## 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

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

When $m \gg r$, we have approximately

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

Next, observe that

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

hence taken together, the calculations above give

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

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

$$
\mathrm{d} X_{i}^{1}(t)=\left(\frac{\alpha}{2} \bar{X}(t)+\frac{\beta}{2} \bar{Y}(t)\right) \mathrm{d} t+\sigma \mathrm{d} W^{i}(t)
$$

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 $\frac{c_{1}-c_{2}}{2} \mathrm{e}^{\frac{\alpha-\beta}{2} t} \rightarrow 0$ 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 $c_{k}=\mathbb{E}_{\text {data }}\left[X^{k}(0) \mid \theta(0)=\theta_{0}\right]$ for $k \in\{1,2\}$ where $\theta(0)$ is all parameters of the neural net at initialization and $\theta_{0}$ 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

$$
\begin{aligned}
\min _{1 \leq i \leq n} \frac{c_{1}-c_{2}}{2} \mathrm{e}^{\frac{\alpha-\beta}{2} t}+\frac{c_{1}+c_{2}}{2} & \mathrm{e}^{\frac{\alpha+\beta}{2} t}-c_{x}+X_{i}^{1}(0)+\sigma W_{i}^{1}(t) \\
& >\max _{1 \leq j \leq n}-\frac{c_{1}-c_{2}}{2} \mathrm{e}^{\frac{\alpha-\beta}{2} t}+\frac{c_{1}+c_{2}}{2} \mathrm{e}^{\frac{\alpha+\beta}{2} t}-c_{2}+X_{j}^{2}(0)+\sigma W_{j}^{2}(t)
\end{aligned}
$$

which is equivalent to

$$
\left(c_{1}-c_{2}\right) \mathrm{e}^{\frac{\alpha-\beta}{2} t}-c_{1}+\min _{1 \leq i \leq n} X_{i}^{1}(0)+\sigma W_{i}^{1}(t)>-c_{2}+\max _{1 \leq j \leq n} X_{j}^{2}(0)+\sigma W_{j}^{2}(t) .
$$

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

## 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 $\boldsymbol{H}$ 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 $\boldsymbol{X}^{k}(m) \in \mathbb{R}^{K}$ 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 $\boldsymbol{X}_{i}^{k}$ should be governed by

$$
\begin{equation*}
\boldsymbol{X}_{i}^{k}(m)-\boldsymbol{X}_{i}^{k}(m-1) \approx h\left[\frac{\partial \boldsymbol{X}_{i}^{k}(m-1)}{\partial \boldsymbol{w}} \frac{\partial \boldsymbol{X}_{J_{m}}^{L_{m}}}{\partial \boldsymbol{w}}\left(\boldsymbol{e}_{L_{m}}-\operatorname{softmax}\left(\boldsymbol{X}_{J_{m}}^{L_{m}}\right)\right)\right] . \tag{B.1}
\end{equation*}
$$

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

$$
\begin{equation*}
\boldsymbol{X}_{i}^{k}(m)-\boldsymbol{X}_{i}^{k}(m-1) \approx \nabla_{\boldsymbol{w}} \boldsymbol{X}_{i}^{k}(m-1) \Delta \boldsymbol{w}(m-1), \tag{B.2}
\end{equation*}
$$

where we observe that

$$
\begin{equation*}
\Delta \boldsymbol{w}(m-1)=-h \frac{\partial L(\boldsymbol{w}(m-1))}{\partial \boldsymbol{w}}=h \frac{\partial \boldsymbol{X}_{J_{m}}^{L_{m} \top}}{\partial \boldsymbol{w}}\left(\boldsymbol{e}_{L_{m}}-\operatorname{softmax}\left(\boldsymbol{X}_{J_{m}}^{L_{m}}\right)\right) . \tag{B.3}
\end{equation*}
$$

However, as equation (B.1) is highly non-linear, $J_{m}$ and $L_{m}$ are random, and the Gram matrix $\frac{\partial \boldsymbol{X}_{i}^{k}(m-1)}{\partial \boldsymbol{w}} \frac{\partial \boldsymbol{X}_{J_{m}}^{L_{m}}{ }^{\top}}{\partial \boldsymbol{w}}$ 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, $\overline{\boldsymbol{X}}^{k}=\mathbb{E} \widetilde{\boldsymbol{X}}^{k}$, where $\widetilde{\boldsymbol{X}}^{k}$ is a random sample, for each class $k$. In the limit of $h \rightarrow 0$ with the identification of time $t=m h$, the above modeling allows us to re-write (B.1) in the per-class mean $\overline{\boldsymbol{X}}^{k}$ and logits from a generic sample $\widetilde{\boldsymbol{X}}^{k}$ in this continuous limit as

$$
\begin{align*}
\mathrm{d} \widetilde{\boldsymbol{X}}_{t}^{k} & \approx \mathbb{E}_{L \sim \operatorname{Unif}([K])}\left[\mathbb{E}_{\widetilde{\boldsymbol{X}} \sim \mathcal{D}_{t}^{L}}\left[\frac{\partial \boldsymbol{X}_{i}^{k}(m-1)}{\partial \boldsymbol{w}} \frac{\partial \widetilde{\boldsymbol{X}}^{\top}}{\partial \boldsymbol{w}}\left(\boldsymbol{e}_{L}-\operatorname{softmax}(\widetilde{\boldsymbol{X}})\right)\right]\right] \mathrm{d} t+\boldsymbol{\Sigma}_{t}^{\frac{1}{2}} \mathrm{~d} \boldsymbol{W}_{t}, \\
& \approx \frac{1}{K} \sum_{L}\left(\left[\mathbb{E}_{\widetilde{\boldsymbol{X}}^{\prime} \sim \mathcal{D}_{t}^{k}, \widetilde{\boldsymbol{X}} \sim \mathcal{D}_{t}^{L}} \frac{\partial \widetilde{\boldsymbol{X}}^{\prime}}{\partial \boldsymbol{w}} \frac{\partial \widetilde{\boldsymbol{X}}^{\top}}{\partial \boldsymbol{w}}\right]\left(\boldsymbol{e}_{L}-\operatorname{softmax}\left(\overline{\boldsymbol{X}}_{t}^{L}\right)\right)\right) \mathrm{d} t+\boldsymbol{\Sigma}_{t}^{\frac{1}{2}} \mathrm{~d} W_{t},  \tag{B.4}\\
& =\frac{1}{K} \sum_{L}\left(\Theta_{k, L}\left(\boldsymbol{e}_{L}-\operatorname{softmax}\left(\overline{\boldsymbol{X}}_{t}^{L}\right)\right)\right) \mathrm{d} t+\boldsymbol{\Sigma}_{t}^{\frac{1}{2}} \mathrm{~d} \boldsymbol{W}_{t} .
\end{align*}
$$

