Semantic Redundancies in Image-Classification Datasets: The 10% You Don't Need

Vighnesh Birodkar* Hossein Mobahi Samy Bengio
Google Research, Mountain View, CA
vighneshb@google.com hmobahi@google.com bengio@google.com

arXiv:1901.11409v1 [cs.CV] 29 Jan 2019

Abstract

Large datasets have been crucial to the success of deep learning models in the recent years, which keep performing better as they are trained with more labelled data. While there have been sustained efforts to make these models more data-efficient, the potential benefit of understanding the data itself, is largely untapped. Specifically, focusing on object recognition tasks, we wonder if for common benchmark datasets we can do better than random subsets of the data and find a subset that can generalize on par with the full dataset when trained on. To our knowledge, this is the *first* result that can find notable redundancies in CIFAR-10 and ImageNet datasets (at least 10%). Interestingly, we observe semantic correlations between required and redundant images. We hope that our findings can motivate further research into identifying additional redundancies and exploiting them for more efficient training or data-collection.

1 Introduction

Large datasets have played a central role in the recent success of deep learning. In fact, the performance of AlexNet [Krizhevsky et al., 2012] trained on ImageNet [Deng et al., 2009] in 2012 is often considered as the starting point of the current deep learning era. Undoubtedly, prominent datasets of ImageNet, CIFAR, and CIFAR-100 [Krizhevsky and Hinton, 2009] have had a crucial role in the evolution of deep learning methods since then; with even bigger datasets like OpenImages [Kuznetsova et al., 2018] and Tencent ML-images [Wu et al., 2019] recently emerging. These developments have led to state-of-the-art architectures such as ResNets [He et al., 2016a], DenseNets [Huang et al., 2017], VGG [Simonyan and Zisserman, 2014], AmoebaNets [Huang et al., 2018], and regularization techniques such as Dropout [Srivastava et al., 2014] and Shake-Shake [Gastaldi, 2017]. However, understanding the properties of these datasets themselves has remained relatively untapped. Limited study along this direction includes [Lin et al., 2018], which proposes a modified loss function to deal with the class imbalance inherent in object detection datasets and [Tobin et al., 2017], which studies modifications to simulated data to help models adapt to the real world, and [Carlini et al., 2018] that demonstrates the existence of prototypical examples and verifies that they match human intuition.

This work studies the properties of ImageNet, CIFAR-10 , and CIFAR-100 datasets from the angle of redundancy. We find that at least 10% of ImageNet and CIFAR-10 can be safely removed by a technique as simple as clustering. Particularly, we identify a certain subset of ImageNet and CIFAR-10 whose removal does not affect the test accuracy when the architecture is trained from scratch on the remaining subset. This is striking, as deep learning techniques are believed to be data hungry [Halevy et al., 2009, Sun et al., 2017]. In fact, recently the work by [Vodrahalli et al., 2018] specifically studying the redundancy of these datasets concludes that there is no redundancy. Our work refutes that claim by providing counter examples.

Contributions. This work resolves some recent misconceptions about the absence of *notable* redundancy in major image classification datasets [Vodrahalli et al., 2018]. We do this by identifying a specific subset, which constitutes above 10% of the training set, and yet its removal causes no drop in the test accuracy. To our knowledge, this is the first time such significant redundancy is shown to exist for these datasets. We emphasize that our contribution is merely to demonstrate the *existence* of such redundancy, but we do not claim any *algorithmic* contributions. However, we hope that our findings can motivate further research into identifying additional redundancies and exploiting them for more efficient training or data-collection. Our findings may also be of interest to active learning community, as it provides an upper-bound on the best performance[1].

2 Related Works

There are approaches which try to prioritize different examples to train on as the learning process goes on such as [Fan et al., 2016] and [Katharopoulos and Fleuret, 2018]. Although these

*Work Done as Google AI Resident.

[1]Suppose we learn about existence of m samples in a dataset of size $n > m$ that can achieve the same test performance as a model trained with all n samples. Then if our active learner cannot reach the full test performance after selecting m samples, we know that there *might* exist a better active learning algorithm, as the ideal subset of size m can achieve full test accuracy.

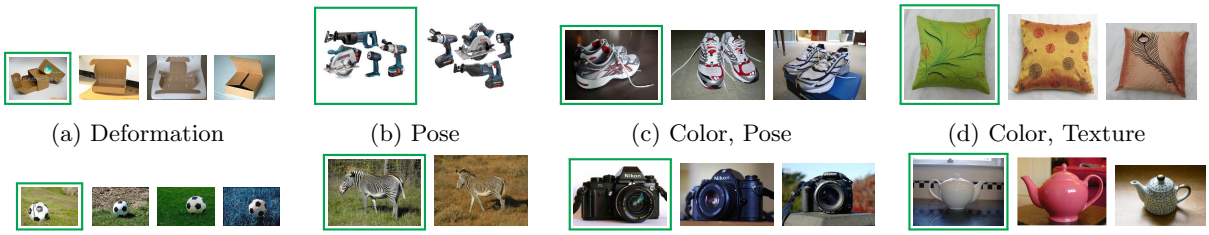

(a) Deformation (b) Pose (c) Color, Pose (d) Color, Texture

(e) Background (f) Background (g) Pose (h) Color, Texture, Pose

Figure 1: Examples of different redundant groups of images from the ImageNet dataset while creating a subset 90% of the size of the full set. In each group, we list the semantic variation considered redundant. The images selected by semantic clustering are highlighted with a green box whereas the rest are discarded with no negative impact on generalization.

techniques involve selecting examples to train on, they do not seek to identify redundant subsets of the data, but rather to sample the full dataset in a way that speeds up convergence.

