

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

S. A. Koohpayegani, A. Tejankar and H. Pirsiavash, "Mean Shift for Self-Supervised Learning," 2021 IEEE/CVF International Conference on Computer Vision (ICCV), 2021, pp. 10306-10315, doi: 10.1109/ICCV48922.2021.01016

<https://doi.org/10.1109/ICCV48922.2021.01016>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Mean Shift for Self-Supervised Learning

Soroush Abbasi Koohpayegani^{1,*} Ajinkya Tejankar^{1,*} Hamed Pirsiavash^{1,2}

¹ University of Maryland, Baltimore County ²University of California, Davis

Abstract

Most recent self-supervised learning (SSL) algorithms learn features by contrasting between instances of images or by clustering the images and then contrasting between the image clusters. We introduce a simple mean-shift algorithm that learns representations by grouping images together without contrasting between them or adopting much of prior on the structure or number of the clusters. We simply “shift” the embedding of each image to be close to the “mean” of the neighbors of its augmentation. Since the closest neighbor is always another augmentation of the same image, our model will be identical to BYOL when using only one nearest neighbor instead of 5 used in our experiments. Our model achieves 72.4% on ImageNet linear evaluation with ResNet50 at 200 epochs outperforming BYOL. Also, our method outperforms the SOTA by a large margin when using weak augmentations only, facilitating adoption of SSL for other modalities. Our code is available here: <https://github.com/UMBCvision/MSF>

1. Introduction

Most current visual recognition algorithms are supervised, meaning that they learn from large scale annotated images or videos. However, in many applications, the annotation process may be expensive, biased, ambiguous, or involve privacy concerns. Self-supervised learning (SSL) algorithms aim to learn rich representations from unlabeled images or videos. Such learned representations can be used along with small annotated data to provide an accurate visual recognition model. We are interested in developing better SSL models using unlabeled images.

Some recent SSL models learn by contrasting between instances of images. They pull different augmentations of an image instance together while pushing them away from other image instances [23, 11]. Some other SSL methods cluster the unlabeled images to a set of clusters with the hope that each cluster will contain semantically similar im-

ages. Then, a model that predicts those clusters learns rich representations similar to supervised learning with labels [9, 52, 10].

These clustering methods also can be considered as contrastive learning since they contrast between different clusters of images. For instance, the SoftMax layer in deep clustering method [9] encourages an image to be assigned to the correct single cluster and not the other clusters.

Also, most clustering algorithms have strong priors on the overall structure of the clusters. For instance, deep clustering (k-means) using Euclidean distance encourages spherical cluster shapes which we believe is unnecessary for the purpose of SSL methods.

Recently, BYOL [22] showed that it is possible to learn rich representations without contrasting between image instances. BYOL [22] works by simply pulling the two views of an image closer without any contrast with other images. The better performance of BYOL [22] compared to MoCo hints that contrasting with other images may be a limiting constraint. For instance, in MoCo [23], since the negative images are sampled randomly, they may be from the same category as the query, resulting in degraded representations. [42] tries to resolve this issue by not treating all negatives equally negative.

Inspired by mean shift clustering, we generalize BYOL to a simple yet effective SSL method where a data point is pulled closer to not only its other augmentations but also the nearest neighbors (NNs) of its augmentation. Unlike DeepCluster [9], SwAV [10], and SeLA [52] that use explicit, mutually exclusive cluster assignment, our method uses mean-shift algorithm that groups similar images together locally without explicit cluster assignment. Moreover, unlike k-means clustering, mean-shift does not have strong priors on the shape, size, or number of the clusters. This makes mean-shift suitable for SSL where such priors are not known. Compared to MoCo [23], SimCLR [11], SwAV [10] and few others, our method never contrasts between different images or even cluster centers.

Since we need a large set of embeddings to search for nearest neighbors, we adopt the memory bank idea [23] to maintain a random set of embeddings. Also, since the

