

An analysis of over-sampling labeled data in semi-supervised learning with FixMatch

Miquel Martí i Rabadán^{*1,2}, Sebastian Bujwid¹, Alessandro Pieropan², Hossein Azizpour¹, and Atsuto Maki¹

¹KTH Royal Institute of Technology, Sweden

²Univrses AB, Sweden

Abstract

Most semi-supervised learning methods over-sample labeled data when constructing training mini-batches. This paper studies whether this common practice improves learning and how. We compare it to an alternative setting where each mini-batch is uniformly sampled from all the training data, labeled or not, which greatly reduces direct supervision from true labels in typical low-label regimes. However, this simpler setting can also be seen as more general and even necessary in multi-task problems where over-sampling labeled data would become intractable. Our experiments on semi-supervised CIFAR-10 image classification using FixMatch show a performance drop when using the uniform sampling approach which diminishes when the amount of labeled data or the training time increases. Further, we analyse the training dynamics to understand how over-sampling of labeled data compares to uniform sampling. Our main finding is that over-sampling is especially beneficial early in training but gets less important in the later stages when more pseudo-labels become correct. Nevertheless, we also find that keeping *some* true labels remains important to avoid the accumulation of confirmation errors from incorrect pseudo-labels.

1 Introduction

Semi-supervised learning has recently established itself as an effective approach to greatly reduce the annotation costs in developing deep learning mod-

els. With labels only for a small portion of the entire dataset, the performance approaches that of fully-supervised learning, e.g. giving competitive results in the context of image classification using less than 1% of labeled samples [3, 22, 18, 23].

Semi-supervised methods have explored different flavours of consistency or perturbation-based regularization [21] combined with entropy minimization [6] or pseudo-labeling [9] by adding one or more extra terms to the training loss computed on the unlabeled samples.

The consistency-based regularization losses aim at exploiting the smoothness assumption [21], i.e. the predictions of a model should not depend on small perturbations in the input such as injecting noise [16, 2, 8], adversarial perturbations [11] or random transformations [17, 8, 3, 22, 18]. Pseudo-labeling and entropy minimization methods enforce instead high-confidence (low-entropy) predictions for unlabeled samples, directly exploiting the low-density assumption [21], i.e. decision boundaries should lie on low-density regions. These approaches are closely related to self-training [13, 25, 1, 23], where pseudo-labels are extracted from a model pre-trained on labeled data and then used for training in a second step or in multiple iterative steps.

Yet, all semi-supervised learning methods face a common issue. Since they bootstrap their own predictions in order to use them as targets, many of these will be incorrect and can lead to so-called *confirmation bias* [20, 1], i.e. fitting incorrect labels that do not represent the true task and thus slowing down or stopping its learning. This is more likely to happen at the start of training since the predictions of the model are not much better than random and the amount of misclassified samples is large.

