

# NEAREST CLASS-CENTER SIMPLIFICATION THROUGH INTERMEDIATE LAYERS

**Ido Ben-Shaul**

Department of Applied Mathematics,  
Tel-Aviv University, Israel  
eBay Research  
ido.benshaul@gmail.com

**Shai Dekel**

Department of Applied Mathematics,  
Tel-Aviv University, Israel  
shaidekel6@gmail.com

## ABSTRACT

Recent advances in neural network theory have introduced geometric properties that occur during training, past the Interpolation Threshold- where the training error reaches zero. We inquire into the phenomena coined *Neural Collapse* in the intermediate layers of the network, and emphasize the innerworkings of Nearest Class-Center Mismatch inside a deepnet. We further show that these processes occur both in vision and language model architectures. Lastly, we propose a Stochastic Variability-Simplification Loss (SVSL) that encourages better geometrical features in intermediate layers, yielding improvements in both train metrics and generalization.

## 1 INTRODUCTION

Several recent works have investigated the dynamics of modern deep neural networks (deepnets) past the point of zero training error Belkin (2021); Nakkiran et al. (2020); Bartlett et al. (2021); Power et al. (2022). The stage at which the training error reaches zero is called the *Interpolation Threshold (IT)*, since at this point the learned network function interpolates between training samples. This is not to be confused with zero-loss, but simply the point where all training samples are correctly classified. The stage of training beyond the IT is coined the *Terminal Phase of Training (TPT)* in Pappayan et al. (2020). It was in this paper that the term *Neural Collapse (NC)* was introduced to describe four interconnected geometrical phenomena that describe the network behavior past the TPT. Let us briefly describe the properties of NC that are most relevant for this paper:

**(NC1) Variability collapse:** As training progresses, the within-class variation of the activations becomes negligible as these activations collapse to their class-means.

**(NC4) Simplification to Nearest Class-Center (NCC):** For a given deepnet activation, the network classifier converges to choosing whichever class has the nearest train class-mean (in standard euclidean distance).

Following this work, the theory of NC has been widely studied Zhu et al. (2021); Rangamani et al. (2021); Lu & Steinerberger (2020); Fang et al. (2021); Galanti et al. (2021); Han et al. (2021). A different line of work Gormez & Koyuncu (2021) has studied Property (NC4) as a method for Early Exits in deepnets.

Various publications have proposed that geometrical properties of intermediate layers shape the way deepnets are trained, and largely affect their successes Cohen et al. (2020); Allen-Zhu & Li (2020); Ben-Shaul & Dekel (2021); Liu & Arik (2020); Alain & Bengio (2017); Baldock et al. (2021). Further, Galanti et al. (2021) have suggested that variability reduction (NC1) appears in intermediate layers.




Figure 1: **STL-10** experiment training procedure. We show the accuracy on both the train and test set, using the different losses: Vanilla Cross-Entropy (**Pink**) and SVSL (**Blue**). The SVSL loss outperforms the vanilla in both metrics. Both models improve in terms of test-performance during the TPT, shown as the shaded area. In the TPT, both models have a constant training accuracy of 1.

It is still early to determine the effects on NC on generalization. Multiple contributions have suggested that intermediate layer NCC simplification promotes large margins and generalization Sokolic et al. (2017); Galanti et al. (2021); Elsayed et al. (2018). On the contrary, Mixon et al. (2020) have shown that it is possible to achieve NC with random labels. The results shown in this work suggest intermediate clustering promotion assists the deepnet in generalization for a non-noisy setting. We delve deeper into the inner workings of NCC Simplification along the lines of Pappayan et al. (2020).

### 1.1 OUR CONTRIBUTIONS

Our contributions can be summarized as follows:

**NCC Simplification in Intermediate Model Layers:** We show that when looking at the NCC Mismatch of deepnet intermediate layers, before and during TPT, there is a beautiful geometric structure that emerges. Namely:

- (i) There is a monotone ordering between NCC mismatch in intermediate layers. The mismatch is lower as the layers gets deeper.
- (ii) NCC Simplification is not only apparent in the final layer of the network, and may propagate back several layers in the network.

