

Better Self-training for Image Classification through Self-supervision

Attaullah Sahito, Eibe Frank, and Bernhard Pfahringer

Department of Computer Science, University of Waikato, Hamilton, New Zealand
a19@students.waikato.ac.nz, {eibe,bernhard}@waikato.ac.nz

Abstract. Self-training is a simple semi-supervised learning approach: Unlabelled examples that attract high-confidence predictions are labelled with their predictions and added to the training set, with this process being repeated multiple times. Recently, self-supervision—learning without manual supervision by solving an automatically-generated pretext task—has gained prominence in deep learning. This paper investigates three different ways of incorporating self-supervision into self-training to improve accuracy in image classification: self-supervision as pretraining only, self-supervision performed exclusively in the first iteration of self-training, and self-supervision added to every iteration of self-training. Empirical results on the SVHN, CIFAR-10, and PlantVillage datasets, using both training from scratch, and Imagenet-pretrained weights, show that applying self-supervision only in the first iteration of self-training can greatly improve accuracy, for a modest increase in computation time.

Keywords: Self-supervised learning · Self-training · Rotational loss.

1 Introduction

Contemporary machine learning based on deep neural networks achieves state-of-the-art results on many tasks, including visual understanding, language modelling, and speech recognition, but often requires a large amount of human-labelled data for such performance. Collecting large amounts of labelled data is expensive and time-consuming. Hence, there is significant interest in algorithms that can harness the benefits of unlabelled data—which is usually plentiful—in addition to the scarce labelled data. Semi-supervised learning (SSL) algorithms are designed for this scenario. Their input comprises labelled examples $L = \{(x_1, y_1), (x_2, y_2), \dots, (x_{|L|}, y_{|L|})\}$ and unlabelled examples $U = \{x'_1, x'_2, \dots, x'_{|U|}\}$, where $x_i, x'_j \in X$, with $i = 1, 2, \dots, |L|$ and $j = 1, 2, \dots, |U|$, and y_i are the labels of x_i , with $y_i \in \{1, 2, 3, \dots, c\}$ and c being the number of classes.

Various approaches have been proposed for semi-supervised learning based on pseudo-labelling [10], autoencoders [16], and generative adversarial networks (GANs) [19]. Self-training is a simple and intuitive generic semi-supervised learning approach. First, a classifier is trained on labelled examples. Then, the trained model is used to predict labels for unlabelled examples—these labels are called “pseudo-labels” as they are not actual ground-truth labels in the original data.

Finally, the unlabelled data with its pseudo-labels are merged with the initially labelled data, and the model is retrained on the merged data. The prediction-retraining loop is normally iterated multiple times. Early work on self-training, for word sense disambiguation in text documents, can be found in [21].

In contrast, self-supervised learning [9] is a fairly recent development in the literature on deep learning. Self-supervised learning is a form of unsupervised learning that works by creating an artificial supervised learning problem based on unlabelled data—a so-called “pretext” or “auxiliary” task—for instance, detecting whether an image has been rotated or not.

In this paper, we empirically compare three different ways of integrating self-supervised learning with self-training. In addition to testing them in a setting with random initial weights and training with cross-entropy loss, we also consider the effect of metric learning losses and transfer learning, as they have shown promise in self-training [15,17,18].

2 Background

Semi-supervised learning approaches have been developed since the 1970s [12], and self-training, in particular, has continued to attract interest. In the deep learning setting, PseudoLabel [10] is an instance of the most basic and straightforward self-training approach, where a neural network model is iteratively trained on initially labelled and pseudo-labelled examples. Although self-training is an appealing and widely-used process, it is worth noting that Arazo et al. [1] found that it can overfit incorrect pseudo-labels, resulting in confirmation bias. Using MixUp [22] and enforcing a minimum number of initially labelled examples to be included in each mini-batch was proposed to reduce this bias.

In contrast, self-supervised learning is a more recent development. It is an unsupervised learning method that trains a model in a supervised fashion on an artificial supervised learning task generated without any human labelling effort. The following papers all define such tasks for image recognition:

- ExemplarCNN [5]: N different classes are generated by applying transformations on random image patches, and the network is trained to predict the correct class of a given patch.
- RotNet [7]: Geometric transformations such as rotations by 0, 90, 180, and 270 degrees are applied to an image, and the neural network is trained to predict the rotation applied to an image.
- Jigsaw Puzzles [13]: For a given image, nine patches are generated, and the network is trained to predict the correct permutation order of the patches for that image.
- Contrastive learning: The network is trained to differentiate between positive and negative samples, where an image and its augmentations are the positive examples, and all other images are the negative ones [2,3].