*Corresponding Author: miquelmr@kth.se

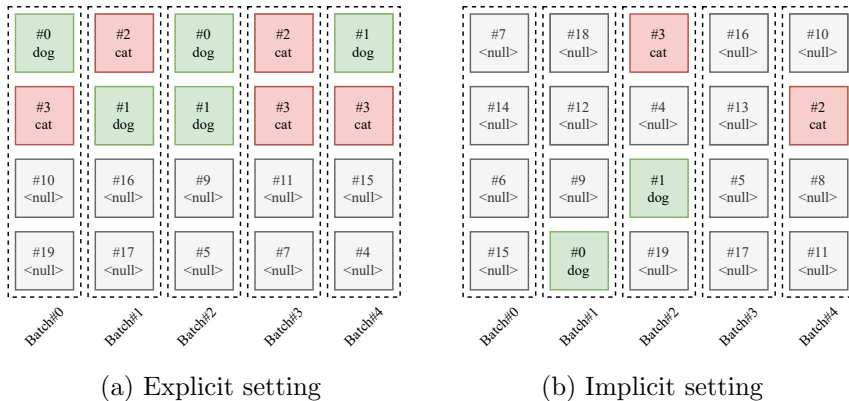


Figure 1: (a) and (b) A toy dataset with 20 samples and a task with 2 classes - dog (green) and cat (red) - with 2 labeled samples per class, and a ratio of labeled samples of 20%. Sampled into 5 training mini-batches of 4 samples each, with over-sampling of labeled data in (a) and uniform sampling in (b). In (a) we *choose* to have an equal number of labeled and unlabeled samples in each mini-batch. Due to this, each labeled sample (red/green, #0-3) is seen 4x more often than each unlabeled sample (gray, #4-19). In (b), each mini-batch can contain different amounts of labeled samples as it is uniformly sampled regardless of labels and each sample is seen exactly once. Labeled samples are seen 2.5x more often in the explicit setting.

Most works try to alleviate this issue by using a ramp-up from zero of the weight for the semi-supervised loss term over training. This greatly reduces the influence of possibly wrong predictions on unlabeled data in the early phases of training. Thus, the only way to learn at the start is from true labels, but since these are scarce there is little supervision available and learning is slow. This in turn makes it more difficult to get useful supervision through the semi-supervised loss.

Repeating labeled samples more often than unlabeled ones during training is a common mini-batch sampling strategy in semi-supervised learning and directly and artificially increases the influence of labeled data over learning [16, 14, 3, 20, 11, 1, 22, 18, 23]. To the best of our knowledge, this mechanism was never well studied with the goal of enhancing the training signal from true labels and it likely just happened as a side effect of common implementation choices when constructing batches from separate datasets for labeled and unlabeled samples.

In this paper, we aim at studying if it is essential to repeat labeled samples more frequently than unlabeled samples or over-sample labeled data for state-of-the-art semi-supervised methods to work well. We call this way of sampling mini-batches the *explicit* semi-supervised learning setting due to its distinction between labeled and unlabeled data in how the data is presented. We compare it to

an alternative setting in which all data is sampled uniformly, regardless of being labeled or not. We call the latter the *implicit* semi-supervised learning setting.¹ Figure 1 illustrates both approaches.

Exploring the simpler implicit setting will highlight the benefits which the common over-sampling of labeled data in the explicit setting brings and which mechanisms allow to achieve these. Understanding the differences between both settings could then shed light on how sampling affects the learning process in semi-supervised methods and could ultimately guide future work into developing more efficient and better-performing methods.

Moreover, there are concrete reasons to avoid the explicit setting. Most importantly, trying to use it for multi-task problems easily leads into a road-block. Samples cannot be said to be labeled or unlabeled in general, since this property can only be assigned to a sample in relation to each task. Thus, samples can be labeled for one or more tasks but unlabeled for others, labeled for no tasks or for all. This leads to different sets of samples labeled for each task, which might or might not overlap.

In Figure 2, we compare an extension of the explicit setting for multi-task scenarios to the implicit setting. One possible direct way to recreate the

¹In the rest of the paper we will use both nomenclatures for the sampling approaches interchangeably.

