

Semi-supervised learning made simple with self-supervised clustering

Anonymous CVPR submission

Paper ID 9110

Abstract

Self-supervised learning models have been shown to learn rich visual representations without requiring human annotations. However, in many real-world scenarios, labels are partially available, motivating a recent line of work on semi-supervised methods inspired by self-supervised principles. In this paper, we propose a conceptually simple yet empirically powerful approach to turn clustering-based self-supervised methods such as SwAV or DINO into semi-supervised learners. More precisely, we introduce a multi-task framework merging a supervised objective using ground-truth labels and a self-supervised objective relying on clustering assignments with a single cross-entropy loss. This approach may be interpreted as imposing the cluster centroids to be class prototypes. Despite its simplicity, we provide empirical evidence that our approach is highly effective and achieves state-of-the-art performance on CIFAR100 and ImageNet.

1. Introduction

In recent years, self-supervised learning became the dominant paradigm for unsupervised visual representation learning. In particular, much experimental evidence shows that augmentation-based self-supervision [3, 8–10, 14–16, 18, 21, 29, 32, 35, 71] can produce powerful representations of unlabeled data. Such models, although trained without supervision, can be naturally used for supervised downstream tasks via simple fine-tuning. However, the most suitable way to leverage self-supervision is perhaps by multi-tasking the self-supervised objective with a custom (possibly supervised) objective. Based on this idea, the community has worked on re-purposing self-supervised methods in other sub-fields of computer vision, as for instance in domain adaptation [25], novel class discovery [31, 77], continual learning [30] and semi-supervised learning [4, 13, 19, 72].

One of the areas that potentially benefits from the advancements in unsupervised representation learning is semi-supervised learning. This is mainly due to the fact

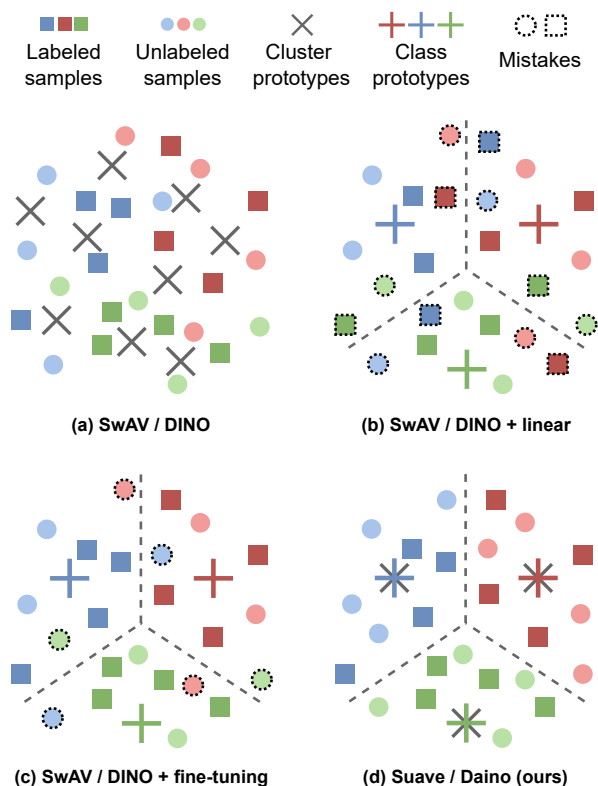


Figure 1. Schematic illustration of the motivation behind the proposed semi-supervised framework. (a) Self-supervised clustering methods like SwAV [15] and DINO [16] compute cluster centers that are not necessarily well aligned with semantic categories, but they do not require labeled data. (b) Adding a linear classifier provides class prototypes, but the labeled (and unlabeled) samples are not always correctly separated. (c) Fine-tuning can help separating labeled data. (d) Our framework learns cluster centroids that are aligned with center prototypes thus correctly separating both labeled and unlabeled data.

that in semi-supervised learning it is crucial to efficiently extract the information in the unlabeled set to improve the classification accuracy on the labeled classes. Indeed, several powerful semi-supervised methods [6, 13, 61, 72] were built upon this idea.

In the self-supervised learning landscape, arguably the

most successful methods belong to the clustering-based family, such as DeepCluster v2 [14], SwAV [15] and DINO [16]. These methods learn representations by contrasting predicted cluster assignments of correlated views of the same image. To avoid collapsed solutions and group samples together, they use simple clustering-based pseudo-labeling algorithms such as k-means and Sinkhorn-Knopp to generate the assignments. A peculiarity of this family of methods is the discretization of the feature space, that allows them to use techniques that were originally developed for supervised learning. Indeed, similar to supervised learning, the cross-entropy loss is adopted to compare the assignments, as they represent probability distributions over the set of clusters.

In this paper, we propose a new approach for semi-supervised learning based on the simple observation that clustering-based methods are amenable to be adapted to a semi-supervised learning setting: the cluster prototypes can be replaced with class prototypes learned with supervision and the same loss function can be used for both labeled and unlabeled data. In practice, semi-supervised learning can be achieved by multi-tasking the self-supervised and supervised objectives. This encourages the network to cluster unlabeled samples around the centroids of the classes in the feature space. By leveraging on these observations we propose a new framework for semi-supervised methods based on self-supervised clustering. We experiment with two instances of that framework: *Suave* and *Daino*, the semi-supervised counterparts of SwAV and DINO. These methods have several favorable properties: i) they are efficient at learning representations from unlabeled data since they are based on the top-performing self-supervised methods; ii) they extract relevant information for the semantic categories associated with the data thanks to the supervised conditioning; iii) they are easy to implement as based on the multi-tasking of two objectives. The motivation behind our proposal is also illustrated in Fig. 1. As shown in the figure, our multi-tasking approach enables to compute cluster centers that are aligned with class prototypes thus correctly separating both labeled and unlabeled data.

Our **contributions** can be summarized as follows:

- We propose a new framework for semi-supervised learning based on multi-tasking of a supervised objective on the labeled data and a clustering-based self-supervised objective on the unlabeled samples;
- We experiment with two representatives of such framework: *Suave* and *Daino*, semi-supervised extensions of SwAV [15] and DINO [16]. These methods, while simple to implement, are powerful and efficient semi-supervised learners;
- Our methods outperform state-of-the-art approaches, often relying on multiple *ad hoc* components, both on

common small scale (CIFAR100) and large scale (ImageNet) benchmarks, setting a new state-of-the-art on semi-supervised learning.

