arXiv:2512.04006v1 [cs.LG] 3 Dec 2025

DIAGONALIZING THE SOFTMAX: HADAMARD INITIALIZATION FOR TRACTABLE CROSS-ENTROPY DYNAMICS

Connall Garrod
Mathematical Institute
University of Oxford
connall.garrod@maths.ox.ac.uk

Jonathan P. Keating
Mathematical Institute
University of Oxford

Christos Thrampoulidis
Department of Electrical and Computer Engineering
University of British Columbia
cthrampo@ece.ubc.ca

ABSTRACT

Cross-entropy (CE) training loss dominates deep learning practice, yet existing theory often relies on simplifications—either replacing it with squared loss or restricting to convex models—that miss essential behavior. CE and squared loss generate fundamentally different dynamics, and convex linear models cannot capture the complexities of non-convex optimization. We provide an in-depth characterization of multi-class CE optimization dynamics beyond the convex regime by analyzing a canonical two-layer linear neural network with standard-basis vectors as inputs—the simplest non-convex extension for which the implicit bias remained unknown. This model coincides with the unconstrained features model used to study neural collapse, making our work the first to prove that gradient flow on CE converges to the neural collapse geometry. We construct an explicit Lyapunov function that establishes global convergence, despite the presence of spurious critical points in the non-convex landscape. A key insight underlying our analysis is an inconspicuous finding: Hadamard Initialization diagonalizes the softmax operator, freezing the singular vectors of the weight matrices and reducing the dynamics entirely to their singular values. This technique opens a pathway for analyzing CE training dynamics well beyond our specific setting considered here.

1 Introduction

1.1 Motivation

Modern overparameterized classification networks exhibit a remarkable geometric structure at the end of training, known as neural collapse (NC) [1]: class features collapse to their respective means, those means form a simplex equiangular tight frame (ETF), and the classifier aligns with this frame. This phenomenon appears consistently across diverse architectures and datasets, despite the absence of any explicit mechanism in the training procedure that enforces it. This raises a central question: *What implicit mechanism in gradient-based optimization produces this precise geometric configuration among the infinitely many solutions that achieve zero training loss?*

This question lies squarely within the implicit bias framework, which seeks to characterize the specific solutions that gradient descent selects when multiple optima exist. For logistic regression, Soudry et al. [2] established that gradient descent converges to the max-margin solution. Subsequent research extended this result in several directions, most notably to deep homogeneous networks [3, 4, 5]. These works prove convergence of gradient descent—or its infinitesimal limit, gradient flow—to KKT points of corresponding non-convex max-margin problems. Yet, in contrast to the convex case, these results are comparatively less explicit: they do not specify *which* KKT point is reached.

For NC, Ji et al. [6] and Thrampoulidis et al. [7] applied this framework to the unconstrained features model (UFM), a theoretical proxy for studying representation geometries in deep networks. They showed that cross-entropy (CE)

gradient flow converges to a KKT point of a max-margin problem whose global minimizer corresponds to NC. Critical gaps remain, however: these results guarantee only convergence to *some* KKT point—which may be a saddle rather than NC—and provide no characterization of the training trajectory or the convergence rate. More broadly, existing implicit bias results remain restricted to three narrow regimes: (i) logistic regression, which lacks depth; (ii) binary classification, which omits multi-class interactions; and (iii) KKT analyses, which cannot distinguish between NC and other critical points. This gap is particularly significant for unregularized training, where the global minimum is attained only at infinity: analyses based solely on asymptotic limits may overlook crucial finite-time dynamics, such as trajectories remaining near non-NC saddles indefinitely or exhibiting non-monotonic behavior that initially moves away from NC. Understanding implicit bias in this setting requires analyzing the *entire trajectory*, not just its limit.

A complementary line of work, pioneered by Saxe et al. [8, 9] and extended by numerous follow-ups, provides tools for such trajectory-level analysis, but only for mean squared error (MSE) loss. Under special initializations that restrict dynamics to singular values, they derived exact solutions for linear networks, showing that singular values evolve independently along sigmoidal curves. This framework was recently applied to the UFM with MSE loss [10]. However, NC is observed in classification networks that use CE loss, which introduces a softmax nonlinearity and couples singular value dynamics in ways that prevent exact solutions. Whether the insights and techniques from MSE can extend to CE—and whether they can rigorously characterize the implicit bias toward NC—has remained an open question.

1.2 Contributions

We provide the first rigorous proof of convergence to NC under multi-class CE training in a non-convex setting. Our key technical innovation is showing that Hadamard Initialization diagonalizes the softmax operator in the UFM. Concretely, for a normalized Sylvester-Hadamard matrix U and any vector a, we prove that:

$$P(U \operatorname{diag}(a) U^T) = U \operatorname{diag}(UP(Ua)) U^T ,$$

where $P(\cdot)$ is the column-wise softmax map. This property enables us to extend the Saxe et al. framework from MSE to CE, reducing its gradient flow to a vector ODE over the $K-1$ non-trivial singular values a of the logit matrix:

$$\frac{da}{dt} = a \odot \left(1 - \tilde{U}^T P(\tilde{U}a)\right) ,$$

where $\odot$ is the entrywise product and $\tilde{U}$ represents the Hadamard modes excluding the trivial constant mode. Unlike the MSE setting, this reduced ODE retains the complex couplings of the softmax: singular values do not evolve independently. Nevertheless, we show the dynamics are sufficiently structured to allow rigorous implicit bias analysis.

Non-monotonic convergence. Prior work has numerically observed monotonic convergence to NC under CE, but we show that this is *not* true in general: convergence can be non-monotonic, and the behavior depends sensitively on the number of classes K. For $K \in \{2, 4\}$, natural distance metrics decrease monotonically, whereas for $K \geq 8$ we construct explicit initializations where these metrics initially *increase* before eventually decreasing. Even more strikingly, we construct metrics that are monotonic even for $K = 8$ but break down for $K \geq 16$. These phenomena are invisible to KKT-based analyses and reveal a subtle, previously unrecognized K-dependence in multi-class dynamics. This complexity is missed by prior implicit bias work, which has focused almost exclusively on the binary case.

Guaranteed convergence to NC via Lyapunov function. Despite the non-monotonicity of natural metrics and the presence of competing saddle points, we prove that gradient flow *always* converges to NC by constructing an explicit Lyapunov function—the KL divergence—that decreases monotonically throughout training. This provides the first rigorous characterization of implicit bias toward NC beyond binary classification. Our results demonstrate that NC is not merely an empirical observation but the inevitable outcome of gradient-based optimization in this setting.

Exponentially slow convergence. As the logit norm grows, the CE dynamics slow down exponentially, causing trajectories with unfavorable initializations to linger near non-NC saddle points for arbitrarily long periods. This plateau behavior is practically significant yet invisible to purely asymptotic analyses.

Role of initialization. To reconcile our non-monotonic examples with the monotonic behavior observed in practice, we show that standard near-zero random initialization schemes naturally bias the dynamics toward the simplex ETF subspace, steering trajectories away from problematic saddle regions. This explains why non-monotonic behavior, while theoretically possible, is rarely observed.

MSE vs CE comparison. To highlight the impact of softmax nonlinearity, we show that MSE and CE dynamics coincide only under restrictive conditions: either in the linearized regime (small norms) or under direct logit optimization (removing depth). Outside these settings, CE exhibits coupled dynamics, non-monotonicity, and prolonged saddle plateaus that do not arise under MSE.

Our theoretical findings are supported by numerical experiments.

Although our analysis focuses on Hadamard Initialization in the unregularized, balanced UFM, the approach extends naturally to deeper networks, regularization, and imbalanced classes—directions we leave for future work. Overall, our results provide a tractable framework for understanding optimization dynamics in non-convex multi-class CE settings, and we hope they motivate the development of similar analytical techniques beyond the UFM.

2 Background

We consider a classification task with K classes and n samples per class. The ith sample of class c is denoted by $x_{ic} \in \mathbb{R}^{d_0}$, with corresponding one-hot encoded label $y_c \in \mathbb{R}^K$. A deep neural network $f_\theta(x) : \mathbb{R}^{d_0} \to \mathbb{R}^K$ models the relationship between inputs and labels. We decompose the network as $f_\theta(x) = W h_{\bar{\theta}}(x)$, where $W \in \mathbb{R}^{K \times d}$ are the last layer weights and $h_{\bar{\theta}} : \mathbb{R}^{d_0} \to \mathbb{R}^d$ is a highly expressive feature map representing the remainder of the network. The parameters $\theta = \{W, \bar{\theta}\}$ are trained by a variant of gradient descent on the CE loss function.

We denote the image of the data under the map $h_{\bar{\theta}}$ by $h_{\bar{\theta}}(x_{ic}) = h_{ic}$, and define the feature matrix H as the collection of all features in class order: $H = [h_{1,1}, ..., h_{n,1}, h_{1,2}, ..., h_{n,K}] \in \mathbb{R}^{d \times Kn}$. Define the class feature means as $\mu_c = \text{Av}_i\{h_{ic}\}$ and the feature mean matrix as $\bar{H} = [\mu_1, \ldots, \mu_K]$. We also define the simplex ETF matrix as

$$S = I_K - \frac{1}{K} 1_K 1_K^T \in \mathbb{R}^{K \times K}.$$

NC then refers to the following phenomenon observed in overparameterized networks as training continues.

Definition 1 (**Neural Collapse**). *A neural network exhibits NC if the final-layer weights and feature matrix satisfy:*

- ***NC1****: Feature vectors collapse onto their class means:* $h_{ic} = \mu_c$.
- ***NC2****: The class means align with the simplex ETF:* $\bar{H}^T \bar{H} \propto S$.
- ***NC3****: The last-layer weights align with the simplex ETF:* $W W^T \propto S$.

The original work by Papyan et al. [1] included a degree of freedom coming from global centering, but in the context of our model this can be dropped from the definition without impacting results. Our analysis focuses on the non-trivial NC2 and NC3 properties, which dictate the geometry of last-layer weights and features.

Unconstrained Features Models. The UFM is a proxy used to study the NC phenomenon. To formally define it, we approximate the feature map $h_{\bar{\theta}}(x)$ as being capable of mapping the training data to arbitrary points in feature space, treating the feature vectors h_{ic} as freely optimized variables. Using the CE loss function with no explicit regularization, our objective becomes to minimize jointly over $W \in \mathbb{R}^{K \times d}$ and $H \in \mathbb{R}^{d \times Kn}$:

$$\mathcal{L}(W, H) = \sum_{c=1}^{K} \sum_{i=1}^{n} \log \Big(1 + \sum_{c' \neq c}^{K} e^{(W h_{ic})_{c'} - (W h_{ic})_c} \Big). \tag{1}$$

Under gradient flow, the parameters evolve according to

$$\frac{dW}{dt} = (Y - P(WH)) H^T, \qquad \frac{dH}{dt} = W^T (Y - P(WH)), \tag{2}$$

where $Y = I_K \otimes 1_n^T$, with $\otimes$ denoting the Kronecker product, and $P(WH)$ is the softmax operation applied columnwise to the logit matrix $Z = WH$ of the network, defined as

$$\big[P(Z)\big]_{ij} = \frac{\exp(Z_{ij})}{\sum_{i'} \exp(Z_{i'j})}.$$

Using the analysis of [3], Ji et al. [6] showed that this gradient flow converges in direction to a KKT point of the following non-convex max-margin problem:

$$(\hat{W}, \hat{H}) \in \text{argmin}_{W,H} \Big\{ \frac{1}{2} \|W\|_F^2 + \frac{1}{2} \|H\|_F^2 \Big\}, \quad \text{subject to } (w_c - w_{c'})^T h_{ic} \geq 1, \forall c \in [K], i \in [n_c], c' \neq c. \tag{3}$$

They further proved that the global minima of this max-margin problem are NC solutions; that is, $\hat{W}$ and $\hat{H}$ satisfy the properties of Definition 1. Moreover, all other critical points are saddles. Thrampoulidis et al. [7] proved that the regularization path of minimizing $\mathcal{L}(W, H)$ with vanishing regularization converges to the global minimizers of Eq. (3). They further characterized these global minima in the imbalanced classification setting, showing how solutions deviate from NC geometry as a function of the imbalance ratio and number of classes.

Despite this progress, the question of implicit bias remains open: whether gradient flow (Eq. (2)) converges to these global minima is unknown even in the balanced case, despite the apparent canonical nature of the setting. To understand this gap, note that the UFM is equivalent to a two-layer linear neural network with first layer H, second layer W, and inputs given by the standard basis vectors in $\mathbb{R}^{Kn}$.

Exact Dynamics with MSE loss. Replacing the CE loss with MSE loss makes the UFM tractable [10]. Saxe et al. [8, 9] introduced such techniques for linear networks trained with MSE loss. Here we show how their methodology applies in the UFM setting with MSE loss. Although their technique is slightly more general, instead being applicable to whitened inputs, the UFM is a canonical instance that as we discussed serves as a meaningful proxy to sufficiently expressive deep neural networks.

Replacing CE loss with MSE loss, the gradient flow equations given in Eq. (2) become:

$$\frac{dW}{dt} = (Y - WH)H^T, \qquad \frac{dH}{dt} = W^T(Y - WH). \tag{4}$$

Let the singular value decomposition of Y be $Y = U\Sigma V^T$, with singular values s_c for $c = 1, \ldots, K$. Restrict initialization to a form of spectral initialization:

$$W = UD_W R^T, \qquad H = RD_H V^T,$$

for some orthogonal matrix $R \in \mathbb{R}^{d\times d}$, with D_W and D_H having their only non-zero diagonal entries being $\alpha_1, ..., \alpha_K$ and $\beta_1, ..., \beta_K$ respectively. Then, at initialization, the label matrix Y and the logit matrix WH are mutually diagonalizable:

$$Y - WH = U(\Sigma - D_W D_H)V^T\,,$$

When put in Eqns. 4, and using $U^TU = V^TV = I$, this implies:

$$\frac{dW}{dt} = U\,(\Sigma - D_W D_H)D_H\,R^T, \qquad \frac{dH}{dt} = R\,D_W(\Sigma - D_W D_H)\,V^T. \tag{5}$$

Thus only the singular values of W and H evolve, while the singular vectors remain fixed. Specifically, from Eq. (5), the singular values evolve according to

$$\frac{d\alpha_i}{dt} = (s_i - \alpha_i\beta_i)\,\beta_i, \qquad \frac{d\beta_j}{dt} = (s_i - \alpha_i\beta_i)\,\alpha_i\,.$$

Note that these equations decouple into K independent pairs of ordinary differential equations, and can be solved to show that the products of the singular values $\alpha_j\beta_j$ follow sigmoidal trajectories from their initial value to s_j. In particular, this implies that the logit singular values $a_i = \alpha_i\beta_i$ converge monotonically to the singular values of the target matrix Y [11]. One can also recover an energy functional corresponding to the dynamics.

3 Hadamard Initialization

Contrasting the gradient flow equations for CE and MSE loss, given in Eqns. (2) and (4), highlights the main difficulty in extending the exact dynamics argument to the CE setting: the presence of the softmax matrix P. In the MSE case, obtaining exact dynamics requires only that Y and WH be simultaneously diagonalizable, which can be enforced through a suitable initialization. In the CE case, however, one must simultaneously diagonalize Y, WH, and P, which demands precise control over how the softmax nonlinearity interacts with the singular value structure of the logit matrix. Because nonlinear functions are generally unnatural operations on matrices, they rarely preserve such structure.

3.1 Diagonalizing the Softmax

Nevertheless, we show that this obstacle can be overcome for CE by initializing with Sylvester Hadamard matrices.

Definition 2 (Sylvester Hadamard Matrix). *The Sylvester Hadamard matrices $\{\Phi_{2^m} : m \in \mathbb{N}\}$ are defined recursively by*

$$\Phi_1 = 1, \quad \Phi_{2^m} = \begin{bmatrix} \Phi_{2^{m-1}} & \Phi_{2^{m-1}} \\ \Phi_{2^{m-1}} & -\Phi_{2^{m-1}} \end{bmatrix} \in \mathbb{R}^{2^m\times 2^m}.$$

Each Φ_{2^m} is a ± 1-valued matrix, whose first row and column consist entirely of ones. Moreover, these matrices are symmetric $\Phi_{2^m}^T = \Phi_{2^m}$ and satisfy the orthogonality relation $\Phi_{2^m}^T\Phi_{2^m} = 2^m I_{2^m}$. Sylvester Hadamard matrices form a special subclass of the more general Hadamard matrices; we leave the study of general matrices to future work.

We will begin by illustrating our technique for the UFM of Eq. (1) in the simplified setting where $n = 1$, meaning one data point per class, before establishing the result for general n. In this case, the label matrix reduces to $Y = I_K$, so its singular vector matrices U and V may be taken as any orthogonal matrix with $U = V$.

We show that by selecting $U = V = \frac{1}{\sqrt{K}}\Phi$, where Φ is a Sylvester Hadamard matrix, mutual diagonalizability can be achieved even in the presence of the softmax matrix.

Theorem 1. *Let $K = 2^m$ for $m \in \mathbb{N}$. Suppose $U = \frac{1}{\sqrt{K}}\Phi$, where Φ is the $K \times K$ Sylvester Hadamard matrix. Denote the columns of U by u_i, $i = 1, \ldots, K$. For a matrix of the form $Z = \sum_i a_i u_i u_i^T$, the corresponding softmax matrix is*

$$P(Z) = \sum_i \nu_i u_i u_i^T,$$

where

$$\nu_i = \frac{\sum_j \Phi_{ij} e^{\frac{1}{K}(\Phi a)_j}}{\sum_j e^{\frac{1}{K}(\Phi a)_j}}.$$

Throughout, we refer to this initialization as *Hadamard Initialization*. The theorem proves that this initialization is preserved under the nonlinear softmax operation. In matrix form, the softmax of a matrix diagonalizable by a Sylvester Hadamard matrix remains diagonalizable in the same basis:

$$P(U \operatorname{diag}(a) U^T) = U \operatorname{diag}(UP(Ua)) U^T, \tag{6}$$

where we recall that the softmax map $P(\cdot)$ acts columnwise. This diagonalization property follows from the interplay of two fundamental properties of the softmax map and of the Sylvester-Hadamard matrices respectively:

1. *Permutation equivariance of the softmax*: For any permutation matrix Π and any vector v, the softmax map satisfies the identity $P(\Pi v) = \Pi P(v)$.
2. *Group structure of Hadamard rows*: The rows of Φ, denoted as $\phi_0, \ldots, \phi_{K-1}$, form a group under the Hadamard product (i.e., elementwise multiplication), governed by the relation $\phi_i \odot \phi_j = \phi_{i \oplus j}$, where $\oplus$ denotes the bitwise XOR operation. A direct consequence of this is a key matrix identity: right-multiplying Φ by a diagonal matrix formed from one of its rows is equivalent to permuting the rows of Φ: $\Phi \operatorname{diag}(\phi_i) = \Pi_i \Phi$, where Π_i is the permutation matrix that maps row j to row $i \oplus j$.

When these properties come together, it ensures that the application of softmax to Z (i.e., the LHS of Eq. (6)) produces a dyadic matrix: a matrix whose entries are of the form $c_{i \oplus j}$. Such matrices are known to be diagonalized by Sylvester-Hadamard matrices [12], which validates the form given in the equation above and proves the theorem. A detailed proof is provided in Section 7.1.

3.2 Tractable CE Dynamics

As a consequence of Theorem 1, the arguments of Saxe et al., detailed in Sec. 2, can be adapted to our setting. Specifically, if the parameter matrices are initialized as $W = UD_W R^T$, $H = RD_H U^T$, for any orthogonal matrix $R \in \mathbb{R}^{d \times d}$, with $U = \frac{1}{\sqrt{K}}\Phi$, then Y, WH, and P are mutually diagonalizable. This ensures that the singular vectors of W and H are fixed, and all time evolution occurs entirely in their singular values. We emphasize that this property is special to our construction: choosing a general orthogonal matrix for U does not yield mutual diagonalizability of the logits and the softmax matrix.

Although the previous argument was presented for $n = 1$, it extends directly to balanced classes of arbitrary size. This is summarized in the following theorem, a proof of which appears in Appendix D.1.

Theorem 2. *Let $K = 2^m$ for $m \in \mathbb{N}$, and let $U = \frac{1}{\sqrt{K}}\Phi$, where Φ is the $K \times K$ Sylvester Hadamard matrix. Consider the UFM described in Eq. (1), with balanced classes of size n. Initialize the parameter matrices as $W = UD_W R^T$, and $H = RD_H V^T$, where $R \in \mathbb{R}^{d \times d}$ is orthogonal, $V = U \otimes Q$ where Q is a right singular matrix of 1_n^T, and the only nonzero singular values of D_W and D_H are $\alpha_1, \ldots, \alpha_K$ and $\beta_1, \ldots, \beta_K$, respectively. Then the gradient flow equations, detailed in Eq. (2), reduce to*

$$\frac{d\alpha_i}{dt} = \frac{\sqrt{n}}{D} b_i \beta_i, \quad \frac{d\beta_i}{dt} = \frac{\sqrt{n}}{D} b_i \alpha_i,$$

where

$$b_i = \sum_{j=1}^{K} [1_K 1_K^T - \Phi]_{ij}\, e^{-\frac{1}{K\sqrt{n}}([1_K 1_K^T - \Phi]a)_j}, \quad D = \sum_{j=1}^{K} e^{-\frac{1}{K\sqrt{n}}([1_K 1_K^T - \Phi]a)_j}, \quad \text{with } a_i = \alpha_i \beta_i. \tag{7}$$

Recall that the first row and column of Φ consist of all ones. Since $[1_K 1_K^T - \Phi]_{1j} = 0$ for all j, it follows that $b_1 = 0$, and therefore α_1, β_1 remain constant. Moreover, since $[1_K 1_K^T - \Phi]_{j1} = 0$ for all j, they do not appear in the evolution equations for the other singular values. Without loss of generality, we set $\alpha_1 = \beta_1 = 0$ and drop these from our vectors. Thus, from now on we view $\alpha, \beta \in \mathbb{R}^{K-1}$ to be vectors containing only the non-trivial singular values, where the indices have been updated to account for the dropping of the trivial singular mode.