An early mention of trying to reduce the training dataset size can be seen in [Ohno-Machado et al., 1998]. Their proposed algorithm splits the training dataset into many smaller training sets and iteratively removes these smaller sets until the generalization performance falls below an acceptable threshold. However, the algorithm relies on creating many small sets out of the given training set, rendering it impractical for modern usage.

[Wei et al., 2015] pose the problem of subset selection as a constrained sub-modular maximization problem and use it to propose an active learning algorithm. The proposed techniques are used by [Kaushal et al., 2018] in the context of image recognition tasks. These drawback however, is that when used with deep-neural networks, simple uncertainty based strategies out-perform the mentioned algorithm.

Another example of trying to identify a smaller, more informative set can be seen in [Lapedriza et al., 2013]. Using their own definition of value of a training example, they demonstrate that prioritizing training over examples of high training value can result in improved performance for object detection tasks. The authors suggest that their definition of training value encourages prototypicality and thus results is better learning.

[Carlini et al., 2018] attempt to directly quantify prototypicality with various metrics and verify that all of them agree with human intuition of prototypicality to various extents. In particular, they conclude that with CIFAR-10 , training on nearly-the-most prototypical examples gives the best performance when using 10% of the training data.

Most recently [Vodrahalli et al., 2018] attempts to find redundancies in image recognition datasets by analyzing gradient magnitudes as a measure of importance. They prioritize examples with high gradient magnitude according to a pre-trained classifier. Their method fails to find redundancies in CIFAR-10 and ImageNet datasets.

Finally, the insights provided by our work may have implications for semi-supervised techniques assessed on notorious image datasets. Currently when evaluated on ImageNet or CIFAR datasets, a fixed-sized subset of the dataset is *randomly* selected according to uniform distribution, and their labels are removed [Ren et al., 2018, Qiao et al., 2018, Tarvainen and Valpola, 2017, Pu et al., 2016, Sajjadi et al., 2016]. This creates a training set with mix of labeled and unlabeled data to be used for assessing semi-supervised learning methods. However, creating the training set by maintain the most informative fraction of the labeled examples may provide new insights about capabilities of semi-supervised methods.

3 Method

3.1 Motivation

In order to find redundancies, it is crucial to analyze each sample in the context of other samples in the dataset. Unlike previous attempts, we seek to measure redundancy by explicitly looking at a dissimilarity measure between samples. In case of there being near-duplicates in the training data, the approach of [Vodrahalli et al., 2018] will not be able to decide between them if their resulting gradient magnitude is high, whereas a dissimilarity measure can conclude that they are redundant if it evaluates to a low value.

3.2 Algorithm

To find redundancies in datasets, we look at the semantic space of a pre-trained model trained on the full dataset. In our case, the semantic representation comes from the penultimate layer of a neural network. To find groups of points which are close by in the semantic space we use Agglomerative Clustering [Defays, 1977]. Agglomerative Clustering assumes that each point starts out as its own cluster initially, and at each step, the pair of clusters which are closest according to the dissimilarity criterion are joined together. Given two images I_1 and I_2, whose latent representations are denoted by vectors $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$. We denote the dissimilarity between $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ by $d(\boldsymbol{x}_1, \boldsymbol{x}_2)$ using the cosine angle between them as follows:

$$d(\boldsymbol{x}_1, \boldsymbol{x}_2) = 1 - \frac{\langle \boldsymbol{x}_1, \boldsymbol{x}_2 \rangle}{\|\boldsymbol{x}_1\| \, \|\boldsymbol{x}_2\|} \, . \tag{1}$$

The dissimilarity between two clusters C_1 and C_2 , $D(C_1, C_2)$ is the maximum dissimilarity between any two of their constituent points:

$$D(C_1, C_2) = \max_{\boldsymbol{x}_1 \in C_1, \boldsymbol{x}_2 \in C_2} d(\boldsymbol{x}_1, \boldsymbol{x}_2) \ . \tag{2}$$

For Agglomerative Clustering, we process points belonging to each class independently. Since the dissimilarity is a pairwise measure, processing each class separately leads to faster computations. We run the clustering algorithm until there are k clusters left, where k is the size of the desired subset. We assume that points inside a cluster belong to the same *redundant group* of images. In each redundant group, we select the image whose representation is closest to the cluster center and discard the rest. Henceforth, we refer to this procedure as semantic space clustering or *semantic clustering* for brevity.

