

All Labels Are Not Created Equal: Enhancing Semi-supervision via Label Grouping and Co-training

Islam Nassar^{1*}, Samitha Herath¹, Ehsan Abbasnejad², Wray Buntine¹, Gholamreza Haffari¹

¹ Dept of Data Science and AI, Faculty of IT, Monash University, Australia

² Australian Institute for Machine Learning, University of Adelaide, Australia

Abstract

Pseudo-labeling is a key component in semi-supervised learning (SSL). It relies on iteratively using the model to generate artificial labels for the unlabeled data to train against. A common property among its various methods is that they only rely on the model's prediction to make labeling decisions without considering any prior knowledge about the visual similarity among the classes. In this paper, we demonstrate that this degrades the quality of pseudo-labeling as it poorly represents visually similar classes in the pool of pseudo-labeled data. We propose SemCo, a method which leverages label semantics and co-training to address this problem. We train two classifiers with two different views of the class labels: one classifier uses the one-hot view of the labels and disregards any potential similarity among the classes, while the other uses a distributed view of the labels and groups potentially similar classes together. We then co-train the two classifiers to learn based on their disagreements. We show that our method achieves state-of-the-art performance across various SSL tasks including 5.6% accuracy improvement on Mini-ImageNet dataset with 1000 labeled examples. We also show that our method requires smaller batch size and fewer training iterations to reach its best performance. We make our code available at <https://github.com/islam-nassar/semco>.

1. Introduction

Deep neural models require large amounts of labeled data to achieve their high performance. This quickly becomes prohibitive and non-scalable especially when labeling data is expensive and/or non practical. Semi-supervised learning (SSL) [5, 33] has hence emerged to explore a diverse set of methods which aim to leverage unlabeled data to enable learning from a smaller set of labeled data.

In the context of image classification, recent methods use unlabeled data to guide learning in different ways.

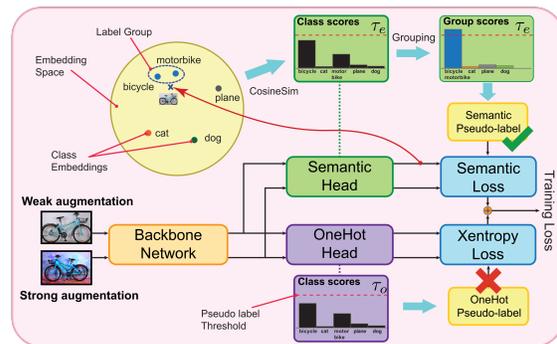


Figure 1: A conceptual diagram of our co-training solution

Some methods primarily focus on consistency regularization [29, 17], where the model is enforced to produce consistent predictions for different perturbed versions of the same unlabeled input image. While others focus on pseudo-labeling [1, 18, 13], where the model is used to produce artificial labels for the unlabeled data that are then used to further train the model. Evidently, combining the two approaches has shown the state-of-the-art results on various image classification tasks [30].

When it comes to pseudo-labeling, a common problem which hinders the SSL performance is the so-called confirmation bias [32]. This takes place when the model reassures its wrong predictions by retraining on them, leading to an accumulation of the error from which the model can not recover. To mitigate this behaviour, some methods use a warm-up phase until the model becomes more reliable [13, 32], or limit the number of pseudo-labeled samples in each mini-batch [1]. Other strategies include using a confidence threshold whereby a sample is only considered for pseudo-labeling if the model is highly confident about its prediction [30, 18]. One property shared by all these methods, however, is that they only rely on the model's output and disregard any prior knowledge about potential similarities among the classes. As we show in Section 2, visually similar classes are expected to confuse the model and therefore get poorly represented in the pseudo-labeled data pool. This fact is even more exacerbated in confidence-

*corresponding author: islam.nassar@monash.edu

based methods [18, 30] as it leads to discarding most of the visually similar samples simply because the model is rarely confident about their predictions. We show that this leads to a class imbalance in the pseudo-labeled pool, and thereby, misguides the training.

In this paper, we demonstrate that by exploiting class labels semantics, we can account for such similarity among the classes. We draw inspiration from few-shot learning methods [9, 41] where we use distributed embeddings to represent class labels. We present two methods to generate label embeddings in a way which encodes a weak prior on the visual similarity among the classes. One such method is based on knowledge graph embeddings [31], while another is based on visual attributes annotation [35]. Having such embeddings provides basis to group the class labels into visually similar concepts and allow considering such grouping while making pseudo-labeling decisions.

The benefit of using label embeddings goes beyond label grouping. Earlier work [9] has shown that using embeddings as training targets (as opposed to one-hot labels) allows the model to map the image features to a more meaningful semantic space, and thereby, enables few shot transfer. In our work, we leverage this idea in a co-training [4] style approach to improve SSL performance. We propose to train two classifiers with the two different views of the class labels, i.e. one-hot and distributed. One of the classifiers makes use of the label grouping during pseudo-labelling, while the other does not. We then allow the two classifiers to learn from their disagreements via a shared consistency regularization loss on the unlabeled data.

We show that our method achieves new state-of-the-art results across five different datasets, while using smaller batch size with fewer training iterations. To summarize, our contributions are:

1. We propose an approach which leverages the semantic similarity among the classes to improve pseudo-labeling quality by addressing the confusion events.
2. We present a co-training-based SSL method which involves two classifiers co-operating via pseudo-labels obtained using their different views of the class label.
3. We show our approach outperforms the state-of-the-art in SSL by a large margin on 5 different datasets including 5.6% on Mini-Imagenet with 1000 labeled point, i.e. 10 labels per class.