Algorithm 1 Self-training using Combined Training.

```

1: Input: Labelled examples  $(x_L, y_L)$ , unlabelled examples  $x_U$ , neural network  $f_\theta$ 
   with parameters  $\theta$ , weight of self-supervised loss  $\lambda_u$ , and selection percentage  $p$ .
2: for each meta-iteration do
3:   for each epoch over  $x_U \cup x_L$  do
4:      $b_L = \text{sample}(x_L, y_L)$ 
5:      $b_U = \text{geometric-transform}(\text{sample}(x_U \cup x_L))$ 
6:     for each mini-batch do
7:        $\mathcal{L} = \mathcal{L}_{SUPER}(f_\theta(b_L)) + \lambda_u \mathcal{L}_{SELF}(f_\theta(b_U))$ 
8:        $\theta = \theta - \nabla_\theta \mathcal{L}$ 
9:     end for
10:  end for
11:   $\text{labels}_U, \text{dist}_U = \text{assign\_labels}(f_\theta, x_U, x_L, y_L)$ 
12:   $x_{new}, y_{new} = \text{select\_top}(\text{dist}_U, \text{labels}_U, p)$ 
13:   $x_L, y_L = \text{concat}((x_L, y_L), (x_{new}, y_{new}))$ 
14:   $x_U = \text{delete\_from}(x_U, x_{new})$ 
15: end for

```

3 Self-training using Self-supervised Learning

Three ways of injecting self-supervision into self-training are proposed below. The specific self-supervision task we use in our experiments is to learn to predict one of six possible geometrical transformations: rotations by 0, 90, 180, or 270 degrees respectively, and horizontal or vertical flips. To investigate the effect of the underlying basic loss function in the different variants of self-training algorithms we consider, we evaluate both cross-entropy loss and triplet loss [20]. When selecting the most confidently predicted examples in self-training, the highest predicted probability is used as confidence when cross-entropy is applied. Triplet loss, being a metric loss, needs a more complex setup based on a nearest-neighbour classifier applied to embeddings generated by the penultimate network layer. The inverse of the distance to the nearest neighbour is used as confidence.

3.1 Combined Training

Combined training (CT) is the most expensive approach we consider: whenever a model is trained in self-training, we jointly train it on both the primary classification task as well as the pretext task, using a combination of self-supervised loss and supervised loss. The supervised and self-supervised branches of the model share the same core convolutional neural network, with only the final layers being different. The loss is a weighted sum of supervised and self-supervised loss:

$$\mathcal{L} = \mathcal{L}_{SUPER}(x_L, y_L) + \lambda_u \mathcal{L}_{SELF}(x_U, y_U), \quad (1)$$

where the hyperparameter $\lambda_u \in [0, 1]$ controls the weight of self-supervised loss.

Algorithm 1 shows pseudocode for each self-training iteration using combined training (CT). For each epoch, all available training examples, i.e., labelled and

Algorithm 2 Self-training using Self-supervised Pretraining.

```

1: Input: Labelled examples  $(x_L, y_L)$ , unlabelled examples  $x_U$ , neural network  $f_\theta$ 
   with parameters  $\theta$ , and selection percentage  $p$ 
2: for each epoch over  $x_U$  do
3:    $b_U = \text{geometric-transform}(\text{sample}(x_U \cup x_L))$ 
4:   for each mini-batch in  $x_U \cup x_L$  do
5:      $\mathcal{L} = \mathcal{L}_{SELF}(f_\theta(b_U))$ 
6:      $\theta = \theta - \nabla_\theta \mathcal{L}$ 
7:   end for
8: end for
9: for each meta-iteration do
10:   $f_\theta = \text{train\_network}(x_L, y_L)$ 
11:   $labels_U, dist_U = \text{assign\_labels}(f_\theta, x_U, x_L, y_L)$ 
12:   $x_{new}, y_{new} = \text{select\_top}(dist_U, labels_U, p)$ 
13:   $x_L, y_L = \text{concat}((x_L, y_L), (x_{new}, y_{new}))$ 
14:   $x_U = \text{delete\_from}(x_U, x_{new})$ 
15: end for

```

unlabelled examples, are employed for self-supervised loss estimation, whereas only the labelled and pseudo-labelled examples are used for supervised loss estimation. The network model is used to assign labels to unlabelled examples, and the top $p\%$ most confidently predicted pseudo-labelled examples are selected and merged with the labelled examples for retraining.

