

Fine-Grained Adversarial Semi-supervised Learning

DANIELE MUGNAI, University of Florence, Italy

FEDERICO PERNICI, University of Florence, Italy

FRANCESCO TURCHINI, University of Florence, Italy

ALBERTO DEL BIMBO, University of Florence, Italy

In this paper we exploit Semi-Supervised Learning (SSL) to increase the amount of training data to improve the performance of Fine-Grained Visual Categorization (FGVC). This problem has not been investigated in the past in spite of prohibitive annotation costs that FGVC requires. Our approach leverages unlabeled data with an adversarial optimization strategy in which the internal features representation is obtained with a second-order pooling model. This combination allows to back-propagate the information of the parts, represented by second-order pooling, onto unlabeled data in an adversarial training setting. We demonstrate the effectiveness of the combined use by conducting experiments on six state-of-the-art fine-grained datasets, which include Aircrafts, Stanford Cars, CUB-200-2011, Oxford Flowers, Stanford Dogs, and the recent Semi-Supervised iNaturalist-Aves. Experimental results clearly show that our proposed method has better performance than the only previous approach that examined this problem; it also obtained higher classification accuracy with respect to the supervised learning methods with which we compared.

CCS Concepts: • **Computing methodologies** → *Object recognition*.

Additional Key Words and Phrases: Fine Grained Visual Categorization, Deep Neural Networks, Semi-supervised Learning, Adversarial Learning

ACM Reference Format:

Daniele Mugnai, Federico Pernici, Francesco Turchini, and Alberto Del Bimbo. 2021. Fine-Grained Adversarial Semi-supervised Learning. *ACM Trans. Multimedia Comput. Commun. Appl.* 1, 1, Article 1 (January 2021), 19 pages. <https://doi.org/10.1145/3485473>

1 INTRODUCTION

Fine-Grained Visual Categorization (FGVC) lies in the continuum between categorization (i.e object classification) and identification (i.e. instance recognition). FGVC is quite subtle and therefore difficult to address with general-purpose object classification methods based on Deep Neural Networks (DNNs) [34]. FGVC is much more challenging than traditional classification tasks due to the inherently subtle intra-class object variability amongst sub-categories. Distinguishing between a cat and a giraffe is easy (i.e. large variability) while, in distinguishing fine-grained classes, typically only a few key features matter as in species of birds [79], dogs [28], flowers [53] or manufacturers and models of cars [32] and aircrafts [46].

Authors' addresses: Daniele Mugnai, University of Florence, Via S. Marta 3, Florence, Italy, daniele.mugnai@unifi.it; Federico Pernici, University of Florence, Via S. Marta 3, Florence, Italy, 50139, federico.pernici@unifi.it; Francesco Turchini, University of Florence, Via S. Marta 3, Florence, Italy, francesco.turchini@unifi.it; Alberto Del Bimbo, University of Florence, Florence, Italy, alberto.delbimbo@unifi.it.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

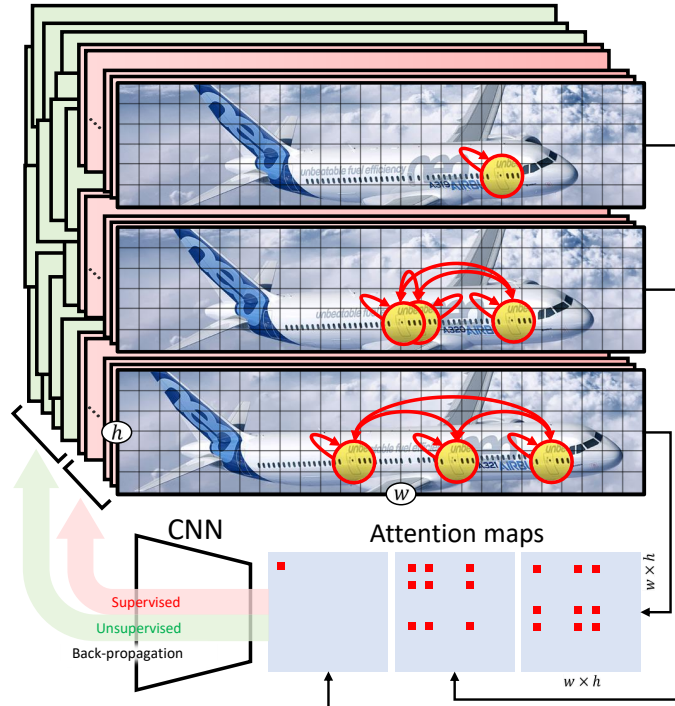


Fig. 1. Illustration of the effect of second-order pooling in Semi-Supervised Fine-Grained Visual Categorization (SSL-FGVC). We show images from three different classes of Airbus aircraft models: A319 (top), A320 (middle) and A321 (bottom). They mainly differ by the number of doors and their position along the fuselage (circles). We propose to take advantage of the long-range attention based part-to-part relationships exploited by second-order pooling and back-propagate this information onto unlabeled data to perform unsupervised structure discovery.

The task becomes significantly more difficult in domains where data is not readily available (e.g., medical images) or domains for which training data is scarce [83]. It is likely that techniques used for representation learning like semi/self-supervised or unsupervised learning that are currently used for visual recognition are not sufficient to significantly improve FGVC. In addition to this, obtaining training data for fine-grained images is prohibitively expensive, as expert knowledge is typically required [12]. In view of these issues, we propose a learning method focusing on the FGVC problem in which labeled data is limited and unlabeled is available.

Recent top performing supervised learning methods have substantially shown that the most successful strategy to FGVC is obtained by identifying, either *explicitly* or *implicitly*, the object parts [14, 30, 52, 92]. The central underlying assumption is that fine-grained information resides within the parts. Many approaches particularly focus on explicitly localizing relevant regions in an image. This is typically achieved by leveraging the extra annotations of bounding box and part annotations (some known datasets provide ground-truth part annotations [3, 79]) to localize regions that provide the most discriminative information. However, in addition to class labeling, the extra human annotations regions are not only difficult to obtain and prohibitively expensive, but can often be error-prone resulting in performance degradation [97]. Methods for unsupervised part detection and mining have been developed [17, 29, 94, 95], however,

these methods pose various challenges, such as missing parts due to occlusions and parts not providing discriminative information. According to this, it remains controversial whether unsupervised detected parts are fully beneficial.

Despite these efforts addressing different aspects regarding extra parts annotation, we are aware of only one published work in the literature addressing FGVC in a Semi-Supervised Learning setting at the image label level [50]. Despite not being investigated, this topic is getting increasing attention and support. In confirmation of this, a dataset¹ for this specific problem has been recently released for the challenge part of the FGVC7 workshop [73] held in conjunction with CVPR2020. The dataset is intended to set out some of the difficulties faced in a practical environment. The competition panel reported that all teams applied the pseudo-label SSL method [36] and that the state-of-the-art Deep-SSL methods [4, 71, 75, 91] provides similar performance but are computationally more expensive.

According to this, we propose an approach that addresses the problem in a *complementary* way in both the SSL setting and the part-based assumption of FGVC. We adopt an adversarial optimization strategy that alternately maximizes the conditional entropy of unlabeled data with respect to the classifier and minimizes it with respect to a second-order feature encoder. This combination allows to *back-propagate* the information of the *parts* captured by the second-order pooling model onto *unlabeled data* in an adversarial training setting as illustrated in Fig. 1. The strategy extends the works in [15, 16], originally proposed for Domain Adaptation, to the specific Semi-Supervised Learning setting. To the best of our knowledge, this is the first approach to leverage adversarial optimization in the specific case of SSL, and we are not aware of previous works that combine SSL with specific strategies exploiting the information of object parts. Although the recent work [50] applies SSL to FGVC datasets, the information of the parts is not explicitly taken into account during the semi-supervised learning process. We have reported a few preliminary baseline evaluations of our investigation in [49].

