

# HIGH-DIMENSIONAL SGD ALIGNS WITH EMERGING OUTLIER EIGENSPACES

**G erard Ben Arous**

Courant Institute  
New York University  
New York, NY, USA  
benarous@cims.nyu.edu

**Reza Gheissari**

Department of Mathematics  
Northwestern University  
Evanston, IL, USA  
gheissari@northwestern.edu

**Jiaoyang Huang**

Department of Statistics and Data Science  
University of Pennsylvania  
Philadelphia, PA, USA  
huangjy@wharton.upenn.edu

**Aukosh Jagannath**

Department of Statistics and Actuarial Science  
Department of Applied Mathematics, and  
Cheriton School of Computer Science  
University of Waterloo  
Waterloo, ON, Canada  
a.jagannath@uwaterloo.ca

## ABSTRACT

We rigorously study the joint evolution of training dynamics via stochastic gradient descent (SGD) and the spectra of empirical Hessian and gradient matrices. We prove that in two canonical classification tasks for multi-class high-dimensional mixtures and either 1 or 2-layer neural networks, the SGD trajectory rapidly aligns with emerging low-rank outlier eigenspaces of the Hessian and gradient matrices. Moreover, in multi-layer settings this alignment occurs per layer, with the final layer’s outlier eigenspace evolving over the course of training, and exhibiting rank deficiency when the SGD converges to sub-optimal classifiers. This establishes some of the rich predictions that have arisen from extensive numerical studies in the last decade about the spectra of Hessian and information matrices over the course of training in overparametrized networks.

## 1 INTRODUCTION

Stochastic gradient descent (SGD) and its many variants, are the backbone of modern machine learning algorithms (see e.g., Bottou (1999)). The training dynamics of neural networks, however, are still poorly understood in the non-convex and high-dimensional settings that are frequently encountered. A common explanation for the staggering success of neural networks, especially when overparameterized, is that the loss landscapes that occur in practice have many “flat” directions and a hidden low-dimensional structure within which the bulk of training occurs.

To understand this belief, much attention has been paid to the Hessian of the empirical risk (and related matrices formed via gradients) along training. This perspective on the training dynamics of neural networks was proposed in LeCun et al. (2012), and numerically analyzed in depth in Sagun et al. (2017a;b). The upshot of these studies was a broad understanding of the spectrum of the Hessian of the empirical risk, that we summarize as follows:

1. It has a *bulk* that is dependent on the network architecture, and is concentrated around 0, becoming more-so as the model becomes more overparametrized;
2. It has (relatively) few *outlier eigenvalues* that are dependent on the data, and evolve non-trivially along training while remaining separated from the bulk;

Since those works, these properties of the Hessian and related spectra over the course of training have seen more refined and large-scale experimentation. Pappayan (2019) found a hierarchical decomposition to the Hessian for deep networks, attributing the bulk, emergent outliers, and a *minibulk* to three different “parts” of the Hessian. Perhaps most relevant to this work, Gur-Ari et al. (2019)

noticed that gradient descent tends to quickly align with a *low-dimensional* outlier subspace of the Hessian matrix, and stay in that subspace for long subsequent times. They postulated that this common low-dimensional structure to the SGD and Hessian matrix may be key to many classification tasks in machine learning. For a sampling of other empirical investigations of spectra of Hessians and information matrices along training, see e.g., Ghorbani et al. (2019); Pappayan (2020); Li et al.; Martin & Mahoney (2019); Cohen et al. (2021); Xie et al. (2023).

From a theoretical perspective, much attention has been paid to the Hessians of deep networks using random matrix theory approaches. Most of this work has focused on the spectrum at a fixed point in parameter space, most commonly at initialization. Early works in the direction include Watanabe (2007); Dauphin et al. (2014). Choromanska et al. (2015) noted similarities of neural net Hessians to spin glass Hessians, whose complexity (exponential numbers of critical points) has been extensively studied see e.g., Auffinger & Ben Arous (2013); Auffinger et al. (2013). More recently, the expected complexity has been investigated in statistical tasks like tensor PCA and generalized linear estimation Ben Arous et al. (2019); Maillard et al. (2020). In Pennington & Worah (2018) and Pennington & Bahri (2017), the empirical spectral distribution of Hessians and information matrices in single layer neural networks were studied at initialization. Liao & Mahoney (2021) studied Hessians of some non-linear models, which they referred to as generalized GLMs, also at initialization.

Under the infinite-width neural tangent kernel limit of Jacot et al. (2018), Fan & Wang (2020) derived the empirical spectral distribution at initialization, and Jacot et al. (2020) studied its evolution over the course of training. In this limit the input dimension is kept fixed compared to the parameter dimension, while our interest in this paper is when the input dimension, parameter dimension, and number of samples all scale together.

An important step towards understanding the evolution of Hessians along training in the high-dimensional setting, is of course understanding the training dynamics themselves. Since the classical work of Robbins & Monro (1951), there has been much activity studying limit theorems for stochastic gradient descent. In the high-dimensional setting, following Saad & Solla (1995a;b), investigations have focused on finding a finite number of functions (sometimes called “observables” or “summary statistics”), whose dynamics under the SGD are asymptotically autonomous in the high-dimensional limit. For a necessarily small sampling of this rich line of work, we refer to Goldt et al. (2019); Veiga et al. (2022); Paquette et al. (2021); Arnaboldi et al. (2023); Tan & Vershynin (2019); Ben Arous et al. (2022). Of particular relevance, it was shown by Damian et al. (2022); Mousavi-Hosseini et al. (2023) that for multi-index models, SGD predominantly lives in the low-dimensional subspace spanned by the ground truth parameters.

A class of tasks whose SGD is amenable to this broad approach is classification of Gaussian mixture models (GMMs). With various losses and linearly separable class structures, the minimizer of the empirical risk landscape with single-layer networks was studied in Mignacco et al. (2020); Loureiro et al. (2021). A well-studied case of a Gaussian mixture model needing a two-layer network is under an XOR-type class structure; the training dynamics of SGD for this task were studied in Refinetti et al. (2021) and Ben Arous et al. (2022) and it was found to have a particularly rich structure with positive probability of convergence to bad classifiers among other degenerate phenomena.

Still, a simultaneous understanding of high-dimensional SGD and the Hessian and related matrices’ spectra along the training trajectory has remained largely open.

## 1.1 OUR CONTRIBUTIONS

In this paper, we study the interplay between the training dynamics (via SGD) and the spectral compositions of the empirical Hessian matrix and an empirical gradient second moment matrix, or simply *G-matrix* (2.2) (similar in spirit to an information matrix) over the course of training. We rigorously show the following phenomenology in two canonical high-dimensional classification tasks with  $k$  “hidden” classes:

1. Shortly into training, the empirical Hessian and empirical G-matrices have  $C(k)$  many outlier eigenvalues, and the SGD itself predominantly lives in their eigenspace. Here  $C(k)$  is explicit and depends on the performance of the classifier to which the SGD converges.

2. In multi-layer settings, this alignment happens within each layer, i.e., the first layer parameters align with the outlier eigenspaces of the corresponding blocks of the empirical Hessian and G-matrices, and likewise for the second layer parameters.
3. This alignment is *not* predicated on success at the classification task: when the SGD converges to a sub-optimal classifier, the empirical Hessian and G matrices have lower rank outlier eigenspaces, and the SGD aligns with those rank deficient spaces.

The first model we consider is the basic example of supervised classification of general  $k$ -component Gaussian mixture models with  $k$  linearly independent classes by a single layer neural network. In Theorem 3.1, we establish alignment of the form of Item 1 above, between each of the  $k$  one-vs-all classifiers and their corresponding blocks in the empirical Hessian and G-matrices. See also the depictions in Figures 3.1–3.2. In order to show this, we show that the matrices have an outlier-minibulk-bulk structure throughout the parameter space, and derive limiting dynamical equations for the trajectory of appropriate summary statistics of the SGD trajectory following the approach of Ben Arous et al. (2022). Importantly, the same low-dimensional subspace is at the heart of both the outlier eigenspaces and the summary statistics. At this level of generality, the SGD can behave very differently within the outlier eigenspace. As an example of the refined phenomenology that can arise, we further investigate the special case where the means are orthogonal; here the SGD aligns specifically with the single largest outlier eigenvalue, which itself has separated from the other  $k - 1$  outliers along training. This is proved in Theorem 3.3 and depicted in Figure 3.3. These results are presented in Section 3.

To demonstrate our results in more complex multi-layer settings, we consider supervised classification of a GMM version of the famous XOR problem of Minsky & Papert (1969). This is one of the simplest models that requires a two-layer neural network to solve. We use a two-layer architecture with a second layer of width  $K$ . As indicated by Item 2 above, in Theorems 4.1–4.2 the alignment of the SGD with the matrices’ outlier eigenspaces occurs within each layer, the first layer having an outlier space of rank two, and the second layer having an outlier space of rank 4 when the dynamics converges to an optimal classifier. This second layer’s alignment is especially rich, as when the model is overparametrized ( $K$  large), its outlier space of rank 4 is not present at initialization, and only separates from its rank- $K$  bulk over the course of training. This can be interpreted as a dynamical version of the well-known spectral phase transition in spiked covariance matrices of Baik et al. (2005): see Figure 4.2 for a visualization. Moreover, the SGD for this problem is known to converge to sub-optimal classifiers with probability bounded away from zero Ben Arous et al. (2022), and we find that in these situations, the alignment still occurs but the outlier eigenspaces into which the SGD moves are rank-deficient compared to the number of hidden classes, 4: see Figure 4.3. These results are presented in Section 4.

## 2 PRELIMINARIES

Before turning to our results, let us begin by introducing the following framework and notation that will be relevant throughout this work. We suppose that we are given data from a distribution  $\mathcal{P}_Y$  over pairs  $\mathbf{Y} = (y, Y)$  where  $y \in \mathbb{R}^k$  is a one-hot “label” vector that takes the value 1 on a class (sometimes identified with the element of  $[k] = \{1, \dots, k\}$  on which it is 1), and  $Y \in \mathbb{R}^d$  is a corresponding feature vector. In training we take as loss a function of the form  $L(\mathbf{x}, \mathbf{Y}) : \mathbb{R}^p \times \mathbb{R}^{k+d} \rightarrow \mathbb{R}_+$ , where  $\mathbf{x} \in \mathbb{R}^p$  represents the network parameter. (As we are studying supervised classification, in both settings this loss will be the usual cross-entropy loss corresponding to the architecture used.)

We imagine we have two data sets, a training set  $(\mathbf{Y}^\ell)_{\ell=1}^M$  and a test set  $(\tilde{\mathbf{Y}}^\ell)_{\ell=1}^{\tilde{M}}$ , all drawn i.i.d. from  $\mathcal{P}_Y$ . Let us first define the stochastic gradient descent trained using  $(\mathbf{Y}^\ell)$ . In order to ensure the SGD doesn’t go off to infinity we add an  $\ell^2$  penalty term (as is common in practice) with Lagrange multiplier  $\beta$ . The (online) stochastic gradient descent with initialization  $\mathbf{x}_0$  and learning rate, or step-size,  $\delta$ , will be run using the training set  $(\mathbf{Y}^\ell)_{\ell=1}^M$  as follows:

$$\mathbf{x}_\ell = \mathbf{x}_{\ell-1} - \delta \nabla L(\mathbf{x}_{\ell-1}, \mathbf{Y}^\ell) - \beta \mathbf{x}_{\ell-1}. \quad (2.1)$$

Our aim is to understand the behavior of SGD with respect to principal subspaces, i.e., outlier eigenvectors, of the empirical Hessian matrix and empirical second moment matrix of the gradient.

This latter matrix is exactly the information matrix when  $L$  is the log-likelihood; in our paper  $L$  is taken to be a cross-entropy loss, so we simply refer to this as the G-matrix henceforth. We primarily consider the empirical Hessian and empirical G-matrices generated out of the test data, namely:

$$\nabla^2 \widehat{R}(\mathbf{x}) = \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} \nabla^2 L(\mathbf{x}, \widetilde{\mathbf{Y}}^\ell), \quad \text{and} \quad \widehat{G}(\mathbf{x}) = \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} \nabla L(\mathbf{x}, \widetilde{\mathbf{Y}}^\ell)^{\otimes 2}. \quad (2.2)$$

(Notice that in our online setting, it is as natural to generate these matrices with test data as with train data. See Remark 5.1 for how our results extend when training data is used.) When the parameter space naturally splits into subsets of its indices (e.g., the first-layer weights and the second-layer weights), for a subset  $I$  of the parameter coordinates, we use subscripts  $\nabla_{I,I}^2 \widehat{R}$  and  $\widehat{G}_{I,I}$  to denote the block corresponding to that subset.

To formalize the notion of alignment between the SGD and the principal directions of the Hessian and G-matrices, we introduce the following language. For a subspace  $B$ , we let  $P_B$  denote the orthogonal projection onto  $B$ ; for a vector  $v$ , we let  $\|v\|$  be its  $\ell^2$  norm; and for a matrix  $A$ , let  $\|A\| = \|A\|_{\text{op}}$  be its  $\ell^2 \rightarrow \ell^2$  operator norm.

**Definition 2.1.** *The alignment of a vector  $v$  with a subspace  $B$  is the ratio  $\rho(v, B) = \|P_B v\|/\|v\|$ . We say a vector  $v$  lives in a subspace  $B$  up to error  $\varepsilon$  if  $\rho(v, B) \geq 1 - \varepsilon$ .*

For a matrix  $A$ , we let  $E_k(A)$  denote the span of the top  $k$  eigenvectors of  $A$ , i.e., the span of the  $k$  eigenvectors of  $A$  with the largest absolute values. We also use the following.

**Definition 2.2.** *We say a matrix  $A$  lives in a subspace  $B$  up to error  $\varepsilon$  if there exists  $M$  such that  $\text{Im}(A - M) \subset B$  with  $\|M\|_{\text{op}} \leq \varepsilon \|A\|_{\text{op}}$ , where  $\|A\|_{\text{op}}$  denotes the  $\ell^2$ -to- $\ell^2$  operator norm.*

### 3 CLASSIFYING LINEARLY SEPARABLE MIXTURE MODELS

We begin by illustrating our results on (arguably) the most basic problem of high-dimensional multi-class classification, namely supervised classification of a  $k$  component Gaussian mixture model with constant variance and linearly independent means using a single-layer network. (This is sometimes used as a toy model for the training dynamics of the last layer of a deep network via the common ansatz that the output of the second-to-last layer of a deep network behaves like a linearly separable mixture of Gaussians: see e.g., the neural collapse phenomenon posited by Pappayan et al. (2020).)

#### 3.1 DATA MODEL

Let  $\mathcal{C} = [k]$  be the collection of classes, with corresponding distinct class means  $(\mu_a)_{a \in [k]} \in \mathbb{R}^d$ , covariance matrices  $I_d/\lambda$ , where  $\lambda > 0$  can be viewed a signal-to-noise parameter, and corresponding probabilities  $0 < (p_a)_{a \in [k]} < 1$  such that  $\sum_{a \in [k]} p_a = 1$ . The number of classes  $k = O(1)$  is fixed (here and throughout the paper  $o(1)$ ,  $O(1)$  and  $\Omega(1)$  notations are with respect to the dimension parameter  $d$ , and may hide constants that are dimension independent such as  $k, \beta$ ).

For the sake of simplicity we take the means to be unit norm. Further, in order for the task to indeed be solvable with the single-layer architecture, we assume that the means are linearly independent, say with a fixed (i.e.,  $d$ -independent) matrix of inner products  $(\overline{m}_{ab})_{a,b} = (\mu_a \cdot \mu_b)_{a,b}$ . Our data distribution  $\mathcal{P}_Y$  is a mixture of the form  $\sum_c p_c \mathcal{N}(\mu_c, I_d/\lambda)$ , with an accompanying class label  $y \in \mathbb{R}^k$ . Namely, if  $Z_\lambda \sim \mathcal{N}(0, I_d/\lambda)$ , our data is given as  $\mathbf{Y} = (y, Y)$  where:

$$y \sim \sum_{a \in [k]} p_a \delta_{\mathbf{1}_a}, \quad \text{and} \quad Y \sim \sum_{a \in [k]} y_a \mu_a + Z_\lambda. \quad (3.1)$$

We perform classification by training a single layer network formed by  $k$  ‘‘all-vs-one’’ classifiers using the cross entropy loss (equivalently, we are doing multi-class logistic regression):

$$L(\mathbf{x}, \mathbf{Y}) = - \sum_{c \in [k]} y_c x^c \cdot Y + \log \sum_{c \in [k]} \exp(x^c \cdot Y), \quad (3.2)$$

where  $\mathbf{x} = (x^c)_{c \in \mathcal{C}}$  are the parameters, each of which is a vector in  $\mathbb{R}^d$ , i.e.,  $\mathbf{x} \in \mathbb{R}^{dk}$ . (Note that we can alternatively view  $\mathbf{x}$  as a  $k \times d$  matrix.)

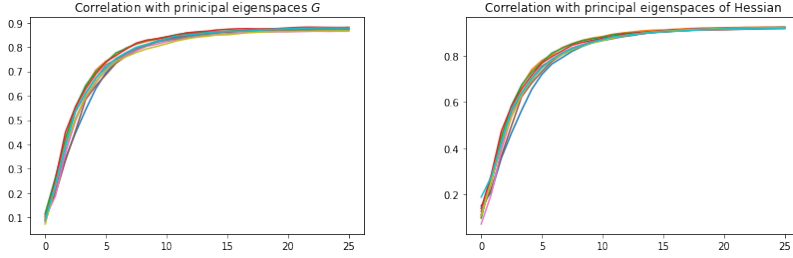


Figure 3.1: The alignment of the SGD trajectory  $\mathbf{x}_\ell^c$  with  $E_k(\nabla_{cc}^2 \widehat{R}(\mathbf{x}_\ell))$  (left) and  $E_k(\widehat{G}_{cc}(\mathbf{x}_\ell))$  (right) over time, for  $c = 1, \dots, k$  (shown in different colors). The  $x$ -axis is rescaled time,  $\ell\delta$ . The parameters here are  $k = 10$  classes in dimension  $d = 1000$  with  $\lambda = 10$ ,  $\beta = 0.01$ , and  $\delta = 1/d$ .

### 3.2 RESULTS AND DISCUSSION

Our first result is that after some linearly many steps, the SGD finds the subspace generated by the outlier eigenvalues of the Hessian and/or G-matrix of the test loss and lives there for future times.

**Theorem 3.1.** *Consider the mixture of  $k$ -Gaussians with loss function from (3.2), and SGD (2.1) with learning rate  $\delta = O(1/d)$ , regularizer  $\beta > 0$ , initialized from  $\mathcal{N}(0, I_d/d)$ . There exists  $\alpha_0, \lambda_0$  such that if  $\lambda \geq \lambda_0$ , and  $\widetilde{M} \geq \alpha_0 d$ , the following hold. For every  $\varepsilon > 0$ , there exists  $T_0(\varepsilon)$  such that for any time horizon  $T_0 < T_f < M/d$ , with probability  $1 - o_d(1)$ ,*

$$\mathbf{x}_\ell^c \text{ lives in } E_k(\nabla_{cc}^2 \widehat{R}(\mathbf{x}_\ell)) \text{ and in } E_k(\widehat{G}_{cc}(\mathbf{x}_\ell)),$$

for every  $c \in [k]$ , up to  $O(\varepsilon + \lambda^{-1})$  error, for all  $\ell \in [T_0\delta^{-1}, T_f\delta^{-1}]$ .

This result is demonstrated in Figure 3.1 which plots the alignment of the training dynamics  $\mathbf{x}_\ell^c$  with the principal eigenspaces of the Hessian and  $G$  for each  $c \in [k]$ . As we see the alignment increases to near 1 rapidly for all blocks in both matrices. This theorem, and all our future results, are stated using a random Gaussian initialization, scaled such that the norm of the parameters is  $O(1)$  in  $d$ . The fact that this is Gaussian is not relevant to the results, and similar results hold for other uninformative initializations with norm of  $O(1)$ .

Theorem 3.1 follows from the following theorem that describes the SGD trajectory, its Hessian and G-matrix (and their top  $k$  eigenspaces), all live up to  $O(1/\lambda)$  error in  $\text{Span}(\mu_1, \dots, \mu_k)$ .

**Theorem 3.2.** *In the setup of Theorem 3.1, the following live in  $\text{Span}(\mu_1, \dots, \mu_k)$  up to  $O(\varepsilon + \lambda^{-1})$  error with probability  $1 - o_d(1)$ :*

1. The state of the SGD along training,  $\mathbf{x}_\ell^c$  for every  $c$ ;
2. The  $b, c$  blocks of the empirical test Hessian,  $\nabla_{bc}^2 \widehat{R}(\mathbf{x}_\ell)$  for all  $b, c \in [k]$ ;
3. The  $b, c$  blocks of the empirical test G-matrix  $\widehat{G}_{bc}(\mathbf{x}_\ell)$  for all  $b, c \in [k]$ .

We demonstrate this result in Figure 3.2.

Inside the low-rank space spanned by  $\mu_1, \dots, \mu_k$ , the training dynamics, Hessian and G-matrix spectra can display different phenomena depending on the relative locations of  $\mu_1, \dots, \mu_k$  and weights  $p_1, \dots, p_k$ . To illustrate the more refined alignment phenomena, let us take as a concrete example  $p_c = \frac{1}{k}$  for all  $c$ , and  $\mu_1, \dots, \mu_k$  orthonormal.

With this concrete choice, we can analyze the limiting dynamical system of the SGD without much difficulty and its relevant dynamical observables have a single stable fixed point, to which the SGD converges in linearly many, i.e.,  $O(\delta^{-1})$ , steps. This allows us to show more precise alignment that occurs within  $\text{Span}(\mu_1, \dots, \mu_k)$  over the course of training.

**Theorem 3.3.** *In the setting of Theorem 3.1, with the means  $(\mu_1, \dots, \mu_k)$  being orthonormal, the estimate  $\mathbf{x}_\ell^c$  has  $\Omega(1)$ , positive, inner product with the top eigenvector of both  $\nabla_{cc}^2 \widehat{R}(\mathbf{x}_\ell)$  and  $\widehat{G}_{cc}(\mathbf{x}_\ell)$  (and negative,  $\Omega(1)$  inner product with the  $k-1$  next largest eigenvectors). Also, the top eigenvector of  $\nabla_{cc}^2 \widehat{R}(\mathbf{x}_\ell)$ , as well as that of  $\widehat{G}_{cc}(\mathbf{x}_\ell)$ , live in  $\text{Span}(\mu_c)$  up to  $O(\varepsilon + \lambda^{-1})$  error.*




Figure 3.2: From left to right: Plot of entries of  $\mathbf{x}_\ell^1$  and the  $k$  leading eigenvectors (in different colors) of  $\nabla_{11}^2 \widehat{R}(\mathbf{x}_\ell)$  and  $\widehat{G}_{11}(\mathbf{x}_\ell)$  respectively at the end of training, namely  $\ell = 50 \cdot d = 25,000$  steps. Here the  $x$ -axis represents the coordinate index. The parameters are the same as in Fig. 3.1 and the means are  $\mu_i = e_{i*50}$ .




Figure 3.3: Left: the eigenvalues (in different colors) of  $\nabla^2 \widehat{R}_{11}(\mathbf{x}_\ell)$  over the course of training. The leading  $k$  eigenvalues are separated from the bulk at all times, and the top eigenvalue, corresponding to  $\mu_1$  separates from the remaining eigenvalues soon after initialization. Right: the inner product of  $\mathbf{x}_\ell^1$  with the means  $\mu_1, \dots, \mu_k$  undergoes a similar separation over the course of training. Parameters are the same as in preceding figures.

Put together, the above three theorems describe the following rich scenario for classification of the  $k$ -GMM. At initialization, and throughout the parameter space in each class-block, the Hessian and G-matrices decompose into a rank- $k$  outlier part spanned by  $\mu_1, \dots, \mu_k$ , and a correction term of size  $O(1/\lambda)$  in operator norm. Furthermore, when initialized randomly, the SGD is not aligned with the outlier eigenspaces, but does align with them in a short  $O(\delta^{-1})$  number of steps. Moreover, when the means are orthogonal, each class block of the SGD  $\mathbf{x}_\ell^c$  in fact correlates strongly with the specific mean for its class  $\mu_c$ , and simultaneously, in the Hessian and G-matrices along training, the eigenvalue corresponding to  $\mu_c$  becomes distinguished from the other  $k - 1$  outliers. We illustrate these last two points in Figure 3.3. We also refer the reader to Section F for further numerical demonstrations. Each of these phenomena appear in more general contexts, and in even richer manners, as we will see in the following section.

## 4 CLASSIFYING XOR-TYPE MIXTURE MODELS VIA TWO-LAYER NETWORKS

With the above discussion in mind, let us now turn to more complex classification tasks that are not linearly separable and require the corresponding network architecture to be multilayer.

### 4.1 DATA MODEL

For our multilayer results, we consider the problem of classifying a 4-component Gaussian mixture whose class labels are in a so-called XOR form. More precisely, consider a mixture of four Gaussians with means  $\mu, -\mu, \nu, -\nu$  where  $\|\mu\| = \|\nu\| = 1$  and, say for simplicity, are orthogonal, and variances  $I_d/\lambda$ . There are two classes, class label 1 for Gaussians with mean  $\pm\mu$ , and 0 for Gaussians with mean  $\pm\nu$ . To be more precise, our data distribution  $\mathcal{P}_Y$  is

$$y \sim \frac{1}{2}\delta_0 + \frac{1}{2}\delta_1 \quad \text{and} \quad Y \sim \begin{cases} \frac{1}{2}\mathcal{N}(\mu, I_d/\lambda) + \frac{1}{2}\mathcal{N}(-\mu, I_d/\lambda) & y = 1 \\ \frac{1}{2}\mathcal{N}(\nu, I_d/\lambda) + \frac{1}{2}\mathcal{N}(-\nu, I_d/\lambda) & y = 0 \end{cases}. \quad (4.1)$$

This is a Gaussian version of the famous XOR problem of Minsky & Papert (1969). It is one of the simplest examples of a classification task requiring a multi-layer network to express a good classifier.

We therefore use a two-layer architecture with the intermediate layer having width  $K \geq 4$  (any less and a Bayes-optimal classifier would not be expressible), ReLu activation function  $g(x) = x \vee 0$  and then sigmoid activation function  $\sigma(x) = \frac{1}{1+e^{-x}}$ . The parameter space is then  $\mathbf{x} = (W, v) \in \mathbb{R}^{Kd+K}$ , where the first layer weights are denoted by  $W = W(\mathbf{x}) \in \mathbb{R}^{K \times d}$  and the second layer weights are denoted by  $v = v(\mathbf{x}) \in \mathbb{R}^K$ . We use the binary cross-entropy loss on this problem,

$$L(\mathbf{x}; \mathbf{Y}) = -yv \cdot g(WY) + \log(1 + e^{v \cdot g(WY)}), \quad (4.2)$$

with  $g$  applied entrywise. The SGD for this classification task was studied in some detail in Refinetti et al. (2021) and Ben Arous et al. (2022), with “critical” step size  $\delta = \Theta(1/d)$ .

## 4.2 RESULTS AND DISCUSSION

We begin our discussion with the analogue of Theorem 3.1 in this setting. As the next theorem demonstrates, in this more subtle problem, the SGD still finds and lives in the principal directions of the empirical Hessian and G-matrices.<sup>1</sup> Here, the principal directions vary significantly across the parameter space, and the alignment phenomenon differs depending on which fixed point the SGD converges to. Moreover, the relation between the SGD and the principal directions of the Hessian and G-matrices can be seen per layer.

**Theorem 4.1.** *Consider the XOR GMM mixture with loss function (4.2) and the corresponding SGD (2.1) with  $\beta \in (0, 1/8)$ , learning rate  $\delta = O(1/d)$ , initialized from  $\mathcal{N}(0, I_d/d)$ . There exist  $\alpha_0, \lambda_0$  such that if  $\lambda \geq \lambda_0$ , and  $\tilde{M} \geq \alpha_0 d$ , the following hold. For every  $\varepsilon > 0$ , there exists  $T_0(\varepsilon)$  such that for any time horizon  $T_0 < T_f < M/d$ , with probability  $1 - o_d(1)$ , for all  $i \in \{1, \dots, K\}$ ,*

1.  $W_i(\mathbf{x}_\ell)$  lives in  $E_2(\nabla_{W_i W_i}^2 \widehat{R}(\mathbf{x}_\ell))$  and in  $E_2(\widehat{G}_{W_i W_i}(\mathbf{x}_\ell))$ , and
2.  $v(\mathbf{x}_\ell)$  lives in  $E_4(\nabla_{vv}^2 \widehat{R}(\mathbf{x}_\ell))$  and  $E_4(\widehat{G}_{vv}(\mathbf{x}_\ell))$ ,

up to  $O(\varepsilon + \lambda^{-1/2})^2$  error, for all  $\ell \in [T_0 \delta^{-1}, T_f \delta^{-1}]$ .

*Remark 1.* The restriction to  $\beta < 1/8$  in Theorems 4.1–4.2 is because when  $\beta > 1/8$  the regularization is too strong for the SGD to be meaningful; in particular, the SGD converges ballistically to the origin in parameter space, with no discernible preference for the directions corresponding to  $\mu, \nu$  as the other directions. The above theorems are still valid there if the notion of error for living in a space from Definition 2.1 were additive instead of multiplicative (i.e., at most  $\varepsilon$  rather than  $\varepsilon \|v\|$ ).

