O}_j^{(t)}[i] \rangle &= D \left(\hat{\mathcal{S}}_{j,j}^{(t)} + \sum_{k \in \mathcal{S}_{\ell(\mathbf{X}[i])}} \hat{\mathcal{S}}_{j,k}^{(t)} + y[i] \sum_{h \neq \ell(\mathbf{X}[i])} \sum_{r \in \mathcal{S}_h} \hat{\mathcal{S}}_{j,r}^{(t)} \delta_r \right) \\ &\geq \Theta \left(e^{\text{polyloglog}(d)} + \Theta(C-1) - \Theta(1) \sum_{h \neq \ell(\mathbf{X}[i])} \sum_{r \in \mathcal{S}_h} \delta_{h,r} \right) \end{aligned} \quad (76)$$

$$\begin{aligned} &\geq \Theta \left(e^{\text{polyloglog}(d)} + \Theta(C-1) - \Theta(qD \log(d)) \right) \\ &\geq e^{\text{polyloglog}(d)}. \end{aligned} \quad (77)$$

Similarly, we also have $y[i] \langle \mathbf{w}^*, \mathbf{O}_j^{(t)}[i] \rangle \leq e^{\text{polyloglog}(d)}$. \square

Lemma E.6. Let $t \in [0, \mathcal{T}]$. Assume that $N = \text{poly}(d)$. With high probability, $\hat{\mathcal{N}}^{(t)} \leq 1/\text{poly}(d)$.

Proof of Lemma E.6. During this time phase, the sigmoid stays constant. Therefore, we have $\mathbb{E}[\hat{\mathcal{N}}^{(t)}] = 0$. Therefore, we apply Hoeffding inequality and Lemma D.6 to get:

$$\hat{\mathcal{N}}^{(t)} \leq \sqrt{8N \log(d) \max_{i \in [N]} M_i^2}, \quad (78)$$

where A_i is a constant such that $\left| \frac{1}{N} \sum_{j \in \mathcal{S}_{\ell(\mathbf{X}[i])}} y[i] \langle \mathbf{w}^*, \mathbf{O}_j^{(t)}[i] \rangle \right| \leq M_i \leq \frac{\text{poly}(D)}{N^2}$. Since $N = \text{poly}(d)$, we finally proved $\hat{\mathcal{N}}^{(t)} \leq 1/\text{poly}(d)$. \square

Lemma E.7. Let $t \leq \mathcal{T}$. The sum of $\hat{\alpha}^{(t)}$'s is bounded as:

$$\sum_{\tau=0}^t \hat{\alpha}^{(\tau)} = t\hat{\alpha}^{(0)} + \Theta(\eta\nu)e^\beta t^2. \quad (79)$$

Proof of Lemma E.7. Let $\tau \in [0, \mathcal{T}]$. We sum the update rule of $\hat{\alpha}^{(t)}$ (Lemma E.1) and obtain: $\hat{\alpha}^{(\tau)} = \hat{\alpha}^{(0)} + \Theta(C\eta\nu)e^\beta \tau$. Summing again this update yields the aimed result. \square

Lemma E.8. Let $(\mathbf{X}, \cdot) \sim \mathcal{D}$ and $j \in [D]$. Assume that $\hat{\alpha}^{(t)} \leq \nu^{1/(p-1)}$. Then, we have:

$$\sigma' \left(D \sum_{k=1}^D \hat{\mathcal{S}}_{j,k}^{(t)} \langle \hat{\mathbf{v}}, \mathbf{X}_k \rangle \right) = \begin{cases} \Theta(\nu) & \text{if } \hat{\alpha}^{(t)} \leq \nu^{1/(p-1)} \\ \Theta(p) \left(D \sum_{k=1}^D \hat{\mathcal{S}}_{j,k}^{(t)} \langle \hat{\mathbf{v}}, \mathbf{X}_k \rangle \right)^{p-1} & \text{otherwise} \end{cases}.$$

Proof of Lemma E.8. We remind that the derivative of the activation function $\sigma'(x) = px^{p-1} + \nu$. We first remark that for all x , $\sigma'(x) \geq \nu$. Besides, we have since $p-1$ is even,

$$px^{p-1} + \nu \leq 3\nu \iff |x| \leq \left(\frac{2\nu}{p} \right)^{1/(p-1)}. \quad (80)$$

In our case, we have $x = D\hat{\alpha}^{(t)} \sum_{k=1}^D \hat{\mathcal{S}}_{j,k}^{(t)} \langle \mathbf{w}^*, \mathbf{X}_k \rangle$. Using Induction Hypothesis B.1, we have $|x| \geq e^{\text{polyloglog}(d)} \alpha^{(t)}$. Therefore, a sufficient condition for (80) to hold is $\hat{\alpha}^{(t)} \leq \nu^{1/(p-1)} e^{-\text{polyloglog}(d)}$. Since $e^{\text{polyloglog}(d)} \ll \text{poly}(d)$, we can simplify this condition as $\hat{\alpha}^{(t)} \leq \nu^{1/(p-1)}$. Proving the second part of the lemma can be done as in the proof of Lemma D.7. \square

E.2 Coupling between the semi-idealized and realistic processes ($t \in [\mathcal{T}, T]$)

In this section, we aim to bound the realistic iterates $\hat{A}_{i,j}^{(t)}$ and $\hat{\mathbf{v}}^{(t)}$ for $t \in [\mathcal{T}, T]$. For this reason, we introduce a "semi-idealized" learning process (subsubsection E.2.1) which may be viewed as a mid-point between the idealized and realistic process. We first bound the iterates in this process. Then, using this process, we show that $\varepsilon_{\mathbf{v}}^{(t)}$ (subsubsection E.2.2) and $\Delta_{\mathbf{A}}^{(t)} := \max_{i \neq j} |\hat{A}_{i,j}^{(t)} - \check{A}_{i,j}^{(t)}|$ (subsubsection E.2.4) stay small. Here, $\check{A}_{i,j}^{(t)}$ is the semi-idealized attention matrix coefficient. Finally, since $\varepsilon_{\mathbf{v}}^{(T)}$ and $\Delta_{\mathbf{A}}^{(T)}$ are small, the final iterates $\hat{\alpha}^{(T)}$ and $\alpha^{(T)}$ are equal (subsubsection E.2.6) and thus, the model fits the labeling function (subsubsection E.2.7).

E.2.1 Defining the semi-idealized process

We define an intermediate learning process that we refer to as the "semi-idealized" process. This process starts at time $t = \mathcal{T}$ involves two parameters: the semi-idealized value vector $\check{\mathbf{v}}$ and semi-idealized attention matrix $\check{\mathbf{A}}$ defined as

- the value vector $\check{\mathbf{v}}$ is fixed and satisfies $\check{\mathbf{v}}^{(t-\mathcal{T})} = \hat{\alpha}^{(t)} \mathbf{w}^*$ for $t \in [\mathcal{T}, T]$.
- $\check{A}_{i,j}^{(t-\mathcal{T})}$ is a trainable parameter and is initialized as $\check{A}_{i,j}^{(0)} = 0$ for $i \neq j$.

Therefore, the only trainable parameter in this process is $\check{\mathbf{A}}$. In the semi-idealized process, we minimize the population risk

$$\min_{\check{\mathbf{A}}} \mathbb{E}_{\mathcal{D}} [\log(1 + e^{-yF(\mathbf{X})})] := \check{\mathcal{L}}(\check{\mathbf{v}}, \check{\mathbf{A}}). \quad (\tilde{\text{P}})$$

We remark that such process present similarities to the idealized case. In particular, it satisfies all the invariance and symmetry properties from Lemma 4.1. We thus define

- $\check{A}_{i,j}^{(t)} = \check{\gamma}^{(t)}$ for all $\ell \in [L]$ and $i, j \in \mathcal{S}_{\ell}$.
- $\check{A}_{i,j}^{(t)} = \check{\rho}^{(t)}$ for all $\ell, m \in [L]$ such that $m \neq \ell$ and $i \in \mathcal{S}_{\ell}$ and $j \in \mathcal{S}_m$.

Therefore, $\check{\gamma}^{(t)}$ and $\check{\rho}^{(t)}$ are respectively updated as in Lemma G.1 and Lemma G.2. We define also the softmax terms

$$\begin{aligned} \check{\Lambda}^{(t)} &= \frac{e^{\beta}}{e^{\beta} + (C-1)e^{\check{\gamma}^{(t)}} + (D-C)e^{\check{\rho}^{(t)}}}, & \check{\Gamma}^{(t)} &= \frac{e^{\check{\gamma}^{(t)}}}{e^{\beta} + (C-1)e^{\check{\gamma}^{(t)}} + (D-C)e^{\check{\rho}^{(t)}}}, \\ \check{\Xi}^{(t)} &= \frac{e^{\check{\rho}^{(t)}}}{e^{\beta} + (C-1)e^{\check{\gamma}^{(t)}} + (D-C)e^{\check{\rho}^{(t)}}}, & \check{G}^{(t)} &= D(\check{\Lambda}^{(t)} + (C-1)\check{\Gamma}^{(t)}). \end{aligned}$$

We finally assume Induction Hypothesis B.1 for this process. This latter can be proved using the same arguments as in subsection D.7.

E.2.2 Realistic dynamics are mainly on $\text{span}(\mathbf{w}^*)$

We previously showed in Lemma E.3 that $\varepsilon_{\mathbf{v}}^{(t)}$ is small in the initial steps. We now show that it stays small during the whole process.

Lemma E.9. *Assume that we run GD on the empirical risk (E) for T iterations with parameters set as in Parametrization 3.1 and the number of samples is $N = \text{poly}(d)$. Then, $\hat{\mathbf{v}}^{(t)}$ mainly lies in $\text{span}(\mathbf{w}^*)$ i.e. for $t \leq T$, $\varepsilon_{\mathbf{v}}^{(t)} \leq 1/\text{poly}(d)$.*

We now proceed to the proof of Lemma E.9. We first characterize the recursion satisfied by $\varepsilon_{\mathbf{v}}^{(t)}$.

Lemma E.10. *Assume that we run GD on the empirical risk (E) for T iterations with parameters set as in Parametrization 3.1. Then, $\varepsilon_{\mathbf{v}}^{(t)}$ satisfies for $t \in [\mathcal{T}, T]$*

$$\varepsilon_{\mathbf{v}}^{(t+1)} \leq \left(1 + \eta(\hat{\alpha}^{(t)})^{p-1} \frac{\text{poly}(D)}{\sqrt{d}}\right) \varepsilon_{\mathbf{v}}^{(t)} + \eta\zeta, \quad \text{where } \zeta = \frac{\text{poly}(D)}{\sqrt{N}}.$$

Proof of Lemma E.10. Let $\mathbf{P} := (\mathbf{I} - \mathbf{w}^* \mathbf{w}^{*\top})$, $\hat{\mathbf{v}}^{(t)} := \hat{\alpha}^{(t)} \mathbf{w}^*$ and $t \in [\mathcal{T}, T]$. The projected update of $\hat{\mathbf{v}}$ satisfies:

$$\begin{aligned} & \|\mathbf{P}\hat{\mathbf{v}}^{(t+1)} - \mathbf{P}\hat{\mathbf{v}}^{(t)}\|_2 \\ & \leq \eta \left\| \frac{1}{N} \sum_{i=1}^N Dy[i] \mathfrak{S}(-y[i] F_{\hat{\mathbf{v}}}(\mathbf{X}[i])) \sum_{m=1}^D \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r^{(i)} \rangle \right)^{p-1} \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \mathbf{P} \mathbf{X}_b[i] \right\|_2 \end{aligned} \quad (81)$$

$$- \mathbb{E} \left[Dy \mathfrak{S}(-y F_{\hat{\mathbf{v}}}(\mathbf{X})) \sum_{m=1}^D \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1} \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \mathbf{P} \mathbf{X}_b \right] \Big\|_2 \quad (82)$$

$$+ \eta \left\| \mathbb{E} \left[Dy \mathfrak{S}(-y F_{\hat{\mathbf{v}}}(\mathbf{X})) \sum_{m=1}^D \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1} \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \mathbf{P} \mathbf{X}_b \right] \right\|_2 \quad (83)$$

$$- \mathbb{E} \left[Dy \mathfrak{S}(-y F_{\hat{\mathbf{v}}}(\mathbf{X})) \sum_{m=1}^D \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1} \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \mathbf{P} \mathbf{X}_b \right] \Big\|_2 \quad (84)$$

$$+ \eta \left\| \mathbb{E} \left[Dy \mathfrak{S}(-y F_{\hat{\mathbf{v}}}(\mathbf{X})) \sum_{m=1}^D \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1} \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \mathbf{P} \mathbf{X}_b \right] \right\|_2 \quad (85)$$

$$- \mathbb{E} \left[Dy \mathfrak{S}(-y F_{\hat{\mathbf{v}}}(\mathbf{X})) \sum_{m=1}^D \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1} \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \mathbf{P} \mathbf{X}_b \right] \Big\|_2. \quad (86)$$

Remark that (86) is equal to zero because $\mathbb{E}[\mathbf{P} \mathbf{X}_u] = 0$ for all $u \in [D]$.

Summand 1: $\|(81) - (82)\|_2$. Using the matrix Hoeffding inequality, we have with high probability,

$\|(81) - (82)\|_2 \leq 16 \sqrt{\log(d) \sum_{i=1}^N M_i^2}$, where

$$\left\| \frac{\eta Dy[i]}{N} \mathfrak{S}(-y[i] F_{\hat{\mathbf{v}}}(\mathbf{X}[i])) \sum_{m=1}^D \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r[i] \rangle \right)^{p-1} \sum_{u=1}^D \hat{S}_{m,u}^{(t)} \mathbf{P} \mathbf{X}_u[i] \right\|_2^2 \leq M_i^2.$$

Induction Hypothesis B.1, Lemma D.11 and $\|\mathbf{P} \mathbf{X}_u[i]\|_2 \leq \sigma^2 d \log(d)$ imply $M_i^2 \leq \frac{\eta^2 \text{poly}(D)}{N^2}$. We deduce that $\|(81) - (82)\|_2 \leq \eta \frac{\text{poly}(D)}{\sqrt{N}}$.

Summand 2: $\|(83) - (84)\|_2$. The function $x \mapsto x^{p-1}$ is $(p-1)M^{p-2}$ Lipschitz on a bounded domain $[0, M]$. We apply this property and $(p-1) \max_{m \in [D]} \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-2} \leq (\hat{\alpha}^{(t)})^{p-2} \text{poly}(D)$ to get:

$$\begin{aligned} & \|(83) - (84)\|_2 \\ & \leq \eta (\hat{\alpha}^{(t)})^{p-2} \text{poly}(D) \mathbb{E} \left[\sum_{m=1}^D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} |\langle \hat{\mathbf{v}}^{(t)} - \hat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle| \cdot \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \|\mathbf{P} \mathbf{X}_b\|_2 \right] \\ & \leq \eta (\hat{\alpha}^{(t)})^{p-2} \text{poly}(D) \mathbb{E} \left[\sum_{m=1}^D \sum_{r=1}^D \hat{S}_{m,r}^{(t)} \varepsilon_{\hat{\mathbf{v}}}^{(t)} |\langle \mathbf{u}^{(t)}, \mathbf{X}_r \rangle| \cdot \sum_{b=1}^D \hat{S}_{m,b}^{(t)} \|\mathbf{P} \mathbf{X}_b\|_2 \right]. \end{aligned} \quad (87)$$

With high probability, $\|\xi_b\|_2 \leq \sigma \sqrt{\log(d)} \leq \sqrt{\log(d)/d}$, $|\langle \mathbf{u}^{(t)}, \xi_r \rangle| \leq \sqrt{\log(d)} \sigma$. We thus get:

$$\|(83) - (84)\|_2 \leq \eta (\hat{\alpha}^{(t)})^{p-2} \frac{\text{poly}(D)}{\sqrt{d}} \varepsilon_{\hat{\mathbf{v}}}^{(t)}. \quad (88)$$

Summand 3: $\|(85) - (86)\|_2$. We have:

$$\begin{aligned} & \|(85) - (86)\|_2 \\ & \leq \eta \left\| \mathbb{E} \left[Dy \sum_{m=1}^D \left(D \sum_{r=1}^D \widehat{S}_{m,r}^{(t)} \langle \widehat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1} \sum_{b=1}^D \widehat{S}_{m,b}^{(t)} \mathbf{P} \mathbf{X}_u \right. \right. \\ & \quad \left. \left. \left[\mathfrak{S} \left(-y \sum_{a=1}^D \left(D \sum_{a'=1}^D \widehat{S}_{a,a'}^{(t)} \langle \widehat{\mathbf{v}}^{(t)}, \mathbf{X}_{a'} \rangle \right)^p \right) - \mathfrak{S} \left(-y \sum_{a=1}^D \left(D \sum_{a'=1}^D \widehat{S}_{a,a'}^{(t)} \langle \widehat{\mathbf{v}}^{(t)}, \mathbf{X}_{a'} \rangle \right)^p \right) \right] \right\|_2. \end{aligned} \quad (89)$$

We apply [Lemma E.12](#) to bound the local change of the sigmoid in (89) which yields:

$$\begin{aligned} & \|(85) - (86)\|_2 \\ & \leq \Theta(\eta D) \mathbb{E} \left[\sum_{m=1}^D \left(D \sum_{r=1}^D \widehat{S}_{m,r}^{(t)} \langle \widehat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1D} \left| D \sum_{a=1}^D \widehat{S}_{a,a'}^{(t)} \langle \widehat{\mathbf{v}}^{(t)} - \widehat{\mathbf{v}}^{(t)}, \mathbf{X}_{a'} \rangle \right| \cdot \sum_{b=1}^D \widehat{S}_{m,u}^{(t)} \|\mathbf{P} \mathbf{X}_b\|_2 \right] \\ & \leq \Theta(\eta D^2) \mathbb{E} \left[\sum_{m=1}^D \left(D \sum_{r=1}^D \widehat{S}_{m,r}^{(t)} \langle \widehat{\mathbf{v}}^{(t)}, \mathbf{X}_r \rangle \right)^{p-1D} \sum_{a=1}^D \sum_{a'=1}^D \widehat{S}_{a,a'}^{(t)} \varepsilon_{\mathbf{v}}^{(t)} |\langle \mathbf{u}^{(t)}, \mathbf{X}_{a'} \rangle| \sum_{b=1}^D \widehat{S}_{m,u}^{(t)} \|\mathbf{P} \mathbf{X}_b\|_2 \right]. \end{aligned} \quad (90)$$

We apply [Induction Hypothesis B.1](#) to bound the softmax terms in (90). Besides, with high probability, we have $\|\boldsymbol{\xi}_b\|_2 \leq \sigma \sqrt{d \log(d)}$, $|\langle \mathbf{u}^{(t)}, \boldsymbol{\xi}_{a'} \rangle| \leq \sqrt{\log(d)} \sigma$. Thus, we have:

$$\|(85) - (86)\|_2 \leq \eta \cdot \text{poly}(D) (\widehat{\alpha}^{(t)})^{p-1} \varepsilon_{\mathbf{v}}^{(t)} \sigma \leq \eta \frac{\text{poly}(D)}{\sqrt{d}} (\widehat{\alpha}^{(t)})^{p-1} \varepsilon_{\mathbf{v}}^{(t)}. \quad (91)$$

We combine the bounds on the three summands to obtain the recursion of $\varepsilon_{\mathbf{v}}^{(t)}$. \square

We now prove [Lemma E.11](#) that gives the final bound on $\varepsilon_{\mathbf{v}}^{(t)}$ for $t \leq T$.

