Towards an Analytical Definition of Sufficient Data

Adam Byerly
Department of Electronic and Electrical Engineering
Brunel University London
Uxbridge, UB8 3PH UK
Department of Computer Science and Information Systems
Bradley University
Peoria, Il, 61615 USA
abyerly@fsmail.bradley.edu

Tatiana Kalganova
Department of Electronic and Electrical Engineering
Brunel University London
Uxbridge, UB8 3PH UK
tatiana.kalganova@brunel.ac.uk

ABSTRACT

We show that, for each of five datasets of increasing complexity, certain training samples are more informative of class membership than others. These samples can be identified *a priori* to training by analyzing their position in reduced dimensional space relative to the classes' centroids. Specifically, we demonstrate that samples nearer the classes' centroids are less informative than those that are furthest from it. For all five datasets, we show that there is no statistically significant difference between training on the entire training set and when excluding up to 2% of the data nearest to each class's centroid.

1 Introduction

The experimental results of [1] demonstrated that for the micro-PCB dataset they introduced, a portion of the data could be excluded from training and models could achieve comparable accuracy when using data augmentation to simulate the particular attributes of the excluded data, particularly when those models used Homogeneous Vector Capsules [2]. The micro-PCB dataset consists of images coded for rotation and perspective and that allowed for data augmentation techniques to simulate known excluded rotations and perspectives. Since most image classification datasets aren't coded in any way beyond class membership, the question naturally arises: When no coding is available, is there a metric that can be extracted directly from the image data that can be acted upon in a useful way?

Since the primary goal in classification is to distinguish between classes, finding metrics that are indicative of those classes, or more specifically, indicative of those classes relative to the other classes, is a corollary goal. The hypothesis that is tested here is that there exists a distance metric that can be leveraged during the training of a convolutional neural network in order to improve or at least maintain testing accuracy by using that metric to *exclude* a portion of the training data.

An obvious candidate for a distance metric is Euclidian distance from the centroid for the class. An obvious problem with Euclidian distance is the curse of dimensionality that results in near uniformity of distance from the centroid among samples when measuring distance in a large number of dimensions. For example, take a small 3-class dataset, consisting of 10 images in each of the classes cat, dog, and truck derived from a Google Image search for each of the 3 terms "cat", "dog" and "truck" as shown in Figure 1. These images are 180 pixels square and made up of 3

color channels. Thus, the dataset is 97,200-dimensional. Fortunately, there exists a class of non-linear dimensionality reduction techniques that can learn to represent high-dimensional data in lower dimensions. These include t-distributed stochastic neighbor embedding (t-SNE) [3] and Uniform Manifold Approximation and Projection (UMAP) [4]. Due to its relative speed (as compared to t-SNE) and strong theoretical foundations, UMAP is quickly becoming one of the most popular non-linear general dimensionality reduction algorithms in use. Figure 2 shows the result of using UMAP to reduce the dog, cat, truck dataset to 2 and 3 dimensions. In this low-dimensional space the Euclidian distances for each sample from each class's centroid take on meaningful differences among themselves.

Figure 1: A Small 3-Class Dataset Consisting of Cats, Dogs, and Trucks
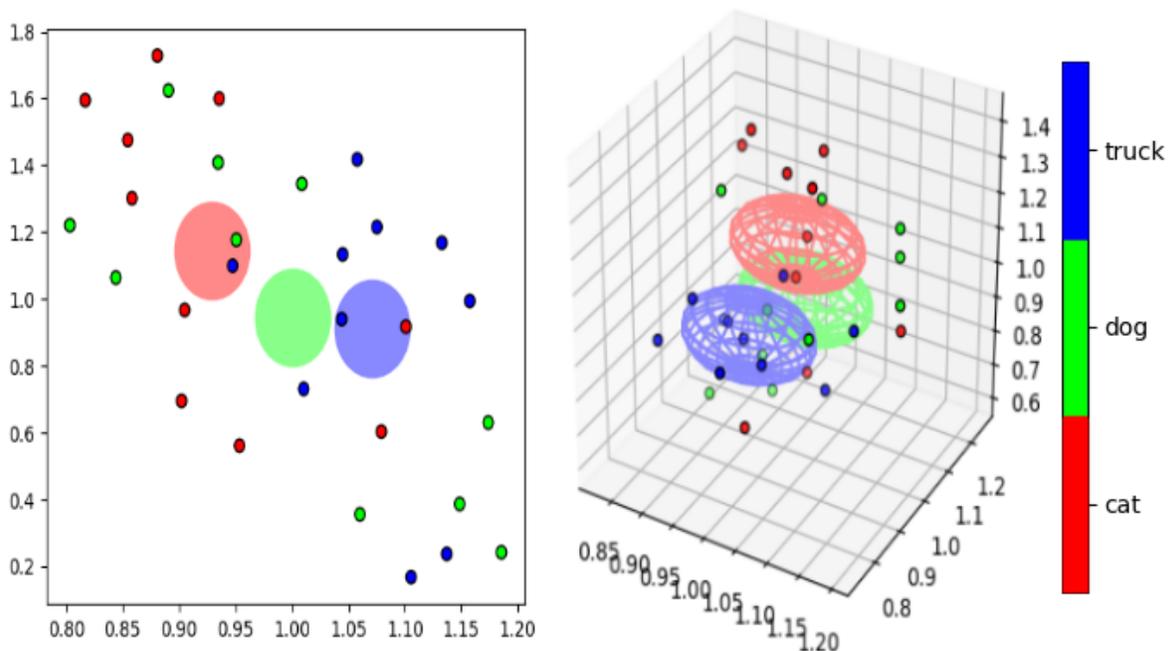

Figure 2: Visualization of 2-Dimensional and 3-Dimensional Reductions of the Cat, Dog, Truck Dataset. The circles (in 2D) and spheres (in 3D) represent the classes' centroids.

Figure 3 shows 2 and 3-dimensional UMAP reductions of the micro-PCB training data. When viewed as a collection of small epsilon balls around each point it is difficult to observe much useful structure. Compare these with Figure 4, which shows a 3-dimensional reduction of the micro-PCB training data produced with UMAP after removing the outliers from each class (those points more than one standard deviation away from their class's centroid) and enclosing the points in a convex hull. Using this visualization method it is possible to observe interesting structure that is not apparent in the visualizations in Figure 3. In the micro-PCB dataset there are three classes of micro-PCBs that are of the same model (Arduino Mega 2560) but produced by different manufacturers. The boards' components and layouts are the same, differing almost entirely by the colors of the substrates and the colors of plastics used for the pins. The

reduction resulted in these classes being tightly clustered together and mostly separate from the remaining classes. A similar tight clustering is exhibited with the two Arduino Unos created by different manufacturers, the two different generations of the Raspberry Pi B+, and the two different Arduino shields.
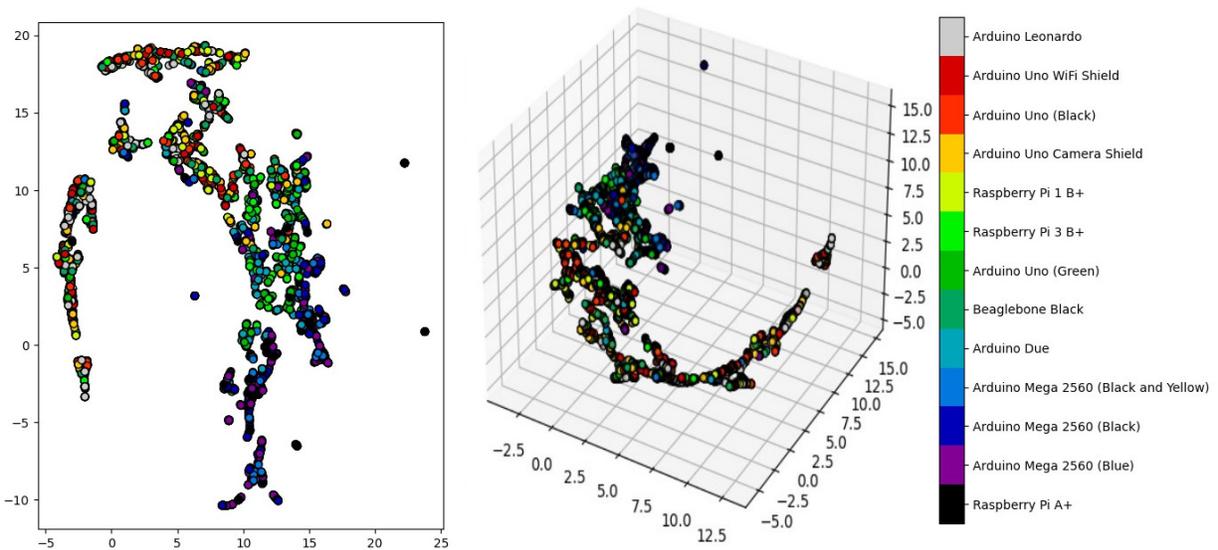

Figure 3: Visualization of 2-Dimensional and 3-Dimensional Reductions of the micro-PCB Dataset
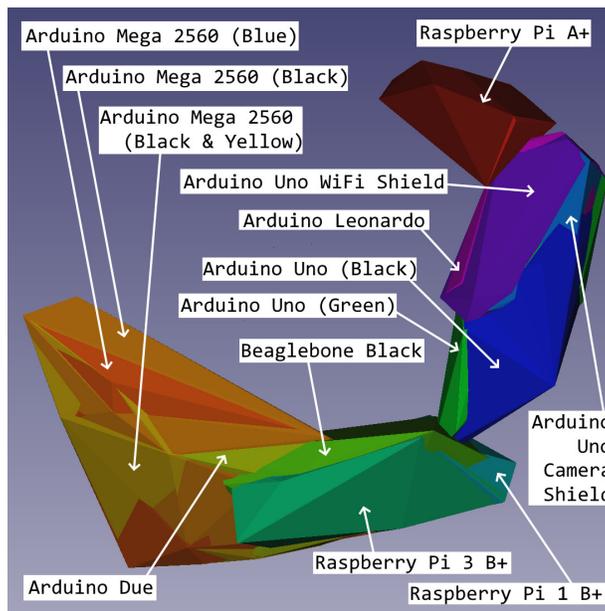

Figure 4: Visualization of 3-Dimensional Reduction of the micro-PCB Dataset (Excluding Outliers)

The ability for this 3-dimensional UMAP reduction to effectively separate the classes such that important semantic information of the micro-PCB dataset emerges in the locations of the samples in the reduced space serves as evidence that a distance metric can indeed be leveraged.