2. Background

We are interested in a K -way semi-supervised image classification problem, where we train a model using batches of both labelled and unlabelled examples. Specifically, each batch comprises labeled examples, $\mathcal{X} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ and unlabeled examples, $\mathcal{U} = \{\mathbf{u}_j\}_{j=1}^{\mu \cdot n}$, where the scalar μ denotes the ratio between the number

of unlabeled and labeled examples in a given batch, and \mathbf{y}_i denotes the one-hot representation of the label. But before we introduce our approach, we begin by reviewing three key concepts underpinning our method.

Consistency Regularization These methods exploit the assumption that predictions for different perturbed versions of a sample should be consistent [19, 29]. One way to operationalise this idea is to produce several augmented versions of a given unlabeled image, then apply a loss to ensure that the predictions for all such versions are consistent. Inspired by recent methods [30, 3], we make use of two types of augmentations, namely, weak augmentations $\mathcal{A}_w(\cdot)$ and strong augmentations $\mathcal{A}_s(\cdot)$, where the notion of intensity relates to how perturbing an augmentation is to an image.

Pseudo-labeling These methods rely on producing synthetic labels for unlabeled data which are then used to re-train the model. Recent alternative variations of pseudo-labeling [18, 3, 30] can broadly be formalized as methods trying to account for unlabeled data by minimizing the following objective,

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{-1}{\mu \cdot n} \sum_{j=1}^{\mu n} \eta_j \log p(y = \hat{y}_j | \mathbf{u}_j, \boldsymbol{\theta}), \quad (1)$$

where $\boldsymbol{\theta}$ represents learnable model parameters, \hat{y} denotes the pseudo-label, and η is an arbitrary function. The choice of \hat{y} and η gives rise to different variations of pseudo-labeling methods [18, 3, 30, 18]. We are particularly interested in confidence-based methods where

$$\hat{y}_j = \arg \max_{y'} p(y = y' | \mathbf{u}_j, \boldsymbol{\theta}) \quad (2)$$

and $\eta_j = \mathbb{1}(p(y = \hat{y}_j | \mathbf{u}_j, \boldsymbol{\theta}) \geq \tau)$,

with $\mathbb{1}$ denoting the Indicator function. In such methods, the unlabeled sample is only retained for pseudo-labeling if the model’s maximum confidence score exceeds a predefined threshold, τ , and the pseudo-label is then selected to be the class with the maximum score¹. This approach mitigates confirmation bias (see Sec. 1) by only retaining high confidence samples. Simultaneously, it encourages entropy minimization [10] whereby the model is encouraged to produce high confidence predictions on the unlabeled data. Recently, Sohn et al. [30] combined such approach with consistency regularization to propose FixMatch, a method which achieves state-of-the-art results on several SSL image classification benchmarks.

Co-training The idea of co-training [4] is to train two models with different views of the data, where each model is trained on the other’s most confident predictions. Given sufficiently diverse views of the data, this approach was shown to improve learning, as it allows the two models to

¹For notation simplicity, we assume here that the $\arg \max$ in Eqn. 2 produces a one-hot probability distribution

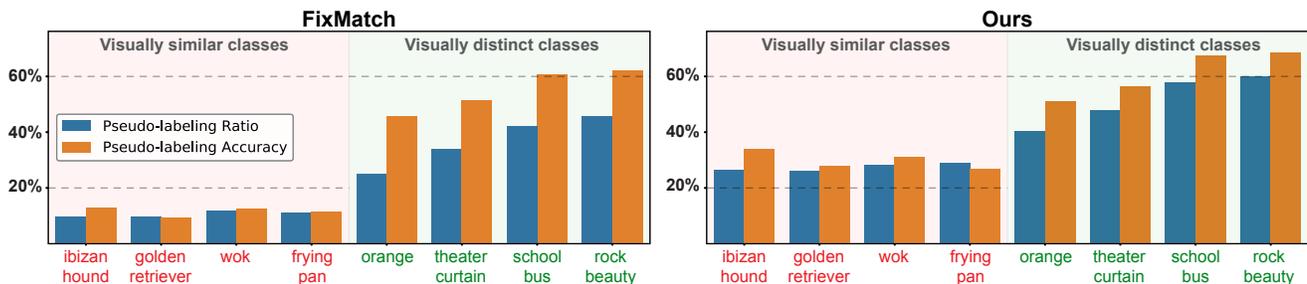


Figure 2: Confidence-based pseudo-labeling comparison between the baseline (left) and our method (right). *Accuracy* values show how much, on average, pseudo-labels for a given class match the true label, while *Ratio* values show the percentage of samples of a given class which are retained for pseudo-labeling (i.e. with confidence score above the threshold). The two metrics are calculated for the 4 most (red) and least (green) visually similar classes over the first 10 epochs of training.

learn based on their disagreements [33]. We adopt a similar strategy, albeit, we use two different views of the label rather than the data. We use the regular one-hot view as well as a distributed view (i.e. label embedding). As we introduced in Sec. 1, using a distributed view of the label grants the ability to map from the image feature space to another meaningful semantic space. This is under the assumption that the label embeddings are learnt in a way that captures semantic similarities among the labels. In [9], authors show how semantic information gleaned from text, in form of word embeddings [21], can be exploited to enable prediction of labels never observed during training. In this work, we combine the above ideas to propose our method.

