
Imitating Deep Learning Dynamics via Locally Elastic Stochastic Differential Equations

Jiayao Zhang (张家耀) Hua Wang (王华) Weijie J. Su (苏炜杰)

University of Pennsylvania

{zjiayao, wanghua, suw}@wharton.upenn.edu

Abstract

Understanding the training dynamics of deep learning models is perhaps a necessary step toward demystifying the effectiveness of these models. In particular, how do data from different classes gradually become separable in their feature spaces when training neural networks using stochastic gradient descent? In this study, we model the evolution of features during deep learning training using a set of stochastic differential equations (SDEs) that each corresponds to a training sample. As a crucial ingredient in our modeling strategy, each SDE contains a drift term that reflects the impact of backpropagation at an input on the features of all samples. Our main finding uncovers a sharp phase transition phenomenon regarding the *intra-class* impact: if the SDEs are *locally elastic* [19] in the sense that the impact is more significant on samples from the same class as the input, the features of the training data become linearly separable, meaning vanishing training loss; otherwise, the features are not separable, regardless of how long the training time is. Moreover, in the presence of local elasticity, an analysis of our SDEs shows that the emergence of a simple geometric structure called the neural collapse of the features. Taken together, our results shed light on the decisive role of local elasticity in the training dynamics of neural networks. We corroborate our theoretical analysis with experiments on a synthesized dataset of geometric shapes and CIFAR-10.

1 Introduction

Deep learning models have achieved significant empirical success over the past decade across a wide spectrum of domains spanning computer vision, natural language processing, and reinforcement learning [31, 43, 48]. Despite these remarkable achievements at the empirical level, there is still much to learn about deep neural networks, as evidenced by the fact that almost all important advances concerning architecture design and optimization for deep learning are based on heuristics, without much input from a theoretical perspective [20, 11, 21, 27].

An important step toward opening these black-box models and unveiling their formidable details is to quantitatively understand the impact of backpropagation in deep learning training. While there has been a continued effort to demystify how simple optimization methods give rise to impressive generalization performance, for example, [49, 26, 4], this is by no means an easy problem, perhaps because of the daunting nonconvex nature of neural networks. Accordingly, for near-term purposes, a more practical approach is to take a phenomenological viewpoint by relating simple empirical patterns to the effectiveness of deep learning models.

In this spirit, we are interested in how data from different classes gradually become separable in their feature space by repetitively calling backpropagation. From a phenomenological viewpoint, this question can be addressed by first analyzing the impact of a single update using a stochastic gradient on the performance of the neural networks. More precisely, imagine that the gradient is evaluated in an image of a cat, how does the hidden representation of another image—say, an image of another

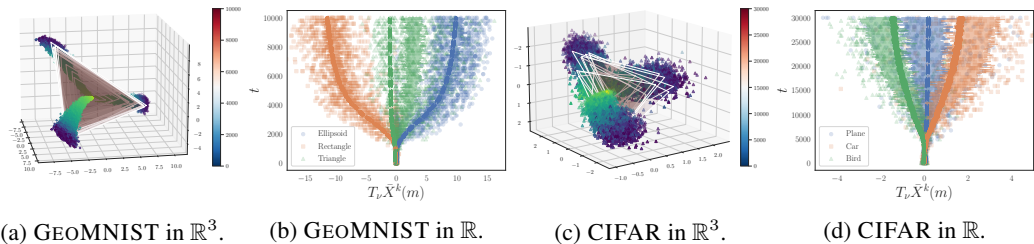


Figure 1: **Separation of features (logits)**. GEOMNIST dataset ((a)—(b)) and CIFAR dataset ((c)—(d)) trained using $K = 3$ classes on a variant of the ALEXNET model. Separation in \mathbb{R} is done by projecting the logits to $\nu \in \mathbb{R}^3$, which is set to be the difference between a pair of class means $\bar{X}^k(T) - \bar{X}^l(T)$ for a large T , where k, l are chosen heuristically. The dashed lines in (b) and (d) are simulated paths from Equation (9) using estimated $\hat{\alpha}(t)$ and $\hat{\beta}(t)$. More details are given in Section 4 and Appendix D.2.

cat or an image of a plane—evolves because of the backpropagation? Recent studies answer this question by introducing a phenomenon called *local elasticity*, which, roughly speaking, means that the impact is generally larger on a similar sample (an image of another cat) than on a dissimilar sample (an image of a plane) [19].

Motivated by the phenomenon of local elasticity, we propose a model that captures the interaction between different training samples during deep learning training using a set of stochastic differential equations (SDEs) that reflect local elasticity in neural networks. Characterizing the *intra-class* and *inter-class* effects is an essential component of our modeling strategy, each of which contains a drift term that imitates the impact of backpropagation on specific training data of all samples.

Our main finding uncovers a sharp phase transition phenomenon regarding the intra-class and inter-class impact; if the SDEs are *locally elastic* in the sense that the impact is more significant on samples from the same class as the input (the intra-class effect is strictly greater than the inter-class effect), the features of the training data are guaranteed to be linearly separable, meaning vanishing training loss; otherwise, the features are not separable, no matter how long the training time is. This result provides convincing theoretical evidence for the presence of *local elasticity* in deep learning [19]. Our model is also quite accurate in simulating the feature dynamics of deep learning. As shown in Figure 1 and detailed in Section 4, the dynamics of the simulated logits are reasonably close to the real dynamics of deep learning on both synthetic and real datasets, indicating a well-suited model for theoretical and practical purposes. Moreover, in the presence of local elasticity, our SDEs also predict the emergence of a simple geometric structure called neural collapse of features [38].

Taken together, our results shed light on the decisive role of local elasticity in the training dynamics of neural networks. We corroborate our theoretical analysis with experiments on a synthetic dataset of geometric shapes, as well as on CIFAR-10. The experimental evidence consistently supported our model, which provides new insights into the dynamics of deep learning training.

1.1 Related Work

Dynamics in Deep Neural Nets. Many properties of linear deep neural nets are relatively well understood, such as the loss landscapes [25], trajectory-based convergence [2, 12], and implicit acceleration [3]. Exact solutions of the training dynamics can be obtained in certain initialization schemes [40, 41, 30]. In the presence of non-linearity, various assumptions are generally made. [39, 17] studied the dynamics of shallow neural nets with non-linearity and the neural tangent kernel (NTK) literature [22, 4, 14] linearizes the network function of an infinitely wide neural net at initialization, which is similar to that of the deep Gaussian process literature [9, 18, 32] (a treatise comparing and contrasting them can be found in [47]). Although as approximations, NTKs are generally used when studying optimization trajectories of neural nets such as [34], which also appears implicitly in many works when studying the optimization trajectories of neural net training [45, 1, 13, 24, 7].

SGD as SDEs in Neural Nets. The study of dynamics or trajectories of weights in deep neural nets via SDEs relies on the more precise characterization of stochasticity [33, 36]. Built on top of

this formalism, [44] studied the trade-off between batch size and learning rate, [23], analyzed factors influencing the quality of local minima, [5] studied the behavior of the SGD near local minima, and [42] studied the effect of learning rates. Although SGD-SDE approximation requires an infinitesimal learning rate, [35] verified that the SDE approximation can be meaningful in practical settings and obtained necessary conditions for the validation of such approximation.

Local Elasticity and Phenomenological Models. Local elasticity is proposed in [19] as a phenomenological approach to reasoning the behaviors of neural networks. This phenomenon has inspired several works on generalization bounds [10] and an improvement on the NTK [6].

2 Binary Separation via LE-SDE

2.1 Setup, Notations and Assumptions

Throughout the paper, we work with the following setup and assumptions. For ease of reading, vectors and matrices are written in boldface and we denote by $[n]$ the set $\{1, \dots, n\}$. When there is no ambiguity, we will write both $X(t)$ and X_t for a continuous-time (possibly stochastic) process.

Classification Problem. Consider a K -class classification problem with $K \geq 2$, with each class having n training examples. We denote by $\mathbf{z}_i^k \in \mathbb{R}^d$ the i -th sample of the k -th class, and $y_i^k \in [K]$ its label, where $i \in [n], k \in [K]$. A neural net is a function $f : \mathbb{R}^d \rightarrow \mathbb{R}^K$ that maps the samples to logits (pre-activation of the softmax).

Feature Vectors. We denote by $\mathbf{X}_i^k(m) \in \mathbb{R}^p$ a p -dimensional feature of the i -th sample in the k -th class learned by the neural net at iteration m . For example, it can be the logits or the output of the second-to-last layer. Assume that the initial values $\mathbf{X}_i^k(0)$ are i.i.d. samples from some distribution for each $i \in [n], k \in [K]$. We use $i, j \in [n]$ as indices for an individual sample, $k, l \in [K]$ for classes, and capital letters J and L to indicate random samples from $\text{Unif}([n])$ and $\text{Unif}([K])$, respectively.

Training Dynamics. We model the training dynamics in neural nets under SGD with an emphasis on *local elasticity*. At the m -th iteration, the J_m -th sample is sampled from the L_m -th class, where $J_m \sim \text{Unif}([n])$ and $L_m \sim \text{Unif}([K])$. Training on $\mathbf{z}_{J_m}^{L_m}$ affects the features of another data sample \mathbf{z}_i^k in the form of

$$\mathbf{X}_i^k(m) - \mathbf{X}_i^k(m-1) = h \cdot E_{k,L_m}(m) \mathbf{X}_{J_m}^{L_m}(m-1) + \sqrt{h} \zeta_i^k(m-1) \quad (1)$$

where $i \in [n], k \in [K]$, h is the step size, and $\zeta_i^k(m)$ is the noise term that is modeled as Gaussian noise. The scalar E_{k,L_m} measures the strength of local elasticity that $\mathbf{z}_{J_m}^{L_m}$ exerts on \mathbf{z}_i^k at iteration m . We assume $\mathbf{X}_i^k(0), \zeta_i^k(m)$ are jointly independent.

Local Elasticity. Clearly, by writing $E_{k,l}(m)$, we assume that this effect depends only on the class l, k , and time m . We write the effect matrix as $\mathbf{E}(m) = (E_{k,l}(m))_{k,l=1}^K$. For ease of exposition, we assume \mathbf{E} only consists of two values $\alpha(m)$ and $\beta(m)$, with $\alpha(m)$ representing the *intra-class* effect and $\beta(m)$ the *inter-class* effect. To this end, we assume the *effective training assumption*, that is, as training progresses, the features become more discriminative: features from the same class are more similar, whereas those from different classes are more distinct, as measured by some similarity measure in the feature space. We also assume that the LE effect is “*proportional*” to the feature $\mathbf{X}_{J_m}^{L_m}(m)$ itself. We generalize this point in Section 3 by introducing a transformation matrix \mathbf{H} on the features.

2.2 Binary LE-SDE

Our construction of equation (1) emphasizes the effect of intra- and inter-class effects on the dynamics of *features*, and thus differs from the usual weight dynamics that is common in the literature. Before deriving the general form of our locally elastic SDE (LE-SDE), we shall familiarize the reader with our model by demonstrating this in the case of binary classification ($K = 2$) with a one-dimensional features ($p = 1$) — the output of the model to be fed into the softmax function, also called the *logit*.

Let the intra-class effect be $E_{11} = E_{22} = \alpha$, and the inter-class effect is $E_{12} = E_{21} = \beta$, both of which are time-independent. Expanding equation (1), for $1 \leq i, j \leq n$ and $m \geq 0$, when we train the

model on the J_m -th training example from the L_m -th class, we have

$$\begin{cases} X_i^1(m) &= X_i^1(m-1) + h \cdot \alpha X_{J_m}^{L_m}(m-1) + \sqrt{h} \cdot \zeta_i^{L_m}(m-1), \\ X_j^2(m) &= X_j^2(m-1) + h \cdot \beta X_{J_m}^{L_m}(m-1) + \sqrt{h} \cdot \zeta_j^{L_m}(m-1). \end{cases}$$

In the limit of $h \rightarrow 0$, we can show that $X_i^k(m)$ approximates some continuous-time stochastic processes $X_i^k(t)$ (under the identification of $t = mh$) governed by the set of stochastic differential equations as follows:

$$dX_i^k(t) = \left(\frac{\alpha}{2} \bar{X}^k(t) + \frac{\beta}{2} \bar{X}^{3-k}(t) \right) dt + \sigma dW_i^k(t), \quad t \geq 0, k \in [K], i \in [n] \quad (2)$$

where $\bar{X}^k(t) := (X_1^k(t) + \dots + X_n^k(t)) / n$, and W_i^k are independent standard Wiener processes. The detailed derivation is given in Appendix A.

Now averaging over i for each k in equation (2), we obtain the following set of two ordinary differential equations (ODEs) governing the per-class means that $\bar{X}^k(t)$ for $k = 1, 2$:

$$d\bar{X}^k(t) = \left(\frac{\alpha}{2} \bar{X}^k(t) + \frac{\beta}{2} \bar{X}^{3-k}(t) \right) dt + \sigma d \frac{W_i^k(t) + \dots + W_i^k(t)}{n}.$$

Taking the limit of $n \rightarrow \infty$, we observe that $\sigma d \frac{W^1(t) + \dots + W^n(t)}{n} \Rightarrow 0$. Thus, the above display converges weakly to the following ODE:

$$\frac{d\bar{X}^k(t)}{dt} = \frac{\alpha}{2} \bar{X}^k(t) + \frac{\beta}{2} \bar{X}^{3-k}(t). \quad (3)$$

With the initial conditions $\mathbb{E}X_i^k(0) = c_k$, the solution to the above ODE is

$$\bar{X}^1(t) = \frac{c_1 - c_2}{2} e^{\frac{\alpha - \beta}{2}t} + \frac{c_1 + c_2}{2} e^{\frac{\alpha + \beta}{2}t}, \quad \bar{X}^2(t) = -\frac{c_1 - c_2}{2} e^{\frac{\alpha - \beta}{2}t} + \frac{c_1 + c_2}{2} e^{\frac{\alpha + \beta}{2}t}.$$

In the finite-sample setting, we may replace \bar{X}^1, \bar{X}^2 in the SDE (2) by their deterministic solutions and obtain

$$\begin{cases} X_i^1(t) &= \frac{c_1 - c_2}{2} e^{\frac{\alpha - \beta}{2}t} + \frac{c_1 + c_2}{2} e^{\frac{\alpha + \beta}{2}t} - c_1 + X_i^1(0) + \sigma W_i^1(t), \\ X_j^2(t) &= -\frac{c_1 - c_2}{2} e^{\frac{\alpha - \beta}{2}t} + \frac{c_1 + c_2}{2} e^{\frac{\alpha + \beta}{2}t} - c_2 + X_j^2(0) + \sigma W_j^2(t). \end{cases}$$

We are now ready to derive the condition under which these $2n$ feature vectors become *asymptotically separable*, that is, $\min_i X_i^1(t) > \max_j X_j^2(t)$ or $\max_i X_i^1(t) < \min_j X_j^2(t)$ as $t \rightarrow \infty$.

Theorem 2.1 (Separation in Binary Classification). *Given the feature vectors $X_i^1(t), X_j^2(t)$ for $i, j \in [n]$, as $t \rightarrow \infty$ and large n ,*

1. *if $\alpha > \beta$, they are asymptotically separable with probability tending to one,*
2. *if $\alpha \leq \beta$, they are asymptotically separable with probability tending to zero.*

This result indicates a sharp phase transition when α is *just* above β , that is, in the regime of *local elasticity*. As long as the intra-class effect is slightly greater than the inter-class effect, separation is guaranteed. This simple model already captures local elasticity and reveals the important role it plays in the *perfect separation* of training samples. We can generalize this model to more realistic settings: when there are multiple classes, when features are high-dimensional, and when the LE matrix \mathbf{E} is time-dependent. In the next section, we discuss each of these three generalizations in more depth.

3 General LE-SDE Model

Now, we consider the general case where $K \geq 2$ and the feature vectors are p -dimensional with $p \geq K$. Inquisitive readers may have already noticed that Theorem 2.1 only asserts the *emergence* of the separation of features, while being inconclusive to their relative orders at separation, that is, *which class converges to where?* This drawback is intrinsic to the toy model as neither intra-class nor inter-class effect identifies different classes. In this section, we introduce the general LE-SDE model that alleviates this difficulty with the help of an extra block matrix \mathbf{H} with the (i, j) -th block $\mathbf{H}_{i,j}$ models how features in the j -th class affect those in the i -th class, which also partially defines how

classes are separated in higher dimensions. In the local elasticity formalism, $\mathbf{H}_{i,j}$ can be viewed as inducing a metric on the feature space under which local elasticity manifests.

As hinted before, in the case of multiple-class features in higher dimensions, we want to guarantee a stronger separation: to know which class converges to where, thus incorporating supervision from label information. For example, when the features are logits (outputs of the neural nets) and the model is trained under the softmax cross-entropy loss, previous work suggests they separate according to specific geometric structures [38]. To this end, we need to adjust the raw feature vectors \mathbf{X}_i^k with a proper transformation that incorporates the label information into the dynamics. This motivates the following modification of the dynamics (1) by adding an extra transformation $\mathbf{H}_{k,L_m} \in \mathbb{R}^{p \times p}$ to the features. For $k \in [K]$, $i \in [n]$, and at iteration m , we have the following:

$$\mathbf{X}_i^k(m) = \mathbf{X}_i^k(m-1) + h \cdot E_{k,L_m}(m) \mathbf{H}_{k,L_m}(m) \mathbf{X}_{J_m}^{L_m}(m-1) + \sqrt{h} \zeta_i^k(m-1). \quad (4)$$

The $\mathbf{H}_{k,L_m}(m)$ term models the LE effect as *proportional* to a linear “transformation” of the features. The dynamics in Equation (1) are special cases when $\mathbf{H}_{k,l}(m) \equiv \mathbf{I}_p$ for all $k, l \in [K]$. By specifying a proper \mathbf{H} , we can overcome the limitation in our toy example of not knowing which class converges to where. We specify interesting choices of \mathbf{H} in Section 3.2.