*Equal contribution

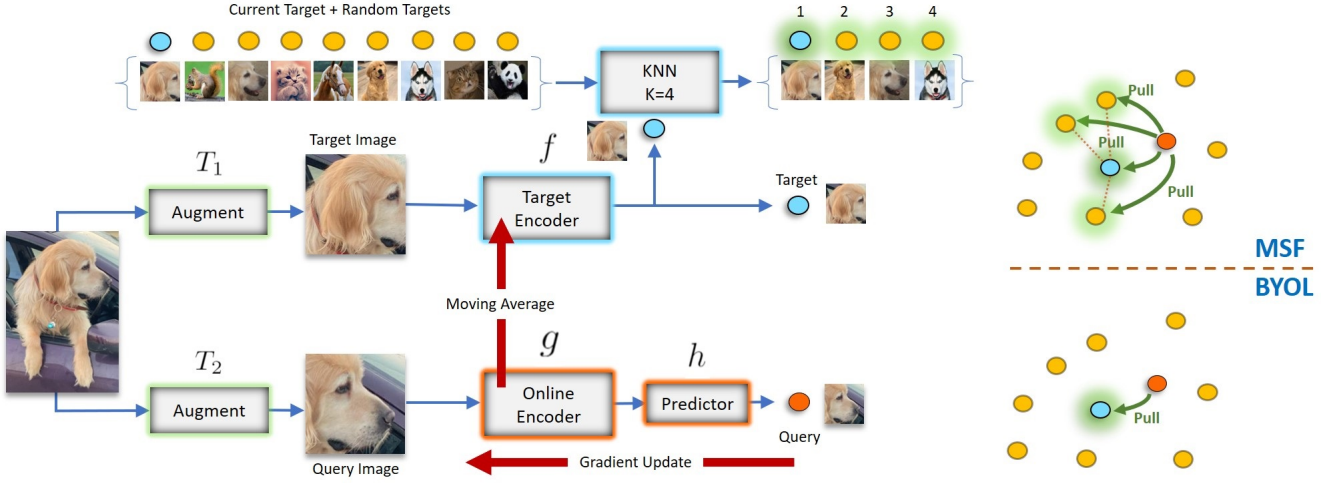


Figure 1: MSF method: Similar to BYOL, we maintain two encoders (“target” and “online”). The online encoder is updated with gradient descent while the target encoder is the moving average of the online encoder. We augment an image twice and feed to both encoders. We add the target embedding to the memory bank and look for its nearest neighbors in the memory bank. Obviously target embedding itself will be the first nearest neighbor. We want to shift the embedding of the input image towards the mean of its nearest neighbors, so we minimize the summation of those distances. Note that our method using only one nearest neighbor is identical to BYOL which pulls different augmentations together without grouping different instances of images. To our knowledge, our method is the first in grouping different instances of images without contrasting between image instances or clusters.

model is evolving over time in the learning process, the old elements in the memory bank will not be valid, so we adopt the momentum idea from [23] to maintain two encoders (“target” and “online”) instead of only one. The online model is updated by the loss and the target model is updated as a moving average of the online model. We feed two different augmentations of an image to these two encoders, then we push the online embedding of the image to be close to the average of nearest neighbors of the target encoding of the image in the target embedding space. Hence, similar to most recent SSL methods, our method also uses the inductive bias that the augmentation should not move the embedding much.

Our experiments show that our method outperforms state-of-the-art methods on various settings. For instance, when trained on unlabeled ImageNet for 200 epochs, it achieves 72.4% linear ImageNet accuracy which is better than BYOL at 200 epochs.

Most recent SSL methods use strong augmentations to improve the accuracy, leading to “augmentation engineering” to improve SSL. However, in many applications, e.g., medical domain, designing such augmentations is not easy and needs extensive domain knowledge. Hence, designing SSL methods that do not rely heavily on large variations of augmentations is interesting. We show that when using only weak augmentations, our method (MSF w/w) outperforms BYOL by a large margin. We hypothesize that NNs act as a proxy for strong augmentations of the query image, so there

is no need for engineering strong augmentations.

2. Method

We are interested in mean-shift clustering, so at each iteration we want to encourage the model to shift the embedding of an image to be closer to the average of its nearest neighbors on a large random set of samples.

Following the notation of BYOL [22], we assume a target encoder f and an online encoder g . Both encoders have the same backbone architecture followed by a projection layer and are initialized equally. The online encoder g is followed by an additional prediction layer h on top of it. The online encoder g and the prediction layer h are updated by back-propagating the loss while the target encoder f is updated by momentum update to be a running average of the online encoder g . Since nearest neighbor needs a large pool of examples, we maintain a first-in-first-out (FIFO) memory bank [23] that includes recent embeddings from the slowly evolving target encoder f .

Given an unlabeled image x , we augment it randomly twice to get $T_1(x)$ and $T_2(x)$. We feed them to encoders and then normalize them with ℓ_2 norm to get $u = \frac{f(T_1(x))}{\|f(T_1(x))\|_2}$ and $v = \frac{h(g(T_2(x)))}{\|h(g(T_2(x)))\|_2}$. We first add u to the memory bank and then, find the k nearest neighbors of u in the memory bank to get a set of embeddings $\{z_j\}_{j=1}^k$. Note that this set includes u itself. Since we know it is another augmentation of the same input image, it should be a good target for v .



Figure 2: Nearest neighbors (NN) of the model at each epoch: For a random query images, we show how the nearest neighbors evolve at the learning time. Initially, NNs are not semantically quite related, but are close in low-level features. The accuracy of 1-NN classifier in the initialization is 1.5% which is 15 times larger than random chance (0.1%). This little signal is bootstrapped in our learning method and results in NNs of the late epochs which are mostly semantically related to the query image. More visualizations can be found in the appendix.

Finally, we minimize the following loss:

$$L = \frac{1}{k} \sum_{j=1}^k \text{dist}(v, z_j)$$

where, $\text{dist}(\cdot, \cdot)$ is the distance metric between two embeddings. We use MSE loss ($\text{dist}(a, b) = \|a - b\|_2^2$) as the distance in our experiments. Minimizing this loss is equivalent to maximizing Cosine similarity as the vectors are already ℓ_2 normalized. The final loss is the summation of the above loss for all input images.

Ideally, we can average the set of nearest neighbors to come up with a single target embedding, but since averaging is dependent on the choice of loss function, we simply minimize the summation of distances. Note that for Euclidean distance, both methods result in identical gradients.

Since u itself is included in our NN search, it will be always the best nearest neighbor. Hence, our method with $k = 1$ will be identical to BYOL which minimizes $\|v - u\|_2^2$ for each image without using a memory bank.