The presence of expectation over non-linearity posed considerate 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 $\boldsymbol{z}^{k}$ from the $k$-th class, in a well-trained neural network, the probability vector associated with $\boldsymbol{z}^{k}$, 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 $e_{k}-\operatorname{softmax}\left(\bar{X}_{t}^{k}\right)$ should be roughly in the direction of $\boldsymbol{e}_{k}-\frac{1}{K} \mathbf{1}$. To approximate this limit by the limit from a linear map, we may use our choice of $\boldsymbol{H}_{k l}=\overline{\boldsymbol{H}}^{l}:=\frac{\boldsymbol{d}_{l} \boldsymbol{d}_{l}^{\top}}{\left\|\boldsymbol{d}_{l}\right\|^{2}}$ for all $k, l \in[K]$, such that in the limit of $t \rightarrow \infty$, we expect

$$
\begin{equation*}
\boldsymbol{e}_{k}-\operatorname{softmax}\left(\overline{\boldsymbol{X}}_{t}^{k}\right) \approx C \overline{\boldsymbol{H}}^{k} \overline{\boldsymbol{X}}_{t}^{k} \tag{B.5}
\end{equation*}
$$

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 $\boldsymbol{H}_{k, l}$ transforms the mean logit $\overline{\boldsymbol{X}}_{t}^{L}$ 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 $X(t) \in \mathbb{R}^{K p}$ (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

$$
\begin{equation*}
F(\widetilde{\boldsymbol{X}}(t), t):=\Theta(t)\left(\left[e_{k}-\sigma\left(\widetilde{\boldsymbol{X}}^{k}(t)\right)\right]_{k=1}^{K}\right) \tag{B.6}
\end{equation*}
$$

where we wrote $[\cdot]$ for vector concatenation and $\Theta(t) \in \mathbb{R}^{(K p) \times(K p)}$ 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 $\varphi(t)=\left(\varphi_{k}(t)\right)_{k=1}^{K} \equiv \bar{X}(t):=\mathbb{E}_{B} X(t)$ where the expectation is taken with respect to the diffusion. Concretely, we have

$$
\begin{equation*}
F(\widetilde{\boldsymbol{X}}(t), t) \approx \widetilde{F}(\widetilde{\boldsymbol{X}}(t), t):=F(\varphi(t), t)+\nabla_{X} F(\varphi(t), t)(\widetilde{\boldsymbol{X}}(t)-\varphi(t)) \tag{B.7}
\end{equation*}
$$

where $\nabla_{X} F$ denotes the Jacobian of $F$ with respect to the spacial variable $\widetilde{\boldsymbol{X}}$. For notation completeness, we introduce

$$
\begin{equation*}
p=\left(p_{k}\right)_{k=1}^{K} \in \mathbb{R}^{K p}, \quad p_{k}:=\sigma\left(\widetilde{\boldsymbol{X}}^{k}(t)\right) \in \mathbb{R}^{p}, \quad k \in[K], \tag{B.8}
\end{equation*}
$$

and similarly

$$
\begin{equation*}
\bar{p}=\left(\bar{p}_{k}\right)_{k=1}^{K} \in \mathbb{R}^{K p}, \quad \bar{p}_{k}:=\sigma\left(\overline{\boldsymbol{X}}^{k}(t)\right) \in \mathbb{R}^{p}, \quad k \in[K] . \tag{B.9}
\end{equation*}
$$

We write the per-class Jacobians as

$$
\begin{equation*}
J_{k k}=J_{k}:=\operatorname{diag}\left(\bar{p}_{k}\right)-\bar{p}_{k} \bar{p}_{k}^{T} \tag{B.10}
\end{equation*}
$$

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

$$
\begin{align*}
\widetilde{F}(\widetilde{\boldsymbol{X}}(t), t) & =\Theta(t)\left(\left[e_{k}-\bar{p}_{k}\right]_{k}+J(t)(\widetilde{\boldsymbol{X}}(t)-\varphi(t))\right)  \tag{B.11}\\
& =\Theta(t)\left(J(t) X(t)+\left[e_{k}-\bar{p}_{k}+J_{k} \varphi_{k}(t)\right]_{k}\right)
\end{align*}
$$

Define $\Psi: \mathbb{R}^{K p} \rightarrow \mathbb{R}^{K p}: z \mapsto\left[e_{k}-\sigma\left(z_{k}\right)\right]_{k}$ and write $\Psi_{k}: \mathbb{R}^{p} \rightarrow \mathbb{R}^{p}$ to be the $k$-th component of $\Psi$, using Taylor's theorem to expand $\Psi(z)$ around $\varphi(t)$ for each $t$, we have

$$
\begin{equation*}
\Psi=\Psi(\varphi)+J(t) \varphi-J(\varphi) z+o(\|z-\varphi\|) \tag{B.12}
\end{equation*}
$$

or

$$
\begin{equation*}
\Psi(\varphi)+J(t) \varphi=\Psi(z)+J(\varphi) z+o(\|z-\varphi\|) \tag{B.13}
\end{equation*}
$$

This implies that

$$
\begin{equation*}
\widetilde{F}=\Theta(t) J(t) \widetilde{\boldsymbol{X}}(t)+\Theta(t) R(t), \quad R(t ; z):=\Psi(z)+J(t) z+o(\|z-\varphi(t)\|) \tag{B.14}
\end{equation*}
$$

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 $z=u:=c \cdot\left[\mathbf{1}_{K} / K\right]_{k=1}^{K}$ 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, $u \in \operatorname{Ker} J(t)$ for all $t$ hence the residue $R(t ; u)=\Psi(u)+o(\|z-\varphi(t)\|)$ is a constant vector (which is colinear with $\boldsymbol{d}=\left(\boldsymbol{d}_{j}\right)_{j=1}^{K}$ defined by $\boldsymbol{d}_{j}=\boldsymbol{e}_{j}-\mathbf{1}_{K} / K$ 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, $\left\|\Psi\left(\varphi_{\infty}\right)\right\| \approx 0$ by construction. Hence the residue $\left.R\left(t ; \varphi_{\infty}\right)=J(t) \varphi_{\infty}+o\left(\| \varphi(t)-\varphi_{\infty}\right) \|\right)$. Here the $o(\cdot)$ term converges to 0 as training progresses, leaving us a term that is asymptotically equivalent to $v=\left(v_{k}\right)_{k=1}^{K}:=$ $J\left(\varphi_{\infty}\right) \varphi_{\infty} \in \mathbb{R}^{K^{2}}$, where $v_{k}=\left[\left(z_{k, i}-\sum_{j=1}^{K} p_{k, j} z_{k, j}\right) p_{i}\right]_{i=1}^{K} \in \mathbb{R}^{K}$. 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 $o\left(\left\|\varphi(t)-\varphi_{\infty}\right\|\right)$ 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 $\left\{\left(\boldsymbol{X}_{i}^{k}\right)_{i \in[n]}\right\}_{k \in[K]}$ is pairwisely separable at time $t$ iffor each pair of $1 \leq k<l \leq K$, there exists a direction $\boldsymbol{\nu}_{k, l}$ such that

$$
\begin{equation*}
\min _{i}\left\langle\boldsymbol{\nu}_{k, l}, \boldsymbol{X}_{i}^{k}(t)\right\rangle>\max _{j}\left\langle\boldsymbol{\nu}_{k, l}, \boldsymbol{X}_{j}^{l}(t)\right\rangle . \tag{C.1}
\end{equation*}
$$