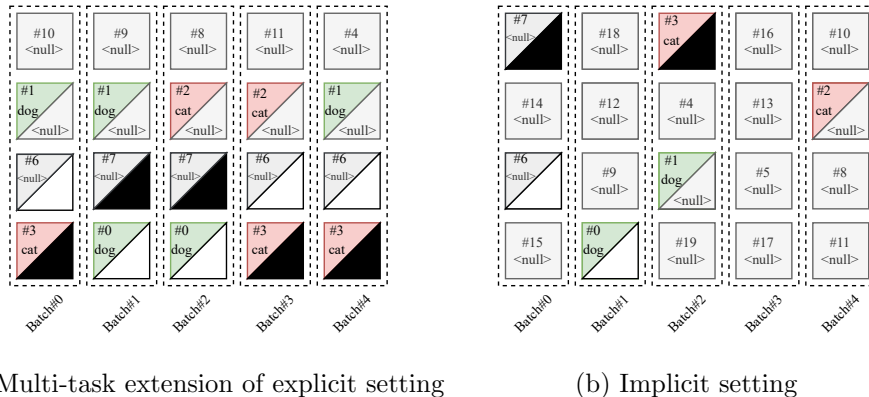


Figure 2: (a) and (b) The labeling of the dataset in Figure 1 is extended with a new task with 2 classes - black and white - with 2 labeled samples per class as well. In (a) we extend the over-sampling of labeled data to the multi-task case so that in each batch there is the same number of labeled and unlabeled samples for each task. In (b) the uniform sampling means the sampling is independent of the labels, just as for the single task case.

explicit setting for multi-task learning is to divide each mini-batch into multiple sub-parts, one for each possible configuration of available task labels. Such a mini-batch can include parts where all samples are labeled or unlabeled, parts with samples labeled only for one or more tasks, etc. Instead, the implicit setting can be used in multi-task problems without any modifications since the sampling is independent of the labels.

The number of sub-parts in a mini-batch constructed using this extension of the explicit setting grows exponentially with the number of tasks T as 2^T if all label configurations are present in the dataset. Therefore, if the number of tasks is large, the number of mini-batch parts can quickly become larger than the total batch size that can be used in practice. Even with fewer tasks, the possible size of each mini-batch part becomes a new hyper-parameter which can be difficult to control since it is limited by the total batch size and the number of parts. Setting the number of tasks to one recovers the original explicit setting, with one part for labeled data for the only task in the problem and one for unlabeled data. Other sampling approaches for semi-supervised multi-task learning problems are likely possible, but our aim is to highlight the issues with the explicit setting that make it not suitable for this class of problems.

To study the impact of different approaches to sampling training mini-batches on semi-supervised performance, we compare the training character-

istics of both the implicit and explicit settings. We run semi-supervised image classification experiments on CIFAR-10 and CIFAR-100 [7], varying the number of available labels. As a semi-supervised method we adopt FixMatch [18], a standard, state-of-the-art method. FixMatch enforces consistency between pseudo-labels obtained from weakly transformed versions of one sample and the predictions obtained from strongly transformed versions of the same sample, incorporating in the same method pseudo-labeling and consistency regularization. We find that both the implicit and explicit settings work for reasonably long training runs, but the latter performs better. This difference is smaller in higher-label regimes or with longer training budgets and we partially attribute it to the low efficiency of learning from true labels at the start of training when using the implicit setting, which becomes more important for shorter training runs.

In sum, this paper (i) emphasizes the importance of the choice of mini-batch sampling strategy in semi-supervised learning, (ii) analyses the differences between the common over-sampling of labeled data and standard uniform sampling throughout the training process using FixMatch, and (iii) demonstrates the viability of uniform sampling with long training budgets, enabling its use in cases where over-sampling labeled data is not possible, such as in multi-task learning.

2 Related work

Most semi-supervised learning methods use some mechanism that minimizes the effect of confirmation bias, either at the earlier stages of training or throughout. One such mechanism directly ramps up the weight for the unsupervised loss term during training, starting from zero. It was originally proposed in Grandvalet and Bengio [6] and is reported to be used in most works [9, 2, 8, 20, 3]. FixMatch does not use the weight ramp-up as it instead sets a threshold to use only confident predictions as pseudo-labels and averages the loss on these over the total number of unlabeled samples instead of over the number of pseudo-labels used. This achieves a similar effect to the ramp-up since it decreases the contribution of the unsupervised loss when the model produces low confidence predictions early in the training.