The main benefit of our adversarial optimization with respect to pseudo-label based methods [36, 50], is that the model can correct its own errors without incurring in wrong classifications that rapidly intensify resulting in confident but erroneous pseudo-labels on the unlabeled data. This soft label assignment may also have implications in those Deep Learning contexts in which data comes from unlabeled video streams [57, 61, 62].

We empirically demonstrate the superiority of our method over many baselines and show the method is safe [82] and [39].

2 RELATED WORK

2.1 Fine Grained Visual Categorization

Recently the accuracy of FGVC benchmarks has been substantially improved using Deep Learning. Despite the complexity of the problem [1], several methods emerged. These methods can be broadly classified in two main categories: global and localization-based [83]. Global methods use the input image as a whole and employ different strategies for pre-training [10], augmentation [31, 76], or pooling [41, 69, 70, 98] to exploit the parts. In contrast, localization-based methods approaches apply sophisticated detection techniques in order to determine the regions of the parts [17, 29, 30, 84, 87, 93–96]. The two categories have shown comparable results.

Recent state-of-the-art methods follow different strategies. The work [30] proposes a layer based on Fisher Vector encoding for aggregating part features, [92] proposes a multi-scale object and part attention approach and the work [87] uses an unsupervised learning scheme to localize informative regions. The work in [52] exploits both large external labeled datasets (the JFT [74] private dataset with 300M images) and large architectures (the AmoebaNet architecture

¹The Semi-Supervised iNaturalist-Aves Dataset: <https://github.com/cvl-umass/semi-inat-2020>

[64] with 550M parameters). It shows that fine-grained labels of JFT achieve better transfer learning for FGVC datasets. In [14] the transformer model [78] evaluated in FGVC benchmarks achieves significant performance improvements. Despite the promising results these approaches strongly rely on supervision, large model architecture or sophisticated part-based models.

The most representative method among the end-to-end feature encoding category, not relying on external part detector or complex augmentation schemes, is Bilinear-CNN [41]. Bilinear-CNN represents an image as a pooled outer product that captures pairwise correlations between the convolutional feature channels and can model part-feature interactions. The method enhances the mid-level learning capability achieving clear performance improvement on FGVC. A similar approach has been concurrently proposed in [25]. These methods have been extended and improved in several different aspects. In [40], the use of Newton-Schulz iteration [22] for efficient and accurate back-propagation is firstly proposed and in [37] strategy is further improved. We follow this approach to implicitly improve the discriminativity of the feature representation in unsupervised structure recovery.

FGVC is also interesting because they are close to the problem of learning features suitable for visual search problems such as image retrieval [7], face recognition [80] and person re-identification [89]. Moreover, fine grained datasets are shown to be less susceptible to catastrophic forgetting in class-incremental learning [47] and preliminary evidence in [8] shows encouraging results in fine-grained continual image retrieval. For the continual re-identification problem, [63] shows that performance does not decrease during learning of multiple datasets. This evidence opens up also to recent methodologies for learning the features so-called compatible [67], based on fixed classifiers [58–60], in which the re-indexing of the gallery is no longer necessary when updating the feature representation.

2.2 Semi-Supervised Learning

There is a wide literature on SSL techniques that we do not discuss here. Comprehensive overviews are provided in [55, 99] and more recently in [56, 77] for Deep SSL methods. SSL algorithms are typically [2, 4, 42, 48, 54, 86] evaluated on small-scale datasets such as CIFAR-10 [33] and SVHN [51]. We are aware of very few examples in the literature where SSL algorithms are evaluated on more challenging datasets such as the FGVC ones. To our knowledge, Mean Teacher [75] currently holds the state-of-the-art result on ILSVRC-2012 [65] when using only 10% of the labels and [50] holds the state state-of-the-art result on FGVC datasets which has been improved in this paper. Recent concurrent works [21, 85] present competitive results on ILSVRC-2012 [65].

A well-known strategy in semi-supervised learning with neural networks follows the concept of pseudo-label [36]. In pseudo-labeling method proposed network predictions are used as labels for unlabeled data: in particular a pseudo-label is assigned to the class which has the maximum predicted probability. This class is considered as if it were the true class. As discussed in [36] this method is equivalent to entropy regularization [18] and the conditional entropy of the class probabilities can be used for a measure of class overlap. By minimizing the entropy for unlabeled data, the overlap of class probability distribution can be reduced. FixMatch [71], one of the best performing methods, combines pseudo-labeling and regularization applying a strong augmentation on unlabeled data. Another recent well performing method is [5] which generates pseudo-labels trained under curriculum labeling.

The main downside of such methods is that the model is unable to correct its own mistakes [4, 71, 75, 91]. We address this issue with a complementary approach based on the so called Gradient Reversal Layer (GRL) firstly introduced in [15] and in [16] for the Unsupervised Domain Adaption (UDA) and the Semi-Supervised Domain Adaption problem (SSDA) respectively. SSL shares similarities with the related topic of Semi-Supervised Domain adaptation (SSDA), in which the test and train samples come from two different distributions [11, 13, 35, 66, 88]. However, standard evaluation of SSL

algorithms typically does not consider this case. Moreover, SSDA requires annotated data in the target distribution, and therefore it is not well suited to fine-grained data due to the cost of expert annotation. In [16] the authors left the extension and the verification of their work to the case of Semi-Supervised setting for future work. In this paper we extend the GRL to the SSL problem and evaluate it with FGVC datasets.

GRL has been used for other tasks mostly inspired by *Domain Adaptation*. Some works applied GRL for the task of object detection [20] and [38] to transfer the knowledge of the strong categories to the weak categories, for the task of image segmentation [26], speech synthesis [95], speech recognition [68] and 3D object reconstruction [27].

2.3 Semi-Supervised FGVC.

The method [50] uses pseudo-labels [36] to handle the unlabeled samples and proposes a novel strategy to select the unlabeled samples that have the highest confidence. Differently from the work presented in this paper, the work [50] uses CNN models with a linear classifier and therefore the learned fine grained structure is not explicitly based on the parts of the objects.

Only very recently, the Semi-Supervised iNaturalist-Aves dataset has been released [73]. The dataset is designed to expose some of the challenges encountered in a realistic setting, such as the fine-grained similarity between classes, significant class imbalance, and domain mismatch between the labeled and unlabeled data. The competition panel reported that all teams applied the pseudo label method [36] and that the state-of-the-art method [91] provides similar performance but is computationally more expensive. It is further reported that other recent state-of-the-art methods [4, 71, 75] do not substantially improve the performance. These results have been confirmed very recently in the evaluation performed in [72]. The experimental evaluation further shows the limitations of existing methods in taking advantage of out-of-class unlabeled data.

The only approach we are aware of using the Semi-Supervised iNaturalist-Aves is [9], in which an ensemble of 8 deep CNN models is learned according to the pseudo label method [36], and the WS-DAN [24]. Many several strategies are further included to improve the final performance of more than 20% points of accuracy. Out-of-class data is also used to cluster classes labels that are used to pre-train the models of the ensemble.

3 CONTRIBUTIONS

- We firstly propose Semi-Supervised Learning to take advantage of objects parts in FGVC to reduce the prohibitive cost of expert annotation of data.
- We demonstrate the effectiveness of our method in comparison with state-of-the-art and carefully tuned approaches. In particular, we improve the state-of-the-art in the Semi-Supervised iNaturalist-Aves dataset from 57.4% to 65.4%.
- We extend and verify the Gradient Reversal Layer (GRL), originally proposed in [16] for Domain Adaptation, to the Semi-Supervised Learning setting.