Lemma E.11. *For all $t \leq T$, $\varepsilon_{\mathbf{v}}^{(t)} \leq O(\frac{\text{poly}(d)}{\sqrt{N}})$. By setting $N = \text{poly}(d)$, $\varepsilon_{\mathbf{v}}^{(t)} \leq \frac{1}{\text{poly}(d)}$.*

Proof of Lemma E.11. We bound $\varepsilon_{\mathbf{v}}^{(t)}$ in the following two regimes: $t \in [\mathcal{T}, \mathcal{T} + \check{\mathcal{T}}_0]$ and $t \in [\mathcal{T} + \check{\mathcal{T}}_0, T]$.

First phase: $t \in [\mathcal{T}, \mathcal{T} + \check{\mathcal{T}}_0]$. Unraveling [Lemma E.10](#) for $t = \mathcal{T}, \dots, \mathcal{T} + \check{\mathcal{T}}_0$ leads to:

$$\begin{aligned} \varepsilon_{\mathbf{v}}^{(t)} & \leq \left[\varepsilon_{\mathbf{v}}^{(\mathcal{T})} + \eta \zeta \check{\mathcal{T}}_0 \right] \prod_{\tau=\mathcal{T}}^{\mathcal{T}+\check{\mathcal{T}}_0} \left(1 + \eta (\widehat{\alpha}^{(\tau)})^{p-1} \frac{\text{poly}(D)}{\sqrt{d}} \right) \\ & \leq \left[\varepsilon_{\mathbf{v}}^{(\mathcal{T})} + \eta \zeta \check{\mathcal{T}}_0 \right] \prod_{\tau=0}^{\check{\mathcal{T}}_0-1} \left(1 + \eta (\widehat{\alpha}^{(\tau)})^{p-1} \frac{\text{poly}(D)}{\sqrt{d}} \right). \end{aligned} \quad (92)$$

[Lemma E.20](#) provides the update of $\widehat{\alpha}^{(t)}$ during this time phase. We thus apply [Lemma K.2](#) to bound the product term in (92).

$$\prod_{\tau=0}^{\check{\mathcal{T}}_0-1} \left(1 + \eta (\widehat{\alpha}^{(\tau)})^{p-1} \frac{\text{poly}(D)}{\sqrt{d}} \right) \leq \left(1 + \frac{\Theta(1)}{C^{2(p-2)} \lambda_0^{p-2}} \right)^{\frac{\eta \text{poly}(D)}{\sqrt{d}}} \leq O(1). \quad (93)$$

Plugging (93) in (92) yields a bound on $\varepsilon_{\mathbf{v}}^{(\mathcal{T}+\check{\mathcal{T}}_0)}$.

$$\varepsilon_{\mathbf{v}}^{(\mathcal{T}+\check{\mathcal{T}}_0)} \leq O(1) (\varepsilon_{\mathbf{v}}^{(\mathcal{T})} + \eta \zeta \check{\mathcal{T}}_0). \quad (94)$$

Second phase: $t \in (\mathcal{T} + \check{\mathcal{T}}_0, T]$. [Lemma E.23](#) shows that $\hat{\alpha}^{(t)}$ gets updated until $t = \check{\mathcal{T}}_2 < T$. Therefore, we have $\varepsilon_{\mathbf{v}}^{(T)} = \varepsilon_{\mathbf{v}}^{(\mathcal{T} + \check{\mathcal{T}}_2)}$. Unraveling [Lemma E.10](#) for $t = \mathcal{T} + \check{\mathcal{T}}_0, \dots, \mathcal{T} + \check{\mathcal{T}}_2$ and using $\hat{\alpha}^{(t)} \leq \frac{\text{polylog}(d)}{C^2 \lambda_0}$ for $t \leq T$ leads to:

$$\varepsilon_{\mathbf{v}}^{(\mathcal{T} + \check{\mathcal{T}}_2)} \leq \left(1 + \frac{\eta \text{polylog}(d) \text{poly}(D)}{\lambda_0^{p-1} \sqrt{d}}\right)^{\check{\mathcal{T}}_2 - \check{\mathcal{T}}_0} \varepsilon_{\mathbf{v}}^{(\mathcal{T} + \check{\mathcal{T}}_0)} + \eta \zeta \frac{\left(1 + \frac{\eta \text{polylog}(d) \text{poly}(D)}{\lambda_0^{p-1} \sqrt{d}}\right)^{\check{\mathcal{T}}_2 - \check{\mathcal{T}}_0} - 1}{\frac{\eta \text{polylog}(d) \text{poly}(D)}{\lambda_0^{p-1} \sqrt{d}}}. \quad (95)$$

Since $\check{\mathcal{T}}_2 - \check{\mathcal{T}}_0 \leq \frac{\lambda_0^p}{\eta e^{\text{polylog}(d)}}$, we have $\left(1 + \frac{\eta \text{polylog}(d) \text{poly}(D)}{\lambda_0^{p-1} \sqrt{d}}\right)^{\check{\mathcal{T}}_2 - \check{\mathcal{T}}_0} \leq O(1)$. Simplifying [\(95\)](#) yields:

$$\varepsilon_{\mathbf{v}}^{(\mathcal{T} + \check{\mathcal{T}}_2)} \leq O\left(\frac{\text{poly}(d)}{\sqrt{N}}\right). \quad (96)$$

[\(96\)](#) implies that we need $N = \text{poly}(d)$ samples to have $\varepsilon_{\mathbf{v}}^{(T)} \leq 1/\text{poly}(d)$. \square

E.2.3 Auxiliary lemmas

In this section, we prove the Lipschitzness of the function appearing in the proof of [Lemma E.10](#).

Lemma E.12. *Let $\psi: \mathbb{R}^d \rightarrow \mathbb{R}$ defined as $\psi(\mathbf{x}) := \mathfrak{S}(-\sum_{m=1}^D x_m^p)$ and $p \geq 3$ be an odd integer. Assume that $x_m^p \geq -1/D$. Then, ψ is p -Lipschitz i.e. for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $|\psi(\mathbf{x}) - \psi(\mathbf{y})| \leq p \|\mathbf{x} - \mathbf{y}\|_1$.*

Proof of Lemma E.12. Let $l \in [D]$. The derivative of ψ with respect to a variable x_l is:

$$\left| \frac{\partial \psi}{\partial x_l} \right| = p \frac{x_l^{p-1} \exp(\sum_{m=1}^D x_m^p)}{(1 + \exp(\sum_{m=1}^D x_m^p))^2} \leq \frac{p x_l^{p-1}}{1 + \exp(\sum_{m=1}^D x_m^p)}. \quad (97)$$

[\(97\)](#) implies a bound on $\|\nabla \psi(\mathbf{x})\|_1$. Indeed, since $\sum_{m=1}^D x_m^p \geq -1$, we have:

$$\|\nabla \psi(\mathbf{x})\|_1 \leq p \frac{\sum_{m=1}^D (y x_m)^{p-1}}{1 + \exp(\sum_{m=1}^D (y x_m)^p)} \leq p. \quad (98)$$

[\(98\)](#) shows that ψ is p -Lipschitz. \square

E.2.4 $\Delta_{\mathbf{A}}^{(t)}$ stays small during the learning process

Here, we bound the gap in attention coefficients between the realistic and semi-idealized cases.

Lemma E.13. *Assume that we run GD on the empirical risk [\(E\)](#) for T iterations with parameters set as in [Parametrization 3.1](#) and the number of samples is $N = \text{poly}(d)$. Then, the attention matrix in the realistic case is very close to the semi-idealized one i.e. for $t \in [\mathcal{T}, T]$,*

$$\Delta_{\mathbf{A}}^{(t)} := \max_{i \neq j} |\hat{A}_{i,j}^{(t)} - A_{i,j}^{(t)}| \leq \frac{1}{\text{poly}(d)}.$$

We now detail the steps to prove [Lemma E.13](#). We first provide the recursion that $\Delta_{\mathbf{A}}^{(t)}$ satisfies.

Lemma E.14. *Assume that we run GD on the empirical risk [\(E\)](#) for T iterations with parameters set as in [Parametrization 3.1](#). Then, the discrepancy $\Delta_{\mathbf{A}}$ satisfies for $t \in (\mathcal{T}, T]$,*

$$\Delta_{\mathbf{A}}^{(t+1)} \leq (1 + \eta R^{(t)}) \Delta_{\mathbf{A}}^{(t)} + \eta \text{poly}(D) \sigma \varepsilon_{\mathbf{v}}^{(t)} + \eta \zeta,$$

where $\Delta_{\mathbf{A}}^{(\mathcal{T})} \leq |\hat{A}_{i,j}^{(\mathcal{T})}|$, $R^{(t)} = O(C)(G^{(t)})^{p-1} \check{\Gamma}^{(t)} (\hat{\alpha}^{(t)})^p$ and $\zeta = \frac{\text{poly}(D)}{\sqrt{N}}$.

Proof of Lemma E.14. In this proof, we maintain the hypothesis that $\Delta_{\mathbf{A}}^{(t)}$ is small. We will eventually prove this statement in Lemma E.15. Let $a, b \in [D]$ such that $a \neq b$. Using GD, $\hat{A}_{a,b}^{(t+1)} - \check{A}_{a,b}^{(t+1-\mathcal{T})}$ satisfies:

$$\begin{aligned} & |\hat{A}_{a,b}^{(t+1)} - \check{A}_{a,b}^{(t+1-\mathcal{T})}| \leq |\hat{A}_{a,b}^{(t)} - \check{A}_{a,b}^{(t-\mathcal{T})}| \\ & + \eta \left| \frac{D}{N} \sum_{i=1}^N y[i] \mathfrak{S}(-y[i] F_{\hat{\mathbf{v}}, \hat{\mathbf{A}}}(\mathbf{X}[i])) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_c[i] \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_b[i] - \mathbf{X}_r[i] \rangle \right| \end{aligned} \quad (99)$$

$$- \mathbb{E} \left[D y \mathfrak{S}(-y F_{\check{\mathbf{v}}, \hat{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \quad (100)$$

$$+ \eta \left| \mathbb{E} \left[D y \mathfrak{S}(-y F_{\hat{\mathbf{v}}, \hat{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \right| \quad (101)$$

$$- \mathbb{E} \left[D y \mathfrak{S}(-y F_{\check{\mathbf{v}}, \hat{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t-\mathcal{T})}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t-\mathcal{T})}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \quad (102)$$

$$+ D \eta \left| \mathbb{E} \left[y \mathfrak{S}(-y F_{\check{\mathbf{v}}, \hat{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t-\mathcal{T})}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t-\mathcal{T})}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \right| \quad (103)$$

$$- \mathbb{E} \left[y \mathfrak{S}(-y F_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t-\mathcal{T})}, \mathbf{X}_c \rangle \right)^{p-1} \check{S}_{a,b}^{(t-\mathcal{T})} \sum_{r \neq b} \check{S}_{a,r}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t-\mathcal{T})}, \mathbf{X}_b - \mathbf{X}_r \rangle \right]. \quad (104)$$

We now bound the three summands above.

Summand 1: |(99) – (100)|. We apply the Hoeffding inequality. With high probability, we have: $|(99) - (100)| \leq \eta \frac{\text{poly}(D)}{\sqrt{N}}$.

Summand 2: |(101) – (102)|. Since $\varepsilon_{\check{\mathbf{v}}}^{(t)}$ is small (Lemma E.11), we can show that:

$$|(101) - (102)| \leq \eta D \mathbb{E} \left[\left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \cdot \sum_{r \neq b} \hat{S}_{a,r}^{(t)} |\langle \mathbf{u}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \right] \varepsilon_{\check{\mathbf{v}}}^{(t)}. \quad (105)$$

With high probability, we have $\langle \mathbf{u}^{(t)}, \boldsymbol{\xi}_b \rangle \leq \sigma \sqrt{\log(d)}$. Using this fact along with

$D \left| \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \cdot \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \right| \leq \text{poly}(D)$, we further bound (105) as:

$$|(101) - (102)| \leq \eta \text{poly}(D) \varepsilon_{\check{\mathbf{v}}}^{(t)} \sigma \sqrt{\log(d)}. \quad (106)$$

Summand 3: |(103) – (104)|. We have the following decomposition.

$$|(103) - (104)|$$

$$\leq \eta \left| \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \hat{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \right| \quad (107)$$

$$- \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \quad (108)$$

$$+ \eta \left| \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \right| \quad (109)$$

$$- \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \quad (110)$$

$$+ \eta \left| \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \right| \quad (111)$$

$$- \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \check{S}_{a,r}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \quad (112)$$

$$+ \eta \left| \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \check{S}_{a,r}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right] \right| \quad (113)$$

$$- \mathbb{E} \left[Dy\mathfrak{S}(-yF_{\check{\mathbf{v}}, \check{\mathbf{A}}}(\mathbf{X})) \left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \check{S}_{a,b}^{(t-\mathcal{J})} \sum_{r \neq b} \check{S}_{a,r}^{(t-\mathcal{J})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \right]. \quad (114)$$

We need to distinguish two sub-cases: $a, b \in \mathcal{S}_\ell$ and $a \in \mathcal{S}_\ell, b \in \mathcal{S}_m$ with $\ell \neq m$.

Subcase 1: $a, b \in \mathcal{S}_\ell$. The proof of [Lemma G.1](#) highlights that when $a, b \in \mathcal{S}_\ell$ the event with largest gradient is event a: " $\ell = \ell(\pi(\mathbf{X}))$ " which happens with probability $1/L$. Therefore, to simplify the calculations, we will only take into account this event. We first bound |(107) – (108)|. We successively apply [Lemma E.12](#) (Lipschitzness of sigmoid) and [Lemma E.17](#) (Lipschitzness of softmax) and get:

$$\begin{aligned} & |(107) - (108)| \\ & \leq \frac{\Theta(\eta D)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{c=1}^D \sum_{c'=1}^D |\check{S}_{c,c'}^{(t-\mathcal{J})} - \hat{S}_{c,c'}^{(t)}| \cdot |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_{c'} \rangle| \sum_{r \neq b} \hat{S}_{a,r}^{(t)} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \\ & \leq \frac{\Theta(\eta D)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{c=1}^D \sum_{c'=1}^D \check{S}_{c,c'}^{(t-\mathcal{J})} |e^{2\Delta_A^{(t)}} - 1| |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_{c'} \rangle| \sum_{r \neq b} \hat{S}_{a,r}^{(t)} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \\ & \leq \frac{\Theta(\eta D)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{c=1}^D \sum_{c'=1}^D \check{S}_{c,c'}^{(t-\mathcal{J})} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_{c'} \rangle| \sum_{r \neq b} \hat{S}_{a,r}^{(t)} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \Delta_A^{(t)} \quad (115) \end{aligned}$$

$$\leq \frac{\Theta(\eta \hat{\alpha}^{(t)} \lambda_0 D^2)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \hat{S}_{a,c}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \Delta_A^{(t)}. \quad (116)$$

where we used $e^{\Delta_A^{(t)}} - 1 \leq 2\Delta_A^{(t)}$ in (115) and [Induction Hypothesis B.1](#) in (116). Using Lipschitz inequalities, we can further expand (116) as a function of the coefficients from $\check{S}^{(t)}$ and $\Delta_A^{(t)}$.