A second mechanism is to over-sample labeled data, i.e. repeating labeled samples more often than unlabeled ones when constructing mini-batches. This is usually implemented by splitting the labeled and unlabeled samples as two separate datasets and sampling the labeled and unlabeled parts of the mini-batch from each. Then, the first set that gets depleted is repeated as many times as needed until the other one is used up. Since the labeled dataset is typically much smaller, this results in the repetition of labeled samples at a higher rate. This approach was already used in Rasmus *et al.* [16]² although without any explicit mention in the paper. Most recent works in semi-supervised learning use this setting, both using consistency-based, pseudo-labeling or self-training methods, but did not study the impact of sampling mini-batches in this way. Some methods set equal sizes for the labeled and unlabeled parts of each training mini-batch [16, 14, 3], but some use a larger unlabeled batch size [20, 11, 1, 22, 18, 23].

Both FixMatch [18] and Noisy Student [23] include ablation experiments on this ratio and conclude that a larger ratio of unlabeled samples in each training mini-batch leads to better performance. Arazo *et al.* [1] is the only work we are aware of that studied the effect of keeping a minimum amount of true labels in each mini-batch in pseudo-labeling, showing that it is indeed an ef-

fective way of reducing confirmation bias. Only Temporal Ensembling [8] and one of the experiments in Mean Teacher [20] use the implicit setting. However, the latter mentions that repeating labeled samples can be beneficial since "the supervised training signal is strong enough early on to train quickly and prevent getting stuck into uncertainty" [5]. All methods relying on the explicit setting can technically be adapted to use the implicit setting instead.

3 Explicit vs. implicit setting

We aim at understanding how essential the explicit setting is for good performance in semi-supervised learning, and whether the implicit one can achieve competitive results. To do so, we will compare the final classification performance of both settings.

3.1 Experimental setting

We compare the explicit and implicit settings head to head when using FixMatch [18], which uses asymmetric perturbations in the input space in a siamese network configuration with two branches. The teacher branch gets weakly augmented samples (only flip-and-shift operations) and generates pseudo-labels from them, filtering out those below a certain confidence threshold. The student branch gets strongly augmented samples via RandAugment [4] and is trained using the standard cross-entropy loss against both true labels (if available) and the pseudo-labels generated by the teacher.

We use the semi-supervised image classification problem on the CIFAR-10 [7] dataset for the main experiments, as it is a standard benchmark in the deep semi-supervised learning literature [14]. We use three different amounts of labeled samples - 40, 250 and 4000 - simulating respectively a problem with extremely limited labeled data at only four labels per class, a scarcely-labeled problem with approximately 0.5% labeled samples, and a semi-supervised learning problem with a larger amount of labeled samples with approximately 9% labeled samples. We reserve 10% of the total samples in the training set for validation and use the rest as training data, which we split into labeled and unlabeled subsets. All splits are done preserving the class balance. The validation set is used for early

²See: <https://github.com/CuriousAI/ladder/blob/5a8daa1760535ec4aa25c20c531e1cc31c76d911/run.py#L74>

stopping and for hyper-parameter tuning based on random search with a budget of 100 trials per experiment. We pick the hyper-parameters leading to the highest validation accuracy on a single seed and data split. We tune the following hyper-parameters for each case: learning rate, batch size, weight of the semi-supervised loss, and the ratio of labeled data in each batch for the explicit setting,. Finally, we report the final classification accuracy values on the test set as the mean and standard deviation over five different random seeds and five different splits of the labeled samples in the training set.

In addition, we perform experiments on CIFAR-100 with 2500 and 10000 labels, following the same setting as for CIFAR-10 but not tuning the hyper-parameters. Instead, we use the ones reported in FixMatch [18] for the explicit setting and our best single guess of a decent set of hyper-parameters for the implicit setting.

For CIFAR-10 experiments, we use a Wide ResNet-28-2 [24] (WRN-28-2) as the base network with a linear classifier after the pooling layer and train it via stochastic gradient descent with Nesterov momentum of 0.9 [15, 12, 19], a cosine learning rate decay schedule [10], and weight decay of 0.0005. For CIFAR-100, we use WRN-28-5 and a smaller weight decay of 0.0001. Evaluation is done on an exponential moving average of the model parameters with decay 0.999 [18], which we found gave more stable results and significantly better performance early in the training.