4 Experiments

We use the ResNet [He et al., 2016a] architecture for all our experiments with the variant described in [He et al., 2016b]. For each dataset, we compare the performance after training on different random subsets to subsets found with semantic clustering. Given a fixed pre-trained model, semantic clustering subsets are deterministic and the only source of stochasticity is due to the random network weight initialization and random mini-batch choices during optimization by SGD.

The semantic space embedding is obtained by pre-training a network on the full dataset. We chose the output after the last average pooling layer as our semantic space representation. All hyperparameters are kept identical during pre-training and also when training with different subset sizes.

As the baseline, we compare against a subset of size k uniformly sampled from the full set. Each class is sampled independently to in order to be consistent with the semantic clustering scheme. Note that random sampling scheme adds an additional source of stochasticity compared to clustering. For both either uniform sampling or cluster based subset selection, we report the mean and standard deviation of the test accuracy of the model trained from scratch using the subset.

4.1 CIFAR-10 & CIFAR-100

We train a 32-layer ResNet for the CIFAR-10 and CIFAR-100 [Krizhevsky and Hinton, 2009] datasets. The semantic representation obtained was a 64-dimensional vector. For both the datasets, we train for 100,000 steps with a learning rate which is cosine annealed [Loshchilov and Hutter, 2016] from 0.1 to 0 with a batch size of 128.

For optimization we use Stochastic Gradient Descent with a momentum of coefficient of 0.9. We regularize our weights by penalizing their ℓ_2 norm with a factor of 0.0001. We found that to prevent weights from diverging when training with subsets of all sizes, warming up the learning rate was necessary. We use linear learning rate warm-up for 2500 steps from 0. We verified that warming up the learning rate performs slightly better than using no warm-up when using the full dataset.

In all these experiments, we report average test accuracy across 10 trials.

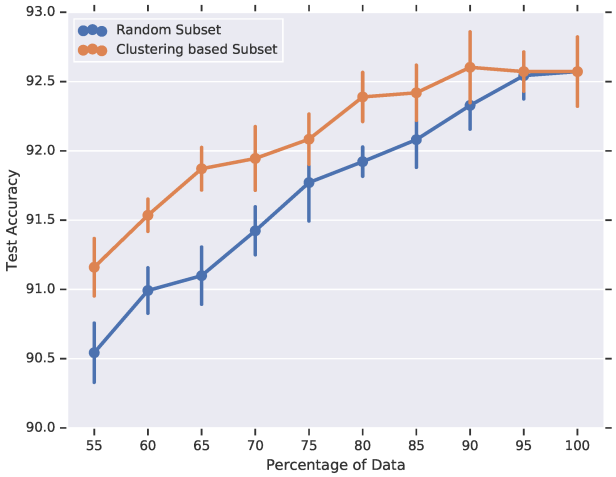

Figure 2: Performance of subsets of varying size on the CIFAR-10 dataset. Each point is an average across 10 trials and the vertical bars denote standard deviation. We see no drop in test accuracy until 10% of the data considered redundant by semantic clustering is removed.

4.1.1 CIFAR-10

We see in the case of the CIFAR-10 dataset in Figure 2 that the same test accuracy can be achieved even after 10% of the training is discarded using semantic clustering. In contrast, training on random subsets of smaller sizes, results in a monotonic drop in performance. Therefore, while we show that at least 10% of the data in the CIFAR-10 dataset is redundant, this redundancy cannot be observed by uniform sampling.

Figure 3 shows examples of images considered redundant with semantic clustering while choosing a subset of 90% size of the full dataset. Each set denotes images the were placed in into the same (redundant) group by semantic clustering. Images in green boxes were retained while the rest were discarded.

Figure 4 shows the number of redundant groups of different sizes for two classes in the CIFAR-10 dataset when seeking a 90% subset. Since a majority of points are retained, most clusters end up containing one element upon termination. Redundant points arise from clustering with two or more elements in them.

4.1.2 CIFAR-100

In the case of the CIFAR-100 dataset, our proposed scheme fails to find redundancies, as is shown in Figure 5, while it does slightly better than random subsets. Both proposed and random methods show a monotonic decrease in test accuracy with decreasing subset size.

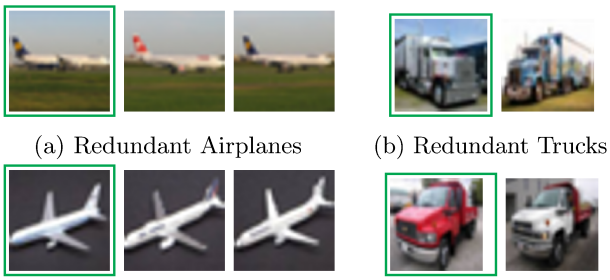

(a) Redundant Airplanes (b) Redundant Trucks

(c) Redundant Airplanes (d) Redundant Trucks

Figure 3: Examples of redundant images in the CIFAR-10 dataset when creating a subset of 90% size of the original set. The figure illustrates similarity between images of each redundant group and variation across different redundant groups. 3a and 3c are two different redundant groups of the class Airplane. 3b and 3d are two different redundant groups from class Truck. In each group, only the images marked with green boxes are kept and the rest, discarded. The discarded images did not lower test accuracy.