**NCC-Simplification is apparent in NLP sequence classification tasks using transformers:** When proposed in Pappayan et al. (2020), the authors show that Neural Collapse appears in several well-known Image Classification models: VGG Simonyan & Zisserman (2015), ResNet He et al. (2016), and DenseNet Huang et al. (2017) on Image Classification datasets. In this paper, we show that NCC simplification is also apparent in Transformer architectures Vaswani et al. (2017), and even more surprisingly, in common NLP tasks. The recent surge in transformer architectures for cross-modal tasks, suggests that there are common behaviors between classic vision architectures, and more recent mechanisms Radford et al. (2021); Dai et al. (2021); Raghu et al. (2021).

**Encouraging Variability Simplification can assist in training and generalization:** We propose a simple intermediate layer variance collapsing loss which we coin the Stochastic Variability Simplification Loss (SVSL). This loss is shown to improve the performance of a wide-variety of tasks by encouraging NCC simplification during training. We show that this loss is able to improve NCC mismatch between intermediate layers, on both the train and test datasets. The different training stages and their respective metrics can be seen for both losses on the STL-10 dataset in Figure 1. The same plot is given for all other datasets in the full manuscript <sup>1</sup>. In all plots throughout this work, the  $x$ -axis represents the epochs during train.

<sup>1</sup>For the full manuscript, please see <https://arxiv.org/pdf/2201.08924.pdf>




Figure 2: NCC Mismatch for Vision datasets: **MNIST**, **Fashion-MNIST** and **CIFAR10**, using both vanilla (**solid**) and SVSL (**dashed**) losses. **Top**: train NCC Mismatch, **Bottom**: test NCC Mismatch. The shaded pink background shows the TPT for the vanilla loss experiment, and the blue for the SVSL. The background is shaded purple at epochs when both experiments are in the TPT.

## 2 PROBLEM SETUP AND BACKGROUND

### 2.1 SUPERVISED CLASSIFICATION

We study the supervised classification setting. In this paper, our experiments include problems in both Vision and NLP- Sequence Classification. In **Image Classification** we are given a training set of  $d := 3 \times W \times H$ - dimensional RGB images, from  $C$  categories. We wish to train a network to differentiate between different image classes. On the other hand, in **Text-Sequence Classification** we are given a training set of sequences. We use standard tokenization techniques to transfer discrete sequences to a continuous euclidean space. Since different sequences may be in different lengths, we pad sequences in our experiments to a certain  $d := \text{MAX\_PAD}$  constant. Similar to the vision setting, we wish to find the appropriate class for each text-sequence (of dimension  $\mathbb{R}^d$ , for  $C$  ground-truth classes).

Let  $g$  represent a deepnet,  $g : \mathbb{R}^d \rightarrow \mathbb{R}^C$  where  $C$  is the number of classes, and network parameters  $\theta$  which are learned through the optimization procedure, or “training”. For the classification setting, the network decision for input  $x \in \mathbb{R}^d$  is defined as  $z = \arg \max_{1 \leq c' \leq C} g(x)_{c'}$ . Let  $L = \{l^{(1)}, \dots, l^{(k)}\}$  represent the set of intermediate layers of network  $g$  (formally defined in Section 5.3), such that  $g := l^{(k)} \circ \dots \circ l^{(1)}$ . We define  $n_j$  as the output dimension of layer  $j$ , such that  $n_0 := d$  and  $n_k := C$ . We note  $g^{(j)}(x)$  as the outputs of the  $j^{\text{th}}$  layer of the network for input sample  $x$ . Using these definitions, the following holds:  $g^{(j)}(x) := l^{(j)} \circ \dots \circ l^{(1)}(x)$ . Given a function  $f$  and a set of indices  $\mathcal{I}$  we introduce the streamlined notation  $f_i := f(x_i)$ .

### 2.2 REPRESENTATION LEARNING

In recent years, the fields of Vision and NLP have both been transformed by representation learning methods in supervised and unsupervised tasks. The main premise is to learn “representations such that similar samples stay close to each other, while dissimilar ones are far apart” Weng (2021). Encouraging clustering in features learned by a deepnet has been a pivotal early idea to improve representations Xie et al. (2016); Tian et al. (2017). In this paper we will show that encouraging clustering of intermediate layers can boost performance. We use the normal convention of calling a sample the *Anchor*, a similar example *Positive* and a dissimilar example as *Negative*. In Vision, representation learning can be approached by using clustering assignments as pseudo-labels Tian et al. (2017), invoking similarity between different augmentations of the same sample Zbontar et al. (2021); Chen et al. (2020); Caron et al. (2021), or even using class labels to wisely pick Positive and Negative items Khosla et al. (2020). In Language Modeling(LM), Masking and Next-Sentence prediction(among others) tasks are used to learn semantically robust representations Devlin et al.