However, $\Delta_{\mathbf{A}}^{(t)}$ is small and we only want terms of order 1 in $\Delta_{\mathbf{A}}^{(t)}$ in (116). Therefore, the only term of order 1 that remains is:

$$\begin{aligned} & |(107) - (108)| \\ & \leq \frac{\Theta(\eta\hat{\alpha}^{(t)}\lambda_0 D^2)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \check{S}_{a,b}^{(t-\mathcal{T})} \sum_{r \neq b} \check{S}_{a,r}^{(t-\mathcal{T})} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \Delta_{\mathbf{A}}^{(t)}. \end{aligned} \quad (117)$$

Bounding the expectation in (117) as in the proof of Lemma G.1 yields $|(107) - (108)| \leq \eta R^{(t)} \Delta_{\mathbf{A}}^{(t)}$. We now bound $|(109) - (110)|$. We therefore apply (Lemma E.16) and get:

$$\begin{aligned} & |(109) - (110)| \quad (118) \\ & \leq \frac{\Theta(\eta D)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} |e^{2(p-1)\Delta_{\mathbf{A}}^{(t)}} - 1| \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \Big| \mathbf{a} \right] \end{aligned} \quad (119)$$

$$\leq \frac{\Theta(\eta D)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \hat{S}_{a,r}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle \Big| \mathbf{a} \right] \Delta_{\mathbf{A}}^{(t)}. \quad (120)$$

where we used $e^{2(p-1)\Delta_{\mathbf{A}}^{(t)}} - 1 \leq 4(p-1)\Delta_{\mathbf{A}}^{(t)}$ in (120). We can further expand (120), keep the terms of first order in $\Delta_{\mathbf{A}}$ and get $|(111) - (112)| \leq \eta R^{(t)} \Delta_{\mathbf{A}}^{(t)}$.

We now bound $|(111) - (112)|$. Using the Lipschitz property of the softmax (Lemma E.17), we have:

$$\begin{aligned} & |(111) - (112)| \\ & \leq \frac{\eta D}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} |e^{2\Delta_{\mathbf{A}}^{(t)}} - 1| \cdot \check{S}_{a,r}^{(t-\mathcal{T})} \cdot |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \\ & \leq \frac{\Theta(\eta D)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \hat{S}_{a,b}^{(t)} \sum_{r \neq b} \check{S}_{a,r}^{(t-\mathcal{T})} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \Delta_{\mathbf{A}}^{(t)}, \end{aligned} \quad (121)$$

where we used $|e^{2\Delta_{\mathbf{A}}^{(t)}} - 1| \leq 4\Delta_{\mathbf{A}}^{(t)}$ in (121). Using the same arguments as above, we obtain $|(111) - (112)| \leq \eta R^{(t)} \Delta_{\mathbf{A}}^{(t)}$.

The bound on $|(113) - (114)|$ can be derived as above. We again use the Lipschitz property of softmax (Lemma E.17) which leads to

$$\begin{aligned} |(113) - (114)| & \leq \frac{\Theta(\eta D)}{L} \mathbb{E} \left[\left(D \sum_{c=1}^D \check{S}_{a,c}^{(t-\mathcal{T})} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_c \rangle \right)^{p-1} \check{S}_{a,b}^{(t-\mathcal{T})} \sum_{r \neq b} \check{S}_{a,r}^{(t-\mathcal{T})} |\langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_b - \mathbf{X}_r \rangle| \Big| \mathbf{a} \right] \Delta_{\mathbf{A}}^{(t)} \\ & \leq \eta R^{(t)} \Delta_{\mathbf{A}}^{(t)}. \end{aligned} \quad (122)$$

Subcase 2: $a \in \mathcal{S}_\ell$ and $b \in \mathcal{S}_m$ with $\ell \neq m$. The proof is analogous to the Subcase 1. We only take into account event a: " $\ell = \ell(\pi(\mathbf{X}))$ and $\delta_j = 0$ " and **Event e:** " $\ell, m \neq \ell(\pi(\mathbf{X}))$ and $\delta_j = 0$ and $\delta_s = 0$ and show that $|(103) - (104)| \leq \eta \tilde{R}^{(t)} \Delta_{\mathbf{A}}^{(t)}$ where $\tilde{R}^{(t)} := O(\hat{\alpha}^{(t)}) (D\hat{\alpha}^{(t)}(\Lambda^{(t)} + (C-1)\Gamma^{(t)}))^{p-1} \Xi^{(t)} \frac{\lambda_0}{D} + O(\hat{\alpha}^{(t)}) (D\hat{\alpha}^{(t)} q D \log(d) \Xi^{(t)})^{p-1} \Xi^{(t)} q D \log(d)$.

Putting all the pieces together. Given the value of the parameters, we know that $\tilde{R}^{(t)} \leq R^{(t)}$ for all $t \in [T]$. Therefore, Summand 3 is bounded as:

$$|(103) - (104)| \leq \eta R^{(t)} \Delta_{\mathbf{A}}^{(t)}. \quad (123)$$

Conclusion. Plugging the bounds on Summands 1, 2 and 3 in the original decomposition of $|\hat{A}_{a,b}^{(t+1)} - \check{A}_{a,b}^{(t+1-\mathcal{T})}|$ yields the bound on $\Delta_{\mathbf{A}}^{(t)}$. The second part of the lemma is obtained using Lemma E.15.

□

Lemma E.15. Let $N = \text{poly}(d)$. Then, for all $t \leq T$, $\Delta_{\mathbf{A}}^{(t)} \leq \frac{1}{\text{poly}(d)}$.

Proof of Lemma E.15. Let $\varepsilon_v > 0$ such that $\varepsilon_v^{(t)} \leq \varepsilon_v$ for $t \in [T]$ – we proved the existence of ε_v in Lemma E.11. We bound $\Delta_{\mathbf{A}}^{(t)}$ when $t \in [\mathcal{T}, \mathcal{T} + \check{\mathcal{T}}_0]$ and $t \in [\mathcal{T} + \check{\mathcal{T}}_0, T]$.

First phase: $t \in [\mathcal{T}, \mathcal{T} + \check{\mathcal{T}}_0]$. Unraveling Lemma E.14 for $t = \mathcal{T}, \dots, \mathcal{T} + \check{\mathcal{T}}_0$ leads to:

$$\begin{aligned} & \Delta_{\mathbf{A}}^{(\mathcal{T} + \check{\mathcal{T}}_0)} \\ & \leq \left(\Delta_{\mathbf{A}}^{(\mathcal{T})} + \eta \check{\mathcal{T}}_0 (\varepsilon_v \text{poly}(D)\sigma + \zeta) \right) \prod_{\tau=\mathcal{T}}^{\mathcal{T} + \check{\mathcal{T}}_0 - 1} \left(1 + \frac{\eta D (\hat{\alpha}^{(\tau)})^p}{L} O(D(\check{\Lambda}^{(t)} + (C-1)\check{\Gamma}^{(t)}))^{p-1} \check{\Gamma}^{(t)} \right). \end{aligned} \quad (124)$$

We now apply Induction Hypothesis B.1 to simplify (124) and get:

$$\Delta_{\mathbf{A}}^{(\mathcal{T} + \check{\mathcal{T}}_0)} \leq \left(\Delta_{\mathbf{A}}^{(\mathcal{T})} + \eta \check{\mathcal{T}}_0 (\varepsilon_v \text{poly}(D)\sigma + \zeta) \right) \prod_{\tau=\mathcal{T}}^{\mathcal{T} + \check{\mathcal{T}}_0 - 1} \left(1 + \frac{O(\eta \lambda_0^p)}{L} (\hat{\alpha}^{(t)})^p \right). \quad (125)$$

We then apply Lemma K.2 to bound the product term in (125). We obtain:

$$\begin{aligned} \Delta_{\mathbf{A}}^{(\mathcal{T} + \check{\mathcal{T}}_0)} & \leq \left(\Delta_{\mathbf{A}}^{(\mathcal{T})} + \eta \check{\mathcal{T}}_0 (\varepsilon_v \text{poly}(D)\sigma + \zeta) \right) \left(1 + \frac{\Theta(1)}{C^{2(p-2)} \lambda_0^{p-2}} \right)^{\frac{\eta \lambda_0^{2p}}{L}} \\ & \leq O\left(\Delta_{\mathbf{A}}^{(\mathcal{T})} + \eta \check{\mathcal{T}}_0 (\varepsilon_v \text{poly}(D)\sigma + \zeta) \right). \end{aligned} \quad (126)$$

Second phase: $t \in [\mathcal{T} + \check{\mathcal{T}}_0, T]$. Unraveling Lemma E.14 for $t = \mathcal{T} + \check{\mathcal{T}}_0, \dots, \mathcal{T} + \check{\mathcal{T}}_2$ and using $\hat{\alpha}^{(t)} \leq \frac{\text{polylog}(d)}{C^2 \lambda_0}$ for $t \leq T$ leads to:

$$\Delta_{\mathbf{A}}^{(\mathcal{T} + \check{\mathcal{T}}_2)} \leq \left(1 + \frac{\eta \text{polylog}(d)}{L} \right)^{\check{\mathcal{T}}_2 - \check{\mathcal{T}}_0} \left(\Delta_{\mathbf{A}}^{(\mathcal{T} + \check{\mathcal{T}}_0) + L \cdot \text{polylog}(d) (\varepsilon_v \text{poly}(D)\sigma + \zeta) \right). \quad (127)$$

Using $\check{\mathcal{T}}_2 - \check{\mathcal{T}}_0 \leq \frac{\lambda_0^p}{\eta e^{\text{polylog}(d)}}$, we have $\left(1 + \frac{\eta \text{polylog}(d)}{L} \right)^{\check{\mathcal{T}}_2 - \check{\mathcal{T}}_0} \leq O(1)$. We thus bound (126) as:

$$\Delta_{\mathbf{A}}^{(\mathcal{T} + \check{\mathcal{T}}_2)} \leq O(\Delta_{\mathbf{A}}^{(\mathcal{T})}) + \text{polylog}(d) \left[\varepsilon_v \sigma + \frac{1}{\sqrt{N}} \right] \left[\frac{1}{C(\hat{\alpha}^{(\mathcal{T})})^{p-1} e^{\text{polyloglog}(d)}} + L \right] \leq O\left(\frac{\nu}{\sqrt{N}} \right). \quad (128)$$

We deduce that setting $N = \text{poly}(d)$ yields $\Delta_{\mathbf{A}}^{(T)} = \Delta_{\mathbf{A}}^{(\mathcal{T} + \check{\mathcal{T}}_2)} \leq 1/\text{poly}(d)$. \square

E.2.5 Auxiliary lemmas

Lemma E.16. Let $\psi: \mathbb{R}^D \rightarrow \mathbb{R}$ defined as $\psi(\mathbf{x}) := (D \sum_{m=1}^D S_m \langle \mathbf{v}, \mathbf{X}_m \rangle)^{p-1}$ where $S_m = (\text{softmax}(x_1, \dots, x_D))_m$, $\{\mathbf{X}_m\}_{m=1}^D$ and \mathbf{v} are fixed vectors and $p \geq 3$ is an odd integer. Then, we have:

$$|\psi(\mathbf{x}) - \psi(\mathbf{y})| \leq \psi(\mathbf{x}) |e^{2(p-1) \max_{c \in [D]} |x_c - y_c|} - 1|.$$

Proof of Lemma E.16. Let $S_c = (\text{softmax}(x_1, \dots, x_D))_c$, $\hat{S}_c = (\text{softmax}(y_1, \dots, y_D))_c$. We have:

$$|\psi(\mathbf{x}) - \psi(\mathbf{y})| \leq \psi(\mathbf{x}) \left| \left[\frac{\sum_{c=1}^D S_c \langle \mathbf{v}, \mathbf{X}_c \rangle}{\sum_{c=1}^D \hat{S}_c \langle \mathbf{v}, \mathbf{X}_c \rangle} \right]^{p-1} - 1 \right| \quad (129)$$

$$\leq \psi(\mathbf{x}) \left| \left[\frac{\sum_{c=1}^D e^{x_c} \langle \mathbf{v}, \mathbf{X}_c \rangle}{\sum_{c=1}^D e^{y_c} \langle \mathbf{v}, \mathbf{X}_c \rangle} \cdot \frac{\sum_{r=1}^D e^{y_r}}{\sum_{r=1}^D e^{x_r}} \right]^{p-1} - 1 \right|, \quad (130)$$

We finally apply the generalized median inequality in (130) and get:

$$\begin{aligned} |\psi(\mathbf{x}) - \psi(\mathbf{y})| &\leq \psi(\mathbf{x}) \left| \left(\max_{c \in [D]} e^{|x_c - y_c|} \cdot \max_{r \in [D]} e^{|x_r - y_r|} \right)^{p-1} - 1 \right| \\ &\leq \psi(\mathbf{x}) \left| e^{2(p-1) \max_{c \in [D]} |x_c - y_c|} - 1 \right|. \end{aligned}$$

□

Lemma E.17 (Lipschitzness of Softmax). *Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^D$. For all $i \in [D]$, we have:*

$$|\text{softmax}(\mathbf{a})_i - \text{softmax}(\mathbf{b})_i| \leq \left| e^{2 \max_{k \in [D]} |b_k - a_k|} - 1 \right| \cdot \text{softmax}(\mathbf{a})_i.$$

Proof of Lemma E.17. Let $i \in [D]$. The difference of softmax is bounded as:

$$\begin{aligned} |\text{softmax}(\mathbf{a})_i - \text{softmax}(\mathbf{b})_i| &\leq \frac{e^{a_i}}{\sum_{j=1}^D e^{a_j}} \cdot \left| \frac{e^{b_i} \sum_{j=1}^D e^{a_j}}{e^{a_i} \sum_{j=1}^D e^{b_j}} - 1 \right| \\ &\leq \frac{e^{a_i}}{\sum_{j=1}^D e^{a_j}} \cdot \left| e^{b_i - a_i} \max_{k \in [D]} e^{a_k - b_k} - 1 \right|, \end{aligned} \quad (131)$$

where we used the median inequality in the last inequality of (131). Since the exponential function is non-decreasing, we deduce:

$$|\text{softmax}(\mathbf{a})_i - \text{softmax}(\mathbf{b})_i| \leq \left| e^{2 \max_{k \in [D]} |b_k - a_k|} - 1 \right| \cdot \text{softmax}(\mathbf{a})_i. \quad (132)$$

□

E.2.6 Dynamics of $\hat{\mathbf{v}}^{(t)}$

Lastly, since $\Delta_{\mathbf{A}}^{(t)}$ remains small and $\hat{\mathbf{v}}^{(t)}$ mainly lies in $\text{span}(\mathbf{w}^*)$, we show that $\hat{\alpha}^{(t)}$ satisfies the same updates as the ideal $\alpha^{(t)}$ (up to some constant factors).

Lemma E.18. *Assume that we run GD on the empirical risk (E) for T iterations with parameters set as in Parametrization 3.1 and the number of samples is $N = \text{poly}(d)$. Then, there exist times $\mathcal{T}, \hat{\mathcal{T}}_0, \hat{\mathcal{T}}_1 > 0$ such that*

1. *Analog of Event I (Lemma D.2): $\hat{\alpha}^{(t+1)} = \hat{\alpha}^{(t)} + \Theta(\eta)(\hat{G}^{(t)})^p (\hat{\alpha}^{(t)})^{p-1}$ for $t \in [\mathcal{T}, \mathcal{T} + \hat{\mathcal{T}}_0]$.*
2. *Analog of Event III (Lemma D.11): $\hat{\alpha}^{(t+1)} = \hat{\alpha}^{(t)} + \Theta(\eta)(\hat{G}^{(t)})^p (\hat{\alpha}^{(t)})^{p-1}$ for $t \in [\mathcal{T} + \hat{\mathcal{T}}_1, T]$.*

Consequently, $\hat{\alpha}^{(t)}$ is non-decreasing and eventually, $\hat{\alpha}^{(T)} = \alpha^{(T)} = \frac{\text{polylog}(d)}{C^2 \lambda_0}$.

These three lemmas imply that at time T , the realistic iterates are very close to the ideal ones. Therefore, they incur nearby test loss and thus the realistic model generalizes. We now proceed to the proof of

In order to analyze the dynamics of $\hat{\mathbf{v}}^{(t)}$, we first show that the gradient (with respect to $\hat{\mathbf{v}}$) in the realistic learning process is very close to the one in the semi-idealized one.