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 $\left\{\left(\boldsymbol{X}_{i}^{k}\right)_{i \in[n]}\right\}_{k \in[K]}$ are asymptotically pairwisely separable iffor each pair of $1 \leq k<l \leq K$, there exist directions $\boldsymbol{\nu}_{k, l}(t)$ such that

$$
\begin{equation*}
\mathbb{P}\left(\min _{i}\left\langle\boldsymbol{\nu}_{k, l}(t), \boldsymbol{X}_{i}^{k}(t)\right\rangle>\max _{j}\left\langle\boldsymbol{\nu}_{k, l}(t), \boldsymbol{X}_{j}^{l}(t)\right\rangle\right) \rightarrow 1 . \tag{C.2}
\end{equation*}
$$

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 $\left\{\left(\boldsymbol{X}_{i}^{k}\right)_{i \in[n]}\right\}_{k \in[K]}$ are asymptotically universally separable if it is asymptotically pairwisely separable, and there exists $\nu$, 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.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 $\left\{\left(\boldsymbol{X}_{i}^{k}\right)_{i \in[n]}\right\}_{k \in[K]}$ are uniformly asymptotically pairwisely separable if for each pair of $1 \leq k<$ $l \leq K$, there exists a direction $\boldsymbol{\nu}_{k, l}$ such that

$$
\begin{equation*}
\liminf _{T \rightarrow \infty} \min _{k, l} \mathbb{P}\left(\min _{i}\left\langle\boldsymbol{\nu}_{k, l}, \boldsymbol{X}_{i}^{k}(t)\right\rangle>\max _{j}\left\langle\boldsymbol{\nu}_{k, l}, \boldsymbol{X}_{j}^{l}(t)\right\rangle, \text { for all } t>T\right)=1 \tag{C.3}
\end{equation*}
$$

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 $\left\{\left(\boldsymbol{X}_{i}^{k}\right)_{i \in[n]}\right\}_{k \in[K]}$ are uniformly asymptotically universally separable if it is uniformly asymptotically pairwisely separable, and there exists $\nu$, 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 (I-model). Let $\boldsymbol{H}_{k, l}=\boldsymbol{I}_{p}$, then the solution to the LE-ODE (6) is given by

$$
\begin{equation*}
\overline{\boldsymbol{X}}(t)=\boldsymbol{c} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}+\left(\mathbf{1}_{K} \otimes \boldsymbol{c}_{0}\right) \mathrm{e}^{\frac{1}{K} A(t)+\frac{K-1}{K} B(t)}, \tag{7}
\end{equation*}
$$

where $\boldsymbol{c}=\left(\boldsymbol{c}_{k}\right)_{k=1}^{K} \in \mathbb{R}^{K p}$ and $\boldsymbol{c}_{0} \in \mathbb{R}^{p}$ are constants with $\sum_{k=1}^{K} \boldsymbol{c}_{k}=\mathbf{0} \in \mathbb{R}^{p}$ and $\overline{\boldsymbol{X}}(0)=\boldsymbol{c}+\boldsymbol{c}_{0}$.

Proof of Proposition 3.2. When $\alpha(t)=\alpha$ and $\beta(t)=\beta$ are constants, $\boldsymbol{E}$ 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 $E=\boldsymbol{Q} \boldsymbol{\Lambda} \boldsymbol{Q}^{-1}$ is given by

$$
\boldsymbol{\Lambda}=\operatorname{diag}\left(\lambda_{1}, \lambda_{0}, \ldots, \lambda_{0}\right) \in \mathbb{R}^{K \times K}, \quad \boldsymbol{Q}=\left[\begin{array}{ccccc}
1 & 1 & 1 & \cdots & 1  \tag{C.4}\\
1 & -1 & 0 & \cdots & 0 \\
1 & 0 & -1 & \cdots & 0 \\
\vdots & & & & \vdots \\
1 & 0 & 0 & \cdots & -1
\end{array}\right]
$$

Hence the eigenpairs of $\boldsymbol{E} \otimes \boldsymbol{I}_{p}$ are $\left(\boldsymbol{e}_{1}-\boldsymbol{e}_{k}\right) \otimes \boldsymbol{e}_{j}$ for $1<k \leq K$ and $1 \leq j \leq p$ with eigenvalue $\lambda_{0}$ and $\left(\mathbf{1}_{K} \otimes \boldsymbol{e}_{j}\right)$ for $1 \leq j \leq p$ with eigenvalue $\lambda_{1}$. Reparameterizing the solution such that the initial values are $\overline{\boldsymbol{X}}^{k}(0)=\boldsymbol{c}_{0}+\boldsymbol{c}_{k}$ for all $k \in[K]$, we have

$$
\begin{equation*}
\overline{\boldsymbol{X}}_{t}=\boldsymbol{c} \mathrm{e}^{\lambda_{0} t}+\left(\mathbf{1}_{K} \otimes \boldsymbol{c}_{0}\right) \mathrm{e}^{\lambda_{1} t} \tag{C.5}
\end{equation*}
$$

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

$$
\begin{equation*}
\overline{\boldsymbol{X}}_{t}=\boldsymbol{c} \boldsymbol{c}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}+\left(\mathbf{1}_{K} \otimes \boldsymbol{c}_{0}\right) \mathrm{e}^{\frac{1}{K} A(t)+\frac{K-1}{K} B(t)} \tag{C.6}
\end{equation*}
$$

where $\sum_{k=1}^{K} \boldsymbol{c}_{k}=\mathbf{0}$.

## C.2.2 Proof of Proposition 3.3

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

$$
\begin{equation*}
\overline{\boldsymbol{X}}(t)=\boldsymbol{c}_{0}+C_{1} \boldsymbol{d} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}+\left(\sum_{l=1}^{K-1} C_{2 l} \boldsymbol{f}_{l}\right) \mathrm{e}^{\frac{1}{K} A(t)+\frac{1}{K(K-1)} B(t)} \tag{9}
\end{equation*}
$$

where $\boldsymbol{f}_{l}$ 's are fixed vectors in $\mathbb{R}^{K^{2}}, \boldsymbol{c}_{0} \in \mathbb{R}^{K^{2}}$ is a constant vector with $K(K-1)$ degrees of freedom, and $C_{1}, C_{2 l} \in \mathbb{R}$, for $l \in[K-1]$ are constants.