A further step of abstraction is to write $\widetilde{\mathbf{X}}^k(m)$ instead of $\mathbf{X}_i^k(m)$, to indicate one generic sample from the distribution $\mathcal{D}^k(m)$ of all the features of class k at iteration m . As in Section 2.2, we can derive the continuous dynamics of equation (4) in the limit of $h \rightarrow 0$ in the same way as equation (2). Similar to writing $\widetilde{\mathbf{X}} = (\widetilde{\mathbf{X}}^k)_{k=1}^K \in \mathbb{R}^{Kp}$ for the concatenation of per-class features, $\bar{\mathbf{X}} = (\bar{\mathbf{X}}^k)_{k=1}^K \in \mathbb{R}^{Kp}$ is the concatenation of per-class mean features. Our model (4) approximates the following SDE with identification $t = mh$ as $h \rightarrow 0$. We term this model the LE-SDE:

$$d\widetilde{\mathbf{X}}_t = \mathbf{M}_t \bar{\mathbf{X}}_t dt + \Sigma_t^{\frac{1}{2}} d\mathbf{W}_t, \quad (5)$$

where \mathbf{W}_t is the standard Wiener process in \mathbb{R}^{Kp} , Σ_t is the covariance matrix, and $\mathbf{M}_t \in \mathbb{R}^{Kp \times Kp}$ is a $K \times K$ block matrix, with each block of size $p \times p$. The (k, l) th block of \mathbf{M}_t is $E_{k,l}(t) \mathbf{H}_{k,l}(t)/K$ when $l \neq k$, and $E_{l,l}(t) \mathbf{H}_{l,l}(t)/K$ when $l = k$. The rationale for dividing K is that we assume that the data are balanced; therefore, each of the K possible classes has an equal chance of being sampled, as proved in Equation (2), where $K = 2$. In Appendix E, we discuss how we can generalize this to model SGD with mini-batches, imbalanced data, and label corruptions.

Taking expectation with respect to the randomness arising from sampling $\widetilde{\mathbf{X}}_t$ from its distribution, the per-class mean $\bar{\mathbf{X}}_t$ satisfies the following system, which we term the LE-ODE:

$$\bar{\mathbf{X}}_t' = \mathbf{M}_t \bar{\mathbf{X}}_t. \quad (6)$$

Under the assumptions in Section 2.1, we define $\gamma(t) = \min \{\alpha(t) - \beta(t), \alpha(t) + (K-1)\beta(t)\}$, $A(t) = \int_0^t \alpha(\tau) d\tau$, $B(t) = \int_0^t \beta(\tau) d\tau$, and $\Gamma(t) = \min \{A(t) - B(t), A(t) + (K-1)B(t)\}$.

3.1 The Separation Theorem

Similar to the discussions in Theorem 2.1, the LE-SDE allows us to derive the separability result for a general K and $p \geq K$. We say the feature vectors $\left\{ (\mathbf{X}_i^k)_{i \in [n]} \right\}_{k \in [K]}$ are *separable* if for any two classes $k \neq l$, there exists a hyperplane in \mathbb{R}^p that linearly separates the features of the two classes. To characterize the separation as in Theorem 2.1, we need conditions on $\alpha(t), \beta(t)$ as therein. Intuitively, when $\gamma(t) \leq 0$, the classes cannot be separated, even in a pairwise manner. Therefore, we focus on a more interesting case when $\gamma(t) > 0$. We now state the following characterization theorem of separability for general LE-SDE dynamics:

Theorem 3.1 (Separation of LE-SDE). *Under our working assumptions in Section 2.1, and in the case of local elasticity (i.e., $\gamma(t) > 0$), assume $\mathbf{H} = (\mathbf{H}_{ij})_{ij}$ is positive semi-definite (PSD) with positive diagonal entries. As $t \rightarrow \infty$, we have¹:*

1. if $\gamma(t) = \omega(1/t)$, the features are separable with probability tending to 1;

¹Here, $\gamma(t) = \omega(1/t)$ stands for $\gamma(t) \gg 1/t$ as $t \rightarrow \infty$. For example, $1/t^{0.5} = \omega(1/t)$ and $(t \ln t)^{-1} = o(1/t)$ as $t \rightarrow \infty$.

2. if $\gamma(t) = o(1/t)$, and the number of per-class-feature n tending to ∞ at an arbitrarily slow rate, the features are asymptotically pairwise separable with probability 0.

This theorem sheds light on the crucial impact of the local elasticity effect for separation in a general case. The proof to Theorem 3.1 as well as discussions on the empirically best ways of choosing the universal direction ν (i.e., a direction that does not depend on the class index) are detailed in Appendix C.2.

3.2 Two Specific Models

We next discuss two specific choices of the \mathbf{H} matrix that allows us to analyze $\widetilde{\mathbf{X}}$ precisely.

3.2.1 Isotropic Feature Learning Model

As a straightforward extension to Section 2.2, we can simply choose $\mathbf{H}_{lk} = \mathbf{I}_p$ to be the identity matrix. This choice of \mathbf{H} is PSD, and thus, we can apply Theorem 3.1 to obtain the conditions for asymptotic separation. In this case, the solution $\widetilde{\mathbf{X}}(t)$ to the LE-ODE can be computed analytically as given in the following proposition.

Proposition 3.2 (l-model). *Let $\mathbf{H}_{k,l} = \mathbf{I}_p$, then the solution to the LE-ODE (6) is given by*

$$\widetilde{\mathbf{X}}(t) = \mathbf{c} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + (\mathbf{1}_K \otimes \mathbf{c}_0) e^{\frac{1}{K}A(t) + \frac{K-1}{K}B(t)}, \quad (7)$$

where $\mathbf{c} = (\mathbf{c}_k)_{k=1}^K \in \mathbb{R}^{Kp}$ and $\mathbf{c}_0 \in \mathbb{R}^p$ are constants with $\sum_{k=1}^K \mathbf{c}_k = \mathbf{0} \in \mathbb{R}^p$ and $\widetilde{\mathbf{X}}(0) = \mathbf{c} + \mathbf{c}_0$.

The derivation of Equation (7) is deferred to Appendix C.2.1. From equation (7), we can easily reconstruct Theorem 3.1 in this special case. The difference between a feature vector from class k and that from class l at time t is given by $(\mathbf{c}_k - \mathbf{c}_l) e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + \Sigma^{1/2}(\mathbf{W}_t^k - \mathbf{W}_t^l)$, provided that the first deterministic term dominates the random second term, thus ensuring separation, which are precisely the conditions specified in Theorem 3.1. We term this model as the *isotropic feature model*, or l-model for short; as the \mathbf{H} matrix has identity matrices as its blocks and consequently the dynamics do not prescribe any preferred directions for each class.

3.2.2 Logits-as-Features Model

An important type of features in neural nets is the *logits*, the outputs of the neural net before the softmax layer. A logit vector (or logits) is K -dimensional, and in this model we identify $\widetilde{\mathbf{X}}^k(t)$ as the logits at time t of a generic sample from the k -th class. In a well-trained neural net, the logits of a learned data instance from the k -th class should have its k -th logit being the largest, and heuristically, the other coordinates should be approximately equal and negative. As we shall detail in Appendix B, the exact dynamics of neural net training pushes the logits $\widetilde{\mathbf{X}}^k$ by its *margin*, $\mathbf{d}_k := \mathbf{e}_k - \text{softmax}(\widetilde{\mathbf{X}}^k)$, which roughly aligns with the direction of $\mathbf{e}_k - \mathbf{1}_p/K$. This suggests us how to choose the metric under which local elasticity acts: we can choose $\mathbf{H}_{l,k}$ such that it always aligns $\widetilde{\mathbf{X}}^k$ in the direction of \mathbf{d}_k , that is,

$$\mathbf{H}_{ij} = \bar{\mathbf{H}}^j := \frac{\mathbf{d}_j \mathbf{d}_j^\top}{\|\mathbf{d}_j\|_2^2} \in \mathbb{R}^{p \times p}, \quad \mathbf{d}_j := \mathbf{e}_j - \frac{1}{K} \mathbf{1}_p \in \mathbb{R}^p, \quad j \in [K]. \quad (8)$$

Roughly speaking, the map $\mathbf{x} \mapsto \bar{\mathbf{H}}^j \mathbf{x}$ projects \mathbf{x} in the direction of \mathbf{d}_j and ideally aligns \mathbf{x} with \mathbf{d}_j after iterative applications; hence, $\bar{\mathbf{H}}^j$ can be viewed as an approximation of the nonlinear transformation in the exact dynamics in the sense that the direction of their stationary point coincides. Furthermore, $\bar{\mathbf{H}}^j$ thus defined has operator norm 1; thus, it does not affect the magnitudes, but only directions. Note that \mathbf{H} does not satisfy the condition in Theorem 3.1 as it is not symmetric; yet the separation theorem can be easily extended in light of the following proposition.

Proposition 3.3 (L-model). *Let \mathbf{H} be the same as in equation (8), then the solution to the LE-ODE (6) is given by*

$$\widetilde{\mathbf{X}}(t) = \mathbf{c}_0 + C_1 \mathbf{d} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + \left(\sum_{l=1}^{K-1} C_{2l} \mathbf{f}_l \right) e^{\frac{1}{K}A(t) + \frac{1}{K(K-1)}B(t)}, \quad (9)$$

where \mathbf{f}_l 's are fixed vectors in \mathbb{R}^{K^2} , $\mathbf{c}_0 \in \mathbb{R}^{K^2}$ is a constant vector with $K(K-1)$ degrees of freedom, and $C_1, C_{2l} \in \mathbb{R}$, for $l \in [K-1]$ are constants.



Figure 2: Samples from GEOMNIST dataset.

The specific form of f_l is not the focus here; equation (9) allows us to prove the statement of Theorem 3.1 under this choice of \mathbf{H} . The proof of Proposition 3.3 is deferred to Appendix C.2.2, where we also provide the analytical solution of f_l 's when $K = 3$.

We term this model as the logits-as-features model, or L-model for short, because it is an elaborate model specifically for logits. In Section 4, we provide concrete demonstrations of our abstract feature vector $\tilde{\mathbf{X}}$ as logits under L-model. We numerically simulated our LE-ODE and compared its predicted dynamics with real deep learning training dynamics. The experimental results provide strong empirical support for the validity of L-model.

3.3 Connection with Neural Collapse

Neural collapse is a recent phenomenological finding on the geometry of the logits learned by deep neural nets at convergence with the cross-entropy loss [38] (see an explanation of neural collapse in [16]). Simply speaking, taking our L-model as an example, with balanced training samples, this model asserts that the logit vectors from different classes at convergence form an equiangular tight frame (ETF). ETFs are the best configuration to spread K unit vectors in an ambient space of p dimensions. Formally, we say a set of vectors $\{\mathbf{s}_i\}_{i=1}^K$ form an ETF in \mathbb{R}^p if they are the columns of a matrix

$$\mathbf{S} = \sqrt{\frac{K}{K-1}} \mathbf{Q} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top \right), \quad (10)$$

where $\mathbf{Q} \in \mathbb{R}^{p \times K}$, and $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_K$. As a direct corollary of Proposition 3.3, when \mathbf{H} is set according to equation (8), we find that our L-model also predicts the existence of neural collapse from the local elasticity point of view.

Proposition 3.4 (Neural Collapse of the LE-ODE). *Under L-model and the same setup as in Theorem 3.1, if $\gamma(t) > 0$ and there exists some $T > 0$ such that $B(t) < 0$ for $t \geq T$, then $\left\{ \bar{\mathbf{X}}^k(t) / \|\bar{\mathbf{X}}^k(t)\| \right\}_{k=1}^K$ forms an ETF as $t \rightarrow \infty$.*

4 Experiments

We perform various experiments to test our theory, where we choose *logits* as our protagonist².

4.1 Setup

Datasets and Models. We perform experiments on a synthesized dataset called GEOMNIST containing $K = 3$ types of geometric shapes (RECTANGLE, ELLIPSOID, and TRIANGLE) and on CIFAR-10 ([28], denoted by CIFAR) with $K \in [2, 3]$ classes. A few samples from GEOMNIST are shown in Figure 2. We vary the number of training samples per class and label pollution ratio p_{err} and use variants of the ALEXNET ([29]) model. More details can be found in the Appendix.

Training Configurations. All models are trained for $T = 10^5$ iterations (for GEOMNIST) or $T = 3 \times 10^5$ iterations (for CIFAR) with a learning rate of 0.005 and a batch size of 1 under the softmax cross-entropy loss. Models on GEOMNIST converged with training and validation losses to zero, and those on CIFAR to validation accuracies greater than 90%.

Estimation Procedures. Each experiment is repeated for $n_{\text{trial}} = 100$ independent runs to estimate $\bar{\mathbf{X}}(t)$. We use both the isotropic feature learning model (Section 3.2.1) and the logits-as-features model (Section 3.2.2), denoted by L-model and l-model respectively, to estimate $\alpha(t)$ and $\beta(t)$. The

²Code for reproducing our experiments is publicly available at github.com:zjiayao/le_sde.git.

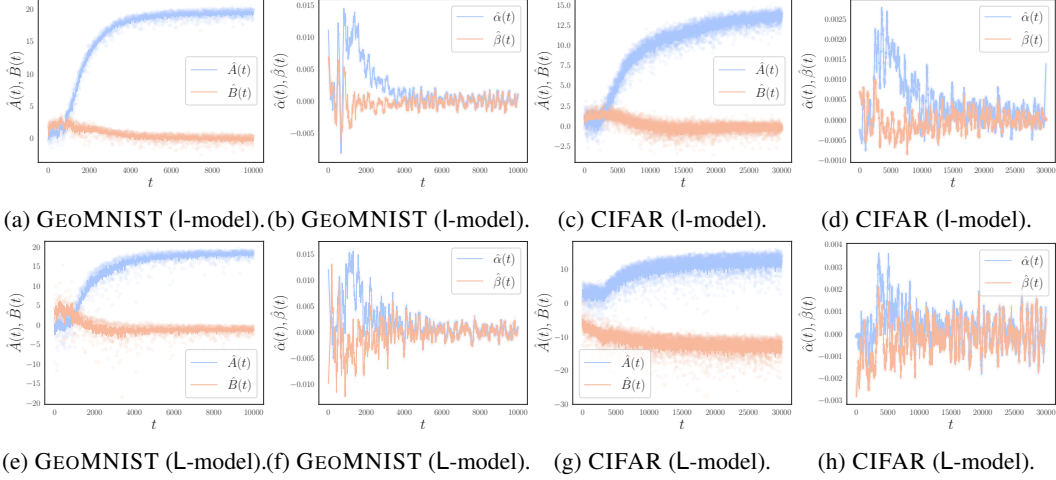


Figure 3: **Estimated $\hat{A}(t)$, $\hat{B}(t)$, $\hat{\alpha}(t)$, and $\hat{\beta}(t)$.** The first row was estimated using l-model and the second L-model; the first two columns are on GEOMNIST and the last two on CIFAR. The first and third rows show $\hat{A}(t)$ and $\hat{B}(t)$ and the other two rows $\hat{\alpha}(t)$ and $\hat{\beta}(t)$.

L-model is used only when $K = 3$. To estimate $\alpha(t)$ and $\beta(t)$, we first estimate $A(t)$ and $B(t)$ by

$$\begin{aligned}
\text{(l-model)} \quad & \begin{cases} \hat{A}(t) &= \text{avg avg}_k \log \left| \frac{\bar{\mathbf{X}}(\bar{\mathbf{X}}^k - \bar{\mathbf{X}}^{k-1})^{K-1}}{\mathbf{c}_0 \mathbf{c}_k^{K-1}} \right|, & \bar{\mathbf{X}}_t &:= \text{avg}_l \bar{\mathbf{X}}_t^l, \\ \hat{B}(t) &= -\text{avg avg}_k \log \left| \frac{\mathbf{c}_0 \bar{\mathbf{X}}^k - \bar{\mathbf{X}}}{\mathbf{c}_k \bar{\mathbf{X}}} \right|, \end{cases} \\
\text{(L-model)} \quad & \begin{cases} \hat{A}(t) &= A'(t) + 2B'(t), & \begin{cases} A'(t) &:= \log \left\langle \bar{\mathbf{X}}^\top \mathbf{v}_1 - 1 \right\rangle, \\ B'(t) &:= \log \left\langle \bar{\mathbf{X}}^\top (\mathbf{v}_2 - \frac{4}{3}\mathbf{v}_1) \right\rangle, \end{cases} \end{cases}
\end{aligned} \tag{11}$$

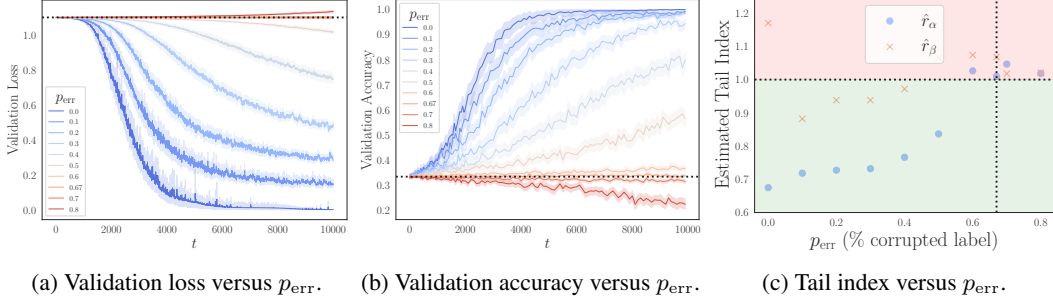
where vector division is interpreted entry-wise. We write avg_l for averaging over the class index, avg for averaging over the coordinates, and define $\langle \mathbf{X} \rangle(t) := \mathbf{X}(t)/\mathbf{X}(0)$. We explain how and why to choose the vectors \mathbf{v}_1 and \mathbf{v}_2 in Appendix D.1. The main idea is to view the eigenvectors of the Kp -by- Kp drift matrix as a concatenation of K vectors of dimension p and construct their linear combinations such that one or more independent components in the solution vanishes. With $A(t)$ and $B(t)$ estimated, we use the Savitzky - Golay filter to obtain $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ through numerical differentiation. For GEOMNIST and CIFAR datasets, we choose window sizes of this filter as 191 and 551, respectively, in Figure 3, and 21 and 21, respectively, in Figure 5.