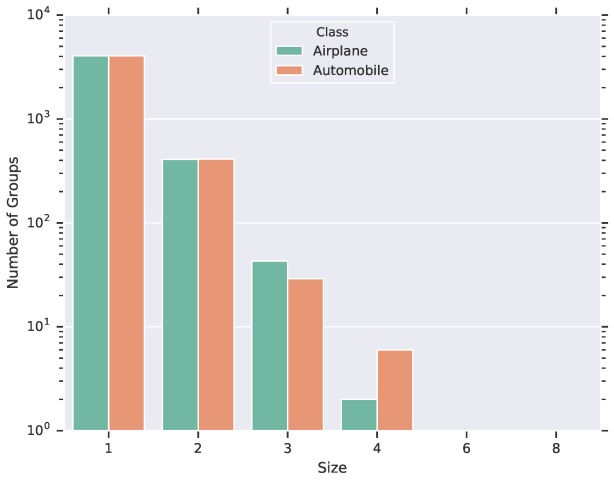

Figure 4: Number of redundant groups of various sizes in the CIFAR-10 dataset when finding a 90% subset for two classes. Note that the y-axis is logarithmic.

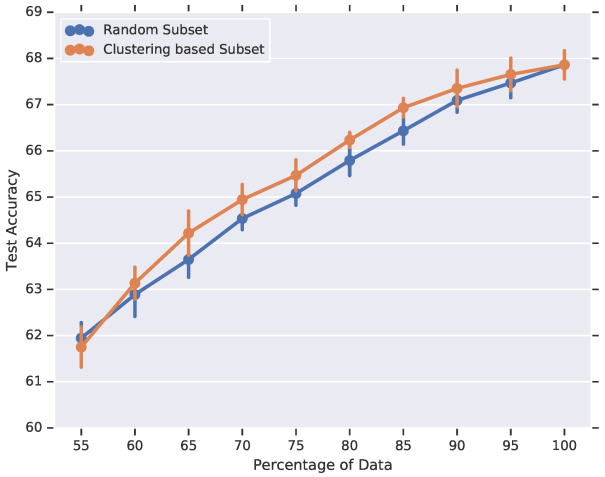

Figure 5: Performance of subsets of varying size on the CIFAR-100 dataset. Each point is an average over 10 trials and the vertical bars denote standard deviation.

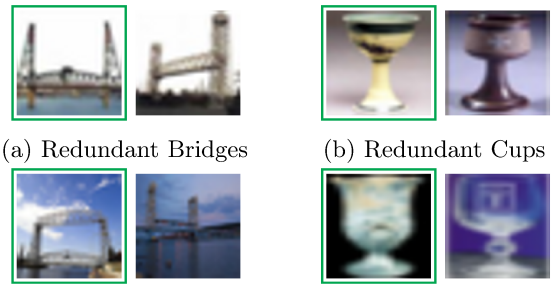

(a) Redundant Bridges (b) Redundant Cups

(c) Redundant Bridges (d) Redundant Cups

Figure 6: Example of variation between images in the same redundant group compared to variation across different redundant groups in the CIFAR-100 dataset. Each column contains a specific class of images. In contrast to Figure 3, the images within each redundant group show much more variations. The groups were found when retaining a 90% subset, and retraining only the selected images (in green boxes) and discarding the rest had a negative impact on test performance.

Figure 6 looks at redundant groups found with semantic clustering to retain 90% of the dataset. As compared to Figure 3, the images within a group show much more semantic variation. Redundant groups in Figure 3 are slight variations of the same object, where as in Figure 6, redundant groups do not contain the same object. We note that in this case the model is not able to be invariant to these semantic changes.

Similar to Figure 4, we plot the number of redundant groups of each size for two classes in CIFAR-100 in Figure 7.

To quantify the semantic variation of CIFAR-100 in relation to CIFAR-10 , we select redundant groups of size two or more, and measure the average dissimilarity(from Equation 1) to the retained sample. We report the average over groups in 3 different classes as well as the entire datasets in Table 1. It is clear that the higher semantic variation in the redundant groups of CIFAR-100 seen in Figure 6 translates to an higher average dissimilarity in Table 1.

4.2 Choice of semantic representation.

To determine the best choice of semantic representation from a pre-trained model, we run experiments after selecting the semantic representation from 3 different layers in the network. Figure 8 shows the results. Here "Start" denotes the semantic representation after the first Convolution layer is a ResNet, "Middle" denotes the representation after the second residual block, and "End" denotes the output of the last average pooling layer. We see that the "End" layer's semantic representation is able to find the largest redundancy.