3. Problem Statement and Motivation

While achieving great results, approaches that rely on pseudo-labeling share a limitation. As Eqn. (1) suggests, they solely rely on the model’s prediction to decide about pseudo-labeling, while disregarding any prior information about possible similarities among the classes. We find in our work that visually similar classes often produce low-confidence predictions, hence are either discarded (for methods which use confidence thresholds such as FixMatch) or confused with others. This leads to class imbalance among the pseudo-labeled instances which potentially misguides SSL training. In Fig. 2 - left, we demonstrate such behaviour by examining the pseudo-labeling statistics of FixMatch method. We use the true labels of the unlabeled data² to calculate the true accuracy of pseudo-labeling for each class. Further, we calculate the ratio of samples retained for pseudo-labeling (i.e. where the classifier confidence exceeds the threshold). We plot these two metrics for the 4 most and least visually similar classes³. We observe that visually similar concepts are chosen less frequently (i.e.

less ratio) for pseudo-labeling and are often mislabeled (i.e. less accuracy) as opposed to visually distinct concepts. Motivated by this observation, we consider the label similarities as a particularly essential prior that is easy to obtain. In the subsequent sections, we will discuss how to obtain and incorporate such a prior for an improved pseudo-labeling.

4. Our Method (SemCo)

We aim to address the issues demonstrated for visually similar classes. We build on top of recent approaches, but we additionally propose to condition pseudo-labeling on our prior knowledge of class similarities. Effectively, we enhance the model by incorporating knowledge about potential confusions based on semantic and visual similarities⁴. To that end, we encode the notion of similarity among the classes using a label embeddings matrix $M \in \mathbb{R}^{K \times d}$ where each row represents a d -dimensional label embedding of class $k \in \{1, \dots, K\}$ ⁵. We further group the labels using a density-based clustering approach such as [6] using a hyperparameter ϵ so that the number of groups are not pre-defined. Subsequently, we obtain Q class groups.

Thereafter, we train two classifiers sharing the same backbone network (see Fig. 1). The *Semantic Classifier* $f_{sc} : \mathbb{R}^{h \times w} \rightarrow \mathbb{R}^d$ maps an input image closer to its corresponding label in the embedding space spanned by the rows of M ; and the *One-Hot Classifier* $f_{oh} : \mathbb{R}^{h \times w} \rightarrow \mathbb{R}^K$ maps input images to a one-hot view of the label. Note that, for brevity, we define the classifiers f_{sc} and f_{oh} to implicitly include the shared backbone network and its parameters.

For each of the two classifiers we minimize a supervised loss on the labeled data and an unsupervised consistency loss on the unlabeled data. Additionally, we add a loss term for co-training to allow the two classifiers to co-operate on pseudo-labeling.

²Note that we have access to the true labels but they are discarded during training to emulate an SSL setting

³We elaborate on how we identify similarity in Sec. 4

⁴We present a probabilistic interpretation of our method in the supplementary material

⁵We defer the discussion on how to obtain M to Sec. 4.5.

4.1. Semantic Classifier

For the supervised loss, we minimize the *cosine loss* between the true label embedding and the predicted label embedding,

$$\mathcal{L}_s^{sc} = \frac{1}{n} \sum_{i=1}^n C(M^T \mathbf{y}_i, f_{sc}(\mathbf{x}_i)), \quad (3)$$

with $C(\mathbf{z}, \mathbf{z}') = 1 - \text{CosineSim}(\mathbf{z}, \mathbf{z}')$.

For the unsupervised loss (see Fig. 3), we draw inspiration from FixMatch [30], where we use a weakly augmented version of the image to obtain a pseudo-label and enforce that against the model’s prediction for a strongly augmented version of the same image. Specifically, for an unlabeled image \mathbf{u}_j , we obtain the predicted embedding for a weakly augmented version of the image: $\mathbf{q}_j = f_{sc}(\mathcal{A}_w(\mathbf{u}_j))$. Then we calculate class scores, $\mathbf{p}_j = p(\mathbf{y}_j|\mathbf{u}_j)$ by normalizing the vector-wise cosine similarity between \mathbf{q}_j and M .

Unlike FixMatch, we consider an unlabeled sample for pseudo labeling if the prediction score for one of the class groups (as opposed to the individual classes) exceeds a pre-defined threshold (τ_e). To elaborate, referring to Fig. 3, due to the visual similarity between “*bicycle*” and “*motorbike*”, the class scores for each of them, individually, is falling below the threshold. However, since they are both identified as “visually similar” based on clustering their embeddings, their scores are added first before applying the threshold. The combined score exceeds the threshold so the sample is retained for pseudo-labeling where the pseudo-label is calculated as the average of the “*bicycle*” and “*motorbike*” embeddings weighted by their normalized class prediction scores.

To put it formally, to obtain the score for a given class label group, we sum the normalized class scores of all its members (where the membership is defined based on clustering M). This gives rise to our group scores \mathbf{g}_j ⁶.