In keeping with this, we define $\Psi \in \mathbb{R}^{K-1 \times K-1}$ by deleting the first column and row of the matrix $1_K 1_K^T - \Phi$, meaning $\Psi = 1_{K-1} 1_{K-1}^T - X$, where $X = \Phi[2:K, 2:K]$ is the core of the Sylvester Hadamard matrix. Note that $-\Psi$ is the Schur complement of the full Hadamard matrix Φ.

It is straightforward to verify that these dynamics correspond to gradient flow on the energy functional

$$E = Kn \log\left(1 + \sum_{i=1}^{K-1} \exp\Big(-\frac{1}{K\sqrt{n}}(\Psi a)_i\Big)\right), \tag{8}$$

where $a_i = \alpha_i \beta_i$. A direct consequence of this variational structure is the continuous symmetry $\alpha_i \to \gamma\alpha_i, \beta_i \to \gamma^{-1}\beta_i$, from which Noether's theorem yields the conserved quantities $\alpha_i^2 - \beta_i^2 = C_i$.

For simplicity, and following common practice in the literature, we set $C_i = 0$, which corresponds to the case $\alpha_i = \beta_i$. This reduction is motivated by the fact that small initialization naturally leads to $C_i \approx 0$ in practice. Under this mild assumption, the dynamics can be reduced to the evolution of the *logit singular values* $a_i = \alpha_i \beta_i$, and our differential equations become the following for $i = 1, ..., K-1$:

$$\frac{da_i}{dt} = \frac{d\alpha_i}{dt}\beta_i + \alpha_i \frac{d\beta_i}{dt} = \frac{\sqrt{n}}{D} b_i(\alpha_i^2 + \beta_i^2) = \frac{2\sqrt{n}}{D} b_i a_i,$$

We will also absorb the $2\sqrt{n}$ factor into the time variable, and absorb a factor of $1/(K\sqrt{n})$ into the a variable, which reduces us to the following dynamics for $a \in \mathbb{R}^{K-1}$:

$$\frac{da_i}{dt} = \frac{1}{D} b_i a_i\,, \tag{9}$$

where

$$b_i = \sum_{j=1}^{K-1} \Psi_{ij} e^{-(\Psi a)_j}, \quad D = 1 + \sum_{j=1}^{K-1} e^{-(\Psi a)_j}. \tag{10}$$

3.3 Properties of the Reduced Dynamics

We provide a detailed analysis of the dynamics in Eq. (9) in Sec. 4. Here, we first establish several basic but important properties of these dynamics that will be used in the subsequent analysis.

On $K = 2^m$. Our construction requires $K = 2^m$ for some $m \in \mathbb{N}$ to exploit the Sylvester-Hadamard structure. Beyond the special case of $K = 2$ (binary classification), the dynamics of the CE UFM have not previously been amenable to analysis. Our construction enables tractable analysis for multiple values of K, capturing genuine multiclass effects. We therefore view the restriction to $K = 2^m$ not as a limitation, but as a modeling choice that enables fine-grained characterization of the dynamics.

Coupled dynamics. Unlike in the MSE case, the terms b_i/D depend on the full set of singular values $a_1, \ldots, a_{K-1}$. The equations therefore do not decouple, and each singular value influences the trajectory of the others. This additional coupling introduces a level of sophistication absent from the MSE setting and underlies many of the departures from its behavior. Because of this lack of decoupling, there are not enough conserved quantities to guarantee closed-form dynamics. Nevertheless, the reduced dynamics in Eq. (9) are sufficiently tractable to characterize the implicit bias and convergence properties, as we demonstrate in the subsequent sections.

Positive singular values. The singular values $a_i(t)$ are monotonically increasing in time. Therefore, if $a_i(0) > 0$, then $a_i(t) > 0$ for all $t \geq 0$. Moreover, if $a_i(0) = 0$, then it remains zero throughout training. We henceforth assume all singular values are initialized to positive values.

Mirror-flow interpretation. Since $a_i(t)$ remains positive for all future times, we can equivalently write the updates (absorbing the K, n-dependent factors) as $\mathrm{d}(\log(a_i))/\mathrm{d}t = -\nabla_{a_i} E(a)$, where $E = E(a)$ is the energy functional in Eq. (8). In this form, the dynamics are viewed as mirror-flow updates with respect to the entropy potential for the energy functional in Eq. (8); see Appendix B for details.

Normalized singular values. From Eq. (9), it is straightforward to show that the norm of the logit singular values (e.g., the L_1-norm $\|a\|_1 = \sum_i a_i$) diverges to infinity. This divergence implies that there are no finite minima of the UFM objective on the optimization surface. We therefore study the dynamics of the *normalized logit singular values* $\hat{a}_i = a_i/\|a\|_1$, which evolve as:

$$\frac{d}{dt}(\hat{a}_i) = \frac{\hat{a}_i}{D}\left[b_i - \bar{b}\right], \quad \text{where } \bar{b} = \sum_{j=1}^{K-1} \hat{a}_j b_j. \tag{11}$$

The updates in Eq. (11) correspond to replicator dynamics in evolutionary game-theory (see, e.g., [13] and references therein), where $\hat{a}_i$ represents the strategy of player i, b_i is that strategy's payoff, and $\bar{b}$ the average payoff. Note that the normalized logit singular value vector $\hat{a}$ lies on the $K-1$ dimensional simplex, since with positive initialization all entries remain positive throughout the dynamics.

4 CE Dynamics and Implicit Bias

4.1 NC as a Candidate Among Stable Directions

Having established simplified dynamics for UFM gradient flow in Sec. 3.2, we now analyze their behavior. In particular, we study their *implicit bias*: among the many directions that achieve zero loss under the energy functional in Eq. (8), which direction do the dynamics converge to? Note that any direction $\hat{a}$ satisfying $(\Psi\hat{a})_i > 0$ for $i \in [K-1]$ can attain arbitrarily small loss by taking $\|a\|_1 \to \infty$. The key question is therefore which of these zero-loss directions gradient flow in Eq. (11) selects.

To gain insight into the structure of potential convergence directions, we first identify the *stable directions*, i.e., those satisfying $\frac{\mathrm{d}\hat{a}}{\mathrm{d}t} = 0$. If the limiting direction satisfies $\hat{a}_i > 0$ for all i (meaning all $K-1$ logit singular values are positive), then the solution is full-rank; otherwise, convergence occurs to a low-rank solution. Among full-rank stable directions, the unique solution is $\hat{a} = \frac{1}{K-1}1_{K-1}$. Indeed, combining $\frac{\mathrm{d}\hat{a}}{\mathrm{d}t} = 0$ with $\hat{a}_i > 0$ for all i implies $b = \bar{b}1_{K-1} \propto 1_{K-1}$, from which a short algebraic argument (see Appendix B.2) yields the claim.

Note that $\hat{a} = \frac{1}{K-1}1_{K-1}$ corresponds exactly to the NC solution of Definition 1. Indeed, equal logit singular values (i.e., $a_i = a_j$ for all $i, j \in [K-1]$) imply that the singular values of W and H are equal as well: $\alpha_i = \beta_i = \alpha_j = \beta_j$ for all $i, j \in [K-1]$. Recall that we have assumed, without loss of generality, that the singular value corresponding to the all-ones eigenspace is zero for both matrices. Together, these conditions yield $W \propto \Phi\,\mathrm{diag}(0, 1, \ldots, 1)R^T$ and $H \propto R\,\mathrm{diag}(0, 1, \ldots, 1)V^T$, which in turn give $WW^T \propto S$ and $H^T H \propto S$, where $S = \frac{1}{K}\Phi\,\mathrm{diag}(0, 1, \ldots, 1)\Phi^T$ is the simplex ETF matrix. This establishes the NC geometry.

4.2 Convergence to Neural Collapse via a Lyapunov Function

Does gradient flow exhibit an implicit preference for the full-rank NC solution $\hat{a} = \frac{1}{K-1}1_{K-1}$ among all stable points? We show that this is indeed the case. Moreover, we establish a stronger result than asymptotic convergence: we identify a Lyapunov function that decreases monotonically along the gradient flow trajectory and explicitly measures the distance from the NC solution.

Theorem 3. *Let $K = 2^m$, where $m > 0$ is a natural number. Let $a(t) \in \mathbb{R}^{K-1}$ be a variable initialized with $a_i(0) > 0$ for all i, that evolves under the time evolution of Eq. (9). Then the metric $D_{\mathrm{KL}}(\frac{1}{K-1}1_{K-1} \,\|\, \hat{a})$, where D_{KL} is the Kullback-Leibler divergence, serves as a Lyapunov function of the flow. Consequently, the metric is decreasing throughout training and we must have $\hat{a} \to \frac{1}{K-1}1_{K-1}$ as $t \to \infty$.*

The proof appears in Section 7.2. This theorem establishes that gradient flow necessarily converges to NC, providing the first rigorous characterization of implicit bias in the CE UFM beyond binary classification.

Our result is notable in two respects.

First, regarding *implicit bias*, existing results typically establish only asymptotic convergence rather than monotonic behavior via a Lyapunov function. Even when restricted to asymptotic convergence, the only prior result applicable to the CE UFM is by Ji and Telgarsky [4], which is limited to binary classification ($K = 2$). However, the binary case is trivial in our framework: the unique stable direction (of rank $K - 1 = 1$) is NC itself. In contrast, for $K \geq 4$, multiple stable directions exist, making the implicit bias question nontrivial. Indeed, as we show in the next subsection, even $K = 4$ versus $K = 8$ versus $K > 8$ exhibit qualitatively different optimization properties: several candidate Lyapunov functions work for smaller K but admit counterexamples for larger K.

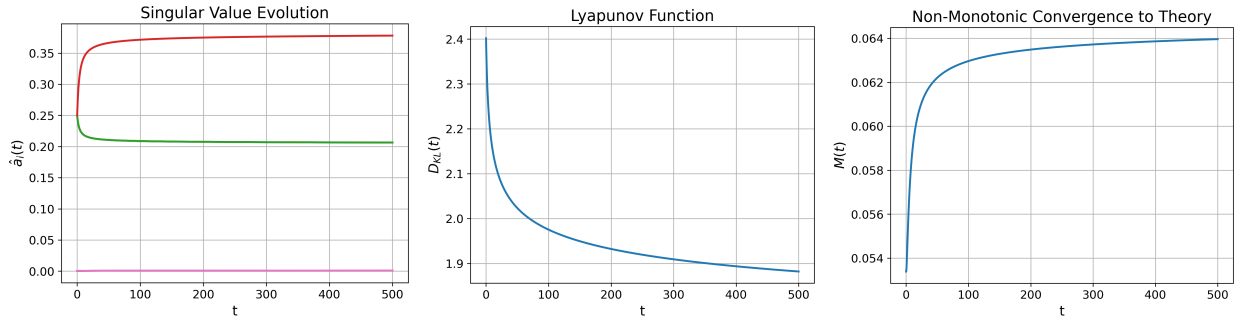

Figure 1: Training of the logit singular values under the evolution given in Eq. (9) with $K = 8$, from the initialization $a(0) = [1, 1, 1, 1, 0.001, 0.001, 0.001]$. **Left:** The evolution of the singular values $\hat{a}_i(t)$. **Middle:** The evolution of the KL divergence $D_{\mathrm{KL}}(\frac{1}{K-1}1_{K-1} \,\|\, \hat{a})$. **Right:** The evolution of the metric M, given in Eq. (12).

Second, regarding *neural collapse*, all previous claims about convergence to NC either used explicit finite regularization [14], analyzed the regularization path [7], or characterized the KKT points of the corresponding non-convex max-margin problem in Eq. (3) [6]. To the best of our knowledge, this is the first proof of NC occurring purely as a result of gradient flow implicit bias under CE loss.

4.3 Non-monotonic Behaviors

The monotonic convergence established in the previous section—via the KL divergence as a Lyapunov function—is a special property. We now show that many other natural distance metrics to NC do not exhibit monotonic behavior under gradient flow.

We begin with the L2-distance metric, which is commonly used in empirical work to quantify distance from NC:

$$M = \frac{1}{2} \sum_{i=1}^{K-1} \left(\hat{a}_i - \frac{1}{K-1} \right)^2 . \tag{12}$$

For $K = 2, 4$, this metric decreases monotonically throughout training. However, this monotonicity breaks down for $K \geq 8$, as demonstrated in the following theorem.

Theorem 4. *Let $K = 2^m$, where $m > 0$ is a natural number. Let $a(t) \in \mathbb{R}^{K-1}$ be a variable initialized with $a_i(0) > 0$ for all i, that evolves under the time evolution of Eq. (9). If $K = 2$ or $K = 4$, then the metric M, defined in Eq. (12), monotonically converges to 0. If $K \geq 8$, then there exist initializations $a(0)$ such that the metric M is not monotonic.*

The proof appears in Appendix D.2. This result reveals a fundamental qualitative difference from the MSE case studied in [8, 11, 9]: here, the coupled singular values can induce non-monotonic behavior in natural distance metrics. Figure 1 illustrates the severity of this phenomenon, showing the evolution of $\hat{a}_i$ and M for an initialization leading to non-monotonic behavior. After a very brief initial decrease, the metric M increases monotonically over the remainder of the considered trajectory, though it must eventually converge to zero.

The role of non-uniform stable points ($K \geq 8$). The mechanism underlying this behavior is as follows: the dynamics first approach a low-rank stable point with non-uniform support—that is, a point where $\hat{a}_i = 0$ for some $i \in [K-1]$ and among the non-zero indices, $\hat{a}_i \neq \hat{a}_j$ for some i, j. Because the trajectory moves toward a non-uniform configuration, metrics such as M measuring uniformity increase. The existence of such non-uniform stable points arises from the structure of the matrix Ψ. Our energy functional (Eq. (8)) depends heavily on the quantities $(\Psi a)_i$, which play the role of margins. Each singular mode $\hat{a}_j$ contributes to enlarging $K/2$ of these margins—precisely those indices $i \in [K-1]$ for which $\Psi_{ij} \neq 0$. As $\|a\|_1$ becomes large, the energy is determined, up to exponentially small terms, by the smallest normalized margins $\arg\min_i (\Psi \hat{a})_i$. While NC maximizes this quantity in the full-rank setting, in rank-constrained regimes we do not necessarily have an equal number of singular values contributing to each margin, making non-symmetric solutions preferable. The case shown in Figure 1 illustrates this behavior. A more detailed explanation appears in Appendix C.

For $K = 4$, all stable points are uniform in their support, preventing such behavior and ensuring a monotonic decrease of M. For $K \geq 8$, when non-monotonicity occurs, the optimization path remains near low-rank stable points for extended periods before transitioning to higher-rank saddles, eventually reaching NC. Similar behavior has been observed empirically by Jacot et al. [15], albeit in settings with finite-norm critical points.

Nuanced optimization behavior with respect to K. Beyond this specific example, the K-dependence of convergence behavior is more nuanced. The dynamics clearly depend on K through the maximum number of rank transitions along trajectories. Theorem 4 captures how the nature of convergence changes qualitatively with K: from monotonic ($K = 4$) to non-monotonic ($K \geq 8$). However, this K-dependence manifests differently across metrics. The reverse KL divergence $D_{\mathrm{KL}}(\hat{a} \,\|\, \frac{1}{K-1} 1_{K-1})$ also admits non-monotonic counterexamples for $K = 8$ (unlike the KL divergence in Theorem 3); see Appendix B.3. Even more remarkably, some metrics exhibit monotonic convergence for $K = 8$ but break down for $K \geq 16$. An example is the logarithm of the ratio between smallest and largest normalized singular values, presented in Appendix B.3. These findings suggest that many other non-trivial behaviors likely emerge as the class number increases.

Exponential slow down. The reason behind the potentially excessive timescale of convergence is that lazy training dynamics emerge in the large norm regime. To see this consider the expressions for b_i/D in terms of the variables $\|a\|_1, \hat{a}$:

$$\frac{b_i}{D} = \frac{1}{1 + \sum_{j=1}^{K-1} e^{-\|a\|_1 (\Psi \hat{a})_j}} \sum_{j=1}^{K-1} \Psi_{ij} e^{-\|a\|_1 (\Psi \hat{a})_j} \,.$$

Clearly as $\|a\|_1$ becomes large, these terms decay exponentially in $\|a\|_1$, and so the derivatives of a_i go to zero at $O(\|a\|_1 e^{-C\|a\|_1})$, for some constant C, in this limit, and the network evolution becomes incredibly slow. In settings where the loss does not take its minimal value at infinity, such as in the MSE case, these slow down effects do not occur.

4.4 Further Comparisons to MSE Dynamics

We have seen that CE dynamics are fundamentally different from MSE dynamics via non-monotonic convergence and slow convergence rates. Here, we further emphasize this difference by identifying the simplifications the CE setting would need to exhibit the simple behavior of MSE.

4.4.1 Linearized Dynamics

The MSE behavior of independent sigmoidal singular value evolution can be recovered in CE dynamics, but only under a severe restriction: the small $\|a\|_1$ regime where linearized dynamics apply. As the following theorem shows, the singular values decouple in this limit and evolve along sigmoidal trajectories exactly as in the MSE case. However, even with small initialization, the system quickly escape this regime, after which coupling effects introduce the complex behaviors discussed in the preceding subsections. The proof appears in Appendix D.3.

Theorem 5. *Let $K = 2^m$, where $m \in \mathbb{N}$. Let $a(t) \in \mathbb{R}^{K-1}$ evolve under the dynamics of Eq. (9), with initialization $a_i(0) > 0$ for all $i = 1, ..., K-1$. To linear order in the scale $\|a\|_1$, the dynamics reduce to*

$$\frac{da_i}{dt} = a_i \left[1 - a_i\right] + O\left(\|a\|_1^3\right).$$

In particular, the linearized dynamics decouple, and can therefore be solved exactly. The resulting trajectories, that are accurate in the linearized regime, are

$$a_i(t) = \frac{a_i(0)}{a_i(0) + (1 - a_i(0))e^{-t}}$$

The key contrast with MSE is that this decoupled behavior persists at all scales for MSE, whereas for CE it is confined to the small $\|a\|_1$ regime. Moreover, the linearized equations imply $a_i/a_j \to 1$ monotonically, pushing the early-stage dynamics toward NC—though the system exits this regime too quickly for this effect alone to guarantee monotonic convergence to NC.

4.4.2 Comparison to Logit Optimization

The second simplification that recovers MSE-like behavior is removing the factorization layer: optimizing logits directly rather than through the product WH. This corresponds to the convex one-layer network analyzed by Saxe et al. [8]. The loss in this setting is

$$\mathcal{L}(Z) = \sum_{c=1}^{K} \sum_{i=1}^{n} \log\Big(1 + \sum_{c' \neq c}^{K} e^{(z_{ic})_{c'} - (z_{ic})_c}\Big), \tag{13}$$

where z_{ic} are the columns of the logit matrix Z arranged in class order. With Hadamard Initialization of the logit matrix Z, we have the following theorem showing the corresponding reduction in the dynamics, and consequently monotonic convergence to NC.

Theorem 6. *Let $K = 2^m$ for $m \in \mathbb{N}$, and let $U = \frac{1}{\sqrt{K}}\Phi$, where Φ is the $K \times K$ Sylvester Hadamard matrix. Consider the optimization problem described in Eq. (13), and initialize the logit matrix as $Z = UD_ZV^T$, with $V = U \otimes Q$ where Q is a right singular matrix of 1_n^T, and D_Z has its nonzero singular values given by $a_1, ..., a_K$. Then the gradient flow equation on this objective, after absorbing a factor of $\frac{1}{K\sqrt{n}}$ into the a variable, a factor of $\frac{1}{K}$ into the t variable, and dropping the trivial singular value, reduces to*

$$\frac{da_i}{dt} = \frac{1}{D}b_i, \qquad i = 1, ..., K-1.$$

With the quantities D and b as defined in Eq. (10). Moreover, the metric M in Eq. (12) decreases monotonically along trajectories of gradient flow.

This result parallels that of Soudry et al. [2] for logistic regression, since NC corresponds to the max-margin solution of the UFM. We see that, when optimizing logits directly, the system converges monotonically to the NC solution—just as in the MSE case studied by Saxe et al. [8, 9]. Taken together, this highlight the role of depth: it is precisely the additional layer that disrupts monotonicity and introduces the more intricate behaviors observed in the CE UFM.

5 Standard Initialization Schemes lead to NC

The evidence presented in Sec. 4.3 raises a natural question: why are such peculiar behaviors not observed in typical training of UFMs [6] or in standard neural networks? To provide an explanation, we first demonstrate that even though we do not generally see monotonic convergence to NC, the metric M must decrease when it has hit a small enough value.

Theorem 7. *Let $K = 2^m$, where $m > 0$ is a natural number. Let $a(t) \in \mathbb{R}^{K-1}$ be a variable initialized with $a_i(0) > 0$ for all i, that evolves under the time evolution of Eq. (9). Then the metric M, defined in Eq. (12), satisfies the following*

$$M \leq \min\left(\frac{1}{8K^2}, \frac{1}{18K^3\|a\|_1^2}\right) \implies \frac{dM}{dt} \leq 0$$

The proof appears in Appendix D.5. This demonstrates that M will decrease for a finite time training trajectory, provided the initialization is sufficiently close to the NC solution.