2. Related works

Self-supervised learning. The self-supervised literature has rapidly become extremely vast [38]. Excluding a recent trend involving denoising autoencoders [59] combined with vision transformers [27] like in MAE [34], the vast majority of existing methods are still relying on multiple views derived from each sample via data augmentation. These augmentation-based methods [3, 8–10, 14–16, 18, 21, 29, 32, 35, 71] are capable to learn rich representations during unsupervised pre-training, achieving performance comparable to their supervised counterparts when fine-tuned with the labels. Basically, these methods enforce correlated views of the same input to have coherent representations in latent space such that the model becomes invariant to the augmentations applied. This corresponds to maximizing the mutual information between views' representation and, in the literature, it has been done using different loss functions.

Contrastive-based methods [8, 18, 35, 37] define a loss based on noise-contrastive estimation [33] as instance discrimination [63] between positive (correlated) and negative views (the remaining samples in the mini-batch). A major drawback of these methods is that they require large mini-batches to have representative negative samples. To overcome this difficulty, consistency-based methods like [21, 32] propose to maximize the cosine similarity of positive pairs without considering negatives, whereas redundancy-reduction-based methods [9, 10, 29, 71] employ principled regularization terms to minimize features redundancy. For instance, in [71] a loss is introduced to minimize the cross-correlation between features of the positive pairs. Similarly, in [9, 10] the features learning process involves different regularization terms including the variance of the embeddings, the covariance and a redundancy reduction objective.

Clustering-based methods [3, 14–16, 45], instead, naturally discretize the latent space via deep clustering. They perform clustering either offline (using the whole dataset), as in DeepCluster v2 [14], PCL [45] and SeLa [3], or online (using mini-batches), as in SwAV [15] and DINO [16], and impose coherent clustering assignments of positive pairs through a cross-entropy. Noticeably, this loss contrasts targets (assignments) and predictions as in the standard supervised learning objective. Nonetheless, surprisingly, to the best of our knowledge there have been no previous attempts in the literature to exploit this favorable property in the semi-supervised scenario. Our work aims to fill this gap.

Semi-supervised learning. Semi-supervised approaches aim to exploit a limited amount of annotations over a large collection of unlabeled data. The most intuitive approach

to this task is perhaps Pseudo-Labels [43], based on self-training via pseudo-labeling: a model trained on labeled data generates categorical pseudo-labels for the unlabeled examples, which will be then integrated into the labeled set for the next model training. However, hard (categorical) labels easily exacerbate the classification bias of the training model, a phenomenon known as *confirmation bias* [2]. To counteract this issue, researchers have shown benefits from soft labels and confidence thresholding [2] as well as from different training strategies like co- and tri-training [17,48,51], model distillation [65], consistency regularization (see below) and model de-biasing [54,62].

Consistency regularization methods operate by introducing additional losses computed on unsupervised samples and enforce consistency of the network output under perturbation of the model and/or the input –augmentation-based similar to self-supervised– [7,42,47,49,53,57,58,64,75]. Recent approaches integrate pseudo-labeling techniques and consistency regularization. FixMatch [55] considers a fixed confidence threshold within a student-teacher architecture, filtering teacher pseudo-labels. [11,12,46] exploit MixUp [74] to improve class boundaries in low-density regions. Other works explore more advanced pseudo-labeling techniques based on adaptive confidence thresholds [66,73] and meta-learning [50], uncertainty estimation [22,60], latent structure regularization [44]. Recently, ConMatch [40] extended prediction consistency with self-supervised features consistency, while SimMatch [76] improved consistency regularization by applying it at both semantic-level and instance-level. Class-aware Contrastive Semi-Supervised Learning (CCSSL) is introduced in [67] to improve the quality of pseudo-labels in presence of unknown (out-of-distribution) classes.

As self-supervised methods have the ability to extract relevant information from unlabeled data, they can be leveraged to tackle the semi-supervised problem. Noticeably, self-supervised pre-training is found beneficial for many consistency regularization methods [13]. Moreover, semi-supervised self-supervised methods exist, which consider multi-tasking [61,72], pre-training distillation [19], exponential moving average normalization [13] and supervised contrastive [4]. Finally, PAWS [6] borrows some principles from self-supervised clustering, but, somewhat similarly to SimMatch [76], it assumes labeled instances to be class representative to which compute the distribution on. However, it does not exploit the multitasking of the self- and semi-supervised objectives nor takes advantage of deep clustering.

3. Clustering-based semi-supervised learning

Clustering-based self-supervision is typically adopted in unsupervised representation learning scenarios where no labeled data is available [14–16]. However, in this work, we

aim to take advantage of the few annotations available in the semi-supervised setting to learn even better representations. Our main intuition is to replace the cluster centroids with class prototypes learned with supervision. In this way, unlabeled samples will be clustered around the class prototypes, guided by the self-supervised clustering-based objective. To this end, we jointly optimize a supervised loss on the labeled data and a self-supervised loss on the unlabeled data. It turns out that using the same loss function (cross-entropy) is a very good choice since it promotes synergy between the two objectives and eases up the implementation. In the following, we formalize self-supervised clustering in Section 3.1. Then, we describe the details of our novel semi-supervised learning framework in Section 3.2 and show its application through two popular self-supervised approaches in Section 3.3.

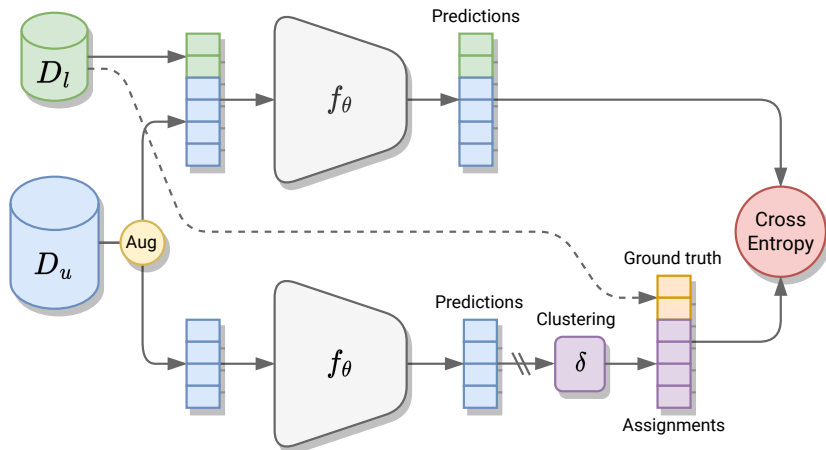
3.1. Clustering-based self-supervised learning

