
Zero-shot Active Learning Using Self Supervised Learning

Abhishek Sinha
Department of Computer Science
Stanford University
a7b23@stanford.edu

Shreya Singh
Department of Computer Science
Stanford University
ssingh16@stanford.edu

1 Introduction

Deep learning algorithms are often said to be data hungry. The performance of such algorithms generally improve as more and more annotated data is fed into the model. While collecting unlabelled data is easier (as they can be scraped easily from the internet), annotating them is a tedious and expensive task. Given a fixed budget available for data annotation, Active Learning helps selecting the best subset of data for annotation, such that the deep learning model when trained over that subset will have maximum generalization performance under this budget.

A common theme in active learning is to start with a random subset of labelled data and then use different strategies to iteratively add new datapoints. We believe that there are three limitations of the current active learning approaches.

- For the current approaches examples are added iteratively in batches, which can make this process difficult to scale. If an organization uses a third party system to do data annotation process, then back and forth iterative communication between them can make the process tedious and difficult to use.
- Active learning approaches are not model agnostic. It means that for any new model, another round of Active learning will have to be run to get the best subset of data given the budget constraints.
- The current approaches start with a random subset of data. Thus the initial data that is collected for labelling is not an optimal subset.

In this work, we aim to propose a new Active Learning approach which is model agnostic as well as one doesn't require an iterative process. We aim to leverage self-supervised learnt features for the task of Active Learning. The benefit of self-supervised learning, is that one can get useful feature representation of the input data, without having any annotation.

2 Related Work

2.1 Active Learning

Two popular directions for Active Learning are -

- Uncertainty Estimation - [6, 10] Select datapoints for which the model is most uncertain about. The idea here is that adding such uncertain data points will help the model better learn about the data, which in turn help it to generalize well.

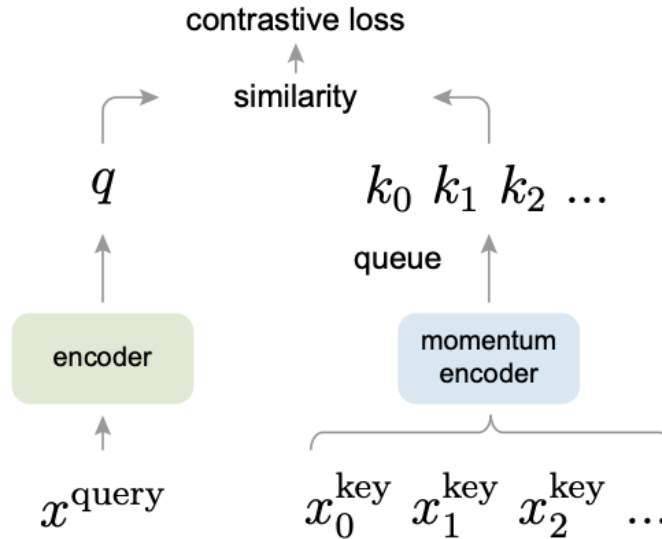


Figure 1: Schematic Overview of Momentum Contrastive Learning Method

- Diversity - [9, 7] - Select datapoints such that the diversity of labelled set is increased. [9] uses core-set/k-centre greedy approach to iteratively sample new datapoints, using features extracted from the model that has been trained over the current pool of labelled data. We will base our algorithm over this approach.

2.2 Self-Supervised Learning

Self-supervised learning [2, 3], is another area of research which aims to do away with the expensive data annotation process. The idea here is to learn good features from the data, which can then be used for other downstream tasks.

Momentum Contrastive Learning [1]. Momentum Contrast (MoCo) trains a visual representation encoder by matching an encoded query q to a dictionary of encoded keys using a contrastive loss. The dictionary keys $\{k_0, k_1, k_2, \dots\}$ are defined on-the-fly by a set of data samples. The dictionary is built as a queue, with the current mini-batch enqueued and the oldest mini-batch dequeued, decoupling it from the mini-batch size. The keys are encoded by a slowly progressing encoder, driven by a momentum update with the query encoder. The learned representations are then evaluated by training a linear classifier using labeled training data and evaluating on a test data. This method enables a large and consistent dictionary for learning visual representations. Figure 1 shows a schematic overview of MoCo.

3 Datasets

We will be performing our experiments over CIFAR-10 dataset [8]. The CIFAR-10 dataset consists of 60000 32x32 colour images in 10 classes, with 6000 images per class. There are 50000 training images and 10000 test images. The dataset is publicly available.

4 Method

We aim to use self-supervised for the task of active learning. Since the features can be extracted without any annotation, these features are agnostic to the model as well as the task which will be used for a given dataset. Further, since these features are pre-computed once, and do not change as a downstream task is performed for the dataset, we can get an ordering in which the data needs to be

annotated from before without going into an iterative process. Specifically, our approach consists of the following steps -

- Train a self-supervised model over the CIFAR-10 dataset, or use an existing self-supervised model trained over the ImageNet dataset [4].
- Get features for all the images in the CIFAR-10 dataset.
- Apply k-centre greedy algorithm to get the ordering in which the dataset needs to be labelled [9].
- Given the budget constraint select the top ordered data-points, annotate them and then train a model for this subset.

In the k-centre greedy approach we start with k randomly selected datapoints as cluster centres. Given a fixed budget B we iteratively add B datapoints by selecting the one which is farthest away from the cluster centres. Once a datapoint is selected, it is also added as a new cluster centre. This is a diversity based Active Learning approach, as we want the selected data-points to be as diverse as possible

5 Experiment Details

We perform our experiments over the CIFAR-10 dataset, and the task is classification. We choose Wide ResNet [11] as the model for this task. We further use CutOut [5] augmentation to improve the generalization ability.