We can now explore how typical initialization schemes, which feature random i.i.d. entries in the parameter matrices, impact the singular values of the logit matrix. Note this differs from our modeling context involving Hadamard matrices. However, empirically, we find that our modeling approximates random dynamics from small initialization; see Sec. 6. We show below that with Gaussian initialization at small scale, the zeroth-order dynamics already bias the system towards the simplex ETF.

Theorem 8. *Consider the UFM defined in Eq. (1) and evolve the parameter matrices under the gradient flow of Eq. (2), with initialization $W(0) = \epsilon d^{-\frac{1}{2}}A_1$, $H(0) = \epsilon d^{-\frac{1}{2}}A_0$, where the matrices A_1, A_0 have entries drawn i.i.d. from a standard Gaussian distribution, and d is the network width. Defining the quantity $t^* = \frac{1}{2\sqrt{n}}\log(\frac{1}{\epsilon})$, and taking the limit $\epsilon \to 0$, the leading order term of $Z(t^*)$ is given by*

$$Z(t^*) \sim \frac{1}{2d\sqrt{n}}\epsilon^2 e^{2t^*\sqrt{n}}\left[SXX^T(S \otimes 1_n^T)\right],$$

where $X \in \mathbb{R}^{K\times d}$ has entries given in distribution by independent standard Gaussians. Hence additionally taking the $d \to \infty$ limit, we recover asymptotically

$$Z(t^*) \sim \frac{1}{2\sqrt{n}}\epsilon^2 e^{2t^*\sqrt{n}}\left(S \otimes 1_n^T\right).$$

The theorem shows that for small initialization, after some period of training in the zeroth order regime, the logit matrix approximates a simplex ETF, with the quality of the approximation improving as initialization becomes smaller and the network width increases. Thus, zeroth-order dynamics naturally drive the logit matrix toward the simplex ETF structure.

Since the linearized dynamics further reinforce this geometry, and the NC solution is stable with an attractive basin at all scales, this naturally leads the trajectory to avoid regions of non-monotonicity that occur around the saddles in the surface. This explains why Ji et al. [6] observe only monotonic behavior: the initialization scheme itself imposes a strong bias toward the simplex ETF, effectively concealing the more intricate dynamics revealed by our analysis.

6 Numerical Experiments

6.1 Setup

We simulate the gradient flow of the UFM (Eq. (1)) using gradient descent with a small step size 0.01. We set the number of classes $K = 4$, the hidden dimension $d = K$, and use (for simplicity) $n = 1$ sample per class. We compare two initialization schemes: (1) *Random Init.:* We initialize W and H with i.i.d. entries drawn from a zero-mean Gaussian distribution with variance $1/K$, scaled by a small $\epsilon = e^{-6}/\sqrt{K} \approx 10^{-3}$. (2) *Hadamard Init.:* We initialize $W = UD_W R^T$ and $H = RD_H V^T$ as described in Theorem 1, $U = V = \frac{1}{\sqrt{K}}\Phi$. Recall from Sec. 3 that the singular value corresponding to the all 1s column of U is not updated by the flow; thus, as in theory, we assume without loss of generality that the respective entry of D_W, D_H is set to zero. To further ensure direct comparison to Random Init., we set the remaining $K-1$ singular values such that $D_W = D_H$ and they are identical to the largest $K-1$ singular values of the logit matrix from its corresponding Random Init. pair. This isolates the *initial* difference between the two schemes to only the singular vectors. Alongside these, and with same initialization, we numerically solve the reduced gradient flow equations for the singular values (Eq. (9)), which we label *ODE*.

All results are averaged over 10 runs, with the shaded regions in our plots representing the standard error of the mean. We track the following metrics as function of iterations: (a) Cross-Entropy Loss, (b) classically used NC alignment metrics (the logit one corresponds to the metric M in Eq. (12)), (c) KL Divergence metric of Thm. 3, (d) singular values, (e) normalized (by L1-norm) singular values, (f) Residual energy outside the signal subspace of Hadamard Init.. See Appendix E for details. We log all metrics at 100 geometrically-spaced iterations (for log-scale visualization).

In Appendix E we present additional experimental ablations on: (a) larger number of classes ($K \geq 4$), (b) non-infinitesimal initialization scale ($\epsilon = 1/\sqrt{K}$), and, (c) overparameterized hidden layer dimension ($d/K > 1$). The latter support the theoretical findings of Theorem 8: overparameterization $d \gg K$ accelerates convergence of the dynamics under infinitesimal random initialization to the simplex ETF geometry.

6.2 Results

Our results are summarized in Fig. 2.

Hadamard Initialization freezes singular vectors. For Hadamard Init., the residual energy in plot (f) remains at its initial zero value for all time. This confirms the central finding of Thm. 1: Hadamard Initialization perfectly diagonalizes the softmax, ensuring the singular vectors of W and H remain fixed throughout training.

Validation of reduced ODE. Plots (d) and (e) validate the theoretical reduction in Thm. 2 of the original matrix ODE (Eq. (1)) to the simpler vector ODE on the logit singular values (Eq. (9)). The empirical singular values from the Hadamard Init. (red) simulation are perfectly aligned with the ODE solution (green dotted line).

Convergence to NC. From both Hadamard and Random Init., the dynamics successfully converge to the NC geometry. This is captured in Plots (b) and (c) by tracking the standard NC convergence metrics from the literature and our Lyapunov function, respectively. Plot (e) also clearly shows the normalized singular values $\hat{a}_i$ for both initializations converging to the target $\frac{1}{K-1} = \frac{1}{3}$, as predicted by Thm. 3, albeit at a slow rate.

Both initializations show a qualitatively similar two-stage behavior:

1. *Alignment / "Burn-in" Stage ($t < 10^3$).* In this first stage, both methods show unchanged *Loss* (plot a) and a (near) plateau in the *alignment metrics* (plots b,c). However, their internal dynamics differ:
   - For Hadamard Init. (red), which starts perfectly aligned (*Residual Energy* in (f) is zero), the *normalized singular values* (e) and *Lyapunov function* (c) remain perfectly flat.
   - For Random Init. (blue), this is an active alignment phase. The *Residual Energy* (f) decays as vectors align to the ETF subspace. This alignment drives a reconfiguration of the *normalized singular values* (e), which are moving but *not* yet converging to their final equilibrium. This reconfiguration, in turn, causes the temporary oscillations seen in the *Lyapunov function* (c).
2. *Convergence Stage ($t > 10^3$)* The starting point of this stage coincides for both methods. Now, the *normalized singular values* (e) for both methods begin their slow convergence toward their final equilibrium. This "learning" phase coincides with a sharp increase in singular values magnitudes (d), a sharp decrease in Loss (a) and the final decrease in all convergence metrics (plots b and c).

The burn-in stage has a natural interpretation for Random Init.: the dynamics first search for a good subspace; once close to the ETF subspace, the parameter magnitudes grow sharply (plot (d)) to saturate the softmax and reduce the

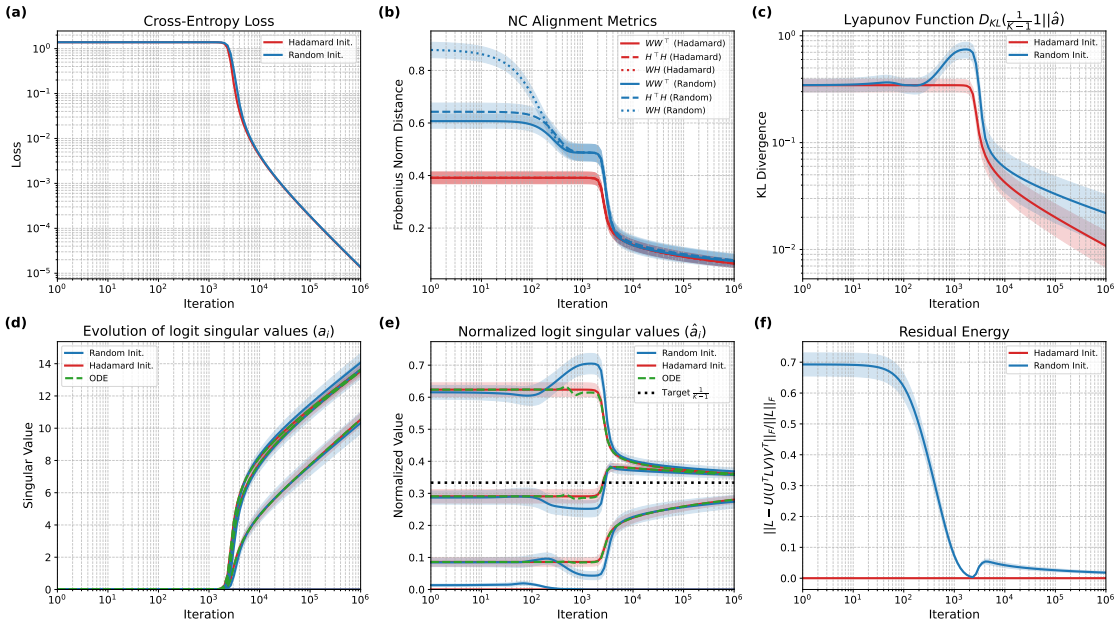

Figure 2: *Comparison of Hadamard Init., Random Init., and the theoretical ODE dynamics.* Plots show metrics vs. training iteration on a log scale. Shaded regions denote standard error over 10 realizations. Sec. 6 for detailed discussion.

CE loss, while the normalized singular values then drift slowly toward equilibrium. A similar "waiting" behavior is visible under Hadamard Init.: even though the subspaces are perfectly aligned from the start, there is initially very slow motion in the singular-value dynamics before the same growth-and-converge pattern sets in. This suggests a long time is needed for the dynamics to "escape" the neighborhood of the saddle point $W = H = 0$ to which the infinitesimal initialization is close to.

Another interesting dynamic, visible in plots (b) and (c), is the difference in monotonicity between the metrics. For Hadamard Init., both the L2-based metric M (plot b) and the Lyapunov (KL) function (plot c) decrease monotonically, which is consistent with Theorem 3 and (for $K = 4$) Theorem 4. For Random Init., however, the story is different: the Lyapunov function (plot c) is non-monotonic, showing oscillations during the alignment phase, even while the L2-based metric M (plot b) is monotonic. As shown in Appendix E, the non-monotonicity of the Lyapunov function for Random Init. persists for $K = 8$ and $K = 16$, while the metric M remains (nearly) monotonic for Random Init. even at larger K. This difference is intuitive: outside the spectral setting we analyze, singular vectors impact metric M but not D_{KL}. Thus M benefits during the alignment phase from singular vector evolution, while D_{KL} does not. Importantly, this does not contradict Theorem 4, which proves the existence of a non-monotonic trajectory. That counterexample requires specific non-uniform initialization, a regime that standard i.i.d. Gaussian initialization naturally avoids (consistent with the conclusions of Sec. 5).

Overall, our findings in the appendix for larger values of $K = 8, 16$ confirm our conclusions about the presence of two-stage dynamics and qualitative consistency between random and Hadamard Init.. However, they also emphasize the intricacies of multi-class CE optimization. We observe that the quantitative approximation of Random Init. by Hadamard Init. deteriorates as K increases (e.g., $K = 16$ in Appendix E). This is yet another departure from the MSE setting [8], where such spectral approximations are typically very accurate. The softmax non-linearity introduces a complex interaction between singular vector alignment and singular value evolution that invites further investigation.

7 Proofs

In this section, we present the proofs of two core theoretical results: that Hadamard initialization diagonalizes the softmax (Theorem 1) and that the KL divergence characterizes convergence to the neural-collapse geometry (Theorem 3).

7.1 Proof of Theorem 1

Let $\phi_i, i \in \{0, 1, \ldots, K-1\}$ be the rows of the $K \times K$ Sylvester Hadamard matrix Φ. The indexing from 0 to $K-1$ is convenient for presenting the following key circular property of Sylvester Hadamard matrices [12, Lemma 3.6].

Fact 1. *For any $i, j \in \{0, \ldots, K-1\}$, it holds that $\phi_i \odot \phi_j = \phi_{i \oplus j}$, where $i \oplus j$ is the bitwise XOR of i and j .*

For convenience, define matrices

$$\Phi_i := \Phi \operatorname{diag}(\phi_i) = \left[\phi_j^T \odot \phi_i^T\right]_{j \in [k]} \quad \text{for all } i \in \{0, \ldots, K-1\}.$$

Note that $\Phi_0 = \Phi$. Using Fact 1, we can express Φ_i in terms of the permutation matrix

$$\Pi_i = \left[e_{i \oplus j}^T\right]_{j \in [k]}, \quad \text{for all } i \in \{0, 1, \ldots, K-1\}$$

as follows:

$$\Phi_i = \Pi_i \Phi, \quad \text{for all } i \in \{0, \ldots, K-1\}. \tag{14}$$

Recall that softmax is applied columnwise. Thus, to prove the advertised identity in Eq. (6), it suffices for any column $i \in \{0, 1, \ldots, K-1\}$ that

$$P(\frac{1}{K}\Phi \operatorname{diag}(a)\, \phi_i) = \frac{1}{K}\Phi \operatorname{diag}\left(\Phi P(\frac{1}{K}\Phi a)\right) \phi_i\,.$$

Using the property $\operatorname{diag}(u)\, v = \operatorname{diag}(v)\, u$ for any u, v on both sides of the equality, and applying Eq. (14), this is equivalent to

$$P(\frac{1}{K}\Pi_i \Phi a) = \frac{1}{K}\Pi_i \Phi \Phi P(\frac{1}{K}\Phi a)\,.$$

But, $\Phi\Phi = \Phi\Phi^T = kI_K$, thus the claim reduces to proving

$$P(\frac{1}{K}\Pi_i \Phi a) = \Pi_i P(\frac{1}{K}\Phi a)\,.$$

In turn, this identity is true because of the permutation equivariance property of softmax, i.e. $P(\Pi u) = \Pi P(u)$ for any permutation matrix Π and vector u. This completes the proof.

7.2 Proof of Theorem 3

The KL divergence $D_{\mathrm{KL}}(p \,\|\, q) = \sum_i p_i \log\left(\frac{p_i}{q_i}\right)$ between two categorical distributions with probabilities p and q is non-negative and equals zero if and only if $p = q$. In our setting, we have:

$$D_{\mathrm{KL}}\left(\frac{1}{K-1} 1_{K-1} \,\|\, \hat{a}\right) = -\log(K-1) - \frac{1}{K-1}\sum_i \log(\hat{a}_i).$$

For the remainder of the proof we denote this quantity by D_{KL}. Differentiating, we obtain

$$\frac{d}{dt}(D_{\mathrm{KL}}) = -\frac{1}{K-1}\sum_i \frac{1}{\hat{a}_i}\frac{d}{dt}(\hat{a}_i) = -\frac{1}{(K-1)D}\left[\left(\sum_i b_i\right) - (K-1)\bar{b}\right], \tag{15}$$

where the second equality follows from substituting Eq. (11).

We now define $x_i = (\Psi a)_i$ and rewrite the derivative in terms of these variables, which recall play the role of margins. First observe that

$$\bar{b} = \sum_i \hat{a}_i b_i = \sum_{ij} \hat{a}_i \Psi_{ij} e^{-(\Psi a)_j} = \frac{1}{\|a\|_1}\sum_{ij} a_i \Psi_{ij} e^{-(\Psi a)_j},$$

which simplifies to

$$\bar{b} = \frac{1}{\|a\|_1}\sum_j x_j e^{-x_j}. \tag{16}$$

Second note

$$\sum_i b_i = \sum_{ij} \Psi_{ij} e^{-(\Psi a)_j}.$$

Using $\sum_i \Psi_{ij} = K$, we obtain

$$\sum_i b_i = K \sum_j e^{-(\Psi a)_j},$$

Moreover, since

$$\sum_i (\Psi \hat{a})_i = \sum_{ij} \Psi_{ij} \hat{a}_j = K \sum_j \hat{a}_j = K \,,$$

we can write

$$\sum_i b_i = \left(\sum_i (\Psi \hat{a})_i \right) \left(\sum_j e^{-(\Psi a)_j} \right) = \frac{1}{\|a\|_1} \left(\sum_i (\Psi a)_i \right) \left(\sum_j e^{-(\Psi a)_j} \right) .$$

Hence,

$$\sum_i b_i = \frac{1}{\|a\|_1} \left(\sum_i x_i \right) \left(\sum_j e^{-x_j} \right) . \tag{17}$$

Substituting Eqns. (16) and (17) into Eq. (15) yields

$$\frac{d}{dt}(D_{\text{KL}}) = \frac{1}{(K-1)D\|a\|_1} \left[(K-1) \sum_j x_j e^{-x_j} - \left(\sum_i x_i \right) \left(\sum_j e^{-x_j} \right) \right] .$$

Now, the bracketed term can be rewritten as $\frac{1}{2} \sum_{ij} (x_i - x_j) \left(e^{-x_i} - e^{-x_j} \right)$, yielding

$$\frac{d}{dt}(D_{\text{KL}}) = \frac{1}{2(K-1)D\|a\|_1} \sum_{ij} (x_i - x_j) \left(e^{-x_i} - e^{-x_j} \right) .$$

For each pair (i, j), the terms $x_i - x_j$ and $e^{-x_i} - e^{-x_j}$ have opposite signs unless $x_i = x_j$. Hence, every summand is non-positive, vanishing only when $x_i = x_j$. Hence the full sum is always non-positive, and only equals zero if x is a uniform vector. Because $D > 0$ and $\|a\|_1 > 0$, the prefactor is strictly positive. Moreover, since Ψ is invertible and has 1_{K-1} as an eigenvector, the condition $x \propto 1_{K-1}$ implies $a \propto 1_{K-1}$, and hence $\hat{a} = \frac{1}{K-1} 1_{K-1}$.

Thus D_{KL} satisfies:

- $D_{\text{KL}} \geq 0$, with equality only when $\hat{a} = \frac{1}{K-1} 1_{K-1}$.
- $\frac{d}{dt}(D_{\text{KL}}) \leq 0$, with equality only at $\hat{a} = \frac{1}{K-1} 1_{K-1}$.

Consequently, D_{KL} decreases monotonically to its minimal value, guaranteeing convergence of $\hat{a}$ to $\frac{1}{K-1} 1_{K-1}$ as measured by this metric.

8 Related Work

Neural Collapse. First identified in deep classification networks by Papyan et al. [1], NC has since become a central focus of theoretical analysis.Researchers have extensively studied the conditions under which NC emerges, primarily using the UFM [16, 17] with explicit L_2 regularization. Within this framework, NC has been studied under a variety of loss functions [18, 16, 6, 19], consistently emerging as a globally optimal structure. Complementary landscape analyses have revealed favorable optimization properties, where the loss surface consists only of global minima and non-degenerate saddles [20, 14, 6]. Beyond the UFM, NC has also been investigated in other theoretical settings [21, 22]. Extensions of NC research have explored its behavior in settings with a large number of classes [23] and imbalanced class distributions [24, 7, 25, 17, 26, 27], where generalized forms of implicit geometries beyond the simplex ETF have been observed starting with [7]. Generalizations from the hard label setting have also been considered for the multilabel setting [28], soft-labeled language setting [29], and multi-class regression [30]. The influence of dataset properties has also been explored [31, 32]. Researchers have further asked whether NC extends beyond the final layer, with evidence coming both from empirical studies [33, 34, 35] and from analyses of deep UFMs [36, 26, 37, 38, 39, 40], as well as beyond the UFM framework altogether [41].

NC has also been connected to a range of practical aspects of deep learning, including adversarial robustness [42], transfer learning [43, 44], generalization [45, 46], and out-of-distribution detection [47, 48]. The impact of network architecture has been studied in ResNets [49, 50] and large language models (LLMs) [51, 52]. For a comprehensive overview of NC and its implications, see the survey by Kothapalli et al. [53].

Implicit Bias. A large body of work studies how gradient descent or gradient flow induces structure in the absence of explicit regularization. For logistic regression, Soudry et al. [2] showed convergence to the max-margin solution. Ji et al. [54] related the implicitly regularized path to explicitly regularized solutions in binary classification, which had been earlier studied by Rosset et al. [55]. Analyses of limiting behavior have since been extended to homogeneous [3, 4] and non-homogeneous networks [56], alternative optimizers [57, 58, 59, 60, 61, 62], large step-sizes [63, 64], as well as, convolutional architectures [5, 65] and transformers [66, 67, 68, 69]. Many of these works focus on binary settings, where no interaction arises between different singular modes. This includes [70], which shows that the singular vectors of the parameter matrices will align. A complementary thread investigates implicit bias in linear networks with MSE loss, examining the effects of gradient descent and network depth [71, 72, 73, 74], including matrix factorization [75, 76, 77, 78], and convergence properties [79]. Owing to the simplicity of MSE, however, these analyses differ substantially from our setting.

Implicit bias has also been studied directly in the unregularized CE UFM. Ji et al. [6] proved convergence to a KKT point of a constrained norm minimization problem. Thrampoulidis et al. [7] characterized the solution found by the regularization path under imbalanced classes. Our results are the first to explicitly characterize the implicit bias of gradient descent dynamics in unregularized settings, and thus provide the first definitive link between implicit bias theory and neural collapse geometries.

