



HAL
open science

How does the degree of novelty impacts semi-supervised representation learning for novel class retrieval?

Quentin Leroy, Olivier Buisson, Alexis Joly

► To cite this version:

Quentin Leroy, Olivier Buisson, Alexis Joly. How does the degree of novelty impacts semi-supervised representation learning for novel class retrieval?. 2023. hal-03871584

HAL Id: hal-03871584

<https://inria.hal.science/hal-03871584>

Preprint submitted on 17 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How does the degree of novelty impacts semi-supervised representation learning for novel class retrieval?

QUENTIN LEROY, Inria, France

OLIVIER BUISSON, Institut National de l'Audiovisuel, France

ALEXIS JOLY, Inria, France

Supervised representation learning with deep networks tends to overfit the training classes and the generalization to novel classes is a challenging question. It is common to evaluate a learned embedding on held-out images of the same training classes. In real applications however, data comes from new sources and novel classes are likely to arise. We hypothesize that incorporating unlabelled images of novel classes in the training set in a semi-supervised fashion would be beneficial for the efficient retrieval of novel-class images compared to a vanilla supervised representation. To verify this hypothesis in a comprehensive way, we propose an original evaluation methodology that varies the degree of novelty of novel classes by partitioning the dataset category-wise either randomly, or semantically, i.e. by minimizing the shared semantics between base and novel classes. This evaluation procedure allows to train a representation blindly to any novel-class labels and evaluate the frozen representation on the retrieval of base or novel classes. We find that a vanilla supervised representation falls short on the retrieval of novel classes even more so when the semantics gap is higher. Semi-supervised algorithms allow to partially bridge this performance gap but there is still much room for improvement.

1 INTRODUCTION

Deep neural networks have been the de facto algorithm for processing and training data for a variety of tasks and data modalities whether it is image [35, 38], text [12, 49], speech [20, 55], video [1–3, 47]. In this study we focus on image data and retrieval tasks. The major impediment is the sheer amount of annotated images necessary to achieve this feat [40]. Several lines of research have been tackling this issue and have brought new ideas helping to reduce the amount of annotation necessary: active learning refers to algorithms that select most effectively the unlabeled samples for an annotator to annotate [42], few-shot learning refers to algorithms that adapt most rapidly to new classes with a few annotated support samples [50, 52], self-supervised learning refers to algorithms that learn from the samples themselves without any annotation [13, 14, 25], semi-supervised learning refers to algorithms that effectively leverage unlabeled data while also training with supervised data [7, 30, 37, 48].

Deep networks provide a strong baseline for image representation in image retrieval applications [5] [4], whether it is visual instance retrieval or more broadly visual category retrieval. Image retrieval, a long tradition in the computer vision community, is the task of retrieving those images in the search corpus that depict the same visual *instance* as the query. Visual object category retrieval evaluates more broadly the retrieving of those images containing the same visual *object* as the query, allowing for a lot more appearance variation between the query and the correct images.

In this paper, we are interested in the retrieval of an object category unseen beforehand in the training corpus used to train the deep embedding. We term this challenging task *novel class retrieval*; it aims to reflect the setting of real-world applications where incoming image data contains visual content novel compared to curated image corpora used for training deep networks. This setting exists in various applications. Take for example a corpus of digitized material of a cultural institution containing a large variety of historical photographs, drawings, paintings and manuscripts of various eras. Choosing a generic embedding for these original materials for efficient retrieval is a challenging question. These

Authors' addresses: Quentin Leroy, Inria, Montpellier, France, quentin.leroy@inria.fr; Olivier Buisson, Institut National de l'Audiovisuel, Bry-sur-Marne, France, obuisson@ina.fr; Alexis Joly, Inria, Montpellier, France, alexis.joly@inria.fr.

novel materials can be queried by users in unanticipated ways. Understanding these new types of material without additional human annotation is a challenging problem.

We focus on the question: What learning algorithm allows for the most efficient retrieval of novel classes? We devise an original methodology for comparing deep embeddings towards the task of novel class retrieval; We partition an image corpus category-wise, keeping a set of base classes for training the deep embedding and leaving aside a set of novel classes for evaluating the embedding. We get a strict disjoint label space between the training classes and the evaluation classes. The base and novel sets are further split into a training and a test set. More importantly, in each instance of our experiments we carry out two kinds of category splits: a *random split* partitioning the labeled space at random and a more carefully crafted *semantic split* that minimizes the information overlap between base and unseen categories to the extent possible (details in section 3).

Our contributions are summarised as follows: 1) We devise an evaluation methodology to evaluate a deep embedding on the task of novel class retrieval, 2) we simulate new visual content by hiding supervision from the deep network and we are able to vary the degree of novelty of novel classes and 3) we compare several deep learning algorithms from different paradigms (supervised, self-supervised, semi-supervised) on the task of novel class retrieval based on three different datasets (CIFAR10, CIFAR100, In).

The remainder of the paper is organized as follows: the next section reviews important lines of research related to our evaluation methodology. In section 3 we detail how we manage to construct a benchmark relevant to novel class retrieval by masking whole categories from the training process and varying the degree of novelty. Next we present the algorithms we implemented and the different datasets and category splits used in the experiments. Finally we present our experiment results and analysis.