We evaluate our approach for different budget constraints, where the budget is defined by the number of annotated/labelled data-points. We show results for budgets of 5k to 25k with an interval of 2500. Thus if the budget is B , we select B data-points, annotate them, train our Wide ResNet using the B data-points and then report the test accuracy of this model.

5.1 Baselines

We choose the following two baseline approaches for comparison -

- **Random** - For each of the budget constraint, randomly select those number of datapoints from the dataset.
- **Core-Set** [9] - Same as our approach, except that it uses a supervised model instead of self-supervised model. Thus, given a subset of annotated datapoints till now, it trains a supervised model over it, and then uses the pre-final layer embeddings as the features of an input. The features are used to perform k-centre greedy approach and then select the new data-points. Since, for any selected subset we need to train a model to get the features, thus the features will be dynamically updated for every iteration of Active learning. This is in contrast to us, where our features are fixed throughout the Active Learning process.

5.2 Active Learning Using Self-Supervised Features (Our Approach)

We use MoCo-v2 [3] and SimCLR [1] to perform self-supervised learning over the CIFAR-10 dataset. We use the following data-augmentation strategies for MoCo-v2 training -

- **Random Grayscale** - Images are randomly converted to grayscale with a probability of 0.2
- **Random Resize Crop** - A crop of scale randomly chosen between 0.08 to 1.0 is made, before resizing the crop again to size 32*32.
- **Random Horizontal Flip** - The image is randomly flipped horizontally with a probability of 0.5.
- **Random Color jitter** - The brightness, hue and saturation are randomly changed with a probability of 0.8.

To test the efficacy of the self supervised learning, we use the linear evaluation protocol [3, 1] where the network is frozen, and only a linear layer is learnt over the labelled trainset. On the linear evaluation protocol, MoCo-v2 achieves a test accuracy of 91.2% whereas SimCLR achieves a test accuracy of 93.6%

5.3 Evaluation

As mentioned above we train different WideResNet models for each of the budget constraints, and then compare the test accuracies of the different algorithms for each budget.

5.4 Results

Figure 2 shows the results for the two baselines as well as our approach over different data splits. As can be seen from the figure, Our approach performs better than the random baseline, but performs slightly worse the Core-Set baseline when the number of labelled samples are less than 20k. For data split sizes of 20k and more, our approach is comparable to the Core-Set baseline.

Our approach although performs comparable or slightly worse than the Core-Set baseline, it has the advantage that the method is non-iterative based and also is model agnostic. Thus, even if instead of Wide ResNet we used any other model for classification, the same sampled data-points could have been used for labelling.




Figure 2: Performance of different approaches

6 Discussion

We plot the frequency of each of the classes amongst the 20k data points selected for labelling in Figures 3 and 4.

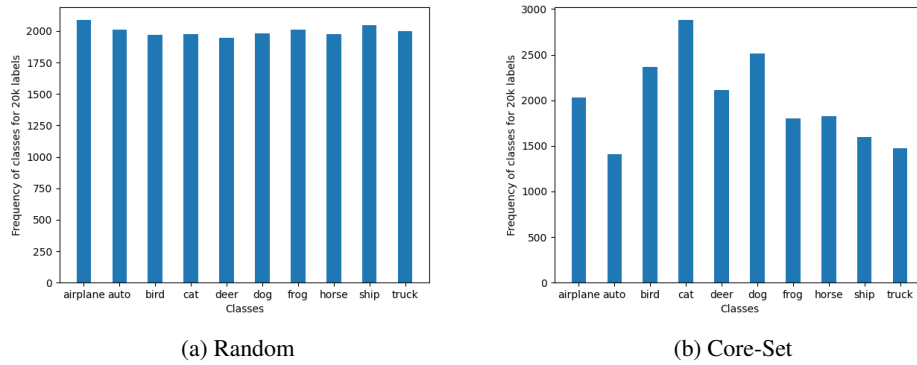


Figure 3: Frequency of classes in the 20k data points selected for labelling for the Baselines




Figure 4: Frequency of classes in the 20k data points selected for labelling for the Baselines

As can be seen in Figure 3a random selection uniformly selects each of the 10 classes. However, the core-set baseline selects examples from cat class more than other as seen in Figure 3b. A similar trend exists for our approaches too in Figures 4a and 4b where the cat class is selected more.

7 Conclusion and Future Work

In this work we proposed a non-iterative and model agnostic method for performing Active Learning using self-supervised features. While being non-iterative our approach is still comparable/slightly worse to the existing state-of-the-art approaches. We showed via frequency histogram, the different classes selected in the labelled dataset, and find similar behaviour as to the state-of-the-art approaches.

In the future we will like to explore other datasets such as CIFAR-100 and ImageNet [4] and also try to further improve our performance with possibly better self-supervised features.

8 Contributions

Abhishek implemented the baselines and did experiments with the MoCo-v2 approach. Shreya did experiments with the SimCLR approach. The project was also advised by Jiaming Song.

References

- [1] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.

- [2] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [3] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [5] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [6] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192. PMLR, 2017.
- [7] Andreas Kirsch, Joost Van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *arXiv preprint arXiv:1906.08158*, 2019.
- [8] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [9] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*, 2017.
- [10] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 93–102, 2019.
- [11] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.