Exact Dynamics. Exact dynamics analyses study deep learning phenomena by reducing gradient flow to evolution equations for the singular values of the parameter matrices and solving these equations explicitly. To date, such works focus exclusively on linear networks trained with MSE loss. The approach was pioneered by Saxe et al. [8, 9] and extended by Gidel et al. [11] to weaker assumptions and to discrete gradient descent. These studies revealed that networks learn the hierarchical structure of singular values sequentially. Subsequent work generalized the technique to alternative initialization schemes [80, 81, 82] and to investigations of the lazy-to-rich regime transition [81, 83]. Recent efforts have sought to broaden the scope beyond linear architectures. Mainali et al. [84] adapted these methods to transformers, though with simplifications including linear attention mechanisms in place of softmax. Similarly, Zhao et al. [10] employed exact dynamics within the UFM to investigate semantic emergence from language data. However, all existing exact dynamics analyses—including these recent extensions—remain limited to MSE loss. Our work breaks this barrier by extending the framework to the CE loss, a more challenging yet practically essential setting.

9 Conclusion

In this work, we introduced a Hadamard Initialization scheme that preserves singular vector structure under the softmax operation, enabling an extension of the exact dynamics framework of Saxe et al. to the CE setting. Our analysis revealed striking differences from prior exact-dynamics results: the softmax nonlinearity couples the singular values, leading to complex dynamics, non-monotonic convergence, and exponential slowdowns. Such features are invisible to existing KKT-based analyses of the CE UFM. We further explained why such complex behaviors are rarely observed in practice, showing that standard initialization schemes naturally steer trajectories away from the regions where these effects manifest. Together, these findings provide the first path-level analysis of implicit bias in the unregularized CE UFM.

The Hadamard Initialization technique offers a tractable tool for studying CE dynamics and opens the door to extending exact-dynamics style analyses beyond linear MSE models. Natural next steps include examining how other results obtained via exact dynamics [11, 80, 81, 82, 83] can be recovered within our framework, as well as developing extensions to deep UFMs, settings with explicit regularization, and linear networks beyond the UFM construction. We hope this work motivates the use of these techniques to further investigate implicit bias and NC in more general, nonlinear architectures.

Acknowledgments

CG is supported by the Charles Coulson Scholarship. The authors also acknowledge support from His Majesty's Government in the development of this research. CT acknowledges generous support by NSERC. CT also thanks Yize Zhao for discussions on very early phases of this project. The authors are grateful to the organizers of the 2024 workshop "Beyond classical regimes in statistical inference and machine learning" in CIMI Toulouse where this collaboration was initiated. For the purpose of Open Access, the authors have applied a CC BY public copyright license to any Author Accepted Manuscript (AAM) version arising from this submission.

References

[1] Vardan Papyan, Xuemei Han, and David L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences of the United States of America*, 117:24652 – 24663, 2020.

[2] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70):1–57, 2018.

[3] Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. *arXiv preprint arXiv:1906.05890*, 2019.

[4] Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 17176–17186. Curran Associates, Inc., 2020.

[5] Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nathan Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Neural Information Processing Systems*, 2018.

[6] Wenlong Ji, Yiping Lu, Yiliang Zhang, Zhun Deng, and Weijie J Su. An unconstrained layer-peeled perspective on neural collapse. In *International Conference on Learning Representations*, 2022.

[7] Christos Thrampoulidis, Ganesh Ramachandra Kini, Vala Vakilian, and Tina Behnia. Imbalance trouble: Revisiting neural-collapse geometry. *Advances in Neural Information Processing Systems*, 35:27225–27238, 2022.

[8] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *CoRR*, abs/1312.6120, 2013.

[9] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116(23):11537–11546, 2019.

[10] Yize Zhao and Christos Thrampoulidis. On the geometry of semantics in next-token prediction. *arXiv preprint arXiv:2505.08348*, 2025.

[11] Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien. Implicit regularization of discrete gradient dynamics in linear neural networks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[12] Kathy J Horadam. *Hadamard matrices and their applications*. Princeton university press, 2012.

[13] Ross Cressman and Yi Tao. The replicator equation and other game dynamics. *Proceedings of the National Academy of Sciences*, 111(supplement_3):10810–10817, 2014.

[14] Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *Advances in Neural Information Processing Systems*, 34:29820–29834, 2021.

[15] Arthur Jacot, François Ged, Berfin Şimşek, Clément Hongler, and Franck Gabriel. Saddle-to-saddle dynamics in deep linear networks: Small initialization training, symmetry, and sparsity. *arXiv preprint arXiv:2106.15933*, 2021.

[16] Dustin G. Mixon, Hans Parshall, and Jianzong Pi. Neural collapse with unconstrained features. *Sampling Theory, Signal Processing, and Data Analysis*, 20, 2020.

[17] Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021.

[18] Jinxin Zhou, Chong You, Xiao Li, Kangning Liu, Sheng Liu, Qing Qu, and Zhihui Zhu. Are all losses created equal: A neural collapse perspective. In *Advances in Neural Information Processing Systems*, volume 35, pages 31697–31710. Curran Associates, Inc., 2022.

[19] X.Y. Han, Vardan Papyan, and David L. Donoho. Neural collapse under MSE loss: Proximity to and dynamics on the central path. In *International Conference on Learning Representations*, 2022.

[20] Jinxin Zhou, Xiao Li, Tianyu Ding, Chong You, Qing Qu, and Zhihui Zhu. On the optimization landscape of neural collapse under MSE loss: Global optimality with unconstrained features. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 27179–27202. PMLR, 17–23 Jul 2022.

[21] Arthur Jacot, Peter Súkeník, Zihan Wang, and Marco Mondelli. Wide neural networks trained with weight decay provably exhibit neural collapse. *arXiv preprint arXiv:2410.04887*, 2024.

[22] Emanuele Zangrando, Piero Deidda, Simone Brugiapaglia, Nicola Guglielmi, and Francesco Tudisco. Neural rank collapse: Weight decay and small within-class variability yield low-rank bias. *arXiv:2402.03991*, 2024.

[23] Jiachen Jiang, Jinxin Zhou, Peng Wang, Qing Qu, Dustin Mixon, Chong You, and Zhihui Zhu. Generalized neural collapse for a large number of classes. *arXiv preprint arXiv:2310.05351*, 2023.

[24] Yibo Yang, Shixiang Chen, Xiangtai Li, Liang Xie, Zhouchen Lin, and Dacheng Tao. Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network? *Advances in neural information processing systems*, 35:37991–38002, 2022.

[25] Tina Behnia, Ganesh Ramachandra Kini, Vala Vakilian, and Christos Thrampoulidis. On the implicit geometry of cross-entropy parameterizations for label-imbalanced data. In *International Conference on Artificial Intelligence and Statistics*, pages 10815–10838. PMLR, 2023.

[26] Hien Dang, Tan Minh Nguyen, Tho Tran, Hung The Tran, Hung Tran, and Nhat Ho. Neural collapse in deep linear networks: From balanced to imbalanced data. In *International Conference on Machine Learning*, 2023.

[27] Wanli Hong and Shuyang Ling. Neural collapse for unconstrained feature model under cross-entropy loss with imbalanced data. *arXiv preprint arXiv:2309.09725*, 2023.

[28] Pengyu Li, Xiao Li, Yutong Wang, and Qing Qu. Neural collapse in multi-label learning with pick-all-label loss. *arXiv preprint arXiv:2310.15903*, 2023.

[29] Yize Zhao, Tina Behnia, Vala Vakilian, and Christos Thrampoulidis. Implicit geometry of next-token prediction: From language sparsity patterns to model representations. *ArXiv*, abs/2408.15417, 2024.

[30] George Andriopoulos, Zixuan Dong, Li Guo, Zifan Zhao, and Keith Ross. The prevalence of neural collapse in neural multivariate regression. *arXiv preprint arXiv:2409.04180*, 2024.

[31] Wanli Hong and Shuyang Ling. Beyond unconstrained features: Neural collapse for shallow neural networks with general data. *arXiv preprint arXiv:2409.01832*, 2024.

[32] Vignesh Kothapalli and Tom Tirer. Kernel vs. kernel: Exploring how the data structure affects neural collapse. *arXiv preprint arXiv:2406.02105*, 2024.

[33] Hangfeng He and Weijie Su. A law of data separation in deep learning. *Proceedings of the National Academy of Sciences of the United States of America*, 120, 2022.

[34] Akshay Rangamani, Marius Lindegaard, Tomer Galanti, and Tomaso A Poggio. Feature learning in deep classifiers through intermediate neural collapse. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 28729–28745. PMLR, 23–29 Jul 2023.

[35] Liam Parker, Emre Onal, Anton Stengel, and Jake Intrater. Neural collapse in the intermediate hidden layers of classification neural networks. *ArXiv*, abs/2308.02760, 2023.

[36] Tom Tirer and Joan Bruna. Extended unconstrained features model for exploring deep neural collapse. In *International Conference on Machine Learning*, 2022.

[37] Connall Garrod and Jonathan P Keating. Unifying low dimensional observations in deep learning through the deep linear unconstrained feature model. *arXiv preprint arXiv:2404.06106*, 2024.

[38] Peter Súkeník, Marco Mondelli, and Christoph H Lampert. Deep neural collapse is provably optimal for the deep unconstrained features model. In *Advances in Neural Information Processing Systems*, volume 36, pages 52991–53024. Curran Associates, Inc., 2023.

[39] Peter Súkeník, Marco Mondelli, and Christoph Lampert. Neural collapse versus low-rank bias: Is deep neural collapse really optimal? *arXiv preprint arXiv:2405.14468*, 2024.

[40] Connall Garrod and Jonathan P. Keating. The persistence of neural collapse despite low-rank bias: An analytic perspective through unconstrained features. *ArXiv*, abs/2410.23169, 2024.

[41] Daniel Beaglehole, Peter Súkeník, Marco Mondelli, and Mikhail Belkin. Average gradient outer product as a mechanism for deep neural collapse. *arXiv preprint arXiv:2402.13728*, 2024.

[42] Jingtong Su, Ya Shi Zhang, Nikolaos Tsilivis, and Julia Kempe. On the robustness of neural collapse and the neural collapse of robustness. *arXiv preprint arXiv:2311.07444*, 2023.

[43] Tomer Galanti, András György, and Marcus Hutter. On the role of neural collapse in transfer learning. *arXiv preprint arXiv:2112.15121*, 2021.

[44] Xiao Li, Sheng Liu, Jinxin Zhou, Xinyu Lu, Carlos Fernandez-Granda, Zhihui Zhu, and Qing Qu. Principled and efficient transfer learning of deep models via neural collapse. *arXiv preprint arXiv:2212.12206*, 2022.

[45] Like Hui, Mikhail Belkin, and Preetum Nakkiran. Limitations of neural collapse for understanding generalization in deep learning. *arXiv preprint arXiv:2202.08384*, 2022.

[46] Peifeng Gao, Qianqian Xu, Yibo Yang, Peisong Wen, Huiyang Shao, Zhiyong Yang, Bernard Ghanem, and Qingming Huang. Towards demystifying the generalization behaviors when neural collapse emerges. *ArXiv*, abs/2310.08358, 2023.

[47] Yingwen Wu, Ruiji Yu, Xinwen Cheng, Zhengbao He, and Xiaolin Huang. Pursuing feature separation based on neural collapse for out-of-distribution detection. *arXiv preprint arXiv:2405.17816*, 2024.

[48] Jiawei Zhang, Yufan Chen, Cheng Jin, Lei Zhu, and Yuantao Gu. Epa: Neural collapse inspired robust out-of-distribution detector. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6515–6519. IEEE, 2024.

[49] Jianing Li and Vardan Papyan. Residual alignment: uncovering the mechanisms of residual networks. *Advances in Neural Information Processing Systems*, 36, 2024.

[50] Peter S'uken'ik, Christoph H. Lampert, and Marco Mondelli. Neural collapse is globally optimal in deep regularized resnets and transformers. *ArXiv*, abs/2505.15239, 2025.

[51] Robert Wu and Vardan Papyan. Linguistic collapse: Neural collapse in (large) language models. *arXiv preprint arXiv:2405.17767*, 2024.

[52] Yize Zhao and Christos Thrampoulidis. On the geometry of semantics in next-token prediction. *ArXiv*, abs/2505.08348, 2025.

[53] Vignesh Kothapalli, Ebrahim Rasromani, and Vasudev Awatramani. Neural collapse: A review on modelling principles and generalization. *Trans. Mach. Learn. Res.*, 2023, 2022.

[54] Ziwei Ji, Miroslav Dudík, Robert E. Schapire, and Matus Telgarsky. Gradient descent follows the regularization path for general losses. In *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 2109–2136. PMLR, 09–12 Jul 2020.

[55] Saharon Rosset, Ji Zhu, and Trevor Hastie. Margin maximizing loss functions. *Advances in neural information processing systems*, 16, 2003.

[56] Yuhang Cai, Kangjie Zhou, Jingfeng Wu, Song Mei, Michael Lindsey, and Peter L Bartlett. Implicit bias of gradient descent for non-homogeneous deep networks. *arXiv preprint arXiv:2502.16075*, 2025.

[57] Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*, pages 1832–1841. PMLR, 2018.

[58] Navid Azizan and Babak Hassibi. Stochastic gradient/mirror descent: Minimax optimality and implicit regularization. *arXiv preprint arXiv:1806.00952*, 2018.

[59] Haoyuan Sun, Kwangjun Ahn, Christos Thrampoulidis, and Navid Azizan. Mirror descent maximizes generalized margin and can be implemented efficiently. *Advances in Neural Information Processing Systems*, 35:31089–31101, 2022.

[60] Scott Pesme, Radu-Alexandru Dragomir, and Nicolas Flammarion. Implicit bias of mirror flow on separable data. *Advances in Neural Information Processing Systems*, 37:114094–114121, 2024.

[61] Chenyang Zhang, Difan Zou, and Yuan Cao. The implicit bias of adam on separable data. *Advances in Neural Information Processing Systems*, 37:23988–24021, 2024.

[62] Chen Fan, Mark Schmidt, and Christos Thrampoulidis. Implicit bias of spectral descent and muon on multiclass separable data. *arXiv preprint arXiv:2502.04664*, 2025.

[63] Mathieu Even, Scott Pesme, Suriya Gunasekar, and Nicolas Flammarion. (s) gd over diagonal linear networks: Implicit bias, large stepsizes and edge of stability. *Advances in Neural Information Processing Systems*, 36:29406–29448, 2023.

[64] Jingfeng Wu, Vladimir Braverman, and Jason D Lee. Implicit bias of gradient descent for logistic regression at the edge of stability. *Advances in Neural Information Processing Systems*, 36:74229–74256, 2023.

[65] Hannah Lawrence, Kristian Georgiev, Andrew Dienes, and Bobak T Kiani. Implicit bias of linear equivariant networks. *arXiv preprint arXiv:2110.06084*, 2021.

[66] Davoud Ataee Tarzanagh, Yingcong Li, Xuechen Zhang, and Samet Oymak. Max-margin token selection in attention mechanism. *Advances in neural information processing systems*, 36:48314–48362, 2023.

[67] Davoud Ataee Tarzanagh, Yingcong Li, Christos Thrampoulidis, and Samet Oymak. Transformers as support vector machines. *arXiv preprint arXiv:2308.16898*, 2023.

[68] Bhavya Vasudeva, Puneesh Deora, and Christos Thrampoulidis. Implicit bias and fast convergence rates for self-attention. *Transactions on Machine Learning Research*, 2024.

[69] Addison Kristanto Julistiono, Davoud Ataee Tarzanagh, and Navid Azizan. Optimizing attention with mirror descent: Generalized max-margin token selection. *arXiv preprint arXiv:2410.14581*, 2024.

[70] Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations*, 2019.

[71] Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 244–253. PMLR, 10–15 Jul 2018.

[72] Can Yaras, P. Wang, Wei Hu, Zhihui Zhu, Laura Balzano, and Qing Qu. The law of parsimony in gradient descent for learning deep linear networks. *ArXiv*, abs/2306.01154, 2023.

[73] Zhenfeng Tu, Santiago Aranguri, and Arthur Jacot. Mixed dynamics in linear networks: Unifying the lazy and active regimes. In *Advances in Neural Information Processing Systems*, volume 37, pages 106059–106104. Curran Associates, Inc., 2024.

[74] Bubacarr Bah, Holger Rauhut, Ulrich Terstiege, and Michael Westdickenberg. Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers. *Information and Inference: A Journal of the IMA*, 11(1):307–353, 2022.

[75] Zhiyuan Li, Yuping Luo, and Kaifeng Lyu. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. *arXiv preprint arXiv:2012.09839*, 2020.

[76] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[77] Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. *Advances in neural information processing systems*, 33:21174–21187, 2020.

[78] Suriya Gunasekar, Blake E. Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nathan Srebro. Implicit regularization in matrix factorization. *2018 Information Theory and Applications Workshop (ITA)*, pages 1–10, 2017.

[79] Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. *arXiv preprint arXiv:1810.02281*, 2018.

[80] Daniel Kunin, Allan Ravent'os, Cl'ementine Carla Juliette Domin'e, Feng Chen, David Klindt, Andrew Saxe, and Surya Ganguli. Get rich quick: exact solutions reveal how unbalanced initializations promote rapid feature learning. *ArXiv*, abs/2406.06158, 2024.

[81] Lukas Braun, Clémentine Dominé, James Fitzgerald, and Andrew Saxe. Exact learning dynamics of deep linear networks with prior knowledge. In *Advances in Neural Information Processing Systems*, volume 35, pages 6615–6629. Curran Associates, Inc., 2022.

[82] Salma Tarmoun, Guilherme Franca, Benjamin D Haeffele, and Rene Vidal. Understanding the dynamics of gradient flow in overparameterized linear models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10153–10161. PMLR, 18–24 Jul 2021.

[83] Cl'ementine Carla Juliette Domin'e, Nicolas Anguita, Alexandra M Proca, Lukas Braun, Daniel Kunin, Pedro A. M. Mediano, and Andrew Saxe. From lazy to rich: Exact learning dynamics in deep linear networks. *ArXiv*, abs/2409.14623, 2024.

[84] Nischal Mainali and Lucas Teixeira. Exact learning dynamics of in-context learning in linear transformers and its application to non-linear transformers. *ArXiv*, abs/2504.12916, 2025.

[85] Mor Shpigel Nacson, Jason Lee, Suriya Gunasekar, Pedro Henrique Pamplona Savarese, Nathan Srebro, and Daniel Soudry. Convergence of gradient descent on separable data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3420–3428. PMLR, 2019.

Table of Contents

1 Introduction **1**

1.1 Motivation . . . 1

1.2 Contributions . . . 2

2 Background **3**

3 Hadamard Initialization **4**

3.1 Diagonalizing the Softmax . . . 4

3.2 Tractable CE Dynamics . . . 5

3.3 Properties of the Reduced Dynamics . . . 6

4 CE Dynamics and Implicit Bias **7**

4.1 NC as a Candidate Among Stable Directions . . . 7

4.2 Convergence to Neural Collapse via a Lyapunov Function . . . 7

4.3 Non-monotonic Behaviors . . . 8

4.4 Further Comparisons to MSE Dynamics . . . 9

4.4.1 Linearized Dynamics . . . 9

4.4.2 Comparison to Logit Optimization . . . 9

5 Standard Initialization Schemes lead to NC **10**

6 Numerical Experiments **11**

6.1 Setup . . . 11

6.2 Results . . . 11

7 Proofs **12**

7.1 Proof of Theorem 1 . . . 12

7.2 Proof of Theorem 3 . . . 13

8 Related Work **14**

9 Conclusion **15**

A Summary of Notations **22**

B Further Theoretical Results **22**

B.1 Logit singular value dynamics as mirror flow . . . 22

B.2 NC is the only full-rank stable direction . . . 23

B.3 Additional examples of non-monotonic behavior . . . 23

B.3.1 Reverse-KL metric . . . 23

B.3.2 Frobenius Distance . . . 23

B.3.3 Log-ratio of max-min singular values . . . 24

C Non-uniform convergence points **25**

D Proofs **26**

D.1 Proof of Theorem 2 26

D.2 Proof of Theorem 4 27

D.3 Proof of Theorem 5 29

D.4 Proof of Theorem 6 30

D.5 Proof of Theorem 7 32

D.6 Proof of Theorem 8 35

E Additional Experiments **36**

E.1 Metrics 36

E.2 Larger K 36

E.3 The effect of initialization scale 37

E.4 The effect of overparameterization 38

E.5 Normalized Gradient Flow 40

A Summary of Notations

Table 1: Summary of Key Notations

Notation	Description	Context / Definition
K	Number of classes.	Assumed to be $K = 2^m$.
n	Number of samples per class.	Assumed to be balanced.
W	Last-layer weight matrix.	$\mathbb{R}^{K\times d}$
H	Feature matrix.	$\mathbb{R}^{d\times Kn}$
Z	Logit matrix.	$Z = WH$
Y	One-hot label matrix.	$Y = I_K \otimes 1_n^T$
$P(Z)$	Softmax matrix.	Softmax applied columnwise to Z.
S	Simplex ETF matrix.	$S = I_K - \frac{1}{K}1_K 1_K^T$
Φ	Sylvester Hadamard matrix.	$\mathbb{R}^{K\times K}$ matrix (Def. 2).
Ψ	Reduced dynamics matrix.	$1_{K-1}1_{K-1}^T - \Phi[2:K, 2:K]$.
U, V	Left/Right singular vector matrices.	$U = V = \frac{1}{\sqrt{K}}\Phi$ for Hadamard Init.
α_i, β_i	Singular values.	Non-trivial singular values of W and H.
a	Logit singular values (vector).	$a = (a_1, \ldots, a_{K-1})$, $a_i = \alpha_i\beta_i$.
$\hat{a}$	Normalized logit singular values.	$\hat{a}_i = a_i/\|a\|_1$.
b, D	Terms in reduced dynamics.	Defined in Eq. 10.
$E(a)$	Energy functional.	Reduced energy functional for a (Eq. 8).
M	L_2-based metric for NC.	Measures distance from $\hat{a}$ to uniform (Eq. 12).
D_{KL}	KL divergence.	Used as a Lyapunov function (Thm. 3).