4 PROBLEM FORMULATION

In Semi-Supervised Learning, in addition to unlabeled data, the learning algorithm is provided with some supervision information, but not necessarily for all examples. In this case, the data is divided in two parts: the set for which labels are available $\mathcal{D}_l = \left\{ \left(x_i^l, y_i^l \right) \right\}_{i=1}^{n_l}$ and the set for which the labels are unknown $\mathcal{D}_u = \left\{ \left(x_i^u \right) \right\}_{i=1}^{n_u}$, where n_l and n_u are the number of examples of the labeled and unlabeled datasets, respectively. It is typically assumed that $X \times Y$ is drawn from an unknown joint probability distribution $p(X, Y)$ and that we observe it through the finite training sample \mathcal{D}_l . The

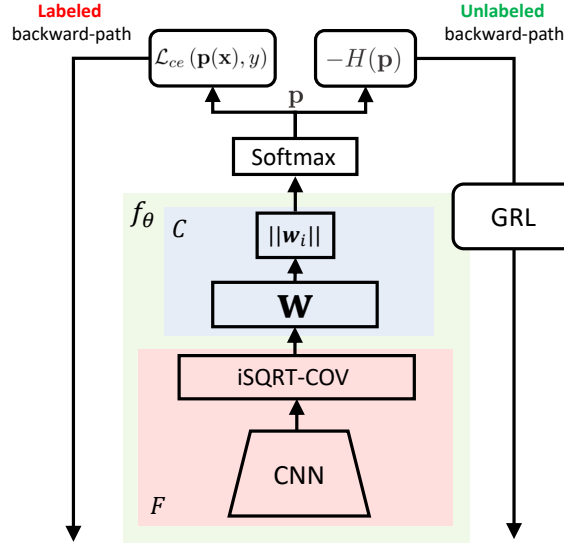


Fig. 2. An overview of the proposed model architecture. The inputs to the network are labeled and unlabeled examples. The model f_θ (light green) consists of the second-order pooling (iSQRT-COV) [81] feature extractor F (light red) and the classifier C having weight vectors \mathbf{w}_i (light blue). C is trained to maximize entropy on unlabeled target whereas F is trained to minimize it. To achieve the adversarial learning, the sign of gradients for entropy loss on unlabeled target examples is flipped by a gradient reversal layer (GRL) [15]. According to this, labeled and unlabeled back-propagation follows two distinct paths.

main goal is to leverage the unlabeled data \mathcal{D}_u to learn a Deep Neural Network model f_θ , with trainable parameters θ , that is more accurate than using the only \mathcal{D}_l . The data \mathcal{D}_u provide additional information about the structure of the data distribution $p(\mathbf{X})$ to better learn the internal feature representation of f_θ . The dependency between $p(\mathbf{X})$ and $p(Y|\mathbf{X})$ is typically established according to the *cluster assumption*, (i.e. data points in the same cluster of $p(\mathbf{X})$ have the same label Y); and *low-density separation*, (i.e. class boundaries of $p(Y|\mathbf{X})$ should lie in an area where $p(\mathbf{X})$ is small) [55]. Due to the low variability of classes in FGVC, these assumptions are quite hard to meet in practice.

4.1 Method Overview

Our base model architecture f_θ for FGVC consists of a special feature extractor F , based on second-order pooling, and a classifier C in which weights \mathbf{W} are normalized to exploit the approximated cosine distance criterion between the classifier prototypes and the features. According to this, in order to classify examples correctly, the normalized direction of a weight vector has to be representative of the features of the corresponding class in term of an angular distance. In this respect, the weight vectors can be regarded as angular estimated prototypes for each class.

Angular classifier prototypes are learned taking into account unlabeled data by exploiting an adversarial entropy optimization. Unlabeled data follow a specific data path for back-propagation that allows to attract the prototypes towards them. As the feature representation takes in account the long-range part-to-part relationships, the information of the parts can be indirectly better back-propagated towards the unlabeled data. The architecture of our method is shown in Fig. 2 and will be detailed in the next Subsection.

4.2 Back-propagation of the Parts onto Unlabeled Data

The goal is to obtain representative classifier prototypes $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_K]$ where K is the number of classes for labeled and unlabeled data and to minimize the distance between the prototypes and the unlabeled examples. We approach the problem in mainly three different fronts: 1) reduce intra-class distance (next paragraph), 2) improve feature discriminativity exploiting part-to-part relationships (Subsection 4.2.1) 3) handling unlabeled data distance between prototypes and features (Subsection 4.2.2).

For labeled data, the general purpose-classification linear classifier already minimizes the distance between features and classifier prototypes. However, as in many other instance recognition tasks (i.e. face recognition, re-identification), features in ideal FGVC are expected to have smaller maximal intra-class distance than minimal inter-class distance under a suitably chosen metric space. The vanilla linear classifier cannot effectively satisfy this criterion [44]. One simple method to enable convolutional neural networks to produce more discriminative features is imposing discriminative constraints on a hypersphere manifold by normalizing the vectors of the classifier weights [43]:

$$\mathbf{w}'_i = \frac{\mathbf{w}_i}{\|\mathbf{w}_i\|} \quad i = 1, 2, \dots, K, \quad \mathbf{w}'_i \in \mathbb{R}^m \quad (1)$$

where, as detailed in the next subsection, $m = \frac{d(d+1)}{2}$ with d , is the channel dimension of the last convolutional layer of the architecture (Fig. 3).




Fig. 3. The $w \times h$ feature channels of dimension d of the last convolutional layer of the CNN architecture are used to compute the covariance matrix. The $\frac{d(d+1)}{2}$ -dimensional values of the upper triangular matrix constitute the internal feature representation vector that allows the model to determine the attention based long-range part-to-part relationships. The forward and backward propagation of the covariance in the adversarial optimization setting of Fig. 2 are computed according to the iSQRT-COV approximated method.

4.2.1 Second-order Pooling Feature Extraction. Here we briefly recall the key elements for feature extraction as proposed in [37]. Let $\mathbf{X} \in \mathbb{R}^{w \times h \times d}$ be the output of the last convolutional layer of an architecture with spatial height h , width w and channel d . The tensor is reshaped into a features matrix \mathbf{X} consisting of $n = wh$ features of d dimension. Then second-order pooling is performed by computing the covariance matrix $\Sigma = \mathbf{X}\bar{\mathbf{X}}^T$ as shown in Fig 3. The covariance matrix Σ is a symmetric positive definite or semidefinite matrix, which can be factorized by EIG/SVD: $\Sigma = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$. Through EIG or SVD, the matrix power can be computed as: $\mathbf{Z} \triangleq \Sigma^\alpha = \mathbf{U}\text{diag}(f(\lambda_1), \dots, f(\lambda_d))\mathbf{U}^T$ where the exponent α is a positive real number. Empirically, the covariance matrix works best when α is 0.5.

Matrix square root computation heavily depends on EIG or SVD operation; however, these methods don't have a fast GPU implementation. To address this problem, [37] proposes a fast method called iterative matrix square root normalization of covariance pooling (iSQRT-COV) that makes use of Newton-Schulz iteration. Let $\mathbf{Y}_0 = \mathbf{A}$ and $\mathbf{Z}_0 = \mathbf{I}$

for $k = 1, \dots, N$, the coupled iteration takes the following form:

$$\begin{aligned} \mathbf{Y}_k &= \frac{1}{2} \mathbf{Y}_{k-1} (\mathbf{3I} - \mathbf{Z}_{k-1} \mathbf{Y}_{k-1}) \\ \mathbf{Z}_k &= \frac{1}{2} (\mathbf{3I} - \mathbf{Z}_{k-1} \mathbf{Y}_{k-1}) \mathbf{Z}_{k-1}. \end{aligned} \quad (2)$$

These equations above can be computed with a single matrix product, which is more suitable for GPU implementation. In order to guarantee the convergence we pre-normalize covariance Σ by its trace or Frobenius norm, i.e., $\mathbf{A} = \frac{1}{\text{tr}(\Sigma)} \Sigma$ or $\frac{1}{\|\Sigma\|_F} \Sigma$. Finally, in order to compensate the data magnitude caused by pre-normalization, the covariance matrix \mathbf{Z} is calculated as:

$$\mathbf{Z} = \sqrt{\text{tr}(\Sigma)} \mathbf{Y}_N \quad \text{or} \quad \mathbf{Z} = \sqrt{\|\Sigma\|_F} \mathbf{Y}_N. \quad (3)$$

The matrix \mathbf{Z} is symmetric and has $\frac{d(d+1)}{2}$ different parameters that are gathered into a vector (i.e. the feature vector) which is the input of final classifier layer of Eq. 1.

