

PUSHING BOUNDARIES: MIXUP’S INFLUENCE ON NEURAL COLLAPSE

Quinn LeBlanc Fisher*, Haoming Meng*, Vardan Papyan
University of Toronto

ABSTRACT

Mixup is a data augmentation strategy that employs convex combinations of training instances and their respective labels to improve the robustness and calibration of deep neural networks. Despite its widespread adoption, the nuanced mechanisms that underpin its success are not entirely understood. The observed phenomenon of Neural Collapse, where the last-layer activations and classifier of deep networks converge to a simplex equiangular tight frame (ETF), provides a compelling motivation to explore whether mixup induces alternative geometric configurations and whether those could explain its success. In this study, we delve into the last-layer activations of training data for deep networks subjected to mixup, aiming to uncover insights into its operational efficacy. Our investigation (code), spanning various architectures and dataset pairs, reveals that mixup’s last-layer activations predominantly converge to a distinctive configuration different than one might expect. In this configuration, activations from mixed-up examples of identical classes align with the classifier, while those from different classes delineate channels along the decision boundary. These findings are unexpected, as mixed-up features are not simple convex combinations of feature class means (as one might get, for example, by training mixup with the mean squared error loss). By analyzing this distinctive geometric configuration, we elucidate the mechanisms by which mixup enhances model calibration. To further validate our empirical observations, we conduct a theoretical analysis under the assumption of an unconstrained features model, utilizing the mixup loss. Through this, we characterize and derive the optimal last-layer features under the assumption that the classifier forms a simplex ETF.

1 INTRODUCTION

Consider a classification problem characterized by an input space $\mathcal{X} = \mathbb{R}^D$ and an output space $\mathcal{Y} := \{0, 1\}^C$. Given a training set $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, with $\mathbf{x}_i \in \mathcal{X}$ denoting the i -th input data point and $\mathbf{y}_i \in \mathcal{Y}$ representing the corresponding label, the goal is to train a model $f_\theta : \mathcal{X} \mapsto \mathcal{Y}$ by finding parameters θ that minimize the cross-entropy loss $\text{CE}(f_\theta(\mathbf{x}_i), \mathbf{y}_i)$ incurred by the model’s prediction $f_\theta(\mathbf{x}_i)$ relative to the true target \mathbf{y}_i , averaged over the training set, $\frac{1}{N} \sum_{i=1}^N \text{CE}(f_\theta(\mathbf{x}_i), \mathbf{y}_i)$.

Papyan, Han, and Donoho (2020) observed that optimizing this loss leads to a phenomenon called Neural Collapse, where the last-layer activations and classifiers of the network converge to the geometric configuration of a simplex equiangular tight frame (ETF). This phenomenon reflects the natural tendency of the networks to organize the representations of different classes such that each class’s representations and classifiers become aligned, equinorm, and equiangularly spaced, providing optimal separation in the feature space. Understanding Neural Collapse is challenging due to the complex structure and inherent non-linearity of neural networks. Motivated by the expressivity of overparametrized models, the *unconstrained features model* (Mixon et al., 2020) and the *layer-peeled model* (Fang et al., 2021) have been introduced to study Neural Collapse theoretically. These mathematical models treat the last-layer features as free optimization variables along with the classifier weights, abstracting away the intricacies of the deep neural network.

*Equal contribution

1.1 MIXUP

Mixup, a data augmentation strategy proposed by Zhang et al. (2017), generates new training examples through convex combinations of existing data points and labels:

$$\mathbf{x}_{ii'}^\lambda = \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_{i'}, \quad \mathbf{y}_{ii'}^\lambda = \lambda \mathbf{y}_i + (1 - \lambda) \mathbf{y}_{i'},$$

where $\lambda \in [0, 1]$ is a randomly sampled value from a predetermined distribution D_λ . Conventionally, this distribution is a symmetric $\text{Beta}(\alpha, \alpha)$ distribution, with $\alpha = 1$ frequently set as the default. The loss associated with mixup can be mathematically represented as:

$$\mathbb{E}_{\lambda \sim D_\lambda} \frac{1}{N^2} \sum_{i=1}^N \sum_{i'=1}^N \text{CE}(f_\theta(\mathbf{x}_{ii'}^\lambda), \mathbf{y}_{ii'}^\lambda). \quad (1)$$

A specific mixup data point, represented as $\mathbf{x}_{ii'}^\lambda$, is categorized as a *same-class* mixup point when $\mathbf{y}_i = \mathbf{y}_{i'}$, and classified as *different-class* when $\mathbf{y}_i \neq \mathbf{y}_{i'}$.

1.2 PROBLEM STATEMENT

Despite the widespread use and demonstrated efficacy of the mixup data augmentation strategy in enhancing the generalization and calibration of deep neural networks, its underlying operational mechanisms remain not well understood. The emergence of Neural Collapse prompts the following question:

Does mixup induce its own distinct configurations in last-layer activations, differing from traditional Neural Collapse? If so, does the configuration contribute to the method’s success?

This study aims to uncover the potential geometric configurations in the last-layer activations resulting from mixup and to determine whether these configurations can offer insights into its success.

1.3 CONTRIBUTIONS

Our contributions in this paper are twofold.

Empirical Study and Discovery We conduct an extensive empirical study focusing on the last-layer activations of mixup training data. Our study reveals that mixup induces a geometric configuration of last-layer activations across various datasets and models. This configuration is characterized by distinct behaviours:

- **Same-Class Activations:** These form a simplex ETF, aligning with their respective classifier.
- **Different-Class Activations:** These form channels along the decision boundary of the classifiers, exhibiting interesting behaviors: Data points with a mixup coefficient, λ , closer to 0.5 are located nearer to the middle of the channels. The density of different-class mixup points increases as λ approaches 0.5, indicating a collapsing behaviour towards the channels.

We investigate how this configuration varies under different training settings and the layer-wise trajectory the features take to arrive at the configuration. Additionally, the configuration offers insight into mixup’s success. Specifically, we measure the calibration induced by mixup and present an explanation for why the configuration leads to increased calibration.

Motivated by our theoretical analysis, we also examine the configuration of the last-layer activations obtained through training with mixup while fixing the classifier as a simplex ETF.

Theoretical Analysis We provide a theoretical analysis of the discovered phenomenon, utilizing an adapted *unconstrained features model* tailored for the mixup training objective. Assuming the classifier forms a simplex ETF under optimality, we theoretically characterize the optimal last-layer activations for all class pairs and for every $\lambda \in [0, 1]$.

1.4 RESULTS SUMMARY

The results of our extensive empirical investigation are presented in Figures 1, 3, 5, 10, and 12. These figures collectively illustrate a consistent identification of a unique last-layer configuration induced by mixup, observed across a diverse range of:

Architectures: Our study incorporated the WideResNet-40-10 (Zagoruyko & Komodakis, 2017) and ViT-B (Dosovitskiy et al., 2021) architectures;

Datasets: The datasets employed included FashionMNIST (Xiao et al., 2017), CIFAR10, and CIFAR100 (Krizhevsky & Hinton, 2009);

Optimizers: We used stochastic gradient descent (SGD), Adam (Kingma & Ba, 2017), and AdamW (Loshchilov & Hutter, 2017) as optimizers.

The networks trained showed good generalization performance and calibration, as substantiated by the data presented in Tables 1 and 2. That is, the values are comparable to those found in other papers (Zhang et al., 2017; Thulasidasan et al., 2020).

Beyond our principal observations, we conducted a counterfactual experiment, the results of which are depicted in Figures 2 and 8. These reveal a notable divergence in the configuration of the last-layer features when mixup is not employed. They also show that for MSE loss, the last-layer activations are convex combinations of the classifiers, which one may expect.

Furthermore, we juxtaposed the findings from our empirical investigation with theoretically optimal features, which were derived from an unconstrained features model and are showcased in Figure 6. To complement these results, we train models using mixup while fixing the classifier as a simplex ETF, and we plot the last-layer features in Figure 7. This yields last-layer features that align more closely with the theoretical features.