Thereafter, we apply our mask to select samples for pseudo-labeling as per,

$$\eta_j^{sc} = \mathbb{1}(\max(\mathbf{g}_j) \geq \tau_e). \quad (4)$$

If a sample is selected for pseudo-labeling, we obtain a pseudo-label embedding ($\hat{\mathbf{y}}_j$)⁷ for such sample as a weighted average of the group members embedding, where we weigh the average based on the original class scores \mathbf{p}_j . Consequently, we apply the loss against the embedding prediction of a strongly augmented version of \mathbf{u}_j as per,

$$\mathcal{L}_u^{sc} = \frac{1}{\mu \cdot n} \sum_{j=1}^{\mu \cdot n} C(\hat{\mathbf{y}}_j, f_{sc}(\mathcal{A}_s(\mathbf{u}_j))) \cdot \eta_j^{sc}. \quad (5)$$

⁶ \mathbf{g}_j is calculated as the inner product of \mathbf{p}_j with the cluster assignment matrix

⁷ $\hat{\mathbf{y}}_j$ is calculated as the inner product of M with the normalised class scores.

4.2. One-Hot Classifier

For the *One-Hot Classifier*, we follow the same procedure as the *Semantic Classifier* with two crucial differences: 1) we use *cross-entropy loss* instead of *cosine loss*, and 2) we don’t apply label grouping before comparing with the confidence threshold. We note here that this classifier operates in a similar way to FixMatch, yet we include the loss equations for completeness. By analogy, the supervised loss is calculated as,

$$\mathcal{L}_s^{oh} = \frac{1}{n} \sum_{i=1}^n H(\mathbf{y}_i, f_{oh}(\mathbf{x}_i)). \quad (6)$$

Here, we use H to represent the cross-entropy loss function. To this end, the unsupervised loss is formulated as,

$$\mathcal{L}_u^{oh} = \frac{1}{\mu \cdot n} \sum_{j=1}^{\mu \cdot n} H(\hat{\mathbf{y}}_j, f_{oh}(\mathcal{A}_s(\mathbf{u}_j))) \cdot \eta_j^{oh}, \quad (7)$$

where $\hat{\mathbf{y}}_j = \arg \max \mathcal{A}_w(\mathbf{u}_j)$; $\eta_j^{oh} = \mathbb{1}(\max(\hat{\mathbf{y}}_j) \geq \tau_o)$. We note that our motivation behind using two different types of loss functions is related to the concept of co-training. An assumption underlying the success of co-training is to ensure that the learners are sufficiently diverse so that they learn better based on their different views [33]. We validate such choice in our ablations (Sec. 5.4)

4.3. Co-training Loss

This loss is meant to enable both classifiers to learn from each other. The intuition is that due to each classifier’s different view of the labels, they will each be confident about different samples of the unlabeled data. We exploit that by retaining a sample for pseudo-labeling if either of the classifiers is confident about its prediction. In case the two classifiers disagree about a sample (i.e. they are both confident about two different labels), the sample is included twice in the loss, once with each of the two pseudo-labels. We experimented with another approach, where in such event, the sample gets discarded but it degraded the performance. We conjecture that it’s because the former approach encourages the two classifiers to be consistent while the latter completely ignores the confusion event. Formally, we define the co-training loss as,

$$\mathcal{L}_{co} = \frac{1}{\mu \cdot n} \sum_{j=1}^{\mu \cdot n} C(M^T \hat{\mathbf{y}}_j, f_{sc}(\mathcal{A}_s(\mathbf{u}_j))) \cdot \eta_j^{oh} + H(\arg \max(\mathbf{p}_j), f_{oh}(\mathcal{A}_s(\mathbf{u}_j))) \cdot \eta_j^{sc} \quad (8)$$

4.4. Total Loss

We now define our final training loss function by combining all five losses (Eqns. 3, and 5 to 8) as per,

$$\mathcal{L}_{total} = \mathcal{L}_s^{sc} + \mathcal{L}_s^{oh} + \lambda_u(\mathcal{L}_u^{sc} + \mathcal{L}_u^{oh}) + \lambda_{co}\mathcal{L}_{co}. \quad (9)$$

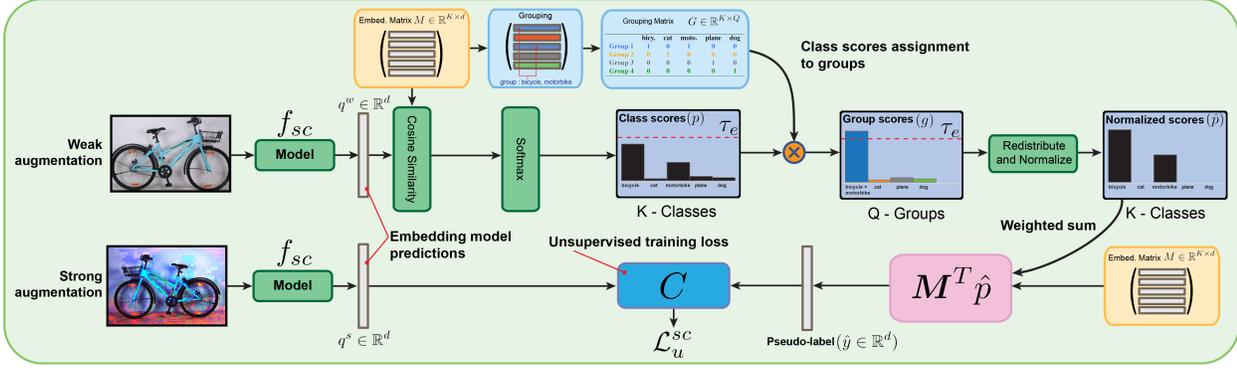


Figure 3: Unsupervised loss for the *Semantic Classifier* - A weakly augmented image is used (upper path) to obtain a predicted embedding, which is then used to obtain class scores. The class scores are summed for each label group (as identified by our grouping method) to obtain group scores. If one of the group scores exceeds the threshold, it is retained for pseudo-labeling. The pseudo-label is then calculated as an average of the group members embeddings weighted by their class scores. The loss is then enforced against the predicted embedding for a strongly augmented image (lower path).