Each training run has a maximum budget for the total number of samples that each model sees during training. We set this budget to be equivalent to 1000 epochs on the full training set, i.e. 45 million samples, after removing 5000 samples from the training set for validation. Therefore we refer to an epoch as one pass through as many samples as the size of the training set, regardless of whether each sample has an associated label or not, or how many times it has been repeated.

Other works [14, 18] have used a fixed number of labeled samples or a fixed number of training steps instead. We believe this approach leads to unfair comparisons between supervised baselines or when using different label ratios in each batch since for higher ratios the models will effectively see many more samples over training. Moreover, comparing the total amount of samples seen aligns better with the amount of compute used for each training

CIFAR-10	40 labels	250 labels	4000 labels	All labels
Supervised	36.77 ± 4.48	59.88 ± .73	87.39 ± .20	96.54 ± .11
FixMatch(E)	83.44 ± 6.68	93.01 ± .58	94.89 ± .16	-
FixMatch(I)	73.90 ± 8.04	91.42 ± .98	94.84 ± .05	-
FixMatch(E) 6x	85.40 ± 2.83	94.40 ± .78	-	-
FixMatch(I) 6x	86.59 ± 4.14	94.00 ± .54	-	-
CIFAR-100	2500 labels	10000 labels	All labels	
Supervised	47.26	67.58	82.38	
FixMatch(E)	62.86	74.61	-	
FixMatch(I)	55.53	71.83	-	

Table 1: Classification accuracy on test data for supervised baselines and FixMatch with both implicit (I) and explicit (E) settings using different amounts of labels. Shown as mean and standard deviation on 5 runs with different seeds and data splits for each CIFAR-10 experiment and a single seed and data split for CIFAR-100. 6x indicates six times longer training budget.

run. Using this definition of training length, MixMatch [3] training runs go through ~ 134 million samples since they use a total batch size of 128 (64 for labeled samples and 64 for unlabeled samples) but they only count the labeled ones towards their definition of total length. FixMatch, using an unlabeled batch size seven times larger, uses a total batch size of 512 and the total number of samples seen during training is ~ 536 million.

We choose to use a smaller budget on purpose to better reflect what would happen in practical applications and to keep contained the computational budget required for our experiments. We noticed that for smaller amounts of labeled data longer training budgets are important, so we include longer experiments in such cases as well.

3.2 Experimental results

Table 1 shows the classification accuracy of our models trained on CIFAR-10 and CIFAR-100 with different amounts of labels when using FixMatch in either the explicit or implicit settings.

We observe that both settings outperform supervised baselines by a significant margin in CIFAR-10. For 4000 labels both settings perform similarly, however we see a larger gap favouring the explicit setting by 1.6 percentage points when only using 250 labels. A significantly larger gap of 9.5 percentage points appears when using 40 labels, although both settings show a much larger standard deviation. Comparing individual runs on each data split

for both settings, we can see that the large standard deviations are caused mainly by the choice of labeled samples in each data split but also that the implicit setting is more sensitive to it.

We hypothesize that this gap is due to the much smaller amount of labeled samples present (on average) in each training mini-batch for the lower-label regimes in the implicit setting, where for example around only one labeled sample appears for every 200 unlabeled samples for the 250-label case. Due to this lack of direct supervision at the start, the model learns at a very slow pace. This is because the model is initially relying only on labeled data as the predictions on unlabeled data are below the confidence threshold. A longer training budget of 6000 epochs improves the performance of both settings. The gap in test accuracy between them reduces over training, and both settings end up performing on par.

We see similar results for CIFAR-100. The gap between settings is larger at lower label regimes, reaching 7.3 percentage points better accuracy for the explicit setting with 2500 labels and 2.8 percentage points with 10000 labels, and both perform better than the supervised baselines.