Table 1: Test accuracy for experiments in Figures 1 and 8.

Network	Dataset	Baseline	Mixup
WideResNet-40-10	FashionMNIST	95.10	94.21
	CIFAR10	96.2	97.30
	CIFAR100	80.03	81.42
ViT-B/4	FashionMNIST	93.71	94.24
	CIFAR10	86.92	92.56
	CIFAR100	59.95	69.83

2 EXPERIMENTS

2.1 MODEL TRAINING

We consider FashionMNIST (Xiao et al., 2017), CIFAR10, and CIFAR100 (Krizhevsky & Hinton, 2009) datasets. Unless otherwise indicated, for all experiments using mixup augmentation, $\alpha=1$ was used, meaning λ was sampled uniformly between 0 and 1. Each dataset is trained on both a Vision Transformer (Dosovitskiy et al., 2021), and a wide residual network (Zagoruyko & Komodakis, 2017). For each network and dataset combination, the experiment with the highest test accuracy is repeated without mixup and is referred to as the “baseline” result. No dropout was used in any experiments. Hyperparameter details are outlined in Appendix B.1.

2.2 VISUALIZATIONS OF LAST-LAYER ACTIVATIONS

For each dataset and network pair, we visualize the last-layer activations for a subset of the training dataset consisting of three randomly selected classes. After obtaining the last-layer activations, they undergo a two-step projection: first onto the classifier for the subset of three classes, then onto a two-dimensional representation of a three-dimensional simplex ETF. A more detailed explanation of the projection can be found in Appendix B.2. The results of this experiment can be seen in Figure 1. Notably, activations from mixed-up examples of the same classes closely align with a simplex ETF structure, whereas those from different classes delineate channels along the decision boundary. Additionally, in certain plots, activations from mixed-up examples of different classes become increasingly sparse as λ approaches 0 and 1. This suggests a clustering of activations towards the channels. When generating plots, we keep the network in train mode ¹. Since ViT does not have batch normalization layers, this difference is not applicable.

¹We choose to have the network in evaluation mode for the WideResNet-40-10 CIFAR100 combination because the batch statistics are highly skewed due to the high number of classes.




Figure 1: **(Visualization of activations outputted by networks trained with mixup)**. Last-layer activations of mixup training data for a randomly selected subset of three classes across various dataset and network architecture combinations trained with mixup. The first row illustrates activations generated by a WideResNet, while the second row showcases activations from a ViT. Each column corresponds to a different dataset. Coloration indicates the type of mixup (same or different class), along with the level of mixup, λ . For each plot, the relevant classifiers are plotted in black.

2.3 COMPARISON OF DIFFERENT LOSS FUNCTIONS

As part of our empirical investigation, we have conducted experiments utilizing Mean Squared Error (MSE) loss instead of cross-entropy, through which we observed in Figure 2 that features of mixed up examples are derived from simple convex combinations of same-class features. Initially, we anticipated a similar uninteresting configuration for cross-entropy; however, our measurements reveal that the resulting geometric configurations are markedly more interesting and complex. Additionally, we compare results in Figure 1 to the baseline (trained without mixup) cross-entropy loss in Figure 2. For the baseline networks, mixup data is loosely aligned with the classifier, regardless of same-class or different-class. The area in between classifiers is noisy and filled with examples where λ is close to 0.5. Additional baseline last-layer activations can be found in Figure 8.




Figure 2: **(Visualization of activations outputted by networks trained with various loss functions)**. Last-layer activations for WideResNet-40-10 trained on the CIFAR10 dataset, subsetted to three randomly selected classes. Projections are generated using the same method as Figure 1. **Left to right:** baseline cross-entropy, MSE mixup, cross-entropy mixup. Colouring indicates mixup type (same-class or different-class), and the level of mixup, λ . Relevant classifiers plotted in black. Additional dataset architecture combinations for baseline cross-entropy are available in appendix D.

2.4 LAYER-WISE TRAJECTORY OF CLS TOKEN

Using the same projection method as in Figure 1, we investigate the trajectory of the CLS token for ViT models. First, we randomly select two CIFAR10 training images. Then, we create a selection of mixed up examples using the respective images. For each mixed up example, we project the path of the CLS token at each layer of the ViT-B/4 network.

Figure 3 presents the results for two images of the same class, and two images of a different class. For different-class mixup, the plot shows that for very small λ , the network first classifies the image as class 1 and only in deeper layers it realizes the image is also partially class 2. Furthermore, the different-class activations in preceding layers appear to be convex combinations of the unmixed activations, akin to their inputs.

The results in Figure 1 suggest that applying mixup to input data enforces a particularly rigid geometric structure on the last-layer activations. Manifold mixup (Verma et al., 2019), a subsequent technique, proposes the mixing of features across various layers of a network. The results in Figure 3 suggest that using regular mixup promotes manifold mixup-like behaviour in previous layers.




Figure 3: **(Projection of CLS token at each layer).** Projections of the CLS token of the mixup of two randomly selected training images for various values of λ . Trajectories start at the origin. Colouring indicates mixup type (same-class or different-class), and the level of mixup, λ .

2.5 CALIBRATION

Thulasidasan et al. (2020) demonstrated that mixup improves calibration for networks. That is, training with mixup causes the softmax probabilities to be in closer alignment with the true probabilities of misclassification. To measure a network’s calibration, we use the expected calibration error (ECE) as proposed by Pakdaman Naeini et al. (2015). The exact definition of ECE can be found in Appendix C. Results for ECE can be found in Table 2. Last-layer activation plots for $\alpha = 0.4$ are available in Figure 9 in Appendix D.

Table 2: CIFAR10 expected calibration error.

Network	Baseline	Mixup ($\alpha = 1.0$)	Mixup ($\alpha = 0.4$)
WideResnet-40-10	0.024	0.077	0.013
ViT-B/4	0.122	0.014	0.019

The configuration presented in Figure 1 sheds light on why mixup improves calibration. Recall, mixup promotes alignment of the model’s softmax probabilities for the training example $\mathbf{x}_{ii'}^\lambda$ with its label $\lambda \mathbf{y}_i + (1 - \lambda) \mathbf{y}_{i'}$. Here, λ acts as a gauge for these softmax probabilities, essentially reflecting the model’s confidence in its predictions. Turning to Figure 4, it becomes therefore evident that as λ nears 0.5, the model’s certainty in its predictions diminishes. This reduction in confidence is manifested geometrically through the spatial distribution of features along the channel. This, in turn, causes an increase in misclassification rates, due to a greater chance of activations erroneously crossing the decision boundary. This simultaneous reduction in confidence and classification accuracy leads to enhanced calibration in the model and is purely attributed to the geometric structure to which the model converged. The above logic holds as we traverse the train mixed up features but we expect some test features to be noisy perturbations of mixed up train features.




Figure 4: **(Diagram showing the relationship between calibration and the configuration).** As λ approaches 0.5, the last-layer activation $h_{ii'}^\lambda$ (black) traverses the blue line of the configuration, leading to less confident predictions. Simultaneously, the variability of the activation (perforated black circle) results in an increase in misclassification due to the probability of being on the incorrect side of the decision boundary (green) increasing.