This theorem is demonstrated in Figure 4.1. There we have plotted the alignment of the rows in the intermediate layer with the space spanned by the top two eigenvectors of the corresponding first-layer blocks of the Hessian and G-matrices, and similarly for the final layer.

As before, the above theorem follows from the following theorem that describes both the SGD trajectory, its Hessian, and its G-matrix, living up to  $O(\varepsilon + \lambda^{-1/2})$  error in their first-layer blocks in  $\text{Span}(\mu, \nu)$  and (applying  $g$  entrywise) in their second layer blocks in

$$\text{Span}(g(W(\mathbf{x}_\ell)\mu), g(-W(\mathbf{x}_\ell)\mu), g(W(\mathbf{x}_\ell)\nu), g(-W(\mathbf{x}_\ell)\nu)).$$

**Theorem 4.2.** *In the setting of Theorem 4.1, up to  $O(\varepsilon + \lambda^{-1/2})$  error with probability  $1 - o_d(1)$ , the following live in  $\text{Span}(\mu, \nu)$ ,*

- The first layer weights,  $W_i(\mathbf{x}_\ell)$  for each  $i \in \{1, \dots, K\}$ ,
- The first-layer empirical test Hessian  $\nabla_{W_i W_i}^2 \widehat{R}(\mathbf{x}_\ell)$  for each  $i \in \{1, \dots, K\}$ ,
- The first-layer empirical test G-matrix  $\widehat{G}_{W_i W_i}(\mathbf{x}_\ell)$  for each  $i \in \{1, \dots, K\}$ ,

and the following live in  $\text{Span}(g(W(\mathbf{x}_\ell)\mu), g(-W(\mathbf{x}_\ell)\mu), g(W(\mathbf{x}_\ell)\nu), g(-W(\mathbf{x}_\ell)\nu))$

<sup>1</sup>As the second derivative of ReLU is only defined in the sense of distributions, care must be taken when studying these objects. This singularity is not encountered by the SGD trajectory almost surely.

<sup>2</sup>In this theorem and in Theorem 4.2, the big- $O$  notation also hides constant dependencies on the initial magnitudes of  $v(\mathbf{x}_0)$ , and on  $T_f$ .




Figure 4.1: (a)–(b) depict the alignment of the first layer weights  $W_i(\mathbf{x}_\ell)$  for  $i = 1, \dots, K$  (in different colors) with the principal subspaces of the corresponding blocks of the Hessian and G-matrices, i.e., with  $E_2(\nabla_{W_i W_i}^2 \hat{R}(\mathbf{x}_\ell))$  and  $E_2(\hat{G}_{W_i W_i}(\mathbf{x}_\ell))$ . (c)–(d) plot the second-layer alignment, namely of  $v(\mathbf{x}_\ell)$  with  $E_4(\nabla_{vv}^2 \hat{R}(\mathbf{x}_\ell))$  and  $E_4(\hat{G}_{vv}(\mathbf{x}_\ell))$ . Parameters are  $d = 1000$ ,  $\lambda = 10$ , and  $K = 20$




Figure 4.2: The eigenvalues (in different colors) of the  $vv$  blocks of the Hessian and G-matrices over time from a random initialization. Initially, there is one outlier eigenvalue due to the positivity of the ReLU activation. Along training, four outlier eigenvalues separate from the bulk, corresponding to the four “hidden” classes in the XOR problem. Parameters are the same as in Figure 4.1.

- The second layer weights  $v(\mathbf{x}_\ell)$ ,
- The second-layer empirical test Hessian  $\nabla_{vv}^2 \hat{R}(\mathbf{x}_\ell)$ ,
- The second-layer empirical test G-matrix  $\hat{G}_{vv}(\mathbf{x}_\ell)$ .

Let us discuss a bit more the phenomenology of the alignment in the second layer. First of all, we observe that the subspace in which the alignment occurs is a random—depending on the initialization and trajectory (in particular, choice of fixed point the SGD converges to)—4-dimensional subspace of  $\mathbb{R}^K$ . Furthermore, if we imagine  $K$  to be much larger than 4 so that the model is overparametrized, at initialization, unlike the 1-layer case studied in Section 3, the Hessian and G-matrices do not exhibit any alignment in their second layer blocks. In particular, the second layer blocks of the Hessian and G-matrices look like (non-spiked) Gaussian orthogonal ensemble and Wishart matrices<sup>3</sup> in  $K$  dimensions at initialization, and it is only over the course of training that they develop 4 outlier eigenvalues as the first layer of the SGD begins to align with the mean vectors and the vectors  $(g(W(\mathbf{x}_\ell)\vartheta))_{\vartheta \in \{\pm\mu, \pm\nu\}}$  in turn get large enough to generate outliers. This crystallization of the last layer around these vectors over the course of training is reminiscent of the neural collapse phenomenon outlined in Pappayan et al. (2020) (see also Han et al. (2022); Zhu et al. (2021)). The simultaneous emergence, along training, of outliers in the Hessian and G-matrix spectra can be seen as a dynamical version of what is sometimes referred to as the BBP transition after Baik et al. (2005) (see also P ech e (2006)). This dynamical transition is demonstrated in Figure 4.2

Finally, we recall that Ben Arous et al. (2022) found a positive probability (uniformly in  $d$ , but shrinking as the architecture is overparametrized by letting  $K$  grow) that the SGD converges to sub-optimal classifiers from a random initialization. When this happens,  $g(W(\mathbf{x}_\ell)\vartheta)$  remains small for the hidden classes  $\vartheta \in \{\pm\mu, \pm\nu\}$  that are not classifiable with SGD output. In those situations, Theorem 4.2 shows that the outlier subspace in the  $vv$ -blocks that emerges will have rank smaller than 4, whereas when an optimal classifier is found it will be of rank 4: see Figure 4.3. Knowing that

<sup>3</sup>These are two of the most classical random matrix models, see Anderson et al. (2010) for more.






Figure 4.3: Evolution of eigenvalues in the  $v$  component of  $G$  over time in rank deficient cases. Here SGD is started from initializations that converge to suboptimal classifiers (this has uniformly positive,  $K$ -dependent, probability under a random initialization). From left to right, the SGD’s classifier varies in the number of hidden classes it discerns, from 1 to 4. There is still a dynamical spectral transition, now with only a corresponding number of emerging outlier eigenvalues.

the classification task entails a mixture of 4 means, this may provide a method for devising a stopping rule for classification tasks of this form by examining the rank of the outlier eigenspaces of the last layer Hessian or  $G$ -matrix. While the probability of such sub-optimal classification is bounded away from zero, the probability exponentially goes to zero as the model is overparametrized via  $K \rightarrow \infty$ , and that gives more chances to allow the second layer SGD, Hessian, and  $G$ -matrices to exhibit full 4-dimensional principal spaces. This serves as a concrete and provable manifestation of the *lottery ticket hypothesis* of Frankle & Carbin (2019).

## 5 OUTLINE AND IDEAS OF PROOF

The proofs of our main theorems break into three key steps.

1. In Sections A–B, we show that the *population* Hessian and  $G$ -matrices, have bulks (and possibly minibulks) that are  $O(1/\lambda)$  in operator norm, and finite  $C(k)$ -rank parts. We can explicitly characterize the low-rank part’s eigenvalues and eigenvectors up to  $O(1/\lambda)$  corrections, as functions of the parameter space, to see exactly where their emergence as outliers occurs depending on the model.
2. In Section C, we analyze the SGD trajectories for the  $k$ -GMM and XOR classification problems. We do this using the limiting effective dynamics theorem proven in Ben Arous et al. (2022) for finite families of summary statistics. We derive ODE limits for these summary statistics, notably for the general  $k$ -GMM which was not covered in that paper: see Theorem C.7. We then pull back the limiting dynamics to finite  $d$  and expand its  $\lambda$ -finite trajectory about its  $\lambda = \infty$  solution. These latter steps involve understanding some of the stability properties of the dynamical system limits of the SGD’s summary statistics.
3. In Section D, we prove concentration for the empirical Hessian and  $G$ -matrices about their population versions, in operator norm, throughout the parameter space. In some related settings, including binary mixtures of Gaussians, Mei et al. (2018) established concentration of the empirical Hessian about the population Hessian uniformly in the parameter space assuming polylogarithmic sample complexity. Our proofs are based on  $\epsilon$ -nets and concentration inequalities for uniformly sub-exponential random variables, together with some twists due, for instance, to non-differentiability of the ReLU function.

*Remark 5.1.* Our results are stated for the empirical Hessian and  $G$ -matrices generated using test data, along the SGD trajectory generated from training data. Since we are considering online SGD, the empirical Hessian and  $G$ -matrices with training data are no more relevant than those generated from test data, but the reader may still wonder whether the same behavior holds. A straightforward modification of our arguments in Section D extends our results for the  $k$ -GMM model to Hessian and  $G$ -matrices generated from train data, if we assume that  $M \gtrsim d \log d$  rather than simply  $M \gtrsim d$ . See the full version for a sketch of this extension. It is an interesting mathematical question to drop this extra logarithmic factor. The extension in the XOR case is technically more involved due to the regularity of the ReLU function. See Section F for numerical demonstrations that the phenomena in the empirical matrices generated from train data are identical to those generated from test data.

## AUTHOR CONTRIBUTIONS

All authors contributed equally to this work.

## ACKNOWLEDGMENTS

The authors thank anonymous referees for their helpful comments. G.B. acknowledges the support of NSF grant DMS-2134216. R.G. acknowledges the support of NSF grant DMS-2246780. The research of J.H. is supported by NSF grants DMS-2054835 and DMS-2331096. A.J. acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canada Research Chairs programme. Cette recherche a été entreprise grâce, en partie, au soutien financier du Conseil de Recherches en Sciences Naturelles et en Génie du Canada (CRSNG), [RGPIN-2020-04597, DGEER-2020-00199], et du Programme des chaires de recherche du Canada.

## REFERENCES

- Greg W Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*. Number 118. Cambridge university press, 2010.
- Luca Arnaboldi, Ludovic Stephan, Florent Krzakala, and Bruno Loureiro. From high-dimensional & mean-field dynamics to dimensionless odes: A unifying approach to sgd in two-layers networks. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 1199–1227. PMLR, 12–15 Jul 2023. URL <https://proceedings.mlr.press/v195/arnaboldi23a.html>.
- Antonio Auffinger and Gérard Ben Arous. Complexity of random smooth functions on the high-dimensional sphere. *Ann. Probab.*, 41(6):4214–4247, 2013. ISSN 0091-1798. doi: 10.1214/13-AOP862.
- Antonio Auffinger, Gérard Ben Arous, and Jiří Černý. Random matrices and complexity of spin glasses. *Comm. Pure Appl. Math.*, 66(2):165–201, 2013. ISSN 0010-3640. doi: 10.1002/cpa.21422.
- Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.
- Gérard Ben Arous, Song Mei, Andrea Montanari, and Mihai Nica. The landscape of the spiked tensor model. *Comm. Pure Appl. Math.*, 72(11):2282–2330, 2019. ISSN 0010-3640. doi: 10.1002/cpa.21861.
- Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. High-dimensional limit theorems for SGD: Effective dynamics and critical scaling. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=Q38D6xxrKHe>.
- Léon Bottou. *On-Line Learning and Stochastic Approximations*. Cambridge University Press, USA, 1999. ISBN 0521652634.
- Anna Choromanska, Mikael Henaff, Michael Mathieu, Gerard Ben Arous, and Yann LeCun. The Loss Surfaces of Multilayer Networks. In Guy Lebanon and S. V. N. Vishwanathan (eds.), *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38 of *Proceedings of Machine Learning Research*, pp. 192–204, San Diego, California, USA, 09–12 May 2015. PMLR. URL <https://proceedings.mlr.press/v38/choromanska15.html>.
- Jeremy Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=jh-rTtvkGeM>.
- Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In Po-Ling Loh and Maxim Raginsky (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp.

- 5413–5452. PMLR, 02–05 Jul 2022. URL <https://proceedings.mlr.press/v178/damian22a.html>.
- Yann N. Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS’14*, pp. 2933–2941, Cambridge, MA, USA, 2014. MIT Press.
- Zhou Fan and Zhichao Wang. Spectra of the conjugate kernel and neural tangent kernel for linear-width neural networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJl-b3RcF7>.
- Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2232–2241. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/ghorbani19b.html>.
- Sebastian Goldt, Madhu Advani, Andrew M Saxe, Florent Krzakala, and Lenka Zdeborová. Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup. *Advances in neural information processing systems*, 32, 2019.
- Guy Gur-Ari, Daniel A. Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace, 2019. URL <https://openreview.net/forum?id=ByeTHsAqtX>.
- X.Y. Han, Vardan Papyan, and David L. Donoho. Neural collapse under MSE loss: Proximity to and dynamics on the central path. In *International Conference on Learning Representations*, 2022. URL [https://openreview.net/forum?id=w1UbdvWH\\_R3](https://openreview.net/forum?id=w1UbdvWH_R3).
- Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL [https://proceedings.neurips.cc/paper\\_files/paper/2018/file/5a4be1fa34e62bb8a6ec6b91d2462f5a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/5a4be1fa34e62bb8a6ec6b91d2462f5a-Paper.pdf).
- Arthur Jacot, Franck Gabriel, and Clement Hongler. The asymptotic spectrum of the hessian of dnn throughout training. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkgscaNYPs>.
- Yann A. LeCun, Léon Bottou, Genevieve B. Orr, and Klaus-Robert Müller. *Efficient Back-Prop*, pp. 9–48. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-35289-8. doi: 10.1007/978-3-642-35289-8\_3. URL [https://doi.org/10.1007/978-3-642-35289-8\\_3](https://doi.org/10.1007/978-3-642-35289-8_3).
- Xinyan Li, Qilong Gu, Yingxue Zhou, Tiancong Chen, and Arindam Banerjee. *Hessian based analysis of SGD for Deep Nets: Dynamics and Generalization*, pp. 190–198. doi: 10.1137/1.9781611976236.22. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611976236.22>.
- Zhenyu Liao and Michael W. Mahoney. Hessian eigenspectra of more realistic nonlinear models. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=o-RYNV01xA8>.
- Bruno Loureiro, Gabriele Sicuro, Cedric Gerbelot, Alessandro Pocco, Florent Krzakala, and Lenka Zdeborova. Learning gaussian mixtures with generalized linear models: Precise asymptotics in high-dimensions. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=j3eGyNMPvh>.

- Antoine Maillard, Gérard Ben Arous, and Giulio Biroli. Landscape complexity for the empirical risk of generalized linear models. In Jianfeng Lu and Rachel Ward (eds.), *Proceedings of The First Mathematical and Scientific Machine Learning Conference*, volume 107 of *Proceedings of Machine Learning Research*, pp. 287–327. PMLR, 20–24 Jul 2020. URL <https://proceedings.mlr.press/v107/maillard20a.html>.
- Charles H. Martin and Michael W. Mahoney. Traditional and heavy tailed self regularization in neural network models, 2019. URL <https://openreview.net/forum?id=SJeFNoRcFQ>.
- Song Mei, Yu Bai, and Andrea Montanari. The landscape of empirical risk for nonconvex losses. *Ann. Statist.*, 46(6A):2747–2774, 12 2018. doi: 10.1214/17-AOS1637.
- Francesca Mignacco, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Dynamical mean-field theory for stochastic gradient descent in gaussian mixture classification. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 9540–9550. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/6c81c83c4bd0b58850495f603ab45a93-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/6c81c83c4bd0b58850495f603ab45a93-Paper.pdf).
- Marvin Minsky and Seymour Papert. An introduction to computational geometry. *Cambridge tiass., HIT*, 479:480, 1969.
- Alireza Mousavi-Hosseini, Sejun Park, Manuela Girotti, Ioannis Mitliagkas, and Murat A Erdogdu. Neural networks efficiently learn low-dimensional representations with SGD. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=6taykzqcPD>.
- Vardan Papyan. Measurements of three-level hierarchical structure in the outliers in the spectrum of deepnet hessians. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5012–5021. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/papyan19a.html>.
- Vardan Papyan. Traces of class/cross-class structure pervade deep learning spectra. *Journal of Machine Learning Research*, 21(252):1–64, 2020. URL <http://jmlr.org/papers/v21/20-933.html>.
- Vardan Papyan, X. Y. Han, and David L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020. doi: 10.1073/pnas.2015509117. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2015509117>.
- Courtney Paquette, Kiwon Lee, Fabian Pedregosa, and Elliot Paquette. SGD in the large: Average-case analysis, asymptotics, and stepsize criticality. In *Conference on Learning Theory*, pp. 3548–3626. PMLR, 2021.
- S. Péché. The largest eigenvalue of small rank perturbations of hermitian random matrices. *Probability Theory and Related Fields*, 134(1):127–173, Jan 2006. ISSN 1432-2064. doi: 10.1007/s00440-005-0466-z.
- Jeffrey Pennington and Yasaman Bahri. Geometry of neural network loss surfaces via random matrix theory. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2798–2806. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/pennington17a.html>.
- Jeffrey Pennington and Pratik Worah. The spectrum of the fisher information matrix of a single-hidden-layer neural network. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL [https://proceedings.neurips.cc/paper\\_files/paper/2018/file/18bb68e2b38e4a8ce7cf4f6b2625768c-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/18bb68e2b38e4a8ce7cf4f6b2625768c-Paper.pdf).

- Maria Refinetti, Sebastian Goldt, Florent Krzakala, and Lenka Zdeborová. Classifying high-dimensional gaussian mixtures: Where kernel methods fail and neural networks succeed. In *International Conference on Machine Learning*, pp. 8936–8947. PMLR, 2021.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *Ann. Math. Statistics*, 22: 400–407, 1951. ISSN 0003-4851. doi: 10.1214/aoms/1177729586.
- David Saad and Sara Solla. Dynamics of on-line gradient descent learning for multilayer neural networks. *Advances in neural information processing systems*, 8, 1995a.
- David Saad and Sara A Solla. On-line learning in soft committee machines. *Physical Review E*, 52(4):4225, 1995b.
- Levent Sagun, Leon Bottou, and Yann LeCun. Eigenvalues of the hessian in deep learning: Singularity and beyond, 2017a. URL <https://openreview.net/forum?id=B186cP9gx>.
- Levent Sagun, Utku Evci, V. Ugur Güney, Yann N. Dauphin, and Léon Bottou. Empirical analysis of the hessian of over-parametrized neural networks. *CoRR*, abs/1706.04454, 2017b. URL <http://arxiv.org/abs/1706.04454>.
- Yan Shuo Tan and Roman Vershynin. Online stochastic gradient descent with arbitrary initialization solves non-smooth, non-convex phase retrieval. *arXiv preprint arXiv:1910.12837*, 2019.
- Rodrigo Veiga, Ludovic Stephan, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Phase diagram of stochastic gradient descent in high-dimensional two-layer neural networks. *arXiv preprint arXiv:2202.00293*, 2022.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Sumio Watanabe. Almost all learning machines are singular. In *2007 IEEE Symposium on Foundations of Computational Intelligence*, pp. 383–388, 2007. doi: 10.1109/FOCI.2007.371500.
- Zeke Xie, Qian-Yuan Tang, Mingming Sun, and Ping Li. On the overlooked structure of stochastic gradients. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=H4GsteoL0M>.
- Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=KRODJAA6pzE>.

## A ANALYSIS OF THE POPULATION MATRICES: 1-LAYER NETWORKS

In this section, we study the Hessian and G-matrix of the *population* loss for the  $k$ -GMM problem whose data distribution and loss were given in (3.1)–(3.2). Specifically, we give the Hessian and G-matrices’  $\lambda$ -large expansion, and showing they have low rank structures with top eigenspaces generated by  $O(1/\lambda)$  perturbations of the mean vectors  $(\mu_1, \dots, \mu_k)$ . In Section D, we will show that the empirical matrices are well concentrated about the population matrices.

### A.1 PRELIMINARY CALCULATIONS AND NOTATION

It helps to first fix some preliminary notation, and give expressions for derivatives of the loss function of (3.2). Differentiating that, we get

$$\nabla_{x_c} L = (-y_c + \frac{\exp(x^c \cdot Y)}{\sum_a \exp(x^a \cdot Y)}) Y, \quad (\text{A.1})$$

the Hessian matrix

$$\nabla_{x_b} \nabla_{x_c} L = \left( \frac{\exp(x^b \cdot Y)}{\sum_a \exp(x^a \cdot Y)} \delta_{bc} - \frac{\exp(x^b \cdot Y) \exp(x^c \cdot Y)}{(\sum_a \exp(x^a \cdot Y))^2} \right) Y \otimes Y, \quad (\text{A.2})$$

and the G-matrix

$$\nabla_{x_b} L \otimes \nabla_{x_c} L = (y_c y_b - \pi_Y(b) y_c - y_b \pi_Y(c) + \pi_Y(c) \pi_Y(b)) Y \otimes Y. \quad (\text{A.3})$$

Given the above, the following probability distribution over the  $[k]$  classes naturally arises, for which we reserve the notation  $\pi_Y(\cdot) \in \mathcal{M}_1([k])$  (which is a probability measure on  $[k]$ ):

$$\pi_Y(c) = \pi_Y(c; \mathbf{x}) := \frac{\exp(x^c \cdot Y)}{\sum_{a \in [k]} \exp(x^a \cdot Y)}. \quad (\text{A.4})$$

Note the dependency of  $\pi_Y$  on the point in parameter space  $\mathbf{x}$ . Since in this section  $\mathbf{x}$  can be viewed as fixed, we will suppress this dependence from the notation. In the rest of this section, we will denote  $Z \sim \mathcal{N}(0, I_d/\lambda)$ . We also denote

$$\langle x^B \rangle = \langle x^B \rangle_{\pi_Y} = \sum_a x^a \pi_Y(a), \quad \langle x^B; x^B \rangle = \langle x^B; x^B \rangle_{\pi_Y} = \sum_a (x^a \otimes x^a) \pi_Y(a) - \langle x \rangle_{\pi_Y}^{\otimes 2}, \quad (\text{A.5})$$

where  $B \sim \pi_Y$ .

We consider the population Hessian  $\nabla^2 \Phi = \nabla^2 \mathbb{E}[L]$ , block by block, using  $\nabla_{bc} \Phi$  to denote the  $d \times d$  block in  $\nabla_x^2 \Phi$  corresponding to  $\nabla_{x_b} \nabla_{x_c} \Phi$ . Then, thanks to (A.2), the  $bc$  block of the population Hessian is of the form

$$\nabla_{bc} \Phi = \mathbb{E}[(\pi_Y(b) \delta_{bc} - \pi_Y(b) \pi_Y(c)) Y \otimes Y].$$

We also consider the population G-matrix  $\Gamma = \mathbb{E}[\nabla L^{\otimes 2}]$ , block by block, using  $\Gamma_{bc}$  to denote the  $d \times d$  block corresponding to  $\mathbb{E}[\nabla_{x_b} L \otimes \nabla_{x_c} L]$ . Then, thanks to (A.3), the  $bc$  block of the population Hessian is of the form

$$\Gamma_{bc} = \mathbb{E}[(y_c y_b - \pi_Y(b) y_c - y_b \pi_Y(c) + \pi_Y(c) \pi_Y(b)) Y \otimes Y].$$

For both population Hessian matrix and G-matrix, We will study the off-diagonal blocks  $a \neq c$  and the diagonal ones  $a = c$  separately.

### A.2 ANALYSIS OF THE POPULATION HESSIAN MATRIX

We now compute exact expressions for the blocks of the population Hessian as  $\lambda$  gets large: see Lemmas A.1–A.2.

We begin by studying the off-diagonal blocks:  $b \neq c$ , for which,

$$\nabla_{bc} \Phi = \mathbb{E}[(\pi_Y(b) \delta_{cb} - \pi_Y(b) \pi_Y(c)) Y^{\otimes 2}] = - \sum_l p_l \mathbb{E}[\pi_{Y_l}(b) \pi_{Y_l}(c) Y_l^{\otimes 2}] \quad (\text{A.6})$$

where  $Y_l = \mu_l + Z$ , i.e., it is distributed like  $Y$  given class choice  $l$ . It helps here to recall the well-known Gaussian integration by parts formula: For  $f$  that is differentiable with derivative of at most exponential growth at infinity,

$$\mathbb{E}f(Z_\lambda)Z_\lambda = \frac{1}{\lambda}\mathbb{E}\nabla f(Z_\lambda)$$

Our goal is to show the following.

**Lemma A.1.** *The off-diagonal blocks of the population Hessian satisfy*

$$\nabla_{bc}\Phi(\mathbf{x}) = \mathbb{E}\left\{\pi_Y(c)\pi_Y(a)\left[\left(\mu_y + \frac{1}{\lambda}[(x^c + x^a) - 2\langle x^B \rangle]\right)^{\otimes 2} + \frac{1}{\lambda}I_d - \frac{2}{\lambda^2}\langle x^B; x^B \rangle\right]\right\}$$

*In particular they are shifts by the identity of a matrix of rank at most  $k^2$ .*

*Proof.* We decompose each term in (A.6) as three terms

$$\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)]\mu_l^{\otimes 2} + 2\text{Sym}(\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)Z] \otimes \mu_l) + \mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)Z^{\otimes 2}] =: \text{(i)} + \text{(ii)} + \text{(iii)}$$

where we've used here multi-linearity of  $\otimes$ . Here  $\text{Sym}(c \otimes b) = (c \otimes b + b \otimes c)/2$ . Let's look at this term-by-term. We leave term (i) as is. For Term (ii), we notice by Gaussian integration by parts,

$$\begin{aligned}\mathbb{E}[\pi_{Y_l}(c)\pi_{Y_l}(b)Z] \otimes \mu_l &= \frac{1}{\lambda}\mathbb{E}[\nabla_Z(\pi_{Y_l}(c)\pi_{Y_l}(b))] \otimes \mu_l \\ &= \frac{1}{\lambda}\mathbb{E}\left[\left((x^b + x^c) - 2\langle x^B \rangle_{\pi_{Y_l}}\right)\pi_{Y_l}(b)\pi_{Y_l}(c)\right] \otimes \mu_l \\ &= \frac{1}{\lambda}(\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)](x^b + x^c) \otimes \mu_l - 2\mathbb{E}_y[\langle x^B \rangle_{\pi_{Y_l}} \pi_{Y_l}(b)\pi_{Y_l}(c)] \otimes \mu_l).\end{aligned}$$

where we recall (A.5) and used the derivative calculation

$$\nabla_Z \pi_{Y_l}(a) = \nabla_Z \frac{\exp[x^a \cdot (\mu_l + Z)]}{\sum_b \exp[x^b \cdot (\mu_l + Z)]} = (x^a - \langle x^B \rangle_{\pi_{Y_l}})\pi_{Y_l}(a). \quad (\text{A.7})$$

This tells us that

$$\text{(ii)} = \frac{2}{\lambda}\text{Sym}\left(\mathbb{E}\left[\left((x^b + x^c) - 2\langle x^B \rangle_{\pi_{Y_l}}\right)\pi_{Y_l}(b)\pi_{Y_l}(c)\right] \otimes \mu_l\right),$$

notice that this is of rank at most 2 and of order  $O(1/\lambda)$ .

Finally for term (iii), we integrate by parts again, to get, for every  $i, j$ ,

$$\begin{aligned}\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)Z_i Z_j] &= \frac{1}{\lambda}\mathbb{E}[\partial_{Z_i}(\pi_{Y_l}(b)\pi_{Y_l}(c))Z_j] + \frac{1}{\lambda}\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)\delta_{ij}] \\ &= \frac{1}{\lambda^2}\mathbb{E}[\partial_{Z_i Z_j}(\pi_{Y_l}(b)\pi_{Y_l}(c))] + \frac{1}{\lambda}\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)\delta_{ij}].\end{aligned} \quad (\text{A.8})$$

so term (iii) is given by

$$\frac{1}{\lambda}\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)I_d] + \frac{1}{\lambda^2}\mathbb{E}[\nabla_Z^2(\pi_{Y_l}(b)\pi_{Y_l}(c))].$$

Examining this second term,

$$\begin{aligned}\mathbb{E}[\partial_{Z_i Z_j}(\pi_{Y_l}(b)\pi_{Y_l}(c))] &= \mathbb{E}\left[\partial_{Z_j}\left\{(x_i^b + x_i^c - 2\langle x_i^B \rangle_{\pi_{Y_l}})\pi_{Y_l}(b)\pi_{Y_l}(c)\right\}\right] \\ &= \mathbb{E}\left[\left(x_i^b + x_i^c - 2\langle x_i^B \rangle_{\pi_{Y_l}}\right)\left((x_j^b + x_j^c) - 2\langle x_j^B \rangle_{\pi_{Y_l}}\right)\pi_{Y_l}(b)\pi_{Y_l}(c)\right] - 2\mathbb{E}[\langle x_i^B; x_j^B \rangle_{\pi_{Y_l}} \pi_{Y_l}(b)\pi_{Y_l}(c)]\end{aligned}$$

where we first used (A.7), then used that

$$\nabla_Z \langle x^B \rangle_{\pi_{Y_l}} = \nabla_Z \frac{\sum_b x^b \exp[x^b \cdot (\mu_l + Z)]}{\sum_b \exp[x^b \cdot (\mu_l + Z)]} = \langle x^B \otimes x^B \rangle_{\pi_{Y_l}} - \langle x^B \rangle_{\pi_{Y_l}}^{\otimes 2} = \langle x^B; x^B \rangle_{\pi_{Y_l}} \quad (\text{A.9})$$

As such,

$$\frac{1}{\lambda^2} \mathbb{E}[\nabla_Z^2(\pi_{Y_l}(b)\pi_{Y_l}(c))] = \frac{1}{\lambda^2} \left\{ \mathbb{E} \left[ (x^b + x^c - 2\langle x^B \rangle_{\pi_{Y_l}})^{\otimes 2} \pi_{Y_l}(b)\pi_{Y_l}(c) \right] - 2\mathbb{E} \left[ \langle x^B; x^B \rangle_{\pi_{Y_l}} \pi_{Y_l}(b)\pi_{Y_l}(c) \right] \right\}$$

This term can be seen to be of rank at most  $k^2$  (each  $\langle x^B \rangle$  is a weighted sum of  $(x^a)_a$ ).