The experiments detailed in this work use the dimensional reduction of each sample as produced by UMAP to determine what samples to exclude during training. These experiments are performed on standard benchmark datasets in order to make the results comparable to other research. It should be noted that with all experiments detailed in this work, all exclusionary methods were applied only to the training set of each dataset and only during training. All testing for all experiments used the entire test set of each dataset.

1.1 Our Contribution

Our contribution is as follows:

1. We demonstrate that training samples in a reduced dimensional space furthest from their class's centroid are more informative than those samples located nearest to the class's centroid, thus providing an *a prior* method for identifying the most informative samples.

2. We demonstrate that reducing the training set size by 2% for all datasets studied produced a statistically insignificant difference in accuracy as compared to when training on all samples. The training set of CIFAR-10 was able to be reduced by 5% and the training set for MNIST was able to be reduced by 10%.

3. We demonstrate that using a 3 dimensional reduction to calculate distances is sufficient relative to 2, 5, and 10 dimensional reductions for all datasets studied.

2 Related Work

k-Nearest Neighbor (kNNs) [5] and other so called "instance-based" methods are also sometimes called "memory-based" because building the model from the training data involves storing the training instances. During evaluation, the sample being evaluated is compared to each stored instance in order to find the best classification for it. Thus, evaluation is an $O(n)$ operation, where n is the number of stored instances. Because the computation cost during evaluation is so high, these algorithms are also sometimes called "lazy learners".

Prior to the relatively recent pivot to large, over-parameterized neural networks for classification, these instance-based methods were some of the most widely-used for classification. As such, a fair amount of research went into reducing the number of training samples that had to be "memorized" both to ease the memory requirements and to speed up evaluation. Early methods for doing this include Condensed Nearest Neighbor (CNN) [6], Selective Nearest Neighbor (SNN) [7], and Edited Nearest Neighbor (ENN) [8]. More recently a slew of other methods have been investigated [9][10][11][12][13].

However, after the success of AlexNet [14], research into classification methods pivoted to large, over-parameterized neural networks. Shortly after that, some research into reducing training dataset size was conducted [15], however, the focus was on acceptable loss of accuracy, rather than maintaining or improving accuracy.

More recently, the focus has shifted to adding training data beyond datasets' canonical training set in order to improve accuracy [16][17][18][19][20][21][22][23][24][25][26][27][28][29][30][31][32].

3 Network Architecture and Training

For all experiments, regardless of the dataset used:

1. The network used for training consisted of a single set of stacked 3×3 convolutions, wherein the first convolutional operation produced 32 feature maps.

2. All subsequent convolutional operations in the network produced an additional 16 feature maps.

3. After all convolutional layers in the network, a set of Unbroken Z-Derived Homogeneous Vector Capsules were used to produce the final classification.

4. Optimization was performed with the Adam optimizer with an initial learning rate of 0.001 that was exponentially decayed every epoch by 0.98.

5. Training proceeded for 300 epochs.

However, because different datasets are formed from images with different sizes, differing number of color channels, and differing complexity of the features present, some slight differences were required depending on the dataset being trained on.

3.1 Training on MNIST and Fashion-MNIST

These datasets [33][34] are composed of single color-channel images comprised of the simplest features relative to the other datasets. They all also use the smallest initial image size of 28×28 pixels. Thus, the network used on these datasets consisted of the fewest convolutional layers (9) and thus the fewest final set of feature maps (160). No padding was used during the convolutional operations, so the final set of feature maps were 10×10. When training this network,

a batch size of 120 was used. Data augmentation used during the training for these datasets was uniform and the same strategy as described in [35].

3.2 Training on CIFAR-10 and CIFAR-100

These datasets [36] are composed of 3 color-channel images comprised of more complex features than MNIST and Fashion-MNIST. Additionally, the images are slightly larger at 32×32 pixels. By using a similar network as was used for MNIST and Fashion-MNIST, but with 2 more convolutional layers and thus 192 feature maps coming out of the final layer, the final set of feature maps were also 10×10. As in the case with MNIST and Fashion-MNIST, when training this network, a batch size of 120 was used. Data augmentation used during the training for these datasets was uniform and the same strategy as used for the experiments in [2].

3.3 Training on Imagenette

This dataset [37] is composed of 3 color-channel images that are substantially larger than CIFAR-10 and CIFAR-100 with more complex features. The raw images vary in size, but were all resized to 299×299. To cope with the larger image size, 15 convolutional layers were used with the first convolutional layer having a stride of 2, and given the addition of 16 feature maps per convolutional operation, this resulted in the presence of 256 feature maps after the final convolution. Max pooling was applied after the fifth and tenth convolutional operations. Using this configuration, the final set of feature maps were 10×10, consistent with all other datasets experimented on. Due to the larger number of parameters required for this network, a batch size of 32 was used, as was dictated by the constraints of available hardware. Data augmentation used during the training for this dataset was the same strategy as used for the experiments in [2].

4 Baseline Results

All subsequent experiments detailed in this chapter involve excluding some subset of the training data during training. In order to understand the impact of those exclusions, a set of baseline experiments was conducted for which all training samples were included for all of the investigated datasets. The results of those experiments are presented in Table 1. For these experiments and all subsequent experiments, five trials of each were conducted.

Table 1: Baseline Results Wherein No Training Samples Were Excluded

Dataset	Accuracy	Standard Deviation
MNIST	99.716%	0.000162481
Fashion-MNIST	93.404%	0.001380724
CIFAR-10	89.146%	0.001518684
CIFAR-100	61.896%	0.001786169
Imagenette	92.390%	0.002333238

5 Data Reduction Strategies

5.1 Experimental Design

For the first set of data reduction experiments detailed in this work, the high-dimensional image data was reduced using UMAP to 3 dimensions. Then, three data reduction strategies for selecting the data to exclude during training were considered. For the first two data reduction strategies, the 3 dimensional centroid for each class was calculated. The *Lateral Exclusion* data reduction strategy excluded samples furthest from the centroid (compare Figure 5 to Figure 6). The *Central Exclusion* data reduction strategy excluded samples nearest to the centroid (compare Figure 5 to Figure 7). The *Random Exclusion* data reduction strategy excluded samples randomly (compare Figure 5 to Figure 8). Then for each data reduction strategy, experiments were conducted excluding 1%, 2%, 5%, 10%, 25%, and 50% of the training data. These experiments were conducted on MNIST, Fashion-MNIST, CIFAR-10, CIFAR-100, and Imagenette.

5.2 Experimental Results

Table 2 shows the results of the experiments for the MNIST dataset. Using the *Random Exclusion* strategy was inferior in all cases. Using the *Lateral Exclusion* strategy produced a mean accuracy of 0.002% greater than when using the
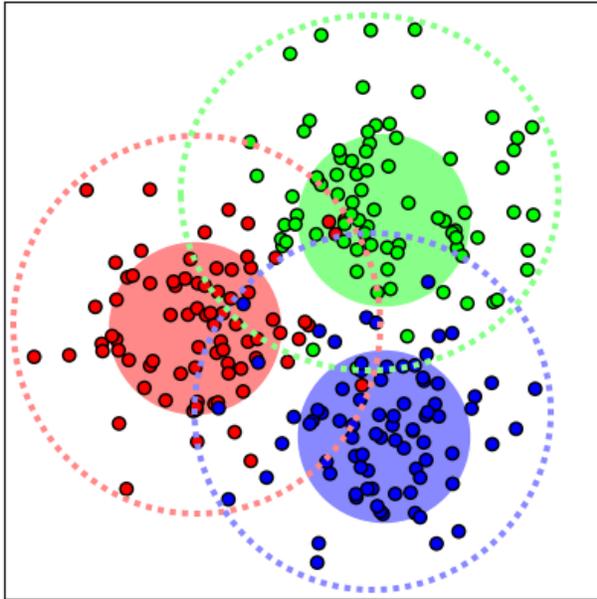
Figure 5: Visualization of 2-Dimensional Data Generated for 3 Relatively Well Separated Classes (All Points)
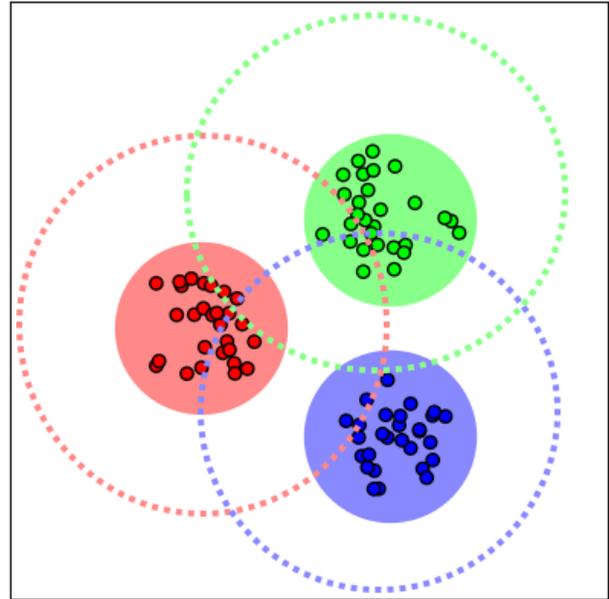