Throughout the paper, we use $\mathbb{R}$ to denote the set of real numbers. I_k represents the $k \times k$ identity matrix, and 1_k denotes an all-ones vector of dimension k. For a given matrix A, A^T is its transpose, and $A[i_1 : i_2, j_1 : j_2]$ indicates the submatrix sliced from rows i_1 to i_2 and columns j_1 to j_2. The operator $\mathrm{diag}(v)$ forms a diagonal matrix using the elements of vector v. We use $\otimes$ for the Kronecker product, $\odot$ for the Hadamard (elementwise) product and $\oplus$ for the bitwise XOR. Finally, $\|\cdot\|_F$, and $\|\cdot\|_1$ denote the Frobenius norm and the L_1 norm, respectively. Some key notation is summarized in Table 1.

B Further Theoretical Results

B.1 Logit singular value dynamics as mirror flow

Recall the dynamics of the logit singular values from Eq. (9) are given as

$$\frac{da_i}{dt} = a_i \cdot \frac{\sum_{j=1}^{K-1} \Psi_{ij} e^{-(\Psi a)_j}}{1 + \sum_{j=1}^{K-1} e^{-(\Psi a)_j}}, \qquad i \in [K-1]. \tag{18}$$

Define the energy functional

$$E(a) = \log\left(1 + \sum_{j=1}^{K-1} \exp\left(-(\Psi a)_j\right)\right).$$

Then, the dynamics in Eq. (18) can be expressed conveniently in vector form as

$$\frac{\mathrm{d}a}{\mathrm{d}t} = -a \odot \nabla_a E(a).$$

For positive initialization, $a > 0$ in the entire trajectory, and thus it holds equivalently that

$$\frac{\mathrm{d}\log(a)}{\mathrm{d}t} = -\nabla_a E(a).$$

But, $\frac{\mathrm{d}\log(a)}{\mathrm{d}t} = \frac{\mathrm{d}\nabla_a \phi(a)}{\mathrm{d}t}$ where $\phi(a) = \sum_{i=1}^{K-1} a_i \log(a_i)$ is the negative entropy potential. Overall, we have expressed the dynamics of the logit singular values in the form of mirror-flow updates with respect to the entropy potential:

$$\frac{\mathrm{d}\nabla_a \phi(a)}{\mathrm{d}t} = -\nabla_a E(a).$$

B.2 NC is the only full-rank stable direction

Given the dynamics of Eq. (11), if we look for a stable direction, meaning $\frac{d\hat{a}_i}{dt} = 0$ for $i = 1, ..., K-1$, that also satisfies full-rank, meaning $\hat{a}_i > 0$ for $i = 1, \dots, K-1$, then this reduces to the condition:

$$b_i = \bar{b}, \quad \text{for } i = 1, ..., K-1.$$

This means that the vector b can be written as $b = \bar{b}1_{K-1}$, which imposes the following condition on the vector a

$$\Psi e^{-(\Psi a)} = \bar{b}1_{K-1}.$$

Then using that Ψ is invertible and $\Psi^{-1}1_{K-1} = \frac{1}{K}1_{K-1}$ gives:

$$a = -\frac{1}{K}\log\left(\frac{1}{K}\bar{b}\right)1_{K-1},$$

and if $a \propto 1_{K-1}$, it implies $\hat{a} = \frac{1}{K-1}1_{K-1}$. Hence, the only full-rank stable direction is $\hat{a} = \frac{1}{K-1}1_{K-1}$.

B.3 Additional examples of non-monotonic behavior

B.3.1 Reverse-KL metric

One can consider the reverse-KL divergence, defined as

$$D_{\text{KL}}\left(\hat{a} \,\|\, \frac{1}{K-1}1_{K-1}\right) = \log(K-1) + \sum_i \hat{a}_i \log(\hat{a}_i).$$

We will denote it by $\tilde{D}_{\text{KL}}$ for this section so as to differentiate it from the KL divergence of the main text. It has derivative

$$\frac{d\tilde{D}_{\text{KL}}}{dt} = \frac{1}{D}\sum_i [b_i - \bar{b}]\hat{a}_i \log(\hat{a}_i).$$

In $K = 8$, a counter-example, demonstrating $\tilde{D}_{\text{KL}}$ is not monotonic decreasing, occurs for

$$a = [1, 1, 1, 1.25, 0.0001, 0.0001, 0.0001],$$

where $\frac{d\tilde{D}_{\text{KL}}}{dt} \approx 0.0037$.

B.3.2 Frobenius Distance

Typically in other works, the rate of convergence to NC is measured through the following metric:

$$\tilde{M} = \left\| \frac{Z}{\|Z\|_F} - \frac{S}{\|S\|_F} \right\|_F^2,$$

which, when expressed in terms of our logit singular values (after absorbing constants as in Eq. (9)), is given by

$$\tilde{M} = \sum_{i=1}^{K-1}\left[\frac{a_i}{\sqrt{\sum_j a_j^2}} - \frac{1}{\sqrt{K-1}}\right]^2.$$

In the main text, to avoid introducing multiple normalizations, we normalized with the L_1 norm, but the metric $\tilde{M}$ displays the same behavior as the metric M. First note that its derivative is given by

$$\frac{d\tilde{M}}{dt} = \frac{2}{D\|Z\|_F^3\sqrt{K-1}}\left[\sum_{ij} a_i a_j b_j (a_j - a_i)\right],$$

where we continue to denote $\sqrt{\sum_j a_j} = \|Z\|_F$. For this to be non-positive we require that:

$$\sum_{ij} a_i a_j b_j [a_j - a_i] \leq 0.$$

A counter-example occurs for $a = [1, 1, 1, 1.25, 0.0001, 0.0001, 0.0001]$, where the LHS of the above equality is ≈ 0.088. Note the existence of this counter-example for $K = 8$ is sufficient for counter-examples to exist for all $K \geq 8$, using the same reasoning as provided for the proof of Theorem 4 in Appendix D.2.

B.3.3 Log-ratio of max-min singular values

Define the following metric:

$$M_{\text{LR}} = \max_i(\log(a_i)) - \min_j(\log(a_j)).$$

This quantity is monotonic decreasing for $K \leq 8$, but non-monotonic for $K \geq 16$.

We first prove this quantity is monotonic decreasing for $K = 8$:

Proposition 1. *Let $K = 8$, and $a(t) \in \mathbb{R}^{K-1}$ be a variable initialized with $a_i(0) > 0$ for all i, that evolves under the time evolution of Eq. (9). Then the metric M_{LR}, is monotonic decreasing.*

Proof:

Recall our dynamics, given in Eq. (9). By multiplying the numerator and denominator by $\exp(2\sum_j a_j)$, we can rewrite this as

$$\frac{d}{dt}(a_i) = \frac{1}{D} b_i a_i,$$

where

$$b_i = \sum_j \Psi_{ij} \exp\left(\left(\left[2(1_{K-1}1_{K-1}^T) - \Psi\right] a\right)_j\right), \quad D = \exp(2\sum_j a_j) + \sum_j \exp\left(\left(\left[2(1_{K-1}1_{K-1}^T) - \Psi\right] a\right)_j\right).$$

Expanding this out directly produces the following expressions for the variables b_i

$$\begin{aligned}
b_1 &= \exp(2(a_2 + a_4 + a_6)) + \exp(2(a_3 + a_4 + a_7)) + \exp(2(a_2 + a_5 + a_7)) + \exp(2(a_3 + a_5 + a_6)) \\
b_2 &= \exp(2(a_1 + a_4 + a_5)) + \exp(2(a_3 + a_4 + a_7)) + \exp(2(a_1 + a_6 + a_7)) + \exp(2(a_3 + a_5 + a_6)) \\
b_3 &= \exp(2(a_2 + a_4 + a_6)) + \exp(2(a_1 + a_4 + a_5)) + \exp(2(a_2 + a_5 + a_7)) + \exp(2(a_1 + a_6 + a_7)) \\
b_4 &= \exp(2(a_1 + a_2 + a_3)) + \exp(2(a_2 + a_5 + a_7)) + \exp(2(a_1 + a_6 + a_7)) + \exp(2(a_3 + a_5 + a_6)) \\
b_5 &= \exp(2(a_2 + a_4 + a_6)) + \exp(2(a_3 + a_4 + a_7)) + \exp(2(a_1 + a_2 + a_3)) + \exp(2(a_1 + a_6 + a_7)) \\
b_6 &= \exp(2(a_1 + a_4 + a_5)) + \exp(2(a_3 + a_4 + a_7)) + \exp(2(a_1 + a_2 + a_3)) + \exp(2(a_2 + a_5 + a_7)) \\
b_7 &= \exp(2(a_2 + a_4 + a_6)) + \exp(2(a_1 + a_4 + a_5)) + \exp(2(a_1 + a_2 + a_3)) + \exp(2(a_3 + a_5 + a_6))
\end{aligned}$$

From this one can note the following properties immediately:

1. There are seven triples of the singular values a_i that appear in the exponents.
2. The triples that feature a_i do not appear in the expression of b_i.
3. Each a_i appears with every other a_j in exactly one triple.
4. Every pair of b_i, b_j share exactly two terms, and differ in two. The two they share are precisely the two not featuring either of their corresponding a_i.

Now suppose that at some time t we have that the largest mode is given by a_i and the smallest by a_j, then the log ratio of these singular values has derivative given by

$$\frac{d}{dt}\left(\log(a_i) - \log(a_j)\right) = \frac{1}{D}[b_i - b_j].$$

Because of the properties of the expressions for the variables b_i described above, the difference $b_i - b_j$ takes the following form:

$$b_i - b_j = x_j x_p x_q + x_j x_r x_s - x_i x_r x_p - x_i x_s x_q,$$

where we have defined $x_i = \exp(2a_i)$ for the purpose of this section. We can alternatively write this expression as

$$b_i - b_j = x_p[x_j x_q - x_i x_r] + x_s[x_j x_r - x_i x_q],$$

and further as

$$b_i - b_j = x_p x_i x_q \left[\frac{x_j}{x_i} - \frac{x_r}{x_q}\right] + x_s x_i x_r \left[\frac{x_j}{x_i} - \frac{x_q}{x_r}\right].$$

Now note that, because x_j is the smallest, and x_i is the largest, we must have:

$$x_j \leq \text{any other choice of } x, \qquad \frac{1}{x_i} \leq \frac{1}{\text{any other choice of } x}.$$

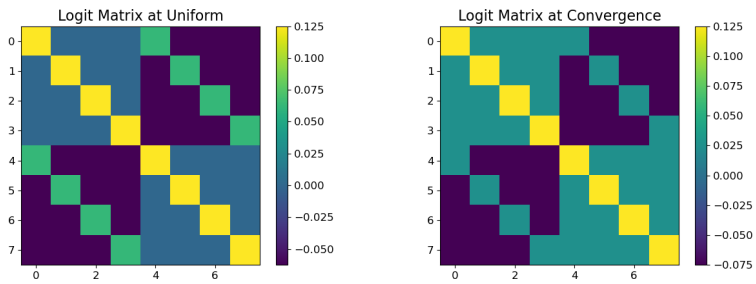

Figure 3: Plots of the normalized logit matrices generated by the vector $a^* = [1, 1, 1, 1, 0, 0, 0]$, which is uniform on its support, and by the non-uniform vector $a_{\text{conv}} = [1, 1, 1, 2, 0, 0, 0]$. The smallest normalized margin achieved by a_{conv} is visibly larger, indicating that this configuration outperforms the uniform one as the logit norm diverges.

This implies that:

$$\frac{x_j}{x_i} \leq \frac{x_r}{x_q}, \qquad \frac{x_j}{x_i} \leq \frac{x_q}{x_r},$$

and hence the difference $b_i - b_j$ is non-positive, and only takes value zero when

$$\frac{x_j}{x_i} = \frac{x_r}{x_q} = \frac{x_q}{x_r},$$

but this can only happen if $x_j = x_i$, meaning the largest and smallest mode must be equal, which implies $\hat{a} = \frac{1}{K-1} 1_{K-1}$. Hence the quantity M_{LR} always has a negative derivative, except when we have a neural collapsed solution.

□

It remains to show a counter-example in $K = 16$. Let $a = [5.01, 0.99, 2, 5, 1, 5, 1, 2, 1, 5, 3, 5, 1, 5, 1]$, then the largest singular value is a_1, and the smallest is a_2, but

$$\frac{d}{dt} \left(\log(a_1) - \log(a_2) \right) \approx 1.8 \times 10^{-13}$$

Hence this metric is not monotonic decreasing for $K = 16$.

C Non-uniform convergence points

Here we provide an expanded discussion of the example shown in Figure 1, in order to build intuition for why non-uniform convergence points can arise in the loss surface for general K. Figure 3 displays the normalized logit matrix with singular-values vector $a^* = [1, 1, 1, 1, 0, 0, 0]$ that lies very close to the initialization used in Figure 1. Since the set of indices for which $a_i \neq 0$ remains stable under the flow, we might expect that the convergence point reached from initialization $a(0) = a^*$ is the minimizer among all configurations with at most the first four indices nonzero.

It is straightforward to show that the following quantity decreases monotonically along the flow:

$$E(a) = \log \left(1 + \sum_{j=1}^{K-1} \exp \left(-\|a\|_1 (\Psi \hat{a})_j \right) \right),$$

where we have chosen to express it in terms of $\hat{a}$ and $\|a\|_1$. This coincides with the energy functional of Eq. (8) up to absorbed constants. In particular, as $\|a\|_1 \to \infty$, the energy is asymptotically:

$$E \simeq \sum_{j = \text{argmin}_u \{(\Psi \hat{a})_u\}} \exp \left(-\|a\|_1 (\Psi \hat{a})_j \right) \Big).$$

Thus, the smallest normalized margin $(\Psi \hat{a})_i$ dominates the energy, and we therefore expect the flow as $t \to \infty$ to maximize this quantity. The normalized margins correspond to the difference between the diagonal and off-diagonal entries of the normalized logit matrix.

Comparing the normalized logit matrix at the uniform configuration a^* with that of the configuration a_{conv}, as shown in Figure 3, we see that a^* has a smaller minimal margin.

Numerically, what occurs in Figure 1 of the main text is that the last three singular values are initially so small that the dynamics treats them effectively as zero. The remaining singular values, therefore, evolve toward the minimizer supported on only the first four indices. This minimizer is not the uniform point, as indicated by its worse margin in Figure 3; instead, it is the configuration a_{conv},in which one singular value is larger than the others. This is the structure toward which the flow in Figure 1 trends over the considered time interval. Because the flow approaches a non-uniform point, most common measures of uniformity increase, and thus fail to act as Lyapunov functions for the dynamics. The KL divergence avoids this issue by effectively accounting for the growth of the three small singular values, which offsets the non-uniform behavior of the remaining ones.

D Proofs

D.1 Proof of Theorem 2

Recall $U = \frac{1}{\sqrt{K}}\Phi$, where Φ is the $K \times K$ Sylvester Hadamard matrix, and that we initialize our parameter matrices as: $W = UD_W R^T$, $H = RD_H V^T$, where $R \in \mathbb{R}^{d\times d}$ is orthogonal, and $V = U \otimes Q$ is orthogonal, with Q being a right singular matrix of 1_n^T, meaning it obeys the relation: $1_n^T = \sqrt{n} e_1^{(n)T} Q^T$, where $e_1^{(n)}$ is the first standard basis vector in $\mathbb{R}^n$. Also recall that $D_W \in \mathbb{R}^{K\times d}$ has its only non-zero singular values given by $\alpha_1, \dots \alpha_K$, and $D_H \in \mathbb{R}^{d\times Kn}$ has its only non-zero singular values given by $\beta_1, \dots, \beta_K$. Note as a consequence of the definitions of U, V, we can write $Y = \sqrt{n} U[I_K, 0_{K\times K(n-1)}]V^T$.

We also have at initialization that the logit matrix is given by

$$Z = UD_W D_H V^T.$$

Note that $D_W \in \mathbb{R}^{K\times d}$ and $D_H \in \mathbb{R}^{d\times Kn}$. Hence, defining $D_Z = D_W D_H \in \mathbb{R}^{K\times Kn}$, and $D_a = \text{diag}(a_1, ..., a_K)$, where $a_i = \alpha_i \beta_i$, this matrix can be written as:

$$D_Z = [D_a, 0_{K\times K(n-1)}] = D_a \otimes e_1^{(n)T}.$$

Hence we can write

$$Z = U\left(D_a \otimes e_1^{(n)T}\right)(U \otimes Q)^T = U\left[(D_a U^T) \otimes (e_1^{(n)T} Q^T)\right].$$

Next, using that $1_n^T = \sqrt{n} e_1^{(n)T} Q^T$, this becomes:

$$Z = \frac{1}{\sqrt{n}} U\left[(D_a U^T) \otimes 1_n^T\right].$$

Then using that this has repeated columns, this becomes:

$$Z = (UD_{\tilde{a}} U^T) \otimes 1_n^T,$$

where we define $\tilde{a} = \frac{1}{\sqrt{n}} a$. Since this is the same matrix as $UD_{\tilde{a}}U^T$, just with repeated columns, we can apply Theorem 1 to calculate the softmax matrix at initialization:

$$P(Z) = (UD_{\tilde{\nu}} U^T) \otimes 1_n^T,$$

where

$$\tilde{\nu}_i = \frac{\sum_j \Phi_{ij} e^{\frac{1}{K}(\Phi\tilde{a})_j}}{\sum_j e^{\frac{1}{K}(\Phi\tilde{a})_j}}.$$

Now consider the gradient flow equations of our model

$$\frac{dW}{dt} = (Y - P)H^T, \qquad \frac{dH}{dt} = W^T(Y - P).$$

Consider the change of variable given by $W = U\tilde{W}$, $H = \tilde{H}V^T$, these have evolution equations given by:

$$\frac{d}{dt}(\tilde{W}) = U^T(Y - P)V\tilde{H}^T, \qquad \frac{d}{dt}(\tilde{H}) = \tilde{W}^T U^T (Y - P)V.$$

Note that, by construction, both the matrices Y, P can be written in SVD form through the matrices U, V

$$Y = \sqrt{n} U D_Y V^T, \qquad P = \sqrt{n} U D_P V^T,$$

where $D_Y = [I_K, 0_{K\times K(n-1)}] \in \mathbb{R}^{K\times Kn}$ and $D_P = [\text{diag}(\tilde{\nu}), 0_{K\times K(n-1)}] \in \mathbb{R}^{K\times Kn}$, and so

$$U^T(Y-P)V = \sqrt{n}[\text{diag}(1_K - \tilde{\nu}), 0_{K\times K(n-1)}],$$

label this matrix as B. Now returning to the evolution equations for $\tilde{H}$ and $\tilde{W}$, we label the rows of $\tilde{W}$ by $\tilde{w}_a$, and the columns of $\tilde{H}$ by $\tilde{h}_a$. These have evolution equations given by:

$$\frac{d}{dt}((\tilde{w}_i)_j) = \sum_x B_{ix}(\tilde{h}_x)_j, \qquad \frac{d}{dt}((\tilde{h}_j)_i) = \sum_x B_{xj}(\tilde{w}_x)_i.$$

Now using that $B_{ix} = \sqrt{n}(1-\tilde{\nu}_i)\delta_{ix}$ gives:

$$\frac{d}{dt}((\tilde{w}_i)_j) = \sqrt{n}(1-\tilde{\nu}_i)(\tilde{h}_i)_j, \qquad \frac{d}{dt}((\tilde{h}_j)_i) = \sqrt{n}(1-\tilde{\nu}_j)(\tilde{w}_j)_i, \tag{19}$$

where in the second of these equations if $j > K$ the derivative is zero.

Note that using our initialization scheme, we have at initialization $\tilde{w}_i = \alpha_i r_i$, $\tilde{h}_i = \beta_i r_i$, for $i = 1, \ldots, K$, and 0 otherwise, where r_i are the columns of R. One observes from Eq. (19) that the derivatives also point in the same direction. Hence this property is maintained at all times, meaning the vectors $\tilde{w}_i, \tilde{h}_i$ only evolve in their coefficients, not in their directions. These coefficients then evolve as:

$$\frac{d}{dt}(\alpha_i) = \sqrt{n}(1-\tilde{\nu}_i)\beta_i, \qquad \frac{d}{dt}(\beta_i) = \sqrt{n}(1-\tilde{\nu}_i)\alpha_i.$$

Translating back to our original parameter matrices W, H, these are the evolution equations for their singular values, and the singular vectors are frozen for all times.

It remains to demonstrate these equations take the exact form in the theorem statement. Note from the expression of $\tilde{\nu}_i$ we have

$$1-\tilde{\nu}_i = \frac{1}{\sum_j e^{\frac{1}{K}(\Phi\tilde{a})_j}}\left[\sum_j (1-\Phi_{ij})e^{\frac{1}{K}(\Phi\tilde{a})_j}\right].$$

Then dividing the numerator and denominator by $\exp(\frac{1}{K}\sum_j \tilde{a}_j)$, gives:

$$1-\tilde{\nu}_i = \frac{1}{\sum_j e^{-\frac{1}{K}([1_K 1_K^T - \Phi]\tilde{a})_j}}\left[\sum_j [1_K 1_K^T - \Phi]_{ij} e^{-\frac{1}{K}([1_K 1_K^T - \Phi]\tilde{a})_j}\right] = \frac{b_i}{D},$$

and hence this gives the desired form, once we recall that $\tilde{a}_i = \frac{1}{\sqrt{n}}a_i$.