Combining all three terms above, we get that the off-diagonal block of the population Hessian matrix is given by the sum over  $y \in [k]$  of

$$\begin{aligned} & \mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)] \mu_l^{\otimes 2} + \frac{2}{\lambda} \left[ \text{Sym} \left( \mathbb{E} \left[ \left( (x^b + x^a) - 2\langle x^B \rangle_{\pi_{Y_l}} \right) \pi_{\pi_{Y_l}}(b)\pi_{\pi_{Y_l}}(c) \right] \otimes \mu_l \right) \right] + \frac{1}{\lambda} \mathbb{E}[\pi_{\pi_{Y_l}}(b)\pi_{\pi_{Y_l}}(c)] I_d \\ & + \frac{1}{\lambda^2} \left( \mathbb{E} \left[ \left( (x^b + x^a) - 2\langle x^B \rangle_{\pi_{Y_l}} \right)^{\otimes 2} \pi_{\pi_{Y_l}}(b)\pi_{\pi_{Y_l}}(c) \right] - 2\mathbb{E} \left[ \langle x^B; x^B \rangle_{\pi_{Y_l}} \pi_{\pi_{Y_l}}(b)\pi_{\pi_{Y_l}}(c) \right] \right). \end{aligned} \quad (\text{A.10})$$

Summing this in  $y$  we get that an off-diagonal block is of the form

$$-(A + \frac{1}{\lambda}B + \frac{1}{\lambda^2}C), \quad (\text{A.11})$$

where

$$\begin{aligned} A &= \sum_l p_l \mathbb{E} \left[ \pi_{\pi_{Y_l}}(c)\pi_{\pi_{Y_l}}(a) \right] \mu_l^{\otimes 2}, \\ B &= \sum_l p_y \mathbb{E}[\pi_{Y_l}(a)\pi_{Y_l}(c)] I_d + 2 \sum_l p_l \text{Sym} \left( \mathbb{E}_l \left[ \left( x^c + x^a - 2\langle x^B \rangle_{\pi_{Y_l}} \right) \pi_{Y_l}(c)\pi_{Y_l}(a) \right] \otimes \mu_l \right), \\ C &= \sum_l p_l \mathbb{E} \left[ \left( x^c + x^a - 2\langle x^B \rangle_{\pi_{Y_l}} \right)^{\otimes 2} \pi_{\pi_{Y_l}}(c)\pi_{\pi_{Y_l}}(a) \right] - 2\mathbb{E} \left[ \langle x^B; x^B \rangle_{\pi_{Y_l}} \pi_{Y_l}(c)\pi_{Y_l}(a) \right]. \end{aligned}$$

In summary to leading order it is  $A$  which is rank  $k$ . To next order it is  $O(1/\lambda)$  and that term is a full rank (identity) plus a rank at most  $2k$  term. To next order it is  $O(1/\lambda^2)$  and this is a covariance-type quantity with respect to the Gibbs probability  $\pi_{Y_y}$ , with rank at most  $k^2$ .

We can group the expression (A.11) further as

$$\sum_l p_l \mathbb{E} \left\{ \pi_{Y_l}(c)\pi_{Y_l}(a) \left[ \left( \mu_l + \frac{1}{\lambda} [x^c + x^a - 2\langle x^B \rangle_{\pi_{Y_l}}] \right)^{\otimes 2} + \frac{1}{\lambda} I_d - \frac{2}{\lambda^2} \langle x^B; x^B \rangle_{\pi_{Y_l}} \right] \right\}$$

This is of the form rank  $k$  plus rank  $k^2$  shifted by the identity (adding one more eigenvalue). Incorporating the average over  $l$  into the expectation, this is exactly the claimed expression.  $\square$

### A.2.1 ON-DIAGONAL BLOCKS

In this section, we study the  $aa$  diagonal blocks

$$\nabla_{aa} \Phi = \mathbb{E} \left[ (\pi_Y(a)(1 - \pi_Y(a))Y^{\otimes 2}) \right] = \sum_l p_l \mathbb{E} \left[ (\pi_{Y_l}(a)(1 - \pi_{Y_l}(a))Y_l^{\otimes 2}) \right]. \quad (\text{A.12})$$

We prove the following large  $\lambda$  expansion.

**Lemma A.2.** *The diagonal  $aa$ -block of the population Hessian  $\nabla_{aa} \Phi$  equals*

$$\begin{aligned} & \mathbb{E} \left[ \pi_Y(a) \left( \mu + \frac{1}{\lambda} (x^a - \langle x^B \rangle_{\pi_Y})^{\otimes 2} - \frac{1}{\lambda^2} \langle x^B; x^B \rangle_{\pi_Y} \right) \right] \\ & - \mathbb{E} \left[ \pi(a)^2 \left( \mu + \frac{2}{\lambda} (x^a - \langle x^B \rangle_{\pi_Y})^{\otimes 2} - \frac{2}{\lambda^2} \langle x^B; x^B \rangle_{\pi_Y} \right) \right] + \frac{1}{\lambda} \mathbb{E}[\pi_Y(a)(1 - \pi_Y(a))] I_d. \end{aligned}$$

*In particular, it is a shift by the identity of a rank at-most  $k^2$  matrix.*

*Proof.* By (A.12), the diagonal  $aa$ -block is given by an average over  $l \in [k]$  of

$$\mathbb{E}[\pi_{Y_l}(a)(1 - \pi_{Y_l}(a))Y^{\otimes 2}] = \mathbb{E}[\pi_{Y_l}(a)(\mu_l + Z)^{\otimes 2}] - \mathbb{E}[\pi_{Y_l}(a)^2(\mu_l + Z)^{\otimes 2}].$$



Note that the second term was exactly what was computed Lemma A.1, setting  $b = c$ . It remains to compute the first. To this end, we proceed as before and for each fixed  $l$ , write the above as

$$\mathbb{E}[\pi_{Y_l}(a)]\mu_l^{\otimes 2} + 2 \text{Sym}(\mathbb{E}[\pi_{Y_l}(a)Z] \otimes \mu_l) + \mathbb{E}[\pi_{Y_l}(a)Z^{\otimes 2}].$$

We will integrate the second and third terms by-parts. By the gradient calculations of (A.7),

$$\mathbb{E}[\pi_{Y_l}(a)Z] = \frac{1}{\lambda} \mathbb{E}[\nabla_Z \pi_{Y_l}(a)] = \frac{1}{\lambda} \mathbb{E}[(x^a - \langle x^B \rangle_{\pi_{Y_l}}) \pi_{Y_l}(a)], \quad (\text{A.13})$$

and by the calculation of (A.9),

$$\begin{aligned} \mathbb{E}[\pi_{Y_l}(a)Z^{\otimes 2}] &= \frac{1}{\lambda} \mathbb{E}[\pi_{Y_l}(a)]I_d + \frac{1}{\lambda^2} \mathbb{E}[\nabla_Z^2 \pi_{Y_l}(a)] \\ &= \frac{1}{\lambda} \mathbb{E}[\pi_{Y_l}(a)]I_d + \mathbb{E}[\pi_{Y_l}(a)((x^a - \langle x^B \rangle_{\pi_{Y_l}})^{\otimes 2} - \langle x^B; x^B \rangle_{\pi_{Y_l}})]. \end{aligned} \quad (\text{A.14})$$

Combining the above expressions yields

$$\mathbb{E}[\pi_{Y_l}(a)Y_l^{\otimes 2}] = \mathbb{E}[\pi_{Y_l}(a)((\mu_l + \frac{1}{\lambda}(x^a - \langle x^B \rangle_{\pi_{Y_l}})^{\otimes 2} + \frac{1}{\lambda}I_d - \frac{1}{\lambda^2}\langle x^B; x^B \rangle_{\pi_{Y_l}})].$$

On the otherhand, Lemma A.1, yields

$$\mathbb{E}[\pi_{Y_l}(a)^2 Y_l^{\otimes 2}] = \mathbb{E}[\pi_{Y_l}(a)^2((\mu_l + \frac{2}{\lambda}(x^a - \langle x^B \rangle_{\pi_{Y_l}})^{\otimes 2} + \frac{1}{\lambda}I_d - \frac{2}{\lambda^2}\langle x^B; x^B \rangle_{\pi_{Y_l}})].$$

Combining these two and averaging over  $l$  yields the desired.  $\square$

### A.3 ANALYSIS OF THE POPULATION G-MATRIX

We now compute an exact expansion of the population G-matrix as  $\lambda$  gets large.

**Lemma A.3.** *For any  $b, c$ , the  $bc$  block of the population G-matrix can be written in the form:*

$$\mathbb{E}[\nabla_{x_b} L \otimes \nabla_{x_c} L] = A + \frac{B}{\lambda} + \frac{C}{\lambda^2},$$

where  $A$  is in the span of  $(\mu_c)_c$  and  $B, C$  have operator norm bounded by 1. In particular,

$$\begin{aligned} A &= \delta_{bc} p_b \mu_b^{\otimes 2} - p_c \mathbb{E}\{\pi_{Y_c}(b)[\mu_c + \frac{1}{\lambda}(x^b - \langle x^B \rangle_{\pi_{Y_c}})]^{\otimes 2}\} - p_b \mathbb{E}\{\pi_{Y_b}(c)[\mu_b + \frac{1}{\lambda}(x^c - \langle x^B \rangle_{\pi_{Y_b}})]^{\otimes 2}\} \\ &\quad + \sum_l p_l \mathbb{E}\{\pi_{Y_l}(b)\pi_{Y_l}(c)[\mu_l - \frac{1}{\lambda}(x^b + x^c - 2\langle x^B \rangle_{\pi_{Y_l}})]^{\otimes 2}\}, \end{aligned}$$

$$B = (\delta_{bc} p_b - (p_b \mathbb{E}[\pi_{Y_b}(c)] + p_c \mathbb{E}[\pi_{Y_c}(b)])) + \sum_l p_l \mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)]I_d,$$

$$C = -p_c \mathbb{E}[\langle x^B; x^B \rangle_{\pi_{Y_c}} \pi_{Y_c}(b)] - p_b \mathbb{E}[\langle x^B; x^B \rangle_{\pi_{Y_b}} \pi_{Y_b}(c)] - 2 \sum_l p_l \mathbb{E}[\langle x^B; x^B \rangle_{\pi_{Y_l}} \pi_{Y_l}(b)\pi_{Y_l}(c)]$$

*Proof.* We recall from (A.3), there are four terms in the  $bc$  block of the G-matrix:

$$\mathbb{E}[\nabla_{x_b} L \otimes \nabla_{x_c} L] =: (i) - (ii) - (iii) + (iv), \quad (\text{A.15})$$

where

$$\begin{aligned} (i) &:= \mathbb{E}[y_c y_b Y \otimes Y], \quad (ii) := \mathbb{E}[\pi_Y(b) y_c Y \otimes Y], \\ (iii) &:= \mathbb{E}[y_b \pi_Y(c) Y \otimes Y], \quad (iv) := \mathbb{E}[\pi_Y(c) \pi_Y(b) Y \otimes Y] \end{aligned}$$

The first term in (A.15) is easy to compute

$$(i) = \mathbb{E}[y_b y_c Y \otimes Y] = \delta_{bc} p_b (\mu_b \otimes \mu_b + I_d / \lambda). \quad (\text{A.16})$$

The second and third terms are similar to each other, we will only compute the second term (ii),

$$\begin{aligned} (ii) &= \mathbb{E}[y_c \pi_Y(b) Y \otimes Y] = p_c \mathbb{E}[\pi_{Y_c}(b)(\mu_c + Z)^{\otimes 2}] \\ &= p_c \mathbb{E}[\pi_{Y_c}(b)\mu_c^{\otimes 2}] + 2p_c \text{Sym} \mathbb{E}[\pi_{Y_c}(b)Z \otimes \mu_c] + p_c \mathbb{E}[\pi_{Y_c}(b)Z^{\otimes 2}] \end{aligned} \quad (\text{A.17})$$

Using the calculations of (A.13) and (A.14), we conclude that

$$(ii) = p_c \mathbb{E}[\pi_{Y_c}(b)(\mu_c^{\otimes 2} + I_d/\lambda)] + \frac{1}{\lambda} 2 \text{Sym}(\mathbb{E}[(x^b - \langle x^B \rangle_{\pi_{Y_c}}) \pi_{Y_c}(b)] \otimes \mu_c) + \frac{2}{\lambda^2} p_c \mathbb{E}[(\langle x^b - \langle x^B \rangle_{\pi_{Y_c}} \rangle^{\otimes 2} - \langle x^B; x^B \rangle_{\pi_{Y_c}}) \pi_{Y_c}(b)] \quad (\text{A.18})$$

For the last term (iv) in (A.15), it has been computed in (A.10) that (iv) is the expectation over  $l$  of

$$\mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)]\mu_l^{\otimes 2} + \frac{1}{\lambda} \left[ 2 \text{Sym}(\mathbb{E}[\langle (x^b + x^c) - 2\langle x^B \rangle_{\pi_{Y_l}} \rangle \pi_{Y_l}(b)\pi_{Y_l}(c)] \otimes \mu_l) \right] + \frac{1}{\lambda} \mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)]I_d + \frac{1}{\lambda^2} \left( \mathbb{E}[\langle (x^b + x^c) - 2\langle x^B \rangle_{\pi_{Y_l}} \rangle^{\otimes 2} \pi_{Y_l}(b)\pi_{Y_l}(c)] - 2\mathbb{E}[\langle x^B; x^B \rangle_{\pi_{Y_l}} \pi_{Y_l}(b)\pi_{Y_l}(c)] \right) \quad (\text{A.19})$$

By plugging (A.16), (A.18), its analogue for (iii), and (A.19) into (A.15), we conclude that the block  $\mathbb{E}[G_{x_b x_c}]$  of the population G-matrix is given by

$$\mathbb{E}[G_{x_b x_c}] = A' + \frac{B'}{\lambda} + \frac{C'}{\lambda^2},$$

where

$$A' = \delta_{bc} p_b \mu_b^{\otimes 2} - p_c \mathbb{E}[\pi_{Y_c}(b)] \mu_c^{\otimes 2} - p_b \mathbb{E}[\pi_{Y_b}(c)] \mu_b^{\otimes 2} + \sum_l p_l \mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)] \mu_l^{\otimes 2},$$

and

$$B' = \left( \delta_{bc} p_b - (p_b \mathbb{E}[\pi_{Y_b}(c)] + p_c \mathbb{E}[\pi_{Y_c}(b)]) + \sum_l p_l \mathbb{E}[\pi_{Y_l}(b)\pi_{Y_l}(c)] \right) I_d - 2p_c \text{Sym}(\mathbb{E}[(x^b - \langle x^B \rangle_{\pi_{Y_c}}) \pi_{Y_c}(b)] \otimes \mu_c) - 2p_b \text{Sym}(\mathbb{E}[(x^c - \langle x^B \rangle_{\pi_{Y_b}}) \pi_{Y_b}(c)] \otimes \mu_b) + \sum_l 2p_l \text{Sym}(\mathbb{E}[\langle (x^b + x^c) - 2\langle x^B \rangle_{\pi_{Y_l}} \rangle \pi_{Y_l}(b)\pi_{Y_l}(c)] \otimes \mu_l),$$

and

$$C' = -p_c \mathbb{E}[\langle (x^b - \langle x^B \rangle_{\pi_{Y_c}}) \rangle^{\otimes 2} - \langle x^B; x^B \rangle_{\pi_{Y_c}}] \pi_{Y_c}(b) - p_b \mathbb{E}[\langle (x^c - \langle x^B \rangle_{\pi_{Y_b}}) \rangle^{\otimes 2} - \langle x^B; x^B \rangle_{\pi_{Y_b}}] \pi_{Y_b}(c) + \sum_l p_l \left( \mathbb{E}[\langle (x^b + x^c) - 2\langle x^B \rangle_{\pi_{Y_l}} \rangle^{\otimes 2} \pi_{Y_l}(b)\pi_{Y_l}(c)] - 2\mathbb{E}[\langle x^B; x^B \rangle_{\pi_{Y_l}} \pi_{Y_l}(b)\pi_{Y_l}(c)] \right).$$

Grouping tensor-squares we obtain the desired decomposition in terms of  $A$ ,  $B$ , and  $C$ .  $\square$

## B ANALYSIS OF POPULATION MATRICES: THE 2-LAYER CASE

In this section we analyze the population Hessian and G-matrices for the 2-layer XOR model, and especially its  $\lambda$  large behavior by viewing it as a perturbation of its  $\lambda = \infty$  value. Specifically, we compute its large  $\lambda$  expansion, and uncover an underlying low-rank structure. We will show that the empirical Hessian and G-matrices concentrate about their population versions in Section D.

### B.1 PRELIMINARY CALCULATIONS

Recall the data model for the 2-layer XOR GMM from (4.1) and the loss function (4.2), with  $\sigma$  being the sigmoid function and  $g$  being ReLU. The 2-layer architecture has intermediate layer width  $K$ , so that the first layer weights  $W$  form a  $K \times d$  matrix, and the second layer weights  $v$  form a  $K$ -vector. Observe that

$$\nabla_v L = -(y - \hat{y})g(WY) \quad \text{and} \quad \nabla_{W_i} L = -(y - \hat{y})g'(W_i \cdot Y)v_i Y,$$

where

$$\hat{y} = \frac{e^{v \cdot g(WY)}}{1 + e^{v \cdot g(WY)}} = \sigma(v \cdot g(WY)). \quad (\text{B.1})$$

The  $vv$ -block of the G-matrix is the  $K \times K$  matrix given by

$$\nabla_v L^{\otimes 2} = (y - \hat{y})^2 g(WY)^{\otimes 2}. \quad (\text{B.2})$$

The  $W_i W_j$  block of the G-matrix corresponding to  $W$  is given by

$$\nabla_{W_i} L \otimes \nabla_{W_j} L = (y - \hat{y})^2 v_i v_j g'(W_i \cdot Y) g'(W_j \cdot Y) Y^{\otimes 2}. \quad (\text{B.3})$$

The per-layer Hessian for the loss on a given sample is then given by

$$\begin{aligned} \nabla_{vv}^2 L &= g(WY)^{\otimes 2} \hat{y}(1 - \hat{y}), \\ \nabla_{W_i W_j}^2 L &= (-\delta_{ij} v_i g''(W_i \cdot Y)(y - \hat{y}) + v_i v_j g'(W_i \cdot Y) g'(W_j \cdot Y) \hat{y}(1 - \hat{y})) Y^{\otimes 2} \end{aligned} \quad (\text{B.4})$$

Note that we may also write the diagonal block in  $W$  of the form,

$$\nabla_{W_i}^2 L = -(y - \hat{y}) \text{diag}(v_i g''(W_i \cdot Y)) \otimes Y^{\otimes 2} + \hat{y}(1 - \hat{y})(v \text{diag}(g'(WY)))^{\otimes 2} \otimes Y^{\otimes 2}$$

We assume that  $W_i \neq 0$  for all  $i$ , a set which we call  $\mathcal{W}_0^c$ . Then recall that for ReLu activation  $g$ , we have  $g'' = \delta_0$  in the sense of distributions. Thus as long as  $W \in \mathcal{W}_0^c$ , we can drop  $g''$  in (B.4) for any finite  $\lambda < \infty$ .

Our aim is to approximate the expected values of (B.4) by their " $\lambda = \infty$ " versions, and show that they are within  $O(\lambda^{-1/2})$  of one another. Let  $\bar{F}(t) = P(G > t)$  where  $G \sim N(0, 1)$  be the tail probability of a standard Gaussian. Let us also introduce the notations for  $\vartheta \in \{\pm\mu, \pm\nu\}$ ,

$$m_i^\vartheta = W_i \cdot \vartheta \quad \text{and} \quad R_{ii} = W_i \cdot W_j \quad \text{for } 1 \leq i, j \leq K.$$

These notations will reappear as summary statistics for the analysis of the SGD in the following sections.

**Lemma B.1.** For  $X \sim \mathcal{N}(\vartheta, I_d/\lambda)$  and  $W$  a  $K \times d$  matrix, we have that

$$\|\mathbb{E}[g(WX)^{\otimes 2}] - g(W\vartheta)^{\otimes 2}\|_{\text{op}} \lesssim_K \psi_2(\vartheta, W, \lambda) \quad (\text{B.5})$$

where

$$\begin{aligned} \psi_2(\vartheta, W, \lambda) &= \max_{1 \leq i, j \leq K} (|m_i^\vartheta|^3 + \lambda^{-3/2} R_{ii}^{3/2})^{1/3} \cdot (|m_j^\vartheta|^3 + \lambda^{-3/2} R_{jj}^{3/2})^{1/3} \\ &\quad \times \left( \bar{F} \left( \sqrt{\frac{\lambda}{R_{ii}}} |m_i^\vartheta| \right) + \bar{F} \left( \sqrt{\frac{\lambda}{R_{jj}}} |m_j^\vartheta| \right) \right)^{1/3} + \frac{W_i \cdot W_j}{\lambda}. \end{aligned} \quad (\text{B.6})$$

*Proof.* By the equivalence of norms in finite dimensional vector spaces, it suffices to control the norm entry-wise at the price of a constant that depends at most on  $K$ .

**Case 1:**  $W_i \cdot \vartheta, W_j \cdot \vartheta \geq 0$ . In this case we have that

$$\mathbb{E}[g(W_i \cdot X)g(W_j \cdot X)] - g(W_i \cdot \vartheta)g(W_j \cdot \vartheta) = -\mathbb{E}[(W_i \cdot \vartheta)(W_j \cdot \vartheta) \mathbb{1}_{W_i \cdot X < 0 \cup W_j \cdot X < 0}] + \frac{W_i \cdot W_j}{\lambda}.$$

The absolute value of the first term is bounded, by Holder's inequality and a union bound, by

$$\mathbb{E}[|W_i \cdot X|^3]^{1/3} \mathbb{E}[|W_j \cdot X|^3]^{1/3} (\mathbb{P}(W_i \cdot X < 0) + \mathbb{P}(W_j \cdot X < 0))^{1/3}.$$

**Case 2:**  $W_i \cdot \vartheta < 0$  (the case  $W_j \cdot \vartheta < 0$  is symmetrical). Here we have

$$\mathbb{E}[g(W_i \cdot X)g(W_j \cdot X)] = \mathbb{E}[(W_i \cdot X)(W_j \cdot X) \mathbb{1}_{W_i \cdot X, W_j \cdot X > 0}],$$

which, by Holder's inequality, is bounded by

$$\mathbb{E}[|W_i \cdot X|^3]^{1/3} \mathbb{E}[|W_j \cdot X|^3]^{1/3} \mathbb{P}(W_i \cdot X > 0)^{1/3}.$$

combining the two cases yields the desired.  $\square$

**Lemma B.2.** There exists a universal  $c > 0$  such that for  $X \sim \mathcal{N}(\vartheta, I_d/\lambda)$  and  $W \in \mathbb{R}^{K \times d}$ ,

$$\mathbb{E}[\|g(WX)^{\otimes 2}\|_{\text{op}}^2]^{1/2} \leq \|W\vartheta\|^4 + 6 \frac{\|W\vartheta\|^2}{\lambda} + \frac{c}{\lambda^2} \|W\|_F^2, \quad (\text{B.7})$$

where  $\|\cdot\|_F$  is the Frobenius norm.

*Proof.* Let  $Z_i = W_i \cdot Z$ . Then

$$\mathbb{E}(W_i \cdot X)^4 = \mathbb{E}[(m_i^\vartheta + Z_i)^4] = (m_i^\vartheta)^4 + \frac{6(m_i^\vartheta)^2}{\lambda} + \frac{c\|W_i\|}{\lambda^2},$$

so that

$$\mathbb{E}[\|g(WX)^{\otimes 2}\|_{\text{op}}^2] = \mathbb{E}[\|g(WX)\|^4] \leq \|W\vartheta\|_4^4 + 6\frac{\|W\vartheta\|^2}{\lambda} + \frac{c}{\lambda^2}\|W\|_F^2,$$

with the bound the following from the fact that  $\|\cdot\|_p \geq \|\cdot\|_q$  for  $p \geq q$ .  $\square$

**Lemma B.3.** For  $X \sim \mathcal{N}(\vartheta, I_d/\lambda)$  and  $v \in \mathbb{R}^K, W \in \mathbb{R}^{K \times d}$  we have

$$\mathbb{E}[(\sigma(v \cdot g(WX))^2 - \sigma(v \cdot g(W\vartheta))^2)^2] \vee \mathbb{E}[(\sigma'(v \cdot g(WX)) - \sigma'(v \cdot g(W\vartheta)))^2] \lesssim_K \psi_3\left(\frac{1}{\lambda}\|v\|_2^2\|W\|_F^2\right), \quad (\text{B.8})$$

where  $\psi_3(x) = x(1+x)$ .

*Proof.* Observe that

$$\begin{aligned} & |\sigma(v \cdot g(WX))^2 - \sigma(v \cdot g(W\vartheta))^2| \\ & \leq \sigma'(v \cdot g(W\vartheta))|v \cdot (g(WX) - g(W\vartheta))| + O(|v \cdot g(WX) - g(W\vartheta)|^2) \quad (\text{B.9}) \\ & \lesssim |v \cdot (g(WX) - g(W\vartheta))| + |v \cdot (g(WX) - g(W\vartheta))|^2, \end{aligned}$$

and similarly for  $|\sigma'(v \cdot g(WX)) - \sigma'(v \cdot g(W\vartheta))|$  where we used that  $\sigma, \sigma', \sigma''$  are all uniformly bounded. By Jensen's inequality, it will suffice to bound the expectation of the quadratic terms. We begin by noting that

$$\|g(WX) - g(W\vartheta)\|^2 \leq \sum_{1 \leq i \leq K} (W_i \cdot Z)^2$$

so that

$$\mathbb{E}[(v \cdot (g(WX) - g(W\vartheta)))^4] \leq \|v\|^4 \mathbb{E}[\|WZ\|^4] \leq \|v\|^4 \frac{1}{\lambda^2} \|W\|_F^4.$$

Consequently, taking the expectation of the right-hand side of (B.9) squared, gives

$$\begin{aligned} \mathbb{E}[(\sigma(v \cdot g(WX))^2 - \sigma(v \cdot g(W\vartheta))^2)^2] & \lesssim \mathbb{E}[(v \cdot (g(WX) - g(W\vartheta)))^2] + \mathbb{E}[(v \cdot (g(WX) - g(W\vartheta)))^4] \\ & \lesssim_K \frac{1}{\lambda} \sqrt{\|v\|_2^4 \|W\|_F^4} + \frac{1}{\lambda^2} \|v\|_2^4 \|W\|_F^4 = \psi_3\left(\frac{1}{\lambda}\|v\|_2^2\|W\|_F^2\right), \end{aligned}$$

and the analogue of it for  $\mathbb{E}[(\sigma'(v \cdot g(WX)) - \sigma'(v \cdot g(W\vartheta)))^2]$  follows similarly.  $\square$

## B.2 POPULATION HESSIAN ESTIMATES

In this section, we study the population Hessian matrices. Our large  $\lambda$  approximations will be uniform over compact sets in parameter space. Thus, we will use  $B$  to denote a ball in parameter space, so that any constant dependencies on the choice of  $B$  are just dependencies on its radius.

**Lemma B.4.** For  $(v, W) \in B$  the first layer block of the population Hessian matrix satisfy

$$\|\mathbb{E}[\nabla_{vv}^2 L] - \frac{1}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma'(v \cdot g(W\vartheta))g(W\vartheta)^{\otimes 2}\|_{\text{op}} \lesssim_{K,B} \psi_2(\vartheta, W, \lambda) + \frac{1}{\lambda}, \quad (\text{B.10})$$

where  $\psi_2$  was defined in (B.6). For  $(v, W) \in B$  with  $W \in \mathcal{W}_0^c$ , the diagonal second layer blocks of the population Hessian satisfy

$$\begin{aligned} & \|\mathbb{E}[\nabla_{W_i W_i}^2 L] - \frac{v_i^2}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma(vg(m_i^\vartheta))(1 - \sigma(vg(m_i^\vartheta)))F\left(\sqrt{\frac{\lambda}{R_{ii}}}m_i^\vartheta\right)\vartheta^{\otimes 2}\|_{\text{op}} \\ & \lesssim \max_{\vartheta \in \{\pm\mu, \pm\nu\}} \left( \frac{1}{\sqrt{\lambda}} \cdot \bar{F}\left(\sqrt{\frac{\lambda}{R_{ii}}}|m_i^\vartheta|\right)^{1/2} + \frac{1}{\lambda} \right). \quad (\text{B.11}) \end{aligned}$$

where we remind the reader that  $m_i^\vartheta = W_i \cdot \vartheta$ ,  $R_{ii} = W_i \cdot W_i$ , and  $F$  and  $\bar{F}$  are the cdf and tail of a standard Gaussian respectively.