3 UNCONSTRAINED FEATURES MODEL FOR MIXUP

3.1 THEORETICAL CHARACTERIZATION OF OPTIMAL LAST-LAYER FEATURES

To study the resulting last-layer features under mixup, we consider an adaptation of the *unconstrained features model* to mixup training. Let $d \geq C - 1$ be the dimension of the last-layer features, $\mathbf{y}_i \in \mathbb{R}^C$ be the one-hot vector in entry i , $\mathbf{W} \in \mathbb{R}^{C \times d}$ be the classifier, and $\mathbf{h}_{ii'}^\lambda \in \mathbb{R}^d$ be the last-layer feature associated with target $\lambda \mathbf{y}_i + (1 - \lambda) \mathbf{y}_{i'}$. Then, adapting Equation 1 to the unconstrained features setting, we consider the optimization problem given by

$$\min_{\mathbf{W}, \mathbf{h}_{ii'}^\lambda} \mathbb{E}_{\lambda \sim D_\lambda} \frac{1}{C^2} \sum_{i=1}^C \sum_{i'=1}^C \left(\text{CE}(\mathbf{W} \mathbf{h}_{ii'}^\lambda, \lambda \mathbf{y}_i + (1 - \lambda) \mathbf{y}_{i'}) + \frac{\lambda \lambda_H}{2} \|\mathbf{h}_{ii'}^\lambda\|_2^2 \right) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2, \quad (2)$$

where $\lambda_W, \lambda_H > 0$ are the weight decay parameters. It is reasonable to consider decay on the features $\mathbf{h}_{ii'}^\lambda$, a practice that is frequently observed in prior work (Zhu et al., 2021; Zhou et al., 2022), due to the implicit decay to the last-layer features from the inclusion of decay in the previous layers’ parameters.

The following theorem characterizes the optimal last-layer features under the assumption that the optimal classifier \mathbf{W} is a simplex ETF, ie.

$$m \sqrt{\frac{C}{C-1}} (\mathbf{I}_C - \frac{1}{C} \mathbf{1}_C \mathbf{1}_C^\top) \mathbf{U}^\top, \quad (3)$$

where $\mathbf{I}_C \in \mathbb{R}^{C \times C}$ is the identity, $\mathbf{1}_C \in \mathbb{R}^C$ is the ones vector, $\mathbf{U} \in \mathbb{R}^{d \times C}$ is a partial orthogonal matrix (satisfying $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_C$), and $m \in \mathbb{R} \setminus \{0\}$ is its multiplier.

Note that we make this assumption as it holds in practice based on our empirical measurements, illustrated in Figure 5.

Theorem 3.1. *Assume that at optimality, \mathbf{W} is a simplex ETF with multiplier m , and denote the i -th row of \mathbf{W} by \mathbf{w}_i . Then, any minimizer of equation 2 satisfies:*

1) **Same-Class:** For all $i = 1, \dots, C$ and $\lambda \in [0, 1]$,

$$\mathbf{h}_{ii}^\lambda = \frac{(1 - C) K}{m^2} \mathbf{w}_i,$$




Figure 5: **(Convergence of classifier to simplex ETF)**. Measurements on the classifier, \mathbf{W} , for each network architecture and dataset combination. **First and third plot:** Coefficient of variation of the classifier norms, $\text{Std}_i(\|\mathbf{w}_i\|_2) / \text{Avg}_i(\|\mathbf{w}_i\|_2)$. **Second and fourth plot:** Standard deviation of the cosines between classifiers of distinct classes, $\text{Std}_{i,i' \neq i}(\langle \mathbf{w}_i, \mathbf{w}_{i'} \rangle / (\|\mathbf{w}_i\|_2 \|\mathbf{w}_{i'}\|_2))$ with $i \neq i'$. As training progresses, measurements indicate that \mathbf{W} is trending toward a simplex ETF configuration.

where $K < 0$ is the unique solution to the equation

$$e^{-CK} - \frac{Cm^2}{(1-C)\lambda_{\mathbf{H}}K} + C - 1 = 0.$$

2) **Different-Class:** For all $i \neq i'$ and $\lambda \in [0, 1]$,

$$\mathbf{h}_{ii'}^\lambda = \frac{(1-C)}{Cm^2} \left((K_\lambda - \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle) \mathbf{w}_i + ((C-1)K_\lambda + \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle) \mathbf{w}_{i'} \right),$$

where $\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle$ is of the form

$$\log \left(e^{K_\lambda} \left(2 - C + \frac{Cm^2}{(1-C)K_\lambda\lambda_{\mathbf{H}}} \pm \frac{1}{2} \sqrt{\left(C - 2 - \frac{Cm^2}{(1-C)K_\lambda\lambda_{\mathbf{H}}} \right)^2 - 4e^{-CK_\lambda}} \right) \right),$$

and $K_\lambda < 0$ satisfies

$$e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle} = \frac{Cm^2}{(1-C)K_\lambda\lambda_{\mathbf{H}}} e^{K_\lambda} \left(\frac{(1-C)\lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle}{Cm^2} + \lambda \right).$$

The proof of Theorem 3.1 can be found in Appendix A.1

Interpretation of Theorem. Theorem 3.1 establishes that, within the framework of our model’s assumptions, the optimal same-class features are independent of λ and align with the classifier as a simplex ETF. In contrast, the optimal features for different classes are linear combinations (depending on λ) of the classifier rows corresponding to the mixed-up targets, governed by the above equations. This is consistent with the observations in Figure 1, where the same-class features consistently cluster at simplex vertices, regardless of the value of λ , while the different-class features dynamically flow between these vertices as λ varies.

In Figure 6, we plot the last-layer features obtained from Theorem 3.1, numerically solving for the values of K and K_λ that satisfy their respective equations.

Similar to the empirical results in Figure 1, the density of different-class mixup points decreases as λ approaches 0 and 1. However, the theoretically optimal features exhibit channels arranged in a hexagonal pattern, differing from the empirical configuration observed in the FashionMNIST and CIFAR10 datasets. In particular, the empirical representations has a more pronounced elongation of different-class features as the mixup parameter λ approaches 0.5. In attempt to understand these differences, we introduce an amplification of these same features in the directions of the classifier rows not corresponding to the mixed-up targets, with increasing amplifications as λ gets closer to 0.5 (details of the amplification function are outlined in Appendix A.2). This results in features that behave more similarly to the empirical outcomes, while achieving a very close (though marginally larger) loss when compared to the true optimal configuration (loss values are indicated below each plot in Figure 6). This demonstrates that the features have some degree of flexibility while remaining in close proximity to the minimum loss.

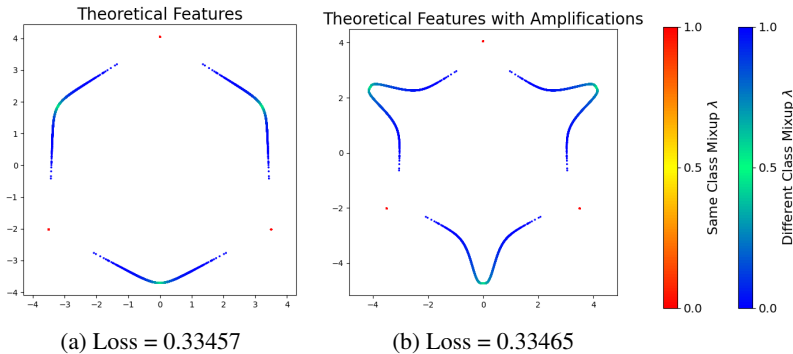


Figure 6: **(Optimal activations from unconstrained features model)**. On the left are optimal last-layer activations obtained from our theoretical analysis. We set $m = 3$, $C = 10$, $d = 100$, and $\lambda_H = 1 \times 10^{-6}$, and randomly sample 5000 different λ values from the Beta(1,1) distribution for a randomly selected subset of three classes. Projections are generated using the same method as depicted in Figure 1. Colouring indicates the mixup type (same-class or different-class), and the level of mixup, λ .

3.2 TRAINING WITH FIXED SIMPLEX ETF CLASSIFIER

To further understand the differences between theoretical activations (Figure 6) and empirical activations (Figure 1), we performed an experiment employing mixup within the training framework detailed in Section 2, but fixing the classifier as a simplex ETF. The resulting last-layer features are visualized in Figure 7.

Prior work (Zhu et al., 2021; Yang et al., 2022; Pernici et al., 2022) have explored the effects of fixing the classifier, but not in the context of mixup. Our observations reveal that when the classifier is fixed as a simplex ETF, the empirical features tend to exhibit a more hexagonal shape in its different-class mixup features, aligning more closely with the theoretically optimal features. Moreover, slightly higher generalization performance is achieved when compared to training with a learnable classifier under the same settings.