Given an unlabeled dataset $\mathcal{D}_u = \{\mathbf{x}_u^1, \dots, \mathbf{x}_u^N\}$, two correlated views of the same input image, $(\mathbf{x}, \hat{\mathbf{x}})$, are generated via data augmentation and embedded through an encoder network f_θ , composed of a backbone g , a projector h and a set of prototypes (or centroids) p implemented as a bias-free linear layer. Performing the forward pass of the backbone and the projector produces two latent representations, $(\mathbf{h}, \hat{\mathbf{h}})$ for the two correlated views $(\mathbf{x}, \hat{\mathbf{x}})$, respectively. Subsequently, the cluster centroids p are used to produce two sets of logits $(\mathbf{p}, \hat{\mathbf{p}})$ corresponding to each of the two latent representations. A softmax function σ can be applied to these logits to obtain probability distributions over the set of clusters, also referred to as cluster assignments.

In principle, the two assignments could be immediately compared in order to encourage the network to output similar predictions for similar inputs (correlated views). However, this may lead to degenerate solutions where all samples are assigned to the same cluster. To avoid collapse, simple clustering techniques are usually employed to embed priors into the cluster assignment process and regularize the training procedure. In practice, we compute:

$$\hat{\mathbf{y}} = \delta(\hat{\mathbf{p}}, \mathbf{c}, \epsilon), \quad (1)$$

where δ is the clustering technique, which takes as input some predicted logits $\hat{\mathbf{p}}$, a context \mathbf{c} and a temperature-like parameter ϵ that usually regulates the entropy of the assignment $\hat{\mathbf{y}}$ (sometimes addressed as pseudo-label). The context \mathbf{c} can be implemented in different flavors depending on the nature of δ . For instance, for offline clustering \mathbf{c} is represented by the features of all the samples in the dataset, while for online clustering the context might contain the features of the current batch or simply a running mean of the overall distribution. The pseudo-label can be categorical (hard) or soft. Empirically soft cluster assignments were shown to yield superior performance [15].



Pseudo-code for Suave and Daino.

```

# x_l, x_u: labeled and unlabeled samples
# label: one-hot ground truth for x_l
def training_step(x_l, x_u, label):
    # correlated views
    x1_u, x2_u = aug(x_u), aug(x_u)

    # forward pass
    p_l = f(x_l)
    p1_u, p2_u = f(x1_u), f(x2_u)

    # compute assignments via clustering
    assign = delta(p2_u.detach())

    # concatenate predictions and targets
    pred = cat(p_l, p1_u)
    target = cat(label, assign)

    # compute loss
    loss = cross_entropy(pred, target)
    return loss

```

Figure 2. Left: overview of the proposed family of semi-supervised methods. Right: PyTorch-like pseudo-code for Suave and Daino.

The output of the clustering technique is then used as a target in the cross-entropy loss:

$$\ell(\mathbf{s}, \hat{\mathbf{y}}) = - \sum_{k=1}^K \hat{y}_k \log(s_k), \quad (2)$$

where $\mathbf{s} = \sigma(\mathbf{p}/\tau)$ has gone through softmax normalization with temperature τ , and K is the number of clusters. Note that a stop-gradient operation is performed, so that the gradient is not propagated through the pseudo-label. It is worth mentioning that the cross-entropy loss in Eq. 2 is asymmetric, but, it can be made symmetric by swapping \mathbf{x} with $\hat{\mathbf{x}}$.

Intuitively, this loss leverages cluster assignments as a proxy to minimize the distance between latent representations $(\mathbf{h}, \hat{\mathbf{h}})$ of augmented views of the same image. As a by-product, the objective also learns a set of cluster prototypes encoded in the last linear layer p , which represent semantic information in the latent space. However, these clusters have no guarantee to be aligned with the true semantic categories represented in the dataset. Nonetheless, the discretization of the feature space is particularly interesting in the context of semi-supervised learning, as described in the following.

3.2. Our semi-supervised learning framework

In the semi-supervised scenario, we assume having access to a partially labeled dataset, $\mathcal{D} = \mathcal{D}_u \cup \mathcal{D}_l$, usually with $|\mathcal{D}_l| \ll |\mathcal{D}_u|$ and \mathcal{D}_l contains a number of C known classes. We propose to exploit clustering-based self-supervised models described in Sec. 3.1 as a base to extract information from unlabeled data, and we extend them to take advantage of the labeled samples.

As mentioned, the main drawback of clustering-based self-supervised methods in the context of semi-supervised

learning is that there are no guarantees that the stochastic optimization process will organize the clusters in the feature space according to the class labels. Indeed, they may be completely misaligned with the actual distribution of the classes. This also potentially hinders the effectiveness of the representations, as some prototypes may encode spurious correlations in the data. An ideal scenario is one where the clusters are centered on the actual class centroids such that the label can be propagated to the unlabeled samples by means of the clustering function. This will generate positive feedback that progressively transfers information from the labeled set into the unlabeled one, thus improving the feature representations learned by the network.

We propose to condition the cluster prototypes to encode the class information by resorting to multi-tasking of the self-supervised and supervised objectives. This can be achieved by optimizing the same loss function in Eq. 2 while replacing the pseudo-label with the ground-truth label when available:

$$\ell(\mathbf{s}, \bar{\mathbf{y}}), \quad \bar{\mathbf{y}} = \begin{cases} \mathbf{y}, & \mathbf{x} \in \mathcal{D}_l \\ \hat{\mathbf{y}}, & \mathbf{x} \in \mathcal{D}_u. \end{cases} \quad (3)$$

Since the linear layer p that contains the prototypes is now shared between the two objectives, we set $K = C$ to have a matching label space, although in principle this is not a hard constraint (as shown in [31]). In a nutshell, in our framework we compute the forward pass for both labeled and unlabeled samples, concatenate the associated predictions and the targets and apply the cross-entropy loss simultaneously as described in Fig. 2. We empirically demonstrate in Sec. 5 that, despite its simplicity, our framework equipped with these design choices is a strong semi-supervised learner.