Proof of Proposition 3.3. Recall that the $(k, l)$-th block of the $\boldsymbol{H}$ matrix is $\overline{\boldsymbol{H}}^{l}:=\boldsymbol{d}_{l} \boldsymbol{d}_{l}^{\top} / \boldsymbol{d}_{l}^{\top} \boldsymbol{d}_{l}$. When $\alpha(t)=\alpha$ and $\beta(t)=\beta$ are constants, we immediately find the matrix $\left(\boldsymbol{E} \otimes \boldsymbol{I}_{K}\right) \circ \boldsymbol{H}$ has an eigenvalue of $\alpha-\beta$ with multiplicity one whose eigenvector is $\boldsymbol{d}$, and an eigenvalue of 0 with multiplicity $K(K-1)$ - since the null spaces of $\boldsymbol{H}_{k l}$ 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 $\boldsymbol{M}_{t}$.

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

$$
\begin{equation*}
\left[w_{1}+w_{2}, 2 w_{1}, 2 w_{2}, 2 w_{3}-w_{4}, w_{3}, w_{4},-w_{5}-2, w_{5}, w_{6}\right]^{\top} \tag{C.7}
\end{equation*}
$$

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

$$
\begin{equation*}
\left[-\xi_{1} w_{1}+\xi_{2}, \xi_{2} w_{1}-\xi_{5}, \xi_{3} w_{1}-\xi_{4},-\xi_{2} w_{1}+\xi_{4}, \xi_{1} w_{1}-\xi_{3},-\xi_{3} w_{1}+\xi_{6}, 1-w_{1}, w_{1}, w_{2}\right]^{\top} \tag{C.8}
\end{equation*}
$$

where $w_{1}$ and $w_{2}$ are free parameters and

$$
\begin{gather*}
\xi_{1}=\frac{2 \alpha+\beta}{3 \beta}, \quad \xi_{2}=\frac{\alpha+2 \beta}{3 \beta}, \quad \xi_{3}=\frac{\alpha-\beta}{3 \beta} \\
\xi_{4}=\frac{\alpha^{2}+\alpha \beta+7 \beta^{2}}{6 \alpha \beta+3 \beta^{2}}, \quad \xi_{5}=\frac{\alpha^{2}+4 \alpha \beta-5 \beta^{2}}{6 \alpha \beta+3 \beta^{2}}, \quad \xi_{6}=\frac{\alpha^{2}-2 \alpha \beta-8 \beta^{2}}{6 \alpha \beta+3 \beta^{2}} . \tag{C.9}
\end{gather*}
$$

## 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 $\boldsymbol{H}=\left(\boldsymbol{H}_{i j}\right)_{i j}$ is positive semi-definite (PSD) with positive diagonal entries. As $t \rightarrow \infty$, we have ${ }^{3}$.

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 $\infty$ 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 $\boldsymbol{H} \in \mathbb{R}^{K p \times K p}$ be a symmetric positive semi-definite (SPD) matrix with positive diagonal entries $\left(d_{i}\right)_{i=1}^{K p}$, recall that we denote by $\boldsymbol{H}_{i j} \in \mathbb{R}^{p \times p}$ the $(i, j)$-th block of $\boldsymbol{H}$ for $i, j \in[K]$. We write $\boldsymbol{M}_{t}=\frac{1}{K}\left(\boldsymbol{E}_{t} \otimes \boldsymbol{I}_{K}\right) \circ \boldsymbol{H}$, where $(\boldsymbol{A} \circ \boldsymbol{B})_{i j}=A_{i j} B_{i j}$ 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 $\boldsymbol{E} \otimes \boldsymbol{I}_{p}$ has an eigenvalue $\lambda_{0}=(\alpha-\beta) / K$ with multiplicity $p(K-1)$ and $\lambda_{1}=(\alpha+(K-1) \beta) / K$ with

[^0]multiplicity $p$. Writing $\lambda_{\min }$ and $\lambda_{\max }$ as the minimum and the maximum of $\left\{\lambda_{0}, \lambda_{1}\right\}$ respectively, by Schur's theorem (Theorem 9.J.2, [37]), we have
\[

$$
\begin{equation*}
\lambda_{\min } \min _{1 \leq i \leq K p} d_{i} \leq \lambda_{i}\left(\boldsymbol{M}_{t}\right) \leq \lambda_{\max } \max _{1 \leq i \leq K p} d_{i} \tag{C.10}
\end{equation*}
$$

\]

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

Note the eigendecomposition of $\boldsymbol{H}$ is $\sum_{i=1}^{K p} \mu_{i} \boldsymbol{u}_{i} \boldsymbol{u}_{i}^{\top}$, and thus the solution to $\overline{\boldsymbol{X}}_{t}^{\prime}=\boldsymbol{M}_{t} \overline{\boldsymbol{X}}$ is

$$
\begin{equation*}
\overline{\boldsymbol{X}}_{t}=\overline{\boldsymbol{X}}_{0}+\sum_{i=1}^{K p} c_{i} \boldsymbol{u}_{i} \mathrm{e}^{\mu_{i} t}, \quad \overline{\boldsymbol{X}}_{0}=\sum_{i=1}^{K p} c_{i} \boldsymbol{u}_{i} \tag{C.11}
\end{equation*}
$$

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

$$
\begin{align*}
\widetilde{\boldsymbol{X}}^{k}(t) & =\widetilde{\boldsymbol{X}}^{k}(0)+\boldsymbol{M}_{t} \overline{\boldsymbol{X}}(t)-\mathbb{E}\left[\widetilde{\boldsymbol{X}}^{k}(0)\right]+\boldsymbol{\Sigma}_{k}^{\frac{1}{2}}(t) \boldsymbol{W}^{k}(t) \\
& =\widetilde{\boldsymbol{X}}^{k}(0)+\sum_{i=1}^{K p} c_{i} \mu_{i} \boldsymbol{u}_{i}^{k} \mathrm{e}^{\mu_{i} t}-\sum_{i=1}^{K p} c_{i} \boldsymbol{u}_{i}^{k}+\boldsymbol{\Sigma}_{k}^{\frac{1}{2}} \boldsymbol{W}^{k}(t), \tag{C.12}
\end{align*}
$$

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

$$
\begin{equation*}
\left\langle\widetilde{\boldsymbol{X}}^{k}(t)-\widetilde{\boldsymbol{X}}^{l}(t), \boldsymbol{\nu}\right\rangle>0 \tag{C.13}
\end{equation*}
$$

with probability ${ }^{4}$ tending to 1 as $t \rightarrow \infty$ for any two classes $k \neq l$. Substituting equation (C.12), we have equivalently

$$
\begin{equation*}
\left\langle\sum_{i=1}^{K p} c_{i}\left(\boldsymbol{u}_{i}^{k}-\boldsymbol{u}_{i}^{l}\right)\left(\mu_{i} \mathrm{e}^{\mu_{i} t}-1\right), \boldsymbol{\nu}\right\rangle>\left\langle\widetilde{\boldsymbol{X}}^{l}(0)-\widetilde{\boldsymbol{X}}^{k}(0), \boldsymbol{\nu}\right\rangle+\left\langle\boldsymbol{\Sigma}_{l}^{\frac{1}{2}} \boldsymbol{W}^{l}(t)-\boldsymbol{\Sigma}_{k}^{\frac{1}{2}} \boldsymbol{W}^{k}(t), \boldsymbol{\nu}\right\rangle \tag{C.14}
\end{equation*}
$$