Based on these results, a possible explanation for the variation in configuration is that during training, the classifier is still being learned and requires several epochs to converge to a simplex ETF, as depicted in Figure 5. During this period, the features may traverse regions that lead to slightly suboptimal loss, as there is flexibility in the features’ structures without much degradation to the loss performance (depicted in Figure 6).

4 RELATED WORK

The success of mixup has prompted many mixup variants, each successful in their own right (Guo et al., 2018; Verma et al., 2019; Yun et al., 2019; Kim et al., 2020). Additionally, various works have been devoted to better understanding the effects and success of the method.

Guo et al. (2018) identified manifold intrusion as a potential limitation of mixup, stemming from discrepancies between the mixed-up label of a mixed-up example and its true label, and they propose a method for overcoming this.




Figure 7: **(Visualization of activations outputted by network trained with mixup fixing the classifier as a simplex ETF)**. Last-layer activations of mixup training data are presented here for a randomly selected subset of three classes. Coloration indicates the type of mixup (same-class or different-class), along with the level of mixup, λ . This model achieves a test accuracy of 97.35%.

In addition to the work by Thulasidasan et al. (2020) on calibration for networks trained with mixup, Zhang et al. (2022) posits that this improvement in calibration due to mixup is correlated with the capacity of the network. Zhang et al. (2021) theoretically demonstrates that training with mixup corresponds to minimizing an upper bound of the adversarial loss.

Chaudhry et al. (2022) delved into the linearity of various representations of a deep network trained with mixup. They observed that representations nearer to the input and output layer exhibit greater linearity compared to those situated in the middle.

Carratino et al. (2022) interprets mixup as an empirical risk minimization estimator employing transformed data, leading to a process that notably enhances both model accuracy and calibration. Continuing on the same path, Park et al. (2022) offers a unified theoretical analysis that integrates various aspects of mixup methods.

Furthermore, Chidambaram et al. (2021) conducted a detailed examination of the classifier optimal to mixup, comparing it with the classifier obtained through standard training.

Recent work has also been devoted to studying the benefits of mixup with feature-learning based analysis by Chidambaram et al. (2023) and Zou et al. (2023). The former considering two features generated from a symmetric distribution for each class and the latter considering a data model with two features of different frequencies, feature noise, and random noise.

The discovery of Neural Collapse by Papayan et al. (2020) has spurred investigations of this phenomenon. Recent theoretical inquiries by Mixon et al. (2020); Fang et al. (2021); Lu & Steinerberger (2020); E & Wojtowysch (2020); Poggio & Liao (2020); Zhu et al. (2021); Han et al. (2021); Tirer & Bruna (2022); Wang et al.; Kothapalli et al. (2022) have delved into the analysis of Neural Collapse employing both the unconstrained features model (Mixon et al., 2020) and the layer-peeled model (Fang et al., 2021). Liu et al. (2023) removes the assumption on the feature dimension and the number of classes in Neural Collapse and presents a Generalized Neural Collapse which is characterized by minimizing intra-class variability and maximizing inter-class separability.

To our knowledge, there has not been any investigation into the geometric configuration induced by mixup in the last layer.

5 CONCLUSION

In conclusion, through an extensive empirical investigation across various architectures and datasets, we have uncovered a distinctive geometric configuration of last-layer activations induced by mixup. This configuration exhibits intriguing behaviors, such as same-class activations forming a simplex equiangular tight frame (ETF) aligned with their respective classifiers, and different-class activations delineating channels along the decision boundary, with varying densities depending on the mixup coefficient. We also examine the layer-wise trajectory that features follow to reach this configuration in the last-layer, and measure the calibration induced by mixup to provide an explanation for why this particular configuration is beneficial for calibration.

Furthermore, we have complemented our empirical findings with a theoretical analysis, adapting the unconstrained features model to mixup. Theoretical results indicate that the optimal same-class features are independent of the mixup coefficient and align with the classifier, while different-class features are dynamic linear combinations of the classifier rows corresponding to mixed-up targets, influenced by the mixup coefficient. Motivated by our theoretical analysis, we also conduct experiments investigating the configuration of the last-layer activations from training with mixup while keeping the classifier fixed as a simplex ETF. We observe that they align more closely with the theoretically optimal features, with slight improvement in test-performance.

These findings collectively shed light on the intricate workings of mixup in training deep networks, emphasizing its role in organizing last-layer activations for improved calibration. Understanding these geometric configurations induced by mixup opens up avenues for further research into the design of data augmentation strategies and their impact on neural network training.

6 ACKNOWLEDGEMENTS

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). This research was enabled in part by support provided by Compute Ontario (<http://www.computeontario.ca/>) and Compute Canada (<http://www.computecanada.ca/>).

REFERENCES

- Luigi Carratino, Moustapha Cissé, Rodolphe Jenatton, and Jean-Philippe Vert. On mixup regularization, 2022.
- Arslan Chaudhry, Aditya Krishna Menon, Andreas Veit, Sadeep Jayasumana, Srikumar Ramalingam, and Sanjiv Kumar. When does mixup promote local linearity in learned representations?, 2022.
- Muthu Chidambaram, Xiang Wang, Yuzheng Hu, Chenwei Wu, and Rong Ge. Towards understanding the data dependency of mixup-style training. *CoRR*, abs/2110.07647, 2021. URL <https://arxiv.org/abs/2110.07647>.
- Muthu Chidambaram, Xiang Wang, Chenwei Wu, and Rong Ge. Provably learning diverse features in multi-view data with midpoint mixup, 2023.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- Weinan E and Stephan Wojtowytsch. On the emergence of tetrahedral symmetry in the final and penultimate layers of neural network classifiers. *arXiv preprint arXiv:2012.05420*, 2020.
- 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.
- Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization, 2018.
- XY 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*, 2021.
- Jang-Hyun Kim, Wonho Choo, and Hyun Oh Song. Puzzle mix: Exploiting saliency and local statistics for optimal mixup, 2020.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- Vignesh Kothapalli, Ebrahim Rasromani, and Vasudev Awatramani. Neural collapse: A review on modelling principles and generalization. *arXiv preprint arXiv:2206.04041*, 2022.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009. URL <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- Weiyang Liu, Longhui Yu, Adrian Weller, and Bernhard Schölkopf. Generalizing and decoupling neural collapse via hyperspherical uniformity gap, 2023.
- Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101, 2017. URL <http://arxiv.org/abs/1711.05101>.
- Jianfeng Lu and Stefan Steinerberger. Neural collapse with cross-entropy loss. *arXiv preprint arXiv:2012.08465*, 2020.
- Dustin G Mixon, Hans Parshall, and Jianzong Pi. Neural collapse with unconstrained features. *arXiv preprint arXiv:2011.11619*, 2020.

- Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1), Feb. 2015. doi: 10.1609/aaai.v29i1.9602. URL <https://ojs.aaai.org/index.php/AAAI/article/view/9602>.
- Vardan Papyan, X. Y. Han, and David L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, sep 2020. doi: 10.1073/pnas.2015509117. URL <https://doi.org/10.1073%2Fpnas.2015509117>.
- Chanwoo Park, Sangdoon Yun, and Sanghyuk Chun. A unified analysis of mixed sample data augmentation: A loss function perspective, 2022.
- Federico Pernici, Matteo Bruni, Claudio Baecchi, and Alberto Del Bimbo. Regular polytope networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(9):4373–4387, September 2022. ISSN 2162-2388. doi: 10.1109/tnnls.2021.3056762. URL <http://dx.doi.org/10.1109/TNNLS.2021.3056762>.
- Tomaso Poggio and Qianli Liao. Explicit regularization and implicit bias in deep network classifiers trained with the square loss. *arXiv preprint arXiv:2101.00072*, 2020.
- Sunil Thulasidasan, Gopinath Chennupati, Jeff Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. On mixup training: Improved calibration and predictive uncertainty for deep neural networks, 2020.
- Tom Tirer and Joan Bruna. Extended unconstrained features model for exploring deep neural collapse. In *international conference on machine learning (ICML)*, 2022.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, Aaron Courville, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states, 2019.
- Peng Wang, Huikang Liu, Can Yaras, Laura Balzano, and Qing Qu. Linear convergence analysis of neural collapse with unconstrained features. In *OPT 2022: Optimization for Machine Learning (NeurIPS 2022 Workshop)*.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017. URL <http://arxiv.org/abs/1708.07747>.
- 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?, 2022.
- Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features, 2019.
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks, 2017.
- Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *CoRR*, abs/1710.09412, 2017. URL <http://arxiv.org/abs/1710.09412>.
- Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup help with robustness and generalization? In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=8yKEo06dKNo>.
- Linjun Zhang, Zhun Deng, Kenji Kawaguchi, and James Zou. When and how mixup improves calibration, 2022.
- Jinxin Zhou, Chong You, Xiao Li, Kangning Liu, Sheng Liu, Qing Qu, and Zhihui Zhu. Are all losses created equal: A neural collapse perspective, 2022.

Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features, 2021.

Difan Zou, Yuan Cao, Yuanzhi Li, and Quanquan Gu. The benefits of mixup for feature learning, 2023.

A THEORETICAL MODEL

A.1 PROOF OF THEOREM 3.1

Our proof uses similar techniques as Yang et al. (2022), but we extend these ideas to the more intricate last-layer features that arise from mixup.

Proof. Assuming that \mathbf{W} is a simplex ETF with multiplier m , our unconstrained features optimization problem in equation 2 becomes separable across λ and i, i' , and so it suffices to minimize

$$L_{ii'}^\lambda = -\lambda \mathbf{y}_i \log \left(\frac{e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle}}{\sum_{k=1}^C e^{\langle \mathbf{w}_k, \mathbf{h}_{ii'}^\lambda \rangle}} \right) - (1-\lambda) \mathbf{y}_{i'} \log \left(\frac{e^{\langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle}}{\sum_{k=1}^C e^{\langle \mathbf{w}_k, \mathbf{h}_{ii'}^\lambda \rangle}} \right) + \frac{1}{2} \lambda_{\mathbf{H}} \|\mathbf{h}_{ii'}^\lambda\|_2^2$$

over each $\mathbf{h}_{ii'}^\lambda$, individually.

$$\begin{aligned} \frac{\partial L_{ii'}^\lambda}{\partial \mathbf{h}_{ii'}^\lambda} &= \mathbf{W}^\top (\mathbf{p} - (\lambda \mathbf{y}_i + (1-\lambda) \mathbf{y}_{i'})) + \lambda_{\mathbf{H}} \mathbf{h}_{ii'}^\lambda && (\mathbf{p} = \text{softmax}(\mathbf{W} \mathbf{h}_{ii'}^\lambda)) \\ &= \sum_{j=1}^C \mathbf{w}_j p_j - \lambda \mathbf{w}_i - (1-\lambda) \mathbf{w}_{i'} + \lambda_{\mathbf{H}} \mathbf{h}_{ii'}^\lambda && (p_j \text{ j-th entry of } \mathbf{p}) \\ &= \sum_{j \neq i, i'} \mathbf{w}_j p_j + (p_i - \lambda) \mathbf{w}_i + (p_{i'} - (1-\lambda)) \mathbf{w}_{i'} + \lambda_{\mathbf{H}} \mathbf{h}_{ii'}^\lambda. \end{aligned}$$

Setting $\frac{\partial L_{ii'}^\lambda}{\partial \mathbf{h}_{ii'}^\lambda} = 0$ gives

$$\sum_{j \neq i, i'} \mathbf{w}_j p_j + (p_i - \lambda) \mathbf{w}_i + (p_{i'} - (1-\lambda)) \mathbf{w}_{i'} + \lambda_{\mathbf{H}} \mathbf{h}_{ii'}^\lambda = 0. \quad (4)$$

Case: $i = i'$

In this case, equation 4 reduces to

$$\sum_{j \neq i} \mathbf{w}_j p_j + (p_i - 1) \mathbf{w}_i + \lambda_{\mathbf{H}} \mathbf{h}_{ii}^\lambda = 0. \quad (5)$$

Taking inner product with $\mathbf{w}_j, j \neq i$ in equation 5 gives

$$\begin{aligned} 0 &= m^2 p_j - \frac{m^2}{C-1} \sum_{k \neq i, j} p_k - \frac{m^2 (p_i - 1)}{C-1} + \lambda_{\mathbf{H}} \langle \mathbf{w}_j, \mathbf{h}_{ii}^\lambda \rangle \\ &= m^2 p_j - \frac{m^2}{C-1} \left(\sum_{k \neq j} p_k - 1 \right) + \lambda_{\mathbf{H}} \langle \mathbf{w}_j, \mathbf{h}_{ii}^\lambda \rangle \\ &= m^2 p_j \left(\frac{C}{C-1} \right) + \lambda_{\mathbf{H}} \langle \mathbf{w}_j, \mathbf{h}_{ii}^\lambda \rangle. \end{aligned}$$

Since $m^2, p_j, \frac{C}{C-1}, \lambda_{\mathbf{H}}$ are all positive, it follows that $\langle \mathbf{w}_j, \mathbf{h}_{ii}^\lambda \rangle < 0$.

For all $j, k \neq i$, we have

$$\frac{\exp(\langle \mathbf{w}_j, \mathbf{h}_{ii}^\lambda \rangle)}{\exp(\langle \mathbf{w}_k, \mathbf{h}_{ii}^\lambda \rangle)} = \frac{p_j}{p_k} = \frac{\langle \mathbf{w}_j, \mathbf{h}_{ii}^\lambda \rangle}{\langle \mathbf{w}_k, \mathbf{h}_{ii}^\lambda \rangle}.$$

Since $\frac{\exp(x)}{x}$ is strictly decreasing on $(-\infty, 0)$, in particular it is injective on $(-\infty, 0)$, so

$$\langle \mathbf{w}_j, \mathbf{h}_{ii}^\lambda \rangle = \langle \mathbf{w}_k, \mathbf{h}_{ii}^\lambda \rangle = K,$$

for some $K < 0$ and $p_j = p_k = p$ for all $j, k \neq i$, where

$$p = \lambda_{\mathbf{H}} \frac{1 - C}{C} \frac{K}{m^2}.$$

Let $S = \sum_{k=1}^C \exp \langle \mathbf{w}_k, \mathbf{h}_{ii}^\lambda \rangle$. Then, $\frac{e^K}{S} = p$ and so

$$S = \frac{e^K}{p} = \frac{C}{1 - C} \cdot \frac{m^2}{K \lambda_{\mathbf{H}}} e^K.$$

Then,

$$\begin{aligned} \mathbf{h}_{ii}^\lambda &= \frac{-1}{\lambda_{\mathbf{H}}} \left(\sum_{j \neq i} \mathbf{w}_j p_j + (p_i - 1) \mathbf{w}_i \right) \\ &= \frac{-1}{\lambda_{\mathbf{H}}} (-p \mathbf{w}_i + (p_i - 1) \mathbf{w}_i) \\ &= \frac{-1}{\lambda_{\mathbf{H}}} (-p \mathbf{w}_i - (C - 1) p \mathbf{w}_i) \\ &= \frac{1}{\lambda_{\mathbf{H}}} C p \mathbf{w}_i \\ &= \frac{(1 - C) K}{m^2} \mathbf{w}_i. \end{aligned}$$