(2019); Liu et al. (2019). We use the notion of representation learning by enforcing low inner-class variability. The methods mentioned only use the penultimate layer of the encoder to penalize the representations. We conjecture that better consistency of class representations in intermediate layers forces representations in final layers to have better geometrical features. Our method does not need to sample pairs (positive or negative), and is simple to implement.

### 3 NEURAL COLLAPSE

We shall now briefly present the particular properties of Neural Collapse that are relevant for this work, as presented in Pappayan et al. (2020); Han et al. (2021). Let  $g$  be a given network and  $\{(x_i, y_i)\}_{i \in \mathcal{I}_{\text{Train}}}, \{(x_i, y_i)\}_{i \in \mathcal{I}_{\text{Test}}}$  be the train and test set resp.

**Definition 1** (Train class-means). *We define the train class means for layer  $l^{(j)}$  and class  $1 \leq c \leq C$  as*

$$\mu_c^{(j)} := \text{Avg}_{i \in \mathcal{I}_{\text{Train}}, y_i=c} \{g_i^{(j)}\}.$$

**Definition 2** (Train within-class covariance). *We define the train within-class covariance for layer  $l^{(j)}$  as*

$$\Sigma_W^{(j)} := \frac{1}{C} \sum_{c=1}^C \text{Avg}_{i \in \mathcal{I}_{\text{Train}}, y_i=c} \{(g_i^{(j)} - \mu_c^{(j)})(g_i^{(j)} - \mu_c^{(j)})^\top\}.$$

Using these definitions, we can formally define properties (NC1) and (NC4) from Pappayan et al. (2020). Let us assume that the network  $g$  may be split into two stages: the ‘‘feature engineering’’ stage  $g^{(k-1)}$  and the final classifier layer,  $l^{(k)}$  such that  $g := l^{(k)} \circ g^{(k-1)}$ .

**Definition 3** (NC1 Variability Collapse).  $\Sigma_W^{(k-1)} \rightarrow 0$ .

**Definition 4** (NC4 Simplification to NCC). *Let:*

$$S := \left\{ i \in \mathcal{I}_{\text{Train}} \mid \arg \max_{1 \leq c \leq C} g(x_i)_c \neq \arg \min_{1 \leq c \leq C} \left\| g^{(k-1)}(x_i) - \mu_c^{(k-1)} \right\|_2 \right\},$$

then  $|S| \rightarrow 0$ , where  $|X|$  is the number of elements in a finite set  $X$ .

In both definitions, the arrow represents the progress with the optimization procedure. Throughout this paper, we make the assumption that the deepnets are of proper capacity to reach the TPT, or in other words to ‘‘fit’’ the data. In the experiments we use large architectures that are highly overparameterized for the given tasks.

## 4 CONTRIBUTIONS

### 4.1 NCC MISMATCH IN INTERMEDIATE LAYERS

The results of Pappayan et al. (2020) show a clear behavior in terms of NCC-Simplification in the penultimate layer (see Definition 4). We further investigate the behavior of intermediate layers in terms of NCC mismatch. Let us define:

**Definition 5** (Layer  $j$  Train NCC mismatch).

$$\Lambda_{\text{Train}}^{(j)} := \frac{1}{N_{\text{Train}}} \left| \left\{ \arg \max_{1 \leq c' \leq C} g(x_i)_{c'} \neq \arg \min_{c'} \left\| g^{(j)}(x_i) - \mu_{c'}^{(j)} \right\|_2 \mid i \in \mathcal{I}_{\text{Train}} \right\} \right|. \quad (1)$$

**Definition 6** (Layer  $j$  Test NCC mismatch).

$$\Lambda_{\text{Test}}^{(j)} := \frac{1}{N_{\text{Test}}} \left| \left\{ \arg \max_{1 \leq c' \leq C} g(x_i)_{c'} \neq \arg \min_{c'} \left\| g^{(j)}(x_i) - \mu_{c'}^{(j)} \right\|_2 \mid i \in \mathcal{I}_{\text{Test}} \right\} \right|. \quad (2)$$