Figure 6: Visualization of 2-Dimensional Data Generated for 3 Relatively Well Separated Classes (*Lateral Exclusion* Visualized)
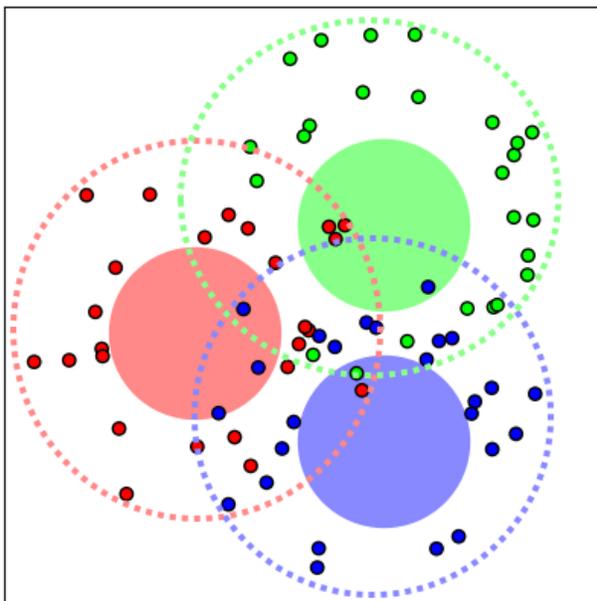

Figure 7: Visualization of 2-Dimensional Data Generated for 3 Relatively Well Separated Classes (*Central Exclusion* Visualized)
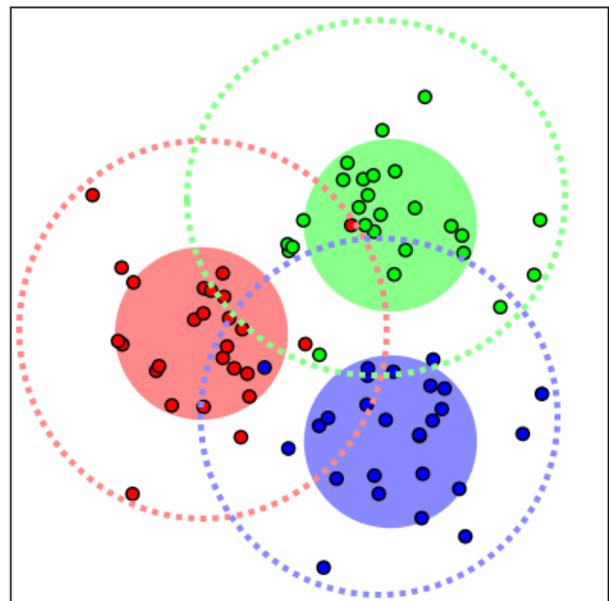
Figure 8: Visualization of 2-Dimensional Data Generated for 3 Relatively Well Separated Classes (*Random Exclusion* Visualized)

Central Exclusion strategy when 2% of the data was excluded, which is not a statistically significant difference. For all other levels of exclusion, using the *Central Exclusion* strategy was statically significantly superior. When excluding 5% of the training data using the *Central Exclusion* strategy, the accuracy achieved was higher than the baseline that included all samples, though not enough trials were conducted to confirm this as statistically significant. When excluding 10%, the accuracy was identical to that of the baseline. When using the *Central Exclusion* strategy, only when excluding 25% or 50% was the accuracy statistically significantly lower than the baseline.

Table 3 shows the results of the experiments for the Fashion-MNIST dataset. In all cases, the *Central Exclusion* strategy proved statistically significantly superior to either of the other two strategies. The *Lateral Exclusion* strategy was superior to the *Random Exclusion* strategy in all but the case of 2% exclusion. However, this superiority was only shown to be statistically significant in the case of the 25% exclusion. There was no statistically significant difference from the baseline when excluding either 1% or 2% of the training data.

The experiments performed on CIFAR-10 (see Table 4) demonstrated the greatest amount of ambiguity among all datasets. All three exclusion strategies outperformed the baseline for exclusions of both 1% and 2%, although not enough trials were conducted to show this to be statistically significant. The *Lateral Exclusion* strategy achieved the highest accuracy when excluding 1%, the *Random Exclusion* strategy achieved the highest accuracy when excluding 10%, and the *Central Exclusion* strategy achieved the highest accuracy in all other cases. The statistically significant differences occurred when excluding 5%, 25%, and 50%. When excluding 5%, the *Central Exclusion* strategy was statistically significantly superior to both of the other two strategies. When excluding 25% and 50%, the *Central Exclusion* strategy was superior to both of the other two strategies, but only statistically significantly superior to the *Lateral Exclusion* strategy.

Table 5 shows the results of the experiments for the CIFAR-100 dataset. In all cases, the accuracy when using the *Central Exclusion* was superior to the other two strategies. The superiority was shown to be statistically significant in the cases of the 10%, 25%, and 50% exclusion experiments. There was no statistically significant difference from the baseline when excluding either 1% or 2% of the training data.

Table 6 shows the results of the experiments for the Imagenette dataset. The *Random Exclusion* strategy achieved a statistically insignificant superior accuracy relative to the other two methods for the 1% and 2% exclusion experiments. For the remainder of the experiments, the *Central Exclusion* strategy achieved a higher accuracy than the other two methods, although this was only statistically significant for 3 out of 4 such experiments (5%, 25%, and 50% exclusion). There was no statistically significant difference from the baseline when excluding 1% the training data.

After examining these five datasets, it can safely be concluded that, in general, the *Central Exclusion* strategy is the superior strategy among the three. In no experiments did either of the other two strategies show a statistically significant superiority to it. Excluding data resulted in equivalent or superior accuracy over the baseline for experiments using 3 of the 5 datasets, including CIFAR-10 when excluding 1%, 2%, or 5%, CIFAR-100 when excluding 1%, and MNIST when excluding 5% or 10%. However, due to the high variance across trials, this was not able to be demonstrated as statistically significant. Similarly, no experiment for any dataset when excluding 1% or 2% was shown to be statistically significantly inferior to the baseline.

Table 2: Data Reduction Strategy Experimental Results — MNIST

Excl. %	Lateral Exclusion		Central Exclusion		Random Exclusion	
	Acc.	Std. Dev.	Acc.	Std. Dev.	Acc.	Std. Dev.
1%	99.702%	1.47×10^{-4}	**99.714%**	1.74×10^{-4}	99.694%	4.90×10^{-5}
2%	**99.714%**	2.42×10^{-4}	99.712%	1.17×10^{-4}	99.696%	2.33×10^{-4}
5%	99.676%	1.62×10^{-4}	**99.720%**	1.67×10^{-4}	99.658%	1.83×10^{-4}
10%	99.658%	1.17×10^{-4}	**99.716%**	1.20×10^{-4}	99.638%	2.32×10^{-4}
25%	99.546%	2.58×10^{-4}	**99.698%**	2.14×10^{-4}	99.526%	2.33×10^{-4}
50%	99.190%	3.03×10^{-4}	**99.686%**	2.65×10^{-4}	99.236%	2.58×10^{-4}

The baseline accuracy, excluding no samples, for this dataset was 99.716%

Table 3: Data Reduction Strategy Experimental Results — Fashion-MNIST

Excl. %	Lateral Exclusion		Central Exclusion		Random Exclusion	
	Acc.	Std. Dev.	Acc.	Std. Dev.	Acc.	Std. Dev.
1%	93.114%	7.89×10^{-4}	**93.346%**	8.71×10^{-4}	93.032%	1.33×10^{-3}
2%	92.896%	6.56×10^{-4}	**93.324%**	1.75×10^{-3}	92.940%	1.66×10^{-3}
5%	92.096%	1.33×10^{-3}	**93.198%**	5.49×10^{-4}	92.084%	7.42×10^{-4}
10%	91.088%	1.34×10^{-3}	**93.184%**	6.47×10^{-4}	91.032%	4.71×10^{-4}
25%	89.096%	9.00×10^{-4}	**92.162%**	9.45×10^{-4}	88.932%	5.38×10^{-4}
50%	85.690%	4.24×10^{-4}	**88.692%**	2.31×10^{-3}	85.558%	1.71×10^{-3}