D.2 Proof of Theorem 4

Recall our choice of metric is given by:

$$M = \frac{1}{2}\sum_{i=1}^{K-1}\left(\hat{a}_i - \frac{1}{K-1}\right)^2.$$

A simple calculation, using the derivatives provided in Eq. (11), gives

$$\frac{dM}{dt} = \frac{1}{D}\sum_i \hat{a}_i^2[b_i - \bar{b}]. \tag{20}$$

First note that $M = 0$ in the $K = 2$ case for all positive initializations, and so trivially is monotonic. We begin by proving monotonicity for $K = 4$.

Note in $K = 4$ we have

$$\Psi = \begin{bmatrix} 2 & 0 & 2 \\ 0 & 2 & 2 \\ 2 & 2 & 0 \end{bmatrix} \implies \Psi a = \begin{bmatrix} 2a_1 + 2a_3 \\ 2a_2 + 2a_3 \\ 2a_1 + 2a_2 \end{bmatrix} \implies \exp(-\Psi a) = \begin{bmatrix} e^{-2a_1-2a_3} \\ e^{-2a_2-2a_3} \\ e^{-2a_1-2a_2} \end{bmatrix}.$$

Hence

$$\frac{1}{D}b = \frac{2}{1+e^{-2a_1-2a_3}+e^{-2a_2-2a_3}+e^{-2a_1-2a_2}}\begin{bmatrix} e^{-2a_1-2a_3}+e^{-2a_1-2a_2} \\ e^{-2a_2-2a_3}+e^{-2a_1-2a_2} \\ e^{-2a_1-2a_3}+e^{-2a_2-2a_3} \end{bmatrix}.$$

Multiplying numerator and denominator by $\exp(2\sum_i a_i)$ gives

$$\frac{1}{D}b = \frac{2}{e^{2(a_1+a_2+a_3)}+e^{2a_1}+e^{2a_2}+e^{2a_3}}\begin{bmatrix} e^{2a_2}+e^{2a_3} \\ e^{2a_1}+e^{2a_3} \\ e^{2a_1}+e^{2a_2} \end{bmatrix}.$$

Writing $x_i = \exp(2a_i)$, this is then

$$\frac{1}{D}b = \frac{2}{x_1x_2x_3+\sum_i x_i}\begin{bmatrix} \sum_i x_i - x_1 \\ \sum_i x_i - x_2 \\ \sum_i x_i - x_3 \end{bmatrix},$$

and then $\bar{b}$ is given by:

$$\frac{1}{D}\bar{b} = \frac{2}{x_1x_2x_3+\sum_i x_i}[\sum_i x_i - \hat{a}\cdot x].$$

Hence using the expression for the derivative of our metric, given in Eq. (20), we arrive at:

$$\frac{dM}{dt} = \frac{2}{x_1x_2x_3+\sum_i x_i}\sum_i \hat{a}_i^2[\hat{a}\cdot x - x_i].$$

The coefficient is strictly positive, and so to argue non-increasing we can drop it. Hence we simply need to show:

$$(\hat{a}\cdot x)\left(\sum_i \hat{a}_i^2\right) - \sum_i \hat{a}_i^2 x_i \leq 0,$$

or equivalently:

$$\sum_i \hat{a}_i^2 \leq \frac{\sum_i \hat{a}_i^2 x_i}{\sum_i \hat{a}_i x_i}.$$

Define the quantities:

$$q_i = \frac{\hat{a}_i x_i}{\sum_i \hat{a}_i x_i},$$

noting that they form a probability distribution. Then our inequality can be viewed as comparing the expectation of $\hat{a}$ over two different probability distributions, one with probabilities $\hat{a}$ and the other with probabilities q.

It is then simple to note that, since for any $\|a\|_1 > 0$, x_i is an increasing function of $\hat{a}_i$, the second distribution places more weight on larger values of $\hat{a}$, and so must be at least as large as the other expectation. They are only equal when there are no larger values of $\hat{a}_i$, meaning $\hat{a} = \frac{1}{K-1}1_{K-1}$. Hence the inequality is satisfied, and the metric is monotonically decreasing.

It remains to show that the metric is not necessarily monotonic decreasing for $K \geq 8$. To show this for $K = 8$, it is sufficient to note that for $a = [1, 1, 1, 1.25, 0.01, 0.01, 0.01]$ the derivative of our metric takes value $\frac{dM}{dt} \approx 7.3 \times 10^{-4}$, and so the metric is increasing for these values. Call this specific counter-example v^*, and denote its corresponding b vector and its weighted average, by b^* and $\bar{b}^*$.

For $K > 8$, define our initialization as $a = [v^*, 0_{K-8}]$, then the derivative of our metric becomes

$$\frac{dM}{dt} = \frac{1}{D}\sum_{i=1}^{7}(\hat{v}_i^*)^2[b_i - \bar{b}], \qquad \text{where } \bar{b} = \sum_{j=1}^{7} v_i^* b_i.$$

Note that only $b_1, ..., b_7$ are necessary for the computation of the sign. We now use the lemma stated and proven at the end of this section, which gives that when a is only non-zero on at most the first $\frac{K}{2} - 1$ entries, then the first $\frac{K}{2} - 1$ entries of the b vector are just equal to double the values they would be in the $\frac{K}{2}$ dimensional case. Applying this iteratively, we get that if $K = 2^m > 8$, then

$$b_i = 2^{m-3}b_i^*, \qquad \bar{b} = 2^{m-3}\bar{b}^*, \qquad \text{for } i = 1, ..., 7,$$

and so

$$\frac{dM}{dt} = \frac{2^{m-3}}{D}\sum_{i=1}^{7}(\hat{v}_i^*)^2[b_i^* - \bar{b}^*], \qquad \text{where } \bar{b} = \sum_{j=1}^{7} v_i^* b_i,$$

but this clearly has the same sign as the counter-example in $K = 8$, and so gives a positive derivative.

This provides a counter-example in any K, but it has entries that are zero, which is not permitted. However if we consider the vector

$$a = [v^*, \epsilon 1_{K-8}],$$

then by continuity, as $\epsilon \to 0$ this tends to the positive value achieved by $a = [v^*, 0_{K-8}]$. Hence for any K there exists a non-zero value of ϵ such that $\frac{dM}{dt} > 0$ for $a = [v^*, \epsilon 1_{K-8}]$. Hence counter-examples exist in all $K \geq 8$.

It only remains to prove the following Lemma.

Lemma 1. *Let $K = 2^m$, and $a^* \in \mathbb{R}^{K-1}$, with $b(a^*) = b^* \in \mathbb{R}^{K-1}$ and $\bar{b}^* = \sum_i \hat{a}_i^* b_i^* \in \mathbb{R}^{K-1}$. Let $a' \in \mathbb{R}^{2K-1}$ be given by $a' = [a^*, 0_K]$, and consider the evolution under the dynamics detailed in Eq. (9), but for dimension $2K$, initialized at a'. Denote the corresponding vectors $b(a'), \bar{b}(a') \in \mathbb{R}^{2K-1}$ by b' and $\bar{b}'$ respectively. Then for $i = 1, ..., K-1$, we have*

$$b_i' = 2b_i^*, \qquad \bar{b}' = 2\bar{b}^*.$$

Proof of Lemma 1:

Let Ψ_{K-1} and Ψ_{2K-1} denote the two matrices derived from the Hadamard matrix, as defined in Sec. 3, with the subscript denoting their dimension. Using the construction in terms of Sylvester Hadamard matrices, as defined in Definition 2, these are related in the following way:

$$\Psi_{2K-1} = \begin{bmatrix} \Psi_{K-1} & 0_{K-1} & \Psi_{K-1} \\ 0_{K-1}^T & 2 & 21_{K-1}^T \\ \Psi_{K-1} & 21_{K-1} & 21_{K-1}1_{K-1}^T - \Psi_{K-1} \end{bmatrix}.$$

Hence

$$\Psi_{2K-1}a' = \begin{bmatrix} \Psi_{K-1} & 0_{K-1} & \Psi_{K-1} \\ 0_{K-1}^T & 2 & 21_{K-1}^T \\ \Psi_{K-1} & 21_{K-1} & 21_{K-1}1_{K-1}^T - \Psi_{K-1} \end{bmatrix} \begin{bmatrix} a^* \\ 0 \\ 0_{K-1} \end{bmatrix} = \begin{bmatrix} \Psi_{K-1}a^* \\ 0 \\ \Psi_{K-1}a^* \end{bmatrix},$$

and so

$$\exp(-\Psi_{2K-1}a) = \begin{bmatrix} \exp(-\Psi_{K-1}a^*) \\ 1 \\ \exp(-\Psi_{K-1}a^*) \end{bmatrix},$$

giving

$$b' = \Psi_{2K-1}\exp(-\Psi_{2K-1}a) = \begin{bmatrix} \Psi_{K-1} & 0_{K-1} & \Psi_{K-1} \\ 0_{K-1}^T & 2 & 21_{K-1}^T \\ \Psi_{K-1} & 21_{K-1} & 21_{K-1}1_{K-1}^T - \Psi_{K-1} \end{bmatrix} \begin{bmatrix} \exp(-\Psi_{K-1}a^*) \\ 1 \\ \exp(-\Psi_{K-1}a^*) \end{bmatrix}.$$

From this it is clear that the top $K-1$ entries of b' are given by the vector $2\Psi_{K-1}\exp(-\Psi_{K-1}a) = 2b^*$. Also note

$$\bar{b}' = \sum_j \hat{a}_j' b_j' = \sum_{j=1}^{K-1} \hat{a}_j^* b_j' = 2\bar{b}^*.$$

D.3 Proof of Theorem 5

Recall our dynamics are given by:

$$\frac{da_i}{dt} = \frac{1}{D}a_i b_i, \qquad \text{where: } b_i = \sum_j \Psi_{ij} e^{-(\Psi a)_j}, \quad D = 1 + \sum_j e^{-(\Psi a)_j}.$$

Suppose that $\|a\|_1 = \epsilon \ll 1$. Then expanding in ϵ gives $e^{-(\Psi a)_j} = 1 - (\Psi a)_j + O(\epsilon^2)$. Also recall that $\sum_j \Psi_{ij} = K$. We hence have:

$$D = 1 + \sum_j [1 - (\Psi a)_j] + O(\epsilon^2)$$

$$= K - K\sum_j a_j + O(\epsilon^2)$$

and hence:

$$D = K(1-\epsilon) + O(\epsilon^2).$$

Similarly we have:

$$b_i = \sum_j \Psi_{ij}(1 - (\Psi a)_j) + O(\epsilon^2)$$
$$= K - (\Psi^2 a)_i + O(\epsilon^2).$$

Note, using the properties of the original underlying Hadamard matrix, we have $\Psi^2 = K(I + 1_{K-1}1_{K-1}^T)$, giving

$$b_i = K - Ka_i - K\epsilon + O(\epsilon^2).$$

Hence, our differential equation for a_i is given by:

$$\frac{da_i}{dt} = a_i\left[K - Ka_i - K\epsilon + O(\epsilon^2)\right]\left[K(1-\epsilon) + O(\epsilon^2)\right]^{-1}$$
$$= a_i\left[1 - a_i - \epsilon + O(\epsilon^2)\right]\left[1 - \epsilon + O(\epsilon^2)\right]^{-1}$$
$$= a_i\left[1 - a_i - \epsilon + O(\epsilon^2)\right]\left[1 + \epsilon + O(\epsilon^2)\right]$$
$$= a_i(1 - a_i) + O(\epsilon^3).$$

Hence, we see that the linear dynamics do decouple. This equation is simple to solve, and gives a sigmoid solution, specifically:

$$a_i(t) = \frac{a_i(0)}{a_i(0) + (1 - a_i(0))e^{-t}}.$$

This solution is only valid as an approximation in the linearized regime.

D.4 Proof of Theorem 6

When optimizing the logits directly, as detailed in Eq. (13), the gradient flow equations are given by:

$$\frac{dZ}{dt} = Y - P.$$

We initialize our matrix Z as $Z = UD_ZV^T$, where $U = \frac{1}{\sqrt{K}}\Phi$, and $V = U \otimes Q$, where Q is a right singular matrix of 1_n^T, meaning it obeys $1_n^T = \sqrt{n}e_1^{(n)T}Q^T$, where $e_1^{(n)}$ is the first standard basis vector in $\mathbb{R}^n$. Also recall D_Z has its only non-zero singular values given by $a_1, \ldots, a_K$. The exact same arguments used in the proof of Theorem 2 then gives that, at initialization, Z can be written as

$$Z = (UD_{\tilde{a}}U^T) \otimes 1_n^T,$$

where $D_{\tilde{a}} = \text{diag}(\tilde{a}_1, \ldots, \tilde{a}_K)$, with $\tilde{a}_i = \frac{1}{\sqrt{n}}a_i$. Then Theorem 1 gives us that

$$P(Z) = (UD_{\tilde{\nu}}U^T) \otimes 1_n^T,$$

where

$$\tilde{\nu}_i = \frac{\sum_j \Phi_{ij}e^{\frac{1}{K}(\Phi\tilde{a})_j}}{\sum_j e^{\frac{1}{K}(\Phi\tilde{a})_j}}.$$

We can now turn to our gradient flow equation, and use the change of variable $Z = U\tilde{Z}V^T$, which evolves as:

$$\frac{d}{dt}\left(\tilde{Z}\right) = U^T(Y - P)V.$$

As in Theorem 2, define $B = U^T(Y - P)V$. Since we have $Y = \sqrt{n}UD_YV^T$ and $P = \sqrt{n}UD_PV^T$, where $D_Y = [I_K, 0_{K\times K(n-1)}] \in \mathbb{R}^{K\times Kn}$ and $D_P = [\text{diag}(\tilde{\nu}), 0_{K\times K(n-1)}] \in \mathbb{R}^{K\times Kn}$, we have:

$$B = \sqrt{n}[\text{diag}(1_K - \tilde{\nu}), 0_{K\times K(n-1)}].$$

Hence the evolution in $\tilde{Z}$ becomes:

$$\frac{d}{dt}\left(\tilde{Z}_{ij}\right) = \sqrt{n}(1-\tilde{\nu}_i)\delta_{ij}.$$

We see that the only entries of $\tilde{Z}$ that evolve are the entries on its diagonal. Consequently if at initialization it takes the form of a singular value matrix, this property continues to hold for all future times. At initialization we have $\tilde{Z} = [\text{diag}(a), 0_{K\times K(n-1)}]$, and so these singular values evolve as:

$$\frac{d}{dt}(a_i) = \sqrt{n}(1-\tilde{\nu}_i).$$

Translating to the original variable Z, this tells us that the singular vectors remain fixed, and all time evolution occurs in the singular values.

Just as in the proof of Theorem 2, using $\tilde{a} = \frac{1}{\sqrt{n}}a$ we can write:

$$(1-\tilde{\nu}_i) = \frac{1}{\sum_j e^{-\frac{1}{K\sqrt{n}}([1_K 1_K^T - \Phi]a)_j}}\left[\sum_j [1_K 1_K^T - \Phi]_{ij} e^{-\frac{1}{K\sqrt{n}}([1_K 1_K^T - \Phi]a)_j}\right].$$

Note that this quantity is zero for $i = 1$, and so a_1 stays at its initial value. It also does not appear in the evolution equations of the other variables, and so we drop it and only consider the non-trivial modes, viewing $a \in \mathbb{R}^{K-1}$ to be the vector containing only the non-trivial singular values.

The evolution equations are then:

$$\frac{d}{dt}(a_i) = \frac{\sqrt{n}}{D}b_i, \qquad \text{where } b_i = \sum_{j=1}^{K-1}\Psi_{ij}e^{-\frac{1}{K\sqrt{n}}(\Psi a)_j}, \quad D = 1 + \sum_{j=1}^{K-1} e^{-\frac{1}{K\sqrt{n}}(\Psi a)_j}.$$

Absorbing constants into a and t via: $a' = \frac{1}{K\sqrt{n}}a$ and $t' = \frac{1}{K}t$, then dropping primes since these re-scalings do not change the phenomenology of the dynamics, reduces us to

$$\frac{d}{dt}(a_i) = \frac{1}{D}b_i, \qquad \text{where } b_i = \sum_{j=1}^{K-1}\Psi_{ij}e^{-(\Psi a)_j}, \quad D = 1 + \sum_{j=1}^{K-1} e^{-(\Psi a)_j}.$$

If we define the variables $\hat{a}_i = a_i/\|a\|_1$, then these have evolution equation given by:

$$\frac{d}{dt}(\hat{a}_i) = \frac{1}{D\|a\|_1}\left[b_i - \hat{a}_i\sum_j b_j\right].$$

Again defining the metric M as

$$M = \frac{1}{2}\sum_i\left[\hat{a}_i - \frac{1}{K-1}\right]^2.$$

This has derivative:

$$\frac{dM}{dt} = \frac{1}{D\|a\|_1}\left[\left(\sum_i \hat{a}_i b_i\right) - \left(\sum_i \hat{a}_i^2\right)\left(\sum_j b_j\right)\right].$$

Now note that, as in the proof of Theorem 3, if we define $x_i = (Ta)_i$, we have:

$$\sum_i \hat{a}_i b_i = \frac{1}{\|a\|_1}\sum_j x_j e^{-x_j},$$

and as was shown in the proof, this satisfies:

$$\frac{1}{\|a\|_1}\sum_j x_j e^{-x_j} \leq \frac{1}{(K-1)\|a\|_1}\left(\sum_i x_i\right)\left(\sum_j e^{-x_j}\right),$$

with equality only when $x \propto 1_{K-1}$. Hence we have that the derivative of M satisfies:

$$\frac{dM}{dt} \leq \frac{1}{D\|a\|_1}\left[\frac{1}{(K-1)\|a\|_1}\left(\sum_i x_i\right)\left(\sum_j e^{-x_j}\right) - \left(\sum_i \hat{a}_i^2\right)\left(\sum_j b_j\right)\right].$$

Then use that $\sum_i b_i = K\sum_i e^{-x_i}$, and $\sum_i x_i = K\|a\|_1$, reducing this to

$$\frac{dM}{dt} \leq \frac{K\sum_j e^{-x_j}}{D\|a\|_1}\left[\frac{1}{(K-1)} - \left(\sum_i \hat{a}_i^2\right)\right].$$

Then note by Cauchy-Schwarz we have $\sum_i \hat{a}_i^2 \geq \frac{1}{K-1}$, with equality only when $\hat{a} \propto 1_{K-1}$, consequently we arrive at

$$\frac{dM}{dt} \leq 0,$$

with equality only when $x \propto 1_{K-1}$ and $\hat{a} \propto 1_{K-1}$. These both only occur when $\hat{a} = \frac{1}{K-1}1_{K-1}$, as argued in the proof of Theorem 3.

Hence the metric M monotonically decreases to zero in this direct logit optimization case.

D.5 Proof of Theorem 7

Recall our metric M, given by:

$$M = \frac{1}{2}\sum_{i=1}^{K-1}\left(\frac{a_i}{\|a\|_1} - \frac{1}{K-1}\right)^2.$$

Throughout the proof we make use of the following properties of $\Psi = 1_{K-1}1_{K-1}^\top - X$:

$$\begin{aligned}
\Psi\, 1_{K-1} &= K\, 1_{K-1}, \\
\Psi^2 &= K\, I_{K-1} + K\, 1_{K-1}1_{K-1}^\top, \\
|\Psi_{ij}| &\leq 2, \qquad \forall\, i,j, \\
\|\psi_u\|_2 &= \sqrt{2K}, \qquad \forall \text{ columns } \psi_u \text{ of } \Psi.
\end{aligned}$$

Additionally, unless otherwise stated, all indices run in $1,\ldots,k-1$ and are dropped for simplifying notation.

We choose to write our singular modes in the following form:

$$a_i = \mu + \mu\epsilon_i, \qquad \sum_i \epsilon_i = 0, \qquad \mu = \frac{1}{K-1}\sum_i a_i\,.$$

In this form M is given by

$$M = \frac{1}{2(K-1)^2}\sum_i \epsilon_i^2\,.$$

Recall its derivative as given in Eq. (20):

$$\frac{dM}{dt} = \frac{1}{D}\sum_i \hat{a}_i^2[b_i - \bar{b}], \qquad \text{where } \bar{b} = \sum_j \hat{a}_j b_j\,.$$

This we can write alternatively as

$$\begin{aligned}
\frac{dM}{dt} &= \frac{1}{D}\left[\left(\sum_i \hat{a}_i^2 b_i\right) - \sum_{ij}\hat{a}_i^2\hat{a}_j b_j\right] = \frac{1}{D\|a\|_1^3}\left[\left(\sum_i a_i^2 b_i\right)\left(\sum_j a_j\right) - \sum_{ij} a_i^2 a_j b_j\right] \\
&= \frac{1}{D\|a\|_1^3}\sum_{ij} a_i^2 a_j (b_i - b_j)\,.
\end{aligned}$$

Now writing this in terms of ϵ and μ gives

$$\frac{dM}{dt} = \frac{1}{D\|a\|_1^3}\sum_{iju}(\mu+\mu\epsilon_i)^2(\mu+\mu\epsilon_j)[\Psi_{iu} - \Psi_{ju}]e^{-\mu(\Psi[1_{K-1}+\epsilon])_u}$$

$$= \frac{e^{-K\mu}}{D(K-1)^3} \sum_{iju} (1+\epsilon_i)^2 (1+\epsilon_j)[\Psi_{iu} - \Psi_{ju}] e^{-\mu(\Psi\epsilon)_u} \,,$$

where we used that $\mu = \frac{1}{K-1}\|a\|_1$, and $\Psi 1_{K-1} = K 1_{K-1}$.