*Proof.* We begin with the estimate on the  $vv$  block. Recall from (B.4) that

$$\mathbb{E}[\nabla_{vv}^2 L] = \mathbb{E}[\sigma'(v \cdot g(WY))g(WY)^{\otimes 2}]. \quad (\text{B.12})$$

By conditioning on the value of  $\vartheta$ , writing  $\mathbb{E}_\vartheta$  for the conditional expectation, it suffices to bound

$$\|\mathbb{E}_\vartheta[\sigma'(v \cdot g(WY))g(WY)^{\otimes 2}] - \sigma'(v \cdot g(W\vartheta))g(W\vartheta)^{\otimes 2}\|_{\text{op}},$$

for each  $\vartheta \in \{\pm\mu, \pm\nu\}$ . To this end, we write

$$\begin{aligned} & \mathbb{E}_\vartheta [\sigma'(v \cdot g(WY))g(WY)^{\otimes 2}] - \sigma'(v \cdot g(W\vartheta))g(W\vartheta)^{\otimes 2} \\ &= \mathbb{E}_\vartheta[\sigma'(v \cdot g(W\vartheta)) \cdot (g(WY)^{\otimes 2} - g(W\vartheta)^{\otimes 2})] + \mathbb{E}_\vartheta[(\sigma'(v \cdot g(WY)) - \sigma'(v \cdot g(W\vartheta)))g(WY)^{\otimes 2}] \\ &=: (i) + (ii) \end{aligned}$$

By uniform boundedness of  $\sigma'$ , we then have by (B.5) that  $\|(i)\|_{\text{op}} \lesssim_K \psi_2(\vartheta, W, \lambda)$ . On the other hand, by (B.7) and (B.8),

$$\|(ii)\|_{\text{op}} \leq (\mathbb{E}[(\sigma'(v \cdot g(WY)) - \sigma'(v \cdot g(W\vartheta)))^2])^{1/2} (\mathbb{E}\|g(WY)\|_{\text{op}}^2)^{1/2} \lesssim_{B,K} \frac{1}{\lambda}.$$

We now turn to the  $W_i W_i$  block of the population Hessian. Recall from (B.4) and the fact that  $W \in \mathcal{W}_0^c$ , that

$$\mathbb{E}[\nabla_{W_i W_i}^2 L] = \mathbb{E}[\hat{y}(1 - \hat{y})v_i^2 g'(W_i \cdot Y)Y^{\otimes 2}],$$

(since  $(g'(x))^2 = g'(x)$ ). We now aim to show the following two bounds, the second of which will give the desired bound of (B.11): for each  $i$ , we have that

$$\|\mathbb{E}[v_i^2 g'(W_i \cdot Y)Y^{\otimes 2}] - \frac{v_i^2}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} F\left(\sqrt{\frac{\lambda}{R_{ii}}} m_i^\vartheta\right) \vartheta^{\otimes 2}\| \lesssim \max_{\vartheta \in \{\pm\mu, \pm\nu\}} \left(\frac{1}{\sqrt{\lambda}} \cdot \bar{F}\left(\sqrt{\frac{\lambda}{R_{ii}}} |m_i^\vartheta|\right)\right)^{1/2} + \frac{1}{\lambda} \quad (\text{B.13})$$

and

$$\begin{aligned} & \|\mathbb{E}[\hat{y}(1 - \hat{y})v_i^2 g'(W_i \cdot Y)Y^{\otimes 2}] - \frac{v_i^2}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma(v \cdot g(W\vartheta))(1 - \sigma(v \cdot g(W\vartheta)))F\left(\sqrt{\frac{\lambda}{R_{ii}}} m_i^\vartheta\right) \vartheta^{\otimes 2}\| \\ & \lesssim \max_{\vartheta \in \{\pm\mu, \pm\nu\}} \left(\frac{1}{\sqrt{\lambda}} \cdot \bar{F}\left(\sqrt{\frac{\lambda}{R_{ii}}} |m_i^\vartheta|\right)\right)^{1/2} + \frac{1}{\lambda}. \quad (\text{B.14}) \end{aligned}$$

Conditioning on  $\vartheta$  and pulling out  $v_i^2$  in (B.13), it suffices to bound

$$\begin{aligned} \mathbb{E}_\vartheta[g'(W_i \cdot Y)Y^{\otimes 2}] &= \mathbb{P}_\vartheta(W_i \cdot Y > 0) \vartheta^{\otimes 2} + 2 \text{Sym} \mathbb{E}_\vartheta[\mathbb{1}_{W_i \cdot Y > 0} Z] \otimes \vartheta + \mathbb{E}_\vartheta[\mathbb{1}_{W_i \cdot Y > 0} Z^{\otimes 2}] \\ &=: (i) + (ii) + (iii). \end{aligned}$$

Term (i) is exactly the term to which we are comparing. In particular

$$\mathbb{P}_\vartheta(W_i \cdot Y > 0) = P(W_i \cdot \vartheta > W_i \cdot Z) = F\left(\sqrt{\frac{\lambda}{R_{ii}}} m_i^\vartheta\right).$$

We thus wish to bound the operator norms of (ii) and (iii). For term (ii), observe that for  $u \in \mathbb{S}^{d-1}$ , since  $\mathbb{E}_\vartheta \langle u, Z \rangle = 0$ , we have that

$$\mathbb{E}_\vartheta[\mathbb{1}_{W_i \cdot Y > 0} \langle u, Z \rangle] = -\mathbb{E}_\vartheta[\mathbb{1}_{W_i \cdot Y < 0} \langle u, Z \rangle].$$

Thus we may bound this, by Cauchy-Schwarz by

$$|\mathbb{E}_m[\mathbb{1}_{W_i \cdot Y > 0} \langle u, Z \rangle]| \leq \mathbb{E}[\langle u, Z \rangle^2]^{1/2} \mathbb{P}_\vartheta((W_i \cdot Y) \text{sgn}(m_i^\vartheta) < 0)^{1/2} = \frac{1}{\sqrt{\lambda}} \cdot \bar{F}\left(\sqrt{\frac{\lambda}{R_{ii}}} |m_i^\vartheta|\right)^{1/2}.$$

It follows, using that  $\|\vartheta\| = 1$ , that

$$\|(ii)\|_{\text{op}} = \sup_u |\langle u, (ii)u \rangle| \leq \frac{2}{\sqrt{\lambda}} \cdot \bar{F}\left(\sqrt{\frac{\lambda}{R_{ii}}} |m_i^\vartheta|\right)^{1/2}.$$



For term (ii), since  $|\sigma| \leq 1$ ,

$$\|(ii)\|_{\text{op}} \leq \|\mathbb{E}_\vartheta[g(WY)^{\otimes 2}] - g(W\vartheta)^{\otimes 2}\|_{\text{op}} \lesssim_K \psi_2(\vartheta, W, \lambda),$$

by (B.5). Combining these two and averaging over  $\vartheta$  yields (B.16).

Recalling (B.3), and conditioning on the mean being  $\vartheta$ , it suffices to control the operator norm of

$$\mathbb{E}_\vartheta[(\hat{y} - y_\vartheta)^2 g'(W_i \cdot Y) Y^{\otimes 2}] - A_\vartheta$$

where  $A_\vartheta$  is the corresponding summand in (B.18). Fix a  $\vartheta$  with  $y_\vartheta = 0$ , the choices being analogous. In this case we wish to bound the difference,

$$\mathbb{E}_\vartheta[\sigma(v \cdot g(WY))^2 g'(W_i \cdot Y) Y^{\otimes 2}] - \sigma(v \cdot g(W\vartheta))^2 F\left(\sqrt{\frac{\lambda}{R_{ii}}} m_i^\vartheta\right) \vartheta^{\otimes 2} = (i) + (ii),$$

where

$$\begin{aligned} (i) &:= \mathbb{E}_\vartheta[(\sigma(v \cdot g(WY))^2 - \sigma(v \cdot g(W\vartheta))^2) g'(W_i \cdot Y) Y^{\otimes 2}], \\ (ii) &:= \sigma(v \cdot g(W\vartheta))^2 \left( \mathbb{E}_\vartheta[g'(W_i \cdot Y) Y^{\otimes 2}] - F\left(\sqrt{\frac{\lambda}{R_{ii}}} m_i^\vartheta\right) \vartheta^{\otimes 2} \right). \end{aligned}$$

By (B.13) and the boundedness of  $\sigma$ ,

$$\|(ii)\|_{\text{op}} \lesssim \max_{\vartheta \in \{\pm\mu, \pm\nu\}} \left( \frac{1}{\sqrt{\lambda}} \cdot \bar{F}\left(\sqrt{\frac{\lambda}{R_{ii}}} |m_i^\vartheta|\right)^{1/2} + \frac{1}{\lambda} \right).$$

For (i), we bound its operator norm by Cauchy–Schwarz, as

$$\|(i)\|_{\text{op}} \leq \mathbb{E}[(\sigma(v \cdot g(W_i \cdot Y))^2 - \sigma(v \cdot g(W_i \cdot \vartheta))^2)^2]^{1/2} \sup_{\|u\|_2 \leq 1} \mathbb{E}_\vartheta[\langle u, Y \rangle^4],$$

where we used that  $|g'(W_i \cdot Y)| \leq 1$ . Since  $\mathbb{E}_\vartheta[\langle u, Y \rangle^4] \lesssim 1 + \frac{1}{\lambda} + \frac{1}{\lambda^2}$  for  $\|u\|_2 \leq 1$ , we have by (B.8) that  $\|(i)\| \lesssim_B \frac{1}{\lambda}$ , from which the desired follows.  $\square$

## C ANALYSIS OF THE SGD TRAJECTORY

Our goal in this section is to show the following results for the SGD for our two classes of classification tasks. The first is our main result on stochastic gradient descent for the mixture of  $k$ -GMM's via a single-layer network.

**Proposition C.1.** *For every  $\epsilon, \beta > 0$ , there exists  $T_0$  such that for all  $T_f > T_0$  and all  $\lambda$  large, the SGD for the  $k$ -GMM with step size  $\delta = O(1/d)$ , initialized from  $\mathcal{N}(0, I_d/d)$ , does the following for all  $c \in [k]$  with probability  $1 - o_d(1)$ :*

1. *There exists  $L(\beta)$  (independent of  $\epsilon, \lambda$ ) such that  $\|\mathbf{x}_\ell^c\| \leq L$  for all  $\ell \in [T_0\delta^{-1}, T_f\delta^{-1}]$ ;*
2. *There exists  $\eta(\beta) > 0$  (independent of  $\epsilon, \lambda$ ) such that  $\mathbf{x}_\ell^c$  is within  $O(\epsilon + \lambda^{-1})$  distance of a point in  $\text{Span}(\mu_1, \dots, \mu_k)$  having  $\|\mathbf{x}^c\| > \eta$ .*

The following is our analogous main result for the XOR GMM with two layer networks.

**Proposition C.2.** *For every  $\epsilon > 0$ , there exists  $T_0$  such that for all  $T_f > T_0$  and all  $\lambda$  large, the SGD for the XOR GMM with width  $K$  fixed, with step-size  $\delta = O(1/d)$ , initialized from  $\mathcal{N}(0, I_d/d)$  in its first layer weights, and  $\mathcal{N}(0, 1)$  i.i.d. in its second layer weights, with probability  $1 - o_d(1)$ , satisfies:*

1. *There is an  $L(K, \beta, v(\mathbf{x}_0))$  (independent of  $\epsilon, \lambda$ ) such that  $\|W_i(\mathbf{x}_\ell)\| \leq L$  for all  $i \leq K$ , and  $v(\mathbf{x}_\ell) \leq L$  for all  $\ell \in [T_0\delta^{-1}, T_f\delta^{-1}]$ ;*
2. *If  $\beta < 1/8$ , there exists  $\eta(K, \beta, v(\mathbf{x}_0)) > 0$  (independent of  $\epsilon, \lambda$ ) such that  $\mathbf{x}_\ell$  has  $|v_i(\mathbf{x}_\ell)| > \eta$  and  $\max_{\vartheta \in \{\pm\mu, \pm\nu\}} W_i(\mathbf{x}_\ell) \cdot \vartheta > \eta$  for all  $\ell \in [T_0\delta^{-1}, T_f\delta^{-1}]$  and  $i \leq K$ . and furthermore, it has  $W_i(\mathbf{x}_\ell)$  lives in  $\text{Span}(\mu, \nu)$  and  $v(\mathbf{x}_\ell)$  lives in  $\text{Span}((g(W(\mathbf{x}_*)\vartheta))_{\vartheta \in \{\pm\mu, \pm\nu\}})$  up to error  $\epsilon + \lambda^{-1/2}$ .*

Note, the dependencies of  $L$  and  $\eta$  on the absolute initial second layer value  $|v(\mathbf{x}_0)|$  are continuous.

Our approach goes by first recalling the main result of Ben Arous et al. (2022) regarding a limit theorem as  $d \rightarrow \infty$  for the trajectories of summary statistics of stochastic gradient descent. We then apply this result to the task of classifying  $k$  Gaussian mixtures, obtaining ballistic limits for the SGD in Theorem C.7 that may be of independent interest. We then analyze the ballistic limits to establish Proposition C.1. Then, we recall the ballistic limits for the XOR Gaussian mixture from Ben Arous et al. (2022) and similarly establish Proposition C.2. Since we are imagining the dimensions of the parameter space and dimension space, and the step size to all be scaling with relation to one another, we use a dummy index  $n$ , in this section to encode their mutual relationship, namely  $d = d_n, p = p_n$  and  $\delta = \delta_n$  and then  $n \rightarrow \infty$ .

Also, to compress notation throughout this section, we will combine the loss function with the regularizer  $\frac{\beta}{2} \|\mathbf{x}\|^2$  so that the stochastic gradient descent updates are indeed making gradient updates with respect to the new loss:

$$\mathbf{x}_\ell = \mathbf{x}_{\ell-1} - \delta \nabla \bar{L}(\mathbf{x}_{\ell-1}, \mathbf{Y}^\ell) \quad \text{where} \quad \bar{L}(\mathbf{x}) = L(\mathbf{x}) + \frac{\beta}{2} \|\mathbf{x}\|^2.$$

### C.1 RECALLING THE EFFECTIVE DYNAMICS FOR SUMMARY STATISTICS

Suppose that we are given a sequence of functions  $\mathbf{u}_n \in C^1(\mathbb{R}^{p_n}; \mathbb{R}^k)$  for some fixed  $k$ , where  $\mathbf{u}_n(x) = (u_1^n(x), \dots, u_k^n(x))$ , and our goal is to understand the evolution of  $\mathbf{u}_n(\mathbf{x}_\ell)$ .

In what follows, let  $H(\mathbf{x}, \mathbf{Y}) = \bar{L}_n(\mathbf{x}, \mathbf{Y}) - \Phi(\mathbf{x})$ , where  $\Phi(\mathbf{x}) = \mathbb{E}[\bar{L}_n(\mathbf{x}, \mathbf{Y})]$ . (Note that since the regularizer term is non-random,  $H$  is the same with  $L$  in place of  $\bar{L}$ .) Throughout the following, we suppress the dependence of  $H$  on  $\mathbf{Y}$  and instead view  $H$  as a random function of  $\mathbf{x}$ , denoted  $H(\mathbf{x})$ . We let  $V(\mathbf{x}) = \mathbb{E}_{\mathbf{Y}}[\nabla H(\mathbf{x}) \otimes \nabla H(\mathbf{x})]$  denote the covariance matrix for  $\nabla H$  at  $\mathbf{x}$ .

**Definition C.3.** A triple  $(\mathbf{u}_n, L_n, P_n)$  is  $\delta_n$ -localizable with localizing sequence  $(E_K)_K$  if there is an exhaustion by compacts  $(E_K)_K$  of  $\mathbb{R}^k$ , and constants  $C_K$  (independent of  $n$ ) such that

1.  $\max_i \sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \|\nabla^2 u_i^n\|_{\text{op}} \leq C_K \cdot \delta_n^{-1/2}$ , and  $\max_i \sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \|\nabla^3 u_i^n\|_{\text{op}} \leq C_K$ ;
2.  $\sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \|\nabla \Phi\| \leq C_K$ , and  $\sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \mathbb{E}[\|\nabla H\|^8] \leq C_K \delta_n^{-4}$ ;
3.  $\max_i \sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \mathbb{E}[\langle \nabla H, \nabla u_i^n \rangle^4] \leq C_K \delta_n^{-2}$ , and  
 $\max_i \sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \mathbb{E}[\langle \nabla^2 u_i^n, \nabla H \otimes \nabla H - V \rangle^2] = o(\delta_n^{-3})$ .

Define the following first and second-order differential operators,

$$\mathcal{A}_n = \sum_i \partial_i \Phi \partial_i, \quad \text{and} \quad \mathcal{L}_n = \frac{1}{2} \sum_{i,j} V_{ij} \partial_i \partial_j. \quad (\text{C.1})$$

Alternatively written,  $\mathcal{A}_n = \langle \nabla \Phi, \nabla \rangle$  and  $\mathcal{L}_n = \frac{1}{2} \langle V, \nabla^2 \rangle$ . Let  $J_n$  denote the Jacobian matrix  $\nabla \mathbf{u}_n$ .

**Definition C.4.** A family of summary statistics  $(\mathbf{u}_n)$  are asymptotically closable for step-size  $\delta_n$  if  $(\mathbf{u}_n, L_n, P_n)$  are  $\delta_n$ -localizable with localizing sequence  $(E_K)_K$ , and furthermore there exist locally Lipschitz functions  $\mathbf{h} : \mathbb{R}^k \rightarrow \mathbb{R}^k$  and  $\Sigma : \mathbb{R}^k \rightarrow \mathbb{R}^{k \times k}$ , such that

$$\sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \|(-\mathcal{A}_n + \delta_n \mathcal{L}_n) \mathbf{u}_n(\mathbf{x}) - \mathbf{h}(\mathbf{u}_n(\mathbf{x}))\| \rightarrow 0, \quad (\text{C.2})$$

$$\sup_{\mathbf{x} \in \mathbf{u}_n^{-1}(E_K)} \|\delta_n J_n V J_n^T - \Sigma(\mathbf{u}_n(\mathbf{x}))\| \rightarrow 0. \quad (\text{C.3})$$

In this case we call  $\mathbf{h}$  the effective drift, and  $\Sigma$  the effective volatility.

For a function  $f$  and measure  $\mu$  we let  $f_* \mu$  denote the push-forward of  $\mu$ . The main result of Ben Arous et al. (2022) was the following limit theorem for SGD trajectories as  $n \rightarrow \infty$ .



**Theorem C.5** (Ben Arous et al., 2022, Theorem 2.2). *Let  $(\mathbf{x}_\ell^{\delta_n})_\ell$  be stochastic gradient descent initialized from  $\mathbf{x}_0 \sim \mu_n$  for  $\mu_n \in \mathcal{M}_1(\mathbb{R}^n)$  with learning rate  $\delta_n$  for the loss  $L_n(\cdot, \cdot)$  and data distribution  $P_n$ . For a family of summary statistics  $\mathbf{u}_n = (u_i^n)_{i=1}^k$ , let  $(\mathbf{u}_n(t))_t$  be the linear interpolation of  $(\mathbf{u}_n(\mathbf{x}_{\lfloor t\delta_n^{-1} \rfloor}^{\delta_n}))_t$ .*

*Suppose that  $\mathbf{u}_n$  are asymptotically closable with learning rate  $\delta_n$ , effective drift  $\mathbf{h}$ , and effective volatility  $\Sigma$ , and that the pushforward of the initial data has  $(\mathbf{u}_n)_* \mu_n \rightarrow \nu$  weakly for some  $\nu \in \mathcal{M}_1(\mathbb{R}^k)$ . Then  $(\mathbf{u}_n(t))_t \rightarrow (\mathbf{u}_t)_t$  weakly as  $n \rightarrow \infty$ , where  $\mathbf{u}_t$  solves*

$$d\mathbf{u}_t = \mathbf{h}(\mathbf{u}_t)dt + \sqrt{\Sigma(\mathbf{u}_t)}d\mathbf{B}_t. \quad (\text{C.4})$$

*initialized from  $\nu$ , where  $\mathbf{B}_t$  is a standard Brownian motion in  $\mathbb{R}^k$ .*

As a result, we can read off in the ballistic regime the following finite- $n$  approximation result.

**Lemma C.6.** *In the setting of Theorem C.5, if the limiting dynamics have  $\Sigma \equiv 0$ , then for every  $T$ , we have with probability  $1 - o_d(1)$ ,*

$$\|\mathbf{u}(\mathbf{x}_{\lfloor \delta^{-1}t \rfloor}) - \mathbf{u}_t\|_{C[0,T]} = o_d(1).$$

## C.2 1-LAYER NETWORKS: BALLISTIC LIMITS

Our first aim in this section is to show that the limit theorem of Theorem C.5 indeed applies to the mixture of  $k$  Gaussians.

Recall the definitions of the data distribution, and the cross-entropy loss for  $x \in \mathbb{R}^{kd}$  (i.e.,  $x_a \in \mathbb{R}^d$  for  $a \in [k]$ ) from (3.2). Recall that  $\bar{m}_{ab} = \mu_a \cdot \mu_b$ , and that in order for the problem to be linearly classifiable (and therefore solvable with this loss) no mean should be the weighted barycenter of the other means, i.e.,  $p_a \mu_a \neq \frac{1}{k} \sum_b p_b \mu_b$ . For this task, the input dimension  $d = n$ , the parameter dimension  $p_n = kn$  and the step-size will be taken to be  $\delta = O(1/n) = O(1/d)$ , and we will tend to use  $d$  rather than the dummy index  $n$ . It will be helpful to write  $Y_a = \mu_a + Z_\lambda$ , and recall for  $a, c \in [k]$ , the Gibbs probability  $\pi_{Y_a}(c) = \pi_{Y_a}(c; \mathbf{x})$  as defined in (A.4).

### C.2.1 THE SUMMARY STATISTICS

We will show that the following family of functions form a set of localizable summary statistics  $\mathbf{u}_n(x) = ((m_{ab}(x))_{a,b}, (\mathbf{R}_{ab}^\perp)_{ab})$ .

$$\begin{aligned} \mathbf{m} &= (m_{ab})_{a,b \in [k]} & \text{where} & & m_{ab} &= x_a \cdot \mu_b \\ \mathbf{R}^\perp &= (\mathbf{R}_{ab}^\perp)_{a,b \in [k]} & \text{where} & & \mathbf{R}_{ab}^\perp &= x_a^\perp \cdot x_b^\perp \end{aligned}$$

where  $x_a^\perp$  denotes the part of  $x_a$  orthogonal to  $\text{Span}(\mu_1, \dots, \mu_k)$ , i.e.,  $x_a^\perp = \mathbf{P}^\perp x_a$ , where we use  $\mathbf{P}^\perp$  to be the projection operator into the orthogonal complement of  $\text{Span}(\mu_1, \dots, \mu_k)$ .

It is not hard to see that the law of the full sequence  $(\pi_{Y_a}(c))_{a,c}$  (and therefore its moments etc...) only depend on  $x$  through  $\mathbf{m}$  and  $\mathbf{R}^\perp$ , and therefore they have no finite  $d$  dependence. The following describes the ODEs obtained by taking the simultaneous  $d \rightarrow \infty$  and  $\delta = O(1/d)$ , limit of the SGD trajectory in its summary statistics.

**Theorem C.7.** *If the step sizes  $\delta = c_\delta/d$ , then the summary statistics  $\mathbf{u}_n = (\mathbf{m}, \mathbf{R}^\perp)$  are  $\delta$ -localizable, and satisfy the following ballistic limit as  $d \rightarrow \infty$ :*

$$\dot{m}_{ab}(t) = p_a \bar{m}_{ab} - \beta m_{ab} - \sum_{c \in [k]} p_c (\bar{m}_{cb} P_a^c + Q_a^{c, \mu_b}), \quad (\text{C.5})$$

$$\dot{\mathbf{R}}_{ab}^\perp(t) = -\beta \mathbf{R}_{ab}^\perp + \sum_{c \in [k]} (p_a Q_a^{c, \mathbf{R}_{bb}^\perp} + p_b Q_b^{c, \mathbf{R}_{aa}^\perp}) - \sum_{c \in [k]} \frac{c_\delta p_c}{\lambda} \mathbb{E}[(\pi_{Y_c}(a) - \mathbf{1}_{a=c})(\pi_{Y_c}(b) - \mathbf{1}_{b=c})], \quad (\text{C.6})$$

where  $P_a^c, Q_a^{c,v}$  are the following Gaussian integrals:

$$P_a^c(\mathbf{m}, \mathbf{R}^\perp) = \mathbb{E}[\pi_{Y_c}(a)] \quad \text{and} \quad Q_a^{c,v}(\mathbf{m}, \mathbf{R}^\perp) = \mathbb{E}[\langle Z_\lambda, v \rangle \pi_{Y_c}(a)] \quad (\text{C.7})$$

and where if  $v \perp \text{Span}(\mu_1, \dots, \mu_k)$  we use the shorthand  $Q_a^{c,r}$  for  $r = \|v\|^2$  since for such  $v$ ,  $Q_a^{c,v}$  only depends on  $v$  through  $\|v\|^2$ . The case where  $\delta = o(1/d)$  is read-off by formally setting  $c_\delta = 0$ .

*Remark 2.* While a priori, it may appear that the quantity  $Q_a^{c,v}$  in (C.7) depends on the dimension  $d$ , if  $v$  is one of  $(\mu_1, \dots, \mu_k)$ , or if it is simply orthogonal to  $\text{Span}(\mu_1, \dots, \mu_k)$ , then  $Q_a^{c,v}$  does not depend on  $d$ . Such cases are the only ones appearing in (C.5)–(C.6). The same can be said for the expectation in the last term of (C.7).

*Proof.* Theorem C.7 will follow from an application of Theorem C.5 to the summary statistics  $\mathbf{m}, \mathbf{R}^\perp$  for the  $k$ -GMM, so we start by verifying that the problem fits the assumptions of  $\delta$ -localizability and asymptotic closability.

*Verifying  $\delta$ -localizability* Our aim is to verify the conditions of  $\delta$  localizability from Definition C.3. Let us begin with some calculations, recalling that the Jacobian matrix  $J = \nabla \mathbf{u}$ . Let  $\nabla_a$  denote the derivative in the  $\mathbb{R}^d$  coordinates corresponding to  $x_a$ . For  $a, b, c \in [k]$ ,

$$\nabla_c m_{ab} = \begin{cases} \mu_b & \text{if } c = a \\ 0 & \text{else} \end{cases}, \quad \text{and} \quad \nabla_c R_{ab}^\perp = \begin{cases} x_a^\perp & a \neq b = c \\ x_b^\perp & b \neq a = c \\ 2x_a^\perp & a = b = c \\ 0 & \text{else} \end{cases}. \quad (\text{C.8})$$

Continuing, all higher derivatives of  $m_{ab}$  are zero. The Hessian of  $R_{ab}^\perp$  is given by

$$\nabla_{ab} m_{cd} = 0 \quad \text{and} \quad \nabla_{ab} R_{ab}^\perp = \begin{cases} \mathbf{P}^\perp & \text{if } a \neq b \\ 2\mathbf{P}^\perp & \text{if } a = b \end{cases}. \quad (\text{C.9})$$

(and other blocks are zero), and higher derivatives of  $R_{ab}^\perp$  are also zero.

Let us also express the loss function as equal in distribution to a random variable whose law depends only on  $\mathbf{m}$  and  $\mathbf{R}^\perp$ . For ease of notation, let

$$R_{ab} = \langle x_a, x_b \rangle \quad \text{and} \quad Y_a = \mu_a + Z_\lambda.$$

Recalling (3.2) and adding in the regularizer, we can write for each fixed  $\mathbf{x}$ ,

$$\bar{L}(\mathbf{x}, \mathbf{Y}) \stackrel{(d)}{=} - \sum_{a \in [k]} y_a (m_{aa} + G_a) + \sum_{b \in [k]} y_b \log \sum_{a \in [k]} e^{m_{ab} + G_a} + \frac{\beta}{2} \sum_{a \in [k]} R_{aa}, \quad (\text{C.10})$$

where  $(G_a)_a$  is a Gaussian vector with covariance matrix  $(R_{ab})$ , and  $y$  is a uniformly drawn 1-hot vector in  $\mathbb{R}^k$ . Observe that the law of  $\bar{L}$  only depends on  $\mathbf{x}$  through the values of the summary statistics.

**Lemma C.8.** *Suppose  $\delta = O(1/d)$ ,  $\lambda > 0$  is fixed (or growing with  $d$ ) and  $\beta > 0$  is fixed. Then the family of summary statistics  $\mathbf{u} = (\mathbf{m}, \mathbf{R}^\perp)$  are  $\delta$ -localizable with balls.*

*Proof.* *Item (1) of  $\delta$ -localizability.* The first part is easily seen to be satisfied by (C.9) since the Hessians of statistics in  $\mathbf{m}$  are all zero, and the Hessian of  $R_{ab}^\perp$  is bounded in operator norm by  $2 + 2k$  using the triangle inequality and  $\|\mu_b\| = 1$ .