Moreover, in the initial stages of the learning, v may be far from u and the other $k - 1$ nearest neighbors may be semantically different from the query image. Since those wrong neighbors are still close to u , the loss will still pull v closer to the neighborhood of u (another augmentation of v). In later stages of learning, when the representation is more mature, the other $k - 1$ neighbors will be semantically related and will contribute to learning since u and v are already closer to each other. Table 2-right and Figure 2 show that the representation improves as the learning progresses.

Strength of augmentation: In most exemplar-based SSL methods, augmentation plays an important role since the main supervision signal is that the augmentation should not change the embedding much. Hence, recent methods, e.g., MoCo v2, SimCLR, and BYOL, use strong augmentations. We believe such aggressive augmentations on the target embeddings u may add randomness to the learning process as some of those augmentation do not look natural, so the nearest neighbors will not be semantically close to the query image. Hence, we use weaker augmentations for the target model to make u and z less noisy while still using strong augmentations for the online model. We refer to this as the weak/strong (“w/s”) variation. This is inspired by [40] which uses weak augmentations in semi-supervised learning. This variation results in almost one point improvement over the regular variation where both encoders use strong augmentations. As shown in Fig. 4 (right), the nearest neighbors are more pure in the “weak only” setting which is consistent with our above intuition. Our experiments show that BYOL also benefits from w/s augmentation to some extent. This is probably due to more robust target encoding. Finally, we explore a weak/weak “w/w” variation where both views are augmented with a weak augmentation.

3. Experiments

We report the results of our self-supervised learning and transfer learning in this section. We use PyTorch library for all experiments.

Mem. Size	kNN Time	kNN GFLOPS	NN	20-NN	Transfer Mean
1M	6.78%	1.05	62.0	64.9	75.5
128K	0.72%	0.13	62.0	65.2	76.3

Table 1: Additional computational cost of finding NNs: Forwarding through each ResNet50 encoder needs 4.14 GFLOPS, so finding NNs adds a small cost. Note that any method like MoCo that uses a memory bank needs this additional cost.

3.1. Self-supervised Learning

Mean Shift (MSF): We use 0.99 for the the momentum of the target encoder, top- $k = 5$, and 1.024M for the memory bank size (which is roughly the same as the size of ImageNet dataset). Our ablation study shows that a memory bank of 128K does not degrade the results. We observe that the added computational cost of NN search is small compared to the overall forward and backward passes. We find that MSF with 128K memory bank size and 512 dimensions for the embedding, uses less than 0.5GB of extra GPU memory for the memory bank and less than 1% of extra computation for finding 5 NNs (see Table 1).

BYOL-asym (baseline): Training SSL methods for more than 200 epochs is not easy due to resource constraints. For instance, training BYOL with ResNet50 for 200 epochs takes roughly 7 days on four RTX 2080-Ti GPUs. Thus, for a fair comparison, we re-implement BYOL in our own framework and call it BYOL-asym. We note and justify the major differences between BYOL-asym and BYOL here. First, we use an asymmetric loss. The original BYOL paper [22] uses symmetric loss which passes each view of the image through both encoders. As a result, the gradient is calculated over $2 \times B$ instances where B is the batch size, so each epoch needs twice computation compared to asymmetric loss. Hence, 200 epochs of BYOL-asym should be compared with 100 epochs of regular BYOL. Second, we use a small batch size of 256 instead of 4096. [22] shows that BYOL works well even with the batch size of 256. Third, we use SGD optimizer instead of LARS. Despite these differences, our implementation works reasonable well compared to reported results in prior work. Our MSF uses the same setup for fairness.

Augmentation: In all of our experiments, “strong” augmentation refers to the augmentation in MoCo v2 [12]. The strong augmentation involves the following stochastic operations: grayscale, color jitter, horizontal flip, and Gaussian blur. The “weak” augmentation is simply a random crop of size 224×224 with area ratio between 0.2 and 1.0 followed by random horizontal flipping with probability 0.5. **MSF w/s** refers to our “weak/strong” variation where the target encoder view is augmented with the weak augmenta-

tion and online encoder view is augmented with the strong augmentation. **MSF w/w** refers to our “weak/weak” variation where both teacher and student views use weak augmentation. **BYOL-asym** and **MSF** use the standard SSL practice of augmenting both views with the strong augmentation.

Architecture: We generally follow [22] for the architectures of both BYOL-asym and MSF. We use the ResNet50 [24] model as backbone in all our experiments. A projection layer (2 layer MLP) is added on top of the backbone. The first layer expands the feature channels from 2048 to 4096. It is followed by BatchNorm and ReLU layers. The final linear layer reduces the feature channels from 4096 to 512. The prediction layer architecture is the same as projection layer except its first layer expands the channels from 512 to 4096. After the pre-training step, online encoder’s backbone is evaluated by removing the projection and prediction layers.

Training: For both BYOL-asym and MSF, we use the SGD (lr=0.05, momentum=0.9, and weight decay=1e-4) optimizer and train for 200 epochs. Learning rate uses cosine scheduler.

3.2. Evaluation on ImageNet

Evaluation on full ImageNet. We evaluate the representations of the pre-trained model by training linear and nearest neighbor (NN) classifiers. We use the code provided by [5] for training both classifiers. A single linear layer is trained on top of a frozen backbone. The features from the backbone are normalized to have unit ℓ_2 norm and then scaled and shifted to have zero mean and unit variance for each dimension. The linear layer is trained with SGD (lr=0.01, epochs=40, batch size=256, weight decay=1e-4, and momentum=0.9). Learning rate is multiplied by 0.1 at 15 and 30 epochs. We use standard supervised ImageNet augmentations [2] during training. For nearest neighbor, we pre-process train and val ImageNet sets with center crop (size 256) augmentation and calculate ℓ_2 normalized embeddings by forwarding through the backbone. We report Top-1 accuracies on ImageNet val set for linear, 1-NN, and 20-NN classifiers in Table 2.