Lemma E.19. *Let $t \in [\mathcal{T}, T]$. With high probability, we have*

$$\|\nabla_{\hat{\mathbf{v}}} \hat{\mathcal{L}}(\hat{\mathbf{v}}^{(t)}, \hat{\mathbf{A}}^{(t)}) - \nabla_{\check{\mathbf{v}}} \check{\mathcal{L}}(\check{\mathbf{v}}^{(t-\mathcal{T})}, \check{\mathbf{A}}^{(t-\mathcal{T})})\|_2 \leq \text{poly}(D) \left(\frac{1}{\sqrt{N}} + \varepsilon_{\check{\mathbf{v}}}^{(t)} + \Delta_{\hat{\mathbf{A}}}^{(t)} \right). \quad (133)$$

By choosing $N = \text{poly}(d)$, we have $\|\nabla_{\hat{\mathbf{v}}} \hat{\mathcal{L}}(\hat{\mathbf{v}}^{(t)}, \hat{\mathbf{A}}^{(t)}) - \nabla_{\check{\mathbf{v}}} \check{\mathcal{L}}(\check{\mathbf{v}}^{(t-\mathcal{T})}, \check{\mathbf{A}}^{(t-\mathcal{T})})\|_2 \leq 1/\text{poly}(d)$.

Proof of Lemma E.19. We have:

$$\|\nabla_{\hat{\mathbf{v}}} \hat{\mathcal{L}}(\hat{\mathbf{v}}, \hat{\mathbf{A}}) - \nabla_{\check{\mathbf{v}}} \check{\mathcal{L}}(\check{\mathbf{v}}, \check{\mathbf{A}})\|_2 \quad (134)$$

$$= \left| \frac{D}{N} \sum_{i=1}^N y[i] \mathfrak{S}(-y[i] F_{\hat{\mathbf{A}}^{(t)}, \hat{\mathbf{v}}^{(t)}}(\mathbf{X}[i])) \sum_{j=1}^D \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_k \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{X}_r, \mathbf{w}^* \rangle \right. \quad (135)$$

$$\left. - D \mathbb{E} \left[y \mathfrak{S}(-y F_{\hat{\mathbf{A}}^{(t)}, \hat{\mathbf{v}}^{(t)}}(\mathbf{X})) \sum_{j=1}^D \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_k \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{X}_r, \mathbf{w}^* \rangle \right] \right| \quad (136)$$

$$+ D \left| \mathbb{E} \left[y \mathfrak{S}(-y F_{\hat{\mathbf{A}}^{(t)}, \hat{\mathbf{v}}^{(t)}}(\mathbf{X})) \sum_{j=1}^D \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_k \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{X}_r, \mathbf{w}^* \rangle \right] \right. \quad (137)$$

$$\left. - \mathbb{E} \left[y \mathfrak{S}(-y F_{\hat{\mathbf{A}}^{(t)}, \check{\mathbf{v}}^{(t)}}(\mathbf{X})) \sum_{j=1}^D \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_k \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{X}_r, \mathbf{w}^* \rangle \right] \right| \quad (138)$$

$$+ D \left| \mathbb{E} \left[y \mathfrak{S}(-y F_{\hat{\mathbf{A}}^{(t)}, \check{\mathbf{v}}^{(t)}}(\mathbf{X})) \sum_{j=1}^D \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_k \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{X}_r, \mathbf{w}^* \rangle \right] \right. \quad (139)$$

$$\left. - \mathbb{E} \left[y \mathfrak{S}(-y F_{\check{\mathbf{A}}^{(t)}, \check{\mathbf{v}}^{(t)}}(\mathbf{X})) \sum_{j=1}^D \sigma' \left(\sum_{k=1}^D \check{S}_{j,k}^{(t)} \langle \check{\mathbf{v}}^{(t)}, \mathbf{X}_k \rangle \right) \sum_{r=1}^D \check{S}_{j,r}^{(t)} \langle \mathbf{X}_r, \mathbf{w}^* \rangle \right] \right| \quad (140)$$

$$+ \varepsilon_v^{(t)} \left| \frac{D}{N} \sum_{i=1}^N y[i] \mathfrak{S}(-y[i] F_{\hat{\mathbf{A}}^{(t)}, \hat{\mathbf{v}}^{(t)}}(\mathbf{X}[i])) \sum_{j=1}^D \sigma' \left(\sum_{k=1}^D \hat{S}_{j,k}^{(t)} \langle \hat{\mathbf{v}}^{(t)}, \mathbf{X}_k \rangle \right) \sum_{r=1}^D \hat{S}_{j,r}^{(t)} \langle \mathbf{X}_r, \mathbf{u}^{(t)} \rangle \right|. \quad (141)$$

We bound each of the terms above using concentration or lipschitz inequalities. Using the same arguments as in the proof of Lemma E.10, we have |(134) – (135)| ≤ poly(D)/√N and |(137) – (138)| ≤ poly(D)ε_v^{(t)}. Using the same steps as in the proof of Lemma E.14, we have |(139)–(140)| ≤ poly(D)Δ_A^{(t)}. Lastly, |(141)| ≤ poly(D)ε_v^{(t)}. Summing up all these terms yields the aimed result. □

Lemma E.19 shows that we can use the gradient from the semi-idealized process to analyze the dynamics of $\hat{\alpha}^{(t)}$ in the real process. Therefore, we can derive similar updates for $\hat{\alpha}^{(t)}$ as in Lemma D.2, Lemma D.8 and Lemma D.11.

Lemma E.20. Let $\check{\mathcal{T}}_0 = \Theta \left(\frac{1}{\eta C (\hat{\alpha}^{(0)})^{p-1}} \right)$. Therefore, $\hat{\alpha}^{(t)}$ is updated as

$$\hat{\alpha}^{(t+1)} = \hat{\alpha}^{(t)} + \Theta(\eta C) (\check{G}^{(t)})^p (\hat{\alpha}^{(t)})^{p-1}.$$

Consequently, $\hat{\alpha}^{(t)}$ is non-decreasing and after $\check{\mathcal{T}}_0$ iterations, we have $\hat{\alpha}^{(t)} \geq \frac{\Omega(1)}{C^2 \lambda_0}$ for $t \geq \check{\mathcal{T}}_0$.

Lemma E.21. Let $\check{\mathcal{T}}_1 = \check{\mathcal{T}}_0 + \Theta \left(\frac{\lambda_0^p}{\eta e^{\text{poly} \log \log(d)}} \right)$. For all $t \in [\check{\mathcal{T}}_1, T]$, we have $\check{\Gamma}^{(t)} \geq \frac{\Omega(\lambda_0)}{D}$.

Lemma E.22. Let $\check{\mathcal{T}}_1 = \check{\mathcal{T}}_0 + \Theta \left(\frac{\lambda_0^p}{\eta e^{\text{poly} \log \log(d)}} \right)$ and $t \in [\check{\mathcal{T}}_1, T]$. $\hat{\alpha}^{(t)}$ updates as

$$\hat{\alpha}^{(t+1)} = \hat{\alpha}^{(t)} + \Theta(\eta C) (\check{G}^{(t)})^p (\hat{\alpha}^{(t)})^{p-1}. \quad (142)$$

Consequently, $\hat{\alpha}^{(t)}$ is non-decreasing and eventually $\hat{\alpha}^{(T)} = \frac{\text{poly} \log(d)}{C^2 \lambda_0}$.

Auxiliary lemma. The following lemma is useful to prove Lemma E.10 and Lemma E.14.

Lemma E.23. The time at which $\hat{\alpha}^{(t)}$ stops increasing is $\check{\mathcal{T}}_2 = \check{\mathcal{T}}_1 + \Theta \left(\frac{\lambda_0^{p-2}}{\eta \text{poly} \log(d)} \right)$.

Proof of Lemma E.23. The result is obtained by applying Lemma K.1 to (142). We have:

$$\check{\mathcal{T}}_2 = \check{\mathcal{T}}_1 + \frac{3(\lambda_0)^{p-2}}{\eta \text{poly} \log(d) e^{\text{poly} \log \log(d)}} + \frac{2^p \lambda_0^p}{e^{\text{poly} \log \log(d)}} \log \log(d).$$

□

E.2.7 The realistic model fits the labeling function

Lemma E.24. *In the realistic case, the model fits the labeling function i.e.*

$$\mathbb{P}_{\mathcal{D}}[f^*(\mathbf{X})F_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) > 0] \geq 1 - o(1). \quad (143)$$

Proof of Lemma E.24. We bound the population risk $\mathcal{L}(\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)})$. We have:

$$\mathcal{L}(\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}) \leq \left| \mathcal{L}(\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}) - \mathcal{L}(\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}) \right| + \mathcal{L}(\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}). \quad (144)$$

Using Lemma E.25, we have $\mathcal{L}(\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}) \leq 1/\text{poly}(d)$. We now bound the first summand in (144) using the 1-Lipschitzness of the logistic function and get:

$$\left| \mathcal{L}(\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}) - \mathcal{L}(\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}) \right| \leq \mathbb{E} \left[\left| F_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) - F_{\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}}(\mathbf{X}) \right| \right] \quad (145)$$

Using Lemma E.26, we have $\mathbb{E} \left[\left| F_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) - F_{\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}}(\mathbf{X}) \right| \right] \leq 1/\text{poly}(d)$. Therefore, we deduce that $\mathcal{L}(\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}) \leq 1/\text{poly}(d)$. Since the 0-1 loss is a convex surrogate, we have

$$\mathbb{P}_{(\mathbf{X}, y) \sim \mathcal{D}}[yF_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) < 0] \leq \mathcal{L}(\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}) \leq 1/\text{poly}(d). \quad (146)$$

We can further expand (146) as in the proof of Lemma D.15 and deduce the aimed result. \square

To prove Lemma E.24, we use the following auxiliary lemma.

Lemma E.25. *After T iterations, the population risk in the semi-idealized case converges i.e. $\check{\mathcal{L}}(\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}) \leq o(1)$.*

Proof of Lemma E.25. The proof is similar to the one of Lemma D.15. \square

Lemma E.26. *For all \mathbf{X} sampled from \mathcal{D} , we have*

$$\left| F_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) - F_{\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}}(\mathbf{X}) \right| \leq 1/\text{poly}(d).$$

Proof of Lemma E.26. We have:

$$\left| F_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) - F_{\mathbf{Q}^{(T)}, \check{\mathbf{v}}^{(T)}}(\mathbf{X}) \right| \leq \left| F_{\hat{\mathbf{A}}^{(T)}, \hat{\mathbf{v}}^{(T)}}(\mathbf{X}) - F_{\hat{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}}(\mathbf{X}) \right| \quad (147)$$

$$+ \left| F_{\hat{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}}(\mathbf{X}) - F_{\check{\mathbf{A}}^{(T)}, \check{\mathbf{v}}^{(T)}}(\mathbf{X}) \right|. \quad (148)$$

We now separately bound (147) and (148).

Bound on (147). Since $x \mapsto x^p + \nu x$ is Lipschitz on a bounded domain, we have:

$$(147) \leq \sum_{m=1}^D \sigma' \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(T)} |\langle \hat{\mathbf{v}}^{(T)}, \mathbf{X}_r \rangle| \right) D \sum_{r=1}^D \hat{S}_{m,r}^{(T)} |\langle \hat{\mathbf{v}}^{(T)} - \check{\mathbf{v}}^{(T)}, \mathbf{X}_r \rangle| \quad (149)$$

$$\leq \varepsilon_{\mathbf{v}}^{(T)} \sum_{m=1}^D \sigma' \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(T)} |\langle \hat{\mathbf{v}}^{(T)}, \mathbf{X}_r \rangle| \right) D \sum_{r=1}^D \hat{S}_{m,r}^{(T)} |\langle \mathbf{u}^{(t)}, \mathbf{X}_r \rangle|, \quad (150)$$

since $\langle \hat{\mathbf{v}}^{(T)}, \mathbf{w}^* \rangle = \langle \check{\mathbf{v}}^{(T)}, \mathbf{w}^* \rangle$. Using Cauchy-Schwarz inequality, (149) simplifies as:

$$(147) \leq \sum_{m=1}^D \sigma' \left(D \sum_{r=1}^D \hat{S}_{m,r}^{(T)} |\langle \hat{\mathbf{v}}^{(T)}, \mathbf{X}_r \rangle| \right)^{p-1} D \sum_{r=1}^D \hat{S}_{m,r}^{(T)} \varepsilon_{\mathbf{v}}^{(T)} \leq \text{poly}(D) \varepsilon_{\mathbf{v}}^{(T)}. \quad (151)$$

Using Lemma E.10, we conclude that (147) $\leq 1/\text{poly}(d)$.

Bound on (148). We again use the Lipschitzness of the power function and get:

$$(148) \leq \sum_{m=1}^D \sigma' \left(D \sum_{r=1}^D \max\{\check{S}_{m,r}^{(T)}, \hat{S}_{m,r}^{(T)}\} |\langle \mathbf{v}^{(T)}, \mathbf{X}_r \rangle| \right) D \sum_{r=1}^D |\hat{S}_{m,r}^{(T)} - \check{S}_{m,r}^{(T)}| \cdot |\langle \hat{\mathbf{v}}^{(T)}, \mathbf{X}_r \rangle|. \quad (152)$$

We apply Lemma E.17 in (152) to get (148) $\leq \text{poly}(D) \Delta_{\mathbf{A}}^{(T)}$. Finally, we apply Lemma E.14 to get (148) $\leq 1/\text{poly}(d)$. \square

F Transfer Learning

In this section, we show that a transformer that has been pre-trained on a structured dataset require a few samples to generalize in a new dataset sharing the same structure.

Theorem 5.1. *Let $\hat{\mathbf{A}}$ be the attention matrix obtained after pre-training as in [subsection 4.2](#). Assume that we run GD for T iterations on $(\tilde{\mathbf{E}})$ to fine-tune the value vector. Using $\tilde{N} \leq \text{polylog}(D)$ samples, the model (\mathbf{T}) transfers to $\tilde{\mathcal{D}}$ i.e. $\mathbb{P}_{\tilde{\mathcal{D}}}[f^*(\mathbf{X})F_{\hat{\mathbf{A}}, \tilde{\mathbf{v}}(T)}(\mathbf{X}) > 0] \geq 1 - o(1)$.*

Proof of Theorem 5.1. Actually, even one step of the update using normalized gradient descent on \mathbf{v} can already achieve test accuracy $\geq 1 - o(1)$. We know that for a datum (\mathbf{X}, y) , the gradient of $L(\mathbf{X})$ with respect to \mathbf{v} is

$$\nabla_{\mathbf{v}} L(\mathbf{X}) = -y \mathfrak{G}(-yF(\mathbf{X})) \sum_{m=1}^D \sigma'(\langle \mathbf{O}_m^{(t)}, \tilde{\mathbf{v}}^{(t)} \rangle) \mathbf{O}_m^{(t)} \quad (153)$$

Since $\tilde{\mathbf{v}}^{(0)} = \mathbf{0}$, we have $F_{\tilde{\mathbf{v}}^{(0)}}(\mathbf{X}) = 0$ and $\sigma'(\langle \mathbf{O}_m^{(0)}, \tilde{\mathbf{v}}^{(0)} \rangle) = \nu$. Thus, the gradient (153) simplifies to

$$\nabla_{\mathbf{v}} L(\mathbf{X}) = -\frac{\nu y}{2} \sum_{j=1}^D \mathbf{X}_j \sum_{m=1}^D \hat{S}_{m,j}^{(t)} \quad (154)$$

By symmetry of the $\hat{S}_{j,m}^{(t)}$, we know that

$$\left| \sum_{m=1}^D \hat{S}_{m,j}^{(t)} - D \right| \leq \frac{D}{\text{poly}(d)}, \quad (155)$$

where $\frac{1}{\text{poly}(d)}$ comes from the ε_Q part in the previous section. Moreover, since the noise and feature noise has mean zero independent of y , we know that there exists some value $c_0 > 0$ (roughly equal to αDC) such that:

$$\mathbb{E}[\nabla_{\mathbf{v}} L(\mathbf{X})] = c_0 \tilde{\mathbf{w}}^*. \quad (156)$$

Now, by standard concentration inequality, we know that for N i.i.d. samples $\mathbf{X}[i], y[i]$, with high probability

$$\nabla_0 := \sum_{i=1}^N \nabla_{\mathbf{v}} L(\mathbf{X}[i]) = c_0 \tilde{\mathbf{w}}^* + \varepsilon_0 \tilde{\mathbf{w}}^* + \chi_0, \quad (157)$$

where ε_0 comes from the feature noise

$$|\varepsilon_0| \leq \frac{c_0 q D \log(d)}{\sqrt{N}}, \quad (158)$$

and χ_0 comes from the noise:

$$\|\chi_0\|_2 \leq \frac{c_0 \sigma \sqrt{d} \sqrt{D}}{\sqrt{N}}. \quad (159)$$

Therefore, if we update using normalized GD:

$$\tilde{\mathbf{v}}^{(1)} = \tilde{\mathbf{v}}^{(0)} + \frac{\nabla_0}{\|\nabla_0\|_2} \quad (160)$$

we have that:

$$\tilde{\mathbf{v}}^{(1)} = c_1 \tilde{\mathbf{w}}^* + \chi_1, \quad (161)$$

where $\chi_1 \perp \tilde{\mathbf{w}}^*$, $\|\chi_1\|_2 = O(1)$ and $c_1 \geq \tilde{\Omega}(1)/\sqrt{D}$. Now, for a new datum \mathbf{X}^{new} with noises $\{\xi_i^{new}\}_{i=1}^D$, we know that w.h.p

$$|\langle \xi_i^{new}, \tilde{\mathbf{v}}^{(1)} \rangle| \leq \frac{\tilde{O}(1)}{\sqrt{d}}, \quad |\langle \tilde{\mathbf{w}}^*, \tilde{\mathbf{v}}^{(1)} \rangle| \geq \frac{\tilde{\Omega}(1)}{\sqrt{D}}. \quad (162)$$