2 RELATED WORK

Novel Class Retrieval. We aim to benchmark training algorithms for the task of retrieving samples of classes unseen during training. It is similar to Novel class discovery [19, 56] where the goal is to cluster the whole set of images of novel classes: instead of clustering we evaluate with retrieval, we do not need to know the number of novel classes or to estimate it. The goal in Few-shot learning [36, 43, 50, 52] is to adapt most rapidly to novel classes with a small set of support supervised samples of those novel classes: in contrast we evaluate class retrieval, that is returning the most images of the class of a given single image query. Metric learning [6, 27, 29, 33, 41, 45] benchmarks usually report performance in retrieval tasks of novel classes; datasets for those benchmarks [27, 45, 51] also have class-disjoint training and testing sets but do not keep a split along the image axis as we do (See Fig. 1); in the testing phase we report retrieval performance on test images of base classes and novel classes among both base and novel-class images while metric learning does not keep test samples of base classes during testing (which prevents from evaluating the performance decrease compared to base classes).

Semi-Supervised Setting. It is often pointed out that labels are expensive to obtain and that is why training data is abundant in unlabeled samples and is lacking labeled samples. Semi-Supervised learning algorithms [7] are often proclaimed to work efficiently in that setting [30, 31, 37, 48], some methods showing impressive performance with a quantity of labels as low as one per class for CIFAR10 [44], or 1% of labeled data for ImageNet [54]. We investigate an overlooked setting where the labeled pool and the unlabeled pool are class-disjoint; meaning that the unlabeled data are known to come from different classes than those seen during training. The authors in [34] noted that this should be part of an effective benchmark for evaluating semi-supervised learning algorithms. They did a set of experiments on

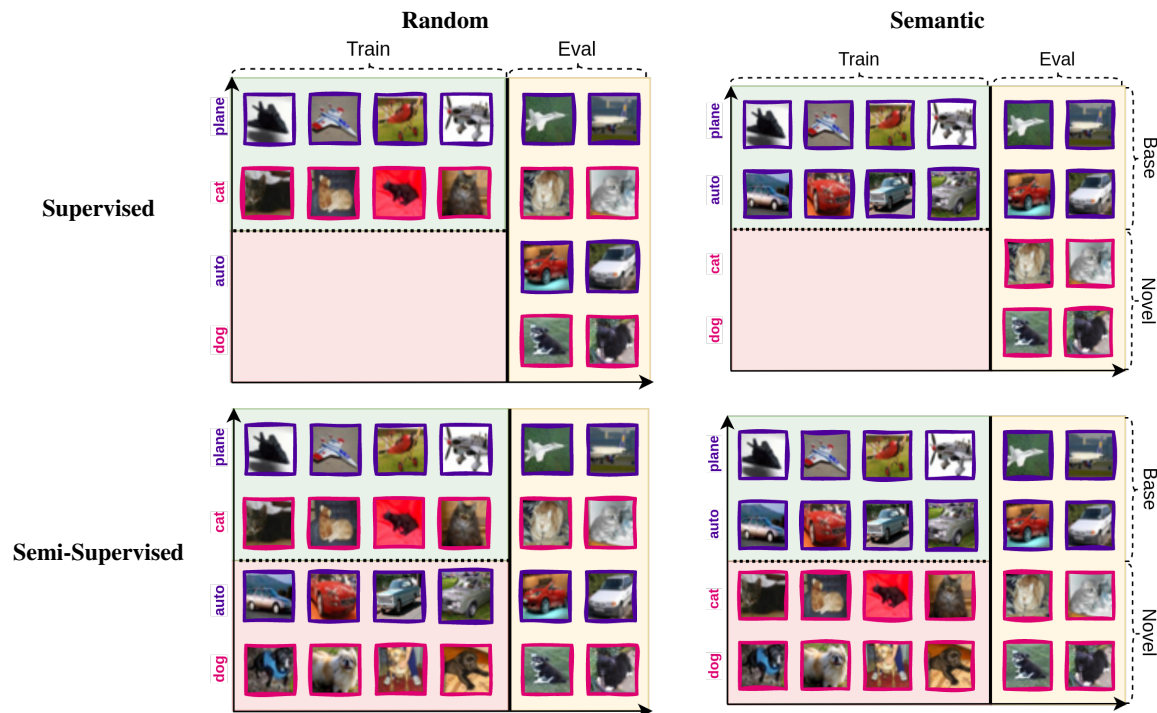


Table 1. Illustration of the two training schemes: We train algorithms with supervision on base classes (top row) and algorithms that incorporate images of the novel classes without supervision (bottom row). When the label space is partitioned randomly (left column **airplane/cat** vs. **automobile/dog**) the semantic gap between base and novel classes is narrower than when the label space is partitioned semantically (right column **airplane/automobile** vs. **cat/dog**).

CIFAR10 where they vary the amount of classes present both in labeled and unlabeled data and showed that performance can decrease in some cases. We take this idea further in the following ways: 1) we keep the labeled and unlabeled pools completely class-disjoint, 2) we experiment on CIFAR10, CIFAR100 and ImageNet100 and 3) we vary the degree of novelty between labeled and unlabeled pools by carrying a random split and a semantic split of the label space.

Partitioning the label space semantically. We experiment novel class retrieval on classification datasets: CIFAR10, CIFAR100 [28] and a subset of ImageNet [40]. For CIFAR100 we can use the super-classes which group together classes that are visually alike: we borrow the splits from [36] to devised the semantic split. For ImageNet we can use the WordNet hierarchy [15]; we chose a subset of the high-level categories used for *tiered*ImageNet [39] (details in Section 5.1).