Evaluation on smaller ImageNet: Similar to [25, 11, 22, 5], we evaluate the pre-trained models on the task of classification with limited ImageNet labels. The training details are the same as above except the training dataset sizes are reduced to 1% and 10% of the train set of ImageNet [39]. The results are reported in Table 3.

3.3. Evaluation on Transfer Learning

Linear classification: Following the procedure outlined in [11, 22], we evaluate the self-supervised pre-trained models for linear classification task on following datasets: Food101 [8], SUN397 [49], CIFAR10 [29], CIFAR100

Method	Ref.	Batch Size	Epochs	Sym. Loss 2x FLOPS	Top-1 Linear	NN	20-NN
Supervised	[4]	256	100	-	76.2	71.4	74.8
Random-init	-	-	-	-	5.1	1.5	2.0
SeLa-v2 [52]	[10]	4096	400	✓	67.2	-	-
SimCLR[11]	[11]	4096	1000	✓	69.3	-	-
SwAV [10]	[10]	4096	400	✓	70.1	-	-
DeepCluster-v2 [9]	[10]	4096	400	✓	70.2	-	-
SimSiam [13]	[13]	256	400	✓	70.8	-	-
MoCo v2 [23]	[13]	256	400	✓	71.0	-	-
MoCo v2 [23]	[12]	256	800	✗	71.1	57.3	61.0
CompRes [†] [5]	[5]	256	1K+130	✗	71.9	63.3	66.8
InvP	[46]	256	800	✗	71.3	-	-
BYOL [22]	[22]	256	300	✓	71.8	-	-
BYOL [22]	[22]	4096	1000	✓	74.3	62.8	66.9
SwAV [‡] [10]	[10]	4096	800	✓	75.3	-	-
SimCLR[11]	[13]	4096	200	✓	68.3	-	-
SwAV [10]	[13]	4096	200	✓	69.1	-	-
MoCo v2 [23]	[13]	256	200	✓	69.9	-	-
SimSiam [13]	[13]	256	200	✓	70.0	-	-
BYOL [22]	[13]	4096	200	✓	70.6	-	-
MoCo v2 [23]	[12]	256	200	✗	67.5	50.9	54.3
CO2 [47]	[47]	256	200	✗	68.0	-	-
BYOL-asym [22]	-	256	200	✗	69.3	55.0	59.2
ISD [42]	[42]	256	200	✗	69.8	59.2	62.0
MSF	-	256	200	✗	71.4	60.6	64.0
MSF w/s	-	256	200	✗	72.4	62.0	64.9
MSF w/s (128K)	-	256	200	✗	72.1	62.0	65.2
SimCLR w/w [11]	[22]	4096	300	✓	40.2	-	-
BYOL w/w [22]	[22]	4096	300	✓	60.1	-	-
MSF w/w	-	256	200	✗	66.3	54.6	57.4

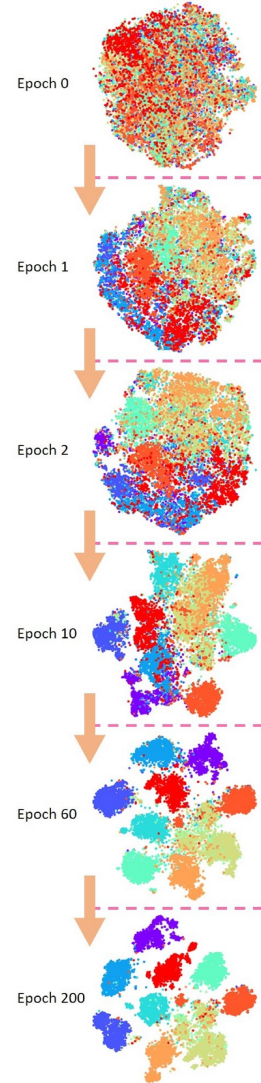


Table 2: Left: Evaluation on full ImageNet: We compare our model on the full ImageNet linear and nearest neighbor benchmarks using ResNet50. We find that given similar computational budget, our models are better than other state-of-the-art methods. Our w/s variation works slightly better than the regular MSF. Interestingly, when using weak augmentations only, our method (MSF w/w) outperforms BYOL and SimCLR with a large margin. We believe this is important in some applications, e.g. medical domain, where augmentation engineering is not easy. Note that methods with symmetric loss are not directly comparable with the other ones as they need to feed each image twice though each encoder. This results in twice the computation for each mini-batch. One may argue that a non-symmetric BYOL with 200 epochs should be compared with symmetric BYOL with 100 epochs only as they use similar amount of computation. Note that symmetric MoCo v2 with 400 epochs is almost the same as asymmetric MoCo v2 with 800 epochs (71.0 vs. 71.1). Note that the accuracy of Random-init for ResNet50 is much lower than AlexNet (14.1% on conv5 layer from [35]). [†]: CompRes is not directly comparable as it uses ResNet50 distilled from a larger SSL teacher model (SimCLR-ResNet50x4). [‡]: SwAV is not comparable as it uses multiple crops together. **Right: Epochwise t-SNE for MSF:** We visualize the ℓ_2 normalized features for 10 random ImageNet classes at certain epochs of MSF training. We find that over the period of training, semantic clusters are formed in the feature space.