*Item (2) of  $\delta$ -localizability.* We begin with the bound on the population loss's norm. By taking expectation of (C.10)

$$\Phi(\mathbf{x}) = - \sum_{a \in [k]} p_a m_{aa} + \sum_{a \in [k]} p_a \mathbb{E} \left[ \log \sum_{b \in [k]} e^{\langle x_b, Y_a \rangle} \right] + \frac{\beta}{2} \sum_{a \in [k]} R_{aa}. \quad (\text{C.11})$$

Taking the derivative of this, we get for each  $c \in [k]$ ,

$$\nabla_c \Phi(\mathbf{x}) = -p_c \mu_c + \sum_{a \in [k]} p_a \mathbb{E}[\pi_{Y_a}(c) Y_a] + \beta \mathbf{x}^c, \quad (\text{C.12})$$

where  $Y_a = \mu_a + Z_\lambda$  and  $\pi$  is as in (A.4). Considering the norm of  $\|\nabla \Phi\|$ , we have

$$\|\nabla \Phi\| \leq k \max_c \|\nabla_c \Phi\| \lesssim k(1 + 1 + \lambda^{-1/2} + \beta \max_a R_{aa})$$

which is bounded by a constant  $C(K)$  for  $\mathbf{m}, \mathbf{R}^\perp$  in a ball of radius  $K$ .

Moving on to the bound on  $\mathbb{E}[\|\nabla H\|^8]$ , first observe using  $\nabla_c H = \nabla_c \bar{L} - \nabla_c \Phi$ , that

$$\nabla_c H = (p_c \mu_c - y_c(\mu_c + Z_\lambda)) + (\pi_Y(c)Y - \mathbb{E}[\pi_Y(c)Y]) \quad (\text{C.13})$$

Taking the norm and the 8'th moment, we use the fact that  $\|u + v\|^8 \lesssim (\|u\|^8 + \|v\|^8)$ , and the fact that  $\pi_{Y_a}(\cdot)$  is a probability mass function and therefore at most 1, to upper bound

$$\max_{c \in [k]} \mathbb{E}[\|\nabla_c H\|^8] \lesssim \sup_a \|\mu_a\|^8 + \mathbb{E}\|Z_\lambda\|^8 \leq 1 + O((d/\lambda)^4).$$

As long as  $\delta_n = O(1/d)$ , since  $\lambda$  is uniformly bounded away from zero, this will be bounded by a constant times  $\delta_n^{-4}$  as required.

*Item (3) of  $\delta$ -localizability.* We next turn to bounding the fourth moments of the directional derivatives, starting with the statistics  $m_{ab}$ :

$$\begin{aligned} \langle \nabla_c H, \nabla_a m_{ab} \rangle &= \langle (p_c - y_c)\mu_c, \mu_b \rangle - \langle y_c Z_\lambda, \mu_b \rangle + p_b \mathbb{E}[\pi_{Y_b}(c)] - y_b \pi_{Y_b}(c) \\ &\quad + \sum_{l \in [k]} \left( p_l \mathbb{E}[\pi_{Y_l}(c) \langle \mu_b, Z_\lambda \rangle] + y_l \pi_{Y_l}(c) \langle \mu_b, Z_\lambda \rangle \right). \end{aligned}$$

Taking the fourth moment, again bounding things by the fourth moments of the terms individually up to a universal constant, then taking an expected value, we see that the first term is bounded by 1, the second by the fourth moment of a Gaussian random variable with variance  $1/\lambda$ , i.e., by  $C/\lambda^2$ , the third and fourth by 1 since  $\pi$  is a probability distribution, and the summands individually have fourth moments bounded by  $C/\lambda^2$  for the same reason. Altogether, we get

$$\mathbb{E}[\langle \nabla_c H, \nabla_a m_{ab} \rangle^4] \lesssim 1 + O(1/\lambda^2),$$

satisfying the requisite bound with room to spare. Turning to the directional derivative in the direction of  $R_{ab}^\perp$ , note that derivatives of  $R_{ab}^\perp$  are orthogonal to  $(\mu_1, \dots, \mu_k)$  so for  $a \neq b$ ,

$$\langle \nabla_c H, \nabla_b R_{ab}^\perp \rangle = \langle \nabla_c H, \nabla_b R_{ba}^\perp \rangle = \left( \mathbb{E}[\pi_Y(c) \langle x_a^\perp, Z_\lambda \rangle] - \pi_Y(c) \langle x_a^\perp, Z_\lambda \rangle \right) - y_c \langle x_a^\perp, Z_\lambda \rangle.$$

Considering the fourth moment of the above, using that  $\pi$  is bounded by 1, and  $\langle x_a^\perp, Z_\lambda \rangle$  is distributed as a Gaussian random variable with variance  $R_{aa}^\perp/\lambda$ , we obtain

$$\mathbb{E}[\langle \nabla_c H, \nabla_b R_{ab}^\perp \rangle^4] \lesssim (R_{aa}^\perp/\lambda)^2,$$

which is bounded by  $C(K)$  while  $R_{aa}^\perp$  is bounded by  $K$ . The diagonal  $a = b$  is the same up to a factor of 2. The last thing to check is the second part of item (3) in  $\delta_n$ -localizability. For this purpose, recall that  $V(x) = \mathbb{E}[\nabla H^{\otimes 2}]$ , and notice from equation C.13 that for the  $k$ -GMM we have  $\nabla_c H \otimes \nabla_d H = ((\pi_Y(c) - y_c)Y - \mathbb{E}[(\pi_Y(c) - y_c)Y]) \otimes ((\pi_Y(d) - y_d)Y - \mathbb{E}[(\pi_Y(d) - y_d)Y]))$

Taking expected values, we get that the  $cd$ -block of  $V$  is given by

$$V_{cd} = \mathbb{E}[\nabla_c H \otimes \nabla_d H] = \text{Cov}\left((\pi_Y(c) - y_c)Y, (\pi_Y(d) - y_d)Y\right), \quad (\text{C.14})$$

where Cov is the covariance matrices associated to the vectors inside it.

For the second part of item (3) in localizability, we only need to consider the statistics  $R_{ab}^\perp$  since the second derivatives of  $m_{ab}$  are all zero. For the ballistic limit it is sufficient to use the bound

$$\mathbb{E}[\langle \nabla^2 R_{ab}^\perp, \nabla_c H \otimes \nabla_d H - V_{cd} \rangle^2] \lesssim \|\nabla^2 R_{ab}^\perp\|_{\text{op}}^2 \cdot \mathbb{E}[\|\nabla_c H\|^4].$$

The first term is bounded by 2 per (C.9). The second can be seen to be bounded via the 8th moment above by  $O((d/\lambda)^2)$  which is  $o(\delta^{-3})$  when  $\delta = O(1/d)$ .  $\square$

*Remark 3.* The reader might notice that there is a lot of room in the bounds above as compared to the allowed thresholds from the  $\delta$ -localizability conditions. The weaker conditions in localizability are to allow taking diffusive limit theorems about saddle points by rescaling the summary statistics by  $d$ -dependent factors, which can help with understand timescales to escape fixed point regions of the limits of Theorem C.7. This was explored in much detail for matrix and tensor PCA in Ben Arous et al. (2022).

*Calculating the drift and corrector* Now that we have verified that the  $\delta$ -localizability conditions apply to the summary statistics  $\mathbf{u} = (\mathbf{m}, \mathbf{R}^\perp)$ , we compute the limiting ODE one gets in the  $d \rightarrow \infty$  limit for these statistics. We will establish individual convergence of  $\mathcal{A}u$  to some  $f_u(\mathbf{u})$  and convergence of  $\delta\mathcal{L}u$  to some  $g_u(\mathbf{u})$  for each  $u \in \mathbf{u}$ , whence  $h$  from Definition C.4 equals  $-f + g$ .

Recall the differential operator  $\mathcal{A}$  from (C.1), the expression for  $\nabla\Phi$  from (C.12), and consider

$$\mathcal{A}m_{ab} = \sum_{c \in [k]} \langle \nabla_c \Phi, \nabla_c m_{ab} \rangle = \langle \nabla_a \Phi, \mu_b \rangle = -\langle p_a \mu_a, \mu_b \rangle + \mathbb{E}[\pi_{Y_c}(a) \langle Y, \mu_b \rangle] + \beta \langle x_a, \mu_b \rangle.$$

Recalling that  $\bar{m}_{ab} = \langle \mu_a, \mu_b \rangle$ , we get

$$\mathcal{A}m_{ab} = -p_a \bar{m}_{ab} + \beta m_{ab} + \sum_{l \in [k]} p_l \left( \bar{m}_{lb} \mathbb{E}[\pi_{Y_l}(a)] + \mathbb{E}[\pi_{Y_l}(a) \langle Z_\lambda, \mu_b \rangle] \right).$$

Notice that the two expected values are Gaussian expectations that only depend on  $x$  through  $(m_{cl})_c$  and  $(R_{lm}^\perp)_{l,m}$ . In particular, we can take the limit as  $d \rightarrow \infty$  to get the limiting drift function

$$f_{m_{ab}}(\mathbf{m}, \mathbf{R}^\perp) = -p_a \bar{m}_{ab} + \beta m_{ab} + \sum_{l \in [k]} p_l \left( \bar{m}_{lb} P_a^l + Q_a^{l, \mu_b} \right).$$

where  $P_a^l$  and  $Q_a^{l, \mu_b}$  are defined per (C.7). The contribution coming from  $\delta\mathcal{L}m_{ab}$  vanishes in the  $d \rightarrow \infty$  limit since the second derivative of  $m_{ab}$  is zero, i.e.,  $g_{m_{ab}} = 0$ .

For the drift function for  $R_{ab}^\perp$ , since  $\langle x_a^\perp, \mu_b \rangle = 0$  for all  $a, b$ , we get

$$\mathcal{A}R_{ab}^\perp = \beta R_{ab}^\perp + \sum_{l \in [k]} (p_a Q_a^{l, x_b^\perp} + p_b Q_b^{l, x_a^\perp}).$$

In particular, we get

$$f_{R_{ab}^\perp} = \beta R_{ab}^\perp + \sum_{l \in [k]} (p_a Q_a^{l, R_{bb}^\perp} + p_b Q_b^{l, R_{aa}^\perp}),$$

using  $Q_a^{l, R_{bb}^\perp} = Q_a^{l, x_b^\perp}$  since  $Q_a^{l, v}$  only depended on  $v$  through its norm when  $v \perp \text{Span}(\mu_1, \dots, \mu_k)$

There is also a contribution from  $\delta\mathcal{L}R_{ab}^\perp$ ; recall that  $\mathcal{L} = \frac{1}{2} \langle V, \nabla^2 \rangle$  and notice that

$$\mathcal{L}R_{ab}^\perp = \frac{1}{2} (\langle V_{ab}, \mathbf{P}^\perp \rangle + \langle V_{ba}, \mathbf{P}^\perp \rangle).$$

Plugging in for  $V_{ab}$  from (C.14), and expanding this out, a calculation yields

$$\begin{aligned} \mathcal{L}R_{ab}^\perp &= \sum_{c \in [k]} p_c \left( \mathbb{E}[\|Z_\lambda\|^2 (\pi_{Y_c}(a) - \mathbf{1}_{a=c})(\pi_{Y_c}(b) - \mathbf{1}_{b=c})] \right. \\ &\quad \left. - \langle \mathbb{E}[Z(\pi_{Y_c}(a) - \mathbf{1}_{a=c}), \mathbb{E}[Z(\pi_{Y_c}(b) - \mathbf{1}_{b=c})]] \rangle \right) + O(1), \end{aligned}$$

where to see that the extra terms are  $O(1)$ , we notice that the inner products of the  $\mu$ 's with each other and  $\mu$ 's with  $Z$ 's are all order 1. By Cauchy–Schwarz, the inner product of the expectations is at most  $O(\sqrt{d/\lambda})$  and will vanish when multiplied by  $\delta = O(1/d)$  and the  $d \rightarrow \infty$  limit is taken; on the other hand, the second moment term is of order  $d/\lambda$ . In particular, if  $\delta = c_\delta/d$ , we get

$$\delta\mathcal{L}R_{ab}^\perp = c_\delta \sum_{c \in [k]} p_c \mathbb{E} \left[ \|d^{-1/2} Z_\lambda\|^2 (\pi_{Y_c}(a) - \mathbf{1}_{a=c})(\pi_{Y_c}(b) - \mathbf{1}_{b=c}) \right] + o(1).$$

We claim that the  $d \rightarrow \infty$  limit of this gives

$$g_{R_{ab}^\perp} = \frac{c_\delta}{\lambda} \sum_{c \in [k]} p_c \mathbb{E} \left[ (\pi_{Y_c}(a) - \mathbf{1}_{a=c})(\pi_{Y_c}(b) - \mathbf{1}_{b=c}) \right] \quad (\text{C.15})$$

Notice that the expectation is bounded by 1 in absolute value and thus the above goes to 0 as  $\lambda \rightarrow \infty$ . In order to show (C.15), let us consider the term coming from  $\pi_{Y_c}(a)\pi_{Y_c}(b)$ , the other terms being

analogous or even easier since the indicator is deterministic. Using the fact that  $\pi$  are probabilities and therefore bounded by 1, note that

$$\left| \mathbb{E}[\|d^{-1/2}Z\|^2 \pi_{Y_c}(a)\pi_{Y_c}(b)] - \frac{1}{\lambda} \mathbb{E}[\pi_{Y_c}(a)\pi_{Y_c}(b)] \right| \leq \mathbb{E}\left[\left|\|d^{-1/2}Z\|^2 - \frac{1}{\lambda}\right|\right]$$

which goes to 0 from the standard fact that for a standard Gaussian vector  $G_d$  in  $\mathbb{R}^d$ , one has  $\mathbb{E}\left[\left(\|\frac{G_d}{\sqrt{d}}\|^2 - 1\right)^2\right] = O(1/d)$ .

The last thing to check is that in the ballistic regime where our summary statistics are not rescaled, the limiting dynamics are indeed an ODE, i.e., there is no limiting stochastic part. Towards that purpose, notice that in any ball of values for  $(\mathbf{m}, \mathbf{R})$ ,

$$\|JVJ^T\| = O(1),$$

using an  $O(1)$  bound on the operator norm of  $V$ , and noticing that  $\|J\|^2$  is bounded by a constant plus the sums of  $R_{ab}^\perp$ . Therefore when multiplied by  $\delta = o(1)$  this vanishes in the limit and therefore the diffusion matrix  $\Sigma$  is identically zero.  $\square$

The following confines this  $\lambda$ -finite dynamical system to a compact set for all times.

**Lemma C.9.** *For every  $\beta > 0$ , there exists  $L(\beta)$  such that for all  $\lambda$  large, the dynamics of Theorem C.7 stays inside the  $\ell^2$ -ball of radius  $L$  for all time.*

*Proof.* Notice that a naive bound on  $Q_a^{d,v}$  from (C.7) is at most  $\sqrt{\|v\|/\lambda}$  since  $\pi_{Y_c}(a)$  is bounded, and  $P_a^d$  is always at most 1. Plugging these bounds into Theorem C.7, together with the definition of  $\bar{m}_{ab}$  and the fact that  $\|\mu_a\| = 1$  for all  $a$ , yields the inequalities for all  $a, b$ ,

$$\begin{aligned} |\dot{m}_{ab}(t) + \beta m_{ab}| &\lesssim 1 + \sqrt{1/\lambda}, \\ |\dot{R}_{aa}^\perp + \beta R_{aa}^\perp| &\lesssim \sqrt{R_{aa}^\perp/\lambda} + \lambda^{-1} + \lambda^{-2}. \end{aligned}$$

By these, for  $\lambda$  larger than a fixed constant, we have  $\dot{R}_{aa}^\perp \leq 1 - \beta R_{aa}^\perp/2$  which Gronwall's inequality ensures will be bounded by  $1 + R_{aa}^\perp(0)$  for all times  $t \geq 0$ . Similarly, we get that for  $\lambda$  sufficiently large,  $|m_{ab}(t)|$  is bounded by  $3 + m_{ab}(0)$ .  $\square$

## C.2.2 THE ZERO-NOISE LIMIT

We now send  $\lambda \rightarrow \infty$ , or simply Taylor expand in the large  $\lambda$  limit, to understand the behavior of the limiting ODE's we derived when  $\lambda$  is large. Let

$$\bar{\pi}_c(a) = \bar{\pi}_c(a; \mathbf{x}) := \frac{e^{m_{ac}}}{\sum_{b \in [k]} e^{m_{bc}}}. \quad (\text{C.16})$$

This is the " $\lambda = \infty$ " value of  $\pi_{Y_c}(a)$ . The aim of this subsection is to establish the following.

**Proposition C.10.** *The  $\lambda \rightarrow \infty$  limit of the ODE system from Theorem C.7 is the following dynamical system:*

$$\dot{m}_{ab}(t) = p_a \bar{m}_{ab} - \beta m_{ab} - \sum_{c \in [k]} p_c \bar{\pi}_c(a) \bar{m}_{cb}, \quad (\text{C.17})$$

$$\dot{R}_{ab}^\perp(t) = -\beta R_{ab}^\perp. \quad (\text{C.18})$$

Moreover, at large finite  $\lambda$ , the difference of the drifts in (C.5)–(C.6) to the above drifts is  $O(\lambda^{-1})$ .

The main thing to prove is the following behavior of integrals of  $\pi_{Y_c}(a)$  as  $\lambda \rightarrow \infty$ .

**Lemma C.11.** *Recalling  $P_a^c$  and  $Q_a^{c,v}$  from (C.7), we have*

$$\begin{aligned} P_a^c &= \bar{\pi}_c(a) + O(1/\lambda). \\ Q_a^{c,v} &= O(\|v\|/\lambda). \end{aligned}$$

*Proof.* By Taylor expanding, we can write

$$\pi_{Y_c}(a) = \bar{\pi}_c(a) + \frac{(x_a \cdot Z)e^{x_a \cdot \mu_c}}{\sum_{b \in [k]} e^{x_b \cdot \mu_c}} - \frac{e^{x_a \cdot \mu_c} (\sum_{b \in [k]} (x_b \cdot Z)e^{x_b \cdot \mu_c})}{(\sum_{b \in [k]} e^{x_b \cdot \mu_c})^2} + O((\max_{b \in [k]} x_b \cdot Z)^2).$$

Taking an expectation of the right-hand side, noting that  $x_b \cdot Z = \mathcal{N}(0, \frac{R_{bb}}{\lambda})$ , everything on the right-hand side after  $\bar{\pi}_d(a)$  is  $O(1/\lambda)$  for  $R_{bb}^\perp$  that is  $O(1)$ .

For  $Q_a^{c,v}$ , using the Gaussian integration by parts formula and (A.7)

$$Q_a^{c,v} = \frac{1}{\lambda} \left( \mathbb{E}[(x_a \cdot v)\pi_{Y_c}(a)] - \mathbb{E}[(x^B \cdot v)_{\pi_{Y_c}} \pi_{Y_c}(a)] \right)$$

Since  $\pi \leq 1$  and  $x_b \cdot v \leq \sqrt{R_{bb}^\perp} \|v\|$ , this is easily seen to be  $O(\|v\| \sqrt{R_{bb}^\perp}/\lambda)$  as claimed.  $\square$

**Proof of Proposition C.10.** For any  $K$ , uniformly over all  $\mathbf{m}, \mathbf{R}$  in a ball of radius  $K$  about the origin, we claim that the limit of the drifts for each of those variables converge to the claimed  $\lambda \rightarrow \infty$  limiting drifts. This is obtained by applying the above lemma to the  $P$  and  $Q$  terms in the drifts in Theorem C.7, and finally the observation that  $\text{Var}(\pi_{Y_c}(a)\pi_{Y_c}(b)) \leq 1$  so that taking  $\lambda \rightarrow \infty$  the last two terms in the drift for  $R_{ab}^\perp$  in Theorem C.7 also vanish.  $\square$

The following gives a quantitative approximation of the ODE by its  $\lambda = \infty$  limit.

**Corollary C.12.** *The trajectories of the ODEs of Theorem C.7 and Proposition C.10 are within distance  $O(t/\lambda)$  of one another.*

*Proof.* Let  $\mathbf{u}$  and  $\tilde{\mathbf{u}}$  be the two solutions to the  $\lambda$  finite and  $\lambda = \infty$  ballistic dynamics respectively. Then, while  $\|\mathbf{u}\|, \|\tilde{\mathbf{u}}\| \leq L$ , we have by Proposition C.10 that

$$\|\dot{\mathbf{u}} - \dot{\tilde{\mathbf{u}}}\| \lesssim_L \lambda^{-1}.$$

Per Lemma C.9, both dynamics remain confined for a large enough  $L(\beta)$  for all times, and therefore integrating the above gives the claim.  $\square$

### C.3 LIVING IN SUBSPACE SPANNED BY THE MEANS

We now wish to show that the SGD trajectory lives in the span of the means. This can be done by showing that on the one hand,  $\mathbf{R}^\perp$  will stay as small as we want after an  $O(1)$  burn-in time, and on the other hand, towards the error being multiplicative in Definition 2.1, for every  $a$ ,  $\mathbf{x}_\ell^a$  needs to be non-negligible.

We first establish that this happens for the  $\lambda = \infty$  dynamics, then pull it back to the dynamics at  $\lambda$  finite but large via Corollary C.12.

**Lemma C.13.** *The solution to the dynamical system of Proposition C.10 is such that for all  $t \geq T_0(\epsilon)$ , it is within distance  $\epsilon$  of a point having  $R_{aa}^\perp = 0$  and  $\max_b |m_{ab}| > c_\beta > 0$  for every  $a \in [k]$ .*

*Proof.* By the expression from Proposition C.10 for the drift of  $R_{aa}^\perp$  for  $a \in [k]$ , the dynamical system has  $R_{aa}^\perp(t) = e^{-\beta t} R_{aa}^\perp(0)$ . In particular, for any  $\epsilon > 0$ , the has  $R_{aa}^\perp(t) < \epsilon$  for all  $t \geq T_0(\epsilon)$ .

We need to show for every  $a$ , in the solution of the dynamical system, some  $(m_{ab})_b$  is bounded away from zero after some small time. Let  $\mathcal{M}_0 = \bigcup_a \bigcap_c \{m_{ac} = 0\}$  be the set of  $(\mathbf{m})$  values we would like to ensure the dynamics stays away from. First observe that the  $\lambda = \infty$  dynamical system of Proposition C.10 is a gradient system for the energy function

$$\mathcal{H}(\mathbf{m}, \mathbf{R}^\perp) = - \sum_{a \in [k]} p_a m_{aa} + \sum_{a \in [k]} p_a \log \sum_{b \in [k]} e^{m_{ab}} + \frac{\beta}{2} (\|\mathbf{m}\|^2 + \|\mathbf{R}^\perp\|^2),$$

so it has no recurrent orbits. It thus suffices to show that for every point in  $\mathcal{M}_0$ , the quantity  $\max_b |m_{ab}|$  has a drift strictly bounded away from zero. If we show that, then the dynamics is guaranteed to leave a  $c_\beta$ -neighborhood of the set  $\mathcal{M}_0$  in a finite time (uniform by continuity considerations and we are already guaranteed that the dynamics stays in a compact set by Lemma C.9).

Consider a point such that  $m_{ab} = 0$  for all  $b \in [k]$ . There, by Proposition C.10,

$$\dot{m}_{ab} = \langle p_a \mu_a - \sum_{c \in [k]} p_c \bar{\pi}_c(a) \mu_c, \mu_b \rangle.$$

We need to show that the maximum over  $b$ , of the absolute values of these, is non-zero. Indeed, if it were 0 for all  $b \in [k]$ , then  $p_a \mu_a = \sum_c p_c \bar{\pi}_c(a) \mu_c$  because the difference of these vectors would be in  $\text{Span}(\mu_1, \dots, \mu_k)$  while being orthogonal to all of  $\mu_1, \dots, \mu_k$ . In turn, however, this is impossible by our assumption that  $\mu_1, \dots, \mu_k$  are linearly independent. Therefore, in the ball of radius  $L(\beta)$  about the origin, for every  $a$ , at least one of the drifts  $\dot{m}_{ab}$  is bounded away from zero uniformly.  $\square$

**Proof of Proposition C.1.** We have shown that the limit dynamics for the summary statistics of the SGD initialized from  $\mathcal{N}(0, I_d/d)$  is the dynamical system of Theorem C.7 initialized from the deterministic initialization  $m_{ab} \sim \delta_0$  and  $R_{ab}^\perp \sim \delta_0$  if  $a \neq b$  and  $\delta_1$  if  $a = b$ .

By Lemma C.9, there exists  $L(\beta)$  such that that dynamical system is confined to a ball of radius at most  $L$  for all time. Since  $\|\mathbf{x}^c\|^2$  is encoded by a smooth function of the summary statistics  $\mathbf{R}_{cc}^\perp(m_{cb})_b$  appearing in the dynamical system of Theorem C.7, this is transferred to the SGD  $\mathbf{x}_\ell$  via Lemma C.6 for a different constant  $L(\beta)$  for  $d$  sufficiently large.

For the second part, by Lemma C.13, the solution to the dynamical system of Theorem thm:k-GMM-ballistic-limit is at distance  $O(\lambda^{-1})$  from the solution of the dynamical system (with the same initialization) of Proposition C.10, for which Lemma C.13 applies. By Lemma C.6, these get pulled back to the summary statistics applied to the SGD itself (at  $d$  sufficiently large depending on  $\epsilon, \lambda$ ). We therefore deduce that for all  $\ell \geq T_0(\epsilon)\delta^{-1}$  steps, the SGD has  $R_{cc}^\perp(\mathbf{x}_\ell) \leq O(\epsilon + \lambda^{-1})$  and has  $\max_b |\langle \mathbf{x}_\ell^c, \mu_b \rangle| > c_\beta > 0$ . These imply the claim using that  $\|\mathbf{x}^c\| \geq \max_b |\langle \mathbf{x}^c, \mu_b \rangle|$  and that  $R_{cc}^\perp(\mathbf{x}^c)$  is by definition the projection of  $\mathbf{x}^c$  orthogonal to  $(\mu_1, \dots, \mu_k)$ .  $\square$

### C.3.1 FIXED POINT ANALYSIS WITH ORTHOGONAL MEANS

The behavior of the ODE system of Proposition C.10 can be sensitive to the relative location of the means  $(\mu_1, \dots, \mu_k)$ . In order to be able to make more precise statements about the alignment of the SGD with specific eigenvectors rather than just living in  $\text{Span}(\mu_1, \dots, \mu_k)$ , we specialize to the case where the means form an orthonormal family of vectors. Here, we can explicitly characterize the fixed points of Proposition C.10.

Namely, in this subsection, assume that  $\bar{m}_{ab} = \mathbf{1}_{a=b}$ . Then, by (C.17) any fixed point must satisfy

$$\beta m_{ab} = p_a \mathbf{1}_{a=b} - p_b \frac{e^{m_{ab}}}{\sum_c e^{m_{cb}}}. \quad (\text{C.19})$$

At a fixed point, the function  $\sum_c m_{cb}$  thus must equal 0. Also, if  $a, c \neq b$  then

$$m_{ab} - m_{cb} = -\frac{p_b}{\beta} \cdot \frac{e^{m_{ab}} - e^{m_{cb}}}{\sum_d e^{m_{ad}}}.$$

Since the function  $x + ce^x$  is strictly increasing, the only solutions to this are at  $m_{ab} = m_{cb}$  (so long as  $p_b > 0$ ). Combining the above two observations, at a fixed point,

$$m_{ab} = -\frac{1}{k-1} m_{bb} \quad \text{for all } a \neq b.$$

Plugging this in to (C.19), we find that at a fixed point

$$m_{bb} = \frac{p_b}{\beta} \left( 1 - \frac{e^{m_{bb}}}{(k-1)e^{-\frac{1}{k-1}m_{bb}} + e^{m_{bb}}} \right).$$

This can easily be seen to have a unique solution, and that solution must have  $m_{bb} \in (0, \frac{p_b}{\beta})$ .

Therefore, as long as  $(p_b)_{b \in [k]}$  are all positive, the dynamical system of Proposition C.10 has a unique fixed point, call it  $\mathbf{u}_*$  at  $(m_{ab})_{a,b}$  as above, and  $R_{ab}^\perp = 0$  for all  $a, b$ . As observed in the proof of Lemma C.13, the dynamical system never leaves a ball of some radius  $L(\beta)$  and is also a gradient system for an energy function  $\mathcal{H}$ . Combining these facts with continuity of the drift functions, it means that for every  $\epsilon > 0$ , there is a  $T_0(\epsilon)$  such that the solution to the SGD gets within distance  $T_0$  of that unique fixed point and stays there for all  $t \geq T_0$ .

Altogether, this leaves us with the following stronger form of Proposition C.1.