3 PROPOSED EVALUATION METHODOLOGY: VARYING THE DEGREE OF NOVELTY

3.1 Random and Semantic Class-disjoint Data Splits

We are given an image set $I = I^T \cup I^E$ partitioned into training I^T and testing I^E sets (Fig. 1). We consider that the label space \mathcal{Y} is partitioned into two parts: a set \mathcal{Y}_b of base classes and a set \mathcal{Y}_n of novel classes. The training set can be partitioned into $I^T = \mathcal{B}^T \cup \mathcal{U}^T$ where $\mathcal{B}^T = \{(I, y); I \in I^T, y \in \mathcal{Y}_b\}$ comprises annotated images of base classes

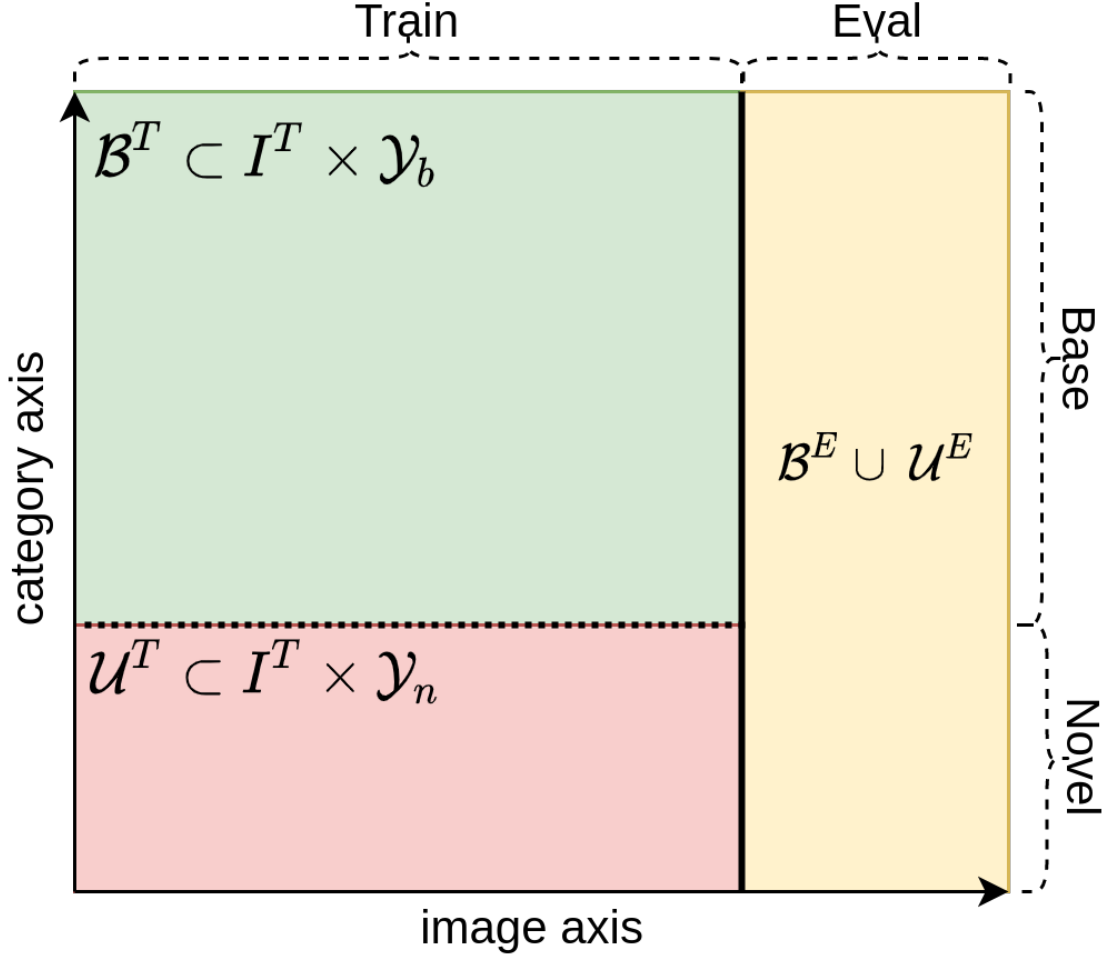


Fig. 1. We partition the label space into $\mathcal{Y} = \mathcal{Y}_b \cup \mathcal{Y}_n$, holding out a set of novel classes \mathcal{Y}_n for evaluation. Contrary to metric learning benchmarks we also keep a train/test partition along the image axis $I = I^T \cup I^E$, allowing for a comparative performance evaluation on base \mathcal{Y}_b and novel \mathcal{Y}_n classes on the test images I^E . We train a various set of algorithms with supervision only on base classes \mathcal{B}^T , and optionally adding images of novel classes of the train partition without supervision \mathcal{U}^T .

and $\mathcal{U}^T = \{(I, y); I \in I^T, y \in \mathcal{Y}_n\}$ comprises unannotated images of novel classes. The test set can be partitioned into $I^E = \mathcal{B}^E \cup \mathcal{U}^E$.

An image I is encoded as $x = \Phi_\theta(I) \in \mathbb{R}^d$, where $\Phi_\theta(\cdot)$ is a neural network whose parameters θ are trained on \mathcal{B}^T and optionally on \mathcal{U}^T . The labels of \mathcal{U}^T are never used for training the network.

In each dataset, two kinds of experiments are carried out involving either a random split or a semantic split. The random split divides randomly the set of classes. The semantic split divides the set of classes in a way that minimizes the information overlap between the base and the novel classes. Table 1 illustrates how the semantic split increases the degree of novelty compared to the random split with classes of CIFAR10. We evaluate the trained network via image retrieval by

running a similarity search among the whole testing set I^T and report R-Precision for queries from base classes \mathcal{B}^E and from novel classes \mathcal{U}^E separately (see Section 5 for the results).

3.2 Evaluation via Image Retrieval: R-Precision

We want to evaluate how well a visual representation transfer to unseen classes without any additional training, contrary to other evaluation methods that allow for additional training on the novel-class labels (linear probes [13], few-shot tasks [16], knn-classifier [53]). Evaluation via clustering requires embeddings to be clustered and is dependent on the clustering algorithm used [33]. We have chosen an evaluation metric that does not require any additional metric nor a choice of a clustering algorithm.