3.2 Self-supervised Pretraining

Pretraining is a common strategy in deep learning: rather than starting with randomly initialised network parameters, we start with parameters obtained by training on a related task. The motivation for pretraining is that the model will learn features that may be useful for the target task. This can be of great importance in semi-supervised learning, where a few labelled and many unlabelled examples are available. It has also been shown that unsupervised pretraining helps neural networks achieve better generalisation [6]. Motivated by this, we consider self-supervision for pretraining on unlabelled examples followed by standard self-training.

Algorithm 2 presents the training procedure for the proposed self-training approach using self-supervised pretraining (SS-Pretrain). At the start of training, the network model is pretrained on all training examples, i.e., labelled and unlabelled examples using self-supervised loss.¹ Following that, simple self-training is applied using supervised loss only on labelled and pseudo-labelled examples.

3.3 Self-training using Single-Step Combined Training

In combined training (see Section 3.1), self-supervised training is applied in each epoch during all self-training iterations. This is rather expensive. Moreover, it

¹ Note that the original labels are ignored in this step.

Algorithm 3 Self-training using Single-step Combined Training.

```

1: Input: Labelled examples  $(x_L, y_L)$ , unlabelled examples  $x_U$ , neural network  $f_\theta$ 
   with parameters  $\theta$ , weight of self-supervised loss  $\lambda_u$ , and selection percentage  $p$ .
2: for each epoch over  $x_U$  do
3:    $b_L = \text{sample}(x_L, y_L)$ 
4:    $b_U = \text{geometric-transform}(\text{sample}(x_U \cup x_L))$ 
5:   for each mini-batch in  $x_U$  do
6:      $\mathcal{L} = \mathcal{L}_{SUPER}(f_\theta(b_L)) + \lambda_u \mathcal{L}_{SELF}(f_\theta(b_U))$ 
7:      $\theta = \theta - \nabla_\theta \mathcal{L}$ 
8:   end for
9: end for
10: for each meta-iteration do
11:    $f_\theta = \text{train\_network}(x_L, y_L)$ 
12:    $\text{labels}_U, \text{dist}_U = \text{assign\_labels}(f_\theta, x_U, x_L, y_L)$ 
13:    $x_{new}, y_{new} = \text{select\_top}(\text{dist}_U, \text{labels}_U, p)$ 
14:    $x_L, y_L = \text{concat}((x_L, y_L), (x_{new}, y_{new}))$ 
15:    $x_U = \text{delete\_from}(x_U, x_{new})$ 
16: end for

```

is conceivable that self-training is no longer beneficial once the model has been sufficiently adapted to the primary task: it may destroy features that are finely tuned to the primary task at hand. Therefore, we also investigate using self-supervision only in the first iteration of self-training. Algorithm 3 shows this self-training approach using single-step combined training (STSSC).

4 Empirical Evaluation

To evaluate the effect of self-supervised learning in self-training, we conduct an extensive comparison using standard benchmark datasets. We always report three different test accuracies: (a) after training on the initially labelled examples, (b) after full self-training, and (c) after training on all labelled training examples available for the benchmark problem. Results obtained for (a) and (c) act as empirical lower and upper bounds for self-training. All results are averaged over three runs with a random selection of initially labelled examples.

4.1 Datasets

We use three benchmark datasets: Street View House Numbers (SVHN), CIFAR-10, and PlantVillage [8]. PlantVillage images have been resized to 96 by 96, 64 by 64, and 32 by 32², respectively, to create three datasets.

Following standard semi-supervised learning practice, a small class-balanced subset of labelled examples is chosen randomly for each dataset. All remaining training examples are used as unlabelled examples. The number of initially labelled examples is 1000 for SVHN, 4000 for CIFAR-10, and 380 for PlantVillage.

² <https://github.com/attaullah/downsampled-plant-disease-dataset>

4.2 Convolutional Architecture

A wide residual network with depth 28 and width 2 (WRN-28-2) is a common choice in semi-supervised classification tasks [14]. In our experiments, we add a fully connected layer at the end of the WRN-28-2 model to produce 64-dimensional embeddings. For all results, to obtain insight into the effect of standard ImageNet pretraining on the outcome of semi-supervised learning, we report test accuracy obtained by running our algorithms by starting with a) randomly initialised network weights and b) ImageNet pretrained weights. For pretraining, downsampled ImageNet [4] images with sizes 32 by 32 are used.