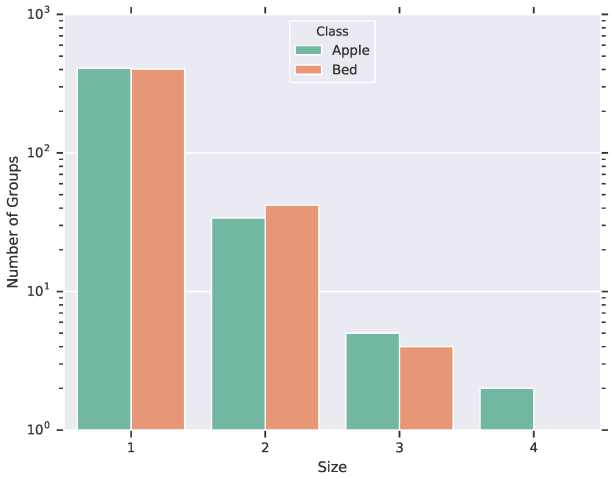

Figure 7: Number of redundant groups of various sizes in the CIFAR-100 dataset when finding a 90% subset for two classes. Note that the y-axis is logarithmic.

Dataset	Class	Average Dissimilarity
CIFAR-10	Airplane	1.73×10^{-3}
	Automobile	1.65×10^{-3}
	Bird	2.22×10^{-3}
	All (mean)	1.84×10^{-3}
CIFAR-100	Apple	6.61×10^{-3}
	Bed	14.16×10^{-3}
	Bowl	20.02×10^{-3}
	All (mean)	13.90×10^{-3}

Table 1: Average dissimilarity to the retained sample across redundant groups (clusters) of size greater than 1. We report the class-wise mean for 3 classes as well as the average over the entire dataset. All clusters were created to find a subset of 90% the size of the full set. We can observe that the average dissimilarity is about an order of magnitude higher for the CIFAR-100 dataset, indicating that there is more variation in the redundant groups.

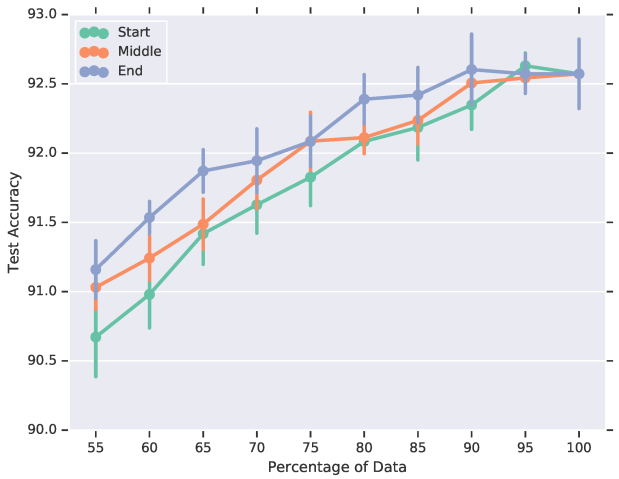

Figure 8: Effectiveness of latent representations from 3 stages in a ResNet at finding redundant subsets.

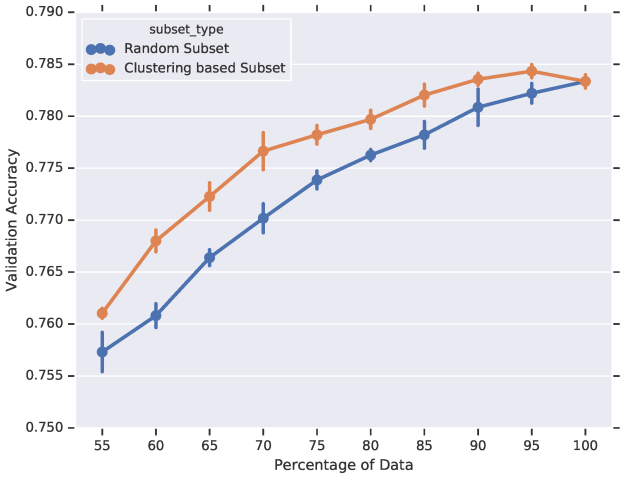

Figure 9: Validation accuracy after training with subsets of various sizes of ImageNet . We plot the average over 5 trials with the vertical bars denoting standard deviation. There is no drop in validation accuracy when 10% of the training data considered redundant by semantic clustering is removed.

4.3 ImageNet

We train a 101-layer ResNet with the ImageNet dataset. It gave us a semantic representation of 2048 dimensions. We use a batch size of 1024 during training and train for $120,000$ steps with a learning rate cosine annealed from 0.4 to 0. Using the strategy from [Goyal et al., 2017], we linearly warm up our learning rate from 0 for 5000 steps to be able to train with large batches. We regularize our weights with ℓ_2 penalty with a factor of 0.0001.

For optimization, we use Stochastic Gradient Descent with a Momentum coefficient of 0.9 while using the Nesterov momentum update. Since the test set is not publicly available we report the average validation accuracy, measured over 5 trials.

The results of training with subsets of varying sizes of ImageNet dataset are shown in Figure 9. Our proposed scheme is able to successfully show that at least 10% of the data can be removed from the training set without any negative impact on the validation accuracy, whereas training on random subsets always gives a drop with decrease in subset size.

Figure 1 shows different redundant groups found in the ImageNet dataset. It is noteworthy that the semantic change considered redundant is different across each group. Figure 11 highlights the similarities between images of the same redundant group and the variation across different redundant groups.

