

Semi-Supervised Learning using Siamese Networks

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. Neural networks have been successfully used as classification models yielding state-of-the-art results when trained on a large number of labeled samples. These models, however, are more difficult to train successfully for semi-supervised problems where small amounts of labeled instances are available along with a large number of unlabeled instances. This work explores a new training method for semi-supervised learning that is based on similarity function learning using a Siamese network to obtain a suitable embedding. The learned representations are discriminative in Euclidean space, and hence can be used for labeling unlabeled instances using a nearest-neighbor classifier. Confident predictions of unlabeled instances are used as true labels for retraining the Siamese network on the expanded training set. This process is applied iteratively. We perform an empirical study of this iterative self-training algorithm. For improving unlabeled predictions, local learning with global consistency [22] is also evaluated.

Keywords: Semi-supervised learning, Siamese networks, Triplet loss, LLGC.

1 Introduction

The modern world generates vast amounts of data and provides many opportunities to exploit it. However, frequently this data is complex, noisy, and lacks obvious structure. Therefore, explicit modeling of, for example, its distribution is too challenging for a human agent. On the other hand, a human can specify an explicit procedure, i.e., an algorithm, for how to construct such a model. Machine learning (ML) is concerned with algorithms that enable computers to learn from data in this way, especially algorithms for prediction. Many ML algorithms need labeled data for such a task, but it is common that fewer labeled data are available than unlabeled ones. Manual labeling is costly and time-consuming. Hence, there is an ever-growing need for ML methods to work with a limited amount of labeled data and also make efficient use of the side information available from unlabeled data. Algorithms designed to do so are known as semi-supervised learning algorithms.

Supervised learning algorithms employ labeled data to predict class labels for unlabeled examples accurately. Unsupervised learning algorithms search for

structure in data, which can then be used as a heuristic to infer labels for these examples, on the basis of assumptions about the structure of data. Semi-Supervised learning (SSL) algorithms lie somewhere between supervised and unsupervised learning. SSL methods are designed to work with labeled $L = \{(x_1, y_1), (x_2, y_2), \dots, (x_{|L|}, y_{|L|})\}$ and unlabeled instances $U = \{x'_1, x'_2, \dots, x'_{|U|}\}$, where X and Y relate to an input space and output space, $x_i, x'_j \in X (i = 1, 2, \dots, |L|, j = 1, 2, \dots, |U|)$ are examples and $y_i \in Y$ are labels of x_i and $Y = \{1, 2, 3, \dots, c\}$, c being the number of classes. Usually, these methods assume a much smaller number of labeled instances than unlabeled ones i.e., $|L| \ll |U|$, because unlabeled instances are more useful when we have a few labeled instances. SSL has proven to be useful especially when we are dealing with anti-causal or confounded problems [15].

Without making any assumptions on how the inputs and outputs are related it is impossible to justify semi-supervised learning as a principled approach [4]. Like the authors in that paper, we make the same three assumptions:

1. If two points x_1, x_2 are close in a high-density region, then their corresponding outputs y_1, y_2 should also be close.
2. If points are in the same structure (referred to as cluster or manifold), they are likely to be of the same class.
3. The decision boundary between classes should lie in a low-density region of input space.

In this work, we will consider a new training method designed to be used with deep neural networks in the semi-supervised learning setting. Instead of the usual approach of learning a direct classification model based on cross-entropy loss, we will use the labeled examples for learning a similarity function between instances, such that instances of the same class are considered similar and those instances belonging to different classes are considered dissimilar. Under this similarity function, which is parameterized by a neural network, the features (embeddings) of labeled examples will be grouped together according to the class labels, in Euclidean space. In addition, we will use these learned embeddings to assign class labels to unlabeled examples. We do this using a simple nearest-neighbor classifier. Following that, confident predictions for unlabeled instances are added to the labeled examples for retraining of the neural network iteratively. In this way, we are able to achieve significant performance improvements over supervised-only training.

2 Related Work

Semi-supervised learning has been under study since the 1970s [12]. Expectation-Maximization (EM) [14] works by labeling unlabeled instances with the current supervised model's best prediction in an iterative fashion (self-learning), thereby providing more training instances for the supervised learning algorithm. Co-training [1] is a similar approach, where two models are trained on two separate subsets of the data features. Confident predictions from one model are then used

as labeled data for the other model. Co-EM [2] combines co-training with EM and achieved better results than either of them. Another, graph-based SSL method, LLGC (Local Learning with Global Consistency) [22], works by propagating labels from labeled to unlabeled instances until labels are stable, maintaining local and global consistency.

There is a substantial amount of literature available on SSL techniques using deep neural network based on autoencoders [16,11], generative adversarial networks (GAN) [18,6,20] and based on regularization [9,17,13]. The Pseudolabel [10] approach is a deep learning version of self-learning with an extra loss from regularization and the reconstruction of a denoising autoencoder.

Our method builds on work investigating similarity metric learning using neural networks. [5] used a network with the contrastive loss for face verification in a supervised fashion. [19] suggested network training to be based on triplets of examples. This work was extended to the semi-supervised paradigm [21] for the image classification task. [7] tries to minimize the sum of cross-entropy and ratio loss between class indicators (sampled from labeled examples for each class) and the intra-class distances of instances calculated based on embeddings.