Here, λ_u and λ_{co} are fixed scalar weights to modulate the contribution of the unsupervised loss and co-training loss, respectively.

4.5. Extracting Label Semantics

In this section, we propose two alternatives to obtain the label embedding matrix M which establishes our prior on the visual similarity among the classes.

Using Knowledge Graphs In cases where the class labels are semantically meaningful, we make use of the ConceptNet knowledge graph [31] together with GloVe [25] and word2vec [21] distributional embeddings as the basis for obtaining the distributed label embeddings. ConceptNet is a multilingual knowledge graph that connects words of natural language with labeled, weighted relations. Since our main goal is to obtain label embeddings which capture visual similarity, we filter the graph to only retain relations which imply such similarity. Specifically, we retain any nodes which share the following relations: *SimilarTo*, *InstanceOf*, *IsA*, *FormOf*, *Synonym*, *EtymologicallyRelatedTo*, *DefinedAs*. A detailed description of such relations and examples thereof can be found in the ConceptNet documentation⁸. On the other hand, GloVe and word2vec are two prominent sets of word embeddings, the former is trained on 840 billion words of the Common Crawl [25], while the latter is trained on 100 billion words of Google News [21]. The two sets capture the distributional similarity among the different words but don't necessarily capture visual similarity. For example, “*cat*” and “*dog*” usually appear in similar contexts (being both animal pets) so they would have a relatively similar GloVe (or word2vec) word embedding even though they are not visually similar. Combining the distributional embeddings with the ConceptNet

⁸<https://github.com/commonsense/conceptnet5/wiki/Relations>

filtered graph allows us to address this problem: we follow a procedure similar to the authors in [31] to retrofit the distributional embeddings with the filtered knowledge graph. Retrofitting [7] is a process which adjusts a word embedding matrix based on a knowledge graph by optimizing an objective function which tries to find for each term a new vector close to the vector's original value but also close to the term neighbours in the graph. Since the retained relations in the graph are those which implies visual similarity, this retrofitting results in a new hybrid set of embeddings which captures distributional similarity but also correlates well with visual similarity. Finally, to handle class labels which are not present in the embeddings vocabulary, we implement a fall-out strategy to find the most reasonable alternative. We provide further description of the retrofitting process and we show a qualitative comparison to demonstrate its effectiveness in the supplementary material. We also provide further details of the fall-out strategy.

Using Class Attributes Annotations In cases where the class labels are not semantically meaningful, a viable alternative is to use manually annotated class attributes. We demonstrate (Sec. 5) that by using attributes annotation of CUB-200 [35] fine-grained dataset as our M matrix, we achieve significant gains against the baseline. Considering that the cost of annotating attributes is expected to be cheaper than annotating data instances, we propose class attributes annotation as a possible alternative.

5. Experiments

To evaluate SemCo, we compare it to various recent SSL baselines on 3 standard benchmarks (CIFAR-10 [15], CIFAR-100 [15], Mini-ImageNet [27]). Further, we experiment on 2 other datasets: CUB-200 [35], to test SemCo on fine-grained tasks; and DomainNet [24], to verify its performance on larger more complex images.

Table 1: Error rates for CIFAR-10, CIFAR-100 and Mini-ImageNet. We report results for two different values of μ - i.e. ratio between unlabeled and labeled data in a mini-batch, for our method and FixMatch. † denotes that the results reported are using the same codebase. * denotes that the result is based on using CNN-13 model. We report the mean and standard deviation across 3 different splits of labeled data for each experiment.

Total Labelled Samples	CIFAR-10		CIFAR-100			Mini-ImageNet		
	250	4000	2500	4000	10000	1000	4000	10000
Pseudo-labeling [18]	49.78±0.43	16.09±0.28	-	-	-	-	-	-
Mean teacher [32]	32.32±2.30	9.19±0.19	-	-	-	-	72.51±0.22	57.55±1.11
UDA [40]	8.82±1.08	4.88±0.18	33.13±0.22	-	24.50±0.25	-	-	-
Label Propagation [13]	-	12.69±0.29*	-	-	-	-	70.29±0.81	57.58±1.47
PLCB [1]	24.81±5.35	6.28±0.30	-	37.55±1.09*	32.15±0.50*	-	56.49±0.51	46.08±0.11
MixMatch† [3]	11.29±0.75	6.24±0.07	39.70±0.27	-	28.59±0.31	60.97±0.31	49.79±0.11	44.27±0.23
FixMatch†($\mu = 3$) [30]	5.78±0.23	4.52±0.01	38.45±0.51	32.22±0.21	28.42±0.09	66.23±1.13	59.73±5.45	44.66±0.12
FixMatch†($\mu = 7$)	4.55±0.12	4.49±0.05	33.64±0.07	31.27±1.30	26.13±0.18	60.97±0.31	49.79±0.11	44.27±0.23
Ours (SemCo)†($\mu = 3$)	5.87±0.31	4.43±0.01	33.80±0.57	29.40±0.18	25.07±0.04	55.35±0.71	46.01±0.93	41.25±0.76
Ours (SemCo)†($\mu = 7$)	5.12±0.27	3.80±0.08	31.93±0.01	28.61±0.23	24.45±0.12	59.35±0.23	49.46±2.20	42.78±0.35