We assess the tail of $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ by a *tail index* defined as $r_\alpha := \sup_s \{s : \lim_{t \rightarrow \infty} \alpha(t) \cdot t^s < \infty\}$ and r_β is defined similarly. We estimate \hat{r}_α by fixing an interval $[T_1, T_2]$ with $T_1 < T_2$ sufficiently large such that we may ignore terms with smaller order and have $\hat{r}_\alpha = 1 - \text{avg}_{T_1 \leq t \leq T_2} \frac{\log \alpha(t)}{\log(1+t)}$, and similarly for \hat{r}_β . We use the estimates from the last 1000 iterations for averaging in our experiments.

4.2 Results

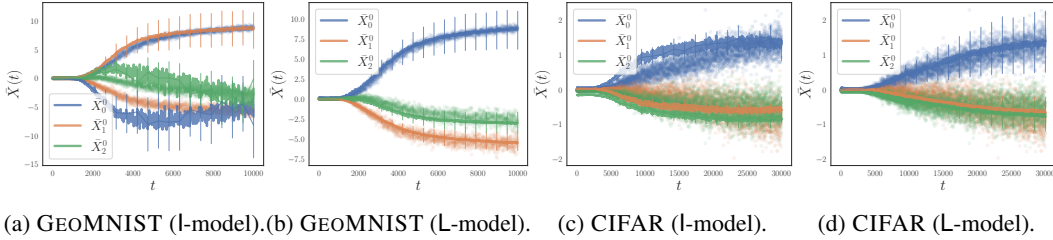
Local Elasticity in Neural Net Training. Local elasticity manifests from our model as the heaviness of the tail of $\gamma(t) = \alpha(t) - \beta(t)$, and in Figure 3, we plot the estimations $\hat{A}(t)$, $\hat{B}(t)$, $\hat{\alpha}(t)$, and $\hat{\beta}(t)$ using both l-model and L-model. We note that (i) The estimations from the two models are visually similar, especially in the late stage of training when t is large; (ii) The major difference lies in the initial stage, where the estimates from L-model behave slightly wilder. This is not surprising because of the effect of the unknown constant offset \mathbf{c}_0 in the L-model; (iii) Both $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ behave similarly on both datasets.

Phase Transition of Separability. Theorem 3.1 states that separation of features under the LE-SDE takes place when $\gamma(t) = \alpha(t) - \beta(t) = \omega(1/t)$, or roughly speaking, when $r_\gamma = \min\{r_\alpha, r_\beta\} < 1$. Although we cannot directly control $\alpha(t)$ and $\beta(t)$, we can bias them by tuning the label corruption ratio p_{err} . When $p_{\text{err}} \approx p_{\text{err}}^* := 2/3$, we are in effect assigning labels completely at random and thus we expect a phase transition of separability should happen around p_{err}^* . This is indeed the story depicted in Figure 4: Figures 4a and 4b show that the validation loss and accuracy for $p \geq p_{\text{err}}$



(a) Validation loss versus p_{err} . (b) Validation accuracy versus p_{err} . (c) Tail index versus p_{err} .

Figure 4: **Phase transition of separability.** (a)—(b) Validation loss and accuracy suggest separation fails for $p_{\text{err}} \geq p_{\text{err}}^* = 2/3$. The dashed line in (a) carries the value at initialization and overlaps with the case where $p_{\text{err}} = 0.6$; the dashed line in (b) is $p_{\text{err}}^* = 2/3$, when labels are assigned completely at random. (c) Tail indices of $\alpha(t)$ and $\beta(t)$. Note that $\gamma(t) = \alpha(t) - \beta(t)$ crosses the horizontal line $r = 1$, entering the non-separable regime (shaded in red) from the separable regime (shaded in green), around the same p_{err} that cross the dashed lines in (a) and (b), as predicted by Theorem 3.1.



(a) GEOMNIST (l-model). (b) GEOMNIST (L-model). (c) CIFAR (l-model). (d) CIFAR (L-model).

Figure 5: **Simulated LE-ODE solutions versus genuine dynamics.** We use $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ estimated from l-model ((a) and (c)) or L-model, ((b) and (d)) and numerically simulate the solution under the L-model. The results were overlaid with true dynamics from neural nets. We note L-model in general imitated true dynamics reasonably well.

are not increasing over time and in Figure 4c we observe the minimum tail index of $\alpha(t)$ and $\beta(t)$ crosses 1 from below around $p_{\text{err}} = 0.6$, entering the non-separable regime (shaded in red) from the separable regime (shaded in green), given in Theorem 3.1.

Simulating DNN Dynamics via LE-ODE. Having estimated $\alpha(t)$ and $\beta(t)$, it is natural to ask, to what capacity can our LE-ODE models recover the real dynamics of deep neural nets? We use the forward Euler method to simulate the L-model using $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ estimated from either l-model or L-model. We choose $K = 3$ and show in Figure 5 the simulated solution (solid line) with error bars depicting one standard deviation over 500 independent runs, overlaying on the real dynamics from DNNs in the background (shaded transparent markers). As the moving average may reduce the magnitudes of $\alpha(t)$ and $\beta(t)$, we rescale the simulated paths such that its first coordinate is approximately equal to the ground truth at convergence. Note that estimations from L-model can faithfully recover the genuine dynamics from neural nets, whereas those from l-model fail, notably in Figure 5a, where the simulated paths preserve the relative magnitude but fail to identify the correct order of three logits.

5 Discussion and Future Works

In this study, we introduce LE-SDE/ODE models that draw inspiration from the local elasticity phenomenon. Conditions for sharp phase transition of separability of features are derived. We also show that once the elasticity strengths $\alpha(t)$ and $\beta(t)$ are well estimated, our model can faithfully simulate the dynamics of neural nets. We outline a few interesting problems for future research while leaving the details in the Appendix. (i) **General LE Matrix.** A similar result as in Theorem 3.1 may be expected for symmetric but no necessarily semi-definite LE matrices $E(t)$. (ii) **Mini-batch Training, Imbalanced Datasets, and Label Corruptions.** Generalizing the drift matrix to $M_t = (E_t \otimes P) \circ H/K$ for a K -by- K doubly stochastic matrix P can be used to model various sampling effects. (iii) **Beyond L-model for Imitating Genuine Dynamics of DNNs.** Although the L-model is shown to be able to mimic the real dynamics reasonably well, we postulate that a more precise model might have its (i, j) -th block encode the other directions other than d_j .

Acknowledgements

This work was supported in part by NSF through CCF-1934876, an Alfred Sloan Research Fellowship, the Wharton Dean’s Research Fund, and ONR Contract N00014-19-1-2620. We would like to thank Dan Roth and the Cognitive Computation Group at the University of Pennsylvania for stimulating discussions and for providing computational resources.

References

- [1] Z. Allen-Zhu, Y. Li, and Z. Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252, 2019.
- [2] S. Arora, N. Cohen, N. Golowich, and W. Hu. A convergence analysis of gradient descent for deep linear neural networks. In *International Conference on Learning Representations*, 2019.
- [3] S. Arora, N. Cohen, and E. Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *International Conference on Machine Learning*, pages 244–253, 2018.
- [4] S. Arora, S. Du, W. Hu, Z. Li, and R. Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332, 2019.
- [5] P. Chaudhari and S. Soatto. Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–10. IEEE, 2018.
- [6] S. Chen, H. He, and W. Su. Label-aware neural tangent kernel: Toward better generalization and local elasticity. In *Advances in Neural Information Processing Systems*, volume 33, pages 15847–15858, 2020.
- [7] Z. Chen, Y. Cao, Q. Gu, and T. Zhang. A generalized neural tangent kernel analysis for two-layer neural networks. *Advances in Neural Information Processing Systems*, 33, 2020.
- [8] J. Cohen, S. Kaur, Y. Li, J. Z. Kolter, and A. Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations*, 2021.
- [9] A. Damianou and N. D. Lawrence. Deep Gaussian processes. In *Artificial intelligence and statistics*, pages 207–215, 2013.
- [10] Z. Deng, H. He, and W. Su. Toward better generalization bounds with locally elastic stability. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 2590–2600, 2021.
- [11] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186. Association for Computational Linguistics, 2019.
- [12] S. Du and W. Hu. Width provably matters in optimization for deep linear neural networks. In *International Conference on Machine Learning*, pages 1655–1664, 2019.
- [13] S. Du, J. Lee, H. Li, L. Wang, and X. Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685, 2019.
- [14] S. S. Du, K. Hou, R. R. Salakhutdinov, B. Póczos, R. Wang, and K. Xu. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [15] S. S. Du, X. Zhai, B. Póczos, and A. Singh. Gradient descent provably optimizes over-parameterized neural networks. In *7th International Conference on Learning Representations*. OpenReview.net, 2019.
- [16] C. Fang, H. He, Q. Long, and W. J. Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 2021.
- [17] S. Goldt, M. Advani, A. M. Saxe, F. Krzakala, and L. Zdeborová. Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [18] T. Hazan and T. Jaakkola. Steps toward deep kernel methods from infinite neural networks. *arXiv preprint arXiv:1508.05133*, 2015.

- [19] H. He and W. Su. The local elasticity of neural networks. In *International Conference on Learning Representations*, 2020.
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [21] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, 2015.
- [22] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [23] S. Jastrzebski, Z. Kenton, D. Arpit, N. Ballas, A. Fischer, Y. Bengio, and A. Storkey. Three factors influencing minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017.
- [24] Z. Ji and M. Telgarsky. Polylogarithmic width suffices for gradient descent to achieve arbitrarily small test error with shallow relu networks. In *International Conference on Learning Representations*, 2020.
- [25] K. Kawaguchi. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [26] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017.
- [27] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015.
- [28] A. Krizhevsky. Learning multiple layers of features from tiny images, 2009.
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25:1097–1105, 2012.
- [30] A. K. Lampinen and S. Ganguli. An analytic theory of generalization dynamics and transfer learning in deep linear networks. In *International Conference on Learning Representations*, 2019.
- [31] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [32] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.
- [33] Q. Li, C. Tai, and E. Weinan. Stochastic modified equations and adaptive stochastic gradient algorithms. In *International Conference on Machine Learning*, pages 2101–2110, 2017.
- [34] Y. Li and Y. Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [35] Z. Li, S. Malladi, and S. Arora. On the validity of modeling sgd with stochastic differential equations (sdes). *arXiv preprint arXiv:2102.12470*, 2021.
- [36] S. Mandt, M. D. Hoffman, and D. M. Blei. Stochastic gradient descent as approximate bayesian inference. *The Journal of Machine Learning Research*, 18(1):4873–4907, 2017.
- [37] A. W. Marshall, I. Olkin, and B. C. Arnold. *Inequalities: Theory of Majorization and its Applications*, volume 143. Springer, second edition, 2011.
- [38] V. Pappas, X. Y. Han, and D. 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.
- [39] D. Saad and S. A. Solla. Dynamics of on-line gradient descent learning for multilayer neural networks. *Advances in Neural Information Processing Systems*, pages 302–308, 1996.
- [40] A. M. Saxe, J. L. McClelland, and S. Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *International Conference on Learning Representations*, 2014.
- [41] A. M. Saxe, J. L. McClelland, and S. Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116(23):11537–11546, 2019.

- [42] B. Shi, W. J. Su, and M. I. Jordan. On learning rates and Schrödinger operators. *arXiv preprint arXiv:2004.06977*, 2020.
- [43] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [44] S. L. Smith, P.-J. Kindermans, and Q. V. Le. Don’t decay the learning rate, increase the batch size. In *International Conference on Learning Representations*, 2018.
- [45] D. Soudry, E. Hoffer, M. S. Nacson, S. Gunasekar, and N. Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- [46] S. Särkkä and A. Solin. *Applied Stochastic Differential Equations*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2019.
- [47] G. Yang. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *arXiv preprint arXiv:1902.04760*, 2019.
- [48] T. Young, D. Hazarika, S. Poria, and E. Cambria. Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, 13(3):55–75, 2018.
- [49] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.

A Derivation of Continuous Dynamics of Binary Case

We will now derive continuous dynamics (2) in the main paper. Let $\mathbb{1}_m = 1$ if class 1 is selected at iteration m and $\mathbb{1}_m = 0$ otherwise. Chaining the dynamic (2) r times, we have

$$\begin{aligned} & X_i^1(m-1+r) - X_i^1(m-1) \\ &= \sum_{q=1}^r \left(\mathbb{1}_{m+q-1} \alpha h X_{I_{m+q-1}}^1(m+q-2) + (1 - \mathbb{1}_{m+q-1}) \beta h X_{I_{m+q-1}}^2(m+q-2) + \zeta_{m+q-2}^i \right). \end{aligned}$$

When $m \gg r$, we have approximately

$$\begin{aligned} & \sum_{q=1}^r \left(\mathbb{1}_{m+q-1} \alpha h X_{I_{m+q-1}}^1(m+q-2) + (1 - \mathbb{1}_{m+q-1}) \beta h X_{I_{m+q-1}}^2(m+q-2) \right) \\ & \approx \sum_{q=1}^r \left(\mathbb{1}_{m+q-1} \alpha h X_{I_{m+q-1}}^1(m-1) + (1 - \mathbb{1}_{m+q-1}) \beta h X_{I_{m+q-1}}^2(m-1) \right) \\ &= \sum_{q=1}^r \mathbb{1}_{m+q-1} \alpha h X_{I_{m+q-1}}^1(m-1) + \sum_{q=1}^r (1 - \mathbb{1}_{m+q-1}) \beta h X_{I_{m+q-1}}^2(m-1) \\ & \approx r \cdot \frac{1}{2} \alpha h \frac{\sum_{i=1}^n X_i^1(m-1)}{n} + r \cdot \frac{1}{2} \beta h \frac{\sum_{j=1}^n X_j^2(m-1)}{n} \\ & \approx \frac{\alpha h r}{2} \bar{X}(m-1) + \frac{\beta h r}{2} \bar{Y}(m-1). \end{aligned}$$

Next, observe that

$$\sum_{q=1}^r \zeta_{m+q-2}^i \sim \mathcal{N}(0, \sigma^2 r h),$$

hence taken together, the calculations above give

$$X_i^1(m-1+r) - X_i^1(m-1) \approx \frac{\alpha h r}{2} \bar{X}(m-1) + \frac{\beta h r}{2} \bar{Y}(m-1) + \mathcal{N}(0, \sigma^2 r h).$$

Writing $\Delta t = r h$ and $t = (m-1)h$, we have

$$X_i^1(t + \Delta t) - X_i^1(t) \approx \frac{\alpha}{2} \bar{X}(t) \Delta t + \frac{\beta}{2} \bar{Y}(t) \Delta t + \sigma \mathcal{N}(0, \Delta t),$$

which is the discretization of

$$dX_i^1(t) = \left(\frac{\alpha}{2} \bar{X}(t) + \frac{\beta}{2} \bar{Y}(t) \right) dt + \sigma dW^i(t).$$

Likewise, we can obtain the dynamics of X_j^2 similarly. We will next prove the separation theorem in binary classification, Theorem 2.1.

Theorem 2.1 (Separation in Binary Classification). *Given the feature vectors $X_i^1(t)$, $X_j^2(t)$ for $i, j \in [n]$, as $t \rightarrow \infty$ and large n ,*

1. if $\alpha > \beta$, they are asymptotically separable with probability tending to one,
2. if $\alpha \leq \beta$, they are asymptotically separable with probability tending to zero.

Proof of Theorem 2.1. Note that whenever $\alpha \leq \beta$, we have $\frac{c_1 - c_2}{2} e^{\frac{\alpha - \beta}{2} t} \rightarrow 0$ as $t \rightarrow \infty$, thus X_i^1 and X_j^2 are interspersed and separation happens with probability tending to zero. This also aligns with our intuition that the intra-class effect should be stronger than its inter-class counterpart.