**Proposition C.14.** *When the means  $(\mu_1, \dots, \mu_k)$  are orthonormal, beyond Proposition C.1, we further have that  $\mathbf{x}$  is at distance  $O(\epsilon + \lambda^{-1})$  of a point  $\mathbf{x}_*$  such that for each  $c$ ,  $\mathbf{x}_*^c$  has positive (bounded away from zero uniformly in  $\epsilon, \lambda$ ) inner product with  $\mu_c$  and negative inner product with  $(\mu_a)_{a \neq c}$ .*

If  $(p_c)_{c \in [k]}$  are assumed to all be equal, then furthermore at  $\mathbf{x}_*$ ,

$$\bar{\pi}_b(a) = \frac{1}{k-1}(1 - \bar{\pi}_c(c)) \quad \text{for all } a, b, c \text{ such that } a \neq b.$$

#### C.4 2-LAYER NETWORKS

We now turn to the analysis of the SGD in the case of multilayer networks for the XOR GMM problem (4.1). In this problem, the input dimension is  $d = n$ , the parameter dimension is  $p = Kd + K$ , and the step-size is again taken to be  $\delta = O(1/d)$ . Consider the following family of  $4K + \binom{K}{2}$  summary statistics of  $\mathbf{x}$ : for  $1 \leq i \leq j \leq K$  and  $\vartheta \in \{\mu, \nu\}$ ,

$$v_i, \quad m_i^\vartheta = W_i \cdot \vartheta, \quad R_{ij}^\perp = W_i^\perp \cdot W_j^\perp, \quad (\text{C.20})$$

where  $W_i^\perp = W_i - \sum_{\vartheta \in \{\mu, \nu\}} m_i^\vartheta \vartheta$ . Use  $\mathbf{u} = (\mathbf{v}, \mathbf{m}^\mu, \mathbf{m}^\nu, \mathbf{R}^\perp)$  for these families.

In Ben Arous et al. (2022) it was shown that this family of summary statistics is  $\delta$ -localizable and asymptotically closable with respect to the loss for the XOR GMM of (4.2), and the following convergence to ODE's was established. For a point  $\mathbf{x} = (v, W) \in \mathbb{R}^{K+Kd}$ , define the quantity

$$\mathbf{A}_i = \mathbb{E}[Y \mathbf{1}_{W_i \cdot Y \geq 0} (-y + \sigma(v \cdot g(WY)))] , \quad (\text{C.21})$$

(where we recall that  $g$  is the ReLU function and  $\sigma$  is the sigmoid function) and let

$$\mathbf{A}_i^\vartheta = \vartheta \cdot \mathbf{A}_i, \quad \mathbf{A}_{ij}^\perp = W_j^\perp \cdot \mathbf{A}_i.$$

Furthermore, let

$$\mathbf{B}_{ij} = \mathbb{E}[\mathbf{1}_{W_i \cdot Y \geq 0} \mathbf{1}_{W_j \cdot Y \geq 0} (-y + \sigma(v \cdot g(WY)))^2] . \quad (\text{C.22})$$

It can be observed that these functions are expressible as functions of  $\mathbf{u}$  alone.

Proposition 5.1 of Ben Arous et al. (2022) proved the following effective ballistic dynamics.

**Proposition C.15.** *Let  $\mathbf{u}_n$  be as in (C.20) and fix any  $\lambda > 0$  and  $\delta = c_\delta/d$ . Then  $\mathbf{u}_n(t)$  converges weakly to the solution of the ODE system  $\dot{\mathbf{u}}_t = -\mathbf{f}(\mathbf{u}_t) + \mathbf{g}(\mathbf{u}_t)$ , initialized from  $\lim_n (\mathbf{u}_n)_* \mu_n$  with*

$$\begin{aligned} f_{v_i} &= m_i^\mu \mathbf{A}_i^\mu(\mathbf{u}) + m_i^\nu \mathbf{A}_i^\nu(\mathbf{u}) + \mathbf{A}_{ii}^\perp(\mathbf{u}) + \beta v_i, & f_{m_i^\mu} &= v_i \mathbf{A}_i^\mu + \beta m_i^\mu, \\ f_{R_{ij}^\perp} &= v_i \mathbf{A}_{ij}^\perp(\mathbf{u}) + v_j \mathbf{A}_{ji}^\perp(\mathbf{u}) + 2\beta R_{ij}^\perp, & f_{m_i^\nu} &= v_i \mathbf{A}_i^\nu + \beta m_i^\nu. \end{aligned}$$

and correctors  $g_{v_i} = g_{m_i^\mu} = g_{m_i^\nu} = 0$ , and  $g_{R_{ij}^\perp} = c_\delta \frac{v_i v_j}{\lambda} \mathbf{B}_{ij}$  for  $1 \leq i \leq j \leq K$ . The case where  $\delta = o(1/d)$  is read-off by formally setting  $c_\delta = 0$ .

##### C.4.1 LARGE $\lambda$ BEHAVIOR

We now wish to investigate the large  $\lambda$  behavior of the dynamical system in Proposition C.15. Our approach to doing this is to give a large  $\lambda$  approximation to the drifts in the above, and then use that to show that the trajectory is close to its  $\lambda = \infty$  version.

The first thing we do is give large  $\lambda$  approximations to the quantities  $\mathbf{A}_i$  and  $\mathbf{B}_{ij}$ . In what follows, we use  $F(x)$  to denote the cumulative distribution function of a standard Gaussian random variable.



**Lemma C.16.** Suppose  $Y \sim \vartheta + Z_\lambda$  for a fixed vector  $\vartheta \in \{\mu, -\mu, \nu, -\nu\}$  and fix a unit vector  $b \in \mathbb{R}^d$  and a vector  $v \in \mathbb{R}^K$ . Then

$$\begin{aligned} & \left| \mathbb{E}[(b \cdot Y) \mathbf{1}_{W_i \cdot Y > 0} \sigma(v \cdot g(WY))] - (b \cdot \vartheta) F\left(m_i^\vartheta \sqrt{\frac{\lambda}{R_{ii}}}\right) \sigma(v \cdot g(m^\vartheta)) \right| \\ & \leq \max_j (1 + v_j) (m_i^\vartheta + \sqrt{\frac{R_{ii}}{\lambda}}) e^{-(m_j^\vartheta)^2 \lambda / (16R_{jj})} + O(\lambda^{-1}). \end{aligned}$$

*Proof.* Let us start with a Taylor expansion of  $\sigma$ : For simplicity, let us use  $\sigma_Y$  and  $\sigma_\vartheta$  to denote  $\sigma(v \cdot g(WY))$  and  $\sigma(v \cdot g(W\vartheta))$  respectively and similarly for  $\sigma'$ . We then have

$$\sigma_Y - \sigma_\vartheta = \sigma'_\vartheta \cdot (v \cdot (g(W\vartheta) - g(WY))) + \sigma''(o) (v \cdot (g(W\vartheta) - g(WY)))^2, \quad (\text{C.23})$$

for some point  $o$  between  $v \cdot g(WY)$  and  $v \cdot g(W\vartheta)$ . Therefore, the left-hand side of the lemma is

$$\begin{aligned} & \mathbb{E} \left[ \mathbf{1}_{W_i \cdot Y > 0} (b \cdot \vartheta) \sigma_\vartheta \right] + \mathbb{E} \left[ \mathbf{1}_{W_i \cdot Y > 0} (b \cdot Z) \sigma_\vartheta \right] + \mathbb{E} \left[ \mathbf{1}_{W_i \cdot Y > 0} (b \cdot Y) \sigma'_\vartheta (v \cdot (g(W\vartheta) - g(WY))) \right] \\ & + \mathbb{E} \left[ \mathbf{1}_{W_i \cdot Y > 0} (b \cdot Y) \sigma''(o) (v \cdot (g(W\vartheta) - g(WY)))^2 \right]. \end{aligned} \quad (\text{C.24})$$

The first term of (C.24) is exactly the term we are comparing to in the left-hand side of the lemma statement. Since  $\mathbb{E}[(b \cdot Z)] = 0$ , the absolute value of the second term of (C.24) is bounded via Cauchy–Schwarz as follows:

$$|\mathbb{E}[\mathbf{1}_{W_i \cdot Y > 0} (b \cdot Z) \sigma_\vartheta]| = |\mathbb{E}[(\mathbf{1}_{W_i \cdot Y > 0} - \mathbf{1}_{m_i^\vartheta > 0}) (b \cdot Z) \sigma_\vartheta]| \leq \frac{\sigma_\vartheta}{\sqrt{\lambda}} F(-|m_i^\vartheta| \sqrt{\lambda/R_{ii}})^{1/2}$$

For the next two terms we can start by rewriting

$$g(W_j \cdot \vartheta) - g(W_j \cdot Y) = (m_j^\vartheta + W_j \cdot Z) \mathbf{1}_{W_j \cdot Y < 0, m_j^\vartheta > 0} - (W_j \cdot Y) \mathbf{1}_{W_j \cdot Y > 0, m_j^\vartheta < 0} - (W_j \cdot Z) \mathbf{1}_{m_j^\vartheta > 0}. \quad (\text{C.25})$$

Using this, the third expectation in (C.24) is

$$\sigma'_\vartheta \mathbb{E} \left[ (b \cdot Y) \mathbf{1}_{W_i \cdot Y > 0} \sum_j v_j \left( (m_j^\vartheta + W_j \cdot Z) \mathbf{1}_{W_j \cdot Y < 0, m_j^\vartheta > 0} - (W_j \cdot Y) \mathbf{1}_{W_j \cdot Y > 0, m_j^\vartheta < 0} - (W_j \cdot Z) \mathbf{1}_{m_j^\vartheta > 0} \right) \right].$$

The first and second terms in the parentheticals are such that the Cauchy–Schwarz inequality can be applied to bound their total contributions by

$$K \sigma'_\vartheta \|v\|_\infty \max_j \left( (m_j^\vartheta)^2 + \frac{R_{jj}}{\lambda} \right)^{1/2} F(-|m_j^\vartheta| \sqrt{\lambda/R_{jj}})^{1/2}.$$

The last term will contribute (up to a net sign)

$$\sigma'_\vartheta v_j \mathbb{E}[(b \cdot \vartheta) (W_j \cdot Z) (\mathbf{1}_{W_i \cdot Y > 0} - \mathbf{1}_{m_i^\vartheta > 0}) \mathbf{1}_{m_j^\vartheta > 0}] + \sigma'_\vartheta v_j \mathbb{E}[(b \cdot Z) (W_j \cdot Z) \mathbf{1}_{W_i \cdot Y > 0} \mathbf{1}_{m_j^\vartheta > 0}].$$

The absolute value of the first term is bounded similarly to an earlier one by Cauchy–Schwarz. The absolute value of the second term is bounded by dropping the indicator functions and applying Cauchy–Schwarz to see that it is  $O(1/\lambda)$  so long as  $\|b\|, \|W_j\| = O(1)$ .

Finally, for the fourth term of (C.24), the square allows us to put absolute values on every term, and immediately apply the Cauchy–Schwarz inequality, to bound its absolute value by

$$\|\sigma''\|_\infty \|v\|_\infty \mathbb{E}[(b \cdot Y)^2]^{1/2} \max_j \mathbb{E}[(g(W_j \cdot \vartheta) - g(W_j \cdot Y))^4]^{1/2}.$$

The fourth moment in the above expression is bounded, up to constant, by the fourth moments of each of the individual terms in (C.25). The first two of those will be at most some constant times  $F(-|m_j^\vartheta| \sqrt{\lambda/R_{jj}})^{1/4}$ . For the last of them, we use  $\mathbb{E}[(W_j \cdot Z)^4]^{1/2} \leq O(1/\lambda)$ .

Combining all of the above bounds, and naively bounding the cdf of the Gaussian via

$$F(-|m_j^\vartheta| \sqrt{\lambda/R_{jj}})^{1/4} \leq e^{-(m_j^\vartheta)^2 \lambda / (16R_{jj})}$$

we arrive at the claimed bound.  $\square$

As a consequence of Lemma C.16, we can deduce that the quantities appearing above satisfy the following large  $\lambda$  behavior.

**Corollary C.17.** *For every  $\mathbf{x}$  such that  $m_i^\mu, m_i^\nu \geq \frac{\log \lambda}{\sqrt{\lambda}}$  for all  $i$ ,*

$$\begin{aligned} m_i^\mu \mathbf{A}_i^\mu &= -\frac{1}{4}g(m_i^\mu)\sigma(-v \cdot g(m^\mu)) - \frac{1}{4}g(-m_i^\mu)\sigma(-v \cdot g(-m^\mu)) + O(\lambda^{-1}) \\ m_i^\nu \mathbf{A}_i^\nu &= \frac{1}{4}g(m_i^\nu)\sigma(v \cdot g(m^\nu)) + \frac{1}{4}g(-m_i^\nu)\sigma(v \cdot g(-m^\nu)) + O(\lambda^{-1}) \\ \mathbf{A}_{ij}^\perp &= O(\lambda^{-1}). \end{aligned}$$

Without the assumption on  $m_i^\mu, m_i^\nu$ , the same holds with  $O(\lambda^{-1})$  replaced by  $O(\lambda^{-1/2})$ , as long as the indicators from  $g(m_i^\vartheta) = m_i^\vartheta \mathbf{1}_{m_i^\vartheta > 0}$  are replaced by the ‘‘soft indicators’’  $F(m_i^\vartheta \sqrt{\lambda/R_{ii}})$ .

*Proof.* We begin with the estimate on  $m_i^\mu \mathbf{A}_i^\mu$ . This quantity can be split into four terms corresponding to whether the mean chosen for  $Y$  is  $\mu, -\mu, \nu, -\nu$ . Namely, if we let  $Y_a \stackrel{d}{=} a + Z_\lambda$ ,

$$\begin{aligned} \mathbf{A}_i^\mu &= \frac{1}{4} \left( \mathbb{E}[(\mu \cdot Y_\mu) \mathbf{1}_{W_i \cdot Y_\mu \geq 0} \sigma(-v \cdot g(WY_\mu))] + \mathbb{E}[(\mu \cdot Y_{-\mu}) \mathbf{1}_{W_i \cdot Y_{-\mu} \geq 0} \sigma(-v \cdot g(WY_{-\mu}))] \right. \\ &\quad \left. + \mathbb{E}[(\mu \cdot Y_\nu) \mathbf{1}_{W_i \cdot Y_\nu \geq 0} \sigma(v \cdot g(WY_\nu))] + \mathbb{E}[(\mu \cdot Y_{-\nu}) \mathbf{1}_{W_i \cdot Y_{-\nu} \geq 0} \sigma(v \cdot g(WY_{-\nu}))] \right) \end{aligned}$$

Now notice that each of the four quantities are of the form of Lemma C.16, with  $b = \mu$ , and  $\vartheta = \mu, -\mu, \nu, -\nu$  respectively (the change of the sigmoid possibly having a negative sign on its argument is realized by switching the sign of  $v$  since that is the only place it appears).

It is easily seen that so long as  $|m_i^\mu| \geq (\log \lambda)/\sqrt{\lambda}$ , then

$$F(m_i^\mu \sqrt{\lambda/R_{ii}}) = \mathbf{1}\{m_i^\mu > 0\} + O(1/\lambda).$$

By a similar bound on the error term on the right of Lemma C.16,

$$\mathbb{E}[(\mu \cdot Y_\mu) \mathbf{1}_{W_i \cdot Y_\mu \geq 0} \sigma(-v \cdot g(WY_\mu))] = \mathbf{1}_{m_i^\mu > 0} \sigma(-v \cdot g(m^\mu)) + O(1/\lambda).$$

and the cases where it is  $Y_\nu$  contribute  $0 + O(\lambda^{-1})$  since  $\mu \cdot \nu = 0$ . Together with the analogous bounds for  $m_i^\nu \mathbf{A}_i^\nu$  and  $\mathbf{A}_{ij}^\perp$ , this gives the first part of the corollary.

If we drop the assumption on  $m_i^\mu$ , we find from the general inequality  $xe^{-x^2L} \leq \frac{C}{\sqrt{L}}$  for some uniform  $C$ , that the errors in the right-hand side of Lemma C.16 become at most  $O(1/\sqrt{\lambda})$  as claimed. Leaving the soft indicator in place, this gives the second part of the corollary.  $\square$

We deduce from the above approximation and a trivial bound of  $O(1/\lambda)$  on  $\mathbf{B}_{ij}$ , a  $\lambda = \infty$  limiting dynamics. However, one has to be slightly careful about the  $\lambda \rightarrow \infty$  limit near the hyperplanes where  $m_i^\mu = 0$  or  $m_i^\nu = 0$ ; Therefore, in what follows we envision a different limit associated to each orthant of the parameter space (associated to the signs of  $m_i^\mu, m_i^\nu$ ). This is reasonable because in each orthant, the  $\lambda = \infty$  dynamical system initialized from that orthant stays in that orthant. The following proposition captures that limiting dynamics (all orthants being written into the below simultaneously, with ambiguities at  $m_i^\mu = 0$  being therefore omitted).

**Proposition C.18.** *In the  $\lambda \rightarrow \infty$  limit, the ODE from Proposition C.15 converges to*

$$\begin{aligned} \dot{v}_i &= \frac{1}{4} \left( g(m_i^\mu)\sigma(-v \cdot g(m^\mu)) + g(-m_i^\mu)\sigma(-v \cdot g(-m^\mu)) \right) \\ &\quad - \frac{1}{4} \left( g(m_i^\nu)\sigma(v \cdot g(m^\nu)) + g(-m_i^\nu)\sigma(v \cdot g(-m^\nu)) \right) - \beta v_i, \\ \dot{m}_i^\mu &= \frac{v_i}{4} \left( \mathbf{1}_{m_i^\mu > 0} \sigma(-v \cdot g(m^\mu)) - \mathbf{1}_{m_i^\mu < 0} \sigma(-v \cdot g(-m^\mu)) \right) - \beta m_i^\mu, \\ \dot{m}_i^\nu &= -\frac{v_i}{4} \left( \mathbf{1}_{m_i^\nu > 0} \sigma(-v \cdot g(m^\nu)) - \mathbf{1}_{m_i^\nu < 0} \sigma(-v \cdot g(-m^\nu)) \right) - \beta m_i^\nu, \end{aligned}$$

and  $\dot{R}_{ij}^\perp = -2\beta R_{ij}^\perp$  for  $1 \leq i \leq j \leq K$ .

Moreover, for  $\lambda$  large, the difference in the drifts above and those of Proposition C.15 are of order  $O(1/\lambda)$  if  $(|m_i^\mu|, |m_i^\nu|)$  are all at least  $\sqrt{(\log \lambda)/\lambda}$ , and  $O(1/\sqrt{\lambda})$  everywhere.

### C.4.2 CONFINING THE SGD TO A COMPACT SET

Similar to the  $k$ -GMM case, our first aim is to confine the SGD to a compact set so long as  $\beta > 0$ .

**Lemma C.19.** *For every  $\beta > 0$ , there exists  $L(\beta, v(\mathbf{x}_0))$  such that for all  $\lambda$  large, the dynamics of Proposition C.15 stays inside the  $\ell^2$ -ball of radius  $L$  for all time.*

*Proof.* We start by considering the evolution of the  $\ell^2$ -norm of all the parameters,  $(v, m^\mu, m^\nu)$ . If we use the shorthand

$$g^\mu = g(m^\mu), \quad g^{-\mu} = g(-m^\mu), \quad \sigma^\mu = \sigma(-v \cdot g^\mu), \quad \sigma^{-\mu} = \sigma(-v \cdot g^{-\mu})$$

and similar quantities with  $\nu$  instead of  $\mu$ , a short calculation shows that we get in the  $\lambda = \infty$  dynamical system of Proposition C.18

$$\|\dot{\mathbf{v}}\|^2 = \frac{1}{2}(v \cdot g^\mu \sigma^\mu + v \cdot g^{-\mu} \sigma^{-\mu}) - \frac{1}{2}(v \cdot g^\nu \sigma^\nu + v \cdot g^{-\nu} \sigma^{-\nu}) - 2\beta \|\mathbf{v}\|^2.$$

Notice that  $\sigma^\mu$  goes to 0 as  $v \cdot g^\mu \rightarrow \infty$ , and moreover,  $x\sigma(-x) \rightarrow 0$  as  $x \rightarrow \infty$ . Therefore, there is a uniform bound of 1 (in fact,  $W(1/e)$  where  $W$  is the product log function) on  $v \cdot g^\mu \sigma^\mu$ . Similar uniform bounds apply to the other three terms, so that in total,

$$\|\dot{\mathbf{v}}\|^2 \leq 2 - 2\beta \|\mathbf{v}\|^2,$$

which implies in particular that its drift is negative when  $\|\mathbf{v}\|^2$  is at least  $L(\beta)$ . Similar arguments go through mutatis mutandis for  $(m^\mu)^2$  and  $(m^\nu)^2$ . These then imply that the drift when  $\lambda$  is finite are similarly negative when  $\|\mathbf{v}\|^2 = L(\beta)$  as long as  $\lambda$  is sufficiently large per the approximation from Proposition C.18. Altogether, this implies that for the dynamical system of Proposition C.15, all of  $\|\mathbf{v}\|$ ,  $\|\mathbf{m}^\mu\|$  and  $\|\mathbf{m}^\nu\|$ , stay bounded by some  $L(\beta)$  for all time.

Using that, if we consider the expression for  $\dot{R}_{ii}^\perp$  and use the boundedness of  $\mathbf{A}_i, \mathbf{B}$ , the drift in Proposition C.15, gives

$$\|\mathbf{R}^\perp\|^2 \leq -\beta \|\mathbf{R}^\perp\|^2 + 4L \|\mathbf{R}^\perp\|,$$

The right-hand side is at most  $C_L - \frac{1}{2}\beta \|\mathbf{R}^\perp\|^2$  for some  $C_L$ , whence by Gronwall,  $\|\mathbf{R}^\perp\|^2$  also stays bounded by some constant  $L'(\beta)$  for all time under the dynamics of Proposition C.15. Altogether these prove the lemma.  $\square$

**Corollary C.20.** *The trajectories of the ODEs of Proposition C.15 and Proposition C.18 are within distance  $O(e^{Ct}/\sqrt{\lambda})$  of one another.*

*Proof.* Let  $\mathbf{u}$  and  $\tilde{\mathbf{u}}$  be the two solutions to the  $\lambda$  finite and  $\lambda = \infty$  ballistic limits respectively. Then

$$\|\dot{\mathbf{u}} - \dot{\tilde{\mathbf{u}}}\| \lesssim_K (1 + L + 2\beta) \|\mathbf{u} - \tilde{\mathbf{u}}\| + O(\lambda^{-1/2}),$$

where we used that  $L$  bounds the norm of  $\mathbf{u}$  and  $\tilde{\mathbf{u}}$  for all times by Lemma C.19, and that the Lipschitz constant of the sigmoid is 1. By Gronwall's inequality, this implies that  $\|\mathbf{u} - \tilde{\mathbf{u}}\| \leq O(e^{Ct}/\sqrt{\lambda})$  as claimed for some  $C$  depending only on  $\beta$ .  $\square$

### C.4.3 LIVING IN THE PRINCIPAL DIRECTIONS NEAR FIXED POINTS

The last thing to conclude is that the fixed points of the  $\lambda = \infty$  dynamical system are indeed living in the principal directions as desired by Theorem 4.2. By the above arguments, for all  $t \geq T_0(\epsilon)$  the dynamics is within distance  $\epsilon$  of one of the fixed points of Proposition C.18. Let us recall the exact locations of those fixed points from Ben Arous et al. (2022).

If  $0 < \beta < 1/8$ , then let  $(I_0, I_\mu^+, I_\mu^-, I_\nu^+, I_\nu^-)$  be any disjoint (possibly empty) subsets whose union is  $\{1, \dots, K\}$ . Corresponding to that tuple  $(I_0, I_\mu^+, I_\mu^-, I_\nu^+, I_\nu^-)$ , is a set of fixed points that have  $R_{ij}^\perp = 0$  for all  $i, j$ , and have

1.  $m_i^\mu = m_i^\nu = v_i = 0$  for  $i \in I_0$ ,
2.  $m_i^\mu = v_i > 0$  such that  $\sum_{i \in I_\mu^+} v_i^2 = \text{logit}(-4\beta)$  and  $m_i^\nu = 0$  for all  $i \in I_\mu^+$ ,

3.  $-m_i^\mu = v_i > 0$  such that  $\sum_{i \in I_\mu^-} v_i^2 = \text{logit}(-4\beta)$  and  $m_i^\nu = 0$  for all  $i \in I_\mu^-$ ,
4.  $m_i^\nu = v_i < 0$  such that  $\sum_{i \in I_\nu^+} v_i^2 = \text{logit}(-4\beta)$  and  $m_i^\mu = 0$  for all  $i \in I_\nu^+$ ,
5.  $-m_i^\nu = v_i < 0$  such that  $\sum_{i \in I_\nu^-} v_i^2 = \text{logit}(-4\beta)$  and  $m_i^\mu = 0$  for all  $i \in I_\nu^-$ .

The following observation is easy to see from the fixed point characterization described above.

*Observation C.21.* Suppose that  $x_*$  is a fixed point amongst the above. Then

- $W(x_*) \in \text{Span}(\mu, \nu)$ ,
- $v(x_*) \in \text{Span}(g^\mu(x_*), g^{-\mu}(x_*), g^\nu(x_*), g^{-\nu}(x_*))$ ,

(with no error).

*Proof.* The first claim that  $W(x_*) \in \text{Span}(\mu, \nu)$  follows from the fact that  $R_{ii}^\perp = 0$  for all  $i$ . Furthermore, if  $I_\nu^+, I_\nu^-$  are empty, then it lives in  $\text{Span}(\mu)$ , and similarly if  $I_\mu^+, I_\mu^-$  are empty, then it lives in  $\text{Span}(\nu)$ .

For the next claim, observe that  $g^\mu(x_*)$  is the vector that is  $m^\mu$  in coordinates belonging to  $I_\mu^+$ , and 0 in all other coordinates.  $g^{-\mu}(x_*)$  is the vector that is  $-m^\mu$  in coordinates belonging to  $I_\mu^-$ , and zero in others.  $g^\nu(x_*)$  is  $m_i^\nu$  for coordinates in  $I_\nu^-$ , zero else, and  $g^{-\nu}(x_*)$  is  $-m_i^\nu$  on  $I_\nu^+$ .

Since the  $I$ -sets are a partition of  $\{1, \dots, K\}$  it is evident that we can express

$$v(x_*) = g^\mu(x_*) + g^{-\mu}(x_*) - g^\nu(x_*) - g^{-\nu}(x_*),$$

implying that indeed  $v(x_*)$  lives in  $\text{Span}(g^\mu(x_*), g^{-\mu}(x_*), g^\nu(x_*), g^{-\nu}(x_*))$   $\square$

It was further argued in (Ben Arous et al., 2022, Section 9.4) that there is a transition at  $\beta = 1/8$  in the regularization, and as long as  $\beta < 1/8$ , with probability 1 with respect to the random initialization, the SGD converges to a fixed point having  $I_0 = \emptyset$ .

**Lemma C.22.** Suppose  $\beta < 1/8$  and consider the initialization

$$v_i(0) \sim \mathcal{N}(0, 1), \quad \text{and} \quad m_i^\mu(0), m_i^\nu(0), R_{ij}^\perp(0) \sim \delta_0 \quad \text{and} \quad R_{ii}^\perp \sim \delta_1, \quad (\text{C.26})$$

where since the dynamical system of Proposition C.10 is defined per orthant, the  $\delta_0$ 's are understood as  $1/2$ - $1/2$  mixtures of  $\delta_{0+}$  and  $\delta_{0-}$ . With probability 1 over this initialization, the dynamical system of Proposition C.10 converges to a fixed point as characterized above, above having  $I_0 = \emptyset$ .

Putting the above together, we can conclude our proof of Proposition C.2.

**Proof of Proposition C.2.** Per Proposition C.15, the limit of the summary statistics of the SGD along training is the solution of that dynamical system initialized from  $m_i^\mu, m_i^\nu, R_{ij}^\perp \sim \delta_0$  for  $i \neq j$  (with equal probability of  $m_i^\mu, m_i^\nu$  being  $\delta_{0+}$  and  $\delta_{0-}$  if we need to distinguish which orthant it is initialized in),  $R_{ii}^\perp \sim \delta_1$ , and  $v_i \sim \mathcal{N}(0, 1)$  i.i.d.

For the first part, notice that the norms  $\|W_i(\mathbf{x})\|^2 = R_{ii}^\perp + (m_i^\mu)^2 + (m_i^\nu)^2$  and  $\|v(\mathbf{x})\|^2 = \sum_{i=1}^K v_i^2$  are expressible in terms of the summary statistics. Therefore, their boundedness follows from Lemma C.19, pulled back to the summary statistics of the SGD per Lemma C.6.

For the second item, by Corollary C.20 and Lemma C.22 together with Observation C.21, for every  $\epsilon > 0$ , there is a  $T_0$  such that for every  $T_f$ , for all  $t \in [T_0, T_f]$ , the dynamical system of Proposition C.15 is within distance  $\epsilon + \lambda^{-1/2}$  of a point  $\mathbf{u}_*$  having  $|v_i(t)| > \eta$  and  $\max\{|m_i^\mu|, |m_i^\nu|\} > \eta$  in and having that its first layer lives in  $\text{Span}(\mu, \nu)$  and its second layer lives in  $\text{Span}((g^\vartheta(\mathbf{u}_*))_{\vartheta \in \{\pm\mu, \pm\nu\}})$ . This is pulled back to the summary statistics applied to the SGD trajectory per Lemma C.6 (with the observation that  $(g^\vartheta(\mathbf{x}))_\vartheta$  are functions of only the summary statistics).  $\square$

## D CONCENTRATION OF HESSIAN AND G-MATRICES

We recall the general forms of the empirical test Hessian matrix  $\nabla^2 \widehat{R}(\mathbf{x})$  and G-matrix  $\widehat{G}(\mathbf{x})$  from (2.2) respectively. Our aim in this section is to establish concentration of those empirical matrices about their population versions throughout the parameter space.