4 Analysis

Training dynamics We analyze in Figure 3 the training curves of models in both settings for 250 and 4000 labels to illustrate the impact of their different training dynamics. We can see that the explicit setting does indeed fit quicker the labeled data; training error, shown in (a, top), and supervised loss, in (b, top), both decrease quickly thanks to more direct supervision from repeating labeled samples. The models start giving confident predictions earlier, as shown in (c, top), and thus the number of pseudo-labels, in (c, bottom), and the semi-supervised loss, in (b, bottom), grow quicker. On the other hand, for the implicit setting, the model takes much longer to fit the training data and to produce high-confidence predictions that will be used as pseudo-labels. Therefore, the semi-supervised loss grows much slower, having no effect for the first 100 epochs while learning is very slow and test accuracy is barely better than random. This clearly points to a lack of direct supervision

from true labels early in training as one cause for the implicit setting lagging behind in terms of final test error with our relatively short training budget with fewer labeled samples.

Moreover, we make use of privileged information (the labels dropped to simulate unlabeled samples) to compute and show in Figure 3 (d) the ratio of correct predictions on unlabeled samples and the ratio of the correct pseudo-labels, i.e. the subset of predictions on unlabeled data whose confidence was above the threshold of 0.95. We see that not only are pseudo-labels used earlier in the explicit setting but are also cleaner, i.e. more of them are correct, from early on and can thus help guide learning of the true classification task better. Also, we can see that the usage of confident predictions on unlabeled data as pseudo-labels is effective since the subset of chosen pseudo-labels has a ratio of correctness consistently higher than that of all predictions on unlabeled data.

This analysis holds for both amounts of labeled samples shown, although the differences are much smaller for the 4000-label experiments. In this last case, the implicit setting has a lower test error earlier during training which we hypothesize might be due to the larger over-fitting to labeled samples in the explicit case, before the loss on pseudo-labels starts having a regularizing effect and the model can generalize to unseen data.

Training without direct supervision We hypothesize that supervision from true labels loses importance when the model can already generate cleaner pseudo-labels, which could explain why either using more labels or longer training budgets reduces the performance gap between settings.

To verify this hypothesis, we run FixMatch with the semi-supervised loss only but starting training from different supervised checkpoints. All models improve their test accuracy at first - the best goes from 87.39% to 92.9% - but their accuracy ends up dropping suddenly to levels close to random while fitting the pseudo-labels. This can be attributed to the accumulation of confirmation errors and the lack of supervision from true labels to correct for these. We see this happen as well for long experiments with the implicit setting and few labels, but with learning recovering after the dip. Thus, we conclude that seeing some labeled samples is still required to guide learning to the true task even when few pseudo-labels are incorrect.