5.1. Datasets

CIFAR-10/100 Both datasets comprise natural images of 10, and 100 classes respectively. Their training set consists of 50k images while the test set consists of 10k images. All the images have a fixed resolution of 32x32. We conduct three different experiments on each of them with varying amounts of labeled data as shown in Table 1.

Mini-ImageNet A subset of the well-known ImageNet [28]. It consists of 100 classes with 600 images per class (84x84 each). We use the same train/test split used by [13] and we create splits for 40 and 100 labeled images per class to enable comparing with the baseline systems. However, we also experiment with 10 images per class to test SemCo in the low data regime.

CUB-200 A fine-grained image classification dataset comprising 11k images from 200 different types of birds annotated with 312 attributes per class. We experiment with 5 and 10 images per class corresponding to 1000 and 2000 total labeled data.

DomainNet The dataset contains 345 classes of images coming from six domains: *Clipart*, *Infograph*, *Painting*, *Quickdraw*, *Real*, and *Sketch*. We report results only on the *Real* domain to evaluate how our method works on larger more complex datasets.

5.2. Experimental Setup

Across all experiments, we follow the standard approach where we randomly select a certain number of samples to represent our labeled set and ignore the labels of the remaining samples and use them to form our unlabeled set. For the standard benchmarks, we compare our results to various existing baselines [18, 32, 3, 40, 13, 1, 30], which employ consistency regularization and/or pseudo-labeling (see Sec. 6). For the two other datasets, we only compare with FixMatch,

being the most similar to our solution and the closest in performance.

Since SemCo bears the most resemblance with [30], [3], and [2], we follow the recommendation of Oliver et al. [23] for a realistic comparison: we integrate the implementation of their methods⁹ into our codebase and use the unified codebase to conduct all the experiments. As for the other baselines, we report the results as mentioned in the original papers, provided that the result is based on the same model architecture we use. We use WideResnet-28-2 [42] for CIFAR-10/100, Resnet-18 [11] for Mini-ImageNet, and Resnet50 for CUB-200 and DomainNet. Additionally, we attach a fully connected layer to the encoder output to act as our *Semantic Classifier* (see Fig. 1). We train our model end-to-end along with the backbone network.

Unless otherwise specified, we use the same hyperparameters for all our experiments. These were tuned on a validation set for a single experiment (CIFAR100 - 2500 labels) and then fixed across all other experiments. In general, we found that our model is not sensitive to the values of λ_u and λ_{co} . Values between 0.5 - 1 all yielded similar performance, albeit, smaller values of λ_u slightly slowed convergence. Further, we found that our model is mostly sensitive to ϵ - the label grouping parameter and hence it was the only parameter we tuned separately for each dataset (see supplements for the full list of hyperparameters).

Since the labels of all datasets are semantically meaningful, we use their retrofitted embeddings (see Sec. 4.5) as targets for our *Semantic Classifier*. The only exception is for CUB-200 where we use human annotated attributes as targets to test our alternative proposal in Sec. 4.5. We start from the 312 dimensional class attributes given in [35] and reduce their dimensionality to 128 using PCA [36]. We, then, use the obtained class attributes matrix as targets for our *Semantic Classifier*.

⁹we don't report results for [2] due to adaptation difficulties

5.3. Main Results

We report standard benchmarks results in Table 1 and CUB-200 and DomainNet results in Table 2. We observe that SemCo outperforms all the baselines with a large margin across the different datasets and amounts of labeled data (except for one case). Notably, SemCo achieves an average error rate of 55.35% on Mini-ImageNet with 1000 labels (i.e. 10 samples per class). This is almost 5.6% improvement compared to the closest baseline. We note here that Mini-ImageNet classes include 13 different species of dogs which share many visual similarities. SemCo grouped 7 of these classes into a single group based on clustering their label embeddings. To understand why the performance degrades on CIFAR10 (250 labeled samples), we looked into the clustering results for CIFAR-10 class embeddings. We observed that the 10 classes were deemed visually distinct by our clustering component, leading to one-member clusters for all 10 classes. The two above results align with our original hypothesis that SemCo is particularly useful when there are visually similar concepts among the classes. Evidently, in such cases, using label grouping in conjunction with our co-training routine helps improving the pseudo-labeling quality. This is also consistent with the pseudo-labeling statistics shown in Fig. 2 - right, where we can see that SemCo significantly improves both the quality and the quantity of pseudo-labeled data.