3.3. Suave and Daino

Our proposed framework described above can convert any clustering-based self-supervised method into a semi-supervised learner, without dropping, adding or replacing any architectural components and reusing the same loss function. We select two representative self-supervised methods to showcase our framework: SwAV [15] and DINO [16], whose semi-supervised extensions we name **Suave** and **Daino**. This choice is motivated by their superior representation learning capabilities and ease of use. In particular, both are online clustering methods, which means that the cluster assignments can be computed on-the-fly without accessing the whole dataset simultaneously. This is a great advantage, especially for large scale datasets. For these reasons, we discard offline clustering methods like DeepCluster [14] and PCL [45], while in principle our approach can be also applied to them.

Suave. SwAV [15], following [3], casts the pseudo-label assignment problem as an instance of the optimal transport problem and proposes the swapped prediction task where the assignment of a view is predicted from the representation of another view. Simply put, SwAV generates pseudo-labels such that each cluster is approximately equally represented in the current batch, preventing the network from falling into degenerate solutions. This is especially convenient as we only need the information in the current batch for clustering. In light of our proposed framework, reusing Eq. 1, in Suave the target for the first sample in the batch can be obtained as follows (the same reasoning can be trivially applied to the other samples in the batch):

$$\hat{\mathbf{y}}_1 = \delta(\hat{\mathbf{p}}_1, [\hat{\mathbf{p}}_2, \dots, \hat{\mathbf{p}}_B], \epsilon) \quad (4)$$

where the context $\mathbf{c} = [\hat{\mathbf{p}}_2, \dots, \hat{\mathbf{p}}_B]$ contains all the logits in the batch except $\hat{\mathbf{p}}_1$. Now we define $\hat{\mathbf{P}} = [\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2, \dots, \hat{\mathbf{p}}_B]$ and $\hat{\mathbf{Y}} = [\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2, \dots, \hat{\mathbf{y}}_B]^\top$, where $\hat{\mathbf{Y}}$ is the matrix that holds the unknown pseudo-labels of the whole batch. The clustering function δ will return the first column of $\hat{\mathbf{Y}}$ that is found by solving:

$$\hat{\mathbf{Y}} = \max_{\mathbf{Y} \in \Gamma} \text{Tr}(\mathbf{Y} \hat{\mathbf{P}}) + \epsilon \text{H}(\mathbf{Y}), \quad (5)$$

where $\epsilon > 0$ is the temperature-like parameter mentioned in Eq. 1, H is the entropy function, Tr is the trace operator, and Γ is the transportation polytope defined as:

$$\Gamma = \{\mathbf{Y} \in \mathbb{R}_+^{C \times B} \mid \mathbf{Y} \mathbf{1}_B = \frac{1}{C} \mathbf{1}_C, \mathbf{Y}^\top \mathbf{1}_C = \frac{1}{B} \mathbf{1}_B\}. \quad (6)$$

These constraints enforce that, on average, each cluster is selected $\frac{B}{C}$ times in each batch, automatically ensuring de-biased assignments. The solution to Eq. 5 is obtained using the Sinkhorn-Knopp algorithm [1, 23].

Daino. DINO [16] aims at further simplifying the pipeline described above. Instead of using optimal transport on the

predictions of the current batch it uses two practical tricks to avoid collapse: a momentum encoder and a pseudo-labeling strategy based on centering and sharpening. A momentum encoder is a “slow” version of the encoder f_θ updated using exponential moving average (EMA). After each gradient step on θ the parameters ϕ of the momentum encoder f_ϕ are updated as follows:

$$\phi \leftarrow \eta \psi + (1 - \eta) \theta, \quad (7)$$

where η is a rate parameter. The rationale behind this choice is that the momentum encoder serves as a teacher producing more stable representations throughout training, improving the optimization process. The teacher is used at every iteration to generate the logits $\hat{\mathbf{p}} = f_\phi(\hat{\mathbf{x}})$ which in turn are input to the clustering function δ in Eq. 1 to obtain the target assignment for Daino:

$$\hat{\mathbf{y}} = \delta(\hat{\mathbf{p}}, \gamma, \epsilon) = \sigma\left(\frac{\hat{\mathbf{p}} - \gamma}{\epsilon}\right), \quad (8)$$

where σ and ϵ are the softmax function and a temperature coefficient (here used to sharpen the distribution) respectively. The context $\mathbf{c} = \gamma$ in this case is a centering vector that approximates and de-biases the overall distribution of the data over the clusters and is also updated using EMA:

$$\gamma \leftarrow \mu \gamma + (1 - \mu) \frac{1}{B} \sum_{b=0}^B \hat{\mathbf{p}}_b, \quad (9)$$

where μ adjusts the rate of the update. In brief, centering prevents one dimension to dominate but encourages high-entropy outputs, while sharpening does the opposite. Empirical evidence shows that this is enough to avoid collapse.

4. Implementation details

Architectures. For large-scale datasets (e.g., ImageNet), we adopt a ResNet50 [36] backbone; while for small-scale datasets (e.g., CIFAR100) we train a Wide ResNet (WRN-28-8) [70]. The convolutional backbone is followed by a projection head consisting of a multi-layer perceptron (with batch normalization in the hidden layers). We use 2 layers as in [6, 15]. After that, we perform L2-normalization and compute the predictions using a bias-free L2-normalized linear layer corresponding to the prototypes. We set the number of prototypes equal to the number of classes of the dataset at hand. Moreover, we perform an online linear evaluation at different depths of the network to identify which layer learns the best representations; we append a detached classification head on top of the backbone, the first and the second layer of projection. We remark that such heads do not impact the efficiency of the training process.

Semi-supervised pre-training. We pre-train our models using LARS [68] optimizer with linear warmup plus cosine learning rate schedule. Also, we adopt weight decay

regularization. The backbone layers of the models can be initialized with self-supervised checkpoints of SwAV and DINO. At each training iteration, we sample a mini-batch composed of unlabeled images¹ and labeled images. Note that we count a training epoch considering a full pass over the unlabeled dataset. We optimize the cross-entropy loss of Eq. 3, re-weighting the labeled and unlabeled terms by their frequency in the batch. Moreover, we soften the supervised targets with label smoothing to mitigate overfitting. For the pseudo-labeling, in Suave, we re-use the same parameters for the Sinkhorn-Knopp as in SwAV; in Daino, we inherit the hyperparameters for the centering and the momentum encoder, whereas we tune the sharpening coefficient. More details are in the supplementary material.