**Our conjectures can now be described as follows**

**Conjecture 7** (Intermediate Layer ordering using NCC mismatch). *There is a clear order between both **train** and **test** NCC mismatch in intermediate layers. The mismatch is lower as the layers get deeper. In the TPT, for  $1 \leq j \leq k$ ,*

$$\Lambda_{\text{Train}}^{(j)} \geq \Lambda_{\text{Train}}^{(j+1)} \quad \text{and} \quad \Lambda_{\text{Test}}^{(j)} \geq \Lambda_{\text{Test}}^{(j+1)}. \quad (3)$$

**Conjecture 8** (NCC mismatch improves in TPT). *At each intermediate layer, both the **train** and **test** NCC mismatch decreases from the IT to the End of Training(EOT).*

$$\Lambda_{\text{Train,IT}}^{(j)} \geq \Lambda_{\text{Train,EOT}}^{(j)} \quad \text{and} \quad \Lambda_{\text{Test,IT}}^{(j)} \geq \Lambda_{\text{Test,EOT}}^{(j)}. \quad (4)$$

#### 4.2 DECREASING NCC MISMATCH USING STOCHASTIC VARIABILITY-SIMPLIFICATION LOSS

In Papyan et al. (2020), the properties shown in Definitions 3 and 4 both act on the penultimate layer of the network. When considering intermediate feature spaces, the property shown in Definition 3 amounts to promoting class clustering. Promoting class clustering can push samples further from decision boundaries between classes, and could therefore increase agreement between the nearest class-center and the classifier. We wish to decrease the NCC mismatch during train, and encourage better clustering through the intermediate layers. Our loss function is proposed as follows:

**Definition 9** (Stochastic Train class-means). *Let  $\mathcal{B} := \{(x_i, y_i)\}_{i \in \mathcal{B}}$ , where  $|\mathcal{B}|$  is the Batch-Size. We define the stochastic train class means for layer  $l^{(j)}$ , batch  $\mathcal{B}$ , and class  $1 \leq c \leq C$  as*

$$\mu_{c,\mathcal{B}}^{(j)} := \text{Avg}_{i \in \mathcal{B}, y_i = c} \{g_i^{(j)}\}.$$

**Definition 10** (Stochastic Variability-Simplification Loss (SVSL)). *Let  $g$  be a deepnet and  $\hat{y}_i = g(x_i)$  for  $(x_i, y_i)$ ,  $i \in \mathcal{I}_{\text{Train}}$ ,  $y_i = c$ ,  $1 \leq c \leq C$ . Let  $\mathcal{B}$  be the batch such that  $i \in \mathcal{B}$ . We also define  $\gamma \in \mathbb{N}$ ,  $1 \leq \gamma < k$  and  $\alpha \in \mathbb{R}_+$  two hyperparameters. The Stochastic Variability-Simplification Loss function is then defined as*

$$\mathcal{L}(\hat{y}_i, y_i) := \text{CE}(\hat{y}_i, y_i) + \eta \sum_{j=\gamma}^k \left\| g^{(j)}(x_i) - \mu_{c,\mathcal{B}}^{(j)} \right\|_2^2. \quad (5)$$

where CE is the well-known Cross-Entropy loss and

$$\eta = \frac{\alpha}{C(k+1-\gamma) |\{i \in \mathcal{B} \mid y_i = c\}|}.$$

The normalizing factor  $\eta$  serves as a mitigating factor in the case of unbalanced batches. It is possible to define the Variability-Collapse in a non-stochastic fashion, by computing the full class-means at layer for every epoch. An example implementation of the SVSL is given in the appendix.

**Using the SVSL, we claim the following behaviors:**

**Conjecture 11** (SVSL improves NCC mismatch). *Using the properly defined hyperparameters  $\alpha, \gamma$ , the Stochastic Variability-Simplification Loss encourages lower **train** and **test** NCC mismatch in intermediate layers. In the TPT, for  $1 \leq j \leq k$ ,*

$$\Lambda_{\text{Train, Vanilla}}^{(j)} \geq \Lambda_{\text{Train, SVSL}}^{(j)} \quad \text{and} \quad \Lambda_{\text{Test, Vanilla}}^{(j)} \geq \Lambda_{\text{Test, SVSL}}^{(j)}. \quad (6)$$

**Conjecture 12** (SVSL can improve test-performance). *The EOT test metrics are improved for all datasets using the SVSL and proper hyperparameter tuning.*