4.2.2 Handling Unlabeled Data with Adversarial Entropy Optimization. The obtained features and the relative prototypes provide improved discriminative power and reduced low intra-class variation, respectively. Yet, the overall goal remains how to include a strategy to minimize the distance between the prototypes and the unlabeled examples. As this cannot be directly achieved with a linear classifier alone, we exploited the adversarial strategy originally proposed in [15, 16]. The adversarial part of our strategy clusters features computed from unlabeled data around the classifier learned prototypes. Therefore, we train the feature extractor F and the classifier C to classify labeled and utilize standard cross-entropy minimization objective to extract discriminative features for the labeled data:

$$\mathcal{L} = -\frac{1}{n_l} \sum_{j=1}^{n_l} \sum_{i=1}^K \hat{y}_j \log p(y_j = i | \mathbf{x}_j), \quad (4)$$

where \hat{y}_j is the true class label for x_j . The *unlabeled data* are used to *maximize* the entropy with respect to the classifier C and to *minimize* the entropy with respect the feature extractor F . The entropy is computed as follow:

$$H = -\frac{1}{n_u} \sum_{j=1}^{n_u} \sum_{i=1}^K p(y_j = i | \mathbf{x}_j) \log p(y_j = i | \mathbf{x}_j). \quad (5)$$

The intuition that high entropy encourages the classifier weight vectors to “move” towards the unlabeled data features is illustrated in Fig. 4(a). This is due to the fact that high entropy, namely the maximization of Eq. 5, tends to achieve a uniform distribution of the softmax output probabilities that consequently encourage each prototype \mathbf{w}_i to be similar to all the unlabeled features. This strategy can be considered as an “*adversarial move*” of the classifier, whose intention is to “*explore*” the representation space driven by the unlabeled data. This exploration is counterbalanced by the effect of low entropy that encourages learning a representation (i.e. the feature extractor F) that clusters labeled and unlabeled features around the estimated prototypes (Fig. 4(b)). Alternating these two opposite “forces” determines a sort of equilibrium in which the discriminative feature extractor F and the classifier C may have better explored the representation space as driven by unlabeled data. This can be considered as a form of soft pseudo-labeling strategy in which soft association can be eventually recovered. This co-optimization can be formulated as an adversarial learning

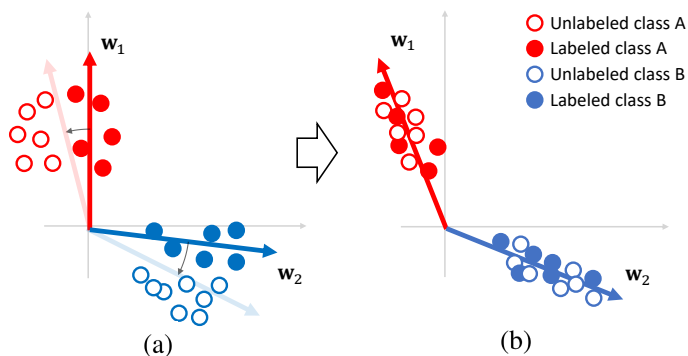


Fig. 4. The adversarial learning intuition between the classifier C and the feature extractor F in a two class learning task. For visualization purposes, features are represented in a two-dimensional space. (a): High entropy between the classifier weight vectors w_1 and w_2 (i.e., the prototypes of the classifier C) and the unlabeled data features encourages the classifier weight vectors to “move” towards the unlabeled data features. (b): This “motion” is counterbalanced by the low entropy which instead learns a representation (i.e. the feature extractor F) that tends to cluster labeled and unlabeled features around the estimated prototypes.

between C and F by weighted summing the two losses of Eq. 4 and Eq. 5 as follows:

$$\hat{\theta}_F = \underset{\theta_F}{\operatorname{argmin}} \mathcal{L} + \lambda H \quad (6)$$

$$\hat{\theta}_C = \underset{\theta_C}{\operatorname{argmin}} \mathcal{L} - \lambda H, \quad (7)$$

where $\lambda > 0$ is the weighting factor between the two losses and θ_C , θ_F are the learnable parameters of the classifier and the feature extractor, respectively. F and C are co-optimized in two steps: in the first step, both F and C are optimized by minimizing the cross-entropy loss on labeled data. In the second step, F and C are optimized in opposite ways on unlabeled data, minimizing the entropy and maximizing the entropy, respectively (the signs of the entropy in the two equations of Eq. 6 and Eq. 7 are opposite). In particular, the negative sign of the entropy H in Eq. 7 changes the minimization problem into a maximization one. In the second step, input data follows the unlabeled path (Fig. 2), on which the classifier and the feature extractor are connected via a Gradient Reversal Layer (GRL) [15]. During forward propagation, GRL acts as an identity transform. During the back-propagation, GRL takes the gradient from the subsequent level, multiplies it by $-\lambda$ and passes it to the preceding layer (Fig. 2). By adding the gradient reversal layer, the training process described above can be achieved through normal model training.

5 EXPERIMENTAL RESULTS

5.1 FGVC Datasets

Evaluation is performed on six datasets: *FGVC Aircraft* [46], *Stanford Cars* [32], *Cub-200-2011* [79], *Stanford Dogs* [28], *Oxford Flowers* [53] and the recently released SSL *iNaturalist-Aves* [73]. *FGVC Aircraft* is made up of 10,000 images of aircrafts, which are organized in four different hierarchies (Models, Variants, Families and Manufacturers). *Stanford Cars* contains 16,185 images of 197 models of cars of various categories like SUVs, coupes, pickups, hatchbacks, station wagons. Each label reports manufacturer, model and production year. *Cub-200-2011* is an extended version of CUB-200 dataset and it is composed of 11,788 images of 200 bird species. The labels of species have been gathered from Wikipedia and images have been retrieved from Flickr. Various metadata like bounding-boxes, attributes and annotations are

available. *Oxford flower 102* [53] is a 102 category dataset consisting of 102 flowers species. Images depict commonly occurring flowers in the United Kingdom. Each class consists of a variable number between 40 and 258 images. The training set consists of 2040 images while the test set consists of 6149 images. The *Stanford Dogs Dataset* [28] contains 120 breeds of dogs from around the world. This dataset has been built using images and annotations from ImageNet for the task of fine-grained categorization. The training set is made up of 12000 images while test set consists of 8500 images. The SSL iNaturalist-Aves dataset is composed of images of 200 different species of birds split as in Tab. 1. We used the 3,959 labeled and the 26,640 unlabeled in-class images to train our model. The dataset also includes 122,208 unlabeled, out-of-class images we did not use.

These datasets are rather complex due to: similarities existing among classes, occlusion, subtle intra-class variations, imbalance set. A summary of the evaluated datasets is shown in Tab. 1.

Table 1. A summary information for the evaluated FGVC datasets. Semi-Supervised iNaturalist-Aves explicitly includes unlabeled data for classes outside the known ones (out-of-class).

Dataset	Classes #	Train #	Test #	Unlabeled (In Class) #	Unlabeled (Out of Class) #
FGVC-Aircraft [46]	100	6667	3333	-	-
CUB-200-2011 [79]	200	5994	5794	-	-
Stanford Cars [32]	196	8144	8041	-	-
Oxford flowers [53]	102	2040	6149	-	-
Stanford Dogs [28]	120	12000	8580	-	-
Semi-Supervised iNaturalist-Aves [73]	200	4000	4000	26640	122208

5.2 Baseline methods

We compare the proposed method with the following baselines:

- **Sup**: Supervised learning with standard linear classifier. No explicit methods for exploiting part-to-part relationships are used and unlabeled data is not considered.
- **Sup-Cov**: This is the second-order pooling (iSQRT-COV feature extractor) [37], trained in a *fully supervised* setting with only cross-entropy loss minimization.
- **Ent-Cov**: This is the second-order pooling (iSQRT-COV feature extractor) [37], trained using both labeled and unlabeled data minimizing both cross-entropy loss and standard entropy minimization [18]. Entropy is calculated on unlabeled data as defined in Eq.5. In this case the adversarial optimization process is not considered.
- **SSLFGC**: This is the SSL for Fine-Grained Classification (SSLFGC) method based on pseudo-label proposed in [50].