Expand the exponential as:

$$\exp(-\mu(\Psi\epsilon)_u) = 1 - \mu(\Psi\epsilon)_u + \sum_{n=2}^{\infty} \frac{1}{n!} \left[-\mu(\Psi\epsilon)_u\right]^n \,.$$

Using $\sum_u [\Psi_{iu} - \Psi_{ju}] = 0$, the first term in $\frac{dM}{dt}$ gives

$$\sum_{iju} (1+\epsilon_i)^2 (1+\epsilon_j)[\Psi_{iu} - \Psi_{ju}] = 0 \,.$$

Further using that $\Psi^2 = K I_{K-1} + K 1_{K-1} 1_{K-1}^T$, the second term becomes

$$\begin{aligned}
-\mu \sum_{iju} (1+\epsilon_i)^2 (1+\epsilon_j)[\Psi_{iu} - \Psi_{ju}](\Psi\epsilon)_u &= -\mu \sum_{ij} (1+\epsilon_i)^2 (1+\epsilon_j)[(\Psi^2\epsilon)_i - (\Psi^2\epsilon)_j] \\
&= -K\mu \sum_{ij} (1+\epsilon_i)^2 (1+\epsilon_j)[\epsilon_i - \epsilon_j] \\
&= -K\mu \big[\sum_i (1+\epsilon_i)^2 \epsilon_i \sum_j (1+\epsilon_j) - \sum_i (1+\epsilon_i)^2 \sum_j (1+\epsilon_j)\epsilon_j \big] \\
&= -K\mu \left[(K-1) \sum_i \epsilon_i^2 + (K-1) \sum_i \epsilon_i^3 - \left(\sum_i \epsilon_i^2 \right)^2 \right] ,
\end{aligned}$$

where in the last line we used $\sum_i \epsilon_i = 0$.

Again because $\sum_i \epsilon_i = 0$ and $\sum_u \Psi_{iu} = K$, we can also simplify the remainder term by using

$$\sum_{ij} (1+\epsilon_i)^2 (1+\epsilon_j)(\Psi_{iu} - \Psi_{ju}) = \sum_i [(K-1)\Psi_{iu}\epsilon_i + (K-1)\Psi_{iu}\epsilon_i^2 - K\epsilon_i^2 - \epsilon_i^2 (\Psi\epsilon)_u] \,.$$

Putting all the above together, our remaining expression for the full derivative is given by

$$\begin{aligned}
\frac{dM}{dt} = &-\frac{K\mu e^{-K\mu}}{D(K-1)^2} \sum_i \epsilon_i^2 - \frac{K\mu e^{-K\mu}}{D(K-1)^2} \sum_i \epsilon_i^3 + \frac{K\mu e^{-K\mu}}{D(K-1)^3} \left(\sum_i \epsilon_i^2 \right)^2 \\
&+ \frac{e^{-K\mu}}{D(K-1)^3} \sum_{iu} \left[(K-1)\Psi_{iu}\epsilon_i + (K-1)\Psi_{iu}\epsilon_i^2 \underbrace{-K\epsilon_i^2 - \epsilon_i^2 (\Psi\epsilon)_u}_{(*)} \right] \left[\sum_{n=2}^{\infty} \frac{1}{n!} [-\mu(\Psi\epsilon)_u]^n \right] .
\end{aligned}$$

We want to show that when M is small enough, this quantity is negative. First let us assume that $|(\Psi\epsilon)_u| \leq K$, in which case we can drop the two terms highlighted as $(*)$, since they contribute a non-positive value, where we used that $e^x - 1 - x \geq 0$. Hence, also dropping a positive multiplicative constant, it remains to show

$$0 \geq -K\mu \sum_i \epsilon_i^2 - K\mu \sum_i \epsilon_i^3 + \frac{K\mu}{K-1} \left(\sum_i \epsilon_i^2 \right)^2 + \sum_{iu} (\Psi_{iu}\epsilon_i + \Psi_{iu}\epsilon_i^2) \left[\sum_{n=2}^{\infty} \frac{1}{n!} [-\mu(\Psi\epsilon)_u]^n \right] .$$

Now we use the following properties to bound the remaining terms:

$$\sum_i \epsilon_i^2 = 2(K-1)^2 M,$$

$$|\sum_i \epsilon_i^3| \leq \left(\sum_i \epsilon_i^2 \right)^{\frac{3}{2}} = 2^{\frac{3}{2}} (K-1)^3 M^{\frac{3}{2}},$$

$$\sum_i \Psi_{iu}\epsilon_i^2 \le 2\sum_i \epsilon_i^2 = 4(K-1)^2 M,$$

$$\sum_i \Psi_{iu}\epsilon_i = (\Psi\epsilon)_u \le |(\Psi\epsilon)_u| \le \sqrt{(\Psi\epsilon)\cdot(\Psi\epsilon)} = \sqrt{2K}(K-1)\sqrt{M}. \tag{21}$$

Note the second property follows from monotonicity of L_p norms, the third uses $\Psi_{iu} \in \{0, 2\}$ for all i, u, and, the last one $\|\psi_u\| = \sqrt{2K}$ for all columns/rows ψ_u of Ψ. This produces the following sufficient condition for a negative derivative:

$$\begin{aligned} 0 \ge & -2K\mu(K-1)^2 M + 2^{\frac{3}{2}} K\mu(K-1)^3 M^{\frac{3}{2}} + 4K\mu(K-1)^3 M^2 \\ & + \sum_u \left[\sqrt{2K}(K-1)\sqrt{M} + 4(K-1)^2 M\right] \underbrace{\left[\sum_{n=2}^{\infty} \frac{1}{n!}[-\mu(\Psi\epsilon)_i]^n\right]}_{R_u} . \end{aligned}$$

We can bound the quantity R_u using $e^x - 1 - x \le \frac{1}{2}x^2 e^{|x|}$, giving

$$R_u \le \frac{1}{2}\mu^2 (\Psi\epsilon)_u^2 \exp(|\mu(\Psi\epsilon)_u|) \le K\mu^2(K-1)^2 M \exp\left(\mu\sqrt{2K}(K-1)\sqrt{M}\right) .$$

where in the second step we used the bound for $(\Psi\epsilon)_u$ from Eq. (21). Hence, it is sufficient to show

$$\begin{aligned} 0 \ge & -2K\mu(K-1)^2 M + 2^{\frac{3}{2}} K\mu(K-1)^3 M^{\frac{3}{2}} + 4K\mu(K-1)^3 M^2 \\ & + \left[\sqrt{2K}(K-1)^2\sqrt{M} + 4(K-1)^3 M\right] K\mu^2(K-1)^2 M \exp\left(\mu\sqrt{2K}(K-1)\sqrt{M}\right) , \end{aligned}$$

where we also picked up an extra factor of $(K-1)$ due to the summation over u.

Now, dropping a factor of $2K(K-1)^2\mu M$, this reduces to

$$0 \ge -1 + 2^{\frac{1}{2}}(K-1)M^{\frac{1}{2}} + 2(K-1)M + \left[\sqrt{\frac{K}{2}}(K-1)^2\sqrt{M} + 2(K-1)^3 M\right] \mu \exp\left(\mu\sqrt{2K}(K-1)\sqrt{M}\right) .$$

Letting $\tilde{M} = 2KM$ and $\tilde{\mu} = (K-1)\mu$, this further reduces to

$$0 \ge -1 + \frac{K-1}{\sqrt{K}}\tilde{M}^{\frac{1}{2}} + \frac{K-1}{K}\tilde{M} + \left[\frac{1}{2}(K-1)\tilde{M}^{\frac{1}{2}} + \frac{(K-1)^2}{K}\tilde{M}\right] \tilde{\mu} e^{\tilde{\mu}\tilde{M}^{\frac{1}{2}}} ,$$

and finally using that $K - 1 \le K$ it is sufficient to show

$$1 \ge \sqrt{K}\tilde{M}^{\frac{1}{2}} + \tilde{M} + \left[\frac{1}{2}\tilde{M}^{\frac{1}{2}} + \tilde{M}\right] K\tilde{\mu} e^{\tilde{\mu}\tilde{M}^{\frac{1}{2}}} .$$

Now assume $\tilde{M}^{\frac{1}{2}} \le \min\left(\frac{1}{2\sqrt{K}}, \frac{1}{3K\tilde{\mu}}\right)$, then we have:

$$\sqrt{K}\tilde{M}^{\frac{1}{2}} + \tilde{M} + \left[\frac{1}{2}\tilde{M}^{\frac{1}{2}} + \tilde{M}\right] K\tilde{\mu} e^{\tilde{\mu}\tilde{M}^{\frac{1}{2}}} \le \frac{1}{2} + \frac{1}{4K} + \left[\frac{1}{2} + \frac{1}{2\sqrt{K}}\right]\frac{1}{3}e^{\frac{1}{3K}} \le \frac{1}{2} + \frac{1}{8} + \left[\frac{1}{2} + \frac{1}{2\sqrt{2}}\right]\frac{1}{3}e^{\frac{1}{6}} \approx 0.961 ,$$

where we used that $K \ge 2$. This quantity is less than one, and hence $\frac{dM}{dt}$ is negative subject to this constraint on $\tilde{M}$. Stated in terms of M and $\|a\|_1$, this constraint is

$$M \le \min\left(\frac{1}{8K^2}, \frac{1}{18K^3\|a\|_1^2}\right) .$$

Recall that we also used $|(\Psi\epsilon)_u| \le K$, previously, but this is a weaker bound than the above (by Eq. (21)) and so holds subject to our condition. Hence we have shown that

$$\frac{dM}{dt} \le 0, \quad \text{when } M \le \min\left(\frac{1}{8K^2}, \frac{1}{18K^3\|a\|_1^2}\right) .$$

D.6 Proof of Theorem 8

We now consider randomly initialized matrices for the UFM model of Eq. (1). Assume that W, H are initialized as $W(0) = \epsilon d^{-\frac{1}{2}} A_1$, $H(0) = \epsilon d^{-\frac{1}{2}} A_0$, so that the logit matrix is given by $Z(0) = \epsilon^2 d^{-1} A_1 A_0$.

We assume that for a large value $t^* \in \mathbb{R}$ the value of ϵ is small enough that for the time window $t \in [0, t^*]$, $W(t) = O(\epsilon)$, $H(t) = O(\epsilon)$, and so $Z(t) = O(\epsilon^2)$, meaning we remain in this lower order approximate regime on that time window. Note by $O(\cdot)$, we mean $\exists C(t^*) > 0$ such that $\|Z(t, \epsilon)\|_F \leq C(t^*)\epsilon^2$ for $t \in [0, t^*]$. We will capture the relationship between t^* and ϵ later in the proof.

Hence for this time window, we have that the softmax is given by:

$$P(t) = \frac{1}{K} 1_K 1_{Kn}^T + P'(t, \epsilon),$$

where $P'(t, \epsilon) = O(\epsilon^2)$. From the gradient flow equations for this model, described in Eq. (2), we have for $t \in [0, t^*]$

$$\frac{dW}{dt} = (S \otimes 1_n^T) H^T - P' H^T, \qquad \frac{dH^T}{dt} = (S \otimes 1_n^T)^T W - P' W.$$

We can reduce this to a second order ODE

$$\frac{d^2 W}{dt^2} = nSW + \underbrace{\frac{d}{dt}(-P' H^T) - (S \otimes 1_n^T) P' W}_{O(\epsilon^3)},$$

where we have the boundary conditions $W(0) = \epsilon d^{-\frac{1}{2}} A_1$, $\frac{dW}{dt}(0) = \epsilon d^{-\frac{1}{2}} (S \otimes 1_n^T) A_0^T + O(\epsilon^3)$. We also note that the terms that feature P' are $O(\epsilon^3)$ on the considered time interval, so should be viewed as perturbative.

This has solution given by

$$W(t) = \epsilon d^{-\frac{1}{2}} \left[\frac{1}{K} 1_K 1_K^T A_1 + \cosh(t\sqrt{n}) S A_1 + \frac{1}{\sqrt{n}} \sinh(t\sqrt{n}) (S \otimes 1_n^T) A_0^T \right] + O(\epsilon^3 t^2),$$

and corresponding solution for H given by

$$H^T(t) = \epsilon d^{-\frac{1}{2}} \left[(I_{Kn} - \frac{1}{n} S \otimes 1_n 1_n^T) A_0^T + \frac{1}{\sqrt{n}} \sinh(t\sqrt{n}) (S \otimes 1_n) A_1 + \frac{1}{n} \cosh(t\sqrt{n}) (S \otimes 1_n 1_n^T) A_0^T \right] + O(\epsilon^3 t^2).$$

We now want to compute $Z(t^*) = W(t^*) H(t^*)$. In this expansion we recover terms of a variety of different orders in t and ϵ. We now define

$$t^* = \frac{1}{2\sqrt{n}} \log\left(\frac{1}{\epsilon}\right)$$

This diverges in the $\epsilon \to 0$ limit, but slowly enough that the logit matrix norm goes to zero, meaning our perturbative terms remain small. By considering the terms in $W(t)$ and $H(t)$, it becomes clear that the leading order term in $Z(t^*) = W(t^*) H(t^*)$ comes from the product of the terms in W, H that are exponential in t^*, giving

$$Z(t^*) = \frac{1}{4d} \epsilon^2 e^{2t^*\sqrt{n}} \left[S A_1 + \frac{1}{\sqrt{n}} (S \otimes 1_n^T) A_0^T \right] \left[\frac{1}{\sqrt{n}} (S \otimes 1_n) A_1 + \frac{1}{n} (S \otimes 1_n 1_n^T) A_0^T \right]^T + \dots,$$

where by $\dots$ we mean terms that vanish faster than leading order. Continuing

$$Z(t^*) = \frac{1}{4d} \epsilon^2 e^{2t^*\sqrt{n}} S \left[A_1 + \frac{1}{\sqrt{n}} (I_K \otimes 1_n^T) A_0^T \right] \left[A_1 + \frac{1}{\sqrt{n}} (I_K \otimes 1_n^T) A_0^T \right]^T (S \otimes \frac{1}{\sqrt{n}} 1_n^T) + \dots.$$

Now define the matrix $X = \frac{1}{\sqrt{2}} [A_1 + (I_K \otimes \frac{1}{\sqrt{n}} 1_n^T) A_0^T]$, this gives

$$Z(t^*) = \frac{1}{2d} \epsilon^2 e^{2t^*\sqrt{n}} S X X^T (S \otimes \frac{1}{\sqrt{n}} 1_n^T) + \dots.$$

Now suppose that the entries of $A_1 \in \mathbb{R}^{K \times d}$ and $A_0^T \in \mathbb{R}^{Kn \times d}$ are initialized with i.i.d. samples from a standard Gaussian. Then the entries of the matrix $\frac{1}{\sqrt{n}} (I_K \otimes 1_n^T) A_0^T \in \mathbb{R}^{K \times d}$ are also i.i.d. samples of a standard Gaussian, and hence the matrix $X \in \mathbb{R}^{K \times d}$ also has entries i.i.d. sampled from a standard Gaussian.

Considering the entries of the matrix $\frac{1}{d}XX^T$ in the large d limit, we find:

$$i = j: \left[\frac{1}{d}(XX^T)\right]_{ii} = \frac{1}{d}\sum_u X_{iu}X_{iu} \to \mathbb{E}[N_1^2] = 1,$$

$$i \neq j: \left[\frac{1}{d}(XX^T)\right]_{ij} = \frac{1}{d}\sum_u X_{iu}X_{ju} \to \mathbb{E}[N_1 N_2] = \mathbb{E}[N_1]\mathbb{E}[N_2] = 0.$$

Where $N_1, N_2 \sim \mathcal{N}(0,1)$.

Hence in the large d limit $\frac{1}{d}XX^T \to I_K$, and we find that

$$\frac{1}{d}SXX^T(S \otimes \frac{1}{\sqrt{n}}1_n^T) \approx \frac{1}{\sqrt{n}}(S \otimes 1_n^T),$$

and so $Z(t^*)$ is well approximated as

$$Z(t^*) = \frac{1}{2\sqrt{n}}\epsilon^2 e^{2t^*\sqrt{n}} S \otimes 1_n^T + \dots,$$

where the higher order terms decay faster in the described limits, and so this gives the asymptotic leading term.

E Additional Experiments

E.1 Metrics

In the numerical simulations discussed in Sec. 6, we track the following metrics.

(a) *Cross-Entropy Loss:* Plots the CE loss in Eq. (1).

(b) *Alignment Metrics:* Tracks the Frobenius norm distance between the nuclear-norm-normalized model components ($WW^T, H^TH, Z = WH$) and the ground-truth Simplex ETF S. For example for the logit matrix $Z = WH$ it plots $\left\|\frac{Z}{\|Z\|_*} - \frac{S}{\|S\|_*}\right\|_F$. For Hadamard Init., where Z and S share singular spaces, this is identical to the (square root) of the metric M in Eq. (12).

(c) *KL Divergence:* Plots the Lyapunov function from Theorem 3, measuring the KL divergence of the normalized logit singular values $\hat{a}$ from the uniform distribution $\frac{1}{K-1}1_{K-1}$.

(d) *Singular Values:* Plots the evolution of the raw singular values $\sigma_i(Z)$ of the logit matrix $Z = WH$. For the `SVD` and `ODE` runs, these are equivalent to the paper's logit singular values a_i.

(e) *Normalized Singular Values:* Plots the L_1-normalized singular values $\hat{\sigma}_i(Z) = \sigma_i(Z)/\sum_j \sigma_j(Z)$. For the `SVD` and `ODE` runs, these are equivalent to the paper's $\hat{a}_i = a_i/\|a\|_1$.

(f) *Residual Energy:* Measures the fraction of the logit matrix Z that lies outside the signal subspace of the Hadamard Init.. Specifically, for $U = V = \Phi(1:K, 2:K)$, it measures $\|Z - U(U^TZV)V^T\|_F / \|Z\|_F$

For efficiency, we track and plot these metrics in 100 logarithmically spaced time instants.

E.2 Larger K

We present the simulation results for $K = 8$ classes in Fig. 4. Many of the core findings from the $K = 4$ case remain:

- *Fixed Vectors:* The Residual Energy (plot f) for the Hadamard Init. (red) is zero for all time, confirming the singular vectors are fixed.
- *ODE Validation:* The ODE (green dashed) perfectly tracks the Hadamard Init. dynamics (red) in plots (d) and (e).
- *Two-Phase Dynamics:* The Random Init. (blue) exhibits a distinct "burn-in" phase (up to $t \approx 5 \times 10^3$) where its singular vectors align with the ETF subspace (plot f) while the loss remains flat (plot a). During this alignment, the normalized singular values of the Hadamard Init. remain flat. It is only after this burn-in phase that the normalized singular values for both Random and Hadamard Initializations begin converging toward the NC equilibrium.

However, note two differences in the $K = 8$ case. First, the alignment of the normalized singular values between the two initializations, even after the burn-in phase, is weaker than in the $K = 4$ case. Second, it is evident from plots (c) and (e) that the convergence to NC is significantly slower for $K = 8$, even for the Hadamard Initialization.

The observations above appear consistent for $K = 16$ as shown in Figure 5.

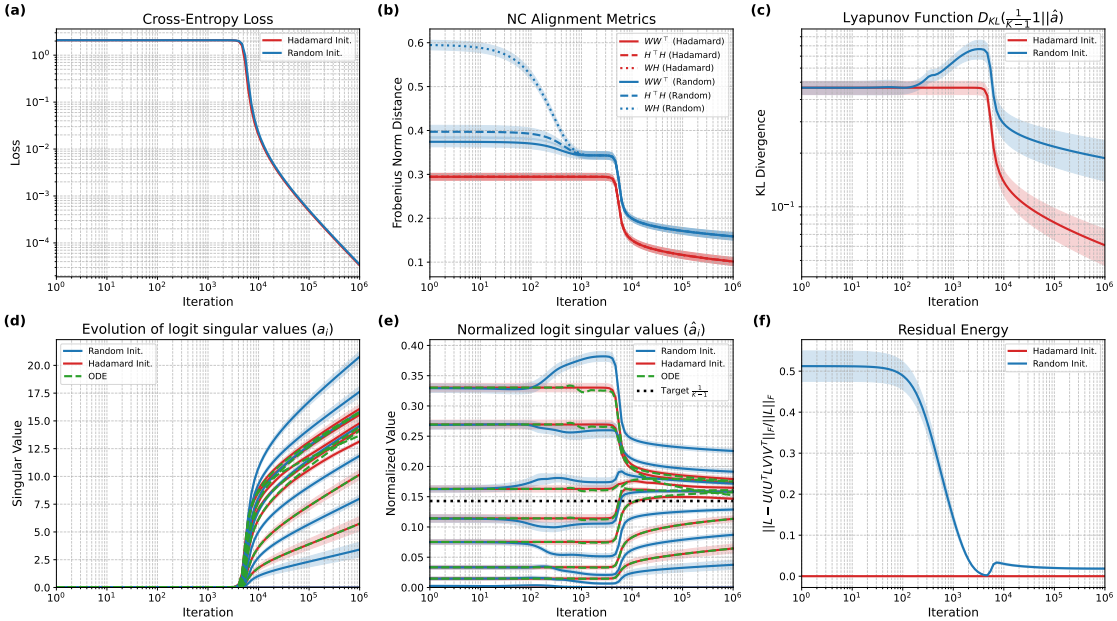

Figure 4: Comparison of Hadamard Init., Random Init., and theoretical ODE for $K = 8$ classes.