Taking inner product with \mathbf{w}_i in equation 5 gives

$$\begin{aligned} 0 &= -\frac{m^2}{C - 1} \sum_{k \neq i} p_k + m^2 (p_i - 1) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{ii}^\lambda \rangle \\ &= m^2 p_i \left(\frac{1}{C - 1} + 1 \right) - m^2 \left(\frac{1}{C - 1} + 1 \right) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{ii}^\lambda \rangle \\ &= m^2 \frac{C}{C - 1} (p_i - 1) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{ii}^\lambda \rangle \\ &= m^2 \frac{C}{C - 1} (-(C - 1)p) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{ii}^\lambda \rangle, \end{aligned}$$

and so

$$\langle \mathbf{w}_i, \mathbf{h}_{ii}^\lambda \rangle = (1 - C)K. \quad (6)$$

By our definition of \mathbf{p} as the softmax applied on $\mathbf{W} \mathbf{h}_{ii}^\lambda$, we have

$$\begin{aligned} e^{\langle \mathbf{w}_i, \mathbf{h}_{ii}^\lambda \rangle} &= S p_i \\ &= \frac{C}{1 - C} \cdot \frac{m^2}{K \lambda_{\mathbf{H}}} e^K (1 - (C - 1)p) \\ &= \frac{C}{1 - C} \cdot \frac{m^2}{K \lambda_{\mathbf{H}}} e^K \left(1 + \frac{((C - 1)^2 \lambda_{\mathbf{H}} K)}{C m^2} \right). \end{aligned}$$

ie.

$$e^{-CK} = \frac{C m^2}{(1 - C) \lambda_{\mathbf{H}} K} + 1 - C.$$

So, K must satisfy $f(K) = 0$, where $f: (-\infty, 0) \rightarrow \mathbb{R}$ is defined by

$$f(x) = e^{-Cx} - \frac{C m^2}{(1 - C) \lambda_{\mathbf{H}} x} - (1 - C),$$

(note that we only consider the domain $(-\infty, 0)$ since we've shown that $K < 0$). We will show that there exists a unique K satisfying these properties.

$$f'(x) = -Ce^{-Cx} + \frac{Cm^2}{(1-C)\lambda_{\mathbf{H}}x^2} < 0,$$

since $-Ce^{-Cx} < 0$ and $\frac{Cm^2}{(1-C)\lambda_{\mathbf{H}}x^2} < 0$ (since all terms in the product are positive except $(1-C) < 0$) for all x . So f is strictly decreasing, and thus it is injective.

$\lim_{x \rightarrow 0^-} f(x) = -\infty$ and $\lim_{x \rightarrow -\infty} f(x) = \infty$, so by continuity of f , there exists $K < 0$ such that $f(K) = 0$, and K is unique by injectivity of f .

Case: $i \neq i'$

Taking inner product with $\mathbf{w}_j, j \neq i, i'$ in equation 4 and using properties of \mathbf{W} as a simplex ETF gives

$$\begin{aligned} 0 &= m^2 p_j - \frac{m^2}{C-1} \sum_{k \neq i, i', j} p_k - \frac{m^2(p_i - \lambda)}{C-1} - \frac{m^2(p_{i'} - (1-\lambda))}{C-1} + \lambda_{\mathbf{H}} \langle \mathbf{w}_j, \mathbf{h}_{i'}^\lambda \rangle \\ &= m^2 p_j - \frac{m^2}{C-1} \left(\sum_{k \neq j} p_k - 1 \right) + \lambda_{\mathbf{H}} \langle \mathbf{w}_j, \mathbf{h}_{i'}^\lambda \rangle \\ &= m^2 p_j \left(1 + \frac{1}{C-1} \right) + \lambda_{\mathbf{H}} \langle \mathbf{w}_j, \mathbf{h}_{i'}^\lambda \rangle \\ &= m^2 p_j \left(\frac{C}{C-1} \right) + \lambda_{\mathbf{H}} \langle \mathbf{w}_j, \mathbf{h}_{i'}^\lambda \rangle. \end{aligned}$$

By the same argument as in the previous case, we get that for all $j, k \neq i, i'$,

$$\langle \mathbf{w}_j, \mathbf{h}_{i'}^\lambda \rangle = \langle \mathbf{w}_k, \mathbf{h}_{i'}^\lambda \rangle = K_\lambda,$$

for some $K_\lambda < 0$ (we will omit the subscript λ for brevity as we are optimizing over each λ individually). Thus $p_j = p_k = p$ for all $j, k \neq i, i'$, where

$$p = \lambda_{\mathbf{H}} \frac{1-C}{C} \frac{K}{m^2}.$$

Let $S = \sum_{k=1}^C \exp \langle \mathbf{w}_k, \mathbf{h}_{i'}^\lambda \rangle$. Then, $\frac{e^K}{S} = p$ and so

$$S = \frac{e^K}{p} = \frac{C}{1-C} \cdot \frac{m^2}{K\lambda_{\mathbf{H}}} e^K. \quad (7)$$

Taking inner product with \mathbf{w}_i in equation 4 gives

$$\begin{aligned} 0 &= \frac{-m^2}{C-1} \sum_{j \neq i, i'} p_j + m^2(p_i - \lambda) - \frac{m^2}{C-1}(p_{i'} - (1-\lambda)) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{i'}^\lambda \rangle \\ &= \frac{-m^2}{C-1}(1 - p_i - p_{i'}) + m^2(p_i - \lambda) - \frac{m^2}{C-1}(p_{i'} - (1-\lambda)) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{i'}^\lambda \rangle \\ &= \frac{m^2}{C-1} p_i + m^2(p_i - \lambda) - \frac{m^2}{C-1} \lambda + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{i'}^\lambda \rangle \\ &= m^2 p_i \left(1 + \frac{1}{C-1} \right) - m^2 \lambda \left(1 + \frac{1}{C-1} \right) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{i'}^\lambda \rangle \\ &= m^2 \left(1 + \frac{1}{C-1} \right) (p_i - \lambda) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{i'}^\lambda \rangle. \end{aligned}$$

So,

$$\frac{Cm^2}{C-1} (p_i - \lambda) + \lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{i'}^\lambda \rangle = 0. \quad (8)$$

Similarly, taking inner product with $\mathbf{w}_{i'}$ in equation 4 gives us

$$\begin{aligned} 0 &= \frac{-m^2}{C-1} \sum_{j \neq i, i'} p_j - \frac{m^2}{C-1} (p_i - \lambda) + m^2 (p_{i'} - (1 - \lambda)) + \lambda_{\mathbf{H}} \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle \\ &= m^2 \left(1 + \frac{1}{C-1} \right) (p_{i'} - (1 - \lambda)) + \lambda_{\mathbf{H}} \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle, \end{aligned}$$

and so

$$\frac{Cm^2}{C-1} (p_{i'} - (1 - \lambda)) + \lambda_{\mathbf{H}} \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle = 0. \quad (9)$$

Summing equation 8 and equation 9 gives

$$\begin{aligned} 0 &= \frac{Cm^2}{C-1} (p_i + p_{i'} - 1) + \lambda_{\mathbf{H}} (\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle + \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle) \\ &= \frac{Cm^2}{C-1} (-(C-2)p) + \lambda_{\mathbf{H}} (\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle + \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle) \\ &= (C-2)\lambda_{\mathbf{H}}K + \lambda_{\mathbf{H}} (\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle + \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle), \end{aligned}$$

and so

$$\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle + \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle = -(C-2)K.$$

Then, using equation 7 and the definition of S gives us

$$\begin{aligned} \frac{C}{1-C} \cdot \frac{m^2}{K\lambda_{\mathbf{H}}} e^K &= S \\ &= \sum_{k=1}^C e^{\langle \mathbf{w}_k, \mathbf{h}_{ii'}^\lambda \rangle} \\ &= (C-2)e^K + e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle} + e^{-(C-2)K - \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle}, \end{aligned}$$

and thus

$$\left(e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle} \right)^2 + e^K \left(C - 2 - \frac{C}{1-C} \cdot \frac{m^2}{K\lambda_{\mathbf{H}}} \right) e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle} + e^{-(C-2)K} = 0. \quad (10)$$