[29], Cars [28], Aircraft [30], Flowers [32], Pets [36], Caltech-101 [20] and DTD [15]. The appendix includes more details on the datasets and training. The results are reported in Table 4. To verify our implementation, we evaluate the official 1000-epoch BYOL weights provided in [1] and compare with the results from [22] in Table 4.

Object detection: Following the procedure outlined in [23], we use Faster-RCNN [38] for the task of object detection on PASCAL-VOC [19]. We use the code provided at [3] with default parameters. All the weights are finetuned on the trainval07+12 set and evaluated on the test07 set. We report an average over 5 runs in Table 5.

Method	Fine-tuned	Epochs	Top-1		Top-5	
			1%	10%	1%	10%
Supervised	✓		25.4	56.4	48.4	80.4
PIRL [31]	✓	800	-	-	57.2	83.8
CO2 [47]	✓	200	-	-	71.0	85.7
SimCLR [11]	✓	1000	48.3	65.6	75.5	87.8
InvP [46]	✓	800	-	-	78.2	88.7
BYOL [22]	✓	1000	53.2	68.8	78.4	89.0
SwAV [†] [10]	✓	800	53.9	70.2	78.5	89.9
MoCo v2 [12]	✗	800	51.5	63.6	77.6	86.1
BYOL [‡] [22]	✗	1000	55.7	68.6	80.0	88.6
CompRes [*] [5]	✗	1K+130	59.7	67.0	82.3	87.5
MoCo v2 [12]	✗	200	43.6	58.4	71.2	82.9
BYOL-asym	✗	200	47.9	61.3	74.6	84.7
ISD [42]	✗	200	53.4	63.0	78.8	85.9
MSF	✗	200	53.5	65.2	78.1	86.4
MSF w/s	✗	200	55.5	66.5	79.9	87.6

Table 3: Evaluation on small labeled ImageNet: We compare our model on the ImageNet 1% and 10% linear evaluation benchmarks for ResNet50. The column “Fine-tuned” refers to whether the full network was fine-tuned or a single linear layer was trained. Given similar computational budgets, both of our models are better than other state-of-the-art methods. We evaluate BYOL and MoCo v2 with our evaluation framework and interestingly, realize that BYOL performs better in linear evaluation compared to finetuning the whole network on the 1% split. We report these numbers for fair comparison. * is ResNet50 compressed from SimCLR-R50x4. † uses a different augmentation strategy than others. ‡ is our evaluation with the official weights [1].

3.4. Ablation study.

Here, we study the effect of MSF hyperparameters and design choices like augmentation strategies, top- k , and memory bank size. We use ResNet50 and train it with ImageNet. In all experiments, we use the default MSF w/s variant and only vary the parameter of interest.

Same view of an instance to both encoders: One may argue that the mean-shift grouping and using different views of the same instance for different encoders are orthogonal ideas, and mean-shift alone might work. We did an experiment by feeding the same augmented view to both encoders ($T1 = T2$) and realized that the model does not learn. It collapses in the first epoch. Hence, we believe using different views is still an important inductive bias.

Effect of k in top- k : This section shows the effect of sampling different top- k nearest neighbors. We use k values from set $\{2, 5, 10, 20, 50\}$. We use $k = 5$ for main experiments, but $k = 10$ improves NN by 0.5 point. Note that setting $k = 1$, makes MSF identical to BYOL. Results are in the Table 6. Additionally, we plot the purity for each experiment in Figure 4. Purity for a single query is the percentage of the samples 2 to k in the top- k nearest neighbors (exclud-

ing u itself) which have the same class as the query. Final purity is calculated by averaging the purities of all samples. One may study the effect of increasing k gradually during iterations as a future extension.




Figure 3: Memory Bank Size: On ImageNet, we do not see an improvement in increasing the memory bank size beyond 128K which needs only 0.5GB of GPU memory.

Size of memory bank: CompRes [5] shows that a large memory bank is important to accurately capture the neighborhood of a random sample in the embedding space. Thus, we vary the size of the memory bank from 256 to 1M to evaluate if larger memory bank can help with more accurate nearest neighbors. Results are in Figure 3. Although our main experiment uses 1M sized memory bank, we find that 128K works equally well. Note that the size of memory bank also depends on the training dataset size.

Comparison of different augmentation strategies: Table 7 shows results for BYOL and MSF with different augmentation strategies. Comparing “s/s” variants with “w/s”, we find that BYOL receives a very small boost from the “w/s” variant while MSF improves consistently by ≈ 1 point on all three benchmarks. We believe this is due to better purity of the nearest neighbors while training (also shown in Fig. 4 (right)). Further, we observe that MSF w/w is significantly better as compared to BYOL w/w. This can be attributed to the nearest neighbors serving as a proxy for strong augmentation.