In each row of Figure 12, we plot two images from a redundant group on the left where the retained image is highlighted in a green box. On the right we display the image closest to each retained image in dissimilarity but excluded from the redundant group. These images were close in semantic space to the corresponding retained images, but were not considered similar enough to be

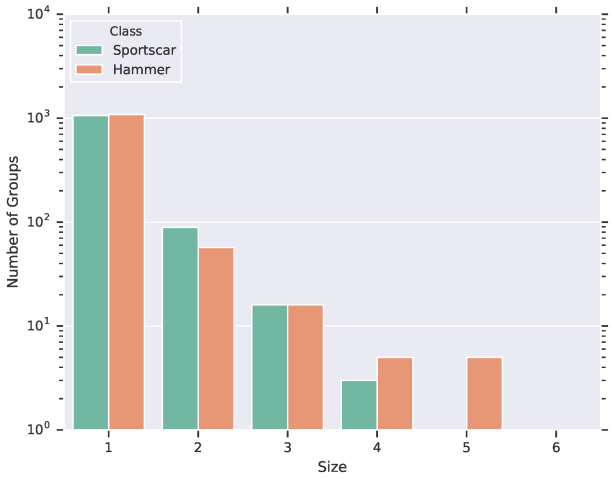

Figure 10: Sizes of redundant groups for the Hammer and Sports-car classes in the ImageNet dataset when finding a 90% subset. Note that the y-axis is logarithmic.

redundant. For example the redundant group in the first row of Figure 12 contains Sedan-like looking red cars. The 2-seater sports car on the right, in spite of looking similar to the cars on the left, was not considered redundant with them.

Figure 10 shows the number of redundant groups of each size when creating a 90% subset. Much akin to Figure 4, a majority of images are not considered redundant and form a group of size 1.

Additional examples of redundancy group on ImageNet is provided in the appendix.

4.4 Implementation Details

We use the open source Tensorflow [Abadi et al., 2016] and tensor2tensor[Vaswani et al., 2018] frameworks to train our models. For clustering, we used the scikit-learn [Pedregosa et al., 2011] library. For the CIFAR-10 and CIFAR-100 experiments we train on a single NVIDIA Tesla P100 GPU. For our ImageNet experiments we perform distributed training on 16 Cloud TPUs.

5 Conclusion

In this work we present a method to find redundant subsets of training data. We explicitly model a dissimilarity metric into our formulation which allows us to find semantically close samples that can be considered redundant. We use an agglomerative clustering algorithm to find redundant groups of images in the semantic space. Through our experiments we are able to show that *at least* 10% of ImageNet and CIFAR-10 datasets are redundant.

We analyze these redundant groups both qualitatively and quantitatively. Upon visual observation, we see that the semantic change considered redundant varies from cluster to cluster. We show examples of a variety of varying attributes in redundant groups, all of which are redundant from the point of view of training the network.

One particular justification for not needing this variation during training could be that the network learns to be invariant to them because of its shared parameters and seeing similar variations in other parts of the dataset.

In Figure 2 and 9, the accuracy without 5% and 10% of the data is slightly higher than that obtained with the full dataset. This could indicate that redundancies in training datasets hamper the optimization process.

For the CIFAR-100 dataset our proposed scheme fails to find any redundancies. We qualitatively compare the redundant groups in CIFAR-100 (Figure 6) to the ones found in CIFAR-10 (Figure 3) and find that the semantic variation across redundant groups is much larger in the former case. Quantitatively this can be seen in Table 1 which shows points in redundant groups of CIFAR-100 are much more spread out in semantic space as compared to CIFAR-10 .

Although we could not find any redundancies in the CIFAR-100 dataset, there could be a better algorithm that could find them. Moreover, we hope that this work inspires a line of work into finding these redundancies and leveraging them for faster and more efficient training.

6 Acknowledgement

We would like to thank colleagues at Google Research for comments and discussions: Thomas Leung, Yair Movshovitz-Attias, Shraman Ray Chaudhuri, Azade Nazi, Serge Ioffe.

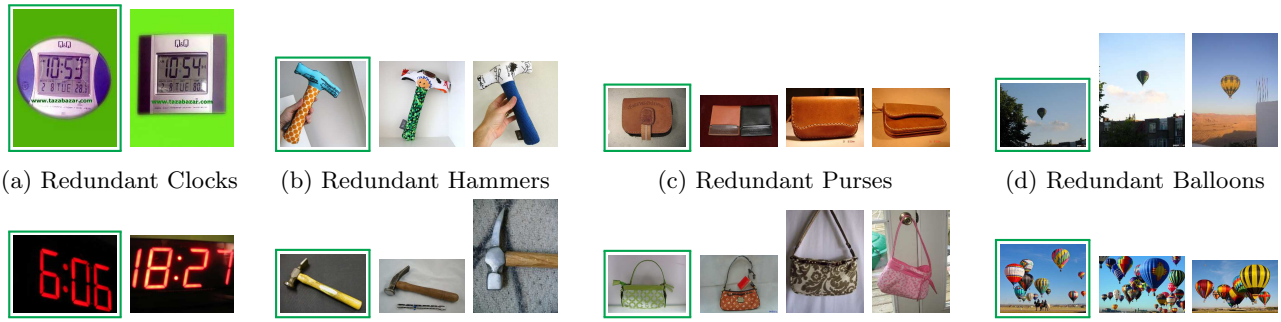

(a) Redundant Clocks (b) Redundant Hammers (c) Redundant Purses (d) Redundant Balloons