We then solve the quadratic equation in $e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle}$ to get

$$e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle} = \frac{-e^K \left(C - 2 - \frac{Cm^2}{(1-C)K\lambda_{\mathbf{H}}} \right) \pm \sqrt{\left(e^K \left(C - 2 - \frac{Cm^2}{(1-C)K\lambda_{\mathbf{H}}} \right) \right)^2 - 4e^{-(C-2)K}}}{2}.$$

So, $\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle$ is of the form

$$\log \left(\frac{-e^K \left(C - 2 - \frac{Cm^2}{(1-C)K\lambda_{\mathbf{H}}} \right) \pm \sqrt{\left(e^K \left(C - 2 - \frac{Cm^2}{(1-C)K\lambda_{\mathbf{H}}} \right) \right)^2 - 4e^{-(C-2)K}}}{2} \right).$$

By our definition of \mathbf{p} as the softmax applied on $\mathbf{W}\mathbf{h}_{ii'}^\lambda$, K must satisfy

$$\begin{aligned} e^{\langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle} &= Sp_i \\ &= \frac{e^K}{p} \\ &= \frac{C}{1-C} \cdot \frac{m^2}{K\lambda_{\mathbf{H}}} e^K \left(\frac{(1-C)\lambda_{\mathbf{H}} \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle}{Cm^2} + \lambda \right). \end{aligned}$$

We have

$$\begin{aligned}
& \sum_{j \neq i, i'} \mathbf{w}_j p_j + (p_i - \lambda) \mathbf{w}_i + (p_{i'} - (1 - \lambda)) \mathbf{w}_{i'} \\
&= p(-\mathbf{w}_i - \mathbf{w}_{i'}) + (p_i - \lambda) \mathbf{w}_i + (p_{i'} - (1 - \lambda)) \mathbf{w}_{i'} \quad (\text{since } \sum_{j=1}^C \mathbf{w}_j = 0) \\
&= (p_i - \lambda - p) \mathbf{w}_i + (p_{i'} - (1 - \lambda) - p) \mathbf{w}_{i'}.
\end{aligned}$$

Substituting this into equation 4, we get

$$\begin{aligned}
\mathbf{h}_{ii'}^\lambda &= \frac{-1}{\lambda_H} ((p_i - \lambda - p) \mathbf{w}_i + (p_{i'} - (1 - \lambda) - p) \mathbf{w}_{i'}) \\
&= \frac{1}{\lambda_H} ((p - (p_i - \lambda)) \mathbf{w}_i + (p - (p_{i'} - (1 - \lambda)))) \mathbf{w}_{i'} \\
&= \frac{(1 - C)K\lambda_H - (1 - C)\lambda_H \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle}{\lambda_H C m^2} \mathbf{w}_i + \frac{(1 - C)K\lambda_H - (1 - C)\lambda_H \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle}{\lambda_H C m^2} \mathbf{w}_{i'} \\
&= \frac{(1 - C)}{C m^2} ((K - \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle) \mathbf{w}_i + (K - \langle \mathbf{w}_{i'}, \mathbf{h}_{ii'}^\lambda \rangle) \mathbf{w}_{i'}) \\
&= \frac{(1 - C)}{C m^2} ((K - \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle) \mathbf{w}_i + ((C - 1)K + \langle \mathbf{w}_i, \mathbf{h}_{ii'}^\lambda \rangle) \mathbf{w}_{i'}),
\end{aligned}$$

□

A.2 AMPLIFICATION OF THEORETICAL FEATURES

In this section we provide additional details of function used to generate the amplified features in Figure 6.

We define $\epsilon(\lambda) = \frac{4}{5} \exp(-20(\lambda - 0.5)^4) - \frac{2}{5}$. Then for the different class features ($i \neq i'$), define the amplified features as $\tilde{\mathbf{h}}_{ii'}^\lambda = \mathbf{h}_{ii'}^\lambda - \epsilon(\lambda) \sum_{j \neq i, i'} \mathbf{w}_j$.

The motivation for the function form of $\epsilon(\lambda)$ is to ensure that it is symmetric about $\lambda = 0.5$, increasing when $\lambda < 0.5$ and decreasing when $\lambda > 0.5$ with its maximum at $\lambda = 0.5$. These properties correspond to a larger amplification as λ approaches 0.5, while preserving symmetry in the amplifications. Note that the exact function $\epsilon(\lambda)$ is not important, but rather that it yields last-layer features that are closer to the empirical results (with elongations), while resulting in just a minor increase in loss. As mentioned in the main text, this implies that the features can have some deviation from the theoretical optimum without much change in the value of the loss.

B EXPERIMENTAL DETAILS

B.1 HYPERPARAMETER SETTINGS

For the WideResNet experiments, we minimize the mixup loss using stochastic gradient descent (SGD) with momentum 0.9 and weight decay 1×10^{-4} . All datasets are trained on a WideResNet-40-10 for 500 epochs with a batch size of 128. We sweep over 10 logarithmically spaced learning rates between 0.01 and 0.25, picking whichever results in the highest test accuracy. The learning rate is annealed by a factor of 10 at 30%, 50%, and 90% of the total training time.

For the ViT experiments, we minimize the mixup loss using Adam optimization (Kingma & Ba, 2017). For each dataset we train a ViT-B with a patch size of 4 for 1000 epochs with a batch size of 128. We sweep over 10 logarithmically spaced learning rates from 1×10^{-4} to 3×10^{-3} and weight decay values from 0 to 0.05, selecting whichever yields the highest test accuracy. The learning rate is warmed up for 10 epochs and is annealed using cosine annealing as a function of total epochs.

B.2 PROJECTION METHOD

For all of the last-layer activation plots, the same projection method is used. First, we randomly select three classes. We denote the centred last-layer activations for said classes by a matrix $H \in$

$\mathbb{R}^{m \times n}$ and the classifier of the network for said classes as $W \in \mathbb{R}^{3 \times m}$. The projection method is then as follows:

1. Calculate $USV^T = \text{SVD}(W^*)$ where W^* is the normalized classifier.
2. Define $Q = UV^T$
3. Let $A \in \mathbb{R}^{2 \times 3}$ be a two dimensional representation of a three dimensional simplex.
4. Compute $X = AQH$ and plot.

C EXPECTED CALIBRATION ERROR

To calculate the expected calibration error, first gather the predictions into M bins of equal interval size. Let B_m be the set of predictions whose confidence is in bin m . We can define the accuracy and confidence of a given bin as

$$\begin{aligned} \text{acc}(B_m) &= \frac{1}{|B_m|} \sum_{i \in B_m} \mathbf{1}(\hat{y}_i = y_i) \\ \text{conf}(B_m) &= \frac{1}{|B_m|} \sum_{i \in B_m} \hat{p}_i \end{aligned}$$

where \hat{p}_i is the confidence of example i . The expected calibration (ECE) is then calculated as

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{conf}(B_m)|$$

D ADDITIONAL LAST-LAYER PLOTS

Here we provide additional plots of last-layer activations. Namely, Figure 8 provides additional baseline last-layer activation plots for mixup data for every architecture and dataset combination in Figure 1. Figure 9 provides last-layer activations for the additional α value in Table 2. Figure 10 shows the evolution of the last-layer activations throughout training. Figure 11 shows last-layer activations for multiple random subsets of three classes. Finally, Figure 12 shows the last-layer activations for a ViT-B/4 trained on CIFAR10 using the AdamW optimizer.