The baseline accuracy, excluding no samples, for this dataset was 93.404%

Table 4: Data Reduction Strategy Experimental Results — CIFAR-10

Excl. %	Lateral Exclusion		Central Exclusion		Random Exclusion	
	Acc.	Std. Dev.	Acc.	Std. Dev.	Acc.	Std. Dev.
1%	**89.332%**	2.83×10^{-3}	89.280%	1.86×10^{-3}	89.318%	1.18×10^{-3}
2%	89.148%	8.93×10^{-4}	**89.206%**	1.21×10^{-3}	89.148%	1.04×10^{-3}
5%	89.054%	1.45×10^{-3}	**89.224%**	1.30×10^{-3}	89.086%	8.09×10^{-4}
10%	88.708%	6.97×10^{-4}	88.620%	1.87×10^{-3}	**88.838%**	1.94×10^{-3}
25%	87.098%	1.44×10^{-3}	**87.930%**	4.00×10^{-4}	87.878%	4.29×10^{-3}
50%	84.044%	1.32×10^{-3}	**85.900%**	1.13×10^{-3}	85.706%	7.40×10^{-3}

The baseline accuracy, excluding no samples, for this dataset was 89.146%

Table 5: Data Reduction Strategy Experimental Results — CIFAR-100

Excl. %	Lateral Exclusion		Central Exclusion		Random Exclusion	
	Acc.	Std. Dev.	Acc.	Std. Dev.	Acc.	Std. Dev.
1%	61.926%	2.37×10^{-3}	**61.940%**	3.72×10^{-3}	61.888%	4.26×10^{-3}
2%	61.766%	2.63×10^{-3}	**61.866%**	1.90×10^{-3}	61.594%	2.44×10^{-3}
5%	61.262%	3.26×10^{-3}	**61.414%**	2.07×10^{-3}	61.102%	3.81×10^{-3}
10%	60.410%	2.72×10^{-3}	**61.044%**	2.53×10^{-3}	60.012%	2.41×10^{-3}
25%	57.440%	2.25×10^{-3}	**59.158%**	2.01×10^{-3}	57.206%	4.07×10^{-3}
50%	50.754%	2.48×10^{-3}	**53.868%**	4.45×10^{-3}	51.334%	3.94×10^{-3}

The baseline accuracy, excluding no samples, for this dataset was 61.896%

Table 6: Data Reduction Strategy Experimental Results — Imagenette

Excl. %	Lateral Exclusion		Central Exclusion		Random Exclusion	
	Acc.	Std. Dev.	Acc.	Std. Dev.	Acc.	Std. Dev.
1%	92.288%	1.46×10^{-3}	92.300%	3.71×10^{-3}	**92.316%**	2.34×10^{-3}
2%	92.022%	2.14×10^{-3}	92.196%	2.28×10^{-3}	**92.342%**	2.19×10^{-3}
5%	91.896%	6.77×10^{-4}	**92.214%**	1.57×10^{-3}	91.994%	8.36×10^{-4}
10%	91.544%	2.91×10^{-3}	**91.844%**	2.91×10^{-3}	91.680%	1.94×10^{-3}
25%	90.422%	2.20×10^{-3}	**90.998%**	2.35×10^{-3}	90.342%	1.78×10^{-3}
50%	87.998%	3.55×10^{-3}	**88.174%**	2.56×10^{-3}	87.914%	5.66×10^{-3}

The baseline accuracy, excluding no samples, for this dataset was 92.390%

6 Dimensions to Reduce to

6.1 Experimental Design

To test the hypothesis that 3-dimensions was the appropriate choice for the dimensional reduction, an additional set of experiments using the *Central Exclusion* data reduction method, and excluding 1%, 2%, 5%, 10%, 25%, and 50% of the training data were conducted. This set of experiments used 2, 5, and 10 dimensional reductions and were conducted on CIFAR-10, CIFAR-100, MNIST, Fashion-MNIST, and Imagenette.

6.2 Experimental Results

Executing 5 trials of each of 6 different amounts of excluded data results in 30 total trials per dataset and number of dimensions being used to determine the exclusion. For each of the 2, 5, and 10 dimensional reductions we compared the mean accuracy achieved across all 30 trials to the 30 trials of the experiments that used a 3 dimensional reduction. Table 7 shows the results of those comparisons. The comparisons showed no statistically significant difference between any paired sets of 30 trials. 5 out of 15 experiments showed a statistically insignificant superiority when using a 3 dimensional reduction, including all 3 comparisons of MNIST and 2 out of 3 comparisons of Imagenette. Although, not reaching a reasonable threshold for statistical significance ($p < 0.05$), the MNIST comparisons had the lowest p-values. At first, this may seem surprising, but upon reflection, it seems reasonable that points UMAP places far from a class's centroid in 3 dimensions would also be likely to be placed far from the class's centroid in 2, 5, or 10 dimensions as well.

Table 7: Comparison of Mean Accuracies for Exclusions Based on Differing Dimensional Reductions

Dataset	Dimensions	Accuracy	3D Accuracy	p-value
MNIST	2	99.7000%	**99.7077%**	0.111615712
MNIST	5	99.7027%	**99.7077%**	0.232822313
MNIST	10	99.6977%	**99.7077%**	0.083150845
Fashion-MNIST	2	**92.3803%**	92.3177%	0.44224238
Fashion-MNIST	5	**92.4217%**	92.3177%	0.405938742
Fashion-MNIST	10	**92.3490%**	92.3177%	0.471922235
CIFAR-10	2	**88.3677%**	88.3600%	0.490392657
CIFAR-10	5	**88.3710%**	88.3600%	0.486238005
CIFAR-10	10	**88.3747%**	88.3600%	0.481346744
CIFAR-100	2	**59.9347%**	59.8817%	0.471617385
CIFAR-100	5	**60.0760%**	59.8817%	0.396706608
CIFAR-100	10	**59.9790%**	59.8817%	0.448318135
Imagenette	2	**91.3237%**	91.2877%	0.463151304
Imagenette	5	91.2153%	**91.2877%**	0.427306126
Imagenette	10	91.2417%	**91.2877%**	0.453380991

7 Distributions of the Dimensional Reductions

UMAP generates numeric values for each dimension of the reduction performed. Visualizations of 2-dimensional and 3-dimensional reductions are usually generated by drawing these points as small epsilon balls around the positions of those numeric values. While these visualizations can provide some sense of where the classes' data are located in space, drawing convex hulls that surround each class's points provides a sharper distinction between the boundaries of the classes in the space. Figure 9 and Figure 10 show the result of drawing these hulls around the 3-dimensional reductions of the full MNIST and Imagenette training data, respectively. These visualizations make it clear that when including all of the training data, there is very little distinction of boundaries in the space.

However, when excluding some of the data furthest from each class's centroid, the classes' hulls start to separate, in some cases, partially and in some cases entirely. The amount of data that must be elided to achieve this is referred to in this work as *Dataset Severability*.

Definition 1 (Dataset Severability). *Dataset Severability* is a qualitative judgment regarding the number of outliers that must be removed from each class in an m-dimensional reduction of the dataset so that structure and/or clustering is able to be observed.

Figure 11 and Figure 12 show well severed classes for visualizations of the MNIST and EMNIST-Digits training data, respectively. In each case, the data within three standard deviations of all dimensions was included, and data outside of these bounds was omitted. This means that high severability is achieved by excluding $\approx$ 0.27% of the outliers for these datasets. Despite consisting of images with the same size and number of color channels (monochromatic) as MNIST and EMNIST-Digits, Fashion-MNIST shows a dissimilar level of severability until all but one standard deviation from the centroid has been excluded (excluding $\approx$ 31.73%).

Especially interesting in the case of the Fashion-MNIST reduction is that there is a readily identifiable semantic difference between each of the four completely severed sets of overlapping hulls. The set of "overlapping" hulls consisting of the single class "trouser" contains the only class in the dataset that is legwear. The set of "overlapping" hulls consisting of the single class "bag" contains the only class in the dataset that is not a type of clothing that covers any part of the body. The set of overlapping hulls consisting of the classes "sneaker", "sandal", and "ankle boot" contains the only classes in the dataset that are footwear. Finally, the set of overlapping hulls consisting of the classes "pullover", "coat", "shirt", "dress", and "t-shirt" contains the only classes in the dataset that primarily cover the torso.

Imagenette (Figure 14), CIFAR-10 (Figure 15), and CIFAR-100 (Figure 16) require all but those samples within half a standard deviation of each class's centroid to be elided before structure emerges (excluding $\approx$ 61.71%). Of these three, the reduction of CIFAR-10 displays the most interesting semantic relationships among the classes' hulls locations in space. The classes on the left of the visualization in Figure 15 are all man-made (specifically vehicles) whereas the classes on the right are all lifeforms. Especially interesting among the lifeforms is that all of the mammals are grouped together with the one amphibian (frog) on the outer edge of the group.

In Figure 17 through Figure 36 the values of each of the 3 dimensions of the reductions for each individual class of the training data for MNIST and Imagenette are plotted. The first thing that can be learned from looking at the MNIST plots is that the reduction for each class produces values significantly different than the others and further, the values in each dimension of each class are tightly grouped in the number line, with the noticeable exception of the third dimension of the class that represents the digit 1. It is worth noting that the stylization of the Hindu-Arabic numeral '1' contains most of its information in 2-dimensions (accounting for translation) thus providing an explanation for the larger variance in the third dimension. The plots of the Imagenette classes, on the other hand, are much more similar to one another and display greater variance in each of the 3 dimensions.