On the other hand, when $\alpha > \beta$, ignoring a null set we may assume $c_1 > c_2$ without loss of generality. To see this, note that by definition $c_k = \mathbb{E}_{\text{data}}[X^k(0) | \theta(0) = \theta_0]$ for $k \in \{1, 2\}$ where $\theta(0)$ is all parameters of the neural net at initialization and θ_0 is a particular realization given the initialization scheme. Here the expectation is taken with respect to the data distribution, and when we ignore a null

set of neural net with respect to the probability measure induced by the initialization scheme, $c_1 \neq c_2$ holds. In other words, this statement can be interpreted as “the expected feature at initialization from the first class is different from that from the second class for a neural net, except possibly on a null set in the space of neural nets with respect to the probability measure induced by the parameter initialization scheme.” It suffices to show that

$$\begin{aligned} \min_{1 \leq i \leq n} \frac{c_1 - c_2}{2} e^{\frac{\alpha - \beta}{2} t} + \frac{c_1 + c_2}{2} e^{\frac{\alpha + \beta}{2} t} - c_x + X_i^1(0) + \sigma W_i^1(t) \\ > \max_{1 \leq j \leq n} -\frac{c_1 - c_2}{2} e^{\frac{\alpha - \beta}{2} t} + \frac{c_1 + c_2}{2} e^{\frac{\alpha + \beta}{2} t} - c_2 + X_j^2(0) + \sigma W_j^2(t), \end{aligned}$$

which is equivalent to

$$(c_1 - c_2) e^{\frac{\alpha - \beta}{2} t} - c_1 + \min_{1 \leq i \leq n} X_i^1(0) + \sigma W_i^1(t) > -c_2 + \max_{1 \leq j \leq n} X_j^2(0) + \sigma W_j^2(t).$$

But the above display happens with probability tending to one provided $\alpha > \beta$, thus completing the proof. \square

B Further Details on Drift Modeling

B.1 Dynamics of the Logits-as-Features Model

This section provides more details on why the construction of \mathbf{H} in Section 3.2.2 is probably a good choice for modeling the dynamics of logits in deep neural nets. Given $K \geq 2$ classes with n training examples per class, the feature vectors are the logits $\mathbf{X}^k(m) \in \mathbb{R}^K$ for all $k \in [K]$, where m is the iteration number. When the neural net is trained under the softmax cross-entropy loss L , at the m -th iteration, if the J_m -th sample from the L_m -th class is sampled, the dynamics of the logits \mathbf{X}_i^k should be governed by

$$\mathbf{X}_i^k(m) - \mathbf{X}_i^k(m-1) \approx h \left[\frac{\partial \mathbf{X}_i^k(m-1)}{\partial \mathbf{w}} \frac{\partial \mathbf{X}_{J_m}^{L_m \top}}{\partial \mathbf{w}} \left(\mathbf{e}_{L_m} - \text{softmax}(\mathbf{X}_{J_m}^{L_m}) \right) \right]. \quad (\text{B.1})$$

The derivation of equation (B.1) is a straightforward computation from the Taylor approximation

$$\mathbf{X}_i^k(m) - \mathbf{X}_i^k(m-1) \approx \nabla_{\mathbf{w}} \mathbf{X}_i^k(m-1) \Delta \mathbf{w}(m-1), \quad (\text{B.2})$$

where we observe that

$$\Delta \mathbf{w}(m-1) = -h \frac{\partial L(\mathbf{w}(m-1))}{\partial \mathbf{w}} = h \frac{\partial \mathbf{X}_{J_m}^{L_m \top}}{\partial \mathbf{w}} \left(\mathbf{e}_{L_m} - \text{softmax}(\mathbf{X}_{J_m}^{L_m}) \right). \quad (\text{B.3})$$

However, as equation (B.1) is highly non-linear, J_m and L_m are random, and the Gram matrix $\frac{\partial \mathbf{X}_i^k(m-1)}{\partial \mathbf{w}} \frac{\partial \mathbf{X}_{J_m}^{L_m \top}}{\partial \mathbf{w}}$ is also time-dependent, direct analyses and simulation of the exact dynamics are difficult. Note that this Gram matrix is also the key element in the NTK literature [22, 15, 2], which is treated as roughly fixed during the lazy training process.

Recall that we consider the per-class mean of logits, $\bar{\mathbf{X}}^k = \mathbb{E} \widetilde{\mathbf{X}}^k$, where $\widetilde{\mathbf{X}}^k$ is a random sample, for each class k . In the limit of $h \rightarrow 0$ with the identification of time $t = mh$, the above modeling allows us to re-write (B.1) in the per-class mean $\bar{\mathbf{X}}^k$ and logits from a generic sample $\widetilde{\mathbf{X}}^k$ in this continuous limit as

$$\begin{aligned} d\widetilde{\mathbf{X}}_t^k &\approx \mathbb{E}_{L \sim \text{Unif}([K])} \left[\mathbb{E}_{\widetilde{\mathbf{X}} \sim \mathcal{D}_t^L} \left[\frac{\partial \mathbf{X}_i^k(m-1)}{\partial \mathbf{w}} \frac{\partial \widetilde{\mathbf{X}}^{\top}}{\partial \mathbf{w}} \left(\mathbf{e}_L - \text{softmax}(\widetilde{\mathbf{X}}) \right) \right] \right] dt + \Sigma_t^{\frac{1}{2}} d\mathbf{W}_t, \\ &\approx \frac{1}{K} \sum_L \left(\mathbb{E}_{\widetilde{\mathbf{X}}' \sim \mathcal{D}_t^k, \widetilde{\mathbf{X}} \sim \mathcal{D}_t^L} \left[\frac{\partial \widetilde{\mathbf{X}}'}{\partial \mathbf{w}} \frac{\partial \widetilde{\mathbf{X}}^{\top}}{\partial \mathbf{w}} \right] \left(\mathbf{e}_L - \text{softmax}(\widetilde{\mathbf{X}}_t^L) \right) \right) dt + \Sigma_t^{\frac{1}{2}} d\mathbf{W}_t, \\ &= \frac{1}{K} \sum_L \left(\Theta_{k,L} \left(\mathbf{e}_L - \text{softmax}(\widetilde{\mathbf{X}}_t^L) \right) \right) dt + \Sigma_t^{\frac{1}{2}} d\mathbf{W}_t. \end{aligned} \quad (\text{B.4})$$

The presence of expectation over non-linearity posed considerable difficulties of using equation (B.4) to analyze neural nets; however, it motivates our LE-SDE model as a linear approximation to it while

explicitly encodes the local elasticity into the dynamics. More precisely, given a sample instance z^k from the k -th class, in a well-trained neural network, the probability vector associated with z^k , i.e., the output from the softmax applied to its logits, should have its k -th entry as the largest and the other entries as roughly equal. Hence $e_k - \text{softmax}(\bar{X}_t^k)$ should be roughly in the direction of $e_k - \frac{1}{K}\mathbf{1}$. To approximate this limit by the limit from a linear map, we may use our choice of $\mathbf{H}_{kl} = \bar{\mathbf{H}}^l := \frac{d_l d_l^T}{\|d_l\|^2}$ for all $k, l \in [K]$, such that in the limit of $t \rightarrow \infty$, we expect

$$e_k - \text{softmax}(\bar{X}_t^k) \approx C \bar{\mathbf{H}}^k \bar{X}_t^k, \quad (\text{B.5})$$

for any $k \in [K]$ and some positive constant C . This is the intuition behind our L-model. To summarize, in our model, the local elasticity matrix E describes the effect $\Theta_{k,L}$, and $\mathbf{H}_{k,l}$ transforms the mean logit \bar{X}_t^L to the direction governed by the supervision.

B.2 Discussions on Linearization

We provide more details of the rationale behind our choice of $M(t)X(t)$, a seemingly linear term, as the surrogate for the non-linear drift in the dynamics given in equation B.4. Our following argument can be extended to any post-activation features. Writing equation B.4 in terms of $X(t) \in \mathbb{R}^{Kp}$ (recall in this case we have $p = K$), the concatenation of K per-class feature vectors $X^k(t) \in \mathbb{R}^p$ for $k \in [K]$, and denoting by $\sigma : \mathbb{R}^K \rightarrow \mathbb{R}^K$ the softmax function for simplicity, we can express the drift term as

$$F(\bar{\mathbf{X}}(t), t) := \Theta(t) \left(\left[e_k - \sigma(\bar{\mathbf{X}}^k(t)) \right]_{k=1}^K \right), \quad (\text{B.6})$$

where we wrote $[\cdot]$ for vector concatenation and $\Theta(t) \in \mathbb{R}^{(Kp) \times (Kp)}$ for the Gram matrix. A commonly used linearization scheme in SDE for non-linear drifts by the filtering community (cf. Chapter 9.1 of [46]) is to linearize F for each t at the mean $\varphi(t) = (\varphi_k(t))_{k=1}^K \equiv \bar{X}(t) := \mathbb{E}_B X(t)$ where the expectation is taken with respect to the diffusion. Concretely, we have

$$F(\bar{\mathbf{X}}(t), t) \approx \tilde{F}(\bar{\mathbf{X}}(t), t) := F(\varphi(t), t) + \nabla_X F(\varphi(t), t) (\bar{\mathbf{X}}(t) - \varphi(t)), \quad (\text{B.7})$$

where $\nabla_X F$ denotes the Jacobian of F with respect to the spacial variable $\bar{\mathbf{X}}$. For notation completeness, we introduce

$$p = (p_k)_{k=1}^K \in \mathbb{R}^{Kp}, \quad p_k := \sigma(\bar{\mathbf{X}}^k(t)) \in \mathbb{R}^p, \quad k \in [K], \quad (\text{B.8})$$

and similarly

$$\bar{p} = (\bar{p}_k)_{k=1}^K \in \mathbb{R}^{Kp}, \quad \bar{p}_k := \sigma(\bar{\mathbf{X}}^k(t)) \in \mathbb{R}^p, \quad k \in [K]. \quad (\text{B.9})$$

We write the per-class Jacobians as

$$J_{kk} = J_k := \text{diag}(\bar{p}_k) - \bar{p}_k \bar{p}_k^T. \quad (\text{B.10})$$

Clearly, the Jacobian $\nabla F(\varphi, t) = J(t)$ can be written as a block-diagonal matrix $J(t) = (J_{kk})_{k=1}^K$ consisting of per-class Jacobians. Now continuing linearization, we can write

$$\begin{aligned} \tilde{F}(\bar{\mathbf{X}}(t), t) &= \Theta(t) \left([e_k - \bar{p}_k]_k + J(t) (\bar{\mathbf{X}}(t) - \varphi(t)) \right) \\ &= \Theta(t) (J(t)X(t) + [e_k - \bar{p}_k + J_k \varphi_k(t)]_k). \end{aligned} \quad (\text{B.11})$$

Define $\Psi : \mathbb{R}^{Kp} \rightarrow \mathbb{R}^{Kp} : z \mapsto [e_k - \sigma(z_k)]_k$ and write $\Psi_k : \mathbb{R}^p \rightarrow \mathbb{R}^p$ to be the k -th component of Ψ , using Taylor's theorem to expand $\Psi(z)$ around $\varphi(t)$ for each t , we have

$$\Psi = \Psi(\varphi) + J(t)\varphi - J(\varphi)z + o(\|z - \varphi\|), \quad (\text{B.12})$$

or

$$\Psi(\varphi) + J(t)\varphi = \Psi(z) + J(\varphi)z + o(\|z - \varphi\|). \quad (\text{B.13})$$

This implies that

$$\tilde{F} = \Theta(t)J(t)\bar{\mathbf{X}}(t) + \Theta(t)R(t), \quad R(t; z) := \Psi(z) + J(t)z + o(\|z - \varphi(t)\|), \quad (\text{B.14})$$

where $R(t; z)$ is the residue that depends on the choice of z around which $\Psi(\varphi)$ is expanded. Note that the first term is a time-varying linear term in $X(t)$. As long as the residue term is negligible, we get exactly the time-varying linear map $M(t)$ in the LE-SDE model.

By choosing different z 's in equation B.14, we can focus on different stages of the real dynamics using the LE-SDE. Two particular choices are of great interest so as to make the residue term vanishing:

- **Around initialization.** Let $z = u := c \cdot [\mathbf{1}_K/K]_{k=1}^K$ be a scaling of vectors of ones where c is some fixed constant. Then each of the K components of $\sigma(u)$ assigns approximately the same probability ($1/K$) for every label. Furthermore, $u \in \text{Ker } J(t)$ for all t hence the residue $R(t; u) = \Psi(u) + o(\|z - \varphi(t)\|)$ is a constant vector (which is colinear with $\mathbf{d} = (\mathbf{d}_j)_{j=1}^K$ defined by $\mathbf{d}_j = \mathbf{e}_j - \mathbf{1}_K/K$ when we discussed the L-model). Note that this approximation works best when the model has not learned much about the data (in the “first stage” as we call it) since $\|u - \varphi(t)\|$ is small in this regime.
- **Around convergence.** Given that the model converges, $\varphi_\infty := \varphi(\infty)$ is finite. Let $z = \varphi_\infty$, under the effective training assumption, $\|\Psi(\varphi_\infty)\| \approx 0$ by construction. Hence the residue $R(t; \varphi_\infty) = J(t)\varphi_\infty + o(\|\varphi(t) - \varphi_\infty\|)$. Here the $o(\cdot)$ term converges to 0 as training progresses, leaving us a term that is asymptotically equivalent to $v = (v_k)_{k=1}^K := J(\varphi_\infty)\varphi_\infty \in \mathbb{R}^{K^2}$, where $v_k = [(z_{k,i} - \sum_{j=1}^K p_{k,j} z_{k,j}) p_i]_{i=1}^K \in \mathbb{R}^K$. Again, under the effective training assumption z_k has its k -th entry $z_{k,k}$ the largest, and p_k has its k -th entry close to 1 while the others to zero. Thus $v_{k,i} \approx 0_K$. We see that in this regime, the approximation $\Theta(t)J(t)X(t)$ is only off by a residue $o(\|\varphi(t) - \varphi_\infty\|)$ that eventually vanishes.

As discussed above, we choose to linearize the drift at convergence instead of around initialization in the L-model given by equation B.5 such that the residue vanishes under the effective training assumption.

C Miscellaneous Proofs

C.1 Various Definitions of Separability

We first give formal definitions of separability which generalize Theorem 2.1 in Section 2.2 to our K -class, p -dimensional feature setting. In the beginning of Section 3.1, we state the definition of separability in natural language, and we formalize the definition therein as follows. It is the most natural definition in terms of linear separation by a hyperplane.

Definition C.1 (Pairwise Separation). *We say the feature vectors $\{(\mathbf{X}_i^k)_{i \in [n]}\}_{k \in [K]}$ is pairwise separable at time t if for each pair of $1 \leq k < l \leq K$, there exists a direction $\nu_{k,l}$ such that*

$$\min_i \langle \nu_{k,l}, \mathbf{X}_i^k(t) \rangle > \max_j \langle \nu_{k,l}, \mathbf{X}_j^l(t) \rangle. \quad (\text{C.1})$$

In the Theorem 3.1, we claim that when $\gamma(t) = \omega(1/t)$, the pairwise separation (Definition C.1) happens with probability tending to 1 as $t \rightarrow \infty$. This is a notion of asymptotic separability stated in Theorem 3.1, yet it is a weaker notion of separation in the following sense: the hyperplane may depend on the classes l, k , and the hyperplane may depend on time t . We state the following asymptotic separable definitions of increasingly stronger guarantees, and will remark on how to obtain them in our proof of Theorem 3.1.

The direct application of Definition C.1 gives us the following (weakest) definition.

Definition C.2 (Asymptotic Pairwise Separation). *We say the feature vectors $\{(\mathbf{X}_i^k)_{i \in [n]}\}_{k \in [K]}$ are asymptotically pairwise separable if for each pair of $1 \leq k < l \leq K$, there exist directions $\nu_{k,l}(t)$ such that*

$$\mathbb{P} \left(\min_i \langle \nu_{k,l}(t), \mathbf{X}_i^k(t) \rangle > \max_j \langle \nu_{k,l}(t), \mathbf{X}_j^l(t) \rangle \right) \rightarrow 1. \quad (\text{C.2})$$

Requiring all the classes to be separable with the same hyperplane give us the following universal separation.

Definition C.3 (Asymptotic Universal Separation). *We say the feature vectors $\{(\mathbf{X}_i^k)_{i \in [n]}\}_{k \in [K]}$ are asymptotically universally separable if it is asymptotically pairwise separable, and there exists ν , such that either $\nu = \nu_{k,l}$ or $\nu = -\nu_{k,l}$, for all $1 \leq k < l \leq K$ in (C.2).*

Note that we allow the universal direction differs in sign for different pair of classes. This is because we do not require separating the k class in any specific order.

We note that the above definitions are in the sense of separation ‘‘in probability’’, which only asserts separation at an arbitrarily fixed large time t . Specifically, it does not guarantee the existence of a fixed direction that can always separate a pair of classes for all sufficiently large t . We now state the almost sure definition of asymptotic pairwise separation, which guarantees the same fixed direction separates a pair of classes for all large enough t . And Theorem 3.1 also holds for such definition.

Definition C.4 (Uniform Asymptotic Pairwise Separation). *We say the feature vectors $\left\{(\mathbf{X}_i^k)_{i \in [n]}\right\}_{k \in [K]}$ are uniformly asymptotically pairwise separable if for each pair of $1 \leq k < l \leq K$, there exists a direction $\boldsymbol{\nu}_{k,l}$ such that*

$$\liminf_{T \rightarrow \infty} \min_{k,l} \mathbb{P} \left(\min_i \langle \boldsymbol{\nu}_{k,l}, \mathbf{X}_i^k(t) \rangle > \max_j \langle \boldsymbol{\nu}_{k,l}, \mathbf{X}_j^l(t) \rangle, \text{ for all } t > T \right) = 1. \quad (\text{C.3})$$

We define uniform asymptotic universal separation the same way as its non-uniform counterpart.

Definition C.5 (Uniform Asymptotic Universal Separation). *We say the feature vectors $\left\{(\mathbf{X}_i^k)_{i \in [n]}\right\}_{k \in [K]}$ are uniformly asymptotically universally separable if it is uniformly asymptotically pairwise separable, and there exists $\boldsymbol{\nu}$, such that either $\boldsymbol{\nu} = \boldsymbol{\nu}_{k,l}$ or $\boldsymbol{\nu} = -\boldsymbol{\nu}_{k,l}$, for all $1 \leq k < l \leq K$ in (C.3).*

The asymptotic inseparability definitions can be similarly stated.

C.2 Proofs in Section 3

Before proving the main result, Theorem 3.1, it is convenient to first prove Propositions 3.2 and 3.3.

C.2.1 Proof of Proposition 3.2

Proposition 3.2 (l-model). *Let $\mathbf{H}_{k,l} = \mathbf{I}_p$, then the solution to the LE-ODE (6) is given by*

$$\bar{\mathbf{X}}(t) = \mathbf{c} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + (\mathbf{1}_K \otimes \mathbf{c}_0) e^{\frac{1}{K}A(t) + \frac{K-1}{K}B(t)}, \quad (7)$$

where $\mathbf{c} = (\mathbf{c}_k)_{k=1}^K \in \mathbb{R}^{Kp}$ and $\mathbf{c}_0 \in \mathbb{R}^p$ are constants with $\sum_{k=1}^K \mathbf{c}_k = \mathbf{0} \in \mathbb{R}^p$ and $\bar{\mathbf{X}}(0) = \mathbf{c} + \mathbf{c}_0$.