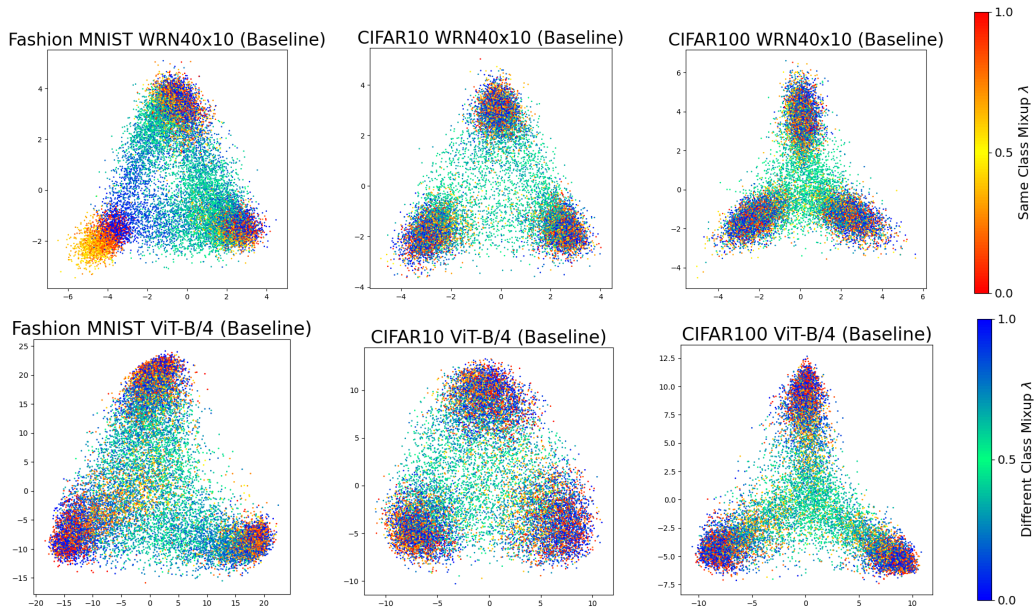


Figure 8: **(Visualization of activations outputted by networks trained without mixup)**. Last-layer activations for a randomly selected subset of three classes of mixup training data for various dataset and network architecture combinations. Projections are generated using the same method as Figure 1. All networks are trained using empirical risk minimization (no mixup). Colouring indicates mixup type (same-class or different-class), and the level of mixup, λ .

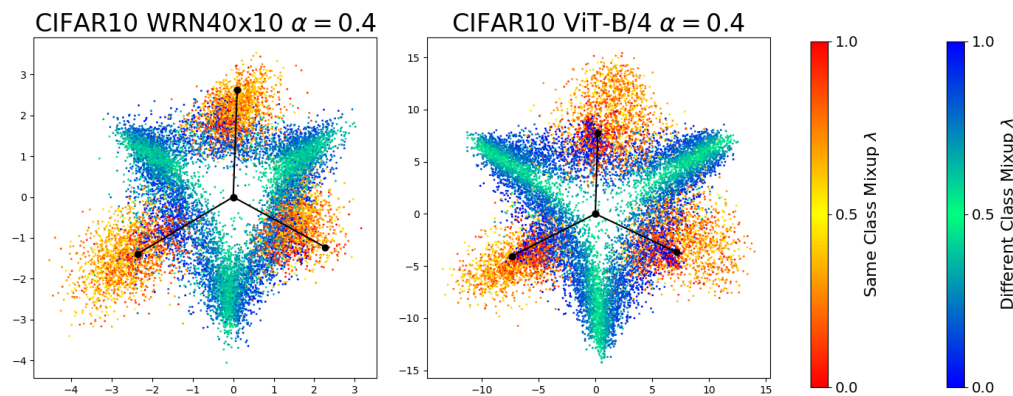


Figure 9: **(Visualization of activations outputted by networks trained with $\alpha = 0.4$)**. Last-layer activations for WideResNet-40-10 and ViT-B/4 trained on the CIFAR10 dataset, subsetted to three randomly selected classes. Projections are generated using the same method as Figure 1. For both cases, $\alpha = 0.4$ is used. Colouring indicates mixup type (same-class or different-class), and the level of mixup, λ . Relevant classifiers plotted in black.

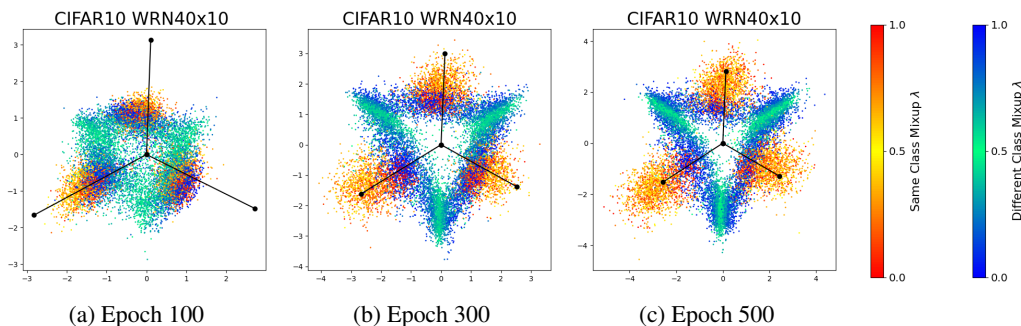


Figure 10: **(Activation convergence during training with mixup)**. Evolution of Last-layer activations for WideResNet-40-10 trained on CIFAR10 through training. Projections are generated in the same manner as Figure 1. Coloration indicates the type of mixup (same-class or different-class), along with the level of mixup, λ . Relevant classifiers plotted in black. As training progresses, different-class mixup points are pushed towards the decision boundary converging to the configuration depicted in Figure 1.

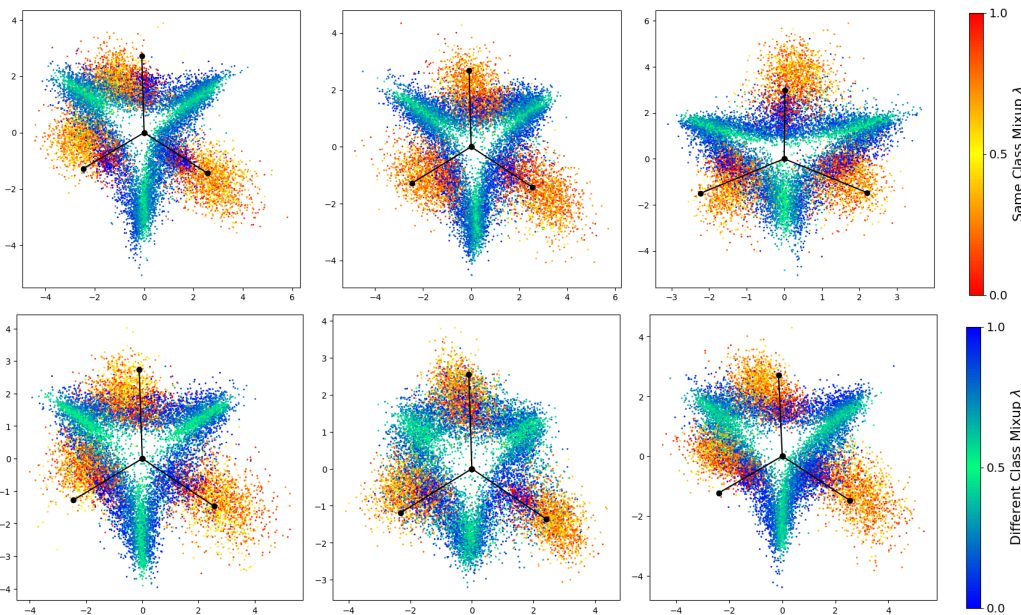


Figure 11: **(Visualization of last-layer activations for multiple subsets of classes)**. Last-layer activations for a randomly selected subsets of three classes of mixup training data for WRN-40-10 trained on CIFAR10. Projections are generated using the same method as Figure 1. Colouring indicates mixup type (same-class or different-class), and the level of mixup, λ . Black indicates relevant classifiers.




Figure 12: **(Visualization of activations outputted by ViT-B trained with mixup using AdamW).** Last-layer activations for ViT-B/4 trained following same training regiment as the ViT’s outlined in section 2 except with the AdamW optimizer. Projection is generated using the same method as Figure 1. Colouring indicates mixup type (same-class or different-class), and the level of mixup, λ . Black indicates relevant classifiers.