4.3 Algorithm Configurations

For SVHN and CIFAR-10, a mini-batch size of 100 is used, while PlantVillage uses a batch size of 64. Adam is used as the optimiser for updating the network parameters with a learning rate of 10^{-3} for randomly initialised weights and 10^{-4} for ImageNet pretrained weights. After the first iteration of self-training, the learning rate is further reduced by a factor of 0.1. Triplet loss is estimated using l_2 -normalised embeddings and a margin value of 1. A 1-nearest-neighbour classifier is employed on the embeddings to compute test accuracy. Self-training is applied for 25 meta-iterations³. For SVHN and CIFAR-10, $p = 5\%$ pseudo-labels are selected from unlabelled examples in each meta-iteration, while PlantVillage uses $p = 2\%$. For self-supervised loss, the weight parameter λ_u is set to 1 to avoid tuning it separately for every single dataset.

4.4 Results: Detailed Evaluation of the Three Algorithms

Table 1 shows test accuracy for CIFAR-10, SVHN, and the three PlantVillage datasets for the WRN-28-2 network using random and pretrained (ImageNet) weights based on combined training (see Section 3.1). Cross-entropy loss is used for the self-supervised component of the combined loss in Equation 1, but we consider both cross-entropy loss and triplet loss for the supervised component. Values in bold highlight the best test accuracy for each dataset for (a) the empirical lower bound provided by training on initial labels (“N-Labelled”), (b) self-training using combined training, and (c) the upper bound provided by training with all labels (“All-Labelled”), across the two loss function configurations and network initialisations considered. Results for triplet loss are obtained by applying a 1-nearest-neighbour classifier to the learned embeddings.

Excluding the results obtained for the lower and upper bounds based on purely supervised training on the CIFAR-10 data with ImageNet pretraining, it is clear that there is no benefit to be had by using metric learning with triplet loss. The results also show that self-training using cross-entropy provides consistent improvements compared to the lower bound, but there is a large gap compared to the results for the upper bound. We can also see that it is generally

³ https://github.com/attaullah/Self-training/blob/master/Self_supervised.md

Table 1: Test accuracy on WRN-28-2 using random and ImageNet pretrained weights for Combined Training (CT).

	Weights	N-Labelled		Self-training		All-Labelled	
		Cross-entropy	Triplet	Cross-entropy	Triplet	Cross-entropy	Triplet
SVHN	Random	86.37 ± 3.57	89.38 ± 1.69	88.10 ± 3.02	79.47 ± 0.82	95.74 ± 0.02	95.07 ± 0.15
	ImageNet	89.62 ± 1.22	88.98 ± 2.14	90.83 ± 1.67	81.96 ± 3.24	95.16 ± 0.16	94.65 ± 0.35
CIFAR-10	Random	80.67 ± 2.95	81.25 ± 1.77	86.04 ± 0.60	81.51 ± 0.73	91.38 ± 0.39	86.97 ± 3.47
	ImageNet	87.00 ± 1.97	90.04 ± 1.22	89.19 ± 0.65	89.11 ± 2.33	93.47 ± 0.08	93.92 ± 0.09
PLANT32	Random	67.46 ± 3.39	55.27 ± 3.77	72.91 ± 3.68	50.99 ± 1.21	97.97 ± 0.64	91.23 ± 0.86
	ImageNet	77.96 ± 0.58	73.38 ± 1.28	78.18 ± 0.09	66.67 ± 0.60	98.38 ± 0.08	96.58 ± 0.48
PLANT64	Random	72.52 ± 3.86	63.86 ± 0.94	83.52 ± 1.84	59.84 ± 0.38	98.52 ± 0.34	91.99 ± 0.30
	ImageNet	78.39 ± 0.87	76.65 ± 1.06	84.98 ± 0.68	69.26 ± 6.11	99.08 ± 0.10	94.67 ± 0.52
PLANT96	Random	75.12 ± 3.91	65.94 ± 1.43	79.94 ± 1.75	44.55 ± 1.70	98.08 ± 0.65	89.47 ± 0.99
	ImageNet	79.07 ± 0.68	79.04 ± 0.67	84.48 ± 1.30	76.26 ± 1.40	99.15 ± 0.12	95.60 ± 1.27

Table 2: Test accuracy on WRN-28-2 using random and ImageNet pretrained weights for SS-Pretrain.