R-Precision. Visual object category retrieval aims to retrieve examples of the same category than a given query from a search set (I^E in our case). A function $R : I \in I^E \mapsto I_R \in \mathcal{P}(I^E)$ maps a query image to a set of similar images. The similarity between two images is evaluated by a dot product in the ℓ_2 -norm normalized feature space (also known as cosine similarity):

$$s(I_i, I_j) = \frac{\Phi_\theta(I_i)}{\|\Phi_\theta(I_i)\|_2} \cdot \frac{\Phi_\theta(I_j)}{\|\Phi_\theta(I_j)\|_2} \quad (1)$$

For an image I , R returns the list of most similar images ranked by the function s in decreasing order. Let us denote the samples of class c by $\mathcal{I}^c = \{(I_j, y_j), I_j \in I^E, y_j = c\}$, and the size of the class by $N^c = |\mathcal{I}^c|$. Image retrieval performance is usually evaluated by Recall@ k [24] (with $k = 1, 5$), in the metric learning literature [45]. We evaluate instead with Precision@ N^c where N^c is the total size of correct hits. The Precision is the proportion of correct hits among the N^c retrieved items. We set the cut-off-rank to be the size of the class query, so that the precision is equal to the recall.

4 ALGORITHMS STUDIED

In this section we present the algorithms we have implemented and tested in our experiments. We experimented with three training schemes. We took four supervised algorithms, two unsupervised algorithms and two semi-supervised algorithms. The architecture of each algorithm are illustrated in Table 2. In the following we briefly introduce each algorithm.

4.1 Supervised Algorithms

Vanilla. This model optimizes the cross-entropy loss on base classes only:

$$\mathcal{L}_V = \sum_{(I,y) \in \mathcal{B}^T} \ell(g(\Phi_\theta(I)), y). \quad (2)$$

Triplet Loss. This model optimizes the triplet loss [41] on base classes only:

$$\mathcal{L}_T = \sum_{(I_a, I_p, I_n) \in \mathcal{T}} [d_{ap} - d_{an} + \alpha]_+, \quad (3)$$

where \mathcal{T} is the set of all triplets mined from \mathcal{B}^T ; α is a margin hyperparameter, $d_{ap} = \|h(\Phi_\theta(I_a)) - h(\Phi_\theta(I_p))\|_2^2$, $d_{an} = \|h(\Phi_\theta(I_a)) - h(\Phi_\theta(I_n))\|_2^2$. In practice we optimize \mathcal{L}_T with SGD and mine *semi-hard* triplets within a batch.

Classification With Triplet Loss. This model jointly optimizes the cross-entropy loss and the triplet loss:

$$\mathcal{L}_{CwT} = \mathcal{L}_V + \mathcal{L}_T. \quad (4)$$

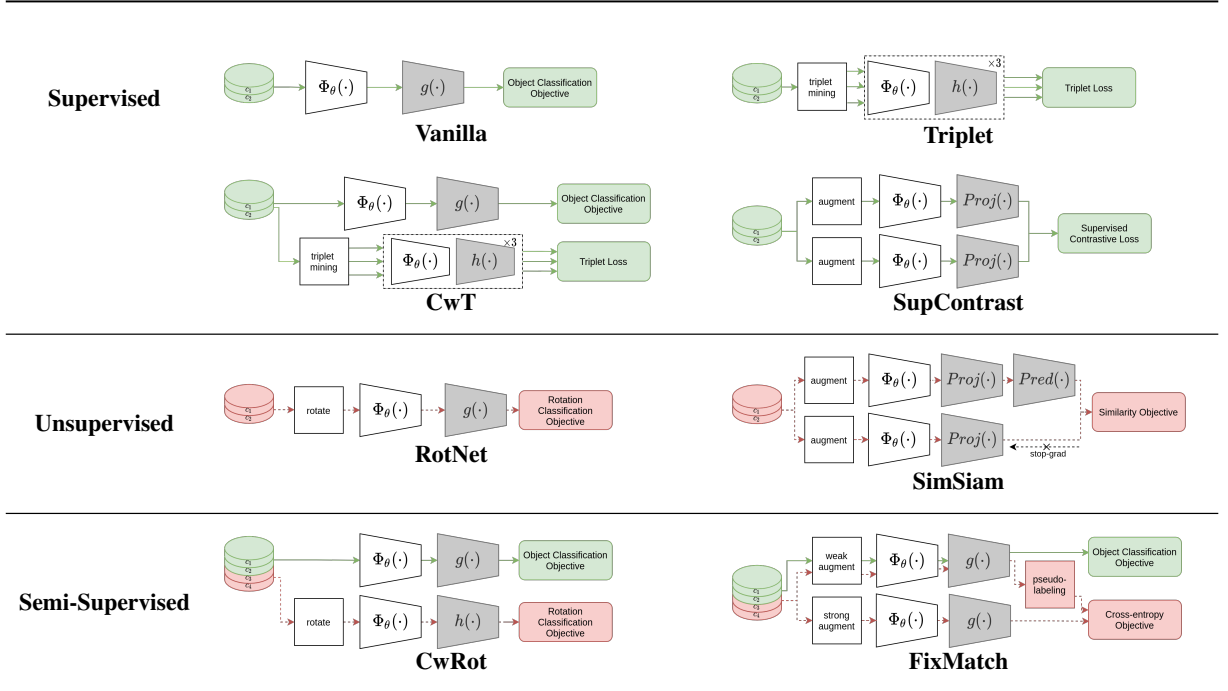


Table 2. Supervised algorithms train a network with images of base classes only (classes c_1, c_2). Unsupervised algorithms train a network with images of base classes without supervision (classes c_1, c_2). Semi-Supervised algorithms train a network with images of base classes with supervision, and images of novel classes without supervision (classes c_1, c_2 and c_3, c_4). At testing time all the heads are discarded, we take as a visual representation the output of the last average pooling layer of the backbone $\Phi_\theta(\cdot)$.