##### 4.2.1 MOTIVATION FOR SVSL

*The Folding Ball Hypothesis* is presented in Chollet (2017) as follows: “Imagine two sheets of colored paper: one red and one blue. Put one on top of the other. Now crumple them together into a small ball. That crumpled paper ball is your input data, and each sheet of paper is a class of data in a classification problem. What a neural network is meant to do is figure out a transformation of the paper ball that would uncrumple it, so as to make the two classes cleanly separable again”. This geometrical notion has been used to try to predict the wellness of such transformations, using their geometrical properties Cohen et al. (2020); Ben-Shaul & Dekel (2021); Alain & Bengio (2017); Montavon et al. (2011). When measuring NCC mismatch during TPT, the network has near 0-training-error. This essentially means that the final feature space (where the inputs to the classifier reside) has near perfect clusters per-class. In Papyan et al. (2020), it is empirically shown that for most deepnets, the penultimate layer has a single cluster for each class. Thus, measuring the **train** NCC mismatch between the  $j$ th feature-space and the classifier is similar to checking the NCC mismatch with the ground-truth labels.

The clustering of feature spaces is an iterative transformation from each layer to the next, where the quality of clustering assists in clustering at the following stage.

Demanding a low NCC mismatch in early layers of the network may be unsatisfiable, as the input samples (e.g. images) cannot necessarily be well clustered with such low capacity (small number of layers). This is the reason we allow the  $\gamma$  hyperparameter to facilitate the earliest layer from which we require the SVSL. Demanding consistency between the NCC and the classifier early in training can interfere with the model learning the proper class predictions, so we leverage between the losses using the  $\alpha$  hyperparameter. A different approach can consist of applying the SVSL only during TPT.

## 5 EXPERIMENT DETAILS

Our experiments aim to demonstrate Conjectures 7,8,11,12 on both Vision and NLP tasks. In Section 5.1 we introduce the datasets that were used. In Section 5.2 we describe the architectures and Section 5.3 goes through the training procedures used.

### 5.1 DATASETS

For **Vision** tasks, we use most of the datasets used in Pappan et al. (2020). Namely: MNIST, FashionMNIST, CIFAR10, CIFAR100, and STL10. Unlike in Pappan et al. (2020), we do not balance the datasets explicitly and keep them as they are. We use mean-std train normalization. In order to get intermediate features, we use PyTorch Hooks Paszke et al. (2019). For the **NLP sequence classification tasks** we use a subset of binary datasets from the GLUE benchmark Wang et al. (2018). We run our experiments on datasets from all three types of tasks: *Single-Sentence Tasks*: CoLA and SST-2, *Similarity and Paraphrase Tasks*: MRPC, and *Inference Tasks*: RTE. All datasets have 2 classes. In order to make all sequences of the same length, both for computing NCC mismatch and maintaining same size features, we pad each of the sequences in all datasets to 32 tokens. Intermediate features are readily given as “hidden\_states” in Wolf et al. (2020).

### 5.2 ARCHITECTURES

**Vision:** For the vision architectures we follow the guidelines set in Pappan et al. (2020). In this paper we use solely the ResNet He et al. (2016) architectures. ResNet18 is used for MNIST, FashionMNIST, and CIFAR10. For CIFAR100 and STL10 the model chosen is the ResNet50 architecture. The layers for the ResNet architecture that are used in the experiments are {Layer1, Layer2, Layer3, Layer4, AvgPool, FC} as implemented in TorchVision Marcel & Rodriguez (2010). **NLP Sequence Classification:** For all sequence classification task we use an Uncased pre-trained BERT Devlin et al. (2019). The layers used for this architecture are the hidden states of the BERT-architecture. We include the embedding-layer in the BERT architecture, and use all hidden-states except the final output layer. In total, we have an initial embedding features (1) and (11) hidden-state layers, for a total of 12 layers in this architecture. In theory, the penultimate layer can also be used in the optimization process.

### 5.3 OPTIMIZATION PROCEDURE

**Vision:** We use the same optimization scheme as in Pappan et al. (2020), using best training hyperparameters as logged, and follow the same training procedure. We train all datasets for 350 Epochs. The Batch-Size for all experiments is 128. All vision experiments are trained using a SGD optimizer as done in the original paper. All SVSL Hyperparameters used are given in Table 2 in Appendix C. We report the top-1-accuracy on the test datasets. We use a threshold of 0.995 for determining the Interpolation Threshold.