	Weights	N-Labelled		Self-training		All-Labelled	
		Cross-entropy	Triplet	Cross-entropy	Triplet	Cross-entropy	Triplet
SVHN	Random	89.51 ± 0.74	86.48 ± 2.00	89.72 ± 0.57	83.73 ± 2.46	95.63 ± 0.46	95.05 ± 0.06
	ImageNet	89.90 ± 0.40	90.08 ± 0.60	89.67 ± 0.10	86.05 ± 0.44	95.90 ± 0.05	95.30 ± 0.16
CIFAR-10	Random	72.02 ± 0.50	75.76 ± 2.11	85.50 ± 0.41	75.51 ± 1.88	90.32 ± 0.94	89.11 ± 0.68
	ImageNet	84.83 ± 1.99	86.98 ± 0.84	91.06 ± 0.47	90.18 ± 0.87	93.48 ± 0.09	92.64 ± 0.70
PLANT32	Random	69.80 ± 1.07	72.03 ± 1.51	68.72 ± 0.30	65.04 ± 1.19	97.14 ± 0.68	96.64 ± 0.07
	ImageNet	72.38 ± 1.22	72.46 ± 1.07	75.21 ± 1.34	69.96 ± 0.05	99.24 ± 0.12	97.77 ± 0.77
PLANT64	Random	73.35 ± 1.76	71.10 ± 1.16	73.45 ± 3.32	67.06 ± 1.42	97.94 ± 0.48	87.76 ± 0.96
	ImageNet	75.64 ± 0.54	77.51 ± 1.05	75.50 ± 0.93	68.37 ± 1.43	99.06 ± 0.27	89.48 ± 0.56
PLANT96	Random	73.29 ± 0.72	70.61 ± 0.83	75.19 ± 0.88	67.18 ± 1.81	98.43 ± 0.05	90.54 ± 1.18
	ImageNet	78.37 ± 1.19	78.74 ± 0.22	79.11 ± 1.24	74.97 ± 0.57	99.47 ± 0.12	93.32 ± 1.35

highly beneficial to use ImageNet weights instead of randomly initialised weights, regardless of dataset, loss function, and learning algorithm. The few observed drops in accuracy for the SVHN data are very small.

Table 2 shows test accuracy for CIFAR-10, SVHN, and PlantVillage32, 64, 96 on WRN-28-2 using random and ImageNet weights after applying SS-Pretrain (see Section 3.2). To apply pretraining based on self-supervision, we run it on all training examples—those that exhibit a ground-truth label as well as those missing one—for 120 epochs. After that, we run 200 epochs of fine-tuning on only the labelled training data, followed by standard self-training with 25 meta-iterations, using either cross-entropy or triplet loss. To provide a fair comparison, we also apply self-supervision in the same manner to obtain initial weights for the upper and lower bounds obtained by purely supervised training.

Values in bold highlight the best test accuracy for each dataset obtained by performing self-supervised pretraining for (a) the empirical lower bound provided by training on initial ground-truth labels (“N-Labelled”), (b) self-training, and (c) the upper bound provided by training with all ground-truth labels (“All-Labelled”), across the two loss function configurations and network initialisations considered. Interestingly, metric learning using triplet loss now often achieves better accuracy than applying cross-entropy loss when using only the initially labelled examples. However, cross-entropy consistently achieves higher test accu-

Table 3: Self-training test accuracy on WRN-28-2 using random and ImageNet pretrained weights for STSSC.

	Weights	Cross-entropy	Triplet
SVHN	Random	95.27 ± 0.20	83.62 ± 0.03
	ImageNet	95.12 ± 0.32	83.47 ± 0.98
CIFAR-10	Random	87.24 ± 0.67	79.36 ± 1.02
	ImageNet	91.64 ± 0.38	90.75 ± 1.48
PLANT32	Random	77.39 ± 1.32	54.44 ± 3.72
	ImageNet	86.72 ± 1.10	65.14 ± 1.85
PLANT64	Random	86.13 ± 0.46	56.23 ± 2.01
	ImageNet	89.48 ± 0.37	74.29 ± 1.27
PLANT96	Random	84.71 ± 1.89	58.14 ± 0.82
	ImageNet	89.95 ± 0.37	74.41 ± 3.75

racy than metric learning when performing self-training and when training using all labelled examples. Comparing self-training to the lower bound, we see improvements when using cross-entropy, but self-training reduces accuracy most of the time when considering metric learning. Networks initialised with ImageNet weights generally perform better than those with randomly initialised weights.