The first head g classifies into the base classes while the additional head h reduces the dimension and is used to compute the triplet loss.

Supervised Contrastive: SupContrast. This model [26] is based on a siamese architecture. It extends self-supervised contrastive methods [9, 18, 21] to the fully supervised-setting: in addition to construct positive pairs with random augmentations of the same image, it also uses the class labels to construct positive pairs. Similar to the triplet loss the contrastive loss enforces positive pairs to attract themselves and negative pairs to repel themselves. We refer the reader to the paper [26] for details.

4.2 Unsupervised Algorithms

Predicting Image Rotations: RotNet. This model [17] is an instance of a self-supervised algorithm [25]. It rotates the input image into four possible rotations, and the head g classifies into the four possible rotations. The model is trained to correctly classify the rotation applied to the input. It is a 4-class classification problem optimized with a cross-entropy loss:

$$\mathcal{L}_R = \sum_{(I, y) \in \mathcal{B}^T} \sum_{r \in \mathcal{R}} \ell(g(\Phi_\theta(I^r)), r), \quad (5)$$

where $\mathcal{R} = \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ is the set of all possible rotations and I^r refers to the corresponding rotated input.

SimSiam. This model [11] is another instance of a self-supervised algorithm [25]. It is based on a siamese architecture and is trained to maximize the similarity between two randomly augmented views of an input image. Despite being simpler than similar methods [9, 18, 21] it shows competitive results for unsupervised visual representation learning [11]. We trained our own models on base classes only. We refer the reader to the paper [11] for details.

4.3 Semi-Supervised Algorithms

FixMatch. This semi-supervised model [44] feeds an unlabeled image to a weak augmentation and a strong augmentation: the weakly augmented input is used for pseudo-labeling and the strongly augmented input is used for computing a cross-entropy loss against the pseudo-label. Note that the head g classifies into the base classes. The pseudo-labels are among the base classes and the unlabeled images that are pseudo-labeled are fed to g that classifies them into the base classes. The pseudo-labeling module is a threshold applied on the probability of the most confident class: only those images that are classified most confidently are passed to the next cross-entropy loss. FixMatch has two main hyperparameters: μ the number of unlabeled images per labeled images and τ the confidence threshold for the pseudo-labeling. We refer the reader to [44] for details.

Classification With Rotation: CwRot. We propose to train a network in a semi-supervised fashion by adding an auxiliary self-supervised loss for unsupervised novel-class samples. We append a second head to the network h tasked to predict the rotation applied to the input image as in the RotNet model [17]. This model optimizes the loss:

$$\mathcal{L}_{CwR} = \sum_{(I,y) \in \mathcal{B}^T} \ell(g(\Phi_\theta(I)), y) + \sum_{r \in \mathcal{R}} \sum_{I \in I^T} \ell(h(\Phi_\theta(I^r)), r). \quad (6)$$

The choice for the auxiliary self-supervised objective was motivated by a series of works that successfully applied it to improve few-shot classification [8, 16, 46], robustness [23], pre-training for novel class discovery [19], image generation [10, 32], semi-supervised learning [54].

This is actually a simplified version of the method proposed in [54], where the consistency regularization and pseudo-labeling components of the training scheme are removed.

5 EXPERIMENTAL RESULTS

5.1 Datasets

We experiment on CIFAR10, CIFAR100 [28] and a subset of ImageNet [40]. For each dataset, we split the label space randomly and semantically. The splits are done once and for all.

CIFAR10. For the random split, we keep 5 classes at random for base classes and the remaining 5 classes for novel classes. We got:

$$\mathcal{Y}_b = \{\text{automobile, bird, deer, frog, ship}\},$$

$$\mathcal{Y}_n = \{\text{airplane, cat, dog, horse, truck}\}.$$

Note that 2 vehicles and 3 animals ended in base classes, and also 2 vehicles and 3 animals ended in novel classes. For the semantic split, we also partitioned the label space evenly into 5 base classes and 5 novel classes; we chose the following

partitioning:

$$\mathcal{Y}_b = \{\text{airplane, automobile, bird, ship, truck}\},$$

$$\mathcal{Y}_n = \{\text{cat, dog, deer, frog, horse}\}.$$

Note that 4 vehicles and 1 animal are in the base classes, while the novel classes consist only of 5 animals and no vehicle.

CIFAR100. CIFAR100 contains 100 classes that can be grouped into 20 super-classes. For the random split, we keep 50 classes at random for base classes and the remaining 50 classes for novel classes. Note that for the random split, the super-classes are spread randomly among the base and the novel classes. For example, the *people* super-class has 3 classes in base classes (baby, girl, man) and 2 classes in novel classes (boy, woman). For the semantic split, we follow the splits of Few-Shot CIFAR100 introduced in [36]. We set the base classes to the train classes and the novel classes to the val classes + test classes. Note that for the semantic split, each super-class is entirely either in the base classes or the novel classes. For example, the *people* super-class has all the 5 classes in novel classes.