We can prove the test accuracy is small using the same proof as in [Lemma D.15](#), where we show that:

$$y\sigma(\langle \tilde{\mathbf{v}}^{(1)}, \mathbf{O}_m^{(1)} \rangle) \geq 0, \quad (163)$$

for $m \in \mathcal{S}_{\ell(\mathbf{X})}$ and it dominates the other $y\sigma(\langle \tilde{\mathbf{v}}^{(1)}, \mathbf{O}_j^{(1)} \rangle)$ for $j \in \mathcal{S}_{\ell}$ with $\ell \neq \ell(\mathbf{X})$. Therefore, we prove that $\mathbb{P}_{\tilde{\mathcal{D}}}[yF(\mathbf{X}) > 0] \geq 1 - o(1)$ which implies the aimed result. \square

Theorem 5.2. *Let $\mathcal{A}: \mathbb{R}^{d \times D} \rightarrow \{-1, 1\}$ be a binary classification algorithm without patch association knowledge. Then, it needs $D^{\Omega(1)}$ training samples to get test error $\leq o(1)$ on $\tilde{\mathcal{D}}$.*

Proof of Theorem 5.2. Let \mathcal{A} be an algorithm. Assume that at training time the algorithm \mathcal{A} has access to $D^{o(1)}$ training data, Since each input X is made of D patches, this means that there exist at least $\Omega(1)$ fraction of $k \in [D]$ such that \mathcal{A} has not seen training samples with $\ell(X) = k$. Consider the following two distributions over $\{-1, 0, 1\}^m$:

1. \mathcal{D}_1 : Sample $z \in \{0, 1\}^m$ where each z_i i.i.d. 1 w.p. $q/2$, -1 w.p. $q/2$ and 0 otherwise.
2. \mathcal{D}_2 : Sample a set \mathcal{S} uniformly at random from $[m]$ of size C , set all $z_i = 1$ for $i \in \mathcal{S}$, and sample other z_j i.i.d. 1 w.p. $q/2$, -1 w.p. $q/2$ and 0 otherwise.

We can easily see that as long as $qm = \text{poly}(C)$, then

$$\mathbf{TV}(\mathcal{D}_1, \mathcal{D}_2) = o(1)$$

This implies that \mathcal{A} must have bad generalization error ($\Omega(1)$) on $\tilde{\mathcal{D}}$. \square

G Gradient descent updates in the idealized process

In this section, we derive the gradient descent updates of $A_{i,j}$ in the idealized learning process.

G.1 Indices in the same set: $i, j \in \mathcal{S}_{\ell}$

Lemma G.1. *Let $T > 0$ be the time where the population loss is at most $o(1)$ and $t \in [0, T]$. Then, $\gamma^{(t)}$ satisfies the update*

$$\gamma^{(t+1)} = \gamma^{(t)} + \eta \Theta(C) \alpha^{(t)} \left(D \alpha^{(t)} [\Lambda^{(t)} + (C-1)\Gamma^{(t)}] \right)^{p-1} \Gamma^{(t)}.$$

Proof of Lemma G.1. Let $\ell \in [L]$ and $i, j \in \mathcal{S}_{\ell}$ with $i \neq j$. The main idea of the proof is to bound the gradient of $L(\mathbf{X})$ with respect to $A_{i,j}$. This gradient is given by [Lemma H.2](#) and is made of two terms: the σ' term and the sum outside σ' . We distinguish the following cases and bound these two terms.

1. $\ell = \ell(\mathbf{X})$. We first bound the outside sum. Using [Lemma K.3](#), we have

$$\left| \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} (1 - y\delta_r) - (D - C) \right| \leq \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} |\delta_r| \leq qD \log(d). \quad (164)$$

Since $(D - C)/2 - qD \log(d) \geq 0$, we rewrite (164) as:

$$\frac{1}{2}(D - C) \leq \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} (1 - y\delta_r) \leq 2(D - C). \quad (165)$$

Regarding the sum inside σ' , we use [Lemma G.3](#) which shows:

$$\Lambda^{(t)} + (C - 1)\Gamma^{(t)} + \Xi^{(t)} \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} \delta_{h,r} = \Theta(1) \left(\Lambda^{(t)} + (C - 1)\Gamma^{(t)} \right). \quad (166)$$

By using (165) and (166), we finally obtain:

$$-\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} = \Theta(D\alpha^{(t)})\mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)}(\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^{p-1} \Gamma^{(t)}(D-C)\Xi^{(t)}. \quad (167)$$

2. $\ell \neq \ell(\mathbf{X})$ and $\delta_s = 0$ for all $s \in \mathcal{S}_\ell$. We first bound the outside sum. Since $\delta_s = 0$, the only non-zero term is the one with factor $\Xi^{(t)}$. Using triangle inequality, we have:

$$\Xi^{(t)} \left| Cy + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} \delta_r \right| \leq \Xi^{(t)} \left(C + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} |\delta_r| \right) = \Xi^{(t)} \kappa(\mathbf{X}). \quad (168)$$

We now bound the sum inside σ' . This sum is actually equal to the outside sum and we can therefore use the bound (168). Therefore, the overall gradient is bounded as:

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)})\mathfrak{S}(-yF(\mathbf{X})) \left(D\Xi^{(t)}\alpha^{(t)}\kappa(\mathbf{X}) \right)^{p-1} \Gamma^{(t)}\Xi^{(t)}\kappa(\mathbf{X}). \quad (169)$$

3. $\ell \neq \ell(\mathbf{X})$ and at least one $\delta_s \neq 0$ and $\delta_i = 0$. We first bound the outside sum. Using Lemma K.3 and $\Lambda^{(t)} + \Gamma^{(t)} + \Xi^{(t)} = 1$, we have:

$$\begin{aligned} & \alpha^{(t)} \left| \Lambda^{(t)}\delta_j + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} (\delta_j - \delta_r) + \Xi^{(t)}C(\delta_j - y) + \Xi^{(t)} \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{m \in \mathcal{S}_h} (\delta_j - \delta_m) \right| \\ & \leq \alpha^{(t)} \left[(C-1)\Gamma^{(t)} + \Xi^{(t)}(C + \Theta(qD)\log(d)) + \mathbf{1}_{\delta_j \neq 0} \right]. \end{aligned} \quad (170)$$

We lastly apply Induction Hypothesis B.1 to show that (170) is less or equal to $\Theta(\alpha^{(t)})$. We now bound the sum inside σ' .

$$\begin{aligned} \alpha^{(t)} \left| \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i,j\}} \delta_r + \Xi^{(t)} \left(Cy + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{m \in \mathcal{S}_h} \delta_m \right) \right| & \leq \alpha^{(t)} \left(\Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i,j\}} |\delta_r| + \Xi^{(t)}\kappa(\mathbf{X}) \right) \\ & \leq \alpha^{(t)} \left((C-1)\Gamma^{(t)} + \Xi^{(t)}\kappa(\mathbf{X}) \right). \end{aligned} \quad (171)$$

The overall bound on the derivative is:

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)})\mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)}((C-1)\Gamma^{(t)} + \Xi^{(t)}\kappa(\mathbf{X})) \right)^{p-1} \Gamma^{(t)}. \quad (172)$$

4. $\ell \neq \ell(\mathbf{X})$ where $\delta_s \neq 0$ for all s . We first bound the outside sum as follows.

$$\begin{aligned} & \alpha^{(t)} \left| \Lambda^{(t)}(\delta_j - \delta_i) + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i,j\}} (\delta_j - \delta_r) + \Xi^{(t)} \left[C(\delta_j - y) + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{m \in \mathcal{S}_h} (\delta_j - \delta_m) \right] \right| \\ & \leq \alpha^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)}\kappa(\mathbf{X}) + \mathbf{1}_{\delta_j \neq 0}). \end{aligned} \quad (173)$$

We now bound the sum inside σ' .

$$\begin{aligned} & \alpha^{(t)} \left| \Lambda^{(t)}\delta_i + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i,j\}} \delta_r + \Xi^{(t)} \left[Cy + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{m \in \mathcal{S}_h} \delta_m \right] \right| \\ & \leq \alpha^{(t)} \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)}\kappa(\mathbf{X}) \right). \end{aligned} \quad (174)$$

We lastly apply Lemma G.3 to show that (174) is bounded by $\Theta(\alpha^{(t)})(\Lambda^{(t)} + (C-1)\Gamma^{(t)})$. Therefore, the overall gradient is bounded as

$$\begin{aligned} \left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| & \leq \Theta(D\alpha^{(t)})\mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)}((C-1)\Gamma^{(t)} + \Lambda^{(t)}) \right)^{p-1} \Gamma^{(t)} \\ & \quad \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)}\kappa(\mathbf{X}) + \mathbf{1}_{\delta_j \neq 0} \right). \end{aligned} \quad (175)$$

Putting all the pieces together. We now bound the derivative of the population loss. Using Tower property and [Lemma I.1](#), we have:

$$\mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] = \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \right] = \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \mid \mathbf{X} \right] \right]. \quad (176)$$

For a *fixed* \mathbf{X} , we now bound the derivative of the loss evaluated in \mathbf{X} . For $i, j \in \mathcal{S}_\ell$, we distinguish the four possible events depending on the randomness of π .

- **Event a:** " $\ell = \ell(\pi(\mathbf{X}))$ " occurs with probability $1/L$.
- **Event b:** " $\ell \neq \ell(\pi(\mathbf{X}))$ and $\delta_s = 0$ for all $s \in \mathcal{S}_\ell$ " occurs with probability $(1-1/L)(1-q)^C$.
- **Event \mathbf{c}_k :** " $\ell \neq \ell(\pi(\mathbf{X}))$ and $\#\{s : \delta_s \neq 0\} = k$ for $1 \leq k \leq C-1$ and $\delta_i = 0$ " occurs with probability $(1-1/L) \binom{C-1}{k} q^k (1-q)^{C-k}$.
- **Event d:** " $\ell \neq \ell(\pi(\mathbf{X}))$ and $\delta_s \neq 0$ for all s " occurs with probability $(1-1/L) \cdot q^C$.

Therefore, the derivative of the loss in \mathbf{X} is:

$$\begin{aligned} \mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \mid \mathbf{X} \right] &= \frac{1}{L} \mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \mid \mathbf{X}, \mathbf{a} \right] \\ &\quad + \left(1 - \frac{1}{L}\right) (1-q)^C \mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \mid \mathbf{X}, \mathbf{b} \right] \\ &\quad + \left(1 - \frac{1}{L}\right) \sum_{k=1}^{C-1} \binom{C-1}{k} q^k (1-q)^{C-k} \mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \mid \mathbf{X}, \mathbf{c}_k \right] \\ &\quad + \left(1 - \frac{1}{L}\right) q^C \mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \mid \mathbf{X}, \mathbf{d} \right]. \end{aligned} \quad (177)$$

Event a is the event that is the most likely to happen. Therefore, we only take into account $\mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \mid \mathbf{a} \right]$ in (177) and obtain:

$$\mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] = \frac{\Theta(D\alpha^{(t)})}{L} \mathbb{E}_{\mathbf{X}} [\mathfrak{S}(-yF(\mathbf{X}))] \left(D\alpha^{(t)}(\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^{p-1} \Gamma^{(t)}. \quad (178)$$

Since the population loss is a $\Omega(1)$ for $t \leq T$, this implies that $\mathbb{E}_{\mathbf{X}}[\mathfrak{S}(-yF(\mathbf{X}))] = \Theta(1)$ ([Lemma D.10](#)). We thus plug (178) in the update of $\gamma^{(t)}$ to obtain the desired result. \square

Corollary G.1. *Let $T > 0$ be the time where the population loss is $o(1)$ and $t \leq T$. Let $G^{(t)} := D(\Lambda^{(t)} + (C-1)\Gamma^{(t)})$. The update of $\gamma^{(t)}$ satisfies:*

$$\gamma^{(t+1)} = \gamma^{(t)} + \Theta(C\eta)(\alpha^{(t)})^p \Gamma^{(t)} (G^{(t)})^{p-1}.$$

Proof of Corollary G.1. [Lemma G.1](#) provides the update rule of $\gamma^{(t)}$.

$$\gamma^{(t+1)} = \gamma^{(t)} + \eta \Theta(C)(\alpha^{(t)})^p \Gamma^{(t)} (G^{(t)})^{p-1}.$$

\square

G.2 Update for $i \in \mathcal{S}_\ell$ and $j \in \mathcal{S}_m$

Lemma G.2. *Let $T > 0$ be the time where the population loss is at most $o(1)$ and $t \in [0, T]$. Then, $\rho^{(t)}$ satisfies the update*

$$\begin{aligned} \left| \frac{\rho^{(t+1)} - \rho^{(t)}}{\eta} \right| &\leq \Theta(C^2 \alpha^{(t)}) \left(D\alpha^{(t)}(\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^{p-1} \Xi^{(t)} \frac{\lambda_0}{D} \\ &\quad + \Theta(\alpha^{(t)}) \left(D\alpha^{(t)} q D \log(d) \Xi^{(t)} \right)^{p-1} \Xi^{(t)} q D \log(d). \end{aligned}$$

Proof of Lemma G.1. Let $\ell, m \in [L]$ such that $\ell \neq m$ and $i \in \mathcal{S}_\ell, j \in \mathcal{S}_m$. The main idea of the proof is to bound the gradient of $L(\mathbf{X})$ with respect to $A_{i,j}$. This gradient is given by Lemma H.2 and is made of two terms: the σ' term and the sum outside σ' . We distinguish the following cases and bound these two terms.

1. $\ell = \ell(\mathbf{X})$ and $\delta_j = 0$. We first bound the outside sum. We apply Lemma G.3 to obtain:

$$\alpha^{(t)} \left| (\Lambda^{(t)} + (C-1)\Gamma^{(t)})y + \Xi^{(t)} \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} \delta_r \right| \leq 2\alpha^{(t)} \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} \right). \quad (179)$$

We now bound the sum inside σ' . This sum is actually equal to the outside sum and we can therefore use the bound (179). Therefore, the overall gradient is bounded as:

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{G}(-yF(\mathbf{X})) \left(D\alpha^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^{p-1} \Xi^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}). \quad (180)$$

2. $\ell = \ell(\mathbf{X})$ and $\delta_j \neq 0$. We first bound the outside sum. We successively apply $\Lambda^{(t)} + \Gamma^{(t)} + \Xi^{(t)} = 1$, Lemma K.3 and Induction Hypothesis B.1 to obtain:

$$\begin{aligned} & \alpha^{(t)} \left| \Lambda^{(t)}(\delta_j - y) + (C-1)\Gamma^{(t)}(\delta_j - y) + \Xi^{(t)} \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} (\delta_j - \delta_r) \right| \\ & \leq \alpha^{(t)} \left[\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} |\delta_r| + \mathbf{1}_{\delta_j \neq 0} \right] \\ & \leq \alpha^{(t)} \left[\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} qD \log(d) + \mathbf{1}_{\delta_j \neq 0} \right] \\ & \leq \alpha^{(t)} \left[\frac{e^{\text{polyloglog}(d)} + \lambda_0}{D} + \Theta(q \log(d)) + \mathbf{1}_{\delta_j \neq 0} \right] = \Theta(\alpha^{(t)}). \end{aligned} \quad (181)$$

We now bound the sum inside σ' . We successively apply Lemma K.3, triangle inequality and Lemma G.3 to obtain:

$$\alpha^{(t)} \left| y(\Lambda^{(t)} + (C-1)\Gamma^{(t)}) + \Xi^{(t)} \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} \delta_r \right| \leq 2\alpha^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}). \quad (182)$$

Thus, we use (181) and (182) to obtain a bound on the derivative.

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{G}(-yF(\mathbf{X})) \left(D\alpha^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^{p-1} \Xi^{(t)}. \quad (183)$$

3. $\ell \neq \ell(\mathbf{X})$ and $\delta_s = 0$ for all s and $m = \ell(\mathbf{X})$. We first bound the outside sum. Since $\Lambda^{(t)} + (C-1)\Gamma^{(t)} + (D-C)\Xi^{(t)} = 1$, we have:

$$\Xi^{(t)} \left| y \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} \right) + \Xi^{(t)} \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} (y - \delta_r) \right| \leq \Xi^{(t)} \left(\Xi^{(t)} \kappa(\mathbf{X}) + \mathbf{1}_{m=\ell(\mathbf{X})} \right). \quad (184)$$

We now bound the sum inside σ' .

$$\Xi^{(t)} \left| y + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} \delta_r(\mathbf{X}) \right| \leq \Xi^{(t)} \kappa(\mathbf{X}). \quad (185)$$

Using (184) and (185), we obtain a bound on the derivative.