By the Gaussian tail bound, the right-hand side of the above display is $O_{\mathbb{P}}\left(C_{0}+C_{1} \sigma_{\max } \sqrt{t}\right)=$ $O_{\mathbb{P}}(\sqrt{t})$ where $C_{0}$ and $C_{1}$ are constants that do not depend on class labels and we assume $\sigma_{\max }=$ $\max _{k \in[K]} \sup _{t}\left\|\boldsymbol{\Sigma}_{k}(t)\right\|_{2}$ is finite. Due to randomization in the initialization, with probability zero $\overline{\boldsymbol{X}}^{k}(0)=\overline{\boldsymbol{X}}^{l}(0)$ and thus

$$
\begin{equation*}
\sum_{i=1}^{K p} c_{i}\left(\boldsymbol{u}_{i}^{k}-\boldsymbol{u}_{i}^{l}\right) \neq \mathbf{0} \tag{C.15}
\end{equation*}
$$

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

$$
\begin{equation*}
c_{j}\left(\boldsymbol{u}_{j}^{k}-\boldsymbol{u}_{j}^{l}\right)\left(\mu_{j} \mathrm{e}^{\mu_{j} t}-1\right) \neq 0 \tag{C.16}
\end{equation*}
$$

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

$$
\begin{equation*}
\exp \left\{\mu_{j} t\right\}>C \sqrt{t} \tag{C.17}
\end{equation*}
$$

holds for some constant $C$ that only depends on $C_{0}, C_{1}$, and $\sigma_{\max }$. 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 $\int_{0}^{t} \mu_{j}(s) \mathrm{d} s$, i.e.,

$$
\begin{equation*}
\exp \left\{\int_{0}^{t} \mu_{j}(s) \mathrm{d} s\right\}>C \sqrt{t} \tag{C.18}
\end{equation*}
$$

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

$$
\begin{equation*}
\exp \{\Gamma(t)\}>C^{\prime} \sqrt{t} \tag{C.19}
\end{equation*}
$$

${ }^{4}$ Here the randomness comes from the dynamics as well as the random sample $\widetilde{\boldsymbol{X}}^{k}$.
for some constant $C^{\prime}$ that only depends on $C_{0}, C_{1}, \sigma_{\max }$, and the specific choice of $\boldsymbol{H}$. Note that to prove equation (C.19), it suffices to require $\gamma(t)$ has a tail that is at least $1 / t$, or

$$
\begin{equation*}
\gamma(t)=\omega\left(\frac{1}{t}\right) \tag{C.20}
\end{equation*}
$$

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

$$
\begin{equation*}
\gamma(t)=o\left(\frac{1}{t}\right) \tag{C.21}
\end{equation*}
$$

as $t \rightarrow \infty$, the probability of separation tends to zero as long as $n \rightarrow \infty$ at an arbitrary rat ${ }^{5}$. 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 $\boldsymbol{\nu}=\boldsymbol{\nu}(k, l)$ exists that ensures separation. We remark that it is possible to remove class-dependence on $\nu$ in our case, and therefore achieve a stronger sense of separability: consider the $(K(K-1) / 2)$-by- $K$ matrix $\Psi$ with $\boldsymbol{u}_{i(k, l)}^{k}-\boldsymbol{u}_{i(k, l)}^{l}$ 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 $\boldsymbol{\nu}$ is equivalent to $\Psi \boldsymbol{\nu} \neq \mathbf{0}$. This is equivalent to that the nullity of $\Psi$ is less than $K$, which is obvious since $\operatorname{rank} \Psi \geq 1$ almost surely by construction.

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 $W_{t}$,

$$
\begin{equation*}
\limsup _{t \rightarrow \infty} \frac{W_{t}}{\sqrt{2 t \log \log t}}=1, \quad \liminf _{t \rightarrow \infty} \frac{W_{t}}{\sqrt{2 t \log \log t}}=-1 \tag{C.22}
\end{equation*}
$$

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 $\left\{\overline{\boldsymbol{X}}^{k}(t) /\left\|\overline{\boldsymbol{X}}^{k}(t)\right\|\right\}_{k=1}^{K}$ forms an ETF as $t \rightarrow \infty$.

[^1]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 $C_{1} \boldsymbol{d}_{k} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}$. As $t \rightarrow \infty$, the unit vector in the direction of $\bar{X}^{k}$ tends to $\boldsymbol{d}_{k}$, and therefore those unit vectors form an ETF, since $\left\{d_{j}\right\}_{j=1}^{K}$ form an ETF.

## 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

$$
\max _{\|\boldsymbol{\nu}\|_{2}=1} \min _{k \neq l \in[K]}\left|\left\langle\boldsymbol{c}_{k}-\boldsymbol{c}_{l}, \boldsymbol{\nu}\right\rangle\right| .
$$

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 $\min _{k \neq l \in[K]}\left|\left\langle\boldsymbol{c}_{k}-\boldsymbol{c}_{l}, \boldsymbol{\nu}\right\rangle\right|$.

Remark. In practice, estimating $\boldsymbol{c}_{0}+\boldsymbol{c}_{l}$ 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 $\left[T_{1}, T_{2}\right]$ with $T_{2}>T_{1} \gg 0$ when the model is almost convergent to estimate $\boldsymbol{\nu}$ as follows:

$$
\begin{equation*}
\max _{\|\boldsymbol{\nu}\|_{2} \leq 1} \min _{k<l} \sum_{t=T_{1}}^{T_{2}}\left|\left\langle\bar{X}^{k}(t)-\bar{X}^{l}(t), \boldsymbol{\nu}\right\rangle\right| . \tag{C.23}
\end{equation*}
$$

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

$$
\begin{equation*}
\boldsymbol{\nu}=\frac{\overline{\boldsymbol{X}}^{k}(T)-\overline{\boldsymbol{X}}^{l}(T)}{\left\|\overline{\boldsymbol{X}}^{k}(T)-\overline{\boldsymbol{X}}^{l}(T)\right\|} \tag{C.24}
\end{equation*}
$$

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 $n_{\mathrm{val}}=\{20,120,400,1200\}$. We also pollute each label class by randomly choosing $p_{\mathrm{err}} \cdot n_{\mathrm{tr}}$ samples to flip the label to another class. (uniform across all other classes). In this setup, we fix $n_{\mathrm{tr}}=n_{\mathrm{val}}=500$ and set $p_{\mathrm{err}} \in\{0.1,0.2, \ldots, 0.8\}$. 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 $r_{\alpha}$ and $r_{\beta}$.
D.1.1 Estimation of Integrated Local Elasticity Strengths $\widehat{A}(t)$ and $\widehat{B}(t)$