ImageNet100. We experiment on a subset of 100 classes of ImageNet [40] which we call ImageNet100 in the following. We keep 16 of the high-level categories among the 34 devised by the authors of *tieredImageNet* [39] using the WordNet hierarchy: 8 categories descending from the *artefact* synset (motor vehicle, craft, durables, garment, musical instrument, game equipment, furnishing, tool), 8 categories descending from the *animal* synset (ungulate, primate, feline, working dog, saurian, aquatic bird, insect, aquatic vertebrate); with 6-7 ImageNet classes per high-level category. For the random split we keep in base classes 3-4 classes for each high-level category at random in order to keep 50 classes, and the novel classes are the remaining 50 classes. For the semantic split, we keep the *artefact* classes for the base classes and the *animal* classes for the novel classes. Unlike *tieredImageNet* we keep all training images in the training, and validation images for testing; like *tieredImageNet* we resize the images to 84x84 resolution.

5.2 Implementation details

All networks share the same ResNet18 [22] backbone. On CIFAR10 and CIFAR100, the input images are 32x32 in resolution, we use a first convolutional layer with stride 1, not followed by a max pooling layer. For ImageNet100, the images are stored on disk at 84x84 resolution and are resized to 224x224 resolution before being fed to the network: this time we used the standard ResNet18 architecture where the first layers reduce the spatial dimension.

We reimplemented from scratch Vanilla, Triplet, RotNet, CwT and CwRot. For Vanilla, Triplet, CwT and CwRot, the networks are optimized with SGD with batch size 128, 0.9 momentum, $1e^{-4}$ weight-decay, and an initial learning rate of 0.1 dropped by 10 every 30 epochs for a total of 100 epochs. For RotNet, we followed the guidelines from the paper [17]. For Triplet and CwT, in all cases we set the margin α to 0.1 we mine *semihard* [41] negative samples, and we set the embedding dimension to 128.

For the other three algorithms we used PyTorch implementations readily available (SimSiam¹, SupContrast², FixMatch³) and made some modifications. We adapted the codebase by modifying the backbone network to match the Resnet18 used for the other baseline algorithms. For SupContrast on ImageNet100 we reduced the number of epochs to 200, decaying the initial learning rate by 10 at epochs 150, 170 and 190. For FixMatch on ImageNet100, we trained for 2^{19} iterations and used $\mu = 5$ unlabeled images per labeled image. We kept the recommended hyperparameters in every

¹<https://github.com/facebookresearch/simsiam>

²<https://github.com/HobbitLong/SupContrast>

³<https://github.com/kekmodel/FixMatch-pytorch>

other case. We would like to emphasize that we did not optimize the hyperparameters for the best performance but instead use sensible defaults, following recommendations of ResNet [22] and original papers of each algorithm.

5.3 Results and analysis

Algo	CIFAR10-Random		CIFAR10-Semantic	
	Base	Novel	Base	Novel
Vanilla	65.147	26.380	68.469	19.758
Triplet	71.479	20.512	74.890	16.760
CwT	72.336	20.391	74.389	15.438
SupContrast	74.881	28.954	79.566	22.726
RotNet	28.476	17.379	27.663	16.732
SimSiam	19.340	19.048	20.518	17.298
CwRot	72.816	37.832	77.680	35.764
FixMatch	57.807	39.238	79.178	27.419

Table 3. R-Precision on base and novel classes of CIFAR10 for a random split (left columns) and a semantic split (right columns). Top rows are supervised algorithms. Middle rows are unsupervised algorithms. Bottom rows are semi-supervised algorithms. Best R-Precision is marked in **bold**.

Algo	CIFAR100-Random		CIFAR100-Semantic	
	Base	Novel	Base	Novel
Vanilla	33.339	15.351	39.778	7.891
Triplet	33.895	10.015	43.745	3.846
CwT	48.656	10.180	57.021	4.131
SupContrast	42.981	14.094	49.000	6.678
RotNet	5.293	3.906	6.149	2.392
SimSiam	5.287	4.521	6.085	3.863
CwRot	34.257	17.525	39.962	9.886
FixMatch	38.771	22.896	47.905	9.444

Table 4. R-Precision on base and novel classes of CIFAR100 for a random split (left columns) and a semantic split (right columns). Top rows are supervised algorithms. Middle rows are unsupervised algorithms. Bottom rows are semi-supervised algorithms. Best R-Precision is marked in **bold**.

Retrieval performance is reported in Table 3 (CIFAR10), Table 4 (CIFAR100) and Table 5 (ImageNet100). Read horizontally we can compare the performance between base and novel classes. In any case the performance is degraded on novel classes, even more so for the semantic split. Read vertically we can compare the performance between the algorithms.

Supervised algorithms. The metric learning algorithms (Triplet and CwT) exhibit a better fit of the base classes to the detriment of a poorer performance on the novel classes compared to Vanilla. The benefit of Triplet over Vanilla on base classes is most noteworthy on CIFAR10 and slight on CIFAR100 and ImageNet100, indicating that the simpler and faster Vanilla algorithm is a better choice on more fine-grained dataset. The contrastive model SupContrast performs the best among the supervised baselines on base classes; on novel classes it also performs the best on CIFAR10 but not

Algo	ImageNet100-Random		ImageNet100-Semantic	
	Base	Novel	Base	Novel
Vanilla	35.202	18.351	35.388	9.931
Triplet	36.680	12.018	27.138	3.905
CwT	48.705	14.090	48.322	4.813
SupContrast	33.495	15.241	34.667	7.738
SimSiam	7.984	7.867	8.359	6.306
RotNet	6.789	7.194	8.469	5.194
CwRot	37.287	20.403	36.373	13.238
FixMatch	38.090	19.322	35.230	11.050

Table 5. R-Precision on base and novel classes of ImageNet100 for a random split (left columns) and a semantic split (right columns). Top rows are supervised algorithms. Middle rows are unsupervised algorithms. Bottom rows are semi-supervised algorithms. Best R-Precision is marked in **bold**.

on CIFAR100 and ImageNet100. The contrastive model SupContrast performs the best on base and novel classes on CIFAR10 but this does not hold on CIFAR100 and ImageNet100.