4. Related Work

Self-supervised learning: The aim of self-supervised learning is to learn representations directly from the data without using any manual data annotation. Specifically, a pretext task is designed based on the inherent structure in the data and a model is trained to solve it. Various pretext tasks have been designed that exploit different structural cues in the data. Here, we focus on following pretext tasks for images: treating each data point as a single class to perform instance level classification [18], predicting the relative location of patches [17, 33], filling up a missing patch in an image [37], predicting a colored image from its grayscale version [53, 54], counting objects in an image [34], predicting the rotation of an image [21], and predict-

Method	Ref.	Epochs	Food 101	CIFAR 10	CIFAR 100	SUN 397	Cars 196	Aircraft	DTD	Pets	Caltech 101	Flowers 102	Mean
Supervised	[22]		72.3	93.6	78.3	61.9	66.7	61.0	74.9	91.5	94.5	94.7	78.9
SimCLR [11]	[22]	1000	72.8	90.5	74.4	60.6	49.3	49.8	75.7	84.6	89.3	92.6	74.0
MoCo v2 [12]	-	800	72.5	92.2	74.6	59.6	50.5	53.2	74.4	84.6	90.0	90.5	74.2
BYOL [22]	[22]	1000	75.3	91.3	78.4	62.2	67.8	60.6	75.5	90.4	94.2	96.1	79.2
BYOL [22]	rep.	1000	75.4	92.7	78.1	62.1	67.1	62.0	76.8	89.8	92.2	95.5	79.2
BYOL-asym [22]	-	200	70.2	91.5	74.2	59.0	54.0	52.1	73.4	86.2	90.4	92.1	74.3
MoCo v2 [12]	-	200	70.4	91.0	73.5	57.5	47.7	51.2	73.9	81.3	88.7	91.1	72.6
MSF	-	200	70.7	92.0	76.1	59.0	60.9	53.5	72.1	89.2	92.1	92.4	75.8
MSF-w/s	-	200	71.2	92.6	76.3	59.2	55.6	53.7	73.2	88.7	92.7	92.0	75.5
MSF-w/s (128K)	-	200	72.3	92.7	76.3	60.2	59.4	56.3	71.7	89.8	90.9	93.7	76.3

Table 4: Linear layer transfer learning evaluation: We compare various SSL methods on transfer tasks by training linear layers. Under similar computational budgets, we show that our models are consistently better or on par with other state-of-the-art methods. Only a single linear layer is trained on top of features. No train time augmentations are used. “rep.” means we have reproduced the results using our evaluation framework for better comparison.




Figure 4: Comparison of purity: Ideally, we want all k nearest neighbors to be from the same category as the input image, so we calculate the percentage of correct neighbors for each input, average it over all images, and call it purity. We exclude the first NN which is identical to u , so the accuracy is over $k-1$ neighbors. We find this metric to be very handy in evaluating the model at the training time since it comes almost for free. (left) shows the purity for different k values with respect to epoch number while (middle) compares the purity of “w/s” variation with the regular “s/s” variation for $k=5$. Purity is higher for “w/s” variation which is consistent with our intuition. In (right), at each epoch of our MSF w/s (top- $k=10$) model, we calculate purity for the target encoder using either only weak (orange) or only strong (blue) augmentations. We see that strong setting has lower purity. This suggests that the stronger augmentation makes the nearest neighbors more noisy. This is aligned with our intuition for using weak augmentation for the target model in w/s variation.

ing the pseudo-labels obtained from clustering [9, 52]. Note that designing the pretext task or the augmentations itself is still manual and needs domain knowledge.

Instance discrimination: Recently, the task of instance discrimination [18] has shown great promise. The basic idea is to treat each image as single class. This is also referred as contrastive learning where positive samples (augmented views of the same instance) are pushed close and away from the negative samples (all other instances). While [18] took a parametric approach for this classification, [48] took a non-parametric approach. The non-parametric approach has seen broad adoption with great results [23, 11, 31, 55, 43, 10]. Two important components of these methods are: memory bank (source of negative samples) and augmentation (constructing positive samples). A simple but effective technique using a momentum encoder

to populate the memory bank is proposed in [23]. A rigorous study of the impact of different augmentations and hyperparameters is conducted in [11]. Improved augmentation strategies are proposed in [10, 11, 44]. Instance discrimination can also be viewed from an information theoretic perspective as the task of maximizing the information between different views of a single image [26, 45, 6].

Consistency regularization: Although negative pairs were thought to be central in preventing the collapse of representations for instance discrimination, [22] proposed a method that does not collapse despite not using any negatives. The objective in [22, 13] simply pulls augmented views of the same image close without any contrast with negative samples. This is also referred to as consistency regularization in the semi-supervised learning framework [41]. Inspired by [22], we propose a more general form of

Method	Ref.	Epochs	AP ₅₀	AP	AP ₇₅
Sup. IN	[13]	-	81.3	53.5	58.8
Scratch	[13]	-	60.2	33.8	33.1
<i>Symmetric loss. 2x FLOPS</i>					
SimCLR	[13]	200	81.8	55.5	61.4
MoCo v2	[13]	200	82.3	57.0	63.3
BYOL	[13]	200	81.4	55.3	61.1
SwAV	[13]	200	81.5	55.4	61.4
SimSiam	[13]	200	82.4	57.0	63.7
<i>Asymmetric loss.</i>					
MoCo v2	[12]	800	82.5	57.4	64.0
InvP	[46]	800	81.8	56.2	61.5
MoCo v2	[12]	200	82.4	57.0	63.6
CO2	[47]	200	82.7	57.2	64.1
BYOL-asym	-	200	81.9	56.8	63.5
MSF	-	200	82.2	56.7	63.4
MSF w/s	-	200	82.2	56.6	63.1

Table 5: Transfer learning to PASCAL VOC object detection: We compare our models on the transfer task of object detection. We find that given a similar computational budget, our method is better than BYOL. The models are trained on the VOC trainval07+12 set and evaluated on the test07 set. We report average over 5 runs.