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

$$
\begin{align*}
& \text { (l-model) }\left\{\begin{array}{l}
\widehat{A}(t)=\operatorname{avg~avg}_{k} \log \left|\frac{\breve{\boldsymbol{X}}\left(\overline{\boldsymbol{X}}^{k}-\widetilde{\boldsymbol{X}}\right)^{K-1}}{\boldsymbol{c}_{0} \boldsymbol{c}_{k}^{K-1}}\right|, \quad \overline{\boldsymbol{X}}_{t}:=\operatorname{avg}_{l} \overline{\boldsymbol{X}}_{t}^{l}, \\
\widehat{B}(t)=-\operatorname{avg} \operatorname{avg}_{k} \log \left|\frac{\frac{c}{0}^{\boldsymbol{c}_{k}} \frac{\overline{\boldsymbol{X}}^{k}-\widetilde{\boldsymbol{X}}}{}}{\tilde{\boldsymbol{X}}}\right|,
\end{array}\right.  \tag{D.1}\\
& (\text { L-model }) \quad\left\{\left.\begin{array}{ll}
\widehat{A}(t) & =A^{\prime}(t)+2 B^{\prime}(t), \\
\widehat{B}(t) & =2\left(B^{\prime}(t)-A^{\prime}(t)\right),
\end{array} \quad\left\{\begin{array}{ll}
A^{\prime}(t) & :=\log \\
B^{\prime}(t) & :=\log
\end{array}\left|\left\langle\overline{\boldsymbol{X}}^{\top} \boldsymbol{v}_{1}-1\right\rangle\right|, \quad \overline{\boldsymbol{X}}^{\top}\left(\boldsymbol{v}_{2}-\frac{4}{3} \boldsymbol{v}_{1}\right)\right\rangle \right\rvert\,,\right.
\end{align*}
$$

where

$$
\begin{equation*}
\boldsymbol{v}_{1}=\frac{1}{4}[1,-1,-1,-1,1,-1,-2,-2,0]^{\top}, \quad \boldsymbol{v}_{2}=\frac{1}{3}[2,-1,-1,-1,2,-1,0,0,0]^{\top} \tag{D.2}
\end{equation*}
$$

We will now explain how it is done.
Estimation in I-model. Recall from Proposition 3.2, the per-class means $\overline{\boldsymbol{X}}_{t}$ solve the LE-ODE (7) under the I-model as

$$
\begin{equation*}
\overline{\boldsymbol{X}}(t)=\boldsymbol{c} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}+\left(\mathbf{1}_{K} \otimes \boldsymbol{c}_{0}\right) \mathrm{e}^{\frac{1}{K} A(t)+\frac{K-1}{K} B(t)}, \tag{D.3}
\end{equation*}
$$

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

$$
\begin{equation*}
\overline{\boldsymbol{X}}_{t}:=\operatorname{avg}_{l} \overline{\boldsymbol{X}}_{t}^{l}=\frac{1}{K} \sum_{l=1}^{K} \overline{\boldsymbol{X}}_{t}^{l}=\boldsymbol{c}_{0} \mathrm{e}^{\frac{1}{K} A(t)+\frac{K-1}{K} B(t)} \in \mathbb{R}^{K} \tag{D.4}
\end{equation*}
$$

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

$$
\begin{equation*}
\overline{\boldsymbol{X}}_{t}^{k}-\overline{\boldsymbol{X}}_{t}=\boldsymbol{c}_{k} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)} \tag{D.5}
\end{equation*}
$$

It follows that

$$
\begin{equation*}
\left|\frac{\boldsymbol{c}_{0}}{\boldsymbol{c}_{k}} \frac{\overline{\boldsymbol{X}}^{k}-\overline{\boldsymbol{X}}}{\overline{\boldsymbol{X}}}\right|=\left|\mathbf{1}_{K} \mathrm{e}^{B(t)}\right|, \quad\left|\frac{\boldsymbol{c}_{0}}{\boldsymbol{c}_{k}} \frac{\overline{\boldsymbol{X}}^{k}-\overline{\boldsymbol{X}}}{\overline{\boldsymbol{X}}}\right|=\left|\mathbf{1}_{K} \mathrm{e}^{A(t)}\right| \tag{D.6}
\end{equation*}
$$

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

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

$$
\begin{equation*}
\overline{\boldsymbol{X}}(t)=\boldsymbol{c}_{0}+C_{1} \boldsymbol{d} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}+\left(\sum_{l=1}^{K-1} C_{2 l} \boldsymbol{f}_{l}\right) \mathrm{e}^{\frac{1}{K} A(t)+\frac{1}{K(K-1)} B(t)} \tag{D.7}
\end{equation*}
$$

where $\boldsymbol{c}_{0}$ is a constant vector with $K(K-1)$ free parameters; $\boldsymbol{f}_{l}$ '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

$$
\begin{equation*}
\boldsymbol{v}_{1}=\frac{1}{4}[1,-1,-1,-1,1,-1,-2,-2,0]^{\top}, \quad \boldsymbol{v}_{2}=\frac{1}{3}[2,-1,-1,-1,2,-1,0,0,0]^{\top} \tag{D.8}
\end{equation*}
$$

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

$$
\begin{equation*}
\boldsymbol{v}_{1}^{\top} \boldsymbol{f}_{j}=0, \quad j=1,2, \quad \boldsymbol{v}_{2}^{\top} \boldsymbol{c}_{0}=0, \quad \mathbf{1}_{K}^{\top} \boldsymbol{d}=0 \tag{D.9}
\end{equation*}
$$

and

$$
\begin{equation*}
\boldsymbol{v}_{1}^{\top} \boldsymbol{c}_{0}=1-\frac{w_{1}+w_{2}+w_{3}}{4}, \quad \boldsymbol{v}_{1}^{\top} \boldsymbol{d}=1, \quad \boldsymbol{v}_{2}^{\top} \boldsymbol{f}_{j}=1, \quad \boldsymbol{v}_{2}^{\top} \boldsymbol{d}=\frac{4}{3} \tag{D.10}
\end{equation*}
$$

Recall that we define $\langle\boldsymbol{X}(t)\rangle:=\boldsymbol{X}(t) / \boldsymbol{X}(0)$, 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

$$
\left\{\begin{array}{l}
\overline{\boldsymbol{X}}_{t}^{\top} \boldsymbol{v}_{1}=C_{1} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}+1-\frac{C_{3}}{4},  \tag{D.11}\\
\overline{\boldsymbol{X}}_{t}^{\top} \boldsymbol{v}_{2}=\frac{3 C_{1}}{4} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}+\left(C_{21}+C_{22}\right) \mathrm{e}^{\frac{1}{K} A(t)+\frac{1}{K(K-1)} B(t)}
\end{array}\right.
$$