Proof of Proposition 3.2. When $\alpha(t) = \alpha$ and $\beta(t) = \beta$ are constants, \mathbf{E} has two distinct eigenvalues: $\lambda_0 = (\alpha - \beta)/K$ with multiplicity $K - 1$ and $\lambda_1 = (\alpha + (K - 1)\beta)/K$ with multiplicity one, and the eigendecomposition $\mathbf{E} = \mathbf{Q}\boldsymbol{\Lambda}\mathbf{Q}^{-1}$ is given by

$$\boldsymbol{\Lambda} = \text{diag}(\lambda_1, \lambda_0, \dots, \lambda_0) \in \mathbb{R}^{K \times K}, \quad \mathbf{Q} = \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & -1 & 0 & \dots & 0 \\ 1 & 0 & -1 & \dots & 0 \\ \vdots & & & & \vdots \\ 1 & 0 & 0 & \dots & -1 \end{bmatrix}. \quad (\text{C.4})$$

Hence the eigenpairs of $\mathbf{E} \otimes \mathbf{I}_p$ are $(\mathbf{e}_1 - \mathbf{e}_k) \otimes \mathbf{e}_j$ for $1 < k \leq K$ and $1 \leq j \leq p$ with eigenvalue λ_0 and $(\mathbf{1}_K \otimes \mathbf{e}_j)$ for $1 \leq j \leq p$ with eigenvalue λ_1 . Reparameterizing the solution such that the initial values are $\bar{\mathbf{X}}^k(0) = \mathbf{c}_0 + \mathbf{c}_k$ for all $k \in [K]$, we have

$$\bar{\mathbf{X}}_t = \mathbf{c} e^{\lambda_0 t} + (\mathbf{1}_K \otimes \mathbf{c}_0) e^{\lambda_1 t}, \quad (\text{C.5})$$

for $\mathbf{c}_0 \in \mathbb{R}^p$ and $\mathbf{c} = (\mathbf{c}_k)_{k=1}^K$ with $\sum_{k=1}^K \mathbf{c}_k = \mathbf{0}$. Based on this solution, it is not difficult to show that the general solution with varying $\alpha(t)$ and $\beta(t)$ is given by

$$\bar{\mathbf{X}}_t = \mathbf{c} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + (\mathbf{1}_K \otimes \mathbf{c}_0) e^{\frac{1}{K}A(t) + \frac{K-1}{K}B(t)}, \quad (\text{C.6})$$

where $\sum_{k=1}^K \mathbf{c}_k = \mathbf{0}$. □

C.2.2 Proof of Proposition 3.3

Proposition 3.3 (L-model). *Let \mathbf{H} be the same as in equation (8), then the solution to the LE-ODE (6) is given by*

$$\bar{\mathbf{X}}(t) = \mathbf{c}_0 + C_1 \mathbf{d} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + \left(\sum_{l=1}^{K-1} C_{2l} \mathbf{f}_l \right) e^{\frac{1}{K}A(t) + \frac{1}{K(K-1)}B(t)}, \quad (9)$$

where \mathbf{f}_l 's are fixed vectors in \mathbb{R}^{K^2} , $\mathbf{c}_0 \in \mathbb{R}^{K^2}$ is a constant vector with $K(K-1)$ degrees of freedom, and $C_1, C_{2l} \in \mathbb{R}$, for $l \in [K-1]$ are constants.

Proof of Proposition 3.3. Recall that the (k, l) -th block of the \mathbf{H} matrix is $\bar{\mathbf{H}}^l := \mathbf{d}_l \mathbf{d}_l^\top / \mathbf{d}_l^\top \mathbf{d}_l$. When $\alpha(t) = \alpha$ and $\beta(t) = \beta$ are constants, we immediately find the matrix $(\mathbf{E} \otimes \mathbf{I}_K) \circ \mathbf{H}$ has an eigenvalue of $\alpha - \beta$ with multiplicity one whose eigenvector is \mathbf{d} , and an eigenvalue of 0 with multiplicity $K(K-1)$ – since the null spaces of \mathbf{H}_{kl} all have dimension $K-1$. We also find it has another eigenvalue $\alpha + \beta/(K-1)$ with multiplicity $K-1$. Therefore, we have the general solution (9) by noting there is an additional $1/K$ factor in the definition of \mathbf{M}_t .

Specifically, when $K=3$, we can explicitly write out the eigenvectors of $(\mathbf{E} \otimes \mathbf{I}_K) \circ \mathbf{H}$ as follows. There is one eigenvector \mathbf{d} corresponding to the eigenvalue $\alpha - \beta$. The eigenvectors corresponding to the eigenvalue 0 take the form of

$$[w_1 + w_2, 2w_1, 2w_2, 2w_3 - w_4, w_3, w_4, -w_5 - 2, w_5, w_6]^\top, \quad (\text{C.7})$$

where $\{w_i\}_{i=1}^6$ are free parameters; the eigenvectors corresponding to the eigenvalue $(\alpha + \beta/(K+1))/K$ take the form of

$$[-\xi_1 w_1 + \xi_2, \xi_2 w_1 - \xi_5, \xi_3 w_1 - \xi_4, -\xi_2 w_1 + \xi_4, \xi_1 w_1 - \xi_3, -\xi_3 w_1 + \xi_6, 1 - w_1, w_1, w_2]^\top, \quad (\text{C.8})$$

where w_1 and w_2 are free parameters and

$$\begin{aligned} \xi_1 &= \frac{2\alpha + \beta}{3\beta}, & \xi_2 &= \frac{\alpha + 2\beta}{3\beta}, & \xi_3 &= \frac{\alpha - \beta}{3\beta}, \\ \xi_4 &= \frac{\alpha^2 + \alpha\beta + 7\beta^2}{6\alpha\beta + 3\beta^2}, & \xi_5 &= \frac{\alpha^2 + 4\alpha\beta - 5\beta^2}{6\alpha\beta + 3\beta^2}, & \xi_6 &= \frac{\alpha^2 - 2\alpha\beta - 8\beta^2}{6\alpha\beta + 3\beta^2}. \end{aligned} \quad (\text{C.9})$$

□

C.2.3 Proof of Theorem 3.1

We are now ready to prove Theorem 3.1. We first recall the theorem statement.

Theorem 3.1 (Separation of LE-SDE). *Under our working assumptions in Section 2.1, and in the case of local elasticity (i.e., $\gamma(t) > 0$), assume $\mathbf{H} = (\mathbf{H}_{ij})_{ij}$ is positive semi-definite (PSD) with positive diagonal entries. As $t \rightarrow \infty$, we have³:*

1. if $\gamma(t) = \omega(1/t)$, the features are separable with probability tending to 1;
2. if $\gamma(t) = o(1/t)$, and the number of per-class-feature n tending to ∞ at an arbitrarily slow rate, the features are asymptotically pairwise separable with probability 0.

Recall that in the main text, we only consider the most natural (pairwise) separation Definition C.1, which corresponds to Definition C.2. After the prove of this theorem, We will also remark how to extend this result to other stronger definitions, especially Definition C.5.

Proof of Theorem 3.1. Let $\mathbf{H} \in \mathbb{R}^{Kp \times Kp}$ be a symmetric positive semi-definite (SPD) matrix with positive diagonal entries $(d_i)_{i=1}^{Kp}$, recall that we denote by $\mathbf{H}_{ij} \in \mathbb{R}^{p \times p}$ the (i, j) -th block of \mathbf{H} for $i, j \in [K]$. We write $\mathbf{M}_t = \frac{1}{K} (\mathbf{E}_t \otimes \mathbf{I}_K) \circ \mathbf{H}$, where $(\mathbf{A} \circ \mathbf{B})_{ij} = A_{ij} B_{ij}$ is the Hadamard product between two matrices of the same size.

First, assume $\alpha(t) = \alpha$ and $\beta(t) = \beta$ are constants, then from Proposition 3.2, we know $\mathbf{E} \otimes \mathbf{I}_p$ has an eigenvalue $\lambda_0 = (\alpha - \beta)/K$ with multiplicity $p(K-1)$ and $\lambda_1 = (\alpha + (K-1)\beta)/K$ with

³Here, $\gamma(t) = \omega(1/t)$ stands for $\gamma(t) \gg 1/t$ as $t \rightarrow \infty$. For example, $1/t^{0.5} = \omega(1/t)$ and $(t \ln t)^{-1} = o(1/t)$ as $t \rightarrow \infty$.

multiplicity p . Writing λ_{\min} and λ_{\max} as the minimum and the maximum of $\{\lambda_0, \lambda_1\}$ respectively, by Schur's theorem (Theorem 9.J.2, [37]), we have

$$\lambda_{\min} \min_{1 \leq i \leq Kp} d_i \leq \lambda_i(\mathbf{M}_t) \leq \lambda_{\max} \max_{1 \leq i \leq Kp} d_i, \quad (\text{C.10})$$

where we write $\mu_i := \lambda_i(\mathbf{M}_t)$ as the i -th largest eigenvalue of \mathbf{M}_t , and $\mathbf{u}_i \in \mathbb{R}^{Kp}$ the corresponding eigenvector and recall that in this case $\mathbf{M}_t = \frac{1}{K}(\mathbf{E} \otimes \mathbf{I}_p) \circ \mathbf{H}$ does not depend on time. We will denote by $\mathbf{u}_i^k \in \mathbb{R}^p$ the k -th block of \mathbf{u}_i for $k \in [K]$, i.e., $(\mathbf{u}_i^k)_j = (\mathbf{u}_i)_{k \times K + j}$.

Note the eigendecomposition of \mathbf{H} is $\sum_{i=1}^{Kp} \mu_i \mathbf{u}_i \mathbf{u}_i^\top$, and thus the solution to $\bar{\mathbf{X}}'_t = \mathbf{M}_t \bar{\mathbf{X}}_t$ is

$$\bar{\mathbf{X}}_t = \bar{\mathbf{X}}_0 + \sum_{i=1}^{Kp} c_i \mathbf{u}_i e^{\mu_i t}, \quad \bar{\mathbf{X}}_0 = \sum_{i=1}^{Kp} c_i \mathbf{u}_i. \quad (\text{C.11})$$

Substituting back this solution to the LE-SDE, we have

$$\begin{aligned} \widetilde{\mathbf{X}}^k(t) &= \widetilde{\mathbf{X}}^k(0) + \mathbf{M}_t \bar{\mathbf{X}}(t) - \mathbb{E}[\widetilde{\mathbf{X}}^k(0)] + \Sigma_k^{\frac{1}{2}}(t) \mathbf{W}^k(t) \\ &= \widetilde{\mathbf{X}}^k(0) + \sum_{i=1}^{Kp} c_i \mu_i \mathbf{u}_i^k e^{\mu_i t} - \sum_{i=1}^{Kp} c_i \mathbf{u}_i^k + \Sigma_k^{\frac{1}{2}} \mathbf{W}^k(t), \end{aligned} \quad (\text{C.12})$$

where $\Sigma_k^{\frac{1}{2}}(t)$ is the covariance for this class. By definition, to prove separation, it suffices to identify a direction $\boldsymbol{\nu}$ such that

$$\langle \widetilde{\mathbf{X}}^k(t) - \widetilde{\mathbf{X}}^l(t), \boldsymbol{\nu} \rangle > 0 \quad (\text{C.13})$$

with probability⁴ tending to 1 as $t \rightarrow \infty$ for any two classes $k \neq l$. Substituting equation (C.12), we have equivalently

$$\left\langle \sum_{i=1}^{Kp} c_i (\mathbf{u}_i^k - \mathbf{u}_i^l) (\mu_i e^{\mu_i t} - 1), \boldsymbol{\nu} \right\rangle > \langle \widetilde{\mathbf{X}}^l(0) - \widetilde{\mathbf{X}}^k(0), \boldsymbol{\nu} \rangle + \langle \Sigma_l^{\frac{1}{2}} \mathbf{W}^l(t) - \Sigma_k^{\frac{1}{2}} \mathbf{W}^k(t), \boldsymbol{\nu} \rangle. \quad (\text{C.14})$$

By the Gaussian tail bound, the right-hand side of the above display is $O_{\mathbb{P}}(C_0 + C_1 \sigma_{\max} \sqrt{t}) = O_{\mathbb{P}}(\sqrt{t})$ where C_0 and C_1 are constants that do not depend on class labels and we assume $\sigma_{\max} = \max_{k \in [K]} \sup_t \|\Sigma_k(t)\|_2$ is finite. Due to randomization in the initialization, with probability zero $\bar{\mathbf{X}}^k(0) = \bar{\mathbf{X}}^l(0)$ and thus

$$\sum_{i=1}^{Kp} c_i (\mathbf{u}_i^k - \mathbf{u}_i^l) \neq \mathbf{0}, \quad (\text{C.15})$$

with probability one. Now equation (C.10) implies that $0 < \lambda_{\min} \min_i d_i \leq \mu_i$, thus when t is sufficiently large, there must exist at least one index j with

$$c_j (\mathbf{u}_j^k - \mathbf{u}_j^l) (\mu_j e^{\mu_j t} - 1) \neq 0 \quad (\text{C.16})$$

as otherwise equation (C.15) is contradicted. Thus separation takes place provided that

$$\exp\{\mu_j t\} > C\sqrt{t} \quad (\text{C.17})$$

holds for some constant C that only depends on C_0, C_1 , and σ_{\max} . The case where $\alpha(t)$ and $\beta(t)$ are time-varying (and so is $\mu_i(t)$) is similar with $\mu_j t$ being replaced by $\int_0^t \mu_j(s) ds$, i.e.,

$$\exp\left\{\int_0^t \mu_j(s) ds\right\} > C\sqrt{t}. \quad (\text{C.18})$$

But equation (C.10) implies that the order of $\mu_j(t)$ is the same as $\gamma(t)$ for all t , hence the above condition is equivalent to

$$\exp\{\Gamma(t)\} > C'\sqrt{t}, \quad (\text{C.19})$$

⁴Here the randomness comes from the dynamics as well as the random sample $\widetilde{\mathbf{X}}^k$.

for some constant C' that only depends on C_0, C_1, σ_{\max} , and the specific choice of \mathbf{H} . Note that to prove equation (C.19), it suffices to require $\gamma(t)$ has a tail that is at least $1/t$, or

$$\gamma(t) = \omega\left(\frac{1}{t}\right) \quad (\text{C.20})$$

as $t \rightarrow \infty$. Since the order of $\mu_1(t) = \max_i \mu_i(t)$ is also the same as $\gamma(t)$ for all t in light of Equation (C.10), whenever

$$\gamma(t) = o\left(\frac{1}{t}\right) \quad (\text{C.21})$$

as $t \rightarrow \infty$, the probability of separation tends to zero as long as $n \rightarrow \infty$ at an arbitrary rate⁵. Note that when $\gamma(t) = \Theta\left(\frac{1}{t}\right)$, the separability depends non-trivially on the constant factors, and the rate of n and t tends to infinity.

Thus we have shown that the order of $\gamma(t)$ characterizes a sharp phase transition in terms of separability. Finally, in the above we proved for each pair of classes k and l , a choice $\nu = \nu(k, l)$ exists that ensures separation. We remark that it is possible to remove class-dependence on ν in our case, and therefore achieve a stronger sense of separability: consider the $(K(K-1)/2)$ -by- K matrix Ψ with $\mathbf{u}_{i(k,l)}^k - \mathbf{u}_{i(k,l)}^l$ as its rows for all $k < l$ where $i(k, l) \in [K]$ is the index of the dominant eigenvalue for the separation between these two classes as discussed above. The existence of a class-independent ν is equivalent to $\Psi\nu \neq \mathbf{0}$. This is equivalent to that the nullity of Ψ is less than K , which is obvious since $\text{rank } \Psi \geq 1$ almost surely by construction. \square

Remark. In the proof of Theorem 3.1, we use a simple consequence of the Gaussian tail bound to derive the critical order for separation of $\gamma(t)$, which is $1/t$. This is correct when we deal with asymptotic pairwise separation (Definition C.2) or asymptotic universal separation (Definition C.3), essentially a law of large numbers result. To obtain the same guarantee for uniform pairwise separation (Definition C.4) or uniform universal separation (Definition C.5), we need to take extra care for the order of $\gamma(t)$. The key observation here is that if we want to guarantee a separation direction that is independent of time t , it is essentially an almost sure statement. To bound the influence of the Wiener process in (C.14), we now need to apply the law of the iterated logarithms. Recall that given a standard Wiener process W_t ,

$$\limsup_{t \rightarrow \infty} \frac{W_t}{\sqrt{2t \log \log t}} = 1, \quad \liminf_{t \rightarrow \infty} \frac{W_t}{\sqrt{2t \log \log t}} = -1, \quad (\text{C.22})$$

hence the conditions in Theorem 3.1 can be generalized as follows:

- if $\gamma(t) = \omega\left(\frac{1}{t \log \log t}\right)$, the features are separable in the sense of uniform asymptotic universal separation;
- if $\gamma(t) = o\left(\frac{1}{t \log \log t}\right)$, the features are *not* separable in the sense of uniform asymptotic universal separation.

Note that the difference between being not uniform asymptotic universally separable and uniform asymptotic universally inseparable, the latter of which happens when $\gamma(t) = o(1/t)$. In practice, the non-uniform version of the theorem is powerful enough, since we do not ask to separate in a pre-specified direction.

C.2.4 Proof of Proposition 3.4.

Proposition 3.4 (Neural Collapse of the LE-ODE). *Under L-model and the same setup as in Theorem 3.1, if $\gamma(t) > 0$ and there exists some $T > 0$ such that $B(t) < 0$ for $t \geq T$, then $\left\{ \bar{\mathbf{X}}^k(t) / \|\bar{\mathbf{X}}^k(t)\| \right\}_{k=1}^K$ forms an ETF as $t \rightarrow \infty$.*

⁵Note that when n is finite, even if two classes of exactly the same mean (i.e. completely intervened) have a non-zero probability of separation: consider $2n$ i.i.d. standard Wiener processes, let n of them be of class 1, and the rest class 2. Then at any time, the probability the two classes are separated is when the largest n belongs to one of the classes, which occurs with probability $\frac{2}{\binom{2n}{n}} > 0$.