	$k=1$	$k=2$	$k=5$	$k=10$	$k=20$	$k=50$
NN	55.8	61.0	62.0	62.5	62.0	61.5
20-NN	59.1	64.2	64.9	65.7	65.4	64.9

Table 6: Effect of k in top- k : Our study shows that MSF is not very sensitive to k . While $k = 10$ performs the best, we report the main results for $k = 5$. Note that setting $k = 1$ makes our method identical to BYOL-asym.

it where the positives can also come from the close neighborhood of a sample grouping similar images together.

Clustering methods: Another class of methods based on clustering have shown promise. The basic idea is to alternate between clustering and learning the representations [50, 51]. This approach was first scaled to large scale pre-training in [17]. A big concern in these methods is to prevent the collapse of all representations into a single cluster. To that end, an optimal transport based formulation of clustering is proposed in [52]. An online clustering algorithm based on the formulation of [52] is proposed in [10].

Clustering and instance discrimination: Clustering methods can be seen as the generalizing the instance discrimination framework. Only the views from the same sample can be positives in instance discrimination [18, 48], but all the members (and their views) of a cluster are positives in clustering based methods [7, 10]. A more flexible middle ground is where the set of positives are based on the

Method	Aug.	Top-1	NN	20-NN
BYOL-asym	s/s	69.3	55.0	59.2
BYOL-asym	w/s	69.5	55.8	59.1
BYOL [†] [22]	w/w	60.1	-	-
MSF	s/s	71.4	60.6	64.0
MSF	w/s	72.4	62.0	64.9
MSF	w/w	66.3	54.6	57.4

Table 7: Comparison augmentations strategies: In s/s, both views are strongly augmented while in w/w they are weakly augmented. w/s refers to weak augmentation for target view and strong augmentation for the online view. w/s improves our method more compared to BYOL. This may be due to more pure nearest neighbors. In w/w setting, MSF is significantly better than BYOL as the nearest neighbors can be a good substitute for strong augmentation. [†]: uses 4096 batch size, 300 epochs, and symmetric loss.

local neighborhood of a sample: top- k nearest neighbors of a sample in [27], nearest neighbors of a sample that are also the members of the same cluster in [55], and top- k graph distance based neighbors in [46]. Our method shares the motivation behind these works: embeddings should be locally clustered around high density regions. We also use top- k nearest neighbors as positives [27, 46], but our method is fundamentally different as there is no notion of negatives in our method. Intuitively, we enforce a more simpler and flexible constraint: move each sample closer to its nearest neighbors in each iteration. This idea is inspired from Mean-Shift Clustering [14, 16] where the cluster assignment for each sample is iteratively updated to be the mean of its nearest neighbors. In contrast to the k -means clustering, Mean-Shift does not make strong assumptions about the shape of clusters.

5. Conclusion

We introduce a simple but effective SSL method based on grouping similar images together in an online fashion. We simply shift the embedding of an image towards the mean of its nearest neighbors. MSF with $k = 1$ is identical to BYOL so MSF can be seen as a generalized form of BYOL. Our extensive experiments show that MSF performs better or on-par compared to state-of-the-art SSL methods on various tasks including ImageNet linear evaluation.

Acknowledgment: This material is based upon work partially supported by the United States Air Force under Contract No. FA8750-19-C-0098, funding from SAP SE, and also NSF grant numbers 1845216 and 1920079. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Air Force, DARPA, or other funding agencies.

References

- [1] Code and weights for byol. <https://github.com/deepmind/deepmind-research/tree/master/byol>. 5, 6
- [2] Official pytorch supervised imagenet training code. <https://github.com/pytorch/examples/blob/master/imagenet/main.py>. 4
- [3] Pytorch implementation of moco: <https://arxiv.org/abs/1911.05722>. <https://github.com/facebookresearch/moco>. 5
- [4] Torchvision models. <https://pytorch.org/docs/stable/torchvision/models.html>. 5
- [5] Soroush Abbasi Koohpayegani, Ajinkya Tejankar, and Hamed Pirsiavash. Compress: Self-supervised learning by compressing representations. *Advances in Neural Information Processing Systems*, 33, 2020. 4, 5, 6
- [6] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15535–15545, 2019. 7
- [7] Miguel A Bautista, Artsiom Sanakoyeu, Ekaterina Tikhoncheva, and Bjorn Ommer. Cliqecnn: Deep unsupervised exemplar learning. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. 8
- [8] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014. 4
- [9] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 132–149, 2018. 1, 5, 7
- [10] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems*, pages 9912–9924. Curran Associates, Inc., 2020. 1, 5, 6, 7, 8
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 1, 4, 5, 6, 7
- [12] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 4, 5, 6, 7, 8
- [13] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020. 5, 7, 8
- [14] Yizong Cheng. Mean shift, mode seeking, and clustering. *IEEE transactions on pattern analysis and machine intelligence*, 17(8):790–799, 1995. 8
- [15] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Computer Vision and Pattern Recognition*, 2014. 5
- [16] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603–619, 2002. 8
- [17] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1422–1430, 2015. 6, 8
- [18] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in neural information processing systems*, pages 766–774, 2014. 6, 7, 8
- [19] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 5
- [20] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Pattern Recognition Workshop*, 2004. 5
- [21] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018. 6
- [22] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020. 1, 2, 4, 5, 6, 7, 8
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 1, 2, 5, 7
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [25] Olivier J Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019. 4
- [26] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2019. 7
- [27] Jiabo Huang, Qi Dong, Shaogang Gong, and Xiatian Zhu. Unsupervised deep learning by neighbourhood discovery. In *International Conference on Machine Learning*, pages 2849–2858. PMLR, 2019. 8
- [28] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3D object representations for fine-grained categorization. In *Workshop on 3D Representation and Recognition*, Sydney, Australia, 2013. 5