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{G}(-yF(\mathbf{X})) \left(D\alpha^{(t)} \kappa(\mathbf{X}) \Xi^{(t)} \right)^{p-1} \Xi^{(t)} \left(\Xi^{(t)} \kappa(\mathbf{X}) + \mathbf{1}_{m=\ell(\mathbf{X})} \right). \quad (186)$$

4. $\ell \neq \ell(\mathbf{X})$ and $\delta_s \neq 0$ for some s and $m = \ell(\mathbf{X})$. We first bound the outside sum. We apply $\Lambda^{(t)} + (C-1)\Gamma^{(t)} + (D-C)\Xi^{(t)} = 1$ and [Lemma G.4](#) to get:

$$\begin{aligned} & \left| \Xi^{(t)} \left[\Lambda^{(t)}(y - \delta_i) + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} (y - \delta_r) + \Xi^{(t)} \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} (y - \delta_r) \right] \right| \\ & \leq \Xi^{(t)} \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} |\delta_r| + \mathbf{1}_{m=\ell(\mathbf{X})} \right) \\ & \leq \Theta(\Xi^{(t)}). \end{aligned} \quad (187)$$

We now bound the sum inside σ' . We apply [Lemma G.3](#) to obtain:

$$\begin{aligned} & \left| \Lambda^{(t)} \delta_i(X) + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} \delta_r(X) + \Xi^{(t)} \left(Cy + \sum_{h \in [L] \setminus \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} \delta_{h,r}(X) \right) \right| \\ & \leq \Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \kappa(\mathbf{X}) \\ & \leq 2 \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} \right). \end{aligned} \quad (188)$$

We combine (187) and (188) and obtain:

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^{p-1} \Xi^{(t)}. \quad (189)$$

5. $\ell, m \neq \ell(\mathbf{X})$ and $\delta_s = 0$ for all s and $\delta_j = 0$. We first bound the outside sum.

$$\Xi^{(t)} \left| Cy + \sum_{h \neq \{\ell, \ell(\mathbf{X})\}} \sum_{r \in \mathcal{S}_h} \delta_r \right| \leq \Xi^{(t)} \kappa(\mathbf{X}). \quad (190)$$

We now bound the sum inside σ' . This sum is actually equal to the outside sum outside and we can therefore use the bound (190). Thus, the derivative is bounded as:

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)} \kappa(\mathbf{X}) \Xi^{(t)} \right)^{p-1} (\Xi^{(t)})^2 \kappa(\mathbf{X}). \quad (191)$$

6. $\ell, m \neq \ell(\mathbf{X})$ and $\delta_s \neq 0$ for some s and $\delta_j = 0$. We first bound the outside sum. We apply [Lemma G.3](#) and obtain:

$$\begin{aligned} & \alpha^{(t)} \left| \Lambda^{(t)} \delta_i + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} \delta_r + \Xi^{(t)} \left[Cy + \sum_{h \neq \{\ell, \ell(\mathbf{X})\}} \sum_{r \in \mathcal{S}_h} \delta_r \right] \right| \\ & \leq \alpha^{(t)} \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \left[C + \sum_{h \neq \{\ell, \ell(\mathbf{X})\}} \sum_{r \in \mathcal{S}_h} |\delta_r| \right] \right) \\ & \leq 2\alpha^{(t)} \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} \right). \end{aligned} \quad (192)$$

We now bound the sum inside the power term. This sum is actually equal to the sum outside the power term and we can therefore use the bound (192). Thus, the derivative is bounded as:

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \right)^{p-1} \Xi^{(t)} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}). \quad (193)$$

7. $\ell, m \neq \ell(\mathbf{X})$ and $\delta_s = 0$ for all s and $\delta_j \neq 0$. We first bound the outside sum.

$$\begin{aligned} & \left| \Xi^{(t)} \left[\Lambda^{(t)} \delta_j + (C-1)\Gamma^{(t)} \delta_j + \Xi^{(t)} \left[C(\delta_j - y) + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} (\delta_j - \delta_r) \right] \right] \right| \\ & \leq \Xi^{(t)} \left(\Xi^{(t)} \kappa(\mathbf{X}) + \mathbf{1}_{\delta_j \neq 0} \right). \end{aligned} \quad (194)$$

We now bound the sum inside σ' .

$$\Xi^{(t)} \left| Cy + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} \delta_r \right| \leq \Xi^{(t)} \kappa(\mathbf{X}). \quad (195)$$

Using (194) and (195), the bound on the derivative is:

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)} \Xi^{(t)} \kappa(\mathbf{X}) \right)^{p-1} \Xi^{(t)} \left(\Xi^{(t)} \kappa(\mathbf{X}) + \mathbf{1}_{\delta_j \neq 0} \right). \quad (196)$$

8. $\ell, m \neq \ell(\mathbf{X})$ and $\delta_s(\mathbf{X}) \neq 0$ for some s and $\delta_j \neq 0$: We first bound the outside sum. We apply Lemma G.4 to get:

$$\begin{aligned} & \left| \Xi^{(t)} \left[\Lambda^{(t)} (\delta_j - \delta_i) + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} (\delta_j - \delta_r) + \Xi^{(t)} \left[C(\delta_j - y) + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} (\delta_j - \delta_r) \right] \right] \right| \\ & \leq \Xi^{(t)} \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \kappa(\mathbf{X}) + \mathbf{1}_{\delta_j \neq 0} \right) \\ & \leq \Theta(\Xi^{(t)}). \end{aligned} \quad (197)$$

We now bound the sum inside the power term. We apply Lemma G.3 and get:

$$\begin{aligned} \left| \Lambda^{(t)} \delta_i + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} \delta_r + \Xi^{(t)} \left(Cy + \sum_{h \neq \{\ell(\mathbf{X}), \ell\}} \sum_{r \in \mathcal{S}_h} \delta_r \right) \right| & \leq \Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \kappa(\mathbf{X}) \\ & \leq \Theta \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} \right). \end{aligned} \quad (198)$$

We plug (197) and (198) to obtain the derivative.

$$\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \leq \Theta(D\alpha^{(t)}) \mathfrak{S}(-yF(\mathbf{X})) \left(D\alpha^{(t)} \left(\Lambda^{(t)} + (C-1)\Gamma^{(t)} \right) \right)^{p-1} \Xi^{(t)}. \quad (199)$$

Putting all the pieces together. We now bound the derivative of the population loss. Using Tower property and Lemma I.1, we have:

$$\mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] = \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_{\pi_1, \pi_2} \left[\frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \middle| \mathbf{X} \right] \right]. \quad (200)$$

For a fixed \mathbf{X} , we now bound the derivative of the loss evaluated in \mathbf{X} . For $i \in \mathcal{S}_\ell$ and $j \in \mathcal{S}_m$, we distinguish the eight possible events depending on the randomness of π .

- **Event a:** " $\ell = \ell(\pi(\mathbf{X}))$ and $\delta_j = 0$ " occurs with probability $(1-q)/L$.
- **Event b:** " $\ell = \ell(\pi(\mathbf{X}))$ and $\delta_j \neq 0$ " occurs with probability q/L .
- **Event c:** " $\ell \neq \ell(\pi(\mathbf{X}))$ and $\delta_s = 0$ for all s and $m = \ell(\pi(\mathbf{X}))$ " occurs with probability $(1-q)^C/L$.
- **Event \mathbf{d}_k :** " $\ell \neq \ell(\pi(\mathbf{X}))$ and $\#\{s : \delta_s \neq 0\} = k$ for $1 \leq k \leq C$ and $m = \ell(\pi(\mathbf{X}))$ " occurs with probability $\binom{C}{k} q^k (1-q)^{C-k}/L$.
- **Event e:** " $\ell, m \neq \ell(\pi(\mathbf{X}))$ and $\delta_j = 0$ and $\delta_s = 0$ for all s " occurs with probability $(1-2/L)(1-q)^{C+1}$.
- **Event \mathbf{f}_k :** " $\ell, m \neq \ell(\pi(\mathbf{X}))$ and $\#\{s : \delta_s \neq 0\} = k$ for $1 \leq k \leq C$ and $\delta_j = 0$ " occurs with probability $(1-2/L)(1-q) \binom{C}{k} q^k (1-q)^{C-k}$.
- **Event g:** " $\ell, m \neq \ell(\pi(\mathbf{X}))$ and $\delta_s = 0$ for all s and $\delta_j \neq 0$ " occurs with probability $(1-2/L)(1-q)^C q$.
- **Event \mathbf{h}_k :** " $\ell, m \neq \ell(\pi(\mathbf{X}))$ and $\#\{s : \delta_s \neq 0\} = k$ for $0 \leq k \leq C$ and $\delta_j \neq 0$ " occurs with probability $(1-2/L) q \binom{C}{k} q^k (1-q)^{C-k}$.

Since events a and e are the ones with highest probabilities, the derivative of the loss is bounded by the expectations conditioned on these events. We have:

$$\begin{aligned} & \mathbb{E}_{\pi_1, \pi_2} \left[\left| \frac{\partial L(\pi(\mathbf{X}))}{\partial A_{i,j}} \right| \right] \\ & \leq \frac{\Theta(D\alpha^{(t)})}{L} \Xi^{(t)} \mathfrak{S}(-yF(\mathbf{X})) (D\alpha^{(t)}(\Lambda^{(t)} + (C-1)\Gamma^{(t)}))^{p-1} (\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \quad (201) \\ & + \Theta(D\alpha^{(t)}) \Xi^{(t)} \mathfrak{S}(-yF(\mathbf{X})) \mathbb{E}_{\pi_1, \pi_2} \left[(D\alpha^{(t)} \kappa(\pi(\mathbf{X})) \Xi^{(t)})^{p-1} \Xi^{(t)} \kappa(\pi(\mathbf{X})) \right] \mathbf{e} \Big]. \end{aligned}$$

We now apply [Induction Hypothesis B.1](#) and [Lemma K.3](#) and finally obtain:

$$\begin{aligned} \mathbb{E}_{\mathbf{X}} \left[\left| \frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right| \right] & \leq \Theta(C^2 \alpha^{(t)}) \mathbb{E}_{\mathbf{X}} [\mathfrak{S}(-yF(\mathbf{X}))] (D\alpha^{(t)}(\Lambda^{(t)} + (C-1)\Gamma^{(t)}))^{p-1} \Xi^{(t)} \frac{\lambda_0}{D} \quad (202) \\ & + \Theta(\alpha^{(t)}) \mathbb{E}_{\mathbf{X}} [\mathfrak{S}(-yF(\mathbf{X}))] (D\alpha^{(t)} q D \log(d) \Xi^{(t)})^{p-1} \Xi^{(t)} q D \log(d). \end{aligned}$$

Since $\mathbb{E}_{\mathbf{X}} [\mathfrak{S}(-yF(\mathbf{X}))] \leq 1$, we thus plug (202) in the update of $\rho^{(t)}$ to obtain the aimed result. \square

Corollary G.2. *Let $T > 0$ be the time where the population loss is $o(1)$ and $t \leq T$. The update of $\rho^{(t)}$ satisfies:*

$$|\rho^{(t+1)}| \leq |\rho^{(t)}| + \eta \cdot \text{polylog}(d) (\alpha^{(t)})^p \left(\frac{1}{D} + \frac{\lambda_0}{D} \Xi^{(t)} (G^{(t)})^{p-1} \right).$$

Proof of Corollary G.2. Using [Lemma G.2](#), $C = \text{polylog}(d)$ and $qD \leq 1$, $\rho^{(t)}$'s update is:

$$|\rho^{(t+1)}| \leq |\rho^{(t)}| + \eta (\alpha^{(t)})^p \text{polylog}(d) \left((D\Xi^{(t)})^{p-1} \Xi^{(t)} + \frac{\lambda_0}{D} (G^{(t)})^{p-1} \Xi^{(t)} \right). \quad (203)$$

Lastly, we apply [Induction Hypothesis B.1](#) to replace $\Xi^{(t)}$ by its value in (203) and thus obtain the aimed result. \square

G.3 Auxiliary lemmas

Lemma G.3. *Let $t > 0$. In the idealized learning process, with high probability, we have:*

$$\Lambda^{(t)} + (C-1)\Gamma^{(t)} + y\Xi^{(t)} \sum_{\ell \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_\ell} \delta_r = \Theta(\Lambda^{(t)} + (C-1)\Gamma^{(t)}).$$

Proof of Lemma G.3. We first bound the sum with factor $\Xi^{(t)}$. Using [Induction Hypothesis B.1](#) and [Lemma K.3](#), we have:

$$\left| y\Xi^{(t)} \sum_{\ell \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_\ell} \delta_r \right| \leq \Theta(q) \log(d) = \frac{1}{D} \text{poly}(C) \log(d). \quad (204)$$

We now bound $\Lambda^{(t)} + (C-1)\Gamma^{(t)}$. Using [Induction Hypothesis B.1](#), we have:

$$\frac{1}{D} \left(e^{\text{polyloglog}(d)} + \Omega(C) \right) \leq (\Lambda^{(t)} + (C-1)\Gamma^{(t)}) \leq \frac{1}{D} \left(e^{\text{polyloglog}(d)} + \lambda_0 \right). \quad (205)$$

Since $\text{poly}(C) \log(d) \leq e^{\text{polyloglog}(d)}/2$, we combine (204) and (205) to get the aimed result. \square

Lemma G.4. *Let $t > 0$. In the idealized learning process, we have with high probability:*

$$\Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} |\delta_r| \leq \Theta(1).$$

Proof of Lemma G.4. We successively apply Lemma K.3 and Induction Hypothesis B.1 to get the desired bound. Indeed, we have:

$$\begin{aligned} \Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} |\delta_r| &\leq \Lambda^{(t)} + (C-1)\Gamma^{(t)} + \Xi^{(t)} qD \log(d) \\ &\leq \frac{1}{D} \left(e^{\text{polyloglog}(d)} + \lambda_0 + qD \log(d) \right) \\ &= \Theta(1). \end{aligned}$$

□

H Gradients

In this section, we present the gradients of the loss \mathcal{L} with respect to \mathbf{v} and $\mathbf{A}_{i,j}$.

Lemma H.1. *Let (\mathbf{X}, y) be a data-point. Then, the gradient of $L(\mathbf{X})$ with respect to \mathbf{v} is:*

$$-\nabla_{\mathbf{v}} L(\mathbf{X}) = Dy \mathfrak{S}(-yF(\mathbf{X})) \sum_{i=1}^D \sigma' \left(D \sum_{r \in [D]} \mathbf{S}_{i,r} \langle \mathbf{v}, \mathbf{X}_r \rangle \right) \sum_{j=1}^D \mathbf{S}_{i,j} \mathbf{X}_j.$$

Lemma H.2. *Let (\mathbf{X}, y) be a data-point and $i, j \in [D]$. The derivative of $L(\mathbf{X})$ with respect to $A_{i,j}$ is:*

$$-\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} = pDy \mathfrak{S}(-yF(\mathbf{X})) \left(D \sum_{r \in [D]} \mathbf{S}_{i,r} \langle \mathbf{v}, \mathbf{X}_r \rangle \right)^{p-1} \sum_{r \neq j} \mathbf{S}_{i,r} \langle \mathbf{v}, \mathbf{X}_j - \mathbf{X}_r \rangle.$$

I Invariance of the problem

I.1 Invariance of the parameters

Lemma 4.2. *For all $t \in [T]$, there exists $\alpha^{(t)} \in \mathbb{R}$ such that $\mathbf{v}^{(t)} = \alpha^{(t)} \mathbf{w}^*$.*

Proof of Lemma 4.2. The proof is by induction. Our induction hypothesis is $\mathbf{v}^{(t)} = \alpha^{(t)} \mathbf{w}^*$ for all $t \geq 0$. For $t = 0$, we know that $\mathbf{v}^{(0)} = \alpha^{(0)} \mathbf{w}^* \in \text{span}(\mathbf{w}^*)$. Assume that $\mathbf{v}^{(t)} = \alpha^{(t)} \mathbf{w}^*$. Let's show that there exists $\alpha^{(t+1)} \in \mathbb{R}$ such that $\mathbf{v}^{(t+1)} = \alpha^{(t+1)} \mathbf{w}^*$. From the update rule, we have:

$$\begin{aligned} \mathbf{v}^{(t+1)} &= \mathbf{v}^{(t)} + \eta \mathbb{E} \left[y \mathfrak{S}(-yF(\mathbf{X})) \sum_{i \in [D]} \sigma'(\langle \mathbf{v}^{(t)}, \mathbf{O}_i^{(t)} \rangle) \sum_{r \in [D]} \mathbf{S}_{i,r} \mathbf{X}_r \right] \\ &= \mathbf{v}^{(t)} + \eta \mathbb{E} \left[y \mathfrak{S}(-yF(\mathbf{X})) \sum_{i \in [D]} \sigma'(\langle \mathbf{v}^{(t)}, \mathbf{O}_i^{(t)} \rangle) \sum_{r \in [D]} \mathbf{S}_{i,r} \mathbb{E}_{\xi}[\mathbf{X}_r] \right], \end{aligned} \quad (206)$$

where \mathbb{E}_{ξ} is the expectation with respect to noise vectors ξ_r . Using the definition of the data distribution and $\mathbb{E}[\xi_r] = 0$, we simplify the update (206):

$$\begin{aligned} \mathbf{v}^{(t+1)} &= \mathbf{v}^{(t)} + \eta \mathbb{E} \left[\mathfrak{S}(-yF(\mathbf{X})) \sum_{i \in [D]} \sigma'(\langle \mathbf{v}^{(t)}, \mathbf{O}_i^{(t)} \rangle) \sum_{r \in \mathcal{S}_{\ell(\mathbf{X})}} \mathbf{S}_{i,r} \right] \mathbf{w}^* \\ &\quad + \eta \mathbb{E} \left[y \mathfrak{S}(-yF(\mathbf{X})) \sum_{i \in [D]} \sigma'(\langle \mathbf{v}^{(t)}, \mathbf{O}_i^{(t)} \rangle) \sum_{h \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_h} \mathbf{S}_{i,r} \delta_r \right] \mathbf{w}^*. \end{aligned} \quad (207)$$

Since $\mathbf{v}^{(t)} = \alpha^{(t)} \mathbf{w}^*$, there exists $\alpha^{(t+1)} \in \mathbb{R}$ such that $\mathbf{v}^{(t+1)} = \alpha^{(t+1)} \mathbf{w}^*$.