Ratio of unlabeled data We observe that SemCo achieves better results with less batch size as opposed to the baseline. As shown in Table 1, we experiment with different values of μ , which defines the ratio between unlabeled and labeled data in each training batch. We find that our method consistently achieves better results even when using less unlabeled data. For example, for CIFAR-100 (4000 labels), we achieve less average error rate (29.4%) with $\mu = 3$ than FixMatch does with $\mu = 7$ (31.27%). More notably, we achieve 13% improvement on Mini-ImageNet (4000 labels) when fixing μ to 3.

Co-training Analysis Further, we investigate the effectiveness of our co-training routine. We use the same experimental setup of capturing pseudo-label metrics (see Fig. 2), but this time, we monitor the rate of disagreement on pseudo-labels among the two classifiers f_{sc} and f_{oh} (i.e. percentage of time the two classifiers are confident about different pseudo-labels for the same unlabeled sample). In Fig. 4 c, we report disagreement curves for the same 8 classes shown in Fig. 2. As the training progresses, we track the pseudo-labeling accuracy (Fig. 4 d) for each of: 1) our classifier ensemble, 2) our one-hot classifier f_{oh} , and 3) the baseline (FixMatch). Note that f_{oh} is using the same method for pseudo-labeling (i.e. confidence threshold on non-grouped labels) as FixMatch. We find that at the beginning of the training, both classifiers highly disagree about

pseudo-labeling, especially for visually similar classes. As the training progresses, we witness a sharp reduction in disagreements coupled with an increase in accuracy for both classifiers. Interestingly, we find that the accuracy of our f_{oh} is considerably higher than FixMatch although both are using the same basis for pseudo-labeling. This demonstrates the success of co-training in making pseudo-labeling consistent and accurate by leveraging the co-operation between the two classifiers.

Convergence Speed In Fig. 4 a and b, we study the convergence plots on Mini-ImageNet and CIFAR-100 (1000 labels). We observe that SemCo achieves the best performance of the baseline with significantly less training iterations. This can be explained through Fig. 4 d: the higher accuracy of pseudo-labeling in the early phase of the training helps better guide the learning and thereby translates to faster convergence.

Table 2: Error rates on CUB-200 dataset and DomainNet Real. Errors are reported based on 1 split for each of the amounts of labeled data. Poor baseline results are omitted.

CUB-200 Method	Total Labeled Samples	
	1000	2000
Supervised baseline	-	70.11
FixMatch	84.35	72.15
Ours (SemCo)	79.44	66.76

DomainNet Real Method	Total Labeled Samples	
	6900	10350
Supervised baseline	47.9	45.2
FixMatch	41.34	39.04
Ours (SemCo)	35.32	32.89

5.4. Ablation

We are interested in isolating the contribution of each of the three key components of SemCo towards the witnessed performance gain.

Label Grouping & Co-training In Table 3, we investigate the effect of label grouping (by controlling our clustering hyperparameter ϵ), and co-training (by toggling λ_{co}). We observe that both components are almost equally important towards the witnessed performance gain. However, co-training seems to provide a slight advantage over label grouping in both experiments.

Label Embeddings as Training Targets We experiment on CIFAR-100 and Mini-ImageNet in another setting where we use the one-hot target for both our classifiers. In such case, the only difference between the two classifiers is that the *Semantic Classifier* implements label grouping while the *One-Hot Classifier* does not. In Table 4, we observe

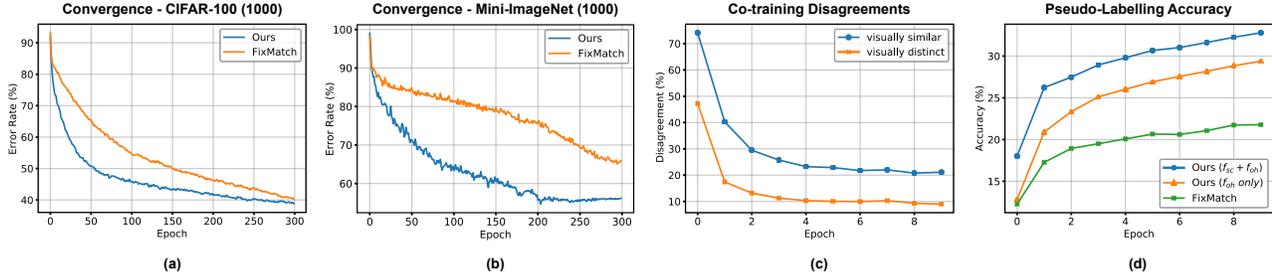


Figure 4: Experimental analysis plots showing: (a,b): Convergence trends of our method and the baseline for CIFAR-100 (a) and Mini-ImageNet (b) with 1000 labeled examples. (c,d): Co-training analysis plots showing the disagreements between our two classifiers for visually similar and distinct classes (c) and the associated pseudo-labeling accuracies (d). The co-training plots are spanning only the first 10 epochs of training.

Table 3: Error Rates for different settings of Co-training and Label Grouping

Label Grouping	Co-training	Mini-ImageNet	CIFAR-100
		1000	2500
✓	✓	55.35	31.93
-	✓	59.60	33.09
✓	-	60.39	33.19
-	-	62.16	34.25

a significant decrease in performance when using the one-hot view for both the classifiers. This strongly supports our hypothesis that co-training with different views of the label does indeed help the learning.