Table 3 shows test accuracy of self-training with STSSC (see Section 3.3) for CIFAR-10, SVHN, and PlantVillage32, 64, 96 on WRN-28-2 using random and ImageNet pretrained weights. The reported results are obtained after one iteration of combined training and a further 24 iterations of regular self-training, using either cross-entropy or triplet loss. Values in bold highlight the best test accuracy for each dataset across the two loss function configurations and network initialisations considered. For the corresponding empirical lower and upper bounds, please consult the results in Table 1. Self-training using cross-entropy outperforms metric learning for all five datasets for both randomly initialised and ImageNet pretrained weights. Moreover, in the majority of the cases, using ImageNet weights achieves higher test accuracy than using random weights. Notably, for the one dataset where random weights yield higher accuracy, the difference in performance is very small.

To further illustrate the effect of starting the training process with ImageNet weights rather than random weights, Figure 1 plots the test accuracy for STSSC over the meta-iterations across three different runs using a) randomly initialised weights and b) ImageNet weights on CIFAR-10 using 4000 labelled examples initially. Cross-entropy loss is used as the loss function throughout. The accuracy curves show definite improvements as more meta-iterations are performed, for both weight initialisations, but the ImageNet version starts from a higher initial accuracy and retains this advantage over the 25 meta-iterations of self-training.

4.5 Comparing Self-training Approaches

Table 4 presents a summary view and comparison of the three approaches to integrating self-supervision into self-training, and, importantly, reports results for

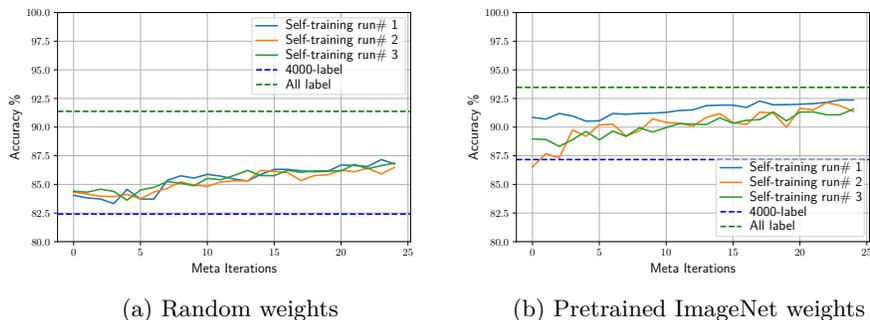


Fig. 1: Comparison of self-training using STSSC on WRN-28-2 for CIFAR-10 using cross-entropy loss.

Table 4: Test accuracy after self-training on WRN-28-2 using various self-supervised settings.

	SVHN	CIFAR-10	PLANT32	PLANT64	PLANT96
Random weights					
Simple	94.28 ± 0.45	83.05 ± 0.99	68.24 ± 2.44	69.36 ± 1.80	76.09 ± 2.35
CT	88.10 ± 3.02	86.04 ± 0.60	72.91 ± 3.68	83.52 ± 1.84	79.24 ± 1.45
SS-Pretrain	89.72 ± 0.57	85.50 ± 0.41	68.72 ± 0.30	73.45 ± 3.32	75.19 ± 0.88
STSSC	95.27 ± 0.20	87.08 ± 0.66	77.39 ± 1.32	86.13 ± 0.46	84.71 ± 1.89
ImageNet weights					
Simple	90.84 ± 0.48	91.62 ± 0.58	87.48 ± 0.84	92.60 ± 1.03	89.65 ± 0.88
CT	90.83 ± 1.67	89.19 ± 0.65	78.18 ± 0.09	84.60 ± 0.77	84.48 ± 1.30
SS-Pretrain	89.67 ± 0.10	91.06 ± 0.47	75.21 ± 1.34	75.50 ± 0.93	79.11 ± 1.24
STSSC	95.12 ± 0.32	91.65 ± 0.43	86.72 ± 1.10	89.48 ± 0.37	89.95 ± 0.37

self-training without using any self-supervised learning—denoted by the name “Simple” in the table. Cross-entropy loss is used for both supervised and self-supervised loss estimation throughout. The results show that STSSC performs best overall, demonstrating the importance of applying self-supervised learning in a judicious manner when combining it with self-training: applying it in all meta-iterations is generally better than simply using it for pretraining, but among the three algorithm variants considered, it is clearly best to only apply it in the first meta-iteration. Compared to the simple self-training baseline without self-supervision, STSSC performs substantially better when starting with randomly initialised weights. However, while the performance also generally improves when using ImageNet weights, and substantially so for the SVHN data, STSSC only wins in three cases, with the simple approach performing best for the two smaller-sized versions of the PlantVillage datasets.