**NLP Sequence Classification:** We follow the default hyperparameters as shown in Wolf et al. (2020) (GLUE finetune example). All experiments are trained using an AdamW Loshchilov & Hutter (2019) optimizer, and the default hyperparameters for 10 epochs. The Batch-Size for all experiments is 8, and the tasks are all binary classification. We report test-accuracy for the datasets: RTE, SST-2, and MRPC, and Matthew’s-Correlation for the CoLA dataset. **SVSL Hyperparameters:** The SVSL parameters are found using a simple baysean optimization scheme (AX-BoTorch Balandat et al. (2020)) for  $\alpha \in [5e - 8, 5e - 5]$  and layers  $\gamma \in \{1, \dots, 11\}$  on the test set. The purpose of these

Table 1: Comparing the Test Metrics of CE-Loss(Vanilla) with Stochastic Variability-Simplification Loss (SVSL) at IT, EOT, and Best-Test-Epoch for both Image and NLP-Sequence Classification Datasets. The metrics are as defined in Section 5.3 in percents. The Matthew’s -Correlation metric for the CoLA dataset is multiplied by a factor of 100. We also note whether the best Test-Metrics are achieved in TPT, for all datasets and methods.

Dataset	IT		EOT		Best Test Epoch			
	Vanilla	SVSL	Vanilla	SVSL	Vanilla	In TPT	SVSL	In TPT
MNIST	<b>99.37</b>	<b>99.36</b>	99.61	<b>99.69</b>	99.65	Yes	<b>99.69</b>	Yes
Fashion MNIST	91.78	<b>93.13</b>	93.82	<b>93.88</b>	93.93	Yes	<b>94.03</b>	Yes
STL10	53.41	<b>55.95</b>	54.11	<b>56.65</b>	54.19	Yes	<b>56.94</b>	Yes
CIFAR10	<b>80.64</b>	80.56	80.96	<b>81.19</b>	80.96	Yes	<b>81.19</b>	Yes
CIFAR100	52.77	<b>53.28</b>	53.31	<b>54.29</b>	53.79	Yes	<b>54.29</b>	Yes
CoLA	51.59	<b>52.91</b>	53.46	<b>55.54</b>	53.95	No	<b>55.54</b>	Yes
RTE	<b>58.84</b>	58.12	55.23	<b>59.57</b>	<b>61.01</b>	No	60.28	Yes
MRPC	70.83	<b>74.26</b>	74.26	<b>75.25</b>	75.00	No	<b>76.71</b>	No
SST-2	87.96	<b>88.42</b>	88.42	<b>88.76</b>	<b>89.22</b>	No	<b>89.22</b>	Yes

experiments is to show the ability to improve the network behavior using the SVSL. Possible future research includes adding the hyperparameters to as part of the network weights. The hyperparameters used are recorded in Table 2. We use a threshold of 0.985 for determining the IT.

## 6 RESULTS

### 6.1 NCC MISMATCH BEHAVIOR IN INTERMEDIATE LAYERS

We wish to demonstrate the Conjectures 7 and 8. The train and test NCC mismatch metrics for the MNIST, Fashion-MNIST, and CIFAR10 datasets are visualized in Figure 2 (Solid Lines). The same metrics for the sequence classification and all Vision experiments are given in the appendix.

### 6.2 VARIABILITY-SIMPLIFICATION LOSS

In this section we wish to demonstrate how using the intermediate-layer SVSL can improve training procedure and generalization. In Section 4.2.1 we describe the underlying logic behind the proposed cost. We advocate that in networks where intermediate NCC mismatch is lower, perform better in the TPT stage. Let us first demonstrate the correctness of Conjecture 11. Figure 2 (Dashed Line) shows the **train** and **test** NCC mismatch of the network using the SVSL with the proposed hyperparameters, for MNIST, FashionMNIST, and CIFAR10. The visualization for the remaining Image datasets is given in the appendix. It is clear that for all datasets, and almost all layers, the NCC mismatch improves when using the SVSL. The same conclusions can be seen for all NLP datasets in the complete manuscript.