Data augmentation. Images are augmented differently based on whether they are unlabeled or labeled. For the unlabeled images, we follow the default self-supervised augmentations of SwAV and DINO, while for the labeled, we adopt lighter augmentation Inception-style (random crop and flip) [56] and color distortion (jittering and greyscale). Note that it is important not to over-distort the labeled images as they are needed to align the clusters and classes.

To boost the self-supervised feature learning, we employ the multi-crop [15] augmentation scheme for unlabeled images. From each input image, we derive two global views from larger and higher-resolution crops and multiple smaller views from tighter crops. This is common practice in self-supervised learning. Similar to SwAV and DINO, we compute a clustering assignment only for the global views and use them as targets for the smaller ones. All the multi-crop views are taken into account when weighting the loss.

Another augmentation technique we empirically found to be useful is based on the combination of CutMix [69] and MixUp [74]. We apply it to both unlabeled (global views only) and labeled images, but separately. For the unlabeled images, we interpolate the learned clustering assignments due to the lack of ground-truth labels. This augmentation allows for shifting decision boundaries to low-density regions of the data [12, 58]. We mix the whole batch at every iteration, but to not over-regularize, we concatenate the mixed images to the current batch, instead of substituting it.

Semi-supervised fine-tuning. Since self-supervised methods are trained with strong augmentations, it is preferable after pre-training to fine-tune them to slightly improve the performance. We discover that a semi-supervised fine-tuning recipe works better than the typical fully supervised fine-tuning e.g., [4]. In practice, we just keep training the model with the same objective (our semi-supervised clustering-based loss) for a few more epochs, while relaxing some of the stronger augmentations adopted during pre-training, i.e., disable multi-crop and color distortions.

¹We sample unlabeled examples from dataset \mathcal{D} , instead of \mathcal{D}_u .

5. Experiments

5.1. Experimental protocol

Datasets. We perform our experiments using two common datasets, i.e., CIFAR100 [41] and ImageNet-1k [26]. In the semi-supervised setting the training set of each of these datasets is split into two subsets, one labeled and one unlabeled. For CIFAR100, which is composed of 50K images equally distributed into 100 classes, we investigate three splits as in [55], retaining 4 (0.8%), 25 (5%), and 100 (20%) labels per class, resulting in total to 400, 2500, and 10000 labeled images, respectively. For ImageNet-1k ($\sim 1.3\text{K}$ images per class, 1K classes), we adopt the same two splits of [19] using 1% and 10% of the labels. In both datasets, we evaluate the performance of our method by computing top-1 accuracy on the respective validation/test sets.

Baselines. We compare our methods, Suave and Daino, with state of the art methods from the semi-supervised literature (see Section 2). In particular, we compare against hybrid consistency regularization methods like SimMatch [76] and ConMatch [40] (and others [46, 60, 73]), and methods that are derived from self-supervised approaches, like PAWS [6] and S⁴L-Rot [72]. We also compare with recent debiasing-based pseudo-labeling methods like DebiasPL [62]. For the sake of fairness, we leave out methods using larger architectures or pre-trained on larger datasets, e.g., DebiasPL with CLIP [52] and SimCLR v2 [19].

5.2. Results

First, we demonstrate the effectiveness of our semi-supervised framework in a small-scale dataset using the CIFAR100 benchmark. Then, we select the best performing model and evaluate it at large-scale on ImageNet-1k (see Section 5.2.1). Finally, we ablate the different components of our models (see Section 5.2.2).

5.2.1 Comparison with the state of the art

Results on CIFAR100. Table 1 shows a comparison between our methods and several semi-supervised approaches in the literature. In particular, we compare to consistency regularization semi-supervised methods (see Section 2) like ConMatch [40] and SimMatch [76], which are the strongest methods on this dataset. First, from the table we observe that both Suave and Daino achieve high performance in all the three splits (400, 2500, 10000). Daino obtains results comparable to the best competitors, while Suave outperforms all the baselines and beats the best methods CCSSL [67] and SimMatch [76] by +2.4%, +1.3%, +1.5% in the three settings, respectively. Overall, these results clearly demonstrate how our clustering-based semi-supervised learning methods achieves state-of-the-art performance without requiring any *ad-hoc* confidence thresh-

Table 1. Comparison with the state-of-the-art on CIFAR100.

Method	Acc@1		
	400	2500	10000
PI-Model [42]	-	42.8	62.1
Mean Teacher [57]	-	46.1	64.2
MixMatch [12]	32.4	60.2	72.2
UDA [64]	53.6	72.3	77.5
ReMixMatch [11]	55.7	72.6	77.0
FixMatch [55]	50.1	71.4	76.8
Dash [66]	55.2	72.8	78.0
CoMatch [44]	60.0	73.0	78.2
Meta Pseudo Labels [50]	55.8	72.3	77.5
FlexMatch [73]	60.1	73.5	78.1
FixMatch+DM [46]	59.8	74.1	79.6
NP-Match [60]	61.1	74.0	78.8
ConMatch [40]	61.1	74.6	-
SimMatch [76]	62.2	74.9	79.4
CCSSL [67]	61.2	75.7	80.1
Daino	61.1	75.2	79.2
Suave	64.6	77.0	81.6

olds for pseudo-labels as in most recent consistency regularization methods. Interestingly, comparing Suave and Daino with their self-supervised counterparts, SwAV and DINO, a remarkable improvement is achieved. In fact SwAV and DINO obtain an accuracy of 64.9 and 66.8 [24], respectively, when linearly evaluated using 100% of the labels after self-supervised pre-training.

Results on ImageNet-1k. We also perform large scale experiments on ImageNet-1k considering the label splits of 1% and 10% as common in previous works [55, 76]. Due to the lack of sufficient computational resources to run this benchmark on two methods (Suave and Daino) at scale, we select only one of the two. The choice falls on Suave for two reasons: i) it was found superior to Daino on CIFAR100 and ii) DINO on ImageNet is most frequently trained with ViTs [27], which requires more memory and finer hyperparameter tuning wrt. ResNets.

We compare against three families of methods, self-supervised-inspired semi-supervised approaches [6, 72], consistency regularization methods [13, 44, 50, 55, 64, 76], and debiasing-based methods [62], and present results in Table 2 and Figure 3. To provide full context to our results, in Table 2, we also report the performance of self-supervised models when simply fine-tuned with the labels.