8 Summary

The experiments we performed show that, for the datasets examined, certain training samples are more informative of class membership than others. These samples can be identified *a priori* to training by analyzing their position in reduced dimensional space relative to the classes' centroids. Specifically, we demonstrated that samples nearer the classes' centroids are less informative than those that are furthest from it. For the five datasets investigated, we have shown that there was no statistically significant difference from the baseline when excluding up to 2% of the data nearest to each class's centroid. For CIFAR-100, superior accuracy was achieved when excluding 1% of the data nearest to each class's centroid. For CIFAR-10, superior accuracy was achieved when excluding 5% of the data nearest to each class's centroid. And for the MNIST dataset, identical accuracy to the baseline was achieved when excluding 10% of the data nearest to each class's centroid.

Additionally, we defined *Dataset Severability* to be a qualitative, yet quantifiable, judgement regarding the separation of classes in a reduced dimensional space. High severability was shown for MNIST and Fashion-MNIST, whereas low severability was shown for CIFAR-10, CIFAR-100, and Imagenette. Those datasets that demonstrated high severability all achieved higher accuracies in the experiments detailed in this work compared to the accuracies of those experiments for the datasets that demonstrated low severability.
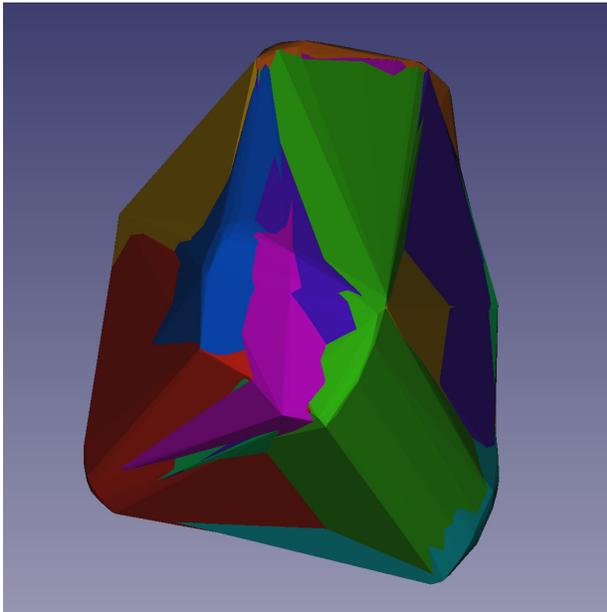
Figure 9: Visualization of 3-Dimensional Reduction of MNIST
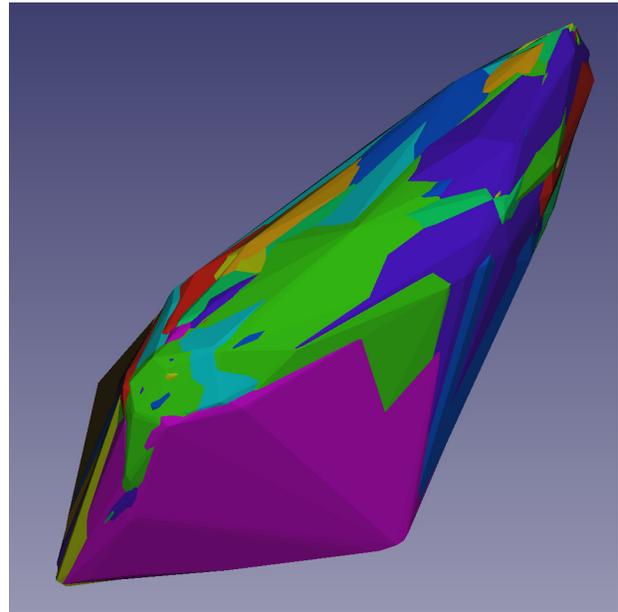
Figure 10: Visualization of 3-Dimensional Reduction of Imagenette
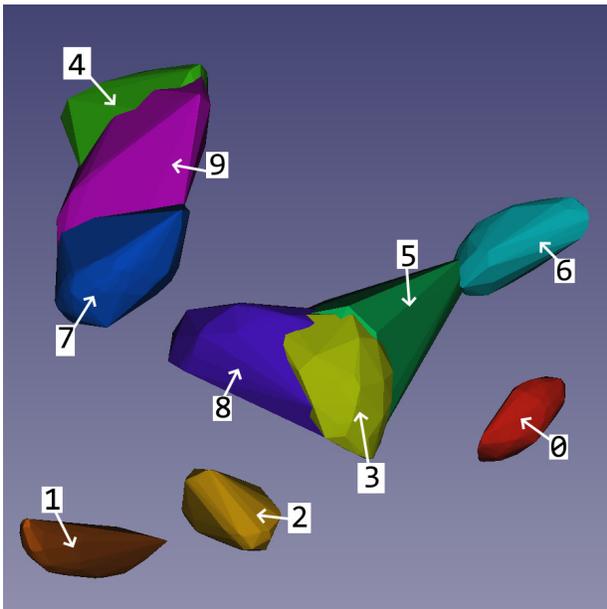
Figure 11: Visualization of 3-Dimensional Reduction of MNIST w/ Classes Labeled (Excluding Outliers)
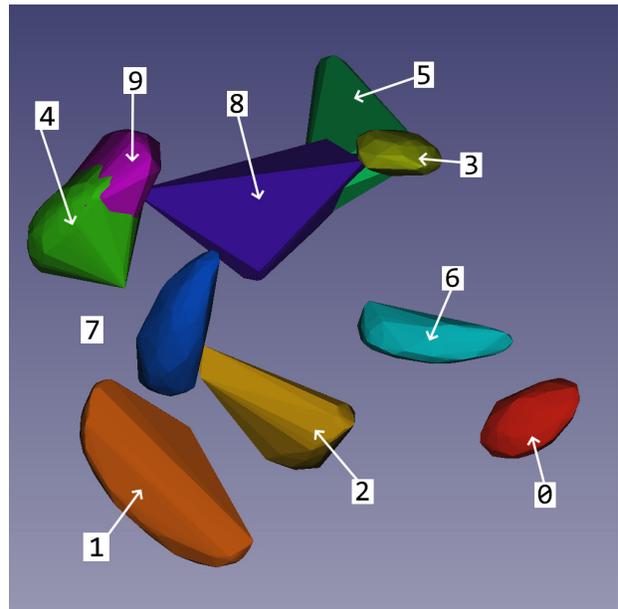
Figure 12: Visualization of 3-Dimensional Reduction of EMNIST-Digits w/ Classes Labeled (Excluding Outliers)
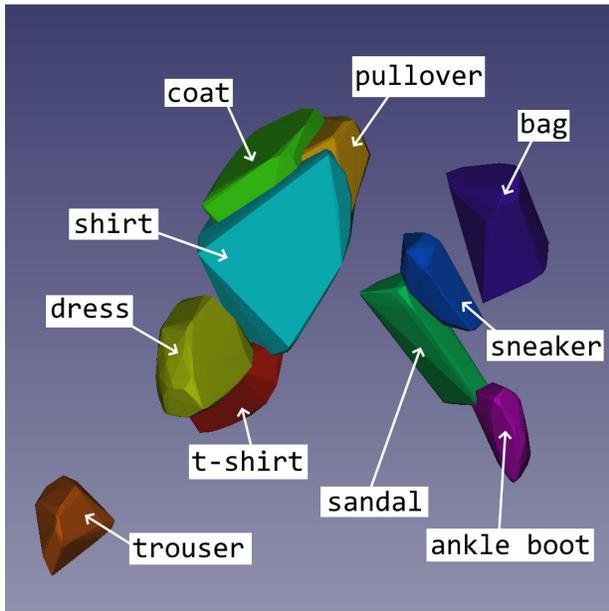
Figure 13: Visualization of 3-Dimensional Reduction of Fashion-MNIST w/ Classes Labeled (Excluding Outliers)
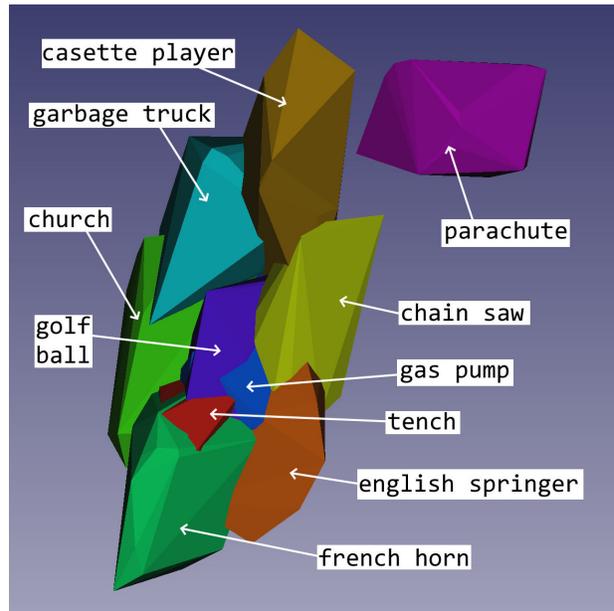
Figure 14: Visualization of 3-Dimensional Reduction of Imagenette w/ Classes Labeled (Excluding Outliers)
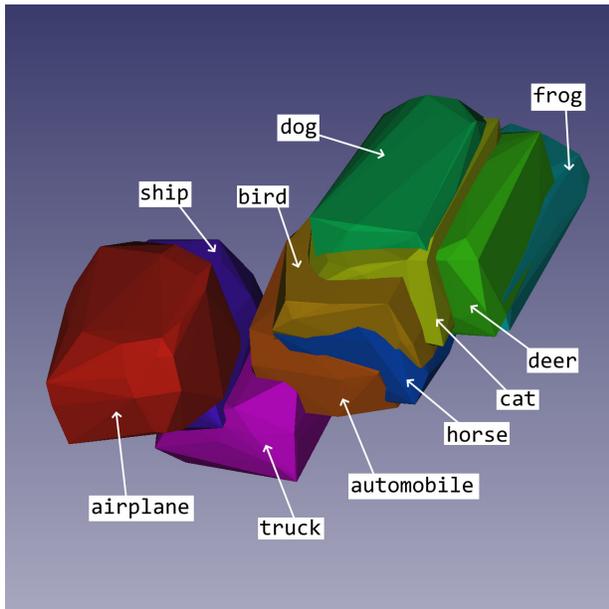
Figure 15: Visualization of 3-Dimensional Reduction of CIFAR-10 w/ Classes Labeled (Excluding Outliers)
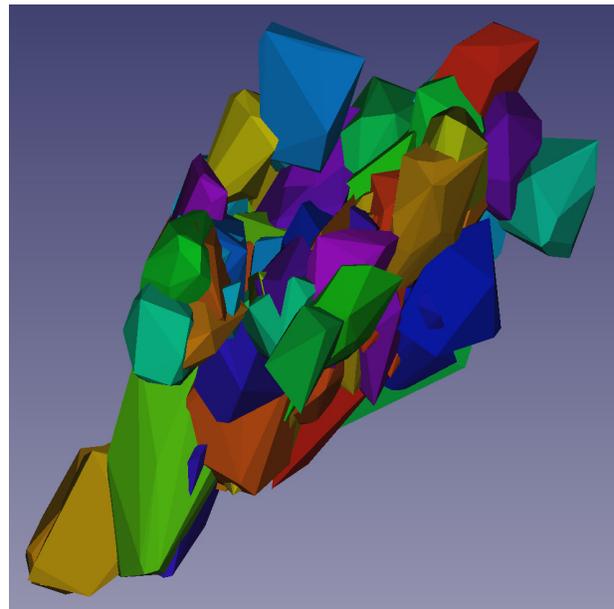
Figure 16: Visualization of 3-Dimensional Reduction of CIFAR-100 w/ Classes Labeled (Excluding Outliers)
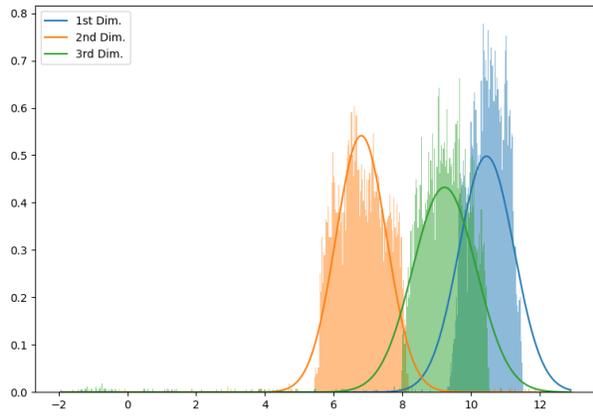
Figure 17: MNIST Reduction Distributions — Class '0'
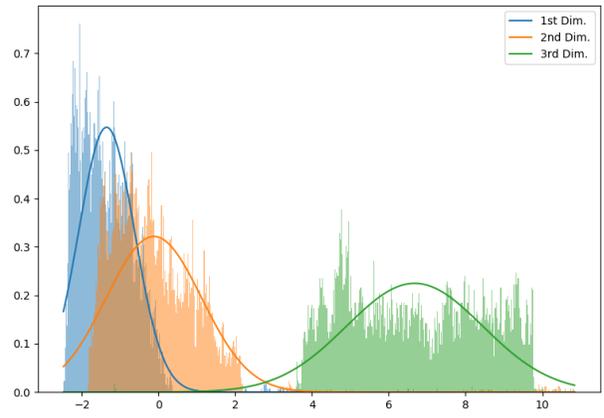
Figure 18: MNIST Reduction Distributions — Class '1'
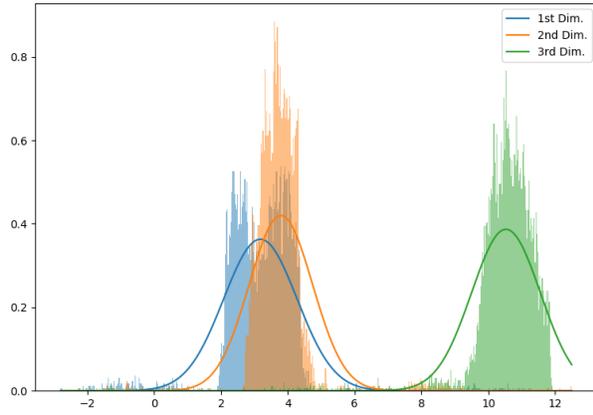
Figure 19: MNIST Reduction Distributions — Class '2'
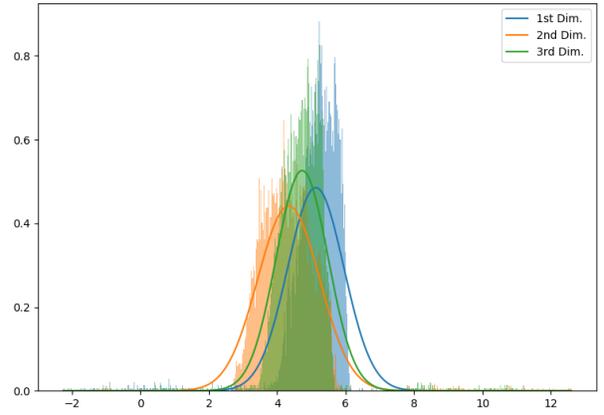
Figure 20: MNIST Reduction Distributions — Class '3'
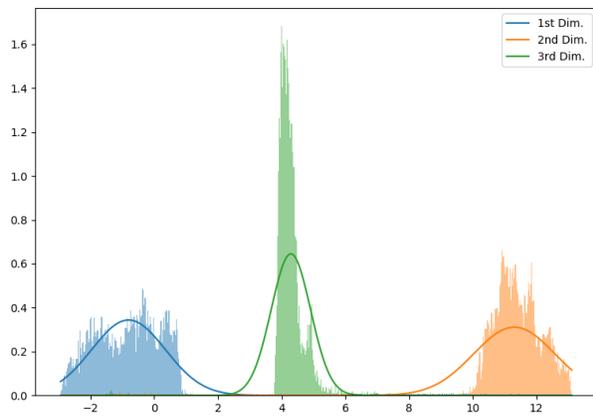