□

Lemma 4.1. *There exist $\beta = \sigma_{\mathbf{A}}$, $\gamma^{(t)}$, $\rho^{(t)} \in \mathbb{R}$ such that for all $t \geq 0$:*

1. for all $i \in [D]$, $A_{i,i}^{(t)} = \beta$.
2. for all $i, j \in [D]$ such that $i, j \in \mathcal{S}_{\ell}$ for some $\ell \in [L]$, $A_{i,j}^{(t)} = \gamma^{(t)}$.

3. for all $i, j \in [D]$ such that $i \in \mathcal{S}_\ell$ and $j \in \mathcal{S}_m$ for some $\ell, m \in [L]$ with $\ell \neq m$, $A_{i,j}^{(t)} = \rho^{(t)}$.

Proof of Lemma 4.1. We initialize $A_{i,i}^{(0)} = \sigma_{\mathbf{A}}$ and do not update $A_{i,i}$, Thus, for all $t \geq 0$, we have $A_{i,i}^{(t)} = \sigma_{\mathbf{A}}$.

The remaining of the proof is by induction. For $i, j \in \mathcal{S}_\ell$, we initialize $A_{i,j}^{(0)} = 0 = \gamma^{(0)}$. For $i \in \mathcal{S}_\ell$ and $j \in \mathcal{S}_m$ with $\ell \neq m$, $A_{i,j}^{(0)} = 0 = \rho^{(0)}$. The induction hypothesis is true for $t = 0$.

We assume that for all $i, j \in \mathcal{S}_\ell$, $A_{i,j}^{(t)} = \gamma^{(t)}$ and for $i \in \mathcal{S}_\ell$ and $j \in \mathcal{S}_m$, $A_{i,j}^{(t)} = \rho^{(t)}$. Let's first prove that $A_{i,j}^{(t+1)} = A_{k,n}^{(t+1)}$ for $i, j \in \mathcal{S}_\ell$ and $k, n \in \mathcal{S}_{\ell'}$. Since the GD update is $A_{i,j}^{(t+1)} = A_{i,j}^{(t)} - \eta \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right]$, it's sufficient to prove that $\mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] = \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{k,n}} \right]$. Let $\pi_1: [L] \rightarrow [L]$, $\pi_2: [C] \rightarrow [C]$ and $\pi = (\pi_1, \pi_2)$ be permutations. From Lemma I.1, we know that \mathbf{X} and $\pi(\mathbf{X})$ have the same distribution which implies $\mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] = \mathbb{E}_{\pi(\mathbf{X})} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right]$. Therefore, we have:

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] \\ &= \mathbb{E}_{\mathbf{X}} \left[y \mathfrak{S}(-yF(\pi(\mathbf{X}))) \sigma' \left(\sum_{h=1}^L \sum_{r \in \mathcal{S}_h} \frac{e^{\langle \mathbf{A}^{(t)} \mathbf{p}_i, \mathbf{p}_r \rangle}}{\sum_{s=1}^D e^{\langle \mathbf{A}^{(t)} \mathbf{p}_i, \mathbf{p}_s \rangle}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(h), \pi_2(r)} \rangle \right) \right. \\ & \quad \left. \frac{e^{\langle \mathbf{A}^{(t)} \mathbf{p}_i, \mathbf{p}_j \rangle}}{\sum_{s=1}^D e^{\langle \mathbf{A}^{(t)} \mathbf{p}_i, \mathbf{p}_s \rangle}} \sum_{k=1}^L \sum_{m \in \mathcal{S}_k} \frac{e^{\langle \mathbf{A}^{(t)} \mathbf{p}_i, \mathbf{p}_m \rangle}}{\sum_{s=1}^D e^{\langle \mathbf{A}^{(t)} \mathbf{p}_i, \mathbf{p}_s \rangle}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(k), \pi_2(m)} - \mathbf{X}_{\pi_1(\ell), \pi_2(j)} \rangle \right] \end{aligned} \quad (208)$$

Using Induction Hypothesis B.1, we simplify (208) as

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}} \left[\frac{\partial \mathcal{L}^{(t)}(\mathbf{X})}{\partial A_{i,j}} \right] \\ &= \mathbb{E}_{\mathbf{X}} \left[y \mathfrak{S}(-yF(\pi(\mathbf{X}))) \sigma' \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(i)} \rangle + \Gamma^{(t)} \sum_{m \in \mathcal{S}_\ell \setminus \{i\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(m)} \rangle \right) \right. \\ & \quad + \Xi^{(t)} \sum_{r \notin \mathcal{S}_\ell} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi(r)} \rangle \left. \Gamma^{(t)} \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(i)} - \mathbf{X}_{\pi_1(\ell), \pi_2(j)} \rangle \right) \right. \\ & \quad \left. + \Gamma^{(t)} \sum_{m \in \mathcal{S}_\ell \setminus \{i\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(m)} - \mathbf{X}_{\pi_1(\ell), \pi_2(j)} \rangle + \Xi^{(t)} \sum_{r \notin \mathcal{S}_\ell} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi(r)} - \mathbf{X}_{\pi_1(\ell), \pi_2(j)} \rangle \right) \right]. \end{aligned} \quad (209)$$

We now set π_1 and π_2 such that $\pi_1(\ell) = \ell'$, $\pi_2(i) = k$ and $\pi_2(j) = n$. Using this choice along with $F(\mathbf{X}) = F(\pi(\mathbf{X}))$ (Lemma I.1), we finally have in (209)

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] \\ &= \mathbb{E}_{\mathbf{X}} \left[y \mathfrak{S}(-yF(\pi(\mathbf{X}))) \sigma' \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', k} \rangle + \Gamma^{(t)} \sum_{m \in \mathcal{S}_{\ell'} \setminus \{i\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', m} \rangle \right) \right. \\ & \quad + \Xi^{(t)} \sum_{r \notin \mathcal{S}_{\ell'}} \langle \mathbf{v}^{(t)}, \mathbf{X}_r \rangle \left. \Gamma^{(t)} \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', k} - \mathbf{X}_{\ell', n} \rangle \right) \right. \\ & \quad \left. + \Gamma^{(t)} \sum_{m \in \mathcal{S}_{\ell'} \setminus \{k\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', m} - \mathbf{X}_{\ell', n} \rangle + \Xi^{(t)} \sum_{r \notin \mathcal{S}_{\ell'}} \langle \mathbf{v}^{(t)}, \mathbf{X}_r - \mathbf{X}_{\ell', n} \rangle \right) \right] \\ &= \mathbb{E}_{\mathbf{X}} \left[\frac{\partial \mathcal{L}^{(t)}(\mathbf{X})}{\partial A_{k,n}} \right]. \end{aligned} \quad (210)$$

Therefore, (210) implies that $A_{i,j}^{(t+1)} = A_{k,n}^{(t+1)}$ thus proving the induction hypothesis. Let's now show that for $i \in \mathcal{S}_\ell, j \in \mathcal{S}_m$ and $k \in \mathcal{S}_{\ell'}, j \in \mathcal{S}_{m'}$, we have $A_{i,j}^{(t+1)} = A_{k,n}^{(t+1)}$. We apply a similar

argument as above. Using [Induction Hypothesis B.1](#), we have:

$$\begin{aligned}
& \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] \\
&= \mathbb{E}_{\mathbf{X}} \left[y \mathfrak{S}(-yF(\pi(\mathbf{X}))) \sigma' \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(i)} \rangle + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(r)} \rangle \right. \right. \\
&\quad \left. \left. + \Xi^{(t)} \sum_{r \notin \mathcal{S}_\ell} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi(r)} \rangle \right) \Xi^{(t)} \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(i)} - \mathbf{X}_{\pi_1(m), \pi_2(j)} \rangle \right. \right. \\
&\quad \left. \left. + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{i\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_1(\ell), \pi_2(r)} - \mathbf{X}_{\pi_1(m), \pi_2(j)} \rangle + \Xi^{(t)} \sum_{r \notin \mathcal{S}_\ell} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi(r)} - \mathbf{X}_{\pi_1(m), \pi_2(j)} \rangle \right) \right]. \tag{211}
\end{aligned}$$

We now set π_1 and π_2 such that $\pi_1(\ell) = \ell'$, $\pi_1(m) = m'$, $\pi_2(i) = k$ and $\pi_2(j) = n$. Using this choice, we finally have in (211)

$$\begin{aligned}
& \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{i,j}} \right] \\
&= \mathbb{E}_{\mathbf{X}} \left[y \mathfrak{S}(-yF(\pi(\mathbf{X}))) \sigma' \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', k} \rangle + \Gamma^{(t)} \sum_{r \in \mathcal{S}_{\ell'} \setminus \{k\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', r} \rangle \right. \right. \\
&\quad \left. \left. + \Xi^{(t)} \sum_{r \notin \mathcal{S}_{\ell'}} \langle \mathbf{v}^{(t)}, \mathbf{X}_r \rangle \right) \Xi^{(t)} \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', k} - \mathbf{X}_{m', n} \rangle \right. \right. \\
&\quad \left. \left. + \Gamma^{(t)} \sum_{r \in \mathcal{S}_{\ell'} \setminus \{k\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\ell', r} - \mathbf{X}_{m', n} \rangle + \Xi^{(t)} \sum_{r \notin \mathcal{S}_{\ell'}} \langle \mathbf{v}^{(t)}, \mathbf{X}_r - \mathbf{X}_{m', n} \rangle \right) \right] \\
&= \mathbb{E}_{\mathbf{X}} \left[\frac{\partial L(\mathbf{X})}{\partial A_{k,n}} \right]. \tag{212}
\end{aligned}$$

Therefore, (212) implies that $A_{i,j}^{(t+1)} = A_{k,n}^{(t+1)}$ thus proving the induction hypothesis. \square

I.2 Invariance by permutation

Lemma I.1. *Let $\pi_1: [L] \rightarrow [L]$ and $\pi_2: [C] \rightarrow [C]$ be two permutations and $\pi = (\pi_1, \pi_2)$. Let $(\mathbf{X}, \cdot) \sim \mathcal{D}$. Then, we have:*

1. *permutation-invariant distribution: \mathbf{X} has the same distribution as $\pi(\mathbf{X})$.*
2. *permutation-invariant model: $F(\mathbf{X}) = F(\pi(\mathbf{X}))$.*

Proof of Lemma I.1. To show that \mathbf{X} and $\pi(\mathbf{X})$ have to same distribution, it is sufficient to show that the items 1 to 6 hold in our definition of the data distribution. We still have that the label y is uniformly sampled on $\{-1, 1\}$. Let $\pi(\mathbf{X}) = (\bar{\mathbf{X}}_1, \dots, \bar{\mathbf{X}}_D)$ where $\bar{\mathbf{X}}_i$ for some j . For 2, the number of tokens is still D after permutation π . For 4, we define the same partition with $\bar{\mathcal{S}}_i = \bar{\mathcal{S}}_{\pi_1(i)}$. Besides, we have $\ell(\bar{\mathbf{X}}) = \pi_1(\ell(\mathbf{X}))$ that is also uniformly sampled on $[L]$ since the permutation on a uniform distribution is also uniform. For $i \in \mathcal{S}_\ell(\bar{\mathbf{X}})$, we have that $\bar{\mathbf{X}}_i$ writes $\mathbf{X}_{\pi_2(k)}$ for some $k \in \mathcal{S}_\ell(\mathbf{X})$ so that we do have $\bar{\mathbf{X}}_i = y\mathbf{w}^* + \xi_i$. The same goes for 6 when $\ell \neq \ell(\bar{\mathbf{X}})$: with $\ell \neq \ell(\bar{\mathbf{X}})$, $\bar{\mathbf{X}}_j = \delta_{\ell,j} y\mathbf{w}^* + \xi_j$, where $\delta_j = 1$ with probability $q/2$, -1 with the same probability and 0 otherwise.

Let \mathbf{X} be a data-point. Using [Lemma 4.1](#), we rewrite $F(\mathbf{X})$ as

$$F(\mathbf{X}) = \sum_{m=1}^D \sigma \left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_m \rangle + \Gamma^{(t)} \sum_{r \in \mathcal{S}_\ell \setminus \{m\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_r \rangle + \Xi^{(t)} \sum_{h \in [L] \setminus \{\ell\}} \sum_{j \in \mathcal{S}_h} \langle \mathbf{v}^{(t)}, \mathbf{X}_j \rangle \right). \tag{213}$$

Let $\pi = (\pi_1, \pi_2)$ be a permutation. For a given $\ell \in [L]$ and $m \in \mathcal{S}_\ell$, assume that $\pi_1(\ell) = \ell$ and $\pi_2(m') = m$. Using [Lemma 4.1](#), we have:

$$\begin{aligned}
F(\pi(\mathbf{X})) &= \sum_{\ell=1}^L \sum_{m \in \mathcal{S}_\ell} \sigma\left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_2(m')} \rangle\right) \\
&\quad + \Gamma^{(t)} \sum_{r \in \mathcal{S}_{\pi_1(\ell')} \setminus \{\pi_2(m')\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_2(r)} \rangle + \Xi^{(t)} \sum_{h' \in [L] \setminus \{\ell'\}} \sum_{j \in \mathcal{S}_{h'}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{\pi_2(j')} \rangle \\
&= \sum_{\ell=1}^L \sum_{m \in \mathcal{S}_\ell} \sigma\left(\Lambda^{(t)} \langle \mathbf{v}^{(t)}, \mathbf{X}_m \rangle\right) \\
&\quad + \Gamma^{(t)} \sum_{r'' \in \mathcal{S}_\ell \setminus \{m\}} \langle \mathbf{v}^{(t)}, \mathbf{X}_{r''} \rangle + \Xi^{(t)} \sum_{j'' \notin \mathcal{S}_\ell} \langle \mathbf{v}^{(t)}, \mathbf{X}_{j''} \rangle \\
&= F(\mathbf{X}).
\end{aligned}$$

□

J Justification of our data distribution

In this section, we justify why the distribution \mathcal{D} ([Assumption 1](#)) is relevant. We first show that linear classifiers poorly generalize ([subsection J.1](#)). We then show that there exists classifiers that generalize without learning patch association ([subsection J.2](#)).