Unsupervised algorithms. RotNet and SimSiam do not show interesting results even on the base classes there were trained on; and we do not draw a consistent conclusion as for which algorithm is best in general. However we note that for the semantic split in some cases RotNet and SimSiam are better than Triplet and CwT on novel classes. For example on ImageNet100-Semantic it is actually better to use an unsupervised method than to use a metric learning method.

Semi-Supervised algorithms. The semi-supervised models mitigate the performance difference between base and novel classes and surpass the supervised baselines with a greater margin on the semantic split than on the random split. We note that CwRot performs consistently better than Vanilla on all datasets on both base and novel classes.

As for FixMatch, it is consistently better than Vanilla on novel classes. It compares more favorably to CwRot on novel classes for the random split than for the semantic split. We argue that it is due to the fact that during training it learns to classify novel-class images into the set of base classes and, because for the semantic split the novel classes are less visually alike to the base classes than for the random split, it less efficiently leverages similar visual patterns between base and novel classes. This comparison between CwRot and FixMatch for the semantic split is an evidence that semi-supervised methods such as FixMatch that rely on the fact that labeled and unlabeled images come from the same set of classes are less efficient in the setting when an entire set of images from a novel class come into play.

5.4 T-SNE embedding visualization.

We show in Table 6 some T-SNE visualization of the embeddings of CIFAR10 for four algorithms. The visualizations support the conclusions from the previous section. In particular, we can see that only the semi-supervised methods succeed in structuring the novel classes in the case of the semantic split (with CwRot providing a better class separation than FixMatch).

6 CONCLUSION

We presented a method to evaluate novel class retrieval. We argued that existing benchmarks for semi-supervised representation learning algorithms lack a setting where unlabeled data are from novel classes.

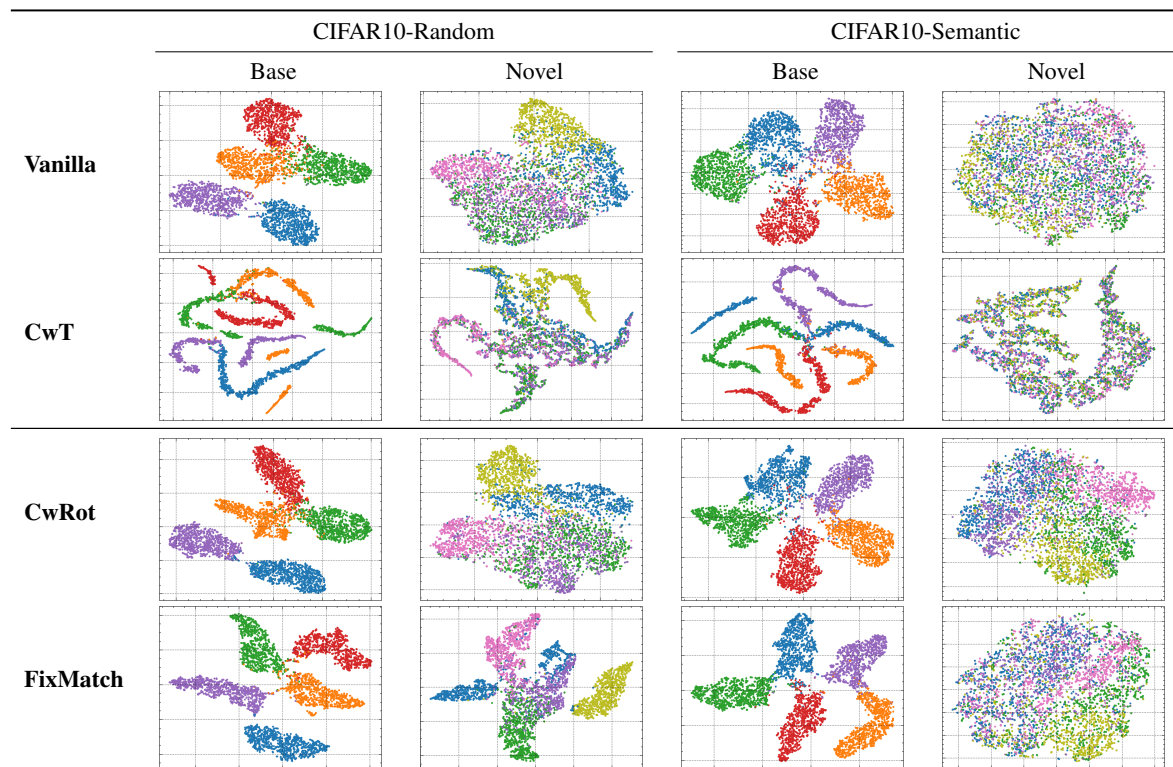


Table 6. T-SNE visualizations of the embeddings of base and novel classes for CIFAR10 for the random split (left columns) and for the semantic split (right columns). The Semi-Supervised algorithms (bottom rows) better separate the novel classes compared to the Supervised algorithms (top rows). The colors encode class membership. Best viewed in color.

We experimented with a variety of representation learning algorithms and showed evidence that semi-supervised learning algorithms mitigate the performance drop on novel classes. Yet there is still much room for improvement for the novel class retrieval task. Semi-supervised algorithms still fall short in retrieving images of novel classes in the random split setting (low-degree of novelty) and even more so in the semantic split setting (higher degree of novelty).