Table 4: Error Rates when using Embedding Targets versus One-Hot Targets for our *Semantic Classifier*, reported on CIFAR-100 and Mini-ImageNet

	Embeddings Target	One-Hot Target
CIFAR-100 (2500)	31.93	33.33
Mini-ImageNet (1000)	55.35	60.33

6. Related Work

Since the seminal “II-model” [26], consistency regularization and pseudo-labelling SSL solutions have seen improvements in the consistency propagation [32, 17], and approaches for generating diverse views [8, 16, 22]. For instance, the Mean-Teacher [32] proposes a teacher model, where its parameters are updated according to an exponential moving average (EMA) rule. Temporal Ensembling [17], maintains an EMA over the predictions for the consistency loss computation. French et al. [8] explore a masking based approach for generating diverse views. The interpolation training given in [34] computes the consistency between interpolated views of unlabelled instances. Miyato et al. [22] explores using adversarial methods for perturbation to create diverse views.

Our solution bears a lot of similarity to FixMatch and

ReMixMatch [30, 2] where the main idea is to use a weakly augmented image to obtain a pseudo-label then enforce it against the model’s prediction for a strongly augmented one. ReMixMatch uses a soft pseudo-label via sharpening, while FixMatch uses a hard label based on confidence. Our method compliments theirs by also conditioning on label semantics while pseudo-labeling.

Using label semantics to benefit learning is not a new idea, prior knowledge from language models [25], graph embeddings [31], and attribute vector representations [35] has helped pushing the performance of computer vision models. In the pioneering work, DeVISE [9] showed distributed label representations derived from unannotated text are helpful for image classification. They also extend their solution to Zero-Shot Learning (ZSL) [39]. Ye et al. [41] proposes distributed labels for Few-shot learning. Such label representations are informative to even generate descriptive representations for classification [38] and has become the backbone representation for ZSL [14, 20]. To this end, literature provides explorations on learning visual-semantic embedding spaces with better discriminative properties [37, 12]. However, to our best knowledge the capacity of such label representations has not been explored for SSL.

7. Conclusion

In this paper, we have introduced a novel semi-supervised learning approach leveraging class label semantics and co-training for more effective and efficient learning. We operationalize this approach for image classification, and demonstrate that it leads to significant gains. We believe the key ingredients of our approach are general and can be extended to supervised and unsupervised learning settings, which we will explore in the future work.

Acknowledgement

This work was partly supported by DARPA’s Learning with Less Labeling (LwLL) program under agreement FA8750-19-2-0501 and by the Australian Government Research Training Program (RTP) Scholarship.

References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020.
- [2] 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. *arXiv preprint arXiv:1911.09785*, 2019.
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 5049–5059, 2019.
- [4] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100, 1998.
- [5] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009.
- [6] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise.
- [7] Manaal Faruqi, Jesse Dodge, Sujay K Jauhar, Chris Dyer, Eduard Hovy, and Noah A Smith. Retrofitting word vectors to semantic lexicons. *arXiv preprint arXiv:1411.4166*, 2014.
- [8] Geoff French, Avital Oliver, and Tim Salimans. Milking cowmask for semi-supervised image classification. *arXiv preprint arXiv:2003.12022*, 2020.
- [9] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, pages 2121–2129, 2013.
- [10] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pages 529–536, 2005.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [12] Yao-Hung Hubert Tsai, Liang-Kang Huang, and Ruslan Salakhutdinov. Learning robust visual-semantic embeddings. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3571–3580, 2017.
- [13] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5070–5079, 2019.
- [14] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3174–3183, 2017.
- [15] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
- [16] Chia-Wen Kuo, Chih-Yao Ma, Jia-Bin Huang, and Zsolt Kira. Featmatch: Feature-based augmentation for semi-supervised learning. *arXiv preprint arXiv:2007.08505*, 2020.
- [17] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.
- [18] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013.
- [19] Yiting Li, Lu Liu, and Robby T Tan. Decoupled certainty-driven consistency loss for semi-supervised learning. *arXiv preprint arXiv:1901.05657*, 2019.
- [20] Shaoteng Liu, Jingjing Chen, Liangming Pan, Chong-Wah Ngo, Tat-Seng Chua, and Yu-Gang Jiang. Hyperbolic visual embedding learning for zero-shot recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [21] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [22] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- [23] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in neural information processing systems*, pages 3235–3246, 2018.
- [24] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019.
- [25] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [26] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Advances in neural information processing systems*, pages 3546–3554, 2015.
- [27] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016.
- [28] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [29] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in neural information processing systems*, pages 1163–1171, 2016.

- [30] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020.
- [31] Robyn Speer, Joshua Chin, and Catherine Havasi. Conceptnet 5.5: An open multilingual graph of general knowledge. *arXiv preprint arXiv:1612.03975*, 2016.
- [32] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017.
- [33] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440, 2020.
- [34] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 3635–3641, 2019.
- [35] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010.
- [36] Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.
- [37] Hao Wu, Jiayuan Mao, Yufeng Zhang, Yuning Jiang, Lei Li, Weiwei Sun, and Wei-Ying Ma. Unified visual-semantic embeddings: Bridging vision and language with structured meaning representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6609–6618, 2019.
- [38] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5542–5551, 2018.
- [39] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning-the good, the bad and the ugly. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4582–4591, 2017.
- [40] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848*, 2019.
- [41] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Learning embedding adaptation for few-shot learning. *arXiv preprint arXiv:1812.03664*, 2018.
- [42] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.