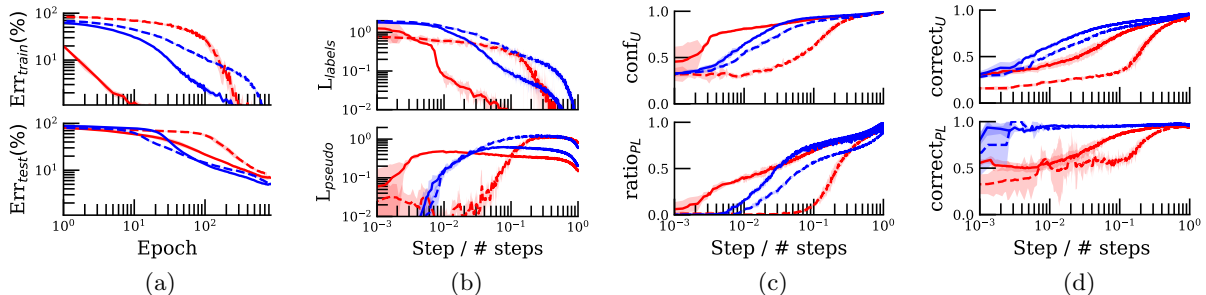


Figure 3: Training curves for CIFAR-10 with 250 labels (red) or 4000 labels (blue) using FixMatch in the explicit (solid line) or implicit (dashed line) settings. (a) Top: training error on labeled samples. Bottom: test error on the unseen test set. Values at the end of each epoch. (b) Top: evolution over training of cross-entropy loss on labeled training data. Bottom: semi-supervised cross-entropy loss between pseudo-labels and predictions on unlabeled data. (c) Top: average confidence of predictions on unlabeled samples in each batch. Bottom: ratio of predictions used as pseudo-labels. (d) Top: Ratio of correct predictions on unlabeled data. Bottom: Ratio of correct pseudo-labels on each batch.

5 Conclusion

Our experiments in semi-supervised image classification show that using the implicit setting incurs a penalty in final accuracy compared to using the explicit setting. How large this gap becomes depends, however, on the amount of labels and training budget. Both settings still outperform supervised baselines by a large margin. The explicit setting is significantly more efficient in the beginning but the difference between both settings is reduced over training and both perform on par when more labeled data is available and for longer training runs. However, seeing labeled samples remains important even in the later stages of training, as it avoids falling into confirmation bias. The more general aspects like the required amount of labels and the most effective way to use them are yet to be studied. Overall, these are important results since the implicit setting is much simpler to use, potentially allowing semi-supervised learning on settings like partially-labeled multi-task problems which would not be possible in the conventional explicit setting.

This analysis was done in the context of FixMatch, which we believe is a good representative of modern semi-supervised deep learning methods combining consistency regularization and pseudo-labeling approaches. Future research will be directed to see if the results also apply to semi-supervised methods that are substantially different.

Finally, we suggest two possible alternative ap-

proaches to mini-batch sampling in the multi-task scenario. First, if there is a large enough amount of samples labeled for all tasks, a version of the explicit setting could be constructed using *fully* labeled samples versus *partially* labeled samples. Second, we could construct a variation of the implicit setting with a guaranteed minimum number of samples labeled for each task so that there is always *some* direct supervision from true labels for each task, in each mini-batch. Studying these and other possible sampling methods will be an interesting direction for future work.

Acknowledgments

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

The support by the Swedish Research Council through grant agreement no. 2016-04022 is also gratefully acknowledged.

The computations were partially enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) at Chalmers Centre for Computational Science and Engineering (C3SE) partially funded by the Swedish Research Council through grant agreement no. 2018-05973.

References

- [1] E. Arazo, D. Ortego, P. Albert, N. E. O'Connor, and K. McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2020. doi: 10.1109/IJCNN48605.2020.9207304.
- [2] P. Bachman, O. Alsharif, and D. Precup. Learning with pseudo-ensembles. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3365–3373, 2014. URL <https://proceedings.neurips.cc/paper/2014/hash/66be31e4c40d676991f2405aaecc6934-Abstract.html>.
- [3] D. Berthelot, N. Carlini, I. J. Goodfellow, N. Papernot, A. Oliver, and C. Raffel. Mixmatch: A holistic approach to semi-supervised learning. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5050–5060, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/1cd138d0499a68f4bb72bee04bbec2d7-Abstract.html>.
- [4] E. D. Cubuk, B. Zoph, J. Shlens, and Q. Le. Randaugment: Practical automated data augmentation with a reduced search space. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/d85b63ef0ccb114d0a3bb7b7d808028f-Abstract.html>.
- [5] CuriousAI. Curiousai/mean-teacher: A state-of-the-art semi-supervised method for image recognition. URL <https://github.com/CuriousAI/mean-teacher/blob/51d20acd65/README.md#tips-for-choosing-hyperparameters-and-other-tuning>.
- [6] Y. Grandvalet and Y. Bengio. Semi-supervised learning by entropy minimization. In *Advances in Neural Information Processing Systems 17 [Neural Information Processing Systems, NIPS 2004, December 13-18, 2004, Vancouver, British Columbia, Canada]*, pages 529–536, 2004. URL <https://proceedings.neurips.cc/paper/2004/hash/96f2b50b5d3613adf9c27049b2a888c7-Abstract.html>.
- [7] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. URL <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- [8] S. Laine and T. Aila. Temporal ensembling for semi-supervised learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=BJ6oOfqge>.
- [9] D.-H. Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3-2, page 896, 2013. URL https://storage.googleapis.com/kaggle-forum-message-attachments/8924/pseudo_label_final.pdf.
- [10] I. Loshchilov and F. Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Skq89Scxx>.