(e) Redundant Clocks (f) Redundant Hammers (g) Redundant Purses (h) Redundant Balloons

Figure 11: This figure highlights semantic similarities between images from the same redundant group and variation seen across different redundant groups of the same class. The redundant groups were found while creating a 90% subset of the ImageNet dataset. Each sub-figure is a redundant group of images according to our algorithm. Each column contains images belonging to the same class, with each row in a column being a different redundant group. For example, the first column contains the *Clock* class. Clocks in 11a are in one group of redundant images whereas clocks in 11e are in another group. From each of the groups in the sub-figures, only the images marked in green boxes are selected by our algorithm and the others are discarded. Discarding these images had no negative impact on validation accuracy.

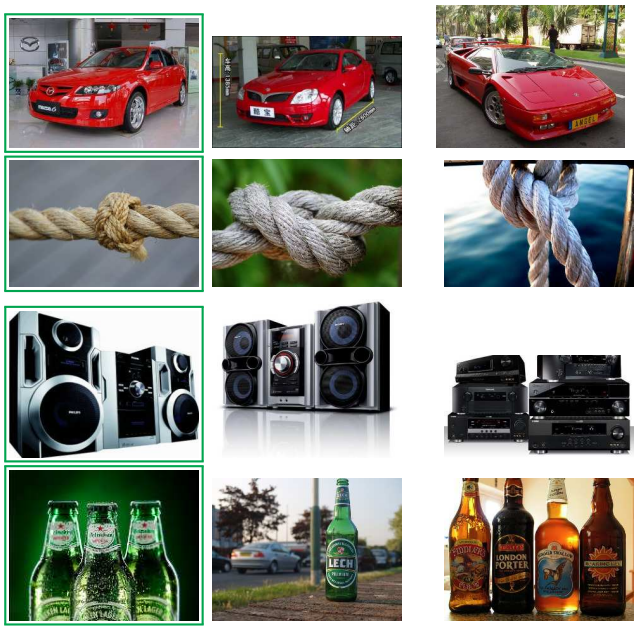

Figure 12: In each row we plot two images from the same redundant group while creating a 90% subset on the left with the retained image highlighted in a green box. On the right we plot the image closest to the retained image in the semantic space but not included in the same redundant group. Note that the image on the right shows a semantic variation which is inconsistent with the one seen in the redundant group.

References

[Abadi et al., 2016] Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283.

[Carlini et al., 2018] Carlini, N., Erlingsson, U., and Papernot, N. (2018). Prototypical examples in deep learning: Metrics, characteristics, and utility. Technical report.

[Defays, 1977] Defays, D. (1977). An efficient algorithm for a complete link method. *The Computer Journal*, 20(4):364–366.

[Deng et al., 2009] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. Ieee.

[Fan et al., 2016] Fan, Y., Tian, F., Qin, T., and Liu, T.-Y. (2016). Neural data filter for bootstrapping stochastic gradient descent. Technical report.

[Gastaldi, 2017] Gastaldi, X. (2017). Shake-shake regularization. *arXiv preprint arXiv:1705.07485*.

[Goyal et al., 2017] Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., and He, K. (2017). Accurate, large minibatch sgd: training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*.

[Halevy et al., 2009] Halevy, A., Norvig, P., and Pereira, F. (2009). The unreasonable effectiveness of data. *IEEE Intelligent Systems*, 24(2):8–12.

[He et al., 2016a] He, K., Zhang, X., Ren, S., and Sun, J. (2016a). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

[He et al., 2016b] He, K., Zhang, X., Ren, S., and Sun, J. (2016b). Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer.

[Huang et al., 2017] Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. (2017). Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269. IEEE.

[Huang et al., 2018] Huang, Y., Cheng, Y., Chen, D., Lee, H., Ngiam, J., Le, Q. V., and Chen, Z. (2018). Gpipe: Efficient training of giant neural networks using pipeline parallelism. *arXiv preprint arXiv:1811.06965*.

[Katharopoulos and Fleuret, 2018] Katharopoulos, A. and Fleuret, F. (2018). Not all samples are created equal: Deep learning with importance sampling. *arXiv preprint arXiv:1803.00942*.

[Kaushal et al., 2018] Kaushal, V., Sahoo, A., Doctor, K., Raju, N., Shetty, S., Singh, P., Iyer, R., and Ramakrishnan, G. (2018). Learning from less data: Diversified subset selection and active learning in image classification tasks. *arXiv preprint arXiv:1805.11191*.

[Krizhevsky and Hinton, 2009] Krizhevsky, A. and Hinton, G. (2009). Learning multiple layers of features from tiny images. Technical report, Citeseer.