- [29] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 4, 5
- [30] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 5
- [31] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. *arXiv preprint arXiv:1912.01991*, 2019. 6, 7
- [32] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, 2008. 5
- [33] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016. 6
- [34] Mehdi Noroozi, Hamed Pirsiavash, and Paolo Favaro. Representation learning by learning to count. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5898–5906, 2017. 6
- [35] Mehdi Noroozi, Ananth Vinjimoor, Paolo Favaro, and Hamed Pirsiavash. Boosting self-supervised learning via knowledge transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9359–9367, 2018. 5
- [36] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar. Cats and dogs. In *Computer Vision and Pattern Recognition*, 2012. 5
- [37] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016. 6
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, volume 28, 2015. 5
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 4
- [40] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [41] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017. 7
- [42] Ajinkya Tejankar, Soroush Abbasi Koohpayegani, Vipin Pillai, Paolo Favaro, and Hamed Pirsiavash. Isd: Self-supervised learning by iterative similarity distillation, 2020. 1, 5, 6
- [43] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 7
- [44] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? In *Advances in Neural Information Processing Systems*, volume 33, pages 6827–6839. Curran Associates, Inc., 2020. 7
- [45] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2018. 7
- [46] Feng Wang, Huaping Liu, Di Guo, and Sun Fuchun. Unsupervised representation learning by invariance propagation. In *Advances in Neural Information Processing Systems*, volume 33, pages 3510–3520. Curran Associates, Inc., 2020. 5, 6, 8
- [47] Chen Wei, Huiyu Wang, Wei Shen, and Alan Yuille. Co2: Consistent contrast for unsupervised visual representation learning. *arXiv preprint arXiv:2010.02217*, 2020. 5, 6, 8
- [48] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018. 7, 8
- [49] Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Computer Vision and Pattern Recognition*, 2010. 4
- [50] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML’16*, page 478–487. JMLR.org, 2016. 8
- [51] Jianwei Yang, Devi Parikh, and Dhruv Batra. Joint unsupervised learning of deep representations and image clusters. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 8
- [52] Asano YM., Rupprecht C., and Vedaldi A. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations*, 2020. 1, 5, 7, 8
- [53] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016. 6
- [54] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1058–1067, 2017. 6
- [55] Chengxu Zhuang, Alex Lin Zhai, and Daniel Yamins. Local aggregation for unsupervised learning of visual embeddings. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6002–6012, 2019. 7, 8

Appendix

Dataset	Classes	Train samples	Val samples	Test samples	Accuracy measure	Test provided
Food101 [8]	101	68175	7575	25250	Top-1 accuracy	-
CIFAR-10 [29]	10	49500	500	10000	Top-1 accuracy	-
CIFAR-100 [29]	100	45000	5000	10000	Top-1 accuracy	-
Sun397 (split 1) [48]	397	15880	3970	19850	Top-1 accuracy	-
Cars [28]	196	6509	1635	8041	Top-1 accuracy	-
Aircraft [30]	100	5367	1300	3333	Mean per-class accuracy	Yes
DTD (split 1) [15]	47	1880	1880	1880	Top-1 accuracy	Yes
Pets [36]	37	2940	740	3669	Mean per-class accuracy	-
Caltech-101 [20]	101	2550	510	6084	Mean per-class accuracy	-
Flowers [32]	102	1020	1020	6149	Mean per-class accuracy	Yes

Table A1: We list the sizes of train, val, and test splits of the transfer datasets. **Test split:** We use the provided test sets for Aircraft, DTD, and Flowers datasets. In case of Sun397, Cars, CIFAR-10, CIFAR-100, Food101, and Pets datasets, we use the provided val set as the hold-out test set. In case of Caltech-101, we use a random split of 30 images per category as the hold-out test set. **Val split:** We use the provided val sets for the datasets DTD and Flowers. For all other datasets, the val set is created by randomly sampling a subset of the train set. In order to be as close to BYOL [22] transfer setup as possible, we use the following val set splitting strategies for each dataset. Aircraft: 20% samples/class. Caltech-101: 5 samples/class. Cars: 20% samples/class. CIFAR-100: 50 samples/class. CIFAR-10: 50 samples/class. Food101: 75 samples/class. Pets: 20 samples/class. Sun397: 10 samples/class.




Figure A1: Nearest neighbors (NN) of the model at each epoch: Similar to Figure 2.




Figure A2: Nearest neighbors (NN) of the model at each epoch: Similar to Figure 2.



Figure A3: Random Clusters: We forward ImageNet training set through our ResNet-50 model and cluster them into 1000 clusters using k-means. We select 30 clusters randomly and show 20 randomly sampled images from each cluster without cherry-picking. Each row corresponds to a cluster. Note that semantically similar images are clustered together.