- [11] T. Miyato, S. I. Maeda, M. Koyama, and S. Ishii. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41: 1979–1993, 8 2019. ISSN 19393539. doi: 10.1109/TPAMI.2018.2858821.
- [12] Y. E. Nesterov. A method for solving the convex programming problem with convergence rate $o(1/k^2)$. *Dokl. Akad. Nauk SSSR*, 269: 543–547, 1983.
- [13] K. Nigam and R. Ghani. Analyzing the effectiveness and applicability of co-training. In *Proceedings of the Ninth International Conference on Information and Knowledge Management, CIKM '00*, page 86–93, New York, NY, USA, 2000. Association for Computing Machinery. ISBN 1581133200. doi: 10.1145/354756.354805.
- [14] A. Oliver, A. Odena, C. Raffel, E. D. Cubuk, and I. J. Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 3239–3250, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/c1fea270c48e8079d8ddf7d06d26ab52-Abstract.html>.
- [15] B. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964. ISSN 0041-5553. doi: 10.1016/0041-5553(64)90137-5.
- [16] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko. Semi-supervised learning with ladder networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3546–3554, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/378a063b8fdb1db941e34f4bde584c7d-Abstract.html>.
- [17] M. Sajjadi, M. Javanmardi, and T. Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1163–1171, 2016. URL <https://proceedings.neurips.cc/paper/2016/hash/30ef30b64204a3088a26bc2e6ecf7602-Abstract.html>.
- [18] K. Sohn, D. Berthelot, N. Carlini, Z. Zhang, H. Zhang, C. Raffel, E. D. Cubuk, A. Kurakin, and C. Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/06964dce9addb1c5cb5d6e3d9838f733-Abstract.html>.
- [19] I. Sutskever, J. Martens, G. E. Dahl, and G. E. Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, pages 1139–1147. JMLR.org, 2013. URL <http://proceedings.mlr.press/v28/sutskever13.html>.
- [20] A. Tarvainen and H. Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In I. Guyon, U. von Luxburg,

- S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 1195–1204, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/68053af2923e00204c3ca7c6a3150cf7-Abstract.html>.
- [21] J. E. van Engelen and H. H. Hoos. A survey on semi-supervised learning. *Machine Learning*, 109:373–440, 2 2020. ISSN 15730565. doi: 10.1007/s10994-019-05855-6.
- [22] Q. Xie, Z. Dai, E. H. Hovy, T. Luong, and Q. Le. Unsupervised data augmentation for consistency training. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/44feb0096faa8326192570788b38c1d1-Abstract.html>.
- [23] Q. Xie, M. Luong, E. H. Hovy, and Q. V. Le. Self-training with noisy student improves imagenet classification. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10684–10695. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPR42600.2020.01070.
- [24] S. Zagoruyko and N. Komodakis. Wide residual networks. In E. R. H. Richard C. Wilson and W. A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, 9 2016. ISBN 1-901725-59-6. doi: 10.5244/C.30.87.
- [25] B. Zoph, G. Ghiasi, T. Lin, Y. Cui, H. Liu, E. D. Cubuk, and Q. Le. Rethinking pre-training and self-training. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/27e9661e033a73a6ad8cefcd965c54d-Abstract.html>.