Proof of Proposition 3.4. By Proposition 3.3, when $\gamma(t) > 0$ and $B(t) < 0$ eventually, the dominating component of \bar{X}_t^k is $C_1 \mathbf{d}_k e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)}$. As $t \rightarrow \infty$, the unit vector in the direction of \bar{X}_t^k tends to \mathbf{d}_k , and therefore those unit vectors form an ETF, since $\{\mathbf{d}_j\}_{j=1}^K$ form an ETF. \square

C.3 Obtaining the Hyperplane

In this subsection we briefly discuss several methods for obtaining the class-independent \mathbb{R}^1 projection $\boldsymbol{\nu}$ asserted by Theorem 3.1. As suggested by the proof, it is not hard to see that a random Gaussian vector $\boldsymbol{\nu}$ satisfies Theorem 3.1 (with probability one) and it can be used to construct the map $T_{\boldsymbol{\nu}}$ that asymptotically separate all classes almost surely. However, the rate of separation will depend on the choice of $\boldsymbol{\nu}$. In practice, we can find a good $\boldsymbol{\nu}$ by solving the following optimization problem

$$\max_{\|\boldsymbol{\nu}\|_2=1} \min_{k \neq l \in [K]} |\langle \mathbf{c}_k - \mathbf{c}_l, \boldsymbol{\nu} \rangle|.$$

Its solution $\boldsymbol{\nu}^*$ is plausible since informally, it is the direction that can separate all the K classes in the “shortest time” with high probability. To see the intuition behind this claim, consider the worst case of the right hand side of equation (C.14), which is $O_{\mathbb{P}}(\sqrt{t})$ and the least t achieving this for all $k \neq l \in [K]$ requires $\boldsymbol{\nu}$ maximizes $\min_{k \neq l \in [K]} |\langle \mathbf{c}_k - \mathbf{c}_l, \boldsymbol{\nu} \rangle|$.

Remark. *In practice, estimating $\mathbf{c}_0 + \mathbf{c}_l$ from several independent trials incur high variance, despite using a larger number of trials. Although the SDE is not time-reversible, since we are mainly interested in the asymptotic behaviour, we can use an interval $[T_1, T_2]$ with $T_2 > T_1 \gg 0$ when the model is almost convergent to estimate $\boldsymbol{\nu}$ as follows:*

$$\max_{\|\boldsymbol{\nu}\|_2 \leq 1} \min_{k < l} \sum_{t=T_1}^{T_2} |\langle \bar{X}^k(t) - \bar{X}^l(t), \boldsymbol{\nu} \rangle|. \quad (\text{C.23})$$

Note that this problem is convex and easy to solve. In practice, we observe directly setting

$$\boldsymbol{\nu} = \frac{\bar{\mathbf{X}}^k(T) - \bar{\mathbf{X}}^l(T)}{\|\bar{\mathbf{X}}^k(T) - \bar{\mathbf{X}}^l(T)\|} \quad (\text{C.24})$$

for a pair $k \neq l$ and a large T can obtain relatively decent separation in \mathbb{R}^1 compared with equation (C.23), which is the case when we construct Figure 1.

D More Details on Experiments

We first recall the setup of our experiments. We generate a dataset consisting of $K = 3$ simple geometric shapes (RECTANGLE, ELLIPSOID, and TRIANGLE) that are rotated to various angles and applied Gaussian blurring, which we conveniently name Geometric-MNIST or GEOMNIST for short. A few samples from GEOMNIST are shown in Figure 2. We use a varying number of training samples per class $n_{\text{tr}} \in \{80, 480, 600, 4800\}$ with the validation sample per class being $n_{\text{val}} = \{20, 120, 400, 1200\}$. We also pollute each label class by randomly choosing $p_{\text{err}} \cdot n_{\text{tr}}$ samples to flip the label to another class. (uniform across all other classes). In this setup, we fix $n_{\text{tr}} = n_{\text{val}} = 500$ and set $p_{\text{err}} \in \{0.1, 0.2, \dots, 0.8\}$. In addition to GEOMNIST, we use CIFAR-10 ([28], denoted by CIFAR) for a more realistic scenario with 5000 training samples and 1000 validation samples per class. We vary the total number of classes $K \in [2, 3]$. Variants of the ALEXNET model ([29]) are used, which consists of two convolutional layers and three fully-connected layers activated by the ReLU function.

D.1 Estimation Procedures

We will discuss here how we estimate several quantities in (4), including local elasticity strengths $\alpha(t)$ and $\beta(t)$ in the l-model and the L-model, and tail indices r_{α} and r_{β} .

D.1.1 Estimation of Integrated Local Elasticity Strengths $\hat{A}(t)$ and $\hat{B}(t)$

Recall that in Equation (11) we give the following formulae for estimating $A(t)$ and $B(t)$:

$$\begin{aligned} \text{(l-model)} \quad & \begin{cases} \hat{A}(t) &= \text{avg avg}_k \log \left| \frac{\tilde{\mathbf{X}}(\tilde{\mathbf{X}}^k - \tilde{\mathbf{X}})^{K-1}}{c_0 c_k^{K-1}} \right|, \\ \hat{B}(t) &= -\text{avg avg}_k \log \left| \frac{c_0 \tilde{\mathbf{X}}^k - \tilde{\mathbf{X}}}{c_k \tilde{\mathbf{X}}} \right|, \end{cases} \quad \tilde{\mathbf{X}}_t := \text{avg}_l \tilde{\mathbf{X}}_t^l, \\ \text{(L-model)} \quad & \begin{cases} \hat{A}(t) &= A'(t) + 2B'(t), \\ \hat{B}(t) &= 2(B'(t) - A'(t)), \end{cases} \quad \begin{cases} A'(t) &:= \log \left\langle \tilde{\mathbf{X}}^\top \mathbf{v}_1 - 1 \right\rangle, \\ B'(t) &:= \log \left\langle \tilde{\mathbf{X}}^\top \left(\mathbf{v}_2 - \frac{4}{3} \mathbf{v}_1 \right) \right\rangle, \end{cases} \end{aligned} \quad (\text{D.1})$$

where

$$\mathbf{v}_1 = \frac{1}{4} [1, -1, -1, -1, 1, -1, -2, -2, 0]^\top, \quad \mathbf{v}_2 = \frac{1}{3} [2, -1, -1, -1, 2, -1, 0, 0, 0]^\top. \quad (\text{D.2})$$

We will now explain how it is done.

Estimation in l-model. Recall from Proposition 3.2, the per-class means $\tilde{\mathbf{X}}_t$ solve the LE-ODE (7) under the l-model as

$$\tilde{\mathbf{X}}(t) = \mathbf{c} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + (\mathbf{1}_K \otimes \mathbf{c}_0) e^{\frac{1}{K}A(t) + \frac{K-1}{K}B(t)}, \quad (\text{D.3})$$

where $\mathbf{c} = (\mathbf{c}_k)_{k=1}^K \in \mathbb{R}^{Kp}$ and $\mathbf{c}_0 \in \mathbb{R}^p$ are constants with $\sum_{k=1}^K \mathbf{c}_k = \mathbf{0}$ and $\tilde{\mathbf{X}}(0) = \mathbf{c} + \mathbf{c}_0$. The specific structure of this solution implies that

$$\tilde{\mathbf{X}}_t := \text{avg}_l \tilde{\mathbf{X}}_t^l = \frac{1}{K} \sum_{l=1}^K \tilde{\mathbf{X}}_t^l = \mathbf{c}_0 e^{\frac{1}{K}A(t) + \frac{K-1}{K}B(t)} \in \mathbb{R}^K, \quad (\text{D.4})$$

and thus for all $k \in [K]$,

$$\tilde{\mathbf{X}}_t^k - \tilde{\mathbf{X}}_t = \mathbf{c}_k e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)}. \quad (\text{D.5})$$

It follows that

$$\left| \frac{\mathbf{c}_0 \tilde{\mathbf{X}}^k - \tilde{\mathbf{X}}}{c_k \tilde{\mathbf{X}}} \right| = \left| \mathbf{1}_K e^{B(t)} \right|, \quad \left| \frac{\mathbf{c}_0 \tilde{\mathbf{X}}^k - \tilde{\mathbf{X}}}{c_k \tilde{\mathbf{X}}} \right| = \left| \mathbf{1}_K e^{A(t)} \right|, \quad (\text{D.6})$$

for all $k \in [K]$. Taking logarithm and averaging over K classes and K coordinate, we have the estimation equation of l-model in equation (11).

Estimation in L-model. Recall from Proposition 3.3, the per-class means $\tilde{\mathbf{X}}_t$ solve the LE-ODE (9) under L-model as

$$\tilde{\mathbf{X}}(t) = \mathbf{c}_0 + C_1 \mathbf{d} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + \left(\sum_{l=1}^{K-1} C_{2l} \mathbf{f}_l \right) e^{\frac{1}{K}A(t) + \frac{1}{K(K-1)}B(t)}, \quad (\text{D.7})$$

where \mathbf{c}_0 is a constant vector with $K(K-1)$ free parameters; \mathbf{f}_l 's are eigenvectors corresponding to the eigenvalue $(A(t) + B(t)/(K-1))/K$. Although exact solutions can be obtained for general K , they are overly complicated thus we will restrict our attention to the case where $K = 3$. Define

$$\mathbf{v}_1 = \frac{1}{4} [1, -1, -1, -1, 1, -1, -2, -2, 0]^\top, \quad \mathbf{v}_2 = \frac{1}{3} [2, -1, -1, -1, 2, -1, 0, 0, 0]^\top, \quad (\text{D.8})$$

from the proof for Proposition 3.3 (Appendix C.2.2), we immediately have

$$\mathbf{v}_1^\top \mathbf{f}_j = 0, \quad j = 1, 2, \quad \mathbf{v}_2^\top \mathbf{c}_0 = 0, \quad \mathbf{1}_K^\top \mathbf{d} = 0, \quad (\text{D.9})$$

and

$$\mathbf{v}_1^\top \mathbf{c}_0 = 1 - \frac{w_1 + w_2 + w_3}{4}, \quad \mathbf{v}_1^\top \mathbf{d} = 1, \quad \mathbf{v}_2^\top \mathbf{f}_j = 1, \quad \mathbf{v}_2^\top \mathbf{d} = \frac{4}{3}. \quad (\text{D.10})$$

Recall that we define $\langle \mathbf{X}(t) \rangle := \mathbf{X}(t)/\mathbf{X}(0)$, with vector-division interpreted as elementwise division, writing $C_3 = w_1 + w_2 + w_3$ with w_i 's being defined in Appendix C.2.2, we have

$$\begin{cases} \tilde{\mathbf{X}}_t^\top \mathbf{v}_1 &= C_1 e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + 1 - \frac{C_3}{4}, \\ \tilde{\mathbf{X}}_t^\top \mathbf{v}_2 &= \frac{3C_1}{4} e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} + (C_{21} + C_{22}) e^{\frac{1}{K}A(t) + \frac{1}{K(K-1)}B(t)}, \end{cases} \quad (\text{D.11})$$

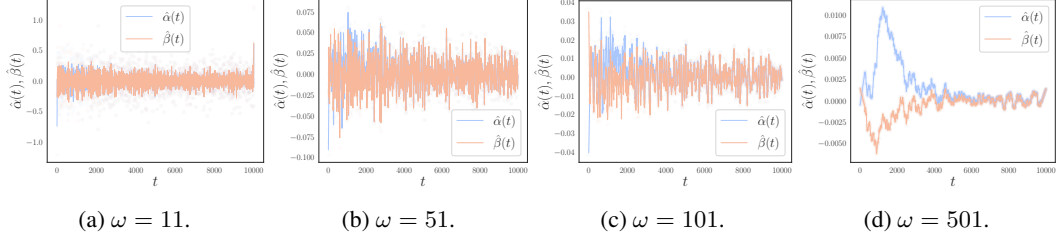


Figure D.1: **Effect of the window size ω in the Savitzky–Golay filter.** These estimations are performed on the GEOMNIST dataset with $p_{\text{err}} = 0$ and $n_{\text{tr}} = 3000$.

and thus

$$\begin{cases} \left\langle \bar{\mathbf{X}}_t^\top \mathbf{v}_1 - 1 \right\rangle &= \frac{C_1 e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} - \frac{C_3}{4}}{C_1 - \frac{C_3}{4}} \\ \left\langle \bar{\mathbf{X}}_t^\top \left(\mathbf{v}_2 - \frac{4}{3}\mathbf{v}_1 \right) \right\rangle &= \frac{(C_{21} + C_{22}) e^{\frac{1}{K}A(t) + \frac{1}{K(K-1)}B(t)} + \frac{3C_3}{16}}{C_{21} + C_{22} + \frac{3C_3}{16}}. \end{cases} \quad (\text{D.12})$$

We define

$$\begin{cases} A'(t) &= \log \left| \left\langle \bar{\mathbf{X}}_t^\top \mathbf{v}_1 - 1 \right\rangle \right|, \\ B'(t) &= \log \left| \left\langle \bar{\mathbf{X}}_t^\top \left(\mathbf{v}_2 - \frac{4}{3}\mathbf{v}_1 \right) \right\rangle \right|, \end{cases} \quad (\text{D.13})$$

when t is sufficiently large such that

$$e^{\frac{1}{K}A(t) - \frac{1}{K}B(t)} \gg \frac{C_3}{4C_1}, \quad e^{\frac{1}{K}A(t) + \frac{1}{K(K-1)}B(t)} \gg \frac{3C_3}{16(C_{21} + C_{22})}, \quad (\text{D.14})$$

we have approximately

$$A'(t) \approx \frac{1}{K}A(t) - \frac{1}{K}B(t), \quad B'(t) \approx \frac{1}{K}A(t) + \frac{1}{K(K-1)}B(t), \quad (\text{D.15})$$

where $K = 3$. Hence $A(t)$ and $B(t)$ can be recovered by

$$\hat{A}(t) = A'(t) + 2B'(t), \quad \hat{B}(t) = 2(B'(t) - A'(t)). \quad (\text{D.16})$$

Although equation (D.16) is only an approximation that is precise only when t is large, we will nonetheless use equation (D.16) for estimation in the L-model for all t .

D.1.2 Estimation of $\alpha(t)$ and $\beta(t)$

Once we have estimates for $A(t)$ and $B(t)$, namely $\hat{A}(t)$ and $\hat{B}(t)$, we may numerically differentiate these estimates to obtain $\hat{\alpha}(t)$ and $\hat{\beta}(t)$. Although the composition of finite difference quotient and moving average yields visibly well results, we shall use the well-established Savitzky–Golay filter for this purpose. There is a window size parameter ω in this filter which roughly corresponds to the window size in moving averages: a smaller ω preserves more fluctuations in the original data and a larger ω smooths the data more. As a rule of thumb, we test on a set of different values for ω and choose one that is both informing and not losing too much detail in our presentations. Specifically, in the main paper, we use $\omega = 191$ for experiments on GEOMNIST and $\omega = 551$ for those on CIFAR. In the case of simulating LE-ODE solutions, we chose $\omega = 21$ to preserve finer details. We show in Figure D.1 estimated $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ trained on GEOMNIST with $p_{\text{err}} = 0$ and $n_{\text{tr}} = 3000$ samples per class under various window sizes $\{11, 51, 101, 301\}$ for the Savitzky–Golay filter. Note that the general trend is not discovered until the window size ω is reasonably large.

D.1.3 Estimation of Tail Index

We are interested in the tail behavior of $\alpha(t)$ and $\beta(t)$. Taking $\alpha(t)$ as an example, suppose

$$\alpha(t) \sim \frac{\alpha_0}{(1+t)^r}, \quad (\text{D.17})$$

for some r , which is the tail index of $\alpha(t)$, we have

$$A(t) = \begin{cases} \frac{\alpha_0(1+t)^{1-r}}{1-r}, & 0 < r < 1, \\ \alpha_0 \log(1+t), & r = 1. \end{cases} \quad (\text{D.18})$$

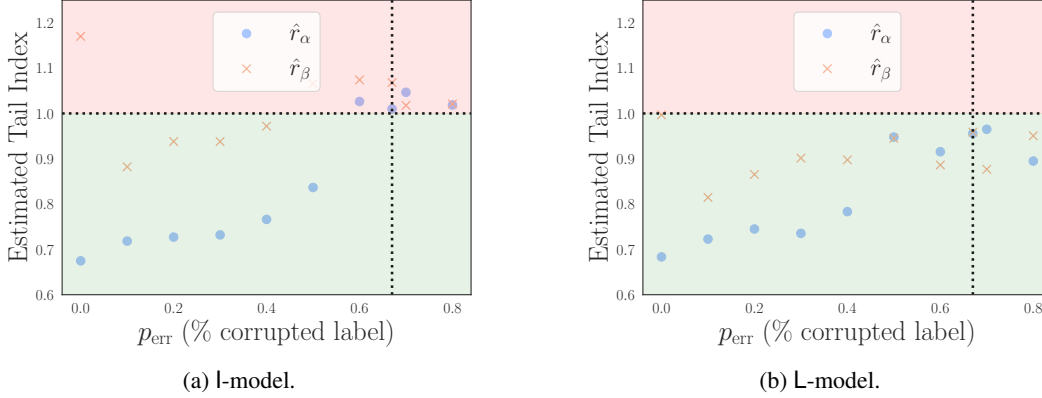


Figure D.2: **Estimated tail indices versus label corruption ratio** p_{err} . The tail indices are estimated using equation (D.19) with $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ estimated via (a) l-model, and (b) L-model. Although the case for the L-model does not exhibit a clear phase transition, we note around $p_{\text{err}} \approx 2/3$, the tail index of $\hat{\beta}(t)$ begins to dominate that of $\hat{\alpha}(t)$.