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

$$
\begin{cases}\left\langle\overline{\boldsymbol{X}}_{t}^{\top} \boldsymbol{v}_{1}-1\right\rangle & =\frac{C_{1} \mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)}-\frac{C_{3}}{4}}{C_{1}-\frac{C_{3}}{4}}  \tag{D.12}\\ \left\langle\overline{\boldsymbol{X}}_{t}^{\top}\left(\boldsymbol{v}_{2}-\frac{4}{3} \boldsymbol{v}_{1}\right)\right\rangle & =\frac{\left(C_{21}+C_{22}\right) \mathrm{e}^{\frac{\mathrm{e}^{K}}{} A(t)+\frac{1}{K(K-1)}} B(t)}{C_{21}+C_{22}+\frac{3 C_{3}}{16}}\end{cases}
$$

We define

$$
\begin{cases}A^{\prime}(t) & =\log \left|\left\langle\overline{\boldsymbol{X}}_{t}^{\top} \boldsymbol{v}_{1}-1\right\rangle\right|  \tag{D.13}\\ B^{\prime}(t) & =\log \left|\left\langle\overline{\boldsymbol{X}}_{t}^{\top}\left(\boldsymbol{v}_{2}-\frac{4}{3} \boldsymbol{v}_{1}\right)\right\rangle\right|\end{cases}
$$

when $t$ is sufficiently large such that

$$
\begin{equation*}
\mathrm{e}^{\frac{1}{K} A(t)-\frac{1}{K} B(t)} \gg \frac{C_{3}}{4 C_{1}}, \quad \mathrm{e}^{\frac{1}{K} A(t)+\frac{1}{K(K-1)} B(t)} \gg \frac{3 C_{3}}{16\left(C_{21}+C_{22}\right)}, \tag{D.14}
\end{equation*}
$$

we have approximately

$$
\begin{equation*}
A^{\prime}(t) \approx \frac{1}{K} A(t)-\frac{1}{K} B(t), \quad B^{\prime}(t) \approx \frac{1}{K} A(t)+\frac{1}{K(K-1)} B(t) \tag{D.15}
\end{equation*}
$$

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

$$
\begin{equation*}
\widehat{A}(t)=A^{\prime}(t)+2 B^{\prime}(t), \quad \widehat{B}(t)=2\left(B^{\prime}(t)-A^{\prime}(t)\right) \tag{D.16}
\end{equation*}
$$

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 $\widehat{A}(t)$ and $\widehat{B}(t)$, we may numerically differentiate these estimates to obtain $\widehat{\alpha}(t)$ and $\widehat{\beta}(t)$. 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 $\omega$ in this filter which roughly corresponds to the window size in moving averages: a smaller $\omega$ preserves more fluctuations in the original data and a larger $\omega$ smooths the data more. As a rule of thumb, we test on a set of different values for $\omega$ 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 $\widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ trained on GEOMNIST with $p_{\mathrm{err}}=0$ and $n_{\mathrm{tr}}=3000$ 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 $\omega$ 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

$$
\begin{equation*}
\alpha(t) \sim \frac{\alpha_{0}}{(1+t)^{r}} \tag{D.17}
\end{equation*}
$$

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

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



Figure D.2: Estimated tail indices versus label corruption ratio $p_{\text {err }}$. The tail indices are estimated using equation (D.19) with $\widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ 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 $\widehat{\beta}(t)$ begins to dominate that of $\widehat{\alpha}(t)$.

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

$$
\begin{equation*}
\frac{\log A(t)}{\log (1+r)}=(1-r)+\frac{\log \alpha_{0}}{\log (1+t)}-\frac{\log (1-r)}{\log (1+t)} \approx 1-r \tag{D.19}
\end{equation*}
$$

and $\widehat{r}_{\alpha}=1-\operatorname{avg}_{t \geq T_{0}} \log A(t) / \log (1+t)$ for some sufficiently large $T_{0}$. We estimate $\widehat{r}_{\beta}$ 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 $\widehat{A}(t), \widehat{B}(t), \widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ versus training time $t$ on both GEOMNIST and CIFAR under both I -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 I-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 $\widehat{A}(t), \widehat{B}(t), \widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ versus training time $t$ on CIFAR under the I-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 $p_{\text {err }}$. 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 $p_{\text {err }}$ 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 $p_{\text {err }}$ using $\alpha(t)$ and $\beta(t)$ estimated under both I-model and L-model in Figure D.2. Although in Figure D.2b there is no sharp phase transition boundary when we increase $p_{\text {err }}$ as Figure D. 2 b does, we observe that at around $p_{\text {err }}=2 / 3$, the estimated tail index of $\widehat{\beta}(t)$ begins to dominate that of $\widehat{\alpha}(t)$, which also supports our Theorem 3.1

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


Figure D.3: Effect of per-class sample size $n$. Estimated $\widehat{A}(t), \widehat{B}(t)$ (the first and the third rows) and $\widehat{\alpha}(t), \widehat{\beta}(t)$ (the second and the fourth rows), under I-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 $\widehat{A}(t), \widehat{B}(t)(($ a) and (c)) and $\widehat{\alpha}(t), \widehat{\beta}(t)(($ b) and (d)) on CIFAR with $K=2((\mathrm{a})-(\mathrm{b}))$ and $K=3((\mathrm{c})-(\mathrm{d}))$, all under I-model. Note that the general trends with different $K$ 's are similar.


Figure D.5: Effect of label corruption ratio $p_{\text {err }}$. Estimated $\widehat{A}(t), \widehat{B}(t)$ (the first and the third rows) and $\widehat{\alpha}(t), \widehat{\beta}(t)$ (the second and the fourth rows), under I-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 $\widehat{A}(t)$ and $\widehat{B}(t), \widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ become indistinguishable when $p_{\text {err }}$ is large, supporting Theorem 3.1
local elasticity strengths $\alpha(t)$ and $\beta(t)$, when adjusted for the evolution of the norm of $\overline{\boldsymbol{X}}^{k}(t)$. As a surrogate, we simply multiply the strengths by the average norm, $\operatorname{avg}_{k}\left\|\overline{\boldsymbol{X}}^{k}(t)\right\|_{2}$ 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 $p_{\text {err }}$, similarly to the unadjusted cases Figure D. 5 .

Simulations of the LE-ODE. With $\alpha(t)$ and $\beta(t)$ estimated (using either the I-model or the Lmodel), 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 $\boldsymbol{\zeta}^{K} \sim \mathcal{N}_{K}\left(\mathbf{0}, \sigma_{k} \boldsymbol{I}_{K^{2}}\right)$ for all $k \in[K]$ with $\sigma_{k}=\left\|\overline{\boldsymbol{X}}^{k}(0)\right\|_{2} / \sqrt{K}$ with $\overline{\boldsymbol{X}}^{k}(0)$ 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-modelwhere the captions indicate which model (I-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 $p_{\text {err }}$ (adjusted for logits norm). Estimated $\widehat{\alpha}(t), \widehat{\beta}(t)$ 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 $\widehat{A}(t)$ and $\widehat{B}(t), \widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ become indistinguishable when $p_{\text {err }}$ is large, supporting Theorem 3.1.