### D.1 HESSIAN AND G-MATRIX: 1-LAYER

We first prove the concentration of the empirical Hessian and G-matrix for the  $k$ -GMM problem about their population versions, which we analyzed in depth in Section A–B. This concentration will be uniform over the entire parameter space. Namely, our aim in this section is to show the following.

**Theorem D.1.** *Consider the  $k$ -GMM data model of (3.1) and sample complexity  $\widetilde{M}/d = \alpha$ . There are constants  $c = c(k)$ ,  $C = C(k)$  (independent of  $\lambda$ ) such that for all  $t > 0$ , the empirical Hessian matrix concentrates as*

$$\sup_{\mathbf{x} \in \mathbb{R}^{kd}} \mathbb{P}(\|\nabla^2(\widehat{R}(\mathbf{x}) - \mathbb{E}[\widehat{R}(\mathbf{x})])\|_{\text{op}} > t) \leq \exp(-[\alpha(t \wedge t^2) - C]d), \quad (\text{D.1})$$

and so does the empirical G-matrix

$$\sup_{\mathbf{x} \in \mathbb{R}^{kd}} \mathbb{P}(\|\widehat{G}(\mathbf{x}) - \mathbb{E}[\widehat{G}(\mathbf{x})]\|_{\text{op}} > t) \leq \exp(-[\alpha(t \wedge t^2) - C]d). \quad (\text{D.2})$$

*Proof.* In the following we fix  $\mathbf{x} \in \mathbb{R}^{kd}$  and simply write  $\widehat{R}(\mathbf{x}), \widehat{G}(\mathbf{x})$  as  $\widehat{R}, \widehat{G}$ . Let  $\widetilde{A} = (\widetilde{Y}^1 \dots \widetilde{Y}^{\widetilde{M}})$  denote the test data matrix and let

$$D_{bc}^{\text{H}} = \text{diag}(\pi_{\widetilde{Y}^\ell}(c)\delta_{bc} - \pi_{\widetilde{Y}^\ell}(c)\pi_{\widetilde{Y}^\ell}(b))_{1 \leq \ell \leq \widetilde{M}}, \quad (\text{D.3})$$

$$D_{bc}^{\text{G}} = \text{diag}((\widetilde{y}_c^\ell \widetilde{y}_b^\ell - \pi_{\widetilde{Y}^\ell}(b)\widetilde{y}_c^\ell - \widetilde{y}_b^\ell \pi_{\widetilde{Y}^\ell}(c) + \pi_{\widetilde{Y}^\ell}(c)\pi_{\widetilde{Y}^\ell}(b))_{1 \leq \ell \leq \widetilde{M}}, \quad (\text{D.4})$$

where  $\pi_Y(c)$  was defined in (A.4). Then  $D_{bc}^{\text{H}}, D_{bc}^{\text{G}}$  are  $\widetilde{M} \times \widetilde{M}$  diagonal matrices for each pair  $bc \in [k]^2$ . We denote the  $(k\widetilde{M}) \times (k\widetilde{M})$  matrices  $\mathbf{D}^{\text{H}} = (D_{bc}^{\text{H}})_{bc}$  and  $\mathbf{D}^{\text{G}} = (D_{bc}^{\text{G}})_{bc}$ , and the  $dk \times k\widetilde{M}$  matrix

$$\widetilde{A}^{\times k} = I_k \otimes \widetilde{A}.$$

With these notations, per (2.2), we can rewrite the Hessian and G-matrices as

$$\nabla^2 \widehat{R} = \frac{1}{\widetilde{M}} \widetilde{A}^{\times k} \mathbf{D}^{\text{H}} (\widetilde{A}^{\times k})^T, \quad (\text{D.5})$$

$$\widehat{G} = \frac{1}{\widetilde{M}} \widetilde{A}^{\times k} \mathbf{D}^{\text{G}} (\widetilde{A}^{\times k})^T. \quad (\text{D.6})$$

To prove that the operator norm of  $\nabla^2(\widehat{R} - \mathbb{E}[\widehat{R}])$  concentrates, we'll use a net argument over the unit ball in  $\mathbb{R}^{dk}$  to show that the following concentrates

$$\sup_{\mathbf{v} \in (\mathbb{R}^d)^k, \|\mathbf{v}\|=1} |\langle \mathbf{v}, \nabla^2(\widehat{R} - \mathbb{E}[\widehat{R}])\mathbf{v} \rangle|, \quad (\text{D.7})$$

where  $\mathbf{v} = (v_c)_{c \in [k]} \in (\mathbb{R}^d)^k$ .

By plugging (D.5) into (D.7), we want a concentration estimate for  $F(\mathbf{v}) = \langle \mathbf{v}, \nabla^2(\widehat{R} - \mathbb{E}[\widehat{R}])\mathbf{v} \rangle$ , which we can rewrite as follows.

$$\begin{aligned} F(\mathbf{v}) &= \frac{1}{\widetilde{M}} \langle \mathbf{v}, \widetilde{A}^{\times k} \mathbf{D}^{\text{H}} (\widetilde{A}^{\times k})^T \mathbf{v} \rangle - \langle \mathbf{v}, \mathbb{E}[\widetilde{A}^{\times k} \mathbf{D}^{\text{H}} (\widetilde{A}^{\times k})^T] \mathbf{v} \rangle \\ &= \sum_{a,b} \frac{1}{\widetilde{M}} \langle v_a, \widetilde{A} D_{ab}^{\text{H}} \widetilde{A}^T v_b \rangle - \langle v_a, \mathbb{E}[\widetilde{A} D_{ab}^{\text{H}} \widetilde{A}^T] v_b \rangle \\ &= \sum_{a,b} \frac{1}{\widetilde{M}} \sum_{\ell} d_{\ell}(a,b) \langle v_a, \widetilde{Y}^{\ell} \rangle \langle v_b, \widetilde{Y}^{\ell} \rangle - \mathbb{E}[d_{\ell}(a,b) \langle v_a, \widetilde{Y}^{\ell} \rangle \langle v_b, \widetilde{Y}^{\ell} \rangle], \end{aligned} \quad (\text{D.8})$$

where  $d_\ell(a, b) = [D_{ab}^H]_{\ell\ell}$ . Recalling (Vershynin, 2018, Section 4.4.1) that for an  $\epsilon$ -net  $\mathcal{N}_\epsilon$  of  $\{\mathbf{v} : \|\mathbf{v}\| = 1\}$ , and any real symmetric matrix  $H$  the following holds

$$\frac{1}{1 - 2\epsilon} \sup_{\mathbf{v} \in \mathcal{N}_\epsilon} |\langle \mathbf{v}, H\mathbf{v} \rangle| \geq \|H\|_{\text{op}} = \sup_{\|\mathbf{v}\|_2=1} |\langle \mathbf{v}, H\mathbf{v} \rangle| \geq \sup_{\mathbf{v} \in \mathcal{N}_\epsilon} |\langle \mathbf{v}, H\mathbf{v} \rangle|.$$

So by a union bound and the  $\epsilon$ -covering number of  $\{\mathbf{v} \in \mathbb{R}^{kd} : \|\mathbf{v}\| = 1\}$ , we have

$$\mathbb{P}\left(\sup_{\|\mathbf{v}\|_2=1} |F(\mathbf{v})| > t\right) \leq |\mathcal{N}_\epsilon| \sup_{\mathbf{v} \in \mathcal{N}_\epsilon} \mathbb{P}(|F(\mathbf{v})| > t/2) \leq (C/\epsilon)^{kd} \mathbb{P}(|F(\mathbf{v})| > t/2), \quad (\text{D.9})$$

so long as  $\epsilon < 1/4$ , say.

To control the last quantity  $\mathbb{P}(|F(\mathbf{v})| > t/2)$ , we notice that  $F(\mathbf{v})$  is a sum of  $O(1)$  many (in fact  $k^2$  many) terms (corresponding to each pair  $a, b$ ) so it suffices by a union bound to control the concentration of each summand. That is, it suffices to understand the concentration of

$$\frac{1}{\widetilde{M}} \sum_{\ell} \left( d_\ell(a, b) \langle v_a, \widetilde{Y}^\ell \rangle \langle v_b, \widetilde{Y}^\ell \rangle - \mathbb{E}[d_\ell(a, b) \langle v_a, \widetilde{Y}^\ell \rangle \langle v_b, \widetilde{Y}^\ell \rangle] \right). \quad (\text{D.10})$$

We recall that the test data

$$\widetilde{Y}^\ell = \sum_{a \in [k]} \widetilde{y}_a^\ell \mu_a + Z_\lambda^\ell := \mu^\ell + Z_\lambda^\ell, \quad (\text{D.11})$$

where  $Z_\lambda^\ell$  are i.i.d.  $\mathcal{N}(0, I_d/\lambda)$ , and  $\mu^\ell \sim \sum_{a \in [k]} p_a \delta_{\mu_a}$ . In this case, note that  $d_\ell(a, b) = \pi_{\widetilde{Y}^\ell}(b) \delta_{ab} - \pi_{\widetilde{Y}^\ell}(b) \pi_{\widetilde{Y}^\ell}(a)$  are uniformly bounded by 2, and that for  $1 \leq \ell \leq \widetilde{M}$

$$\langle v_a, \widetilde{Y}^\ell \rangle \stackrel{d}{=} \langle v_a, \mu^\ell \rangle + \langle v_a, Z_\lambda^\ell \rangle, \quad (\text{D.12})$$

are i.i.d. sub-Gaussian (with norm  $\mathcal{O}(1)$ ) for fixed  $v_a$ , since  $\|v_a\|_2 \leq 1$  and  $\|\mu_c\|_2 = 1$  for all  $c$ , and  $\lambda \geq 1$  say. Thus, for all  $a, b$ , the products  $(d_\ell(a, b) \langle v_a, \widetilde{Y}^\ell \rangle \langle v_b, \widetilde{Y}^\ell \rangle)_\ell$  are i.i.d. uniformly sub-exponential random variables. As such, (D.10) is a sum of i.i.d. centered uniformly sub-exponential random variables, and Bernstein's inequality yields that there exists a small constant  $c = c(k)$

$$\mathbb{P}(|F(\mathbf{v})| > t/2) \leq \exp(-c\widetilde{M}(t \wedge t^2)). \quad (\text{D.13})$$

It follows from plugging (D.13) into (D.9) and taking  $\epsilon = 1/8$ , that

$$\mathbb{P}\left(\sup_{\|\mathbf{v}\|_2=1} |F(\mathbf{v})| > t/2\right) \leq \exp(-c\widetilde{M}(t \wedge t^2) + Ckd),$$

which yields the concentration bound (D.1) for empirical Hessian matrix, by noticing that  $\widetilde{M}/d = \alpha$ .

For the proof of the concentration bound (D.2) for empirical G-matrix, thanks to (D.6), we can define  $F(\mathbf{v})$  (as in (D.8)) using  $\mathbf{D}^G$  instead of  $\mathbf{D}^H$ . Noticing that the entries  $\mathbf{D}_{ab}^H$  are bounded by 4, the concentration of this new  $F(\mathbf{v})$  follows from Bernstein's inequality by the same argument. Then the concentration estimates of this new  $F(\mathbf{v})$  together with a epsilon-net argument gives (D.2).  $\square$

## D.2 HESSIAN AND G-MATRIX: 2-LAYER GMM MODEL

For the reader's convenience, we recall the empirical Hessian and G-matrix from the beginning of Section B. Our main aim in this subsection is to prove the following analogue of Theorem D.1 for the XOR GMM with a 2-layer network. There is a small problem with differentiating the ReLU activation twice exactly at zero, so let us define the set  $\mathcal{W}_0^c = \bigcap_{i \leq K} \{W_i \neq 0\}$ .

**Theorem D.2.** *We consider the 2-layer XOR GMM model from (4.1)–(4.2) with sample complexity  $\widetilde{M}/d = \alpha$ . For any  $L$ , there are constants  $c = c(K, L)$ ,  $C = C(K, L)$  such that for all  $t > 0$ , the empirical Hessian matrix concentrates as*

$$\sup_{\substack{\|(v, W)\| \leq L \\ W \in \mathcal{W}_0^c}} P(\|\nabla^2(\widehat{R}(v, W) - \mathbb{E}[\widehat{R}(v, W)])\|_{\text{op}} > t) \leq \exp\{-[c\alpha(t \wedge t^2) - C]d\}, \quad (\text{D.14})$$

and the empirical G-matrix concentrates as

$$\sup_{\|(v, W)\| \leq L} P(\|\widehat{G}(v, W) - \mathbb{E}[\widehat{G}(v, W)]\|_{\text{op}} > t) \leq \exp\{-[c\alpha(t \wedge t^2) - C]d\}. \quad (\text{D.15})$$

*Proof.* We begin by recalling some expressions. Letting  $\hat{y}^\ell = \sigma(v \cdot g(W\tilde{Y}^\ell))$  as in (B.1), by (B.4), we have

$$\begin{aligned}\nabla_{vv}^2 \hat{R} &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} \hat{y}^\ell (1 - \hat{y}^\ell) g(W\tilde{Y}^\ell)^{\otimes 2}, \\ \nabla_{W_i W_j}^2 \hat{R} &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (\delta_{ij} v_i g''(W_i \cdot \tilde{Y}^\ell) (\hat{y}^\ell - \tilde{y}^\ell) + v_i v_j \hat{y}^\ell (1 - \hat{y}^\ell) g'(W_i \cdot \tilde{Y}^\ell) g'(W_j \cdot \tilde{Y}^\ell) (\tilde{Y}^\ell)^{\otimes 2}), \\ \nabla_{v W_j}^2 \hat{R} &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} ((\tilde{y}^\ell - \hat{y}^\ell) g'(W_j \cdot \tilde{Y}^\ell) \mathbf{e}_j + v_j \hat{y}^\ell (1 - \hat{y}^\ell) g'(W_j \cdot \tilde{Y}^\ell) g(W \cdot \tilde{Y}^\ell)) \otimes \tilde{Y}^\ell,\end{aligned}\tag{D.16}$$

and by (B.2)–(B.3), we have

$$\begin{aligned}\hat{G}_{vv} &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (\hat{y}^\ell - \tilde{y}^\ell)^2 g(W\tilde{Y}^\ell) \otimes g(W\tilde{Y}^\ell) \\ \hat{G}_{W_i W_j} &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (\hat{y}^\ell - \tilde{y}^\ell)^2 v_i v_j g'(W_i \cdot \tilde{Y}^\ell) g'(W_j \cdot \tilde{Y}^\ell) \tilde{Y}^\ell \otimes \tilde{Y}^\ell, \\ \hat{G}_{v W_j} &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (\hat{y}^\ell - \tilde{y}^\ell)^2 v_j g'(W_j \cdot \tilde{Y}^\ell) g(W \cdot \tilde{Y}^\ell) \otimes \tilde{Y}^\ell.\end{aligned}\tag{D.17}$$

Recalling the data distribution from (4.1), so long as  $W \in \mathcal{W}_0^c$ , almost surely, all the coordinates of  $WY$  are nonzero: as long as  $\lambda < \infty$ ,

$$\mathbb{P}(\exists i, W_i \tilde{Y}^\ell = 0) = 0, \quad \text{for all } W \in \mathcal{W}_0^c.\tag{D.18}$$

For any  $W \in \mathcal{W}_0^c$ , we thus can ignore the second derivative term  $g''(W_i \cdot \tilde{Y}^\ell)$  in  $\nabla_{W_i W_j}^2 \hat{R}$ , so a.s.

$$\nabla_{W_i W_j}^2 \hat{R} = \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (v_i g'(W_i \cdot \tilde{Y}^\ell)) (v_j g'(W_j \cdot \tilde{Y}^\ell)) \hat{y}^\ell (1 - \hat{y}^\ell) (\tilde{Y}^\ell)^{\otimes 2}.\tag{D.19}$$

With these calculations in hand, the proof is similar to that of Theorem D.1, we will only emphasize the main differences. Let  $B$  be the ball of radius  $L$  in parameter space. Fix  $(v, W) \in B$  and simply write  $\hat{R}(v, W), \hat{G}(v, W)$  as  $\hat{R}, \hat{G}$ . Our goal is to prove concentration of the operator norm

$$\sup_{\substack{(\mathbf{a}, \mathbf{u}) \in \mathbb{R}^K \times (\mathbb{R}^d)^K \\ \|\mathbf{a}, \mathbf{u}\| = 1}} \langle (\mathbf{a}, \mathbf{u}), (\nabla^2(\hat{R} - \mathbb{E}[\hat{R}])(\mathbf{a}, \mathbf{u})) \rangle,$$

where  $\mathbf{a} = (a_1, a_2, \dots, a_K)$  and  $\mathbf{u} = (u_1, u_2, \dots, u_K) \in (\mathbb{R}^d)^K$ . As in the proof of Theorem D.1, specifically (D.9), by an epsilon-net argument, we only need prove a concentration estimate for the following quantity

$$F(\mathbf{a}, \mathbf{u}) = \langle (\mathbf{a}, \mathbf{u}), \nabla^2(\hat{R} - \mathbb{E}[\hat{R}])(\mathbf{a}, \mathbf{u}) \rangle,\tag{D.20}$$

individually per  $(\mathbf{a}, \mathbf{u}) : \|\mathbf{a}, \mathbf{u}\| = 1$ . The inner product on the right-hand side of (D.20) splits into  $(K+1)^2$  terms corresponding the terms in (D.16) and (D.19):

$$\begin{aligned}\langle \mathbf{a}, \nabla_{vv}^2 \hat{R} \mathbf{a} \rangle &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} \hat{y}^\ell (1 - \hat{y}^\ell) \langle \mathbf{a}, g(W\tilde{Y}^\ell) \rangle^2, \\ \langle u_i, \nabla_{W_i W_j}^2 \hat{R} u_j \rangle &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (v_i g'(W_i \cdot \tilde{Y}^\ell)) (v_j g'(W_j \cdot \tilde{Y}^\ell)) \hat{y}^\ell (1 - \hat{y}^\ell) \langle u_i, \tilde{Y}^\ell \rangle \langle u_j, \tilde{Y}^\ell \rangle, \\ \langle \mathbf{a}, \nabla_{v W_j}^2 \hat{R} u_j \rangle &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (a_j g'(W_j \cdot \tilde{Y}^\ell) (\tilde{y}^\ell - \hat{y}^\ell) + (v_j \hat{y}^\ell (1 - \hat{y}^\ell) g'(W_j \cdot \tilde{Y}^\ell) g(W \cdot \tilde{Y}^\ell))) \langle \tilde{Y}^\ell, u_j \rangle.\end{aligned}\tag{D.21}$$

The concentration bound for  $F(\mathbf{a}, \mathbf{u})$  follows from concentration bounds for each of the above terms about their respective expected values. Each such term minus its expected value will be a sum of  $\widetilde{M}$  i.i.d. centered random variables. It remains to show that the summands are uniformly (in  $\lambda$  and  $B_L$ ) sub-exponential. Then the concentration of  $F(\mathbf{a}, \mathbf{u})$  follows from Bernstein's inequality as in (D.13).

To that end, we notice that  $(\widehat{y}^\ell)(1 - \widehat{y}^\ell)$ ,  $\widetilde{y}^\ell$  are bounded by 1, and  $v_i, v_j$  and  $\|g'\|_\infty$  are all bounded by a  $C(L)$ . By the same argument as in the proof of Theorem D.1, we have that  $\langle u_i, \widetilde{Y}^\ell \rangle$  and  $\langle u_j, \widetilde{Y}^\ell \rangle$  are uniformly sub-Gaussian.

Next we show that  $\langle g(W\widetilde{Y}^\ell), \mathbf{a} \rangle$  is also sub-Gaussian. In law, conditionally on the mean choice among  $\pm\mu, \pm\nu$ , we have  $W\widetilde{Y}^\ell \stackrel{d}{=} W(\pm\mu + Z_\lambda)$  or  $W\widetilde{Y}^\ell \stackrel{d}{=} W(\pm\nu + Z_\lambda)$ . Since  $(\mathbf{a}, \mathbf{u}) \in B$  and  $\|\mu\| = \|\nu\| = 1$ , we have by Cauchy-Schwarz that  $\|W\mu\|, \|W\nu\| \leq L$ . Then we can write

$$\langle g(W\widetilde{Y}^\ell), \mathbf{a} \rangle = \langle g(WZ_\lambda), \mathbf{a} \rangle + \langle g(W\widetilde{Y}^\ell) - g(WZ_\lambda), \mathbf{a} \rangle.$$

By the uniform Lipschitz continuity of  $g$ , we have

$$|\langle g(W\widetilde{Y}^\ell) - g(WZ_\lambda), \mathbf{a} \rangle| \leq \|\mathbf{a}\|(\|W\nu\| + \|W\mu\|) \leq 2.$$

We also notice that if  $Z_1 = \sqrt{\lambda}Z_\lambda$ , then

$$\nabla_{Z_1} \langle g(WZ_\lambda), \mathbf{a} \rangle = \frac{1}{\sqrt{\lambda}} \sum_{i=1}^K a_i g'(W_i Z_\lambda) W_i \leq \frac{1}{\sqrt{\lambda}} \|g'\|_\infty \|\mathbf{a}\|_2 \|W\|_{2 \rightarrow 2} \lesssim \frac{L}{\sqrt{\lambda}},$$

implying that it is a uniformly  $L$ -Lipschitz function of standard Gaussians, from which it follows that  $g(WZ_\lambda)$  is uniformly sub-Gaussian (see (Vershynin, 2018, Section 5.2.1)). For the expectation we have

$$\begin{aligned} |\mathbb{E}[\langle g(WZ/\sqrt{\lambda}), \mathbf{a} \rangle]| &\leq |\mathbb{E}[\langle g(\mathbf{0}), \mathbf{a} \rangle]| + |\mathbb{E}[\langle g(WZ/\sqrt{\lambda}) - g(\mathbf{0}), \mathbf{a} \rangle]| \\ &\leq |g(0)|\sqrt{K}\|\mathbf{a}\|_2 + \|g'\|_\infty \mathbb{E}[\|WZ/\sqrt{\lambda}\|_2] \|\mathbf{a}\|_2 \\ &\leq |g(0)|\sqrt{K}\|\mathbf{a}\|_2 + \frac{1}{\lambda} \|g'\|_\infty \|\mathbf{a}\|_2 \sqrt{\text{Tr}[WW^\top]} \leq C(K, L). \end{aligned}$$

We conclude that  $\langle g(W\widetilde{Y}^\ell), \mathbf{a} \rangle$  is sub-Gaussian with norm bounded by  $\mathcal{O}(1)$  (depending only on  $K, L$ .) As a consequence, each summand in (D.21) is sub-exponential, with norm uniformly bounded by  $\mathcal{O}(1)$  (depending only on  $K, L$ .) Bernstein's inequality gives that there exists a small constant  $c = c(K, L)$

$$\mathbb{P}(|F(\mathbf{a}, \mathbf{u})| > t/2) \leq \exp(-c\widetilde{M}(t \wedge t^2)). \quad (\text{D.22})$$

It follows from (D.22) and the epsilon-net argument over the unit ball of  $(\mathbf{a}, \mathbf{u}) \in \mathbb{R}^K \times \mathbb{R}^{Kd}$  as in (D.9), that

$$\mathbb{P}\left(\sup_{\|(\mathbf{a}, \mathbf{u})\|=1} |F(\mathbf{a}, \mathbf{u})| > t\right) \leq \exp(-c\widetilde{M}(t \wedge t^2) - CK^2d),$$

which yields the concentration bound (D.14) for empirical Hessian matrix, by noticing that  $\widetilde{M}/d = \alpha$ .

For the proof of the concentration bound (D.15) for the empirical G-matrix, thanks to (D.17), we can define  $F(\mathbf{v})$  (as in (D.20)) using  $\widehat{G}$  instead of  $\nabla^2 \widehat{R}$ . Then we need concentration bounds for the following quantities

$$\begin{aligned} \langle \mathbf{a}, \widehat{G}_{vv} \mathbf{a} \rangle &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (\widehat{y}^\ell - \widetilde{y}^\ell)^2 \langle \mathbf{a}, g(W\widetilde{Y}^\ell) \rangle^2 Y^\ell, \\ \langle u_i, \widehat{G}_{W_i W_j} u_j \rangle &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (\widehat{y}^\ell - \widetilde{y}^\ell)^2 v_i v_j g'(W_i \cdot \widetilde{Y}^\ell) g'(W_j \cdot \widetilde{Y}^\ell) \langle u_i, \widetilde{Y}^\ell \rangle \langle u_j, \widetilde{Y}^\ell \rangle, \\ \langle \mathbf{a}, \widehat{G}_v u_j \rangle &= \frac{1}{\widetilde{M}} \sum_{\ell=1}^{\widetilde{M}} (\widehat{y}^\ell - \widetilde{y}^\ell)^2 v_j g'(W_j \cdot \widetilde{Y}^\ell) \langle \mathbf{a}, g(W \cdot \widetilde{Y}^\ell) \rangle \langle u_j, \widetilde{Y}^\ell \rangle. \end{aligned} \quad (\text{D.23})$$



By the same argument as that after (D.21), each summand in (D.23) is uniformly (with constant only depending on  $K, L$ ) sub-exponential. The concentration of this new  $F(\mathbf{a}, \mathbf{u})$  follows from Bernstein's inequality by the same argument. Then the concentration estimates of this new  $F(\mathbf{a}, \mathbf{u})$  together with a epsilon-net argument gives (D.15).  $\square$

## E PROOFS OF MAIN THEOREMS

In this section we put together the ingredients we have established in the preceding sections to deduce our main theorems, Theorems 3.1–4.2.

### E.1 1-LAYER NETWORK FOR MIXTURE OF $k$ GAUSSIANS

We begin with the theorems for classification of the  $k$ -GMM with a single layer network.

**Proof of Theorem 3.2.** We prove the alignment (up to multiplicative error  $O(\epsilon + \lambda^{-1})$ ) of the SGD trajectory, the Hessian, and the G-matrix, with  $\text{Span}(\mu_1, \dots, \mu_k)$  one at a time.

The results for the SGD were exactly the content of item (2) of Proposition C.1. Namely, that item tells us that the part of  $\mathbf{x}_\ell^c$  orthogonal to  $\text{Span}(\mu_1, \dots, \mu_k)$  has norm at most  $O(\epsilon + \lambda^{-1})$  while  $\|\mathbf{x}_\ell^c\| \geq \eta - O(\epsilon + \lambda^{-1}) \geq \eta/2$  for an  $\eta$  independent of  $\epsilon, \lambda$ . Absorbing  $\eta^{-1}$  into the big- $O$ , this is exactly the definition of living in a subspace per Definition 2.1.

We recall from Lemma A.1, for  $b \neq c$ , the  $bc$ -block of the population Hessian is given by

$$\mathbb{E}[\nabla_{cb}^2 \widehat{R}(\mathbf{x})] = \sum_{l \in [k]} p_l \Pi_{cb}^l \mu_l^{\otimes 2} + \mathcal{E}_{cb}^R, \quad (\text{E.1})$$

where  $\Pi_{cb}^l = \mathbb{E}[\pi_{Y_l}(c)\pi_{Y_l}(b)]$  for  $Y_l = \mu_l + Z_\lambda$ , and where  $\mathcal{E}_{cb}^R$  is a matrix with operator norm bounded by  $\mathcal{O}(1/\lambda)$ . By Lemma A.2, the  $aa$  diagonal block of the population Hessian is given by

$$\mathbb{E}[\nabla_{aa}^2 \widehat{R}(\mathbf{x})] = \sum_{l \in [k]} p_l \Pi_{aa}^l \mu_l^{\otimes 2} + \mathcal{E}_{aa}^R, \quad (\text{E.2})$$

where  $\Pi_{aa}^l = \mathbb{E}[\pi_{Y_l}(a)(1 - \pi_{Y_l}(a))]$  and  $\mathcal{E}_{aa}^R$  is a matrix with norm bounded by  $\mathcal{O}(1/\lambda)$ .

The two expressions (E.1) and (E.2), together with the concentration of the empirical Hessian matrix Theorem D.1 implies that blocks of the empirical Hessian concentrate around low rank matrices; i.e., for each  $\mathbf{x}$ , we have

$$\begin{aligned} \nabla_{cb}^2 \widehat{R}(\mathbf{x}) &= \sum_{1 \leq l \leq k} p_l \Pi_{cb}^l \mu_l^{\otimes 2} + \mathcal{E}_{cb}^R + (\nabla_{cb}^2 \widehat{R}(\mathbf{x}) - \nabla_{cb}^2 \mathbb{E}[\widehat{R}(\mathbf{x})]) \\ &=: \sum_{1 \leq l \leq k} p_l \Pi_{cb}^l \mu_l^{\otimes 2} + \widehat{\mathcal{E}}_{cb}^R. \end{aligned} \quad (\text{E.3})$$

where for each fixed  $\mathbf{x}$ , except with probability  $e^{-(c\alpha\epsilon^2 - C)d}$  we have  $\|\widehat{\mathcal{E}}_{cb}^R\| \lesssim \epsilon + \frac{1}{\lambda}$ . Since the test and training data are independent of one another, this also means for any realization of the SGD trajectory  $(\mathbf{x}_\ell)_{\ell \leq T_f \delta^{-1}}$ , and so long as  $T_f \delta^{-1} = e^{o(d)}$ , by a union bound we get

$$\mathbb{P}\left(\max_{\ell \leq T_f \delta^{-1}} \|\nabla_{cb}^2 \widehat{R}(\mathbf{x}_\ell) - \sum_{l \in [k]} p_l \Pi_{cb}^l(\mathbf{x}_\ell) \mu_l^{\otimes 2}\| \geq C(\epsilon + \lambda^{-1})\right) = o_d(1). \quad (\text{E.4})$$

where we've put in the  $\mathbf{x}_\ell$  to emphasize the dependence of  $\Pi_{cb}^l$  on the location in parameter space.