Figure 21: MNIST Reduction Distributions — Class '4'
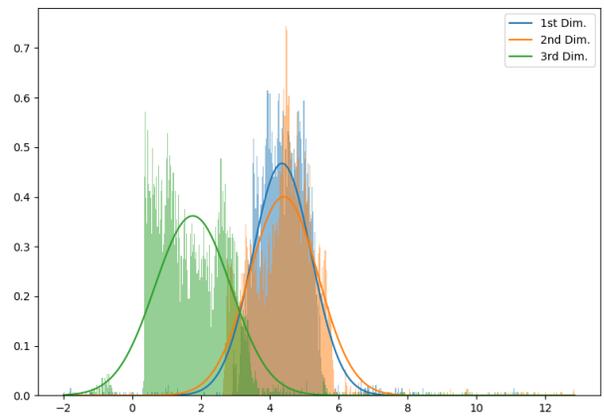
Figure 22: MNIST Reduction Distributions — Class '5'
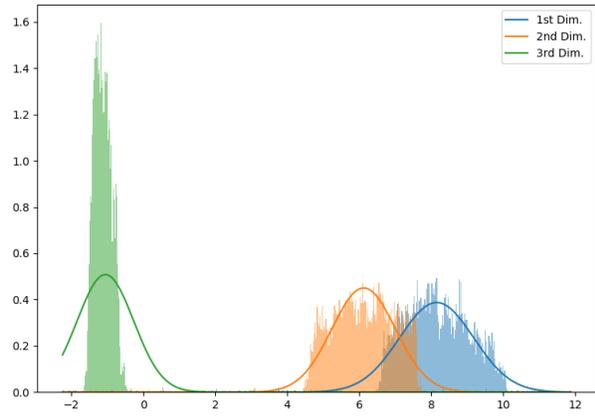
Figure 23: MNIST Reduction Distributions — Class '6'
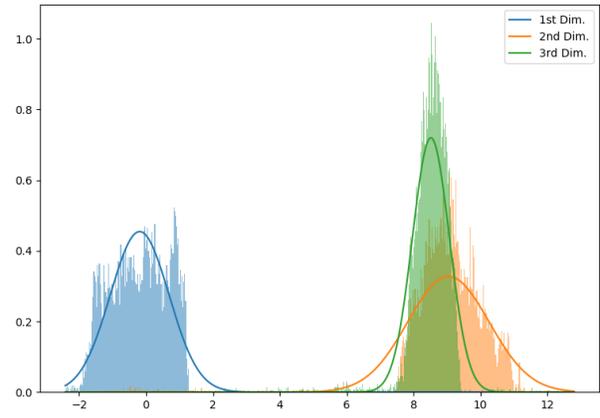
Figure 24: MNIST Reduction Distributions — Class '7'
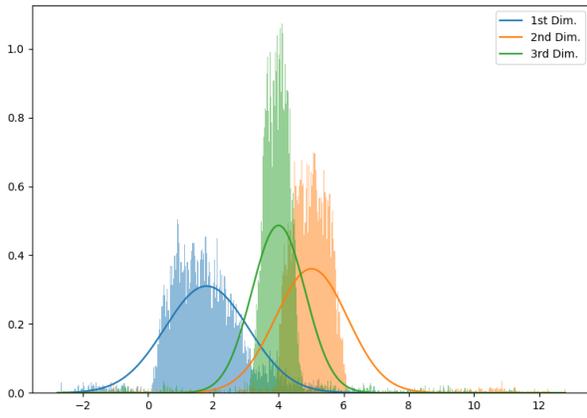
Figure 25: MNIST Reduction Distributions — Class '8'
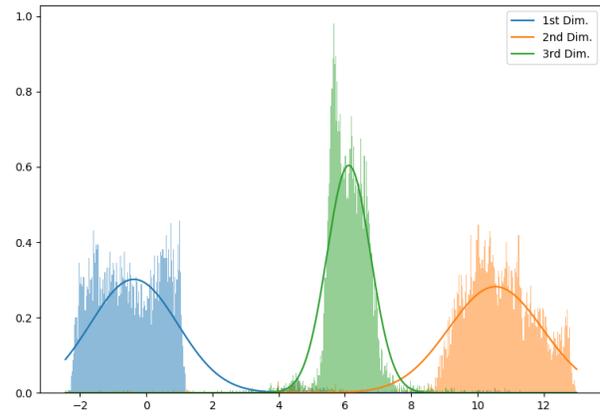
Figure 26: MNIST Reduction Distributions — Class '9'
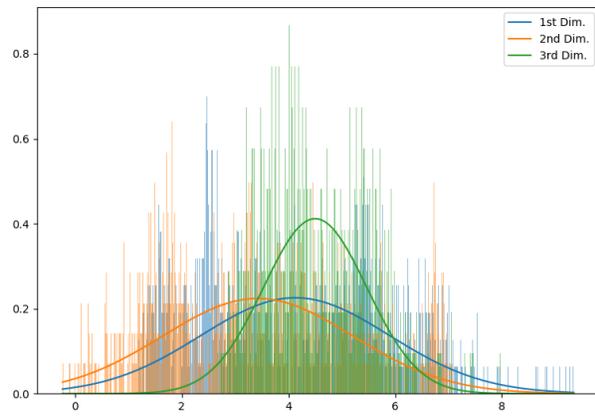
Figure 27: Imagenette Reduction Distributions — Class 'Tench'
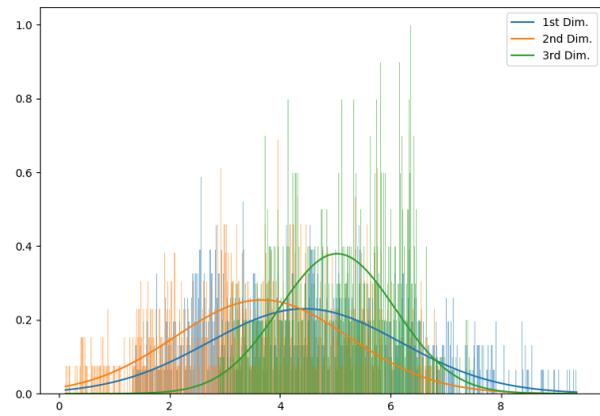
Figure 28: Imagenette Reduction Distributions — Class 'English Springer'
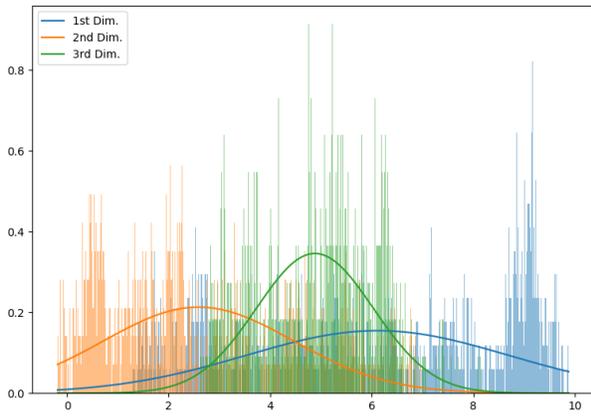
Figure 29: Imagenette Reduction Distributions — Class 'Cassette Player'
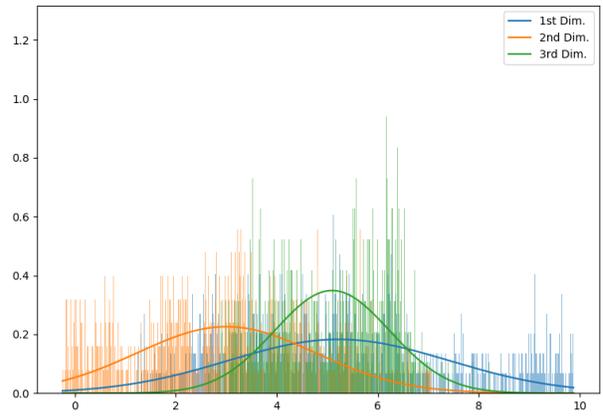
Figure 30: Imagenette Reduction Distributions — Class 'Chain Saw'
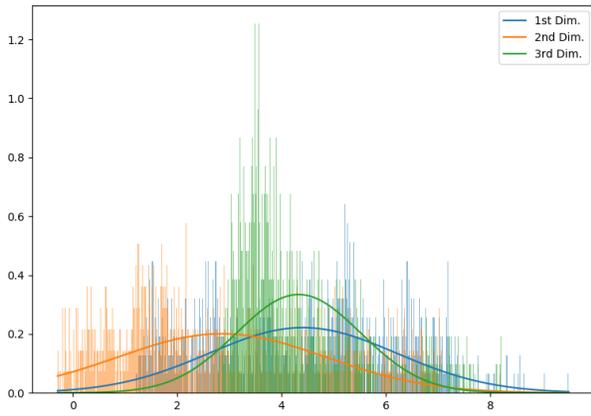