Hence with t being sufficiently large, we can estimate r using

$$\frac{\log A(t)}{\log(1+r)} = (1-r) + \frac{\log \alpha_0}{\log(1+t)} - \frac{\log(1-r)}{\log(1+t)} \approx 1-r, \quad (\text{D.19})$$

and $\hat{r}_\alpha = 1 - \text{avg}_{t \geq T_0} \log A(t) / \log(1+t)$ for some sufficiently large T_0 . We estimate \hat{r}_β similarly. Although this estimator suffers from large bias when the true model has a constant offset, i.e., when $\alpha(t) \sim \alpha_1 + \alpha_0/(1+t)^r$, we choose it over other estimators based on $\alpha(t) \sim \alpha_1 + \alpha_0/(1+t)^r$ as it is simpler and it is directly based on the integrated local elasticity strength $A(t)$ without the need to perform numerical differentiation beforehand.

D.2 More Results from Experiments

Effects of Training Sample Size n . We show in Figure D.3 estimated $\hat{A}(t)$, $\hat{B}(t)$, $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ versus training time t on both GEOMNIST and CIFAR under both l-model and L-model with various per-class sample size $n = n_{\text{tr}} \in \{80, 480, 600, 4800\}$. We choose the window size $\omega = 151$ when applying the Savitzky–Golay filter. Note that the estimates under both l-model and L-model do not vary significantly under different choices of N and share similar trends: (i) $\beta(t)$ is dominated by $\alpha(t)$; (ii) $\alpha(t)$ has an initial increasing stage and a second stage converging to the vicinity around zero. This is expected since our theory is independent of n once n is reasonably large.

Effects of Number of Classes K . In Figure D.4 we show the estimated $\hat{A}(t)$, $\hat{B}(t)$, $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ versus training time t on CIFAR under the l-model with number of classes $K \in \{2, 3\}$. We note that in both cases, generally speaking, the shapes of $A(t)$ and $B(t)$ are similar: $A(t)$ increases while $B(t)$ decreases, suggesting this behavior is more general and is likely independent of the number of classes. On the other hand, note that the number of classes affects the scale of $A(t)$ and $B(t)$.

Effects of Label Corruption Ratio p_{err} . In Section 3 we have shown that separability depends on the tail behavior of $\alpha(t) - \beta(t)$. In Figure D.5 we demonstrate this more directly. Here we intentionally choose a large window size of $\omega = 351$ to highlight the trends of $\alpha(t)$ and $\beta(t)$. Note that $A(t)$ and $B(t)$ (and consequently $\alpha(t)$ and $\beta(t)$) become more mixed and indistinguishable as we increase the label corruption ratio p_{err} from 0 (no corrupted label) to 0.8. This also reaffirms the important rule played by the local elasticity strengths $\alpha(t)$ and $\beta(t)$ in terms of separability.

We also plot the estimated tail indices versus label corruption ratio p_{err} using $\alpha(t)$ and $\beta(t)$ estimated under both l-model and L-model in Figure D.2. Although in Figure D.2b there is no sharp phase transition boundary when we increase p_{err} as Figure D.2b does, we observe that at around $p_{\text{err}} = 2/3$, the estimated tail index of $\hat{\beta}(t)$ begins to dominate that of $\hat{\alpha}(t)$, which also supports our Theorem 3.1.

Local Elasticity Strengths Adjusted for Logits Norms In our LE-SDE/ODE model l-model and L-model, each block of the H matrix has operator norm 1. Hence it is of interests to inspect the

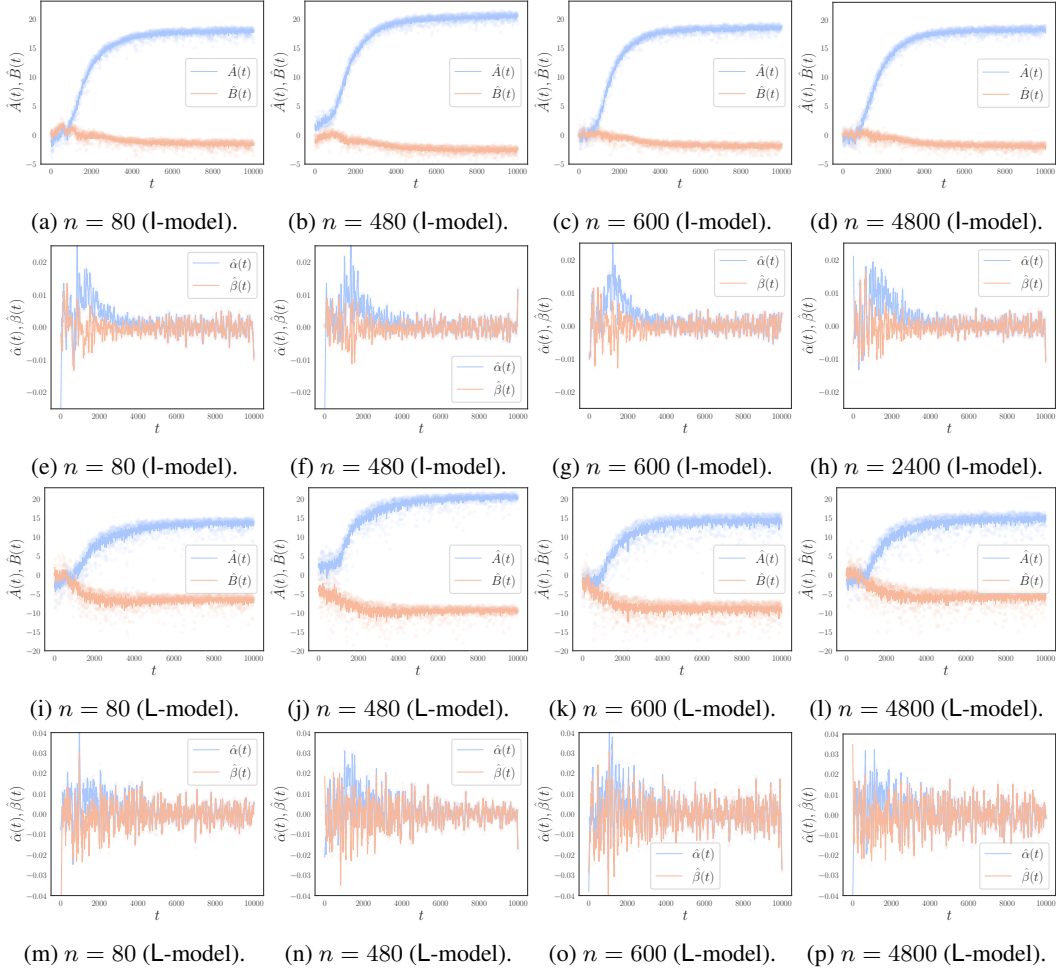


Figure D.3: **Effect of per-class sample size n .** Estimated $\hat{A}(t)$, $\hat{B}(t)$ (the first and the third rows) and $\hat{\alpha}(t)$, $\hat{\beta}(t)$ (the second and the fourth rows), under l-model (the first two rows) and L-model (the last two rows). Note that n appears to be insignificant in determining the shapes and magnitudes of the curves.

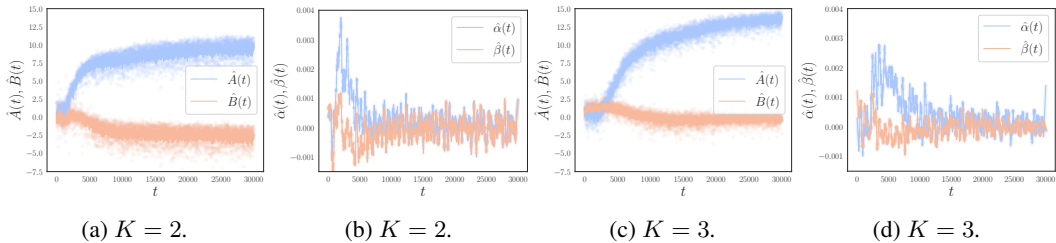


Figure D.4: **Effect of number of classes K .** Estimated $\hat{A}(t)$, $\hat{B}(t)$ ((a) and (c)) and $\hat{\alpha}(t)$, $\hat{\beta}(t)$ ((b) and (d)) on CIFAR with $K = 2$ ((a)-(b)) and $K = 3$ ((c)-(d)), all under l-model. Note that the general trends with different K 's are similar.

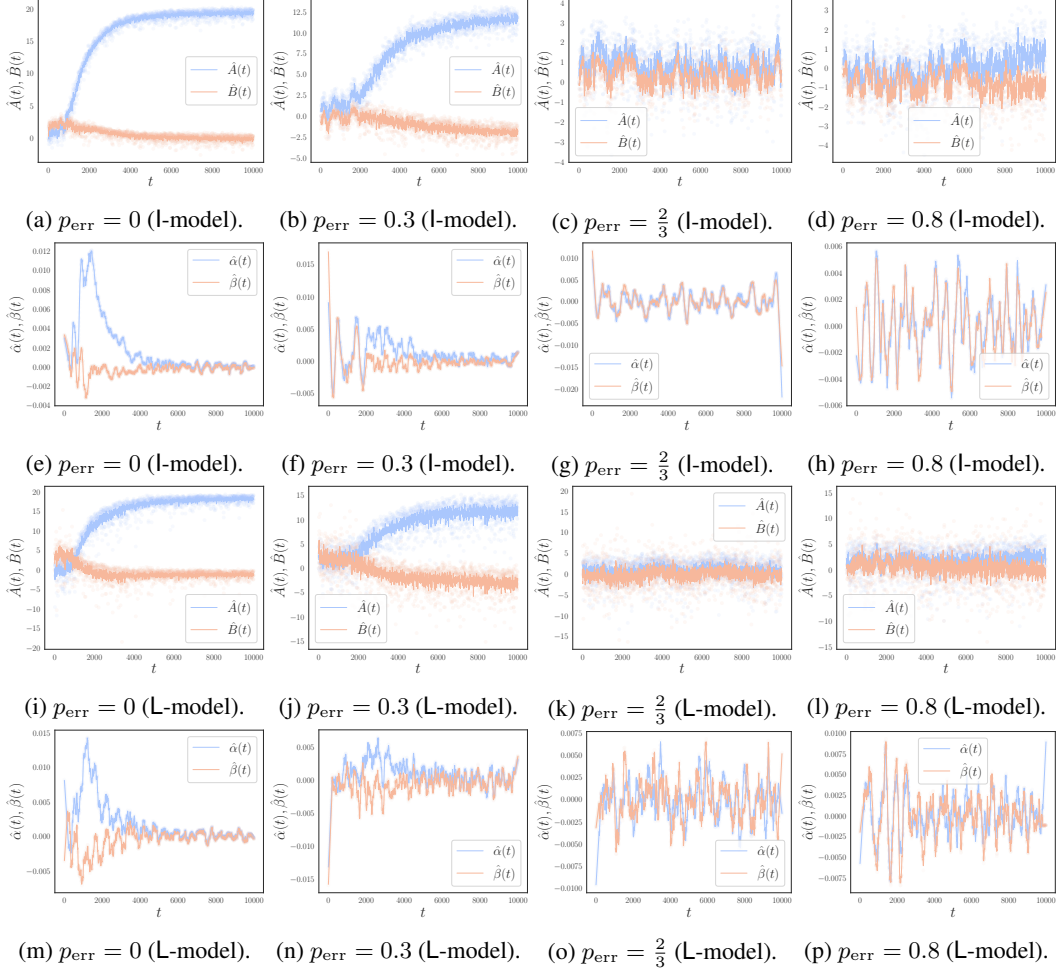


Figure D.5: **Effect of label corruption ratio** p_{err} . Estimated $\hat{A}(t)$, $\hat{B}(t)$ (the first and the third rows) and $\hat{\alpha}(t)$, $\hat{\beta}(t)$ (the second and the fourth rows), under l-model (the first two rows) and L-model (the last two rows) on GEOMNIST under various label corruptions ratios $p_{\text{err}} \in \{0, 0.3, 2/3, 0.8\}$ are shown. Note $\hat{A}(t)$ and $\hat{B}(t)$, $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ become indistinguishable when p_{err} is large, supporting Theorem 3.1.

local elasticity strengths $\alpha(t)$ and $\beta(t)$, when adjusted for the evolution of the norm of $\bar{\mathbf{X}}^k(t)$. As a surrogate, we simply multiply the strengths by the average norm, $\text{avg}_k \left\| \bar{\mathbf{X}}^k(t) \right\|_2$ and show the results in Figure D.6. Again we use a large window size $\omega = 351$ in the Savitzky–Golay filter to highlight the general trends. We observe that when $p_{\text{err}} = 0$, the general trend is similar to the unadjusted versions: $\alpha(t)$ has an initial increasing stage, and then converges to the vicinity of zero. We also note that $\alpha(t)$ and $\beta(t)$ become more indistinguishable as we increase p_{err} , similarly to the unadjusted cases Figure D.5.

Simulations of the LE-ODE. With $\alpha(t)$ and $\beta(t)$ estimated (using either the l-model or the L-model), we can simulate the LE-ODE under either l-model or L-model. In our setup, the initial value for the k th class is set to be $\zeta^K \sim \mathcal{N}_K(\mathbf{0}, \sigma_k \mathbf{I}_{K^2})$ for all $k \in [K]$ with $\sigma_k = \left\| \bar{\mathbf{X}}^k(0) \right\|_2 / \sqrt{K}$ with $\bar{\mathbf{X}}^k(0)$ sampled from simulations in DNNs. Empirically, we find that simulations under the l-model does not generate faithful trajectories compared with the ground truth (genuine dynamics from simulations on deep neural nets), as shown in Figure D.7. Hence in Figure D.8, we only show the case when the simulation is done under the L-model where the captions indicate which model (l-model or L-model) the estimation of $\alpha(t)$ and $\beta(t)$ is performed. Here $\bar{X}_i^k(t) \in \mathbb{R}$ denotes the i -th logit from

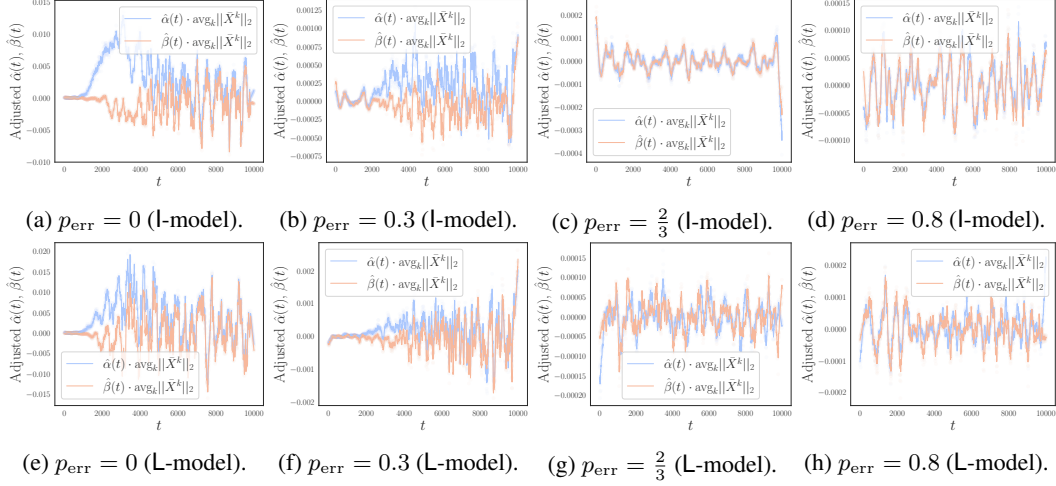


Figure D.6: **Effect of label corruption ratio p_{err} (adjusted for logits norm).** Estimated $\hat{\alpha}(t)$, $\hat{\beta}(t)$ under l-model (the first two rows) and L-model (the last two rows) on GEOMNIST under various label corruptions ratios $p_{\text{err}} \in \{0, 0.3, 2/3, 0.8\}$ are shown. All quantities were multiplied by the average norm of per-class means in each iteration, i.e., Note $\hat{A}(t)$ and $\hat{B}(t)$, $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ become indistinguishable when p_{err} is large, supporting Theorem 3.1.

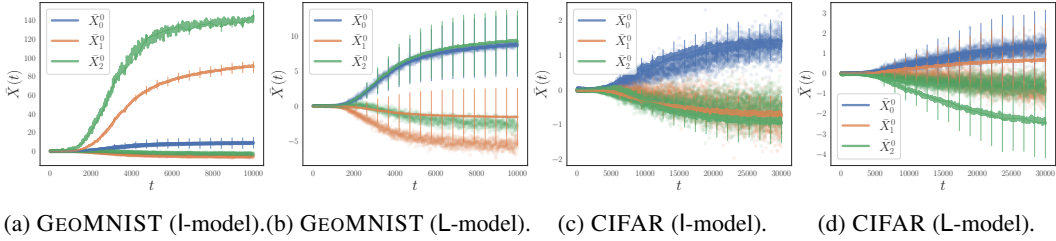


Figure D.7: **Simulated LE-ODE solutions under the l-model versus genuine dynamics.** We use $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ estimated from l-model ((a) and (c)) or L-model ((b) and (d)) and numerically simulate the solution under the l-model. The results were overlaid with true dynamics from neural nets. Note that the results are uniformly bad, indicating l-model is not expressive enough to capture the genuine dynamics in deep neural nets.

the per-class mean logits vector of the k -th class; as explained in Appendix B, a well-trained model should have the k -th logit being the largest among all $i \in [K]$ in \bar{X}_i^k when t is sufficiently large. Since the estimation of $\alpha(t)$ and $\beta(t)$ relies on numerical differentiation, which smooths the data and reduces their magnitudes, we manually align the k -th logit from the k -th per-class mean vector, \bar{X}_i^k , with that from the ground truth. All simulations are performed with $K = 3$ and for $N = 500$ trials with the initial data being Gaussian random vectors with zero mean and identity covariance such that the norm at initialization is approximately equal to that from the ground truth.