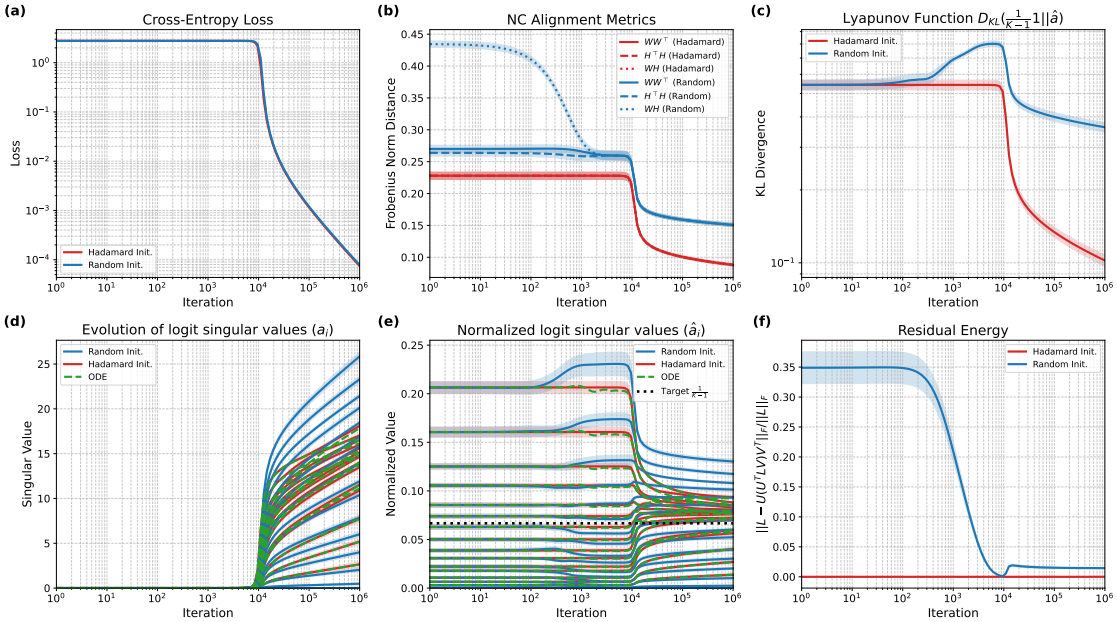

Figure 5: Comparison of Hadamard Init., Random Init., and theoretical ODE for $K = 16$ classes.

E.3 The effect of initialization scale

Thus far we used small initializations: recall from Sec. 6 that W and H are initialized with i.i.d. Gaussian entries of variance $1/K$ and then scaled by $\epsilon = e^{-\delta}$ with $\delta = 6$. For MSE, following Saxe et al. [9], this vanishing initialization, yields dynamics similar to spectral initialization; our previous CE experiments were the analogous setting. For completeness, we also report results under a standard scale, i.e., $\epsilon = 1$ ($\delta = 0$). Figures 6 and 7 show $K = 4$ and $K = 8$, respectively. The two-stage dynamics seen with infinitesimal initialization persist, but with quantitative changes: e.g., the initial "burn-in" shortens to $t \sim 5 \times 10^2$. Notably, while the alignment of logit singular values to their NC equilibrium (measured by NC metrics and KL divergence) remains comparable to the infinitesimal case, the residual

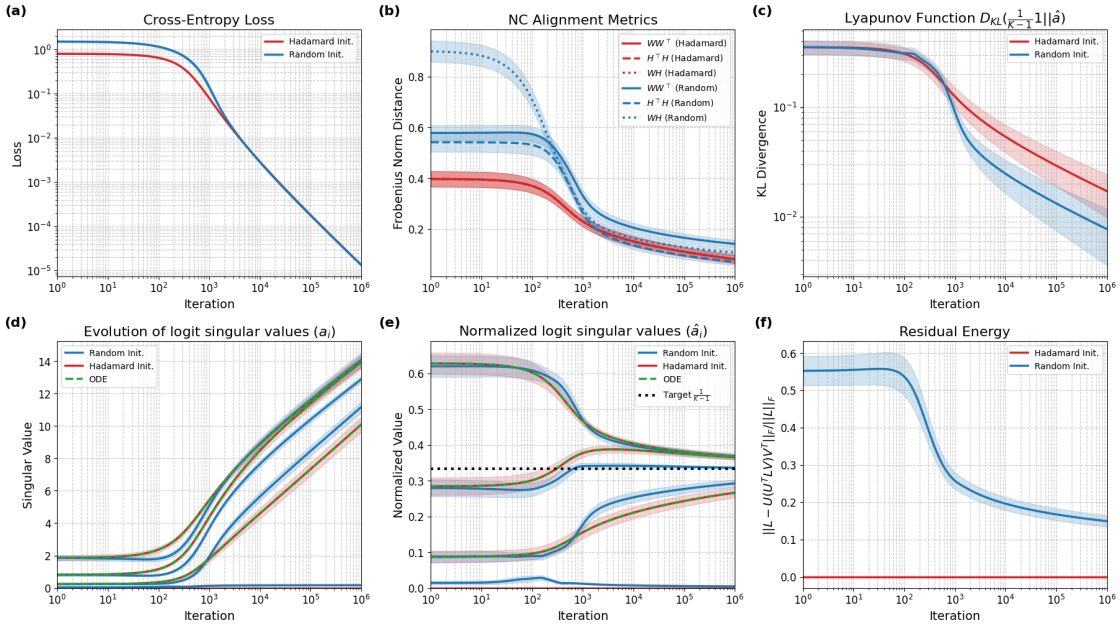

Figure 6: Same as Figure 2 for $K = 4$ but now with standard initialization scale (rather than infinitesimal).

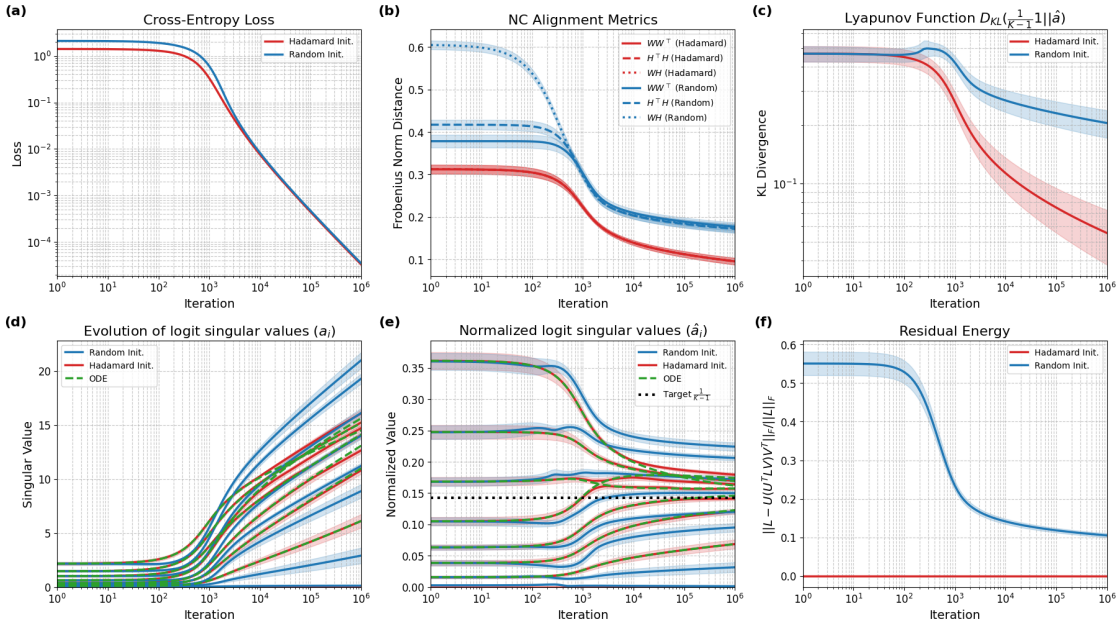

Figure 7: Same as Figure 4 for $K = 8$ but now with standard initialization scale (rather than infinitesimal).

energy is markedly higher: subspace alignment from Random Initialization to the simplex ETF subspace proceeds much more slowly.

E.4 The effect of overparameterization

In the experiments thus far, we have set the UFM's hidden dimension to $d = K$. Since the simplex ETF subspace has dimension $K - 1$, this is essentially the minimal choice ensuring that the simplex ETF structure is attainable. For completeness, we examine the overparameterized regime with $d > K$.

For Hadamard Initialization, our theory predicts that d should not affect the dynamics. Figures 8 and 9, showing overparameterization factors $d/K = 5$ and 10 respectively, confirm this prediction.

However, overparameterization benefits Random Initialization: while the loss plateau persists, the burn-in phase (where parameter magnitudes remain small) now exhibits rapid convergence of normalized singular values toward equilibrium. Interestingly, this occurs despite the residual energy decaying at roughly the same rate as in the $d = K$ case, though it remains closer to zero during the second phase, avoiding the slight increase observed without overparameterization.

This behavior aligns with the predictions of Theorem 8: Under infinitesimal random initialization, zero-th order UFM dynamics already drive the solution to simplex ETF geometry.

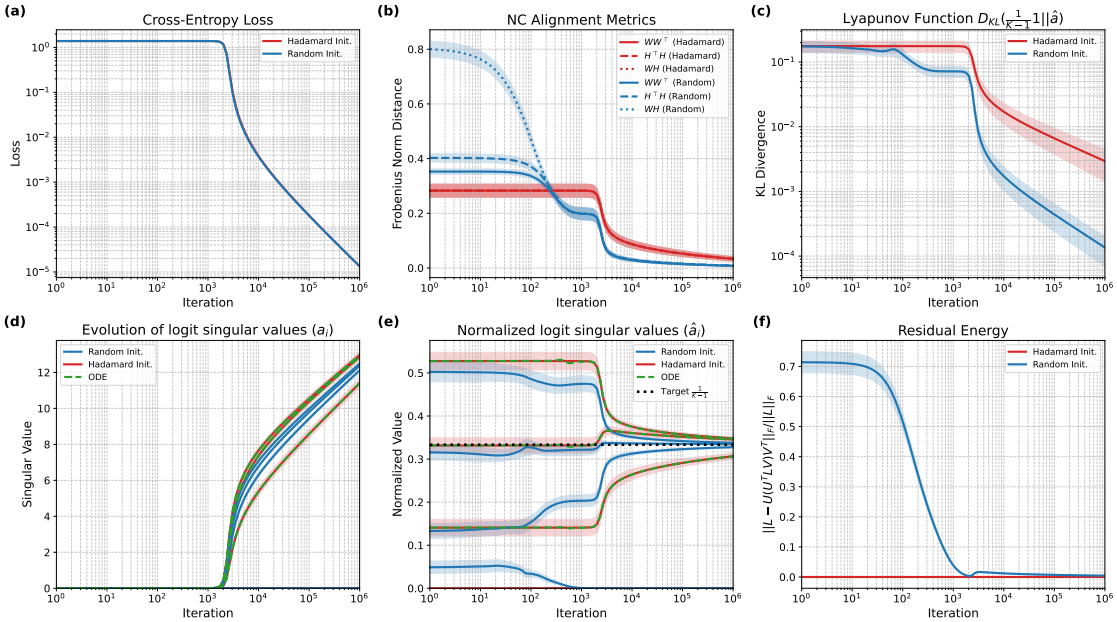

Figure 8: Same as Figure 2 for $K = 4$ but now with hidden dimension $d = 5 \cdot K$.

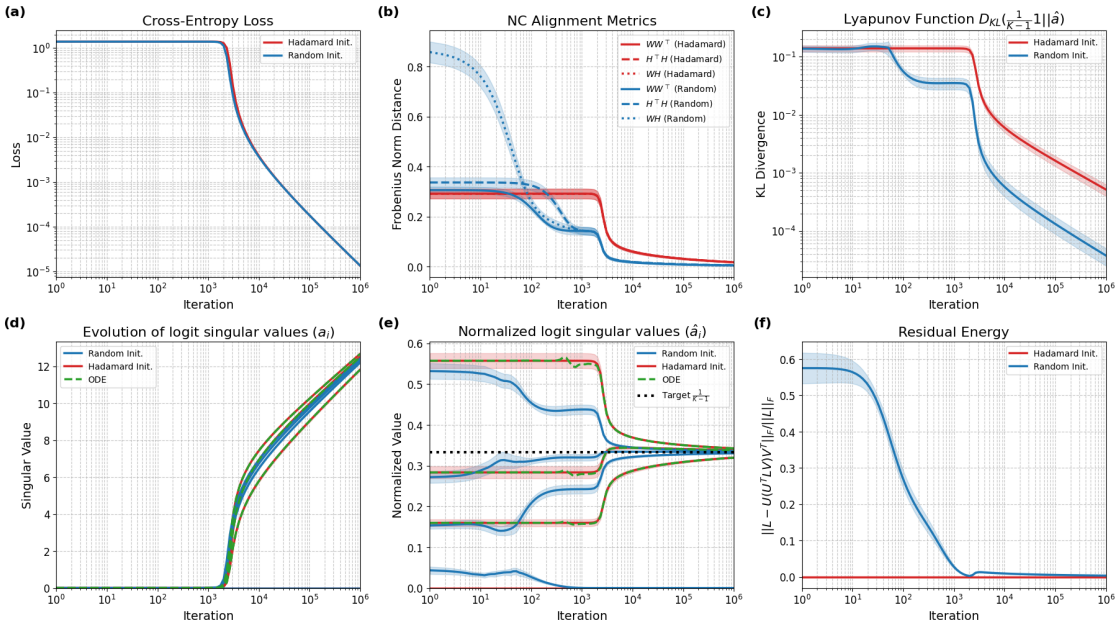

Figure 9: Same as Figure 2 for $K = 8$ but now with hidden dimension $d = 10 \cdot K$.

E.5 Normalized Gradient Flow

We have seen that gradient flow converges to the NC solution (with all normalized singular values $\hat{a}_i$ equal), but convergence can be very slow. The counterexamples constructed for metric M in Sec. 4.3—specifically Figure 1—illustrate this. For the initialization shown, metric M increases, but Theorem 3 guarantees it must eventually decrease to 0 as the flow converges to equilibrium. However, this occurs extremely slowly. To numerically verify this behavior, we implement an accelerated version of the reduced flow in Eq. (2), the *normalized* gradient flow:

$$\frac{da}{dt} = \frac{F(a)}{\|F(a)\|_2} \qquad \text{where } F(a) = a \odot \nabla_a E(a) = \frac{a \odot b}{D}. \tag{22}$$

For convex models, L_2-normalized flow accelerates optimization without modifying convergence to the same max-margin solution as standard flow [85]. We leave proving this property in our non-convex setting to future work.

Unfortunately, direct implementation of the updates in Eq. (22) encounters numerical issues: The gradients of the energy functional $E(a)$ involve terms that can overflow when margins become large during training. This is particularly worrisome for normalized flow, which accelerates margin growth. To implement gradient flow stably, we rewrite gradient terms by factoring out dominant exponential terms and subtracting the minimum margin, an idea inspired by [3, App. I].

First note that given the gradient components $F_i(a) = a_i b_i / D$, the denominator D cancels entirely:

$$\frac{da_i}{dt} = \frac{F_i(a)}{\|F(a)\|_2} = \frac{a_i b_i / D}{\sqrt{\sum_k (a_k b_k / D)^2}} = \frac{a_i b_i}{\|a \odot b\|_2}.$$

We now apply a stability trick to the vector b itself. Let $m_j = (\Psi a)_j$ be the margins and $m_{\min} = \min_j m_j$. We define a numerically stable rescaled vector $\hat{b}$ as:

$$\hat{b}_i := b_i \cdot e^{m_{\min}} = \left(\sum_{j=1}^{K-1} \Psi_{ij} e^{-m_j} \right) e^{m_{\min}} = \sum_{j=1}^{K-1} \Psi_{ij} e^{-(m_j - m_{\min})}. \tag{23}$$

Substituting $b_i = \hat{b}_i e^{-m_{\min}}$ back into the normalized dynamics, the large exponential terms cancel:

$$\frac{da_i}{dt} = \frac{a_i (\hat{b}_i e^{-m_{\min}})}{\|a \odot (\hat{b} e^{-m_{\min}})\|_2} = \frac{a_i \hat{b}_i}{\|a \odot \hat{b}\|_2}.$$

This gives the final, stable formulation for the normalized dynamics:

$$\frac{da_i}{dt} = \frac{a_i \sum_{j=1}^{K-1} \Psi_{ij} e^{-(m_j - m_{\min})}}{\sqrt{\sum_k \left(a_k \sum_{j=1}^{K-1} \Psi_{kj} e^{-(m_j - m_{\min})} \right)^2}} \tag{24}$$

In this form, all exponents are non-positive, with at least one exponent in the sums being exactly zero. This is numerically stable for large values of a.

Figure 10 compares the normalized flow (solid lines) with regular gradient flow (dashed lines) for the non-monotonic initialization from Figure 1. The acceleration is dramatic: while regular flow requires time scales beyond $t = 10^6$ to approach the NC equilibrium $\hat{a}_i = 1/7$ (top right), normalized flow achieves this by path length $\tau \approx 10^3$. Both the Lyapunov function D_{KL} (bottom left) and metric M (bottom right) converge orders of magnitude faster under normalized flow.

The non-monotonic behavior of metric M observed in Figure 1 remains visible: normalized flow (solid blue) exhibits a similar but compressed trajectory where M first increases before eventually decreasing. Notably, the singular values (top left) become extremely large when the decrease in M begins. Without the stable implementation above, we would have encountered numerical overflow at this stage. This confirms that normalized flow preserves the qualitative dynamics while making convergence tractable on practical time scales.

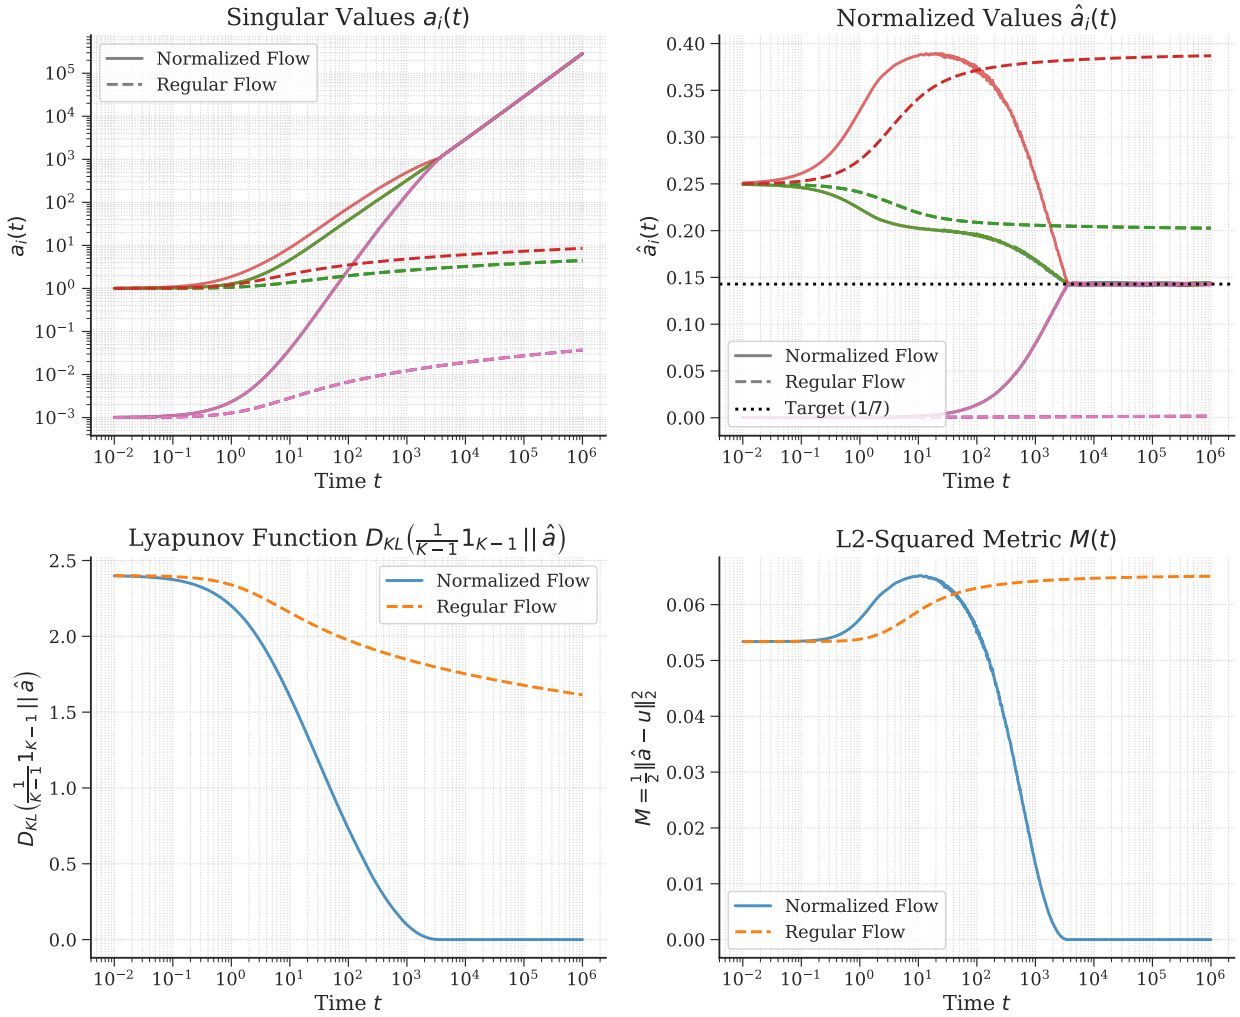

Figure 10: **Normalized Flow vs. Regular ODE Flow** ($K = 8$)**.** A direct comparison of the normalized flow (solid lines) and the regular stable ODE flow (dashed lines) for the highly non-uniform initialization $a(0) = [1, 1, 1, 1, 0.001, 0.001, 0.001]$.