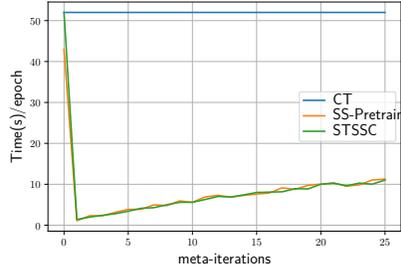


Fig. 2: Time spent on each epoch over meta-iterations of self-training approaches using self-supervised learning.

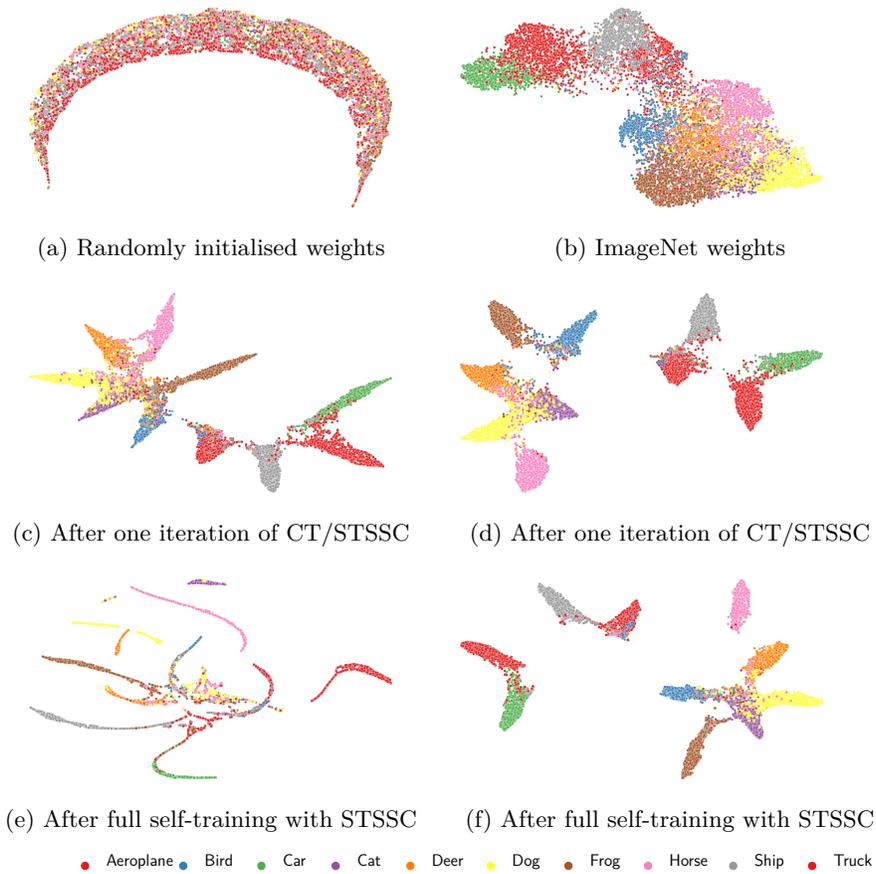


Fig. 3: UMAP visualisation of CIFAR10 test set embeddings obtained from WRN-28-2 using random weights [left: (a),(c),(e)] and ImageNet weights [right: (b),(d),(f)] by applying STSSC training.

Figure 2 compares the time needed for each epoch over the self-training meta-iterations for all three self-supervision-based self-training approaches, i.e., CT, SS-Pretrain, and STSSC. To estimate time, WRN-28-2 is trained on CIFAR-10 with 4000-labelled examples initially for 25 meta-iterations of self-training. Time spent by each epoch for CT is constant and higher than for SS-Pretrain and STSSC. After the first iteration, SS-Pretrain and STSSC need very little time compared to CT.