We shall show the validity of Conjecture 12. Table 1 compares the test-performance of the vanilla Cross-Entropy (CE) loss with that of the SVSL on all datasets. This comparison is done at the IT, EOT, and also at the best Test-Epoch. We see that SVSL outperforms the vanilla CE at almost all stages of training. We also see that most datasets reach their best Test-Scores during TPT. Even when using the best possible Test Epoch, the SVSL loss achieves better or as-good results in all but one dataset. When the best Test Epoch is not achieved in the TPT, the scores achieved at the best Test-Epoch are comparable to the ones achieved at EOT. Practitioners in the field often look at regions of near-zero training-error, and use a validation set to choose the proper early-stopping criterion. This stage is formally given as the TPT, and hence a convincing method is to look at the performance mainly in this region. In these tasks we use the testing set as a proxy for the validation set. We see that even when allowing ourselves to look at all Test-Scores, the SVSL still achieves better performance on an array of tasks. All training graphs with both losses are given for the Vision datasets in Figures 1 and in the appendix. In practice, one may use a hold-out validation/cross-validation set to choose the best epoch and achieve similar results to the maximal points in the plots.

Table 2: SVSL-Hyperparameters used for **Image** and **NLP sequence-classification** tasks, optimized using BoTorch Balandat et al. (2020).

	$\alpha$	$\gamma$
<b>MNIST</b>	$1e - 5$	Layer 1
<b>F-MNIST</b>	$1e - 5$	Layer 1
<b>STL10</b>	$1e - 5$	Layer 1
<b>CIFAR10</b>	$4e - 6$	AvgPool
<b>CIFAR100</b>	$1e - 5$	Layer 3
<b>CoLA</b>	$5e - 6$	1
<b>RTE</b>	$1e - 7$	11
<b>MRPC</b>	$1e - 7$	11
<b>SST-2</b>	$4.132e - 5$	10

## 7 DISCUSSION AND POSSIBLE APPLICATIONS

We point out two additional perspectives when looking at the results given in this work:

- (i) **First-Layer NCC mismatch may suggest at the dataset “difficulty”:** When looking at NCC mismatch in early layers of the network, there is an interesting thought experiment that can be suggested. On one end, if the beginning layers have a very low NCC mismatch during the TPT - this means that the network is already achieving very good class clustering early on. The earlier this happens, the less capacity the model has to achieve this clustering. When looking at Figure 2 (MNIST) - we see that in the first layer, we already have a mismatch of  $\sim 0.16$ . Since MNIST is a very simple task, this might be intuitive. However, when looking at Figure 2 (CIFAR10) - we see that the first layer only reaches 0.6 mismatch. This again, is intuitive as CIFAR10 is a hard task and we would not expect a few layers to be enough to properly cluster the features. This can be shown in several of the graphs along this paper. Perhaps this notion of thinking can aid in defining a concept of “dataset difficulty” for a certain model architecture.
- (ii) **NCC-Collapse may be useful for efficient inference in large models:** In most experiments shown in the paper, the NCC-Collapse does not happen solely in the penultimate model layer. In fact, in some architectures the collapse propagates a few layers back in the network. Suppose that for a trained network, the collapse occurs from all layers after layer  $j$ . This means that in order to get the prediction of the model on a new sample, we only need run a forward pass up to the  $j$ -th layer, and find the nearest train class-means (which needs to be computed once). In very deepnets, this can result in more efficient inference time.

We hope that these points and others shown in this paper can encourage researchers to explore further the geometrical phenomena in intermediate layers.

## 8 CONCLUSION

In this paper, we expand the notion of NCC-Mismatch as proposed in Papyan et al. (2020). We describe how looking at intermediate layers of the network can assist in understanding the geometric phenomena that is Neural Collapse. This paper further expands these notions to NLP tasks, and shows common structure in the different modalities. We also show how encouraging inner-layer class-center consistency can assist in the training and generalization. We hope further research using these methods can continue to enrich the study in deepnets and their training paradigms. We further our discussion into possible usecases of results brought forward in the appendix.

## ACKNOWLEDGEMENTS

We would like to thank Gilad Fuchs and Tomer Galanti for illuminating discussions during the preparation of this manuscript.