5.3 Evaluation on FGVC Datasets

As one of the aim of these experiments is to compare our method with the semi-supervised method described in [50], we follow the datasets splits reported in the paper. According to this, the training set is partitioned into 75% for training and 25% for validation, while the entire test set is used as unlabeled data. Our learning configuration includes two convolutional architectures of substantially different expressive power: the ResNet50 and the ResNet101 that have 25 and 44 million parameters, respectively. Both architectures are pre-trained on the ILSVRC-2012 ImageNet dataset [65]. Images of any dataset are resized to 448x448. Data augmentation is performed according to basic image transformations as random crop, center crop and random horizontal flip. The proposed method is trained with SGD optimization and the




Fig. 5. Accuracy as a function of the model parameter λ evaluated for all the datasets. Values around $\lambda = 0.025$ achieve best accuracy for three of the datasets: *Oxford Flowers*, *Stanford Dogs* and *SS iNaturalist-Aves*. *FGVC Aircrafts* achieves the maximum accuracy around $\lambda = 0.2$ while *Stanford Dogs* and *Stanford Cars* at $\lambda = 0.05$.

Table 2. Comparison with other baselines on the Datasets reported in Tab. 1. We report the standard classification accuracy, following the dataset splitting described in section 5.3. Overall, our method largely outperforms other baselines for all evaluated datasets and for both architectures we used, with some exceptions as *Flowers* (where results are similar) and *Dogs*.

Method	Dataset						#Params
	Aircrafts	CUB-200-2011	Cars	Flowers	Dogs	SS iNat-Aves	
Sup	87.38	85.50	88.35	94.05	74.49	52.7	ResNet50 (25M)
Sup-Cov [37]	88.59	85.70	91.48	95.06	85.22	64.7	ResNet50 (25M)
Ent-Cov	84.45	85.00	87.30	93.38	83.83	65.1	ResNet50 (25M)
Ours w/o Cov	88.10	85.80	91.42	94.96	77.40	50.50	ResNet50 (25M)
Ours	91.71	87.04	92.74	95.49	85.13	65.4	ResNet50 (25M)
SSLFGC [50]	-	82.72	93.71	96.72	91.67	-	Incept. ResNetV2 (56M)
Ours	91.74	87.33	93.83	96.74	87.28	69.85	ResNet101 (44M)

learning rate is set to 0.0012 for both the CNN architectures and for the second-order pooling (iSQRT-COV algorithm). The learning rate for the classifier is set to 0.003 following the implementation of [37] and the batch size is set to 10. To balance efficiency and accuracy in the number of iterations N in Eq. 2, we follow the same setting described in the original works [37] [40] and set N to 5 in all of our experiments. Training is performed for 50k iterations and the validation in early stopping to prevent overfitting and for evaluating the sensitivity of the λ hyper-parameter.

In particular, the sensitivity of the hyper-parameter λ of Eq. 6 and Eq. 7 is selected by grid-search in the interval $[0.025, 1]$. As shown in Fig. 5 the resulted optimal values vary according to the specific dataset. Values around $\lambda = 0.025$ achieves the best accuracy for three of the datasets: *Oxford Flowers*, *Stanford Dogs* and *SS iNaturalist-Aves*. *FGVC Aircrafts* achieves the maximum accuracy around $\lambda = 0.2$ while *Stanford Dogs* and *Stanford Cars* at $\lambda = 0.05$. For the subsequent evaluations each dataset uses the value λ which obtained the best accuracy.

As for $\lambda = 0$ the proposed method optimizes the standard cross-entropy loss, it performs standard supervised learning. According to this, it is evident that the performance of the proposed method exploits favourably the unlabeled data in all the datasets when $\lambda > 0$. This because all the curves in Fig. 5 increase in a neighbourhood of the origin.

We compare with the baselines presented in Subsection 5.2 for all the datasets described in Subsection 5.1 and report all the results in Tab. 2. As it can be noticed, in comparison with Sup-Cov and Ent-Cov baselines, our approach is overall more effective, especially in *FGVC Aircrafts*, *CUB-200-2011* and *Cars*. Comparable results are obtained in *Flowers*. Conversely, SSLFGC obtains higher accuracy on *Dogs* dataset by a clear margin. The results of SSLFGC are taken from Tab. 5 of [50] which reports the results considering the top-k most confident pseudo-labels images added to the training set. We choose their best results with $k = 10$. In this specific case SSLFGC is better than the current supervised learning state-of-the-art method [100] achieving an accuracy of 90.3%. However, it is worth to notice that the SSLFGC approach exploits a CNN architecture with more parameters. Tab.2 also includes our approach without the Second-Order-Pooling, indicated as Ours w/o Cov, showing a clear detrimental effect on the overall performance. For the case of the iNaturalist-Aves dataset in which there is a large number of unlabeled data the improvement in exploiting the parts through the second order pooling is about 15%.

Comparative feature evaluation according to t-SNE [45] is shown in Fig. 7 for the Sup-Cov supervised baseline and our SSL method for three of the six datasets. The figure shows that the method effectively leverages the structure of unsupervised data to obtain more discriminative features.

5.4 Safety Evaluation

Semi-supervised classification methods could have lower performance than their supervised counterparts in some situations [82] and [39]. It is therefore desirable to design secure semi-supervised classification methods which do not perform worse than their supervised counterparts. According to this, we evaluated the proposed method varying the amount of available unlabeled data using the same FGVC datasets split we used in the previous section. In this section we evaluate the *label rate* (i.e., the amount of unlabeled data used for training). When the label rate = 0, no unlabeled data is used for training (this is the case of the Sup-Cov baseline). When the label rate = 1, the entire test set is used as unlabeled data. As shown in Fig. 6(a) the classification accuracy in all FGVC datasets increases as the label rate increases. Although performance degrades slightly at the beginning, on average the trend is that of positive growth. The proposed approach performs better at label rate 1 than at label rate 0 for all the evaluated datasets. This result is apparently obvious but yet not to be taken for granted in SSL [82]. Further possible causes of the slight degradation (i.e. not strict growth) at the beginning may also be found in the different levels of sensitivity of the algorithms while varying datasets [54]. To this end, we further evaluated in Fig.5(b) the behavior of our method in the absence of the adversarial learning strategy (i.e. [18] with the second order pooling that we called Ent-Cov). As evidenced by the figure, the adversarial strategy significantly improves the overall capability of dealing with unlabeled data. For Aircrafts and Cars, for example, Ent-Cov is not able to take advantage of the unlabeled data while for the other datasets our method significantly improves the performance. Contrarily, the Ent-Cov method doesn't efficiently leverage unlabeled data; the accuracy test for Flowers, Cars, Aircraft accuracy decreases as the number of unlabeled data grows up. For Dogs and Birds datasets, the trend of accuracy is similar to our method.