4.6 Visualising Embeddings

Embeddings generated for CIFAR-10 by STSSC are visualised in Figure 3, using UMAP [11] to produce two-dimensional projections. The embeddings are obtained at the start of training: (a) and (b); after one meta-iteration: (c) and (d); and after 25 meta-iterations of self-training: (e) and (f). Configurations (a), (c), and (e) use randomly initialised weights, (b), (d), and (f) use ImageNet weights. Self-training clearly generates well-separated clusters.

5 Conclusion

In this paper, we have shown that self-supervision can substantially improve self-training for image classification problems. Self-supervision was applied in three different ways. The best performance was achieved by using self-supervision only in the first iteration of self-training, with cross-entropy loss and initial weights obtained from ImageNet. There are a number of directions for future research, including more complex scenarios such as self-supervised pretraining followed by a limited number of combined training meta-iterations, followed by further regular self-training. Alternatively, one could consider the use of different and potentially more challenging self-supervision tasks.

References

1. Arazo, E., Ortego, D., Albert, P., O’Connor, N.E., McGuinness, K.: Pseudo-labelling and confirmation bias in deep semi-supervised learning. In: 2020 International Joint Conference on Neural Networks (IJCNN). pp. 1–8. IEEE (2020)
2. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International Conference on Machine Learning. pp. 1597–1607. PMLR (2020)
3. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.E.: Big self-supervised models are strong semi-supervised learners. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Advances in Neural Information Processing Systems, December 6–12, 2020, virtual (2020)
4. Chrabaszcz, P., Loshchilov, I., Hutter, F.: A downsampled variant of ImageNet as an alternative to the CIFAR datasets. arXiv preprint arXiv:1707.08819 (2017)
5. Dosovitskiy, A., Fischer, P., Springenberg, J.T., Riedmiller, M., Brox, T.: Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(9), 1734–1747 (2015)

6. Erhan, D., Courville, A., Bengio, Y., Vincent, P.: Why does unsupervised pretraining help deep learning? In: Proceedings of the 13th International Conference on Artificial Intelligence and Statistics. pp. 201–208. JMLR Workshop and Conference Proceedings (2010)
7. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: Int. Conference on Learning Representations (2018)
8. Hughes, D.P., Salathé, M.: An open access repository of images on plant health to enable the development of mobile disease diagnostics through machine learning and crowdsourcing. CoRR (2015), <http://arxiv.org/abs/1511.08060>
9. Jing, L., Tian, Y.: Self-supervised visual feature learning with deep neural networks: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence (2020)
10. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on Challenges in Representation Learning, ICML. vol. 3, p. 2 (2013)
11. McInnes, L., Healy, J., Saul, N., Großberger, L.: UMAP: Uniform Manifold Approximation and Projection. Journal of Open Source Software **3**(29), 861 (2018)
12. McLachlan, G.J.: Iterative reclassification procedure for constructing an asymptotically optimal rule of allocation in discriminant analysis. Journal of the American Statistical Association **70**(350), 365–369 (1975)
13. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European Conf. on Computer Vision. pp. 69–84. Springer (2016)
14. Oliver, A., Odena, A., Raffel, C., Cubuk, E.D., Goodfellow, I.J.: Realistic evaluation of deep semi-supervised learning algorithms. In: Proc. of the 32nd Int. Conference on Neural Information Processing Systems. pp. 3239–3250 (2018)
15. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Learning and transferring mid-level image representations using convolutional neural networks. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1717–1724 (2014)
16. Rasmus, A., Berglund, M., Honkala, M., Valpola, H., Raiko, T.: Semi-supervised learning with ladder networks. In: Advances in Neural Information Processing Systems. pp. 3546–3554 (2015)
17. Sahito, A., Frank, E., Pfahringer, B.: Semi-supervised learning using Siamese networks. In: Liu, J., Bailey, J. (eds.) AI 2019: Advances in Artificial Intelligence. pp. 586–597. Springer International Publishing, Cham (2019)
18. Sahito, A., Frank, E., Pfahringer, B.: Transfer of pretrained model weights substantially improves semi-supervised image classification. In: Australasian Joint Conference on Artificial Intelligence. pp. 433–444. Springer (2020)
19. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training GANs. In: Advances in Neural Information Processing Systems. pp. 2234–2242 (2016)
20. Schroff, F., Kalenichenko, D., Philbin, J.: FaceNet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 815–823 (2015)
21. Yarowsky, D.: Unsupervised word sense disambiguation rivalling supervised methods. In: 33rd Annual Meeting of the Association for Computational Linguistics. pp. 189–196 (1995)
22. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimisation. In: International Conference on Learning Representations (2018)