Figure 31: Imagenette Reduction Distributions — Class 'Church'
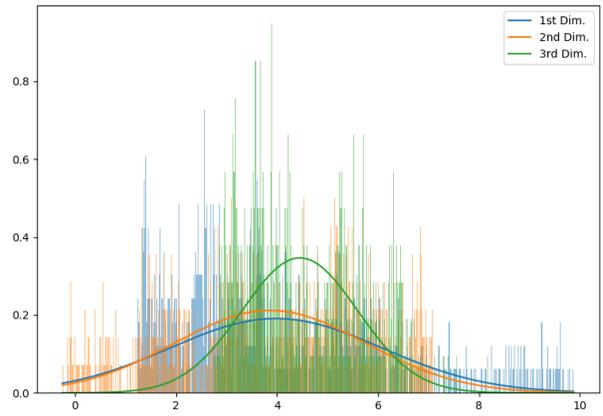
Figure 32: Imagenette Reduction Distributions — Class 'French Horn'
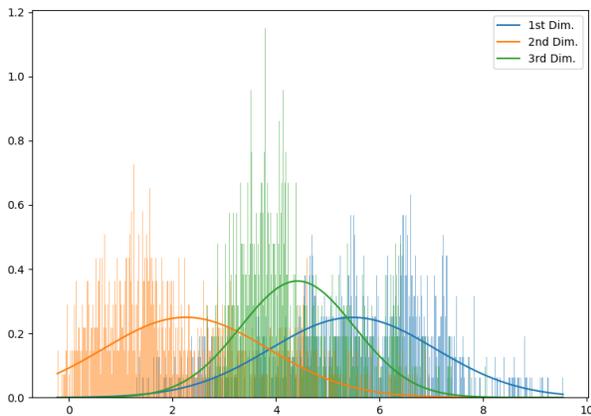
Figure 33: Imagenette Reduction Distributions — Class 'Garbage Truck'
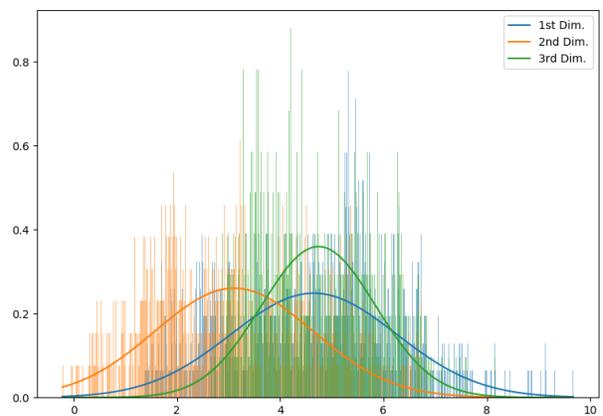