REFERENCES

- [1] Jean-Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-Supervised MultiModal Versatile Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A Video Vision Transformer. *arXiv:2103.15691 [cs]*, 2021.
- [3] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. SoundNet: Learning Sound Representations from Unlabeled Video. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [4] Artem Babenko and Victor Lempitsky. Aggregating Deep Convolutional Features for Image Retrieval. *arXiv:1510.07493 [cs]*, 2015.
- [5] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. Neural Codes for Image Retrieval. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2014.
- [6] Sean Bell and Kavita Bala. Learning visual similarity for product design with convolutional neural networks. *ACM Transactions on Graphics (TOG)*, 2015.
- [7] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien, editors. *Semi-Supervised Learning*. Adaptive Computation and Machine Learning. MIT Press, 2006.

- [8] Da Chen, Yuefeng Chen, Yuhong Li, Feng Mao, Yuan He, and Hui Xue. Self-Supervised Learning For Few-Shot Image Classification. *Proc. IEEE Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations. In *Proc. of the International Conf. on Machine Learning (ICML)*, 2020.
- [10] Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-Supervised GANs via Auxiliary Rotation Loss. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [11] Xinlei Chen and Kaiming He. Exploring Simple Siamese Representation Learning. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805 [cs]*, 2019.
- [13] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised Visual Representation Learning by Context Prediction. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2015.
- [14] Carl Doersch and Andrew Zisserman. Multi-task Self-Supervised Visual Learning. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2017.
- [15] Christiane Fellbaum. *WordNet_ An Electronic Lexical Database*. MIT Press, 1998.
- [16] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Boosting Few-Shot Visual Learning with Self-Supervision. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019.
- [17] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised Representation Learning by Predicting Image Rotations. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2018.
- [18] 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, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [19] Kai Han, Sylvestre-Alvise Rebuffi, Sébastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. AutoNovel: Automatically Discovering and Learning Novel Visual Categories. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2021.
- [20] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. Deep Speech: Scaling up end-to-end speech recognition. *arXiv:1412.5567 [cs]*, 2014.
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Representation Learning. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using Self-Supervised Learning Can Improve Model Robustness and Uncertainty. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [24] H Jégou, M Douze, and C Schmid. Product Quantization for Nearest Neighbor Search. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2011.
- [25] Longlong Jing and Yingli Tian. Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2019.
- [26] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised Contrastive Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [27] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3D Object Representations for Fine-Grained Categorization. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV) Workshops*, 2013.
- [28] Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Technical report, University of Toronto, 2009.
- [29] Brian Kulis. Metric Learning: A Survey. *Foundations and Trends® in Machine Learning*, 2013.
- [30] Samuli Laine and Timo Aila. Temporal Ensembling for Semi Supervised Learning. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2016.
- [31] Junnan Li, Caiming Xiong, and Steven Hoi. CoMatch: Semi-supervised Learning with Contrastive Graph Regularization. *arXiv:2011.11183 [cs]*, Mar. 2021.
- [32] Mario Lucic, Michael Tschannen, Marvin Ritter, Xiaohua Zhai, Olivier Bachem, and Sylvain Gelly. High-Fidelity Image Generation With Fewer Labels. In *Proc. of the International Conf. on Machine Learning (ICML)*, 2019.
- [33] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A Metric Learning Reality Check. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020.
- [34] Avital Oliver, Augustus Odena, Colin Raffel, Ekin D. Cubuk, and Ian J. Goodfellow. Realistic Evaluation of Deep Semi-Supervised Learning Algorithms. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [35] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. Learning and Transferring Mid-level Image Representations Using Convolutional Neural Networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [36] Boris N. Oreshkin, Pau Rodriguez, and Alexandre Lacoste. TADAM: Task dependent adaptive metric for improved few-shot learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [37] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V Le. Meta Pseudo Labels. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021.

- [38] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. CNN Features Off-the-Shelf: An Astounding Baseline for Recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2014.
- [39] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. Meta-Learning for Semi-Supervised Few-Shot Classification. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2018.
- [40] 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)*, 2015.
- [41] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A Unified Embedding for Face Recognition and Clustering. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [42] Burr Settles. Active Learning Literature Survey. Technical report, University of Wisconsin–Madison, 2010.
- [43] Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical Networks for Few-shot Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [44] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [45] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep Metric Learning via Lifted Structured Feature Embedding. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [46] Jong-Chyi Su, Subhansu Maji, and Bharath Hariharan. When Does Self-supervision Improve Few-shot Learning? In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2019.
- [47] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A Joint Model for Video and Language Representation Learning. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019.
- [48] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv:1703.01780 [cs, stat]*, 2017.
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *NIPS*, 2017.
- [50] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching Networks for One Shot Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [51] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, California Institute of Technology, 2011.
- [52] Yaqing Wang, Quanming Yao, James Kwok, and Lionel M. Ni. Generalizing from a Few Examples: A Survey on Few-Shot Learning. *ACM Computing Surveys*, 2019.
- [53] Zhirong Wu, Yuanjun Xiong, Stella Yu, and Dahua Lin. Unsupervised Feature Learning via Non-Parametric Instance-level Discrimination. *arXiv:1805.01978 [cs]*, 2018.
- [54] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4L: Self-Supervised Semi-Supervised Learning. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019.
- [55] Yu Zhang, James Qin, Daniel S. Park, Wei Han, Chung-Cheng Chiu, Ruoming Pang, Quoc V. Le, and Yonghui Wu. Pushing the Limits of Semi-Supervised Learning for Automatic Speech Recognition. *arXiv:2010.10504 [cs, eess]*, 2020.
- [56] Zhun Zhong, Linchao Zhu, Zhiming Luo, Shaozi Li, Yi Yang, and Nicu Sebe. OpenMix: Reviving Known Knowledge for Discovering Novel Visual Categories in An Open World. *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020.