Suave obtains performances that are comparable with state-of-the-art methods on ImageNet-1k. In particular, when compared against related methods (e.g., DebiasPL, SimMatch, FixMatch-EMAN, CoMatch) with similar batch-size and number of epochs, it obtains the best score (+0.6%p) on the 10% setting and the third best score on the 1% setting, -1.0%p from the best baseline Sim-

Table 2. Comparison with the state-of-the-art on ImageNet-1k. All the models reported use ResNet-50 unless differently specified. In the first and the second column, we indicate within brackets whether a model is initialized from a self-supervised checkpoint and the number of epochs of that pre-training. For brevity, we refer to FixMatch as FM. *SimMatch does not report the batch size, and its value is inferred from the public repository. † refers to epochs on labeled data.

Method	Epochs	Batch size		Acc@1	
		Unlab.	Lab.	10%	1%
<i>with similar batch size and number of epochs</i>					
S ⁴ L-Rotation [72]	200 [†]	256	256	61.4	-
UDA [55]	~480	15360	512	68.8	-
FM-DA (MoCo v2) [44]	(800) 400	640	160	72.2	59.9
PAWS [6]	100	256	1680	70.2	-
CoMatch (MoCo v2) [44]	(800) 400	640	160	73.7	67.1
FM-EMAN (MoCo-EMAN) [13]	(800) 300	320	64	74.0	63.0
SimMatch [76]	400	320*	64*	74.4	67.2
DebiasPL (MoCo-EMAN) [62]	(800) 50	640	128	-	65.3
	(800) 200	640	128	-	66.5
Suave	(100) 100	256	128	73.6	63.8
	(200) 100	256	128	74.3	65.0
	(800) 100	256	128	75.0	66.2
<i>with larger batch size or number of epochs</i>					
Meta Pseudo Labels [50]	~800	2048	2048	73.9	-
PAWS [6]	100	4096	6720	73.9	63.8
	200	4096	6720	75.0	66.1
	300	4096	6720	75.5	66.5
DebiasPL (MoCo-EMAN) [62]	(800) 300	1280	256	-	67.1
<i>self-supervised pre-training with fine-tuning ([‡]linear evaluation)</i>					
MoCo v2 [20]	(800)	256	-	66.1	49.8
SimCLR v2 [18]	(1000)	4096	-	68.4	57.9
BYOL [32]	(1000)	2048	-	68.8	53.2
MoCo-EMAN [13]	(800)	256	-	68.1	57.4
SwAV [13]	(800)	4096	-	70.2	53.9
NNCLR [28]	(1000)	4096	-	70.2	56.4
Barlow Twins [71]	(1000)	2048	-	69.7	55.0
FNC [37]	(1000)	4096	-	71.1	63.7
DINO (ViT-S) [‡] [16]	(800)	1024	-	72.2	64.5
MSN (ViT-S) [‡] [5]	(800)	1024	-	-	67.2

Match [76]. Importantly, other approaches like PAWS [4] require larger batches to obtain comparable results. This aspect is even more evident by looking at Fig. 3 (bottom), where PAWS significantly underperforms with respect to our method when decreasing the batch size. This behavior can be ascribed to the fact that PAWS adopts a k-nearest neighbor approach on the labeled instances to generate the assignment vectors (pseudo-labels), and thus it requires a sufficiently high number of labeled examples to well represent the class distributions. In contrast, Suave generates the assignments using the learnable cluster/class prototypes, which are a fixed number independent of the batch size.

Figure 3 (top) compares Suave with the best methods in Table 2, highlighting the fast convergence of our method. Performing 100 epochs of semi-supervised pre-training is enough to achieve comparable results to other related methods like FixMatch-EMAN, PAWS, and SimMatch, which require at least twice the semi-supervised epochs. It is worth




Figure 3. Comparison with related works on the training efficiency in terms of pre-training epochs (top) and size of the mini-batches (bottom). Plotted results refer to the ImageNet 10% split.

noting that Suave naturally benefits from self-supervised initialization, as it basically shares the objective with the self-supervised counterpart SwAV. Indeed, by looking at Table 2 we can observe that better SwAV checkpoints lead to higher accuracy. At the same time, the difference between Suave (SwAV-100) and Suave (SwAV-800) is relatively small, 1.4% and 2.4% on 10% and 1%, respectively.

5.2.2 Ablation study

Here, we assess the importance of the main components of our method: (i) the multi-task training objective, (ii) the adopted fine-tuning strategy, and (iii) the quality of representations at different depth of the network.

In Table 3, we demonstrate the importance of our multi-tasking strategy. Suave significantly outperforms vanilla SwAV, even in its improved version, i.e., SwAV (repro) obtained with the fine-tuning protocol borrowed from PAWS [6]. More importantly, when comparing Suave to SwAV+CT (SwAV where the self-supervised loss is multi-tasked with a supervised contrastive loss [39]) we still observe that our method is far superior. This clearly indicates that it is better to exploit the available labels to condition the prototypes rather than using them to sample positive instances for contrastive training.

Table 4 provides evidence that fine-tuning Suave with our semi-supervised fine-tuning strategy (see Section 4) is more effective than adopting the classical fully supervised

Table 3. Impact of our multi-task training strategy. Legend: UPT = “unsupervised pre-training”, MT = “multi-tasking” and SL = “same loss for all training phases”.

Method	UPT	MT	SL	Epochs	Acc@1	
					10%	1%
SwAV	✓		✓	800	70.2	53.9
SwAV (repro)	✓		✓	800	72.3	57.0
SwAV+CT [4]	✓	✓		(400)30	70.8	-
Suave	✓	✓	✓	(800)100	75.0	66.2

Table 4. Impact of our semi-supervised fine-tuning strategy.

Method	Labels (%)	Semi-supervised pre-training	Supervised fine-tuning	Semi-supervised fine-tuning
Suave	1%	64.1	64.8	66.2
	10%	73.0	74.8	75.0

Table 5. Online linear evaluation on projection layers.

Method	Labels (%)	Backbone	Projection layer 1	Projection layer 2
Suave	1%	65.2	66.2	65.1
	10%	74.2	75.0	74.8

recipe. We believe that feeding the model with unlabeled data during the fine-tuning phase allows the model to better fit the real distribution. As shown in the table, the semi-supervised fine-tuning recipe enables a larger improvement over the performance after semi-supervised pre-training.

Finally, Table 5 shows how the quality of the learned representations changes at different depths of the projector. It turns out that regardless of the label split, we obtain the most discriminative representations at the first layer of the projection. This is consistent with what was observed by other works that use similar augmentations (e.g., PAWS [6]). We ascribe this behaviour to the fact that the last layer tries to build as much invariance as possible to strong augmentations, which hurts the classification accuracy.