Figure 34: Imagenette Reduction Distributions — Class 'Gas Pump'
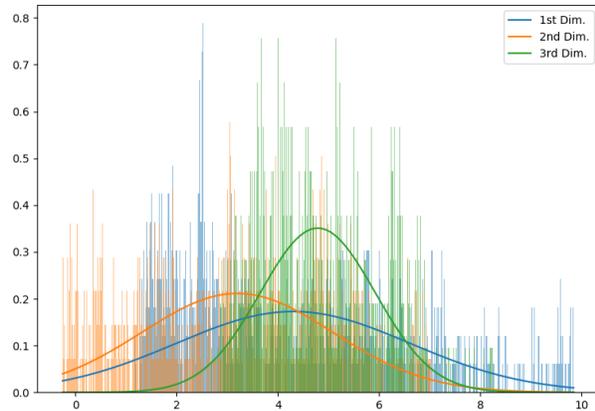
Figure 35: Imagenette Reduction Distributions — Class 'Golf Ball'
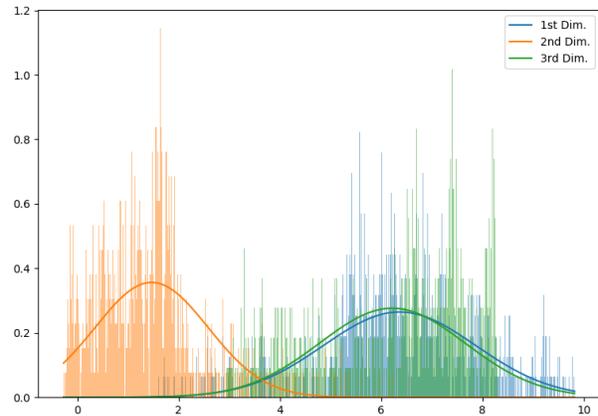
Figure 36: Imagenette Reduction Distributions — Class 'Parachute'

References

[1] Adam Byerly, Tatiana Kalganova, and Anthony J. Grichnik. "On the Importance of Capturing a Sufficient Diversity of Perspective for the Classification of Micro-PCBs". In: *Intelligent Decision Technologies*. Vol. 238. Springer Singapore, 2021, pp. 209–219. ISBN: 978-981-16-2765-1.

[2] Adam Byerly and Tatiana Kalganova. "Homogeneous Vector Capsules Enable Adaptive Gradient Descent in Convolutional Neural Networks". In: *IEEE Access* 9 (2021), pp. 48519–48530. DOI: doi:10.1109/ACCESS.2021.3066842.

[3] Laurens van der Maaten and Geoffrey Hinton. "Visualizing Data using t-SNE Laurens". In: *Journal of Machine Learning Research* 9 (2008), pp. 2579–2605.

[4] Leland McInnes, John Healy, and James Melville. "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction". In: (2018).

[5] T. Cover and P. Hart. "Nearest neighbor pattern classification". In: *IEEE Transactions on Information Theory* 13.1 (1967), pp. 21–27. DOI: 10.1109/TIT.1967.1053964.

[6] P. Hart. "The condensed nearest neighbor rule (Corresp.)" In: *IEEE Transactions on Information Theory* 14.3 (1968), pp. 515–516. DOI: doi:10.1109/TIT.1968.1054155.

[7] G. Ritter et al. "An algorithm for a selective nearest neighbor decision rule (Corresp.)" In: *IEEE Transactions on Information Theory* 21.6 (1975), pp. 665–669. DOI: doi:10.1109/TIT.1975.1055464.

[8] Dennis L. Wilson. "Asymptotic Properties of Nearest Neighbor Rules Using Edited Data". In: *IEEE Transactions on Systems, Man, and Cybernetics* SMC-2.3 (1972), pp. 408–421. DOI: doi:10.1109/TSMC.1972.4309137.

[9] D Randall Wilson and Tony R Martinez. "Reduction Techniques for Instance-Based Learning Algorithms". In: *Machine Learning* 38 (2000), pp. 257–286.

[10] María Teresa Lozano Albalate. "Data reduction techniques in classification processes". PhD thesis. 2007.

[11] Fernando Vázquez, J. Salvador Sánchez, and Filiberto Pla. "A Stochastic Approach to Wilson's Editing Algorithm". In: *Pattern Recognition and Image Analysis*. 2005, pp. 35–42. ISBN: 978-3-540-32238-2.

[12] Chien-Hsing Chou, Bo-Han Kuo, and Fu Chang. "The Generalized Condensed Nearest Neighbor Rule as A Data Reduction Method". In: *18th International Conference on Pattern Recognition (ICPR'06)*. Vol. 2. 2006, pp. 556–559. DOI: doi:10.1109/ICPR.2006.1119.

[13] Stefanos Ougiaroglou and Georgios Evangelidis. "Efficient dataset size reduction by finding homogeneous clusters". In: *Balkan Conference in Informatics (BCI)*. 2012, pp. 168–173. ISBN: 9781450312400. DOI: doi:10.1145/2371316.2371349.

[14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: *NIPS 2012 - 25th Conference on Neural Information Processing Systems*. 2012, pp. 1097–1105. DOI: doi:10.1145/3065386.

[15] Mohammad Amin Shayegan and Saeed Aghabozorgi. "A new dataset size reduction approach for PCA-based classification in OCR application". In: *Mathematical Problems in Engineering* (2014). ISSN: 15635147. DOI: doi:10.1155/2014/537428.

[16]   Alexey Dosovitskiy et al. "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale". In: *Ninth International Conference on Learning Representations (ICLR)*. 2020.

[17]   Alexander Kolesnikov et al. "Big Transfer (BiT): General Visual Representation Learning". In: *16th European Conference on Computer Vision*. 2020.

[18]   Hugo Touvron et al. "Fixing the train-test resolution discrepancy: FixEfficientNet". In: *Advances in Neural Information Processing Systems* 32 (2019).

[19]   Hieu Pham et al. "Meta Pseudo Labels". In: (2020). URL: http://arxiv.org/abs/2003.10580.

[20]   Qizhe Xie et al. "Self-training with noisy student improves imagenet classification". In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 2020, pp. 10684–10695. DOI: doi:10.1109/CVPR42600.2020.01070.

[21]   Pierre Foret et al. "Sharpness-Aware Minimization for Efficiently Improving Generalization". In: *Ninth International Conference on Learning Representations (ICLR)*. 2020.

[22]   Xiaohua Zhai et al. "Scaling Vision Transformers". In: (2021). URL: http://arxiv.org/abs/2106.04560.

[23]   Carlos Riquelme et al. "Scaling Vision with Sparse Mixture of Experts". In: (2021). URL: http://arxiv.org/abs/2106.05974.

[24]   Michael S. Ryoo et al. "TokenLearner: What Can 8 Learned Tokens Do for Images and Videos?" In: (2021). URL: http://arxiv.org/abs/2106.11297.

[25]   Chao Jia et al. "Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision". In: *Proceedings of the 38th International Conference on Machine Learning (PMLR)*. 2021.

[26]   Xiaoyi Dong et al. "CSWin Transformer: A General Vision Transformer Backbone with Cross-Shaped Windows". In: (2021). URL: http://arxiv.org/abs/2107.00652.

[27]   Ze Liu et al. "Swin Transformer: Hierarchical Vision Transformer using Shifted Windows". In: *The International Conference on Computer Vision (ICCV)*. 2021.

[28]   Zihang Dai et al. "CoAtNet: Marrying Convolution and Attention for All Data Sizes". In: (2021). URL: http://arxiv.org/abs/2106.04803.

[29]   Haiping Wu et al. "CvT: Introducing Convolutions to Vision Transformers". In: *The International Conference on Computer Vision (ICCV)*. 2021.

[30]   Mingxing Tan and Quoc V. Le. "EfficientNetV2: Smaller Models and Faster Training". In: (2021). URL: http://arxiv.org/abs/2104.00298.

[31]   Ilya Tolstikhin et al. "MLP-Mixer: An all-MLP Architecture for Vision". In: (2021). URL: http://arxiv.org/abs/2105.01601.

[32]   Andrew Brock et al. "High-Performance Large-Scale Image Recognition Without Normalization". In: (2021). URL: http://arxiv.org/abs/2102.06171.

[33]   Yann LeCun, Corinna Cortes, and CJ Burges. "MNIST handwritten digit database". In: *ATT Labs [Online]*. 2 (2010). URL: http://yann.lecun.com/exdb/mnist.

[34]   Han Xiao, Kashif Rasul, and Roland Vollgraf. *Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms*. 2017. arXiv: 1708.07747 [cs.LG].

[35]   Adam Byerly, Tatiana Kalganova, and Ian Dear. "No Routing Needed Between Capsules". In: *Neurocomputing* 463 (2021), pp. 545–553. ISSN: 0925-2312. DOI: https://doi.org/10.1016/j.neucom.2021.08.064.

[36]   Alex Krizhevsky. *Learning Multiple Layers of Features from Tiny Images*. Tech. rep. 2009.

[37]   Jeremy Howard. *Imagenette*. 2018. URL: https://github.com/fastai/imagenette/ (visited on 01/21/2022).