## REFERENCES

- Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. 2017.
- Zeyuan Allen-Zhu and Yuanzhi Li. Backward feature correction: How deep learning performs deep learning. *ArXiv*, abs/2001.04413, 2020.
- Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020. URL <http://arxiv.org/abs/1910.06403>.
- Robert J. N. Baldock, Hartmut Maennel, and Behnam Neyshabur. Deep learning through the lens of example difficulty. *ArXiv*, abs/2106.09647, 2021.
- Peter L. Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta Numerica*, 30:87 – 201, 2021.
- Mikhail Belkin. Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation. *Acta Numerica*, 30:203 – 248, 2021.
- Ido Ben-Shaul and Shai Dekel. Sparsity-probe: Analysis tool for deep learning models. *ArXiv*, abs/2105.06849, 2021.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *ArXiv*, abs/2104.14294, 2021.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. *ArXiv*, abs/2002.05709, 2020.
- Francois Chollet. *Deep Learning with Python*. Manning Publications Co., USA, 1st edition, 2017. ISBN 1617294438.
- Uri Cohen, SueYeon Chung, Daniel D. Lee, and Haim Sompolsky. Separability and geometry of object manifolds in deep neural networks. *Nature Communications*, 11:746, February 2020. doi: 10.1038/s41467-020-14578-5.
- Zihang Dai, Hanxiao Liu, Quoc V. Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. *ArXiv*, abs/2106.04803, 2021.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*, abs/1810.04805, 2019.
- Gamaleldin F. Elsayed, Dilip Krishnan, Hossein Mobahi, Kevin Regan, and Samy Bengio. Large margin deep networks for classification. In *NeurIPS*, 2018.
- 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 of the United States of America*, 118, 2021.
- Tomer Galanti, András Gyorgy, and Marcus Hutter. On the role of neural collapse in transfer learning. *ArXiv*, abs/2112.15121, 2021.
- Alperen Gormez and Erdem Koyuncu. Class means as an early exit decision mechanism. *ArXiv*, abs/2103.01148, 2021.
- X. Y. Han, Vardan Papyan, and David L. Donoho. Neural collapse under mse loss: Proximity to and dynamics on the central path. *ArXiv*, abs/2106.02073, 2021.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.

- Gao Huang, Zhuang Liu, and Kilian Q. Weinberger. Densely connected convolutional networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2261–2269, 2017.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *ArXiv*, abs/2004.11362, 2020.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692, 2019.
- Yu-Han Liu and Sercan O. Arik. Explaining deep neural networks using unsupervised clustering. *ArXiv*, abs/2007.07477, 2020.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- Jianfeng Lu and Stefan Steinerberger. Neural collapse with cross-entropy loss. *ArXiv*, abs/2012.08465, 2020.
- Sébastien Marcel and Yann Rodriguez. Torchvision the machine-vision package of torch. In *Proceedings of the 18th ACM International Conference on Multimedia*, MM '10, pp. 1485–1488, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781605589336. doi: 10.1145/1873951.1874254. URL <https://doi.org/10.1145/1873951.1874254>.
- Dustin G. Mixon, Hans Parshall, and Jianzong Pi. Neural collapse with unconstrained features. *ArXiv*, abs/2011.11619, 2020.
- Grégoire Montavon, Mikio L. Braun, and Klaus-Robert Müller. Kernel analysis of deep networks. *J. Mach. Learn. Res.*, 12:2563–2581, 2011.
- Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*, 2021, 2020.
- Varдан 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 of the United States of America*, 117:24652 – 24663, 2020.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- Alethea Power, Yuri Burda, Harrison Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *ArXiv*, abs/2201.02177, 2022.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021.
- Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *ArXiv*, abs/2108.08810, 2021.
- Akshay Rangamani, Mengjia Xu, Andrzej Banburski, Qianli Liao, Tomaso, and Poggio. Dynamics and neural collapse in deep classifiers trained with the square loss. 2021.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.
- Jure Sokolic, Raja Giryes, Guillermo Sapiro, and Miguel R. D. Rodrigues. Robust large margin deep neural networks. *IEEE Transactions on Signal Processing*, 65:4265–4280, 2017.

- Kai Tian, Shuigeng Zhou, and Jihong Guan. Deepcluster: A general clustering framework based on deep learning. In *ECML/PKDD*, 2017.
- Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *ArXiv*, abs/1706.03762, 2017.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *ArXiv*, abs/1804.07461, 2018.
- Lilian Weng. Contrastive representation learning. *lilianweng.github.io/lil-log*, 2021. URL <https://lilianweng.github.io/lil-log/2021/05/31/contrastive-representation-learning.html>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- Junyuan Xie, Ross B. Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. *ArXiv*, abs/1511.06335, 2016.
- Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021.
- Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *ArXiv*, abs/2105.02375, 2021.