6. Conclusion

We have presented a novel approach for semi-supervised learning. Our framework leverages from clustering-based self-supervised methods and adopts a multi-task objective to combine a supervised loss with the cross-entropy loss typically adopted for clustering assignments in [15,16]. Despite its simplicity, we demonstrate that our approach is highly effective, setting a new state of the art for semi-supervised learning on CIFAR-100 and ImageNet. Future works will assess the potential of our framework with other backbones, e.g., to evaluate the performance of Daino with transformers. We also believe that other problems, such as unsupervised domain adaptation, may benefit from the findings of this work. We also leave this to future investigation.

References

- [1] Elad Amrani, Leonid Karlinsky, and Alex Bronstein. Self-supervised classification network. In *European Conference on Computer Vision*, pages 116–132. Springer, 2022. 5
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E O’Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2020. 3
- [3] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations*, 2020. 1, 2, 5
- [4] Mahmoud Assran, Nicolas Ballas, Lluís Castrejon, and Michael Rabbat. Supervision accelerates pre-training in contrastive semi-supervised learning of visual representations. *arXiv preprint arXiv:2006.10803*, 2020. 1, 3, 6, 7, 8
- [5] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *European Conference on Computer Vision*, pages 456–473. Springer, 2022. 7
- [6] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael Rabbat. Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 1, 3, 5, 6, 7, 8
- [7] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. There are many consistent explanations of unlabeled data: Why you should average. In *International Conference on Learning Representations*, 2019. 3
- [8] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *Advances in neural information processing systems*, 32, 2019. 1, 2
- [9] Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-invariance-covariance regularization for self-supervised learning. In *International Conference on Learning Representations*, 2022. 1, 2
- [10] Adrien Bardes, Jean Ponce, and Yann LeCun. Vicregl: Self-supervised learning of local visual features. In *NeurIPS*, 2022. 1, 2
- [11] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. In *International Conference on Learning Representations*, 2020. 3, 7
- [12] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32, 2019. 3, 6, 7
- [13] Zhaowei Cai, Avinash Ravichandran, Subhansu Maji, Charles Fowlkes, Zhuowen Tu, and Stefano Soatto. Exponential moving average normalization for self-supervised and semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 194–203, 2021. 1, 3, 7
- [14] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, 2018. 1, 2, 3, 5
- [15] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020. 1, 2, 3, 5, 6, 8
- [16] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. 1, 2, 3, 5, 7, 8
- [17] Dongdong Chen, Wei Wang, Wei Gao, and Zhi-Hua Zhou. Tri-net for semi-supervised deep learning. In *Proceedings of twenty-seventh international joint conference on artificial intelligence*, pages 2014–2020, 2018. 3
- [18] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 1, 2, 7
- [19] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33:22243–22255, 2020. 1, 3, 6
- [20] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 7
- [21] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 1, 2
- [22] Yanbei Chen, Xiatian Zhu, Wei Li, and Shaogang Gong. Semi-supervised learning under class distribution mismatch. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3569–3576, 2020. 3
- [23] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013. 5
- [24] Victor Guilherme Turrissi da Costa, Enrico Fini, Moin Nabi, Nicu Sebe, and Elisa Ricci. solo-learn: A library of self-supervised methods for visual representation learning. *J. Mach. Learn. Res.*, 23:56–1, 2022. 7
- [25] Victor G Turrissi da Costa, Giacomo Zara, Paolo Rota, Thiago Oliveira-Santos, Nicu Sebe, Vittorio Murino, and Elisa Ricci. Dual-head contrastive domain adaptation for video action recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1181–1190, 2022. 1
- [26] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255, 2009. 6
- [27] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, 918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