J.1 Generalized linear models poorly generalize

Theorem 3.1. *Let \mathcal{D} be as in [Assumption 1](#). Let $g(\mathbf{X}) = \phi\left(\sum_{j=1}^D \langle \mathbf{w}_j, \mathbf{X}_j \rangle\right)$ be a generalized linear model. Then, g does not fit the labeling function i.e. $\mathbb{P}[f^*(\mathbf{X})g(\mathbf{X}) \leq 0] \geq 1/8$.*

Proof of [Theorem 3.1](#). For every data point \mathbf{X} , consider $\Delta(\mathbf{X}) := \sum_{j \in [D]} \delta_j$, it is very easy to see that for every integer p , as long as $\Pr[\Delta = p] = \Omega(1/\text{polylog}(d))$, we have that:

$$\frac{\Pr[\Delta = p]}{\Pr[\Delta = p - 2C]} = 1 - o(1) \quad (214)$$

Consider two independently sampled data points, \mathbf{X}, \mathbf{X}' with label 1, -1 respectively, consider the event when $\Delta(\mathbf{X}) = p - 2C, \Delta(\mathbf{X}') = p$ and all the noises ξ_i, ξ'_i of \mathbf{X}, \mathbf{X}' satisfies $\xi_i = \xi'_i$, then we know that

$$\sum_{j \in [D]} \langle \mathbf{w}_j, \mathbf{X}_j \rangle = \sum_{j \in [D]} \langle \mathbf{w}_j, \mathbf{X}'_j \rangle \quad (215)$$

By Eq (214) we also know that the density of \mathbf{X} and \mathbf{X}' under the data-generation distribution satisfies

$$p(\mathbf{X}) = (1 \pm o(1))p(\mathbf{X}')$$

Now, we know that

$$P_0 := \mathbb{P}[f^*(\mathbf{X})g(\mathbf{X}) \leq 0 \mid y(\mathbf{X}) = 1] = 2 \int_{\mathbf{X}} 1_{f^*(\mathbf{X})g(\mathbf{X}) \leq 0} p(\mathbf{X}, y(\mathbf{X}) = 1) d\mathbf{X} \quad (216)$$

$$= 2 \int_{\mathbf{X}} 1_{f^*(\mathbf{X})g(\mathbf{X}) \leq 0} p(\mathbf{X}', y(\mathbf{X}') = -1) d\mathbf{X} \pm o(1) \quad (217)$$

$$= 2 \int_{\mathbf{X}} 1_{f^*(\mathbf{X})g(\mathbf{X}') \leq 0} p(\mathbf{X}', y(\mathbf{X}') = -1) d\mathbf{X} \pm o(1) \quad (218)$$

$$= 2 \int_{\mathbf{X}'} 1_{f^*(\mathbf{X}')g(\mathbf{X}') \geq 0} p(\mathbf{X}', y(\mathbf{X}') = -1) d\mathbf{X}' \pm o(1) \quad (219)$$

$$= 1 - \Pr[f^*(\mathbf{X}')g(\mathbf{X}') < 0 \mid y(\mathbf{X}') = -1] \pm o(1) \quad (220)$$

$$\geq 1 - \Pr[f^*(\mathbf{X}')g(\mathbf{X}') \leq 0 \mid y(\mathbf{X}') = -1] \pm o(1) \quad (221)$$

$$= 1 - \Pr[f^*(\mathbf{X})g(\mathbf{X}) \leq 0 \mid y(\mathbf{X}) = -1] \pm o(1) \quad (222)$$

Therefore, $\mathbb{P}[f^*(\mathbf{X})g(\mathbf{X}) \leq 0 \mid y(\mathbf{X}) = 1] \geq \frac{1}{2} - o(1)$. \square

J.2 Classifiers fitting the labelling function without patch association

Theorem 3.2. *Let \mathcal{D} be defined as in [Assumption 1](#). There exists a (one-layer) transformer \mathcal{M} so that $\mathbb{P}[f^*(\mathbf{X})\mathcal{M}(\mathbf{X}) \leq 0] = d^{-\omega(1)}$ but for all $\ell \in [L]$, $i \in \mathcal{S}_\ell$, $\text{Top}_C \{ \langle \mathbf{p}_i^{(\mathcal{M})}, \mathbf{p}_j^{(\mathcal{M})} \rangle \}_{j=1}^D \cap \mathcal{S}_\ell = \emptyset$.*

Proof of [Theorem 3.2](#). We can consider a transformer in our setting, whose weights are defined as: $\mathbf{v} = \mathbf{w}^*$, $A_{i,j} = \beta > 0$ for $i, j \in \mathcal{S}_\ell$. $A_{i,j} = 2\beta$ for $i \in \mathcal{S}_\ell, j \in \mathcal{S}_{\ell+1}$ (We denote $\mathcal{S}_{L+1} = \mathcal{S}_1$). For a sufficiently large β , it is easy to check that $\mathbb{P}[f^*(\mathbf{X})\mathcal{M}(\mathbf{X}) \leq 0] = d^{-\omega(1)}$ but for all $\ell \in [L]$, $i \in \mathcal{S}_\ell$, $\text{Top}_C \{ \langle \mathbf{p}_i^{(\mathcal{M})}, \mathbf{p}_j^{(\mathcal{M})} \rangle \}_{j=1}^D = \mathcal{S}_{\ell+1} \cap \mathcal{S}_\ell = \emptyset$. \square

K Technical lemmas

In this section, we present the technical lemmas used in the paper.

K.1 Tensor Power Method

Lemma K.1. *Let $\{z^{(t)}\}_{t \geq 0}$ be a positive sequence defined by the following recursions*

$$\begin{cases} z^{(t+1)} \geq z^{(t)} + m(z^{(t)})^k \\ z^{(t+1)} \leq z^{(t)} + M(z^{(t)})^k \end{cases},$$

where $z^{(0)} > 0$ is the initialization, $k > 1$ is an integer and $m, M > 0$. Let $v > 0$ such that $z^{(0)} \leq v$. Then, the time \mathcal{T} such that $z^{(t)} \geq v$ for all $t \geq \mathcal{T}$ is:

$$\mathcal{T} = \frac{3}{m(z^{(0)})^{k-1}} + \frac{2^{k+1}M}{m} \left\lceil \frac{\log(v/z^{(0)})}{\log(2)} \right\rceil.$$

Proof of [Lemma K.1](#). Let $n \in \mathbb{N}^*$. Let T_n be the time where $z^{(t)} \geq 2^n z^{(0)}$. This time exists because $z^{(t)}$ is a non-decreasing sequence. We want to find an upper bound on this time. We start with the case $n = 1$. By summing the recursion, we have:

$$z^{(T_1)} \geq z^{(0)} + m \sum_{s=0}^{T_1-1} (z^{(s)})^k. \quad (223)$$

We use the fact that $z^{(s)} \geq z^{(0)}$ in [\(223\)](#) and obtain:

$$T_1 \leq \frac{z^{(T_1)} - z^{(0)}}{m(z^{(0)})^k}. \quad (224)$$

Now, we want to bound $z^{(T_1)} - z^{(0)}$. Using again the recursion and $z^{(T_1-1)} \leq 2z^{(0)}$, we have:

$$z^{(T_1)} \leq z^{(T_1-1)} + M(z^{(T_1-1)})^k \leq 2z^{(0)} + 2^k M(z^{(0)})^k. \quad (225)$$

Combining (224) and (225), we get a bound on T_1 .

$$T_1 \leq \frac{1}{m(z^{(0)})^{k-1}} + \frac{2^k M}{m}. \quad (226)$$

Now, let's find a bound for T_n . Starting from the recursion and using the fact that $z^{(s)} \geq 2^{n-1}z^{(0)}$ for $s \geq T_{n-1}$ we have:

$$z^{(T_n)} \geq z^{(T_{n-1})} + m \sum_{s=T_{n-1}}^{T_n-1} (z^{(s)})^k \geq z^{(T_{n-1})} + 2^{k(n-1)} m (z^{(0)})^k (T_n - T_{n-1}). \quad (227)$$

On the other hand, by using $z^{(T_{n-1})} \leq 2^n z^{(0)}$ we upper bound $z^{(T_n)}$ as follows.

$$z^{(T_n)} \leq z^{(T_{n-1})} + M(z^{(T_{n-1})})^k \leq 2^n z^{(0)} + 2^{kn} M(z^{(0)})^k. \quad (228)$$

Besides, we know that $z^{(T_{n-1})} \geq 2^{n-1}z^{(0)}$. Therefore, we upper bound $z^{(T_n)} - z^{(T_{n-1})}$ as

$$z^{(T_n)} - z^{(T_{n-1})} \leq 2^{n-1}z^{(0)} + 2^{kn} M(z^{(0)})^k. \quad (229)$$

Combining (227) and (229) yields:

$$T_n \leq T_{n-1} + \frac{1}{2^{(k-1)(n-1)} m (z^{(0)})^{k-1}} + \frac{2^k M}{m}. \quad (230)$$

We now sum (230) for $n = 2, \dots, n$, use (226) and obtain:

$$T_n \leq T_1 + \frac{2}{m(z^{(0)})^{k-1}} + \frac{2^k M n}{m} \leq \frac{3}{m(z^{(0)})^{k-1}} + \frac{2^k M(n+1)}{m} \leq \frac{3}{m(z^{(0)})^{k-1}} + \frac{2^{k+1} M n}{m}. \quad (231)$$

Lastly, we know that n satisfies $2^n z^{(0)} \geq v$ which implies $n = \left\lceil \frac{\log(v/z_0)}{\log(2)} \right\rceil$ in (231). \square

Lemma K.2. Let $\{z^{(t)}\}_{t \geq 0}$ be a positive sequence defined by the following recursions

$$\begin{cases} z^{(t+1)} \geq z^{(t)} + m(z^{(t)})^k \\ z^{(t+1)} \leq z^{(t)} + M(z^{(t)})^k \end{cases},$$

where $z^{(0)} > 0$, $k > 1$ is an integer and $m, M > 0$. Let $v > 0$ such that $z^{(0)} \leq v$ and \mathcal{T} be the time such that $z^{(t)} \geq v$ for all $t \geq \mathcal{T}$. Assume that $\frac{Av^2}{m} \ll 1$. Then, we have for $\kappa \in \{1, 2\}$:

$$\prod_{\tau=0}^{\mathcal{T}-1} (1 + A(z^{(\tau)})^{k-1+\kappa}) \leq (1 + v^{k-1})^{\frac{2^{\kappa+1} M A v^\kappa \log(v/z^{(0)})}{m}}.$$

Proof of Lemma K.2. Let $n \in \mathbb{N}^*$ and let T_n be the time such that $z^{(t)} \geq 2^n z^{(0)}$ for $t \geq T_n$. Starting from the recursion, we have:

$$\log(z^{(t+1)}) \geq \log(z^{(t)}) + \log(1 + m(z^{(t)})^{k-1}). \quad (232)$$

Since $z^{(t)}$ is a non-decreasing sequence, (232) satisfies:

$$(z^{(t+1)})^\kappa \log(z^{(t+1)}) \geq (z^{(t)})^\kappa \log(z^{(t)}) + (z^{(t)})^\kappa \log(1 + m(z^{(t)})^{k-1}). \quad (233)$$

We now sum (233) for $t = T_{n-1}, \dots, T_n$ and get:

$$\frac{A}{m} (z^{(T_n)})^\kappa \log(z^{(T_n)}) \geq \frac{A}{m} (z^{(T_{n-1})})^\kappa \log(z^{(T_{n-1})}) + \sum_{t=T_{n-1}}^{T_n-1} \frac{A(z^{(t)})^\kappa}{m} \log(1 + m(z^{(t)})^{k-1}). \quad (234)$$

Since $\frac{A(z^{(t)})^\kappa}{m} \leq \frac{A(2^n z^{(0)})^\kappa}{m} \leq \frac{Av^\kappa}{m} \ll 1$, we have $(1 + m(z^{(t)})^{k-1})^{\frac{A(z^{(t)})^\kappa}{m}} \geq 1 + A(z^{(t)})^{k-1+\kappa}$. We thus lower bound (234) as:

$$\frac{A}{m}(z^{(T_n)})^\kappa \log(z^{(T_n)}) \geq \frac{A}{m}(z^{(T_{n-1})})^\kappa \log(z^{(T_{n-1})}) + \sum_{t=T_{n-1}}^{T_n-1} \log(1 + A(z^{(t)})^{k-1+\kappa}). \quad (235)$$

On the other hand, by using $z^{(T_{n-1})} \leq 2^n z^{(0)}$ and $z^{(T_n)} \leq 2^{n+1} z^{(0)}$, we have the following upper bound.

$$\begin{aligned} & \frac{A}{m}(z^{(T_n)})^\kappa \log(z^{(T_n)}) \\ & \leq \frac{A}{m}(z^{(T_n)})^\kappa \log(z^{(T_{n-1})}) + \frac{MA}{m}(z^{(T_n)})^\kappa \log(1 + (z^{(T_{n-1})})^{k-1}) \\ & \leq \frac{A \cdot (2^{n+1} z^{(0)})^\kappa}{m} \log(2^n z^{(0)}) + \frac{MA \cdot (2^{n+1} z^{(0)})^\kappa}{m} \log(1 + (2^n z^{(0)})^{k-1}). \end{aligned} \quad (236)$$

Since $z^{(T_{n-1})} \geq 2^{n-1} z^{(0)}$, (236) is finally bounded as:

$$\begin{aligned} & \frac{A}{m} \left((z^{(T_n)})^\kappa \log(z^{(T_n)}) - (z^{(T_{n-1})})^\kappa \log(z^{(T_{n-1})}) \right) \\ & \leq \frac{\Theta(A) \cdot (2^n z^{(0)})^\kappa}{m} \log(2^n z^{(0)}) + \frac{MA \cdot (2^{n+1} z^{(0)})^\kappa}{m} \log(1 + (2^n z^{(0)})^{k-1}). \end{aligned} \quad (237)$$

We combine (235) and (237) to obtain:

$$\begin{aligned} \sum_{t=T_{n-1}}^{T_n-1} \log(1 + m(z^{(t)})^k) & \leq \frac{\Theta(A) \cdot (2^n z^{(0)})^\kappa}{m} \log(2^n z^{(0)}) \\ & \quad + \frac{MA \cdot (2^{n+1} z^{(0)})^\kappa}{m} \log(1 + (2^n z^{(0)})^{k-1}). \end{aligned} \quad (238)$$

We now sum (238) and get:

$$\sum_{t=0}^{T_n-1} \log(1 + m(z^{(t)})^k) \leq n(2^n z^{(0)})^\kappa \left(\frac{\Theta(A)}{m} \log(2^n z^{(0)}) + \frac{2^\kappa MA}{m} \log(1 + (2^n z^{(0)})^{k-1}) \right). \quad (239)$$

We replace $2^n z^{(0)}$ by v and n by $\log(v/z^{(0)})$ in (239) to get the aimed result. \square

K.2 Probabilistic lemmas

Lemma K.3. Let $\{\delta_r\}_{r=1}^{D-C}$ be i.i.d. random variables such that with probability q $\delta_r = \pm 1$ and zero otherwise. Then, with probability at least $1 - 1/\text{poly}(d)$, we have:

$$\sum_{\ell \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_\ell} |\delta_r| \leq \Theta(q(D-C)) \log(d).$$

Proof of Lemma K.3. First, note that $|\delta_r|$ is a Bernoulli random variable with parameter q . Therefore, $\sum_{\ell \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_\ell} |\delta_r|$ is a binomial random variable $\mathcal{B}(D-C, q)$. Therefore, we apply a Chernoff bound to obtain:

$$\mathbb{P} \left[\sum_{\ell \neq \ell(\mathbf{X})} \sum_{r \in \mathcal{S}_\ell} |\delta_r| \geq (1 + \varepsilon)(D-C)q \right] \leq \exp \left(-\frac{(D-C)q\varepsilon^2}{3} \right). \quad (240)$$

Setting $\varepsilon = \sqrt{\frac{3}{(D-C)q} \log(\text{poly}(d))}$ in (240) yields the desired result. \square

Lemma K.4. Let $\{\delta_r\}_{r=1}^C$ be i.i.d. random variables such that with probability q $\delta_r = \pm 1$ and zero otherwise. Then, with probability at least $1 - 1/\text{poly}(d)$, we have:

$$\sum_{r=1}^C |\delta_r| \leq O(1).$$

Proof of Lemma K.4. Let $k \in \mathbb{N}$ and $\Delta := \sum_{r=1}^C |\delta_r|$. The tail bound is bounded as:

$$\mathbb{P}[\Delta \geq k] = \sum_{j=k}^C \binom{C}{j} q^j (1-q)^{C-j} \leq q^k \sum_{j=k}^C \binom{C}{j} \leq q^k \sum_{j=0}^C \binom{C}{j} = 2^C q^k. \quad (241)$$

We want to find k such that $2^C q^k \leq 1/\text{poly}(d)$ which implies $k \leq \log(d)/\log(D) \leq O(1)$. \square

K.3 Logarithmic inequalities

Lemma K.5. *Let $a \in \mathbb{R}$ such that $C_- \leq a \leq C_+$, where $C_+, C_- > 0$. Let $p \geq 3$ be an odd integer. Then, the following inequality holds:*

$$\frac{0.1}{C_+} \log(1 + \exp(-a^p)) \leq \frac{a^{p-1}}{1 + \exp(a^p)} \leq \frac{10}{C_-} \log(1 + \exp(-a^p)).$$

Proof of Lemma K.5. We first remark that:

$$\frac{a^{p-1}}{1 + \exp(a^p)} = \frac{a^p}{a(1 + \exp(a^p))}. \quad (242)$$

Upper bound. We upper bound (242) by applying $a \geq C_-$:

$$\frac{a^{p-1}}{1 + \exp(a^p)} \leq \frac{a^p}{C_-(1 + \exp(a^p))} \quad (243)$$

We obtain the final bound by applying Lemma K.6 to (243).

Lower bound. We lower bound (242) by using $a \leq C_+$:

$$\frac{a^{p-1}}{1 + \exp(a^p)} \geq \frac{a^p}{C_+(1 + \exp(a^p))} \quad (244)$$

We obtain the final bound by applying Lemma K.6 to (244). \square

Lemma K.6 (Connection between derivative and loss). *Let $x > 0$. Then, we have:*

$$0.1 \log(1 + \exp(-x)) \leq \mathfrak{G}(-x) \leq 10 \log(1 + \exp(-x)) \quad (245)$$

Lemma K.7. *Let $x, y > 0$. Assume that $y \leq x$. Then, we have:*

$$\log(1 + xy) \leq (1 + y) \log(1 + x).$$