We train our network based on triplets of images and use the triplet margin loss [19]. We found this to perform better than the contrastive loss or the ratio loss in our experiments, while the network is trained in a self-learning fashion. For improving intermediate predictions, we use LLGC [22] in order to get better labels for unlabeled instances in subsequent iterations. Although triplet networks and LLGC are not new, this is the first attempt, to our knowledge, of combining these two approaches for semi-supervised learning.

3 Siamese Networks

Siamese networks [3] are neural networks that are particularly efficient when we have a large number of classes and a few labeled instances per class. Siamese networks can be thought of multiple networks with identical copies of the same function, with the same weights. They can be employed for training a similarity function given labeled data. Fig. 1 shows a simple network architecture based on convolutional (CONV) and max-pooling (MP) layers. An input example is passed to the network for computing the embeddings. Different losses are used for training Siamese networks, such as contrastive loss, margin-based loss, and triplet loss. Network parameters are updated according to the loss calculated on embeddings.

3.1 Triplet Loss

The triplet loss [19] has been used for face recognition. A triplet’s anchor example a , positive example p , and negative example n are provided as a training example to the network for getting corresponding embeddings. During optimisation of the network parameters, we draw all possible triplets from labeled examples based on class labels. For each mini-batch used in stochastic gradient descent, all valid




Fig. 1: Network Architecture

triplets(i, j, k) are selected where $labels[i] = labels[j], i \neq j$ and $labels[i] \neq labels[k]$. Then the loss is calculated according to the following equation using the Euclidean distance $d(.,.)$ between the embedded examples:

$$\mathcal{L} = \max(d(a, p) - d(a, n) + m, 0) \quad (1)$$

where m is the so-called "margin" and constitutes a hyperparameter.

As illustrated in Fig. 2, the triplet loss attempts to push away the embedded negative example n from the embedded anchor example a based on a given margin m and the given positive example p . Depending on the location of the negative example with respect to the anchor and the positive example, it is possible to distinguish between hard negative examples, semi-hard negative examples, and easy negative examples. The latter are effectively ignored during optimisation because they yield the value zero for the loss.




Fig. 2: Triplet loss

3.2 Self-learning using Siamese networks

In the first iteration of our semi-supervised learning approach, to be able to label (some of) the unlabeled examples instances, the Siamese network is trained

on labeled examples only, using triplet loss. Then the standard nearest neighbor classifier is used to predict labels for the unlabeled examples and a fixed percentage p of unlabeled examples is chosen based on their distance to the labeled instances and added to the set of labeled examples for the next iteration. Throughout, embedded data is used to calculate distances. For more details see the pseudo-code in Listing 1.

Algorithm 1 Proposed approach based on Siamese self-training

```

1: Input: Labeled examples  $(x_L, y_L)$ , Unlabeled examples  $x_U$ , number of meta-
   iterations  $i$  and selection percentage  $p$ 
2: for 1 to  $i$  do
3:    $train\_siamesenetwork(x_L, y_L)$ 
4:    $embed_U = siamesenetwork(x_U)$ 
5:    $embed_L = siamesenetwork(x_L)$ 
6:    $labels_U, dist_U = KNN(embed_U, embed_L, y_L)$ 
7:    $sorted\_dist_U, sorted\_labels_U = sort(dist_U, labels_U)$ 
8:    $x_{new}, y_{new} = select\_top(sorted\_dist_U, sorted\_labels_U, p)$ 
9:    $x_L, y_L = concat((x_L, y_L), (x_{new}, y_{new}))$ 
10:   $x_U = delete\_from(x_U, x_{new})$ 
11: end for

```

4 Local Learning with Global Consistency (LLGC)

We also investigate local learning with global consistency [22] in addition to the nearest-neighbor classifier. LLGC works by propagating label information to the neighbors of an example. The goal of LLGC is to predict labels for unlabeled instances. The algorithm initializes a matrix $Y_{n \times c}$ to represent label information, where $Y_{ij} = 1$ if example i is labeled as j , and otherwise $Y_{ij} = 0$. We implement a little variation here for the unlabeled examples: instead of using $Y_{ij} = 0$ for all j when i is unlabeled, we use predicted labels obtained with the nearest-neighbour classifier after training the Siamese network.

LLGC is based on calculating an adjacency matrix. This adjacency matrix is then used to establish a matrix S that is applied to update the label probabilities for the unlabeled examples. The adjacency matrix is calculated using Eq. 2 by employing embeddings $f(x_i)$ and $f(x_j)$ for each pair of two examples x_i and x_j , obtained from the Siamese network. The parameter σ is a hyper-parameter.

$$W_{ij} = \begin{cases} e^{-\sigma \times |f(x_i) - f(x_j)|^2}, & \text{if } i \neq j \\ 0 & \text{if } i = j. \end{cases} \quad (2)$$

The matrix S is computed as:

$$S = D^{-1/2} \times W \times D^{-1/2} \quad (3)$$

where D is a diagonal matrix: $D_i = \sum_{j=1}^n W_{ij}$. The initial matrix of label probabilities is set to $F(0) = Y$, and the probabilities are updated by:

$$F(t+1) = S.F(t) \times \alpha + (1 - \alpha) \times Y \quad (4)$$

where $\alpha \in [0, 1)$ is a hyper-parameter for controlling the propagation of label information. The above operation is repeated till convergence. Finally, labels for the unlabeled instances are calculated as:

$$y