We observe the following: **(i) Both l-model and L-model approximately preserve the relative magnitude between different logits.** We find that the ratios between converging values of sample paths (dashed lines), $\lim_{t \rightarrow \infty} \bar{X}_i^k(t) / \bar{X}_j^k(t)$ for $i \neq j$, are roughly equal to the ground truth. This indicates that both models can capture the relative magnitudes of logits in real dynamics. **(ii) The l-model fails to identify the correct class.** After manual alignment, we note the l-model does not always yield faithful results, meaning the largest logit from the k -th class, $\max_{l \in [K]} \bar{X}_l^k(t)$ for large t , is not necessarily k . This is not surprising though, since the l-model itself does not differentiate features from different classes, and $\alpha(t)$ and $\beta(t)$ thus estimated fails to honor the interactions between different classes. **(iii) The L-model is able to identify the correct class while obliviously recovers the trajectories of incorrect classes.** On the other hand, simulated trajectories from the L-model faithfully recover the trajectories of correct classes $\arg\max_{l \in [K]} \bar{X}_l^k$ for each $1 \leq k \leq K = 3$. However, for any $k \in [K]$, we note that the trajectories of incorrect classes (i.e., $\{j : j \neq k\}$) are sometimes mismatched. This is because that the L-model, by construction, only uses the information from the correct class, i.e., $H_{i,j}$

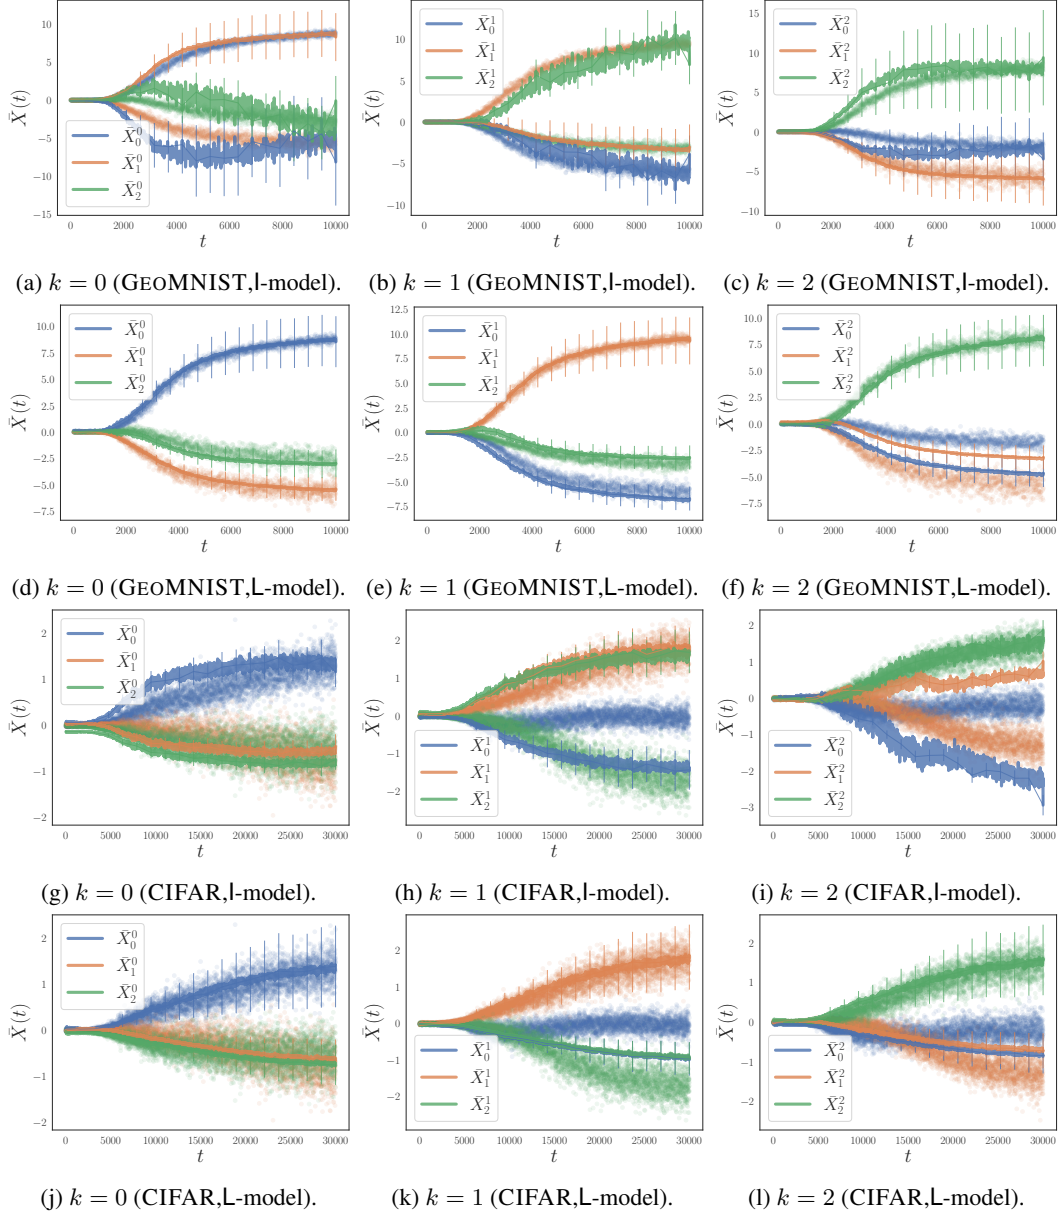
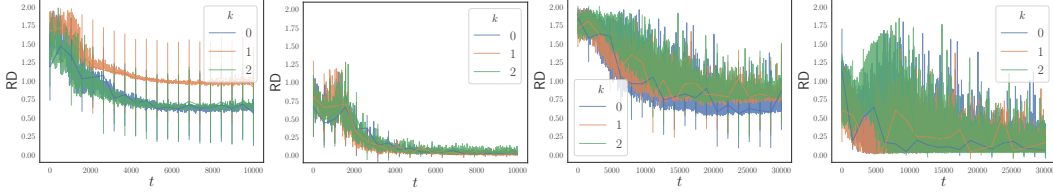


Figure D.8: **Simulated LE-ODE paths under the L-model versus genuine dynamics.** The simulation is done under the L-model with estimated $\hat{\alpha}(t)$ and $\hat{\beta}(t)$ from the l-model (the first and the third rows) and the L-model (the second and the fourth rows), on the GEOMNIST (the first two rows) and CIFAR (the last two rows). We show the trajectories for each class $k \in [3]$, where $\bar{X}_l^k(t)$ denotes the path of the l -th logit of the k -th class for $k, l \in [3]$.



(a) GEOMNIST (l-model). (b) GEOMNIST (L-model). (c) CIFAR (l-model). (d) CIFAR (L-model).

Figure D.9: **Relative difference RD_k between genuine and simulated dynamics.** The RD is computed according to equation D.21 (the lower the better). Note that the L-model performs better than l-model throughout training and better captures the later stages of the training (indicated by decreasing RD), supporting our discussions in Appendix B.2.

is set to be $d_j d_j^\top / d_j^\top d_j$ while not specifying other directions. A model that is capable of identifying incorrect classes needs necessarily more information on those classes. We postulate that a better model might be

$$\mathbf{H}_{i,j} = p_j \frac{d_j d_j^\top}{d_j^\top d_j} + \sum_{l \neq j}^K p_l \frac{d_l d_l^\top}{d_l^\top d_l}, \quad (\text{D.20})$$

with $p_j > p_l$ for all $l \neq j$. We leave explorations along this direction in future works.

Residue of LE-ODE Simulations. Under the same experiment setup, we also visualize the residue of using LE-ODE to imitate the genuine dynamics of neural nets, as shown in Figure D.9. We measure the goodness-of-fit via relative difference (RD) defined for each class $k \in [K]$ as

$$RD_k(t) := \frac{\|\bar{\mathbf{X}}^k(t) - \bar{\mathbf{Y}}^k(t)\|_{\mathbf{H}^k}}{\left(\|\bar{\mathbf{X}}^k(t)\|_2 + \|\bar{\mathbf{Y}}^k(t)\|_2\right) / 2}, \quad (\text{D.21})$$

where $\bar{\mathbf{X}}(t)$ and $\bar{\mathbf{Y}}(t)$ are genuine and simulated trajectories, respectively, and $\|\cdot\|_{\mathbf{H}}$ denotes the norm induced by the matrix \mathbf{H} that is used to define models (i.e., the identity matrix for the l-model and $\bar{\mathbf{H}}$ in equation 8 for the L-model). This choice normalizes the difference under the similarity defined by \mathbf{H} (which we care the most) and ranges from 0 to 2 (the lower the better). From the results we note that: **(i) The L-model performs consistently better than the l-model throughout training.** We note that the RD under the L-model are overall smaller and there is no significant differences across classes. **(ii) The L-model is better suited for capturing later stages of training.** This can be seen from a decreasing trend of the RD under the L-model, which corroborates our discussions in Appendix B.2. In particular, the approximation becomes better as training progresses (indicated by a decreasing RD). However, the performance of the L-model around initialization is still commendable. **(iii) The L-model is not perfect.** Although RD under L-model is small in the terminal stage (of the order 10^{-2} to 10^{-1}), it is non-zero, and it has a higher RD in the early stage of training. This indicates that non-dominant directions are also important for the LE-ODE to capture the remainder of the feature similarity. A possible avenue for future research in this regard is discussed in equation D.20.

D.3 The Two-Stage Behavior of Logits Evolution

An interesting observation that can be made when going through the experiments is the emergence of a two-stage behavior in many quantities. Specifically: (i) the training loss and validation loss do not decrease at a perceivable rate in the first few hundreds (or thousands) iterations; then they begin to drop at a relatively fast rate until convergence; (ii) the local elasticity strength $\alpha(t)$ increases at the initial stage, then drops, which is also manifested by the behavior of $A(t)$, which resembles roughly a sigmoidal curve; (iii) the magnitudes of $\text{avg}_{k \in [K]} \|\bar{\mathbf{X}}^k\|_2$ also resembles that of $A(t)$ (not shown), which has a fast growing stage and a converging stage.

We coin this seemingly generic phenomenon as the *two-stage behavior* that consists of a *de-randomization* stage and an *amplification* stage, and demonstrate the supporting experiments in Figure D.10. Here we trained on GEOMNIST with $K = 3$, where each triangle is a hyperplane spanned by the per-class mean logits vectors, $\bar{\mathbf{X}}^k(t) \in \mathbb{R}^3$ for $k \in [3]$. In Figure D.10a we plot those

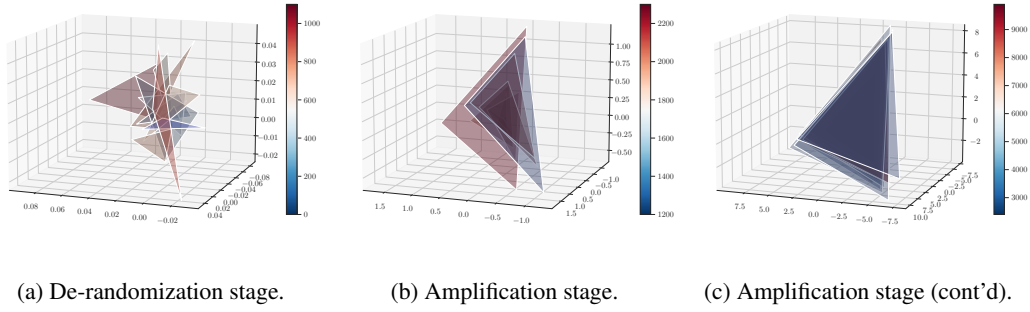


Figure D.10: **The hyperplanes formed by $\{\bar{X}_t^k\}_{k=1}^3$ with colors corresponding to iterations.** (a) In the first 1200 iterations, the validation loss does not drop perceptibly, and the hyperplanes are visibly chaotic. (b)-(c) Starting from around the 1200-th iteration, the validation loss drops, and the hyperplanes exhibit converging behavior.

hyperplanes for the first 1200 iterations; Figure D.10b iterations from 1200 to 2400; and Figure D.10c the remaining iterations. We observe that in the first 1200 iterations, the hyperplanes are “chaotic” in that their behavior is highly dependent on specific initialization values. In this de-randomization stage, the supervision guides the dynamics to identify the correct and *deterministic* (c.f. Proposition 3.4) direction for the separation of features from random initialization, thus the name *de-randomization*. With the correct direction being identified (around iteration 1200), the losses begin to drop at a relatively fast speed and the hyperplane remains approximately the same throughout the training while the magnitude of logits increases, which pushes the classes to be more discriminative and further drives down the losses, hence the *amplification* stage.

We believe a more precise characterization of $\alpha(t)$ and $\beta(t)$ and potentially generalization to the \mathbf{E} matrix would likely help us to study this two-stage behavior more rigorously and potentially answer the interesting questions such as *how long does the first stage take on average?* *How does local elasticity ($\alpha(t)$ and $\beta(t)$) affect this behavior?* We leave these questions for future works.

E Future Work and Extensions

General LE Matrix. Throughout the paper, we have modeled the LE matrix \mathbf{E} as $(\alpha(t) - \beta(t))\mathbf{I}_K + \beta(t)\mathbf{1}_K\mathbf{1}_K^\top$. Although it depends on t , this model falls short when we move into the more realistic realm where the inter-class and intra-class effects are dependent on the class labels. When \mathbf{E} is SPD, under the same assumptions as in Theorem 3.1, due to a theorem by Schur (Theorem 9.B.1, [37]), we know the eigenvalues of $(\mathbf{E} \otimes \mathbf{I}_K) \circ \mathbf{H}$ are bounded within $\lambda_{\min}(\mathbf{E}) \min_i H_{ii}$ and $\lambda_{\max}(\mathbf{E}) \max_i H_{ii}$, where $\{H_{ii}, i \in [Kp]\}$ is the diagonal entries of the matrix \mathbf{H} , hence we expect a very similar result in this case as in Theorem 3.1. In the more general case where \mathbf{E} is symmetric but not necessarily semi-definite, a more precise analysis on the spectrum of $(\mathbf{E} \otimes \mathbf{I}_K) \circ \mathbf{H}$ is needed, though the proof framework would not be too different.

Mini-batch Training, Imbalanced Datasets, and Label Corruptions. As discussed in Section 3, we can incorporate mini-batches and imbalanced datasets in our model easily. Taking imbalanced datasets as an example, recall that each block of \mathbf{M}_t in equation (5) takes the form of $E_{k,l}\mathbf{H}_{k,l}/K$, where $1/K$ signifies that each class has the same $1/K$ probability of being sampled during any iteration in training. This can be generalized by changing the (k, l) -th block to $E_{k,l}\mathbf{H}_{k,l} \cdot p_l$ for $\sum_l p_l = 1$, where p_l is the probability of class l being sampled. More succinctly, instead of defining $\mathbf{M}_t = (\mathbf{E}_t \otimes \mathbf{I}_K) \circ \mathbf{H}/K$, we let $\mathbf{M}_t = (\mathbf{E}_t \otimes \mathbf{P}) \circ \mathbf{H}/K$ for a K -by- K doubly stochastic matrix \mathbf{P} that models this sampling effect. In the same vein, we can also model the case when the data are polluted by corrupted labels. Let $\mathbf{P} = (p_{k,l})_{k,l} \in \mathbb{R}^{K \times K}$ with $p_{k,l}$ representing the probability of a sample from class k mis-labelled as class l . A well-defined model needs further assumption on the structure of \mathbf{P} and we leave this theoretical modeling to future work.

Covariance Structures and Fine-Grained Analyses. Although our model encompasses a covariance term in the LE-SDE model, we do not explicitly use its structure. Nonetheless, as indicated from the proof of Theorem 3.1 (cf. Appendix C.2.3), the relative magnitude of the covariance to the

drift term (i.e., the local elasticity effect) affects the separation when $\gamma(t) = \Theta(1/t)$. However, when the order of $\gamma(t)$ is guaranteed to be strictly above or below $1/t$ as $t \rightarrow \infty$, the covariance affects the separation only through the constant factor for the separation rate. That said, a more precise analysis of covariance would by all means facilitate fine-grained analyses at the edge of separation.

Beyond L-model for Imitating Genuine Dynamics of DNNs. We show in Section 4 that using estimates of $\alpha(t)$ and $\beta(t)$, the L-model can be used to imitate the genuine dynamics of DNNs. As is shown in more detail in Appendix D.2, we note that although simulations under the L-model are already superior to those under the l-model in that the correct classes are identified, the L-model sometimes still fails to identify the correct trajectories for the incorrect classes. This is not very surprising though, as the supervision from the labels only affects L-model though the \mathbf{H} matrix, whose (i, j) -th block is defined as $\bar{\mathbf{H}}^j = \mathbf{d}_j \mathbf{d}_j^\top / \mathbf{d}_j^\top \mathbf{d}_j$ where $\mathbf{d}_j = \mathbf{e}_j - \mathbf{1}_K / K$ — which only encodes information about the correct class. We postulate that a more precise model might be to assign the (i, j) -th block of \mathbf{H} as

$$\mathbf{H}_{i,j} = p_j^j \bar{\mathbf{H}}^j + \sum_{l \neq j}^K p_l^j \bar{\mathbf{H}}^l, \quad p_j^j > p_l^j, \quad \sum_{l=1}^K p_l^j = 1, \quad \forall l \neq j, \quad j \in [K]. \quad (\text{E.1})$$

The Two-Stage Behavior. In Section 4, we observed a clear two-stage behavior of our numerical simulation of our LE-SDE, as well as in real deep learning dynamics. The first stage is a *de-randomization* stage, which gradually eliminates the effect of random initialization and searches for the correct directions to be separated. The second stage is an *amplification* stage, where the model amplifies the magnitudes of the features in those directions. As can be seen from Appendix D.2, these empirical observations naturally lead to many interesting questions: Is this a universal phenomenon in deep learning, and what are the conditions to guarantee entering the second stage? Can our LE-SDE predict such a two-stage phenomenon theoretically? What is the role of local elasticity in this transition? We leave the investigation of these questions to future works.