5.5 Evaluation on the Semi-Supervised iNaturalist-Aves Dataset

This dataset is significantly harder to evaluate using the protocol we followed to compare with [50]. This is due to the fact that the ratio between the number of labeled images and the unlabeled ones, except for *Oxford Flowers*, is greater than one. We compare with Sup-Cov [37] as described in the previous section and with the six different SSL methods evaluated in [72]: Pseudo-Labeling [36], Curriculum Pseudo-Labeling, [5], FixMatch [71], Self-Training [72], MoCo (Momentum Contrast) [19] and MoCo + Self-Training [6]. Specifically, the Self-training baseline initially trains a teacher




Fig. 6. Test accuracy for all FGVC datasets varying the rate number of unlabeled data (i.e. the label rate). The x-axis reports the label rate of test set used as unlabeled dataset. (a): Our method. (b): Ent-Cov. The proposed method improves performance as the number of the unlabeled data increases.

model with only labeled data, then transfers the knowledge to a student model by distillation [23] using both labeled and unlabeled data. The MoCo + Self-Training performs a self-supervised pre-training with MoCo then removes the final MLP layers and adds a classification layer that is trained with labeled data. The results of the comparison are shown in Tab. 3. Our method shows state-of-the-art performance with respect to the baselines with an accuracy result of 69.85% and 65.4% with the ResNet101 and ResNet50 architecture, respectively. The gain in classification accuracy from FixMatch (i.e. the best performing algorithm) using the same CNN backbone is 8% (from 57.4% to 65.4%). As evidenced from the table the increase in classification accuracy is mostly due to the second order pooling layer and secondly by the adversarial strategy. This also confirms the results on the Sec. 5.3.

One of the participants to the FGVC7-competition covered their training details in [9]. As no implementation is available, a fair comparison with our method is difficult to be conducted. The method employs an ensemble of eight CNN architectures for a total of about 500M parameters learned with batch size of 128 elements and several sophisticated training tricks (for example high-resolution images and advanced data augmentation like CutMix [90]), which are difficult to reproduce and which considerably improve the performance of about 20% points. ImageNet based pretrained models are further improved by clustering with unlabeled out-of-class data.

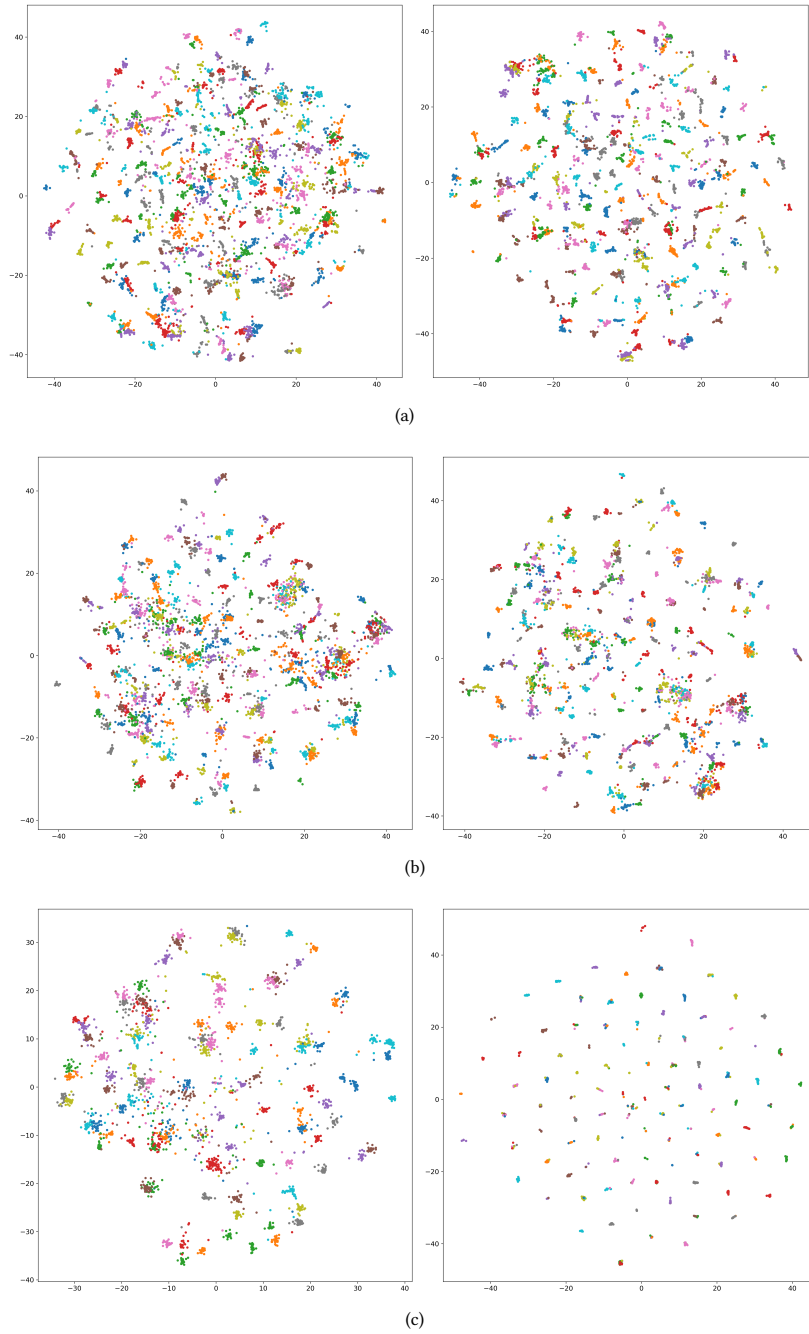


Fig. 7. Feature visualization with t-SNE. Qualitative comparison between features learned by the Sup-Cov baseline (*left* column) and the method proposed in this paper (*right* column) for the: (a) Cars, (b) Birds and (c) Aircrafts datasets. As evidenced by the figures, the features learned by the proposed method are more discriminative.

Table 3. Results on the Semi-Supervised iNaturalist-Aves Dataset (FGVC7 challenge). Our method achieves a significant improvement by leveraging unsupervised data.

Method	Accuracy	#Params
Pseudo-Label [36]	54.4	ResNet50 (25M)
Curriculum Pseudo-Label [5]	53.4	ResNet50 (25M)
FixMatch [71]	57.4	ResNet50 (25M)
Self-Training [72]	55.5	ResNet50 (25M)
MoCo [19]	55.5	ResNet50 (25M)
MoCo + Self-Training [6]	52.7	ResNet50 (25M)
Sup [72]	52.7	ResNet50 (25M)
Sup-Cov [37]	64.7	ResNet50 (25M)
Ours w/o Cov	50.5	ResNet50 (25M)
Ours	65.4	ResNet50 (25M)
Ours	69.85	ResNet101 (44M)

5.6 Ablation of Mini-batch Composition

At every iteration two mini-batches are sampled: one from the labeled data and the other from the unlabeled data. The two mini-batches are used to evaluate the objective in Eq. 6 and Eq. 7, as shown in Fig. 2, they follow two distinct backward path during back-propagation. We investigated the sensitivity of the relative and absolute size of the two mini-batches. Table 4 shows the effect of varying the relative composition of the two mini-batches on the accuracy evaluated on the six FGVC datasets. As evidenced by the table, the configuration for the two mini-batches achieving the best average accuracy is with 10 labeled and 10 unlabeled examples (i.e. 86.25%). Keeping the amount of labeled examples constant while increasing and decreasing of 5 examples the unlabeled mini-batch does not improve the average accuracy over the datasets. In both cases we obtained similar values of the average accuracy: 85.62 and 85.87. As shown in the same table, increasing the size to 20 elements for each mini-batch does not improve the accuracy. This result suggests that there exists an inherent trade-off between the information carried by the two different mini-batches.

Table 4. Sensitivity of the mini-batch composition. Best performance is achieved when: 1) the mini-batches for the two distinct paths contain the same number of examples and 2) when their size is 10.