Figure D.7: Simulated LE-ODE solutions under the l-model versus genuine dynamics. We use $\widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ 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 I-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 I-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 I-model fails to identify the correct class. After manual alignment, we note the I-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 oblivious to incorrect classes. On the other hand, simulated trajectories from the L-model faithfully recover the trajectories of correct classes $\operatorname{argmax}_{l \in[K]} \bar{X}_{l}^{k}$ 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., $\boldsymbol{H}_{i, j}$


Figure D.8: Simulated LE-ODE paths under the L-model versus genuine dynamics. The simulation is done under the L-model with estimated $\widehat{\alpha}(t)$ and $\widehat{\beta}(t)$ 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 $k, l \in[3]$.


Figure D.9: Relative difference $\mathbf{R D}_{k}$ 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 I-model throughout training and better captures the later stages of the training (indicated by decreasing RD), supporing our discussions in Appendix B.2.
is set to be $\boldsymbol{d}_{j} \boldsymbol{d}_{j}^{\top} / \boldsymbol{d}_{j}^{\top} \boldsymbol{d}_{j}$ 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

$$
\begin{equation*}
\boldsymbol{H}_{i, j}=p_{j} \frac{\boldsymbol{d}_{j} \boldsymbol{d}_{j}^{\top}}{\boldsymbol{d}_{j}^{\top} \boldsymbol{d}_{j}}+\sum_{l \neq j}^{K} p_{l} \frac{\boldsymbol{d}_{l} \boldsymbol{d}_{l}^{\top}}{\boldsymbol{d}_{l}^{\top} \boldsymbol{d}_{l}} \tag{D.20}
\end{equation*}
$$

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

$$
\begin{equation*}
\operatorname{RD}_{\mathrm{k}}(t):=\frac{\left\|\overline{\boldsymbol{X}}^{k}(t)-\overline{\boldsymbol{Y}}^{k}(t)\right\|_{\boldsymbol{H}^{k}}}{\left(\left\|\overline{\boldsymbol{X}}^{k}(t)\right\|_{2}+\left\|\overline{\boldsymbol{Y}}^{k}(t)\right\|_{2}\right) / 2} \tag{D.21}
\end{equation*}
$$

where $\overline{\boldsymbol{X}}(t)$ and $\overline{\boldsymbol{Y}}(t)$ are genuine and simulated trajectories, respectively, and $\|\cdot\|_{\boldsymbol{H}}$ denotes the norm induced by the matrix $\boldsymbol{H}$ that is used to define models (i.e., the identity matrix for the l-model and $\overline{\boldsymbol{H}}$ in equation 8 for the L-model). This choice normalizes the difference under the similarity defined by $\boldsymbol{H}$ (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 I-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 $\operatorname{avg}_{k \in[K]}\left\|\overline{\boldsymbol{X}}^{k}\right\|_{2}$ 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 derandomization 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


Figure D.10: The hyperplanes formed by $\left\{\bar{X}_{t}^{k}\right\}_{k=1}^{3}$ with colors corresponding to iterations. (a) In the first 1200 iterations, the validation loss does not drop perceivably, 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.
spanned by the per-class mean logits vectors, $\overline{\boldsymbol{X}}^{k}(t) \in \mathbb{R}^{3}$ for $k \in[3]$. In Figure D.10a we plot those 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 $\boldsymbol{E}$ 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 $\boldsymbol{E}$ as $(\alpha(t)-\beta(t)) \boldsymbol{I}_{K}+$ $\beta(t) \mathbf{1}_{K} \mathbf{1}_{K}^{\top}$. 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 $\boldsymbol{E}$ 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 $\left(\boldsymbol{E} \otimes \boldsymbol{I}_{K}\right) \circ \boldsymbol{H}$ are bounded within $\lambda_{\min }(\boldsymbol{E}) \min _{i} H_{i i}$ and $\lambda_{\max }(\boldsymbol{E}) \max _{i} H_{i i}$, where $\left\{H_{i i}, i \in[K p]\right\}$ is the diagonal entries of the matrix $\boldsymbol{H}$, hence we expect a very similar result in this case as in Theorem 3.1. In the more general case where $\boldsymbol{E}$ is symmetric but not necessarily semi-definite, a more precise analysis on the spectrum of $\left(\boldsymbol{E} \otimes \boldsymbol{I}_{K}\right) \circ \boldsymbol{H}$ 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 $\boldsymbol{M}_{t}$ in equation (5) takes the form of $E_{k, l} \boldsymbol{H}_{k, l} / K$, 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 $E_{k, l} \boldsymbol{H}_{k, l} \cdot p_{l}$ for $\sum_{l} p_{l}=1$, where $p_{l}$ is the probability of class $l$ being sampled. More succinctly, instead of defining $\boldsymbol{M}_{t}=\left(\boldsymbol{E}_{t} \otimes \boldsymbol{I}_{K}\right) \circ \boldsymbol{H} / K$, we let $\boldsymbol{M}_{t}=\left(\boldsymbol{E}_{t} \otimes \boldsymbol{P}\right) \circ \boldsymbol{H} / K$ for a $K$-by- $K$ doubly stochastic matrix $\boldsymbol{P}$ that models this sampling effect. In the same vein, we can also model the case when the data are polluted by corrupted labels. Let $\boldsymbol{P}=\left(p_{k, l}\right)_{k, l} \in \mathbb{R}^{K \times K}$ 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 $\boldsymbol{P}$ 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 I-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 $\boldsymbol{H}$ matrix, whose $(i, j)$-th block is defined as $\overline{\boldsymbol{H}}^{j}=\boldsymbol{d}_{j} \boldsymbol{d}_{j}^{\top} / \boldsymbol{d}_{j}^{\top} \boldsymbol{d}_{j}$ where $\boldsymbol{d}_{j}=\boldsymbol{e}_{j}-\mathbf{1}_{K} / K-$ 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 $\boldsymbol{H}$ as

$$
\begin{equation*}
\boldsymbol{H}_{i, j}=p_{j}^{j} \overline{\boldsymbol{H}}^{j}+\sum_{l \neq j}^{K} p_{l}^{j} \overline{\boldsymbol{H}}^{l}, \quad p_{j}^{j}>p_{l}^{j}, \quad \sum_{l=1}^{K} p_{l}^{j}=1, \quad \forall l \neq j, \quad j \in[K] . \tag{E.1}
\end{equation*}
$$

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 derandomization 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.


[^0]:    ${ }^{3}$ Here, $\gamma(t)=\omega(1 / t)$ 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 $t \rightarrow \infty$.

[^1]:    ${ }^{5}$ Note 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 $2 n$ 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 $\frac{2}{\binom{2 n}{n}}>0$.