[Krizhevsky et al., 2012] Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.

[Kuznetsova et al., 2018] Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Malloci, M., Duerig, T., et al. (2018). The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*.

[Lapedriza et al., 2013] Lapedriza, A., Pirsiavash, H., Bylinskii, Z., and Torralba, A. (2013). Are all training examples equally valuable? *arXiv preprint arXiv:1311.6510*.

[Lin et al., 2018] Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. (2018). Focal loss for dense object detection. *IEEE transactions on pattern analysis and machine intelligence*.

[Loshchilov and Hutter, 2016] Loshchilov, I. and Hutter, F. (2016). Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.

[Ohno-Machado et al., 1998] Ohno-Machado, L., Fraser, H. S., and Ohrn, A. (1998). Improving machine learning performance by removing redundant cases in medical data sets. In *Proceedings of the AMIA Symposium*, page 523. American Medical Informatics Association.

[Pedregosa et al., 2011] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830.

[Pu et al., 2016] Pu, Y., Gan, Z., Henao, R., Yuan, X., Li, C., Stevens, A., and Carin, L. (2016). Variational autoencoder for deep learning of images, labels and captions. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29*, pages 2352–2360. Curran Associates, Inc.

[Qiao et al., 2018] Qiao, S., Shen, W., Zhang, Z., Wang, B., and Yuille, A. L. (2018). Deep co-training for semi-supervised image recognition. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XV*, pages 142–159.

[Ren et al., 2018] Ren, M., Ravi, S., Triantafillou, E., Snell, J., Swersky, K., Tenenbaum, J. B., Larochelle, H., and Zemel, R. S. (2018). Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations*.

[Sajjadi et al., 2016] Sajjadi, M., Javanmardi, M., and Tasdizen, T. (2016). Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29*, pages 1163–1171. Curran Associates, Inc.

[Simonyan and Zisserman, 2014] Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

[Srivastava et al., 2014] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.

[Sun et al., 2017] Sun, C., Shrivastava, A., Singh, S., and Gupta, A. (2017). Revisiting unreasonable effectiveness of data in deep learning era. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 843–852. IEEE.

[Tarvainen and Valpola, 2017] Tarvainen, A. and Valpola, H. (2017). Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 1195–1204. Curran Associates, Inc.

[Tobin et al., 2017] Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. In *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*, pages 23–30. IEEE.

[Vaswani et al., 2018] Vaswani, A., Bengio, S., Brevdo, E., Chollet, F., Gomez, A. N., Gouws, S., Jones, L., Kaiser, Ł., Kalchbrenner, N., Parmar, N., et al. (2018). Tensor2tensor for neural machine translation. *arXiv preprint arXiv:1803.07416*.

[Vodrahalli et al., 2018] Vodrahalli, K., Li, K., and Malik, J. (2018). Are all training examples created equal? an empirical study. *CoRR*, abs/1811.12569.

[Wei et al., 2015] Wei, K., Iyer, R., and Bilmes, J. (2015). Submodularity in data subset selection and active learning. In *International Conference on Machine Learning*, pages 1954–1963.

[Wu et al., 2019] Wu, B., Chen, W., Fan, Y., Zhang, Y., Hou, J., Huang, J., Liu, W., and Zhang, T. (2019). Tencent ml-images: A large-scale multi-label image database for visual representation learning. *arXiv preprint arXiv:1901.01703*.

A Appendix

Each row is a redundant group of images. The left most image is retained in each row for the 90% subset.

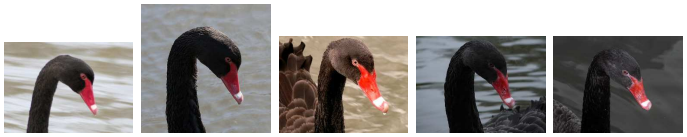

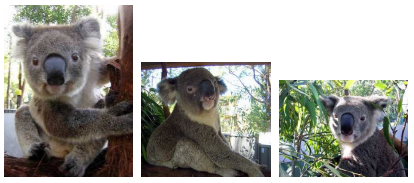

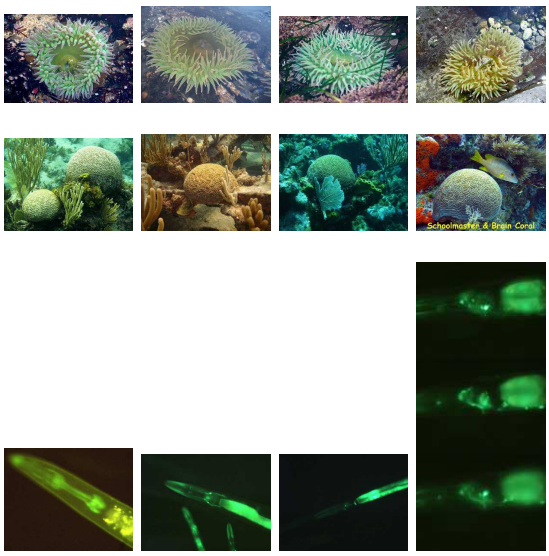
Schoolmaster & Brain Coral