Moreover, we recall from item (1) of Proposition C.1, that there exists a constant  $L$  such that for all  $\lambda$  large, with probability  $1 - o_d(1)$ , the SGD trajectory has  $\|\mathbf{x}_\ell\| \leq L$  for all  $T_0 \delta^{-1} \leq \ell \leq T_f \delta^{-1}$ . It follows by definition of  $\pi$  that there exists a constant  $c = c(L) > 0$ , such that the coefficients in (E.1) and (E.2) are lower bounded:  $p_l \Pi_{cb}^l(\mathbf{x}_\ell), p_l \Pi_{aa}^l(\mathbf{x}_\ell) \geq c$  for all  $T_0 \delta^{-1} \leq \ell \leq T_f \delta^{-1}$ . Thus the first sum  $\sum_{1 \leq l \leq k} p_l \Pi_{cb}^l \mu_l^{\otimes 2}$  on the righthand side of (E.3) is positive definite, and its norm is lower bounded by  $c$  (uniformly in  $\epsilon, \lambda$ ). Together with (E.4), we conclude that the  $b, c$  blocks of the test

Hessian  $\nabla_{bc}^2 \widehat{R}(\mathbf{x}_\ell)$  live in  $\text{Span}(\mu_1, \mu_2, \dots, \mu_k)$  up to error  $\mathcal{O}(\epsilon + \lambda^{-1})$ . Namely, this is because it satisfies Definition 2.2 with the choice of  $M = \widehat{\mathcal{E}}_{cb}^R$ , after absorbing  $c^{-1}$  into the big- $\mathcal{O}$ .

By the same argument, thanks to Lemma A.3, the  $bc$  block of the population G-matrix is given by

$$\delta_{bc} p_b \mu_b^{\otimes 2} - p_c \mathbb{E}[\pi_{Y_c}(b)] \mu_c^{\otimes 2} - p_b \mathbb{E}[\pi_{Y_b}(c)] \mu_b^{\otimes 2} + \sum_l p_l \mathbb{E}[\pi_{Y_l}(b) \pi_{Y_l}(c)] \mu_l^{\otimes 2} + \mathcal{E}_{bc}^G, \quad (\text{E.5})$$

where  $\mathcal{E}_{bc}^R$  is a matrix with norm bounded by  $\mathcal{O}(1/\lambda)$ . The expression (E.5), together with the concentration of empirical G-matrix Theorem D.1, and a union bound over  $\ell \leq T_f \delta^{-1}$ , implies that blocks of the empirical G-matrix concentrate around

$$\nabla_{bc}^2 \widehat{G}(\mathbf{x}_\ell) = \delta_{bc} p_b \mu_b^{\otimes 2} - p_c \mathbb{E}[\pi_{Y_c}(b)] \mu_c^{\otimes 2} - p_b \mathbb{E}[\pi_{Y_b}(c)] \mu_b^{\otimes 2} + \sum_l p_l \mathbb{E}[\pi_{Y_l}(b) \pi_{Y_l}(c)] \mu_l^{\otimes 2} + \widehat{\mathcal{E}}_{bc}^G(\mathbf{x}_\ell), \quad (\text{E.6})$$

where except with probability  $T_f \delta^{-1} e^{-(c\alpha\epsilon^2 - C)d}$  we have  $\|\widehat{\mathcal{E}}_{ca}^G(\mathbf{x}_\ell)\| \lesssim \epsilon + \frac{1}{\lambda}$  for all  $\ell \leq T_f \delta^{-1}$ . Namely, the analogue of (E.4) will apply to  $\nabla_{bc}^2 \widehat{G}$  about the low-rank part of the above. In order to deduce the claimed alignment of Theorem 3.2 it remains to show that each block of the matrix in (E.6) (minus  $\widehat{\mathcal{E}}_{bc}^G$  has operator norm bounded away from zero uniformly in  $\epsilon, \lambda$ . Towards this, recall that we can work on the event that  $\|\mathbf{x}_\ell\| \leq L$ . In the on-diagonal blocks, we can rewrite the part of (E.6) that is not  $\widehat{\mathcal{E}}_{bb}^G$  as

$$p_b \mathbb{E}[(1 - \pi_{Y_b}(b))^2] \mu_b^{\otimes 2} + \sum_{l \neq b} p_l \mathbb{E}[\pi_{Y_l}(b)^2] \mu_l^{\otimes 2}. \quad (\text{E.7})$$

This matrix is positive definite and since  $\|\mathbf{x}_\ell\| \leq L$ , there exists  $c(L)$  such that the coefficients of  $\mu_c$  are all bounded away from zero by  $c$ , so that the norm of the above is bounded from below by zero for all  $(\mathbf{x}_\ell)_{T_0 \delta^{-1} \leq \ell \leq T_f \delta^{-1}}$ .

Thus by C.1, the coefficients  $p_b \mathbb{E}[(1 - \pi_{Y_b}(b))^2]$ ,  $p_l \mathbb{E}[\pi_{Y_l}(b)^2]$  in (E.7) are lower bounded. Thus (E.7) is positive definite, there exists a constant  $c > 0$ , such that its norm is lower bounded by  $c$ .

For  $b \neq c$ , we can lower bound the operator norm of the matrix

$$-p_c \mathbb{E}[\pi_{Y_c}(b)(1 - \pi_{Y_c}(c))] \mu_c^{\otimes 2} - p_b \mathbb{E}[\pi_{Y_b}(c)(1 - \pi_{Y_b}(b))] \mu_b^{\otimes 2} + \sum_{l \neq b, c} p_l \mathbb{E}[\pi_{Y_l}(b) \pi_{Y_l}(c)] \mu_l^{\otimes 2}, \quad (\text{E.8})$$

while  $\|\mathbf{x}_\ell\| \leq L$  using that if  $k > 2$  then the last sum contributes some positive portion outside of  $\text{Span}(\mu_c, \mu_b)$  which can be used to lower bound the operator norm by some  $c(L)$  and if  $k = 2$ , then the first two terms are the only two, are negative definite, and have coefficients similarly bounded away from zero by some  $-c(L)$ .  $\square$

**Proof of Theorem 3.1.** The theorem follows from Theorem 3.2 together with the observation that the matrices in (E.1) and (E.5) that are not the error portion  $\mathcal{E}_{cb}^R$  and  $\mathcal{E}_{bc}^G$  respectively, are of rank  $k$  since they are sums of  $k$  rank-1 matrices and as explained in the above proof, each of their eigenvalues are bounded away from zero uniformly in  $\epsilon, \lambda$ , (in a manner depending only on  $L$ ).  $\square$

**Proof of Theorem 3.3.** The claims regarding the SGD are proved in Proposition C.14.

It remains to prove alignment of the top eigenvectors of the  $cc$ -blocks of the Hessian and G-matrices with  $\mu_c$ . Recall from (E.3), the decomposition of the empirical Hessian matrix

$$\nabla_{aa}^2 \widehat{R}(\mathbf{x}) =: \sum_{1 \leq l \leq k} p_l \Pi_{aa}^l \mu_l^{\otimes 2} + \widehat{\mathcal{E}}_{aa}^R, \quad (\text{E.9})$$

where  $\Pi_{aa}^l = \mathbb{E}[\pi_{Y_l}(a)(1 - \pi_{Y_l}(a))]$ , and with probability  $e^{-(c\alpha\epsilon^2 - C)d}$  we have  $\|\widehat{\mathcal{E}}_{ca}^R\| \lesssim \epsilon + \frac{1}{\lambda}$ .

By our assumption that  $\mu_1, \mu_2, \dots, \mu_k$  are orthonormal, thus the first part in the decomposition (E.9) can be viewed as an orthogonal decomposition.  $\nabla_{aa}^2 \widehat{R}(\mathbf{x})$  is a perturbation of  $\sum_{1 \leq l \leq k} p_l \Pi_{aa}^l \mu_l^{\otimes 2}$  by  $\widehat{\mathcal{E}}_{aa}^R$ . In turn, by the same reasoning as in Lemma C.11,

$$\Pi_{aa}^l = \bar{\Pi}_{aa}^l + \mathcal{O}(\lambda^{-1}) \quad \text{where} \quad \bar{\Pi}_{aa}^l = \mathbb{E}[\bar{\pi}_l(a)(1 - \bar{\pi}_l(a))],$$

for  $\bar{\pi}_l(a) = e^{mat} / \sum_b e^{mbi}$  as in (C.16). By Lemma C.19, there exists  $c(\beta) > 0$  such that for all  $(\mathbf{x}_\ell)_{T_0\delta^{-1} \leq \ell \leq T_f\delta^{-1}}$ , all the coefficients  $p_l \Pi_{aa}^l$  in (E.9) are lower bounded:  $p_l \Pi_{aa}^l \geq c$ . Thus  $\sum_{1 \leq l \leq k} p_l \Pi_{aa}^l \mu_l^{\otimes 2}$  has  $k$  positive eigenvalues,  $\{p_l \Pi_{aa}^l\}_{1 \leq l \leq k}$ , and each of them is lower bounded by  $c$ . The associated eigenvectors are given by  $\{\mu_l\}_{1 \leq l \leq k}$ . We furthermore claim that the one corresponding to  $\mu_a^{\otimes 2}$  is separated from the others uniformly in  $\epsilon, \lambda$ . This follows from the fact that we derived in Proposition C.14 that  $\mathbf{x}_\ell$  is within  $O(\epsilon + \lambda^{-1})$  distance of a point  $\mathbf{x}_*$  such that  $\bar{\pi}_b(a) = \frac{1}{k-1}(1 - \bar{\pi}_c(c))$  as long as  $a \neq b$ . This implies that  $\bar{\pi}_l(a)$  for  $l \neq a$  is closer to 0 than  $\bar{\pi}_a(a)$  is close to 1, whence  $\bar{\pi}_a(a)(1 - \bar{\pi}_a(a)) > \bar{\pi}_l(a)(1 - \bar{\pi}_l(a))$  by an amount that is uniform in  $\epsilon, \lambda$ . This ensures that along  $(\mathbf{x}_\ell)$  the largest eigenvector in  $\sum_l \Pi_{aa}^l \mu_a^{\otimes 2}$  is the one with eigenvector  $\mu_a$  and the next  $k - 1$  are those corresponding to  $(\mu_l)_{l \neq a}$ . Altogether, by eigenvector stability, the top eigenvector of  $\nabla_{aa}^2 \widehat{R}(\mathbf{x}_\ell)$  lives in  $\text{Span}(\mu_a)$  up to error  $\mathcal{O}(\epsilon + \lambda^{-1})$  and the next  $k - 1$  all live in  $\text{Span}((\mu_l)_{l \neq a})$  up to an error  $\mathcal{O}(\epsilon + 1/\lambda)$  for all  $\ell \in [T_0\delta^{-1}, T_f\delta^{-1}]$ .

The statement for the empirical G-matrix is established similarly by using (E.6) as input. The part that is different from the above is to see that its top eigenvector in the  $aa$ -block is the one corresponding approximating by  $\mu_a$  and its next  $k - 1$  are those approximating  $(\mu_l)_{l \neq a}$ . For that, recall the expression (E.7) and use that since  $(1 - \bar{\pi}_a(a)) > \bar{\pi}_l(a)$ , the coefficient of  $\mu_a^{\otimes 2}$  in the  $aa$ -block is strictly larger than the coefficients of  $\mu_l$  for  $l \neq a$ .  $\square$

## E.2 2-LAYER NETWORK FOR XOR GAUSSIAN MIXTURE

We now turn to proving our main theorems for the XOR Gaussian mixture model with a 2-layer network of width  $K$ .

**Proof of Theorem 4.2.** We prove the alignment individually for each of the SGD, the Hessian matrix, and the G-matrix. For the SGD trajectory, the claimed alignment was exactly the content of item (2) in Proposition C.2.

Letting  $F$  denote the cdf of a standard Gaussian, thanks to Lemma B.4,

$$\mathbb{E}[\nabla_{vv}^2 L] = \frac{1}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma'(v \cdot g(W\vartheta)) g(W\vartheta)^{\otimes 2} + \mathcal{E}_{vv}^R, \quad (\text{E.10})$$

$$\mathbb{E}[\nabla_{W_i W_i}^2 L] = \frac{v_i^2}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} F\left(m_i^\vartheta \sqrt{\frac{\lambda}{R_{ii}}}\right) \vartheta^{\otimes 2} + \mathcal{E}_{W_i W_i}^R, \quad (\text{E.11})$$

where the two error matrices satisfy  $\|\mathcal{E}_{vv}^R\|, \|\mathcal{E}_{W_i W_i}^R\| = \mathcal{O}(1/\sqrt{\lambda})$ . The two expressions (E.10) and (E.11), together with the concentration of empirical Hessian matrix Theorem D.2 imply that for every  $L$ , every fixed  $\mathbf{x} = (v, W)$  with  $W \notin \mathcal{W}_0^c$ , the blocks of the empirical Hessian matrix satisfy

$$\nabla_{vv}^2 \widehat{R}(v, W) = \frac{1}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma'(v \cdot g(W\vartheta)) g(W\vartheta)^{\otimes 2} + \widehat{\mathcal{E}}_{vv}^R, \quad (\text{E.12})$$

$$\nabla_{W_i W_i}^2 \widehat{R}(v, W) = \frac{v_i^2}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} F\left(m_i^\vartheta \sqrt{\frac{\lambda}{R_{ii}}}\right) \vartheta^{\otimes 2} + \widehat{\mathcal{E}}_{W_i W_i}^R, \quad (\text{E.13})$$

where except with probability  $e^{-(c\alpha\epsilon^2 - C)d}$ , we have

$$\|\widehat{\mathcal{E}}_{vv}^R(v, W)\|_{\text{op}}, \|\widehat{\mathcal{E}}_{W_i W_i}^R(v, W)\|_{\text{op}} \lesssim \epsilon + \frac{1}{\sqrt{\lambda}}. \quad (\text{E.14})$$

Using independence of the test and training data, recalling from item (1) of Proposition C.2 that the SGD stays confined to a ball of radius  $L$  in parameter space for all  $\ell \leq T_f\delta^{-1}$ , and from item (2) of Proposition C.2 that  $\|W_i\| > 0$  for all  $i$  for all  $T_0\delta^{-1} \leq \ell \leq T_f\delta^{-1}$ , taking a union bound over

$\ell \leq T_f \delta^{-1}$ , we get for  $\alpha > \alpha_0$ ,

$$\mathbb{P}\left(\max_{T_0 \delta^{-1} \leq \ell \leq T_f \delta^{-1}} \|\nabla_{vv}^2 \widehat{R}(\mathbf{x}_\ell) - \frac{1}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma'(v \cdot g(W\vartheta)) g(W\vartheta)^{\otimes 2}\|_{\text{op}} \geq C(\varepsilon + \lambda^{-1/2})\right) = o(1). \quad (\text{E.15})$$

$$\mathbb{P}\left(\max_{T_0 \delta^{-1} \leq \ell \leq T_f \delta^{-1}} \|\nabla_{W_i W_i}^2 \widehat{R}(\mathbf{x}_\ell) - \frac{v_i^2}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} F\left(m_i^\vartheta \sqrt{\frac{\lambda}{R_{ii}}}\right) \vartheta^{\otimes 2}\|_{\text{op}} \geq C(\varepsilon + \lambda^{-1/2})\right) = o(1). \quad (\text{E.16})$$

where in the quantity that the blocks of the empirical Hessian are being compared to,  $v, W$  evaluated along the SGD trajectory  $\mathbf{x}_\ell$ .

It remains to show that the low-rank matrices in (E.15)–(E.16) have some operator norm uniformly bounded away from zero to deduce the alignment of the empirical test Hessian with the claimed vectors. By item (2) of Proposition C.2, there exists some constant  $c > 0$  uniform in  $\lambda$ , such that with probability  $1 - o_d(1)$ , for all  $\ell \leq T_f \delta^{-1}$  the SGD  $\mathbf{x}_\ell$  is such that

$$\max_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma'(v \cdot g(W\vartheta)) \|g(W\vartheta)\| = \max_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma'\left(\sum_i v_i g(m_i^\vartheta)\right) \|g(W\vartheta)\| > c, \quad (\text{E.17})$$

Thus the deterministic matrix in (E.15)  $\sum_{\vartheta \in \{\pm\mu, \pm\nu\}} \sigma'(v \cdot g(W\vartheta)) g(W\vartheta)^{\otimes 2}$  is positive definite and its norm is lower bounded away from 0 for all  $(\mathbf{x}_\ell)_{\ell \leq T_f \delta^{-1}}$ . Together with (E.15), we conclude that the second-layer test Hessian  $\nabla_{vv}^2 \widehat{R}(\mathbf{x}_\ell)$  live in  $\text{Span}(g(W(\mathbf{x}_\ell) \cdot \vartheta)_{\vartheta \in \{\pm\mu, \pm\nu\}})$ .

For (E.16), by item (2) of Proposition C.2, there exists  $c > 0$  (independent of  $\epsilon, \lambda$ ) such that with probability  $1 - o_d(1)$ , for every  $i$ , it must be the case that

$$\max_{\vartheta \in \{\pm\mu, \pm\nu\}} \frac{v_i^2}{4} F\left(m_i^\vartheta \sqrt{\frac{\lambda}{R_{ii}}}\right) > c \quad \text{for all } (\mathbf{x}_\ell)_{\ell \leq T_f \delta^{-1}}. \quad (\text{E.18})$$

Then the deterministic matrix in (E.16) is positive definite and has norm uniformly lower bounded away from 0 for all  $\mathbf{x}_\ell$  for  $\ell \leq T_f \delta^{-1}$ . Together with (E.16), we conclude that the first-layer test Hessian  $\nabla_{W_i W_i}^2 \widehat{R}(\mathbf{x}_\ell)$  lives in  $\text{Span}(\mu, \nu)$ .

By the same argument, thanks to Lemma B.5, together with the concentration of the empirical G-matrix from Theorem D.2, the blocks of the empirical G-matrix concentrate around low rank matrices

$$\widehat{G}_{vv}(v, W) = \frac{1}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} (\sigma(v \cdot g(W\vartheta)) - y_\vartheta)^2 g(W\vartheta)^{\otimes 2} + \widehat{\mathcal{E}}_{vv}^G, \quad (\text{E.19})$$

$$\widehat{G}_{W_i W_i}(v, W) = v_i^2 A + \widehat{\mathcal{E}}_{W_i W_i}^G,$$

where recalling (B.18),

$$A = \frac{1}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} (y_\vartheta - \sigma(v \cdot g(W\vartheta)))^2 F\left(m_i^\vartheta \sqrt{\frac{\lambda}{R_{ii}}}\right) \vartheta^{\otimes 2}, \quad (\text{E.20})$$

and for every  $\mathbf{x} = (v, W)$  of norm at most  $L$ , except with probability  $e^{-(c\alpha\varepsilon^2 - C)d}$  we have  $\|\widehat{\mathcal{E}}_{vv}^G\|, \|\widehat{\mathcal{E}}_{W_i W_i}^G\| \lesssim \varepsilon + \frac{1}{\sqrt{\lambda}}$ . Union bounding over the  $T_f \delta^{-1}$  points along the trajectory of the SGD, as in the lead-up to (E.15)–(E.16), we get

$$\mathbb{P}\left(\max_{\ell \leq T_f \delta^{-1}} \|\widehat{G}_{vv}(\mathbf{x}_\ell) - \frac{1}{4} \sum_{\vartheta \in \{\pm\mu, \pm\nu\}} (\sigma(v \cdot g(W\vartheta)) - y_\vartheta)^2 g(W\vartheta)^{\otimes 2}\|_{\text{op}} \geq C(\varepsilon + \lambda^{-1/2})\right) = o(1). \quad (\text{E.21})$$

$$\mathbb{P}\left(\max_{\ell \leq T_f \delta^{-1}} \|\widehat{G}_{W_i W_i}(\mathbf{x}_\ell) - v_i^2 A\|_{\text{op}} \geq C(\varepsilon + \lambda^{-1/2})\right) = o(1). \quad (\text{E.22})$$

where in the quantity that the blocks of the empirical G-matrix are being compared to,  $v, W$  are evaluated along the SGD trajectory  $\mathbf{x}_\ell$ .

We recall from item (1) of Proposition C.2 the SGD stays inside the  $\ell^2$ -ball of radius  $L$  for all time. Thus for  $\mathbf{x} = \mathbf{x}_\ell$ , the coefficients  $(\sigma(v \cdot g(W\vartheta)) - y_\vartheta)^2$ ,  $\sigma(v \cdot g(W\vartheta))^2$ ,  $(1 - \sigma(v \cdot g(W\vartheta)))^2$  in (E.19) and (E.20) are lower bounded away from 0. By the same argument as for the test Hessian matrix, we conclude that the second-layer test G-matrix  $\widehat{G}_{vv}(\mathbf{x}_\ell)$  lives in  $\text{Span}(g(W(\mathbf{x}_\ell)\vartheta)_{\vartheta \in \{\pm\mu, \pm\nu\}})$  and its first layer lives in  $\text{Span}(\mu, \nu)$ , up to error  $O(\epsilon + \lambda^{-1/2})$ .  $\square$

**Proof of Theorem 4.1.** The first layer alignment follows from Theorem 4.2 together with the observation that the part of (E.13) besides  $\widehat{\mathcal{E}}_{W_i W_i}^R$  is a rank-2 matrix with eigenvectors  $\mu, \nu$  and having both eigenvectors bounded away from zero uniformly in  $(\epsilon, \lambda)$ ; this latter fact comes from the observation that the sum of the two cdf's  $F(a) + F(-a) = 1$  always, and  $v_i$  being bounded away from zero per part (2) of Proposition C.2. A similar behavior ensures the same for the G-matrix's first layer top two eigenvectors per (E.19)–(E.20) and a uniform lower bound on the sigmoid function over all possible parameters in a ball of radius  $L$ , as guaranteed by part (1) of Proposition C.2.

For the alignment of the second layer, it follows from Theorem 4.2 together with the following. First observe that the part of (E.12) that is not  $\widehat{\mathcal{E}}_{vv}^R$  is a rank-4 matrix with eigenvectors  $(g(W\vartheta))_{\vartheta \in \{\pm\mu, \pm\nu\}}$ . To reason that all of the corresponding eigenvectors are uniformly (in  $\epsilon, \lambda$ ) bounded away from zero, use the uniform lower bound on  $\sigma'$  while the parameters are in a ball of radius  $L$  about the origin as promised by part (1) of Proposition C.2, together with the uniform lower bound, for each  $i$ , on one of  $(W_i \cdot \vartheta)_{\vartheta \in \{\pm\mu, \pm\nu\}}$  which holds for the SGD after  $T_0(\epsilon)$  per the proof of part (2) of Proposition C.2.  $\square$

## F ADDITIONAL FIGURES

### F.1 ADDITIONAL FIGURES FOR K-GMM

In this section, we include more figures depicting the spectral transitions for the  $k$ -GMM model.

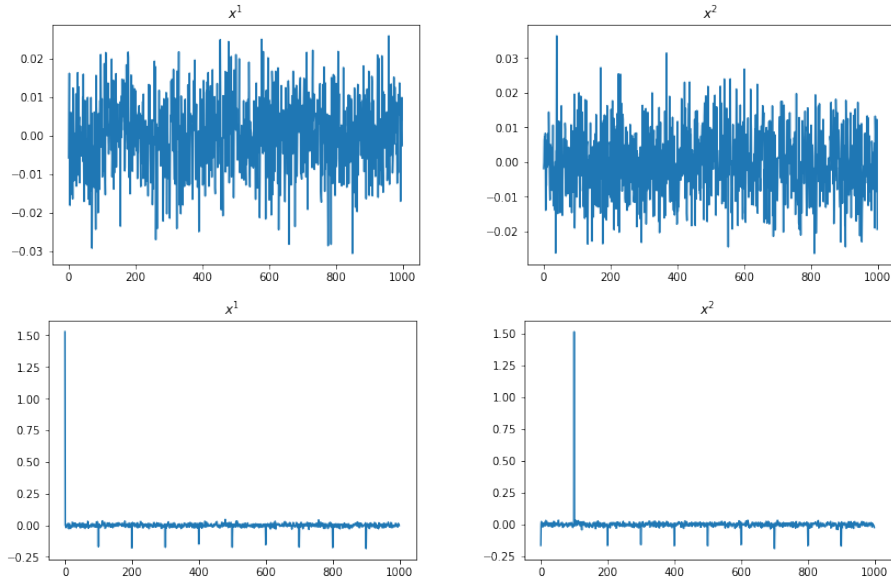


Figure F.1:  $x^1$  and  $x^2$  at initial and final training times. Initially they are random and then they correlate (and anti-correlate) as per Theorem 3.2. Parameters are the same as in Figure 3.1.




Figure F.2: Further illustration of dynamical BBP from Fig. 3.3. Here we see that there is initially two components of the spectrum, then as time evolves, the top eigenvalue and the next 9 separate from each other.

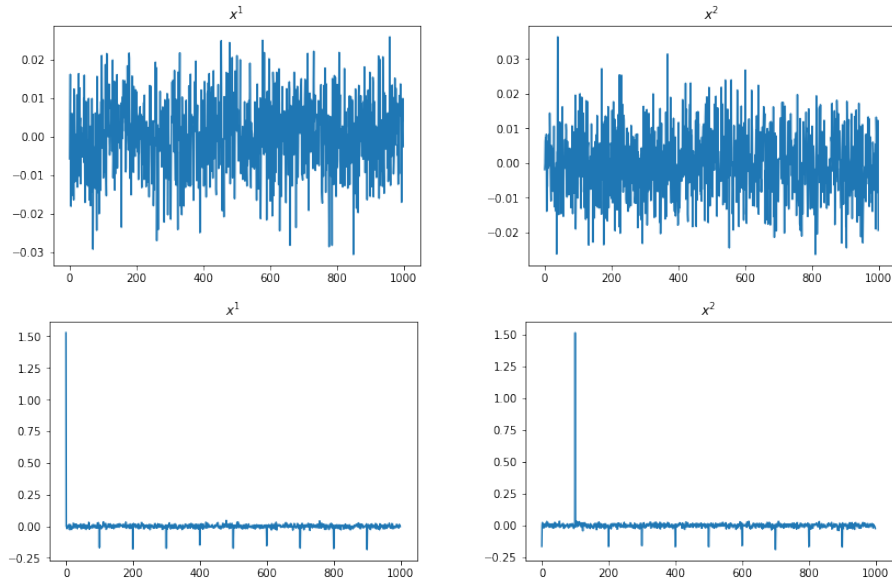


Figure F.3:  $x^1$  and  $x^2$  at initial and final training times. Initially they are random and then they correlate (and anti-correlate) as per Theorem 3.2. Parameters are the same as in Figure 3.1.

## F.2 TRAINING DATA

We include here a few plots of the preceding discussion run with training data, as opposed to test data. In the following, the plots are generated using the Hessian of the empirical risk computed over the entire training data set and similarly the G-matrix generated using the entire training data set, as opposed to the test data as in earlier plots.

### F.2.1 K-COMPONENT MIXTURE MODEL




Figure F.4: Parameters are the same as in Fig. 3.1




Figure F.5: Training data analogue of Fig. 3.3. the eigenvalues of  $\nabla^2 \widehat{R}_{11}(\mathbf{x}_\ell)$  over the course of training (but with training data. The leading  $k$  eigenvalues are separated from the bulk at all times, and the top eigenvalue, corresponding to  $\mu_1$  separates from the remaining eigenvalues soon after initialization. Right: the inner product of  $\mathbf{x}_\ell^1$  with the means  $\mu_1, \dots, \mu_k$  undergoes a similar separation over the course of training. Parameters are the same as in preceding figures.

### F.2.2 XOR MIXTURE MODEL




Figure F.6: Training data analogue of Fig. 4.1 and (b) depict the alignment of the first layer weights  $W_i(\mathbf{x}_\ell)$  with the principal subspaces of the corresponding blocks of the Hessian and G-matrices formed using training data, i.e., with  $E_2(\nabla_{W_i W_i}^2 \widehat{R}(\mathbf{x}_\ell))$  and  $E_2(\widehat{G}_{W_i W_i}(\mathbf{x}_\ell))$ . (c) and (d) plot the second-layer alignment, namely of  $v(\mathbf{x}_\ell)$  with  $E_4(\nabla_{vv}^2 \widehat{R}(\mathbf{x}_\ell))$  and  $E_4(\widehat{G}_{vv}(\mathbf{x}_\ell))$ . Parameters are  $d = 500$ ,  $\lambda = 10$ , and  $K = 20$




Figure F.7: Training data analogue of Fig. 4.2 The eigenvalues of the  $vv$  blocks of the Hessian and G-matrices over time from a random initialization. Along training, the top four eigenvalues separate from the bulk, and four outlier eigenvalues emerge, corresponding to the four “hidden” classes in the XOR problem. Parameters are the same as in Figure 4.1.