#labeled/#unlabeled	Dataset						Avg	#Params
	Aircrafts	CUB	Cars	Flowers	Dogs	SS iNat-Aves		
10/5	90.72	86.92	92.38	93.99	84.96	64.80	85.62	ResNet50 (25M)
10/10	91.71	87.04	92.74	95.49	85.13	65.40	86.25	ResNet50 (25M)
10/15	91.26	86.92	92.56	92.96	85.68	65.85	85.87	ResNet50 (25M)
20/20	90.75	86.47	91.95	94.00	86.38	66.75	86.05	ResNet50 (25M)

6 CONCLUSION

We firstly introduced a Semi-Supervised Learning (SSL) method which combines ideas and components from adversarial entropy optimization and second-order pooling. Our main goal was reducing the prohibitive annotation cost of Fine-Grained Visual Categorization (FGVC) according to the SSL setting, a combination of problems which has not been investigated in the past. Through extensive experiments on FGVC datasets, including the very recent *Semi-Supervised iNaturalist-Aves*, we found that the proposed method exhibits significantly improved performance compared to some baselines we included and to the only previous work that examined this problem. In almost all FGVC datasets we evaluated, we obtained state-of-the-art performance. Moreover, we firstly introduced the Gradient Reversal Layer (GRL), one of the key elements of adversarial entropy optimization to the SSL setting. This extension and its validation were left as future work in the original study [15, 16].

Acknowledgment: This work was partially supported by the European Commission under European Horizon 2020 Programme, grant number 951911 - AI4Media. This research was also partially supported by Leonardo Finmeccanica S.p.A.

REFERENCES

- [1] Connor Anderson, Matt Gwilliam, Adam Teuscher, Andrew Merrill, and Ryan Farrell. 2020. Facing the Hard Problems in FGVC. *arXiv preprint arXiv:2006.13190* (2020).
- [2] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. 2018. There Are Many Consistent Explanations of Unlabeled Data: Why You Should Average. In *International Conference on Learning Representations*.
- [3] Thomas Berg, Jiongxin Liu, Seung Woo Lee, Michelle L Alexander, David W Jacobs, and Peter N Belhumeur. 2014. Birdsnap: Large-scale fine-grained visual categorization of birds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. 2019. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*.
- [5] Paola Cascante-Bonilla, Fuwen Tan, Yanjun Qi, and Vicente Ordonez. 2021. Curriculum Labeling: Revisiting Pseudo-Labeling for Semi-Supervised Learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 8 (May 2021), 6912–6920.
- [6] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. Big Self-Supervised Models are Strong Semi-Supervised Learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- [7] Wei Chen, Yu Liu, Weiping Wang, Erwin Bakker, Theodoros Georgiou, Paul Fieguth, Li Liu, and Michael S Lew. 2021. Deep image retrieval: A survey. *arXiv preprint arXiv:2101.11282* (2021).
- [8] Wei Chen, Yu Liu, Weiping Wang, Tinne Tuytelaars, Erwin M. Bakker, and Michael S. Lew. 2020. On the Exploration of Incremental Learning for Fine-grained Image Retrieval. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press.
- [9] Cheng Cui, Zhi Ye, Yangxi Li, Xinjian Li, Min Yang, Kai Wei, Bing Dai, Yanmei Zhao, Zhongji Liu, and Rong Pang. 2020. Semi-Supervised Recognition under a Noisy and Fine-grained Dataset. *arXiv preprint arXiv:2006.10702* (2020).
- [10] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. 2018. Large scale fine-grained categorization and domain-specific transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [11] Hal Daumé III, Abhishek Kumar, and Avishek Saha. 2010. Frustratingly easy semi-supervised domain adaptation. In *Proceedings of the 2010 Workshop on Domain Adaptation for Natural Language Processing*.
- [12] Jia Deng, Jonathan Krause, and Li Fei-Fei. 2013. Fine-grained crowdsourcing for fine-grained recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [13] Jeff Donahue, Judy Hoffman, Erik Rodner, Kate Saenko, and Trevor Darrell. 2013. Semi-supervised domain adaptation with instance constraints. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale.
- [15] Yaroslav Ganin and Victor Lempitsky. [n.d.]. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*.
- [16] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* (2016).

- [17] Weifeng Ge, Xiangru Lin, and Yizhou Yu. 2019. Weakly supervised complementary parts models for fine-grained image classification from the bottom up. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [18] Yves Grandvalet and Yoshua Bengio. 2005. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*.
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [20] Zhenwei He and Lei Zhang. 2019. Multi-adversarial faster-rcnn for unrestricted object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6668–6677.
- [21] Olivier J Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. 2019. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272* (2019).
- [22] Nicholas J Higham. 2008. *Functions of matrices: theory and computation*. SIAM.
- [23] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [24] Tao Hu, Honggang Qi, Qingming Huang, and Yan Lu. 2019. See better before looking closer: Weakly supervised data augmentation network for fine-grained visual classification. *arXiv preprint arXiv:1901.09891* (2019).
- [25] Catalin Ionescu, Orestis Vantzos, and Cristian Sminchisescu. 2015. Matrix backpropagation for deep networks with structured layers. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [26] Mehran Javanmardi and Tolga Tasdizen. 2018. Domain adaptation for biomedical image segmentation using adversarial training. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, 554–558.
- [27] Hiroharu Kato and Tatsuya Harada. 2019. Learning view priors for single-view 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9778–9787.
- [28] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. 2011. Novel Dataset for Fine-Grained Image Categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*.
- [29] Dimitri Korsch, Paul Bodesheim, and Joachim Denzler. 2019. Classification-specific parts for improving fine-grained visual categorization. In *German Conference on Pattern Recognition*.
- [30] Dimitri Korsch, Paul Bodesheim, and Joachim Denzler. 2020. End-to-end Learning of a Fisher Vector Encoding for Part Features in Fine-grained Recognition. *arXiv preprint arXiv:2007.02080* (2020).
- [31] Jonathan Krause, Benjamin Sapp, Andrew Howard, Howard Zhou, Alexander Toshev, Tom Duerig, James Philbin, and Li Fei-Fei. 2016. The unreasonable effectiveness of noisy data for fine-grained recognition. In *European Conference on Computer Vision*.
- [32] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*.
- [33] Alex Krizhevsky et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [34] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*.
- [35] Abhishek Kumar, Avishek Saha, and Hal Daume. 2010. Co-regularization based semi-supervised domain adaptation. In *Advances in neural information processing systems*.
- [36] Dong-Hyun Lee. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML 2013 Workshop : Challenges in Representation Learning (WREPL), Atlanta, Georgia, USA*.
- [37] Peihua Li, Jiangtao Xie, Qilong Wang, and Zilin Gao. 2018. Towards faster training of global covariance pooling networks by iterative matrix square root normalization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [38] Yan Li, Junge Zhang, Kaiqi Huang, and Jianguo Zhang. 2018. Mixed supervised object detection with robust objectness transfer. *IEEE transactions on pattern analysis and machine intelligence* 41, 3 (2018), 639–653.
- [39] Yu-Feng Li and De-Ming Liang. 2019. Safe semi-supervised learning: a brief introduction. *Frontiers of Computer Science* 13, 4 (2019), 669–676.
- [40] Tsung-Yu Lin and Subhransu Maji. 2017. Improved Bilinear Pooling with CNNs. In *Proceedings of the British Machine Vision Conference (BMVC)*. BMVA Press.
- [41] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. 2015. Bilinear cnn models for fine-grained visual recognition. In *Proceedings of the IEEE international conference on computer vision*.
- [42] Bin Liu, Zhirong Wu, Han Hu, and Stephen Lin. 2019. Deep metric transfer for label propagation with limited annotated data. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*.
- [43] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. 2017. Sphreface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [44] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. 2016. Large-margin softmax loss for convolutional neural networks. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*.
- [45] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* (2008).
- [46] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151* (2013).
- [47] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. 2020. Class-incremental learning: survey and performance evaluation. *arXiv preprint arXiv:2010.15277* (2020).