- 972 Mostafa Dehghani, Matthias Minderer, Georg Heigold, Syl-
973 vain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is
974 worth 16x16 words: Transformers for image recognition at
975 scale. In *International Conference on Learning Representations*, 2021. 2, 7
976
- [28] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre
977 Sermanet, and Andrew Zisserman. With a little help from my
978 friends: Nearest-neighbor contrastive learning of visual rep-
979 resentations. In *Proceedings of the IEEE/CVF International
980 Conference on Computer Vision*, pages 9588–9597, 2021. 7
981
- [29] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto,
982 and Nicu Sebe. Whitening for self-supervised representation
983 learning. In *International Conference on Machine Learning*,
984 pages 3015–3024. PMLR, 2021. 1, 2
985
- [30] Enrico Fini, Victor G Turrisi da Costa, Xavier Alameda-
986 Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-
987 supervised models are continual learners. In *Proceedings of
988 the IEEE/CVF Conference on Computer Vision and Pattern
989 Recognition*, 2022. 1
990
- [31] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun
991 Zhong, Moin Nabi, and Elisa Ricci. A unified objective for
992 novel class discovery. In *Proceedings of the IEEE/CVF In-
993 ternational Conference on Computer Vision*, 2021. 1, 4
994
- [32] Jean-Bastien Grill, Florian Strub, Florent Althé, Corentin
995 Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch,
996 Bernardo Avila Pires, Zhaohan Guo, Mohammad Ghesh-
997 laghi Azar, et al. Bootstrap your own latent-a new approach
998 to self-supervised learning. *Advances in neural information
999 processing systems*, 33:21271–21284, 2020. 1, 2, 7
1000
- [33] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive
1001 estimation: A new estimation principle for unnormalized
1002 statistical models. In *Proceedings of the thirteenth inter-
1003 national conference on artificial intelligence and statistics*,
1004 pages 297–304, 2010. 2
1005
- [34] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr
1006 Dollár, and Ross Girshick. Masked autoencoders are scalable
1007 vision learners. In *Proceedings of the IEEE/CVF Conference
1008 on Computer Vision and Pattern Recognition*, pages 16000–
1009 16009, 2022. 2
1010
- [35] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross
1011 Girshick. Momentum contrast for unsupervised visual rep-
1012 resentation learning. In *Proceedings of the IEEE/CVF confer-
1013 ence on computer vision and pattern recognition*, pages
1014 9729–9738, 2020. 1, 2
1015
- [36] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun.
1016 Deep residual learning for image recognition. In *Proceed-
1017 ings of the IEEE conference on computer vision and pattern
1018 recognition*, pages 770–778, 2016. 5
1019
- [37] Tri Huynh, Simon Kornblith, Matthew R Walter, Michael
1020 Maire, and Maryam Khademi. Boosting contrastive self-
1021 supervised learning with false negative cancellation. In *Pro-
1022 ceedings of the IEEE/CVF Winter Conference on Applica-
1023 tions of Computer Vision*, pages 2785–2795, 2022. 2, 7
1024
- [38] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki
1025 Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey
on contrastive self-supervised learning. *Technologies*, 9(1):2,
2020. 2
- [39] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna,
1026 Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and
1027 Dilip Krishnan. Supervised contrastive learning. *Advances
1028 in Neural Information Processing Systems*, 2020. 8
1029
- [40] Jiwon Kim, Youngjo Min, Daehwan Kim, Gyuseong Lee,
1030 Junyoung Seo, Kwangrok Ryoo, and Seungryong Kim. Con-
1031 match: Semi-supervised learning with confidence-guided
1032 consistency regularization. In *European Conference on
1033 Computer Vision*, 2022. 3, 6, 7
1034
- [41] Alex Krizhevsky, Geoffrey Hinton, et al. *Learning multiple
1035 layers of features from tiny images*. Toronto, ON, Canada,
1036 2009. 6
1037
- [42] Samuli Laine and Timo Aila. Temporal ensembling for semi-
1038 supervised learning. In *International Conference on Learn-
1039 ing Representations*, 2017. 3, 7
1040
- [43] Dong-Hyun Lee et al. Pseudo-label: The simple and effi-
1041 cient semi-supervised learning method for deep neural net-
1042 works. In *Workshop on challenges in representation learn-
1043 ing, ICML*, 2013. 3
1044
- [44] Junnan Li, Caiming Xiong, and Steven CH Hoi. Comatch:
1045 Semi-supervised learning with contrastive graph regulariza-
1046 tion. In *Proceedings of the IEEE/CVF International Confer-
1047 ence on Computer Vision*, pages 9475–9484, 2021. 3, 7
1048
- [45] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Pro-
1049 totypical contrastive learning of unsupervised representa-
1050 tions. In *International Conference on Learning Representa-
1051 tions*, 2021. 2, 5
1052
- [46] Zicheng Liu, Siyuan Li, Ge Wang, Cheng Tan, Lirong Wu,
1053 and Stan Z Li. Decoupled mixup for data-efficient learning.
1054 *arXiv preprint arXiv:2203.10761*, 2022. 3, 6, 7
1055
- [47] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken
1056 Nakae, and Shin Ishii. Distributional smoothing with virtual
1057 adversarial training. In *International Conference on Learn-
1058 ing Representations*, 2016. 3
1059
- [48] Islam Nassar, Samitha Herath, Ehsan Abbasnejad, Wray
1060 Buntine, and Gholamreza Haffari. All labels are not created
1061 equal: Enhancing semi-supervision via label grouping and
1062 co-training. In *Proceedings of the IEEE/CVF Conference
1063 on Computer Vision and Pattern Recognition*, pages 7241–
1064 7250, 2021. 3
1065
- [49] Sungrae Park, JunKeon Park, Su-Jin Shin, and Il-Chul Moon.
1066 Adversarial dropout for supervised and semi-supervised
1067 learning. In *Proceedings of the AAAI conference on artifi-
1068 cial intelligence*, volume 32, 2018. 3
1069
- [50] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V Le. Meta
1070 pseudo labels. In *Proceedings of the IEEE/CVF Conference
1071 on Computer Vision and Pattern Recognition*, 2021. 3, 7
1072
- [51] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan
1073 Yuille. Deep co-training for semi-supervised image recogni-
1074 tion. In *Proceedings of the european conference on computer
1075 vision (eccv)*, pages 135–152, 2018. 3
1076
- [52] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya
1077 Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry,
1078 Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learn-
1079 ing transferable visual models from natural language super-
vision. In *International Conference on Machine Learning*,
pages 8748–8763. PMLR, 2021. 6

- [53] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems*, 29, 2016. 3
- [54] Hugo Schmutz, Olivier Humbert, and Pierre-Alexandre Mattei. Don't fear the unlabelled: safe deep semi-supervised learning via simple debiasing. *arXiv preprint arXiv:2203.07512*, 2022. 3
- [55] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 3, 6, 7
- [56] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015. 6
- [57] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 3, 7
- [58] Vikas Verma, Kenji Kawaguchi, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. In *Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI*, pages 3635–3641, 2019. 3, 6
- [59] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(12), 2010. 2
- [60] Jianfeng Wang, Thomas Lukasiewicz, Daniela Massiceti, Xiaolin Hu, Vladimir Pavlovic, and Alexandros Neophytou. Np-match: When neural processes meet semi-supervised learning. In *International Conference on Machine Learning*, pages 22919–22934. PMLR, 2022. 3, 6, 7
- [61] Xiao Wang, Daisuke Kihara, Jiebo Luo, and Guo-Jun Qi. Enaet: A self-trained framework for semi-supervised and supervised learning with ensemble transformations. *IEEE Transactions on Image Processing*, 30:1639–1647, 2020. 1, 3
- [62] Xudong Wang, Zhirong Wu, Long Lian, and Stella X Yu. Debaised learning from naturally imbalanced pseudo-labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3, 6, 7
- [63] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018. 2
- [64] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33, 2020. 3, 7
- [65] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. 3
- [66] Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. Dash: Semi-supervised learning with dynamic thresholding. In *International Conference on Machine Learning*. PMLR, 2021. 3, 7
- [67] Fan Yang, Kai Wu, Shuyi Zhang, Guannan Jiang, Yong Liu, Feng Zheng, Wei Zhang, Chengjie Wang, and Long Zeng. Class-aware contrastive semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14421–14430, 2022. 3, 6, 7
- [68] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. 5
- [69] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019. 6
- [70] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 5
- [71] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. 1, 2, 7
- [72] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1476–1485, 2019. 1, 3, 6, 7
- [73] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34:18408–18419, 2021. 3, 6, 7
- [74] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. 3, 6
- [75] Liheng Zhang and Guo-Jun Qi. Wcp: Worst-case perturbations for semi-supervised deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3912–3921, 2020. 3
- [76] Mingkai Zheng, Shan You, Lang Huang, Fei Wang, Chen Qian, and Chang Xu. Simmatch: Semi-supervised learning with similarity matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14471–14481, 2022. 3, 6, 7
- [77] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10867–10875, 2021. 1