- [48] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* (2018).
- [49] Daniele Mugnai, Federico Pernici, Francesco Turchini, and Alberto Del Bimbo. 2021. Soft Pseudo-labeling Semi-Supervised Learning Applied to Fine-Grained Visual Classification. In *Pattern Recognition. ICPR International Workshops and Challenges: Virtual Event, January 10–15, 2021, Proceedings, Part IV*. Springer International Publishing, 102–110.
- [50] Obed Tettey Nartey, Guowu Yang, Jinzhao Wu, and Sarpong Kwadwo Asare. 2019. Semi-Supervised Learning for Fine-Grained Classification with Self-Training. *IEEE Access* (2019).
- [51] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. 2011. Reading digits in natural images with unsupervised feature learning. (2011).
- [52] Jiquan Ngiam, Daiyi Peng, Vijay Vasudevan, Simon Kornblith, Quoc V Le, and Ruoming Pang. 2018. Domain adaptive transfer learning with specialist models. *arXiv preprint arXiv:1811.07056* (2018).
- [53] Maria-Elena Nilsback and Andrew Zisserman. 2008. Automated Flower Classification over a Large Number of Classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*.
- [54] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. 2018. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in neural information processing systems*.
- [55] Chapelle Olivier, Scholkopf Bernhard, and Zien Alexander. 2006. Semi-supervised learning. In *IEEE Transactions on Neural Networks*.
- [56] Yassine Ouali, Céline Hudelot, and Myriam Tami. 2020. An Overview of Deep Semi-Supervised Learning. *arXiv preprint arXiv:2006.05278* (2020).
- [57] Federico Pernici, Federico Bartoli, Matteo Bruni, and Alberto Del Bimbo. 2018. Memory based online learning of deep representations from video streams. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2324–2334.
- [58] F. Pernici, M. Bruni, C. Baecchi, and A. D. Bimbo. 2021. Regular Polytope Networks. *IEEE Transactions on Neural Networks and Learning Systems* (2021), 1–15. <https://doi.org/10.1109/TNNLS.2021.3056762>
- [59] Federico Pernici, Matteo Bruni, Claudio Baecchi, and Alberto Del Bimbo. 2019. Maximally Compact and Separated Features with Regular Polytope Networks.. In *CVPR Workshops*. 46–53.
- [60] Federico Pernici, Matteo Bruni, Claudio Baecchi, Francesco Turchini, and Alberto Del Bimbo. 2020. Class-incremental Learning with Pre-allocated Fixed Classifiers. In *25th International Conference on Pattern Recognition, ICPR 2020, Milan, Italy, January 10-15, 2021*. IEEE Computer Society.
- [61] Federico Pernici, Matteo Bruni, and Alberto Del Bimbo. 2020. Self-supervised on-line cumulative learning from video streams. *Computer Vision and Image Understanding* 197 (2020), 102983.
- [62] Federico Pernici and Alberto Del Bimbo. 2017. Unsupervised incremental learning of deep descriptors from video streams. In *2017 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, 477–482.
- [63] Nan Pu, Wei Chen, Yu Liu, Erwin M. Bakker, and Michael S. Lew. 2021. Lifelong Person Re-Identification via Adaptive Knowledge Accumulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 7901–7910.
- [64] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. 2019. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*.
- [65] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* (2015).
- [66] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. 2019. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [67] Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. 2020. Towards backward-compatible representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6368–6377.
- [68] Yusuke Shinohara. 2016. Adversarial Multi-Task Learning of Deep Neural Networks for Robust Speech Recognition.. In *Interspeech*. San Francisco, CA, USA, 2369–2372.
- [69] Marcel Simon, Yang Gao, Trevor Darrell, Joachim Denzler, and Erik Rodner. 2017. Generalized orderless pooling performs implicit salient matching. In *Proceedings of the IEEE international conference on computer vision*.
- [70] Marcel Simon, Erik Rodner, Trevor Darrell, and Joachim Denzler. 2018. The whole is more than its parts? from explicit to implicit pose normalization. *IEEE transactions on pattern analysis and machine intelligence* (2018).
- [71] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685* (2020).
- [72] Jong-Chyi Su, Zezhou Cheng, and Subhansu Maji. 2021. A realistic evaluation of semi-supervised learning for fine-grained classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12966–12975.
- [73] Jong-Chyi Su and Subhansu Maji. 2021. The Semi-Supervised iNaturalist-Aves Challenge at FGVC7 Workshop.
- [74] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*.
- [75] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*.
- [76] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. 2019. Fixing the train-test resolution discrepancy. In *Advances in Neural Information Processing Systems*.

- [77] Jesper E Van Engelen and Holger H Hoos. 2020. A survey on semi-supervised learning. *Machine Learning* (2020).
- [78] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*.
- [79] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. The caltech-ucsd birds-200-2011 dataset. (2011).
- [80] Mei Wang and Weihong Deng. 2021. Deep face recognition: A survey. *Neurocomputing* 429 (2021), 215–244. <https://doi.org/10.1016/j.neucom.2020.10.081>
- [81] Q. Wang, J. Xie, W. Zuo, L. Zhang, and P. Li. 2020. Deep CNNs Meet Global Covariance Pooling: Better Representation and Generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
- [82] Yunyun Wang and Songcan Chen. 2013. Safety-aware semi-supervised classification. *IEEE transactions on neural networks and learning systems* (2013).
- [83] Xiu-Shen Wei, Jianxin Wu, and Quan Cui. 2019. Deep learning for fine-grained image analysis: A survey. *arXiv preprint arXiv:1907.03069* (2019).
- [84] Tianjun Xiao, Yichong Xu, Kuiyuan Yang, Jiaying Zhang, Yuxin Peng, and Zheng Zhang. 2015. The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [85] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. 2019. Unsupervised Data Augmentation for Consistency Training. *arXiv preprint arXiv:1904.12848* (2019).
- [86] I. Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. 2019. Billion-scale semi-supervised learning for image classification. *arXiv:1905.00546* [cs.CV]
- [87] Ze Yang, Tiange Luo, Dong Wang, Zhiqiang Hu, Jun Gao, and Liwei Wang. 2018. Learning to navigate for fine-grained classification. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [88] Ting Yao, Yingwei Pan, Chong-Wah Ngo, Houqiang Li, and Tao Mei. 2015. Semi-supervised domain adaptation with subspace learning for visual recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [89] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. 2021. Deep learning for person re-identification: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021).
- [90] Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Seong Joon Oh, Youngjoon Yoo, and Junsuk Choe. 2008. CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. 6022–6031.
- [91] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. 2019. S4: Self-supervised semi-supervised learning. In *Proceedings of the IEEE international conference on computer vision*.
- [92] Fan Zhang, Guisheng Zhai, Meng Li, and Yizhao Liu. 2020. Three-branch and Mutil-scale learning for Fine-grained Image Recognition (TBMSL-Net). *arXiv preprint arXiv:2003.09150* (2020).
- [93] Han Zhang, Tao Xu, Mohamed Elhoseiny, Xiaolei Huang, Shaoting Zhang, Ahmed Elgammal, and Dimitris Metaxas. 2016. Spda-cnn: Unifying semantic part detection and abstraction for fine-grained recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [94] Jian Zhang, Runsheng Zhang, Yaping Huang, and Qi Zou. 2019. Unsupervised part mining for fine-grained image classification. *arXiv preprint arXiv:1902.09941* (2019).
- [95] Lianbo Zhang, Shaoli Huang, Wei Liu, and Dacheng Tao. 2019. Learning a mixture of granularity-specific experts for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [96] Ning Zhang, Jeff Donahue, Ross Girshick, and Trevor Darrell. 2014. Part-based R-CNNs for fine-grained category detection. In *European conference on computer vision*. Springer.
- [97] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. 2017. Learning multi-attention convolutional neural network for fine-grained image recognition. In *Proceedings of the IEEE international conference on computer vision*.
- [98] Heliang Zheng, Jianlong Fu, Zheng-Jun Zha, and Jiebo Luo. 2019. Learning deep bilinear transformation for fine-grained image representation. In *Advances in Neural Information Processing Systems*.
- [99] Xiaojin Zhu, Zoubin Ghahramani, and John D Lafferty. 2003. Semi-supervised learning using gaussian fields and harmonic functions. In *Proceedings of the 20th International conference on Machine learning (ICML-03)*.
- [100] Peiqin Zhuang, Yali Wang, and Yu Qiao. 2020. Learning Attentive Pairwise Interaction for Fine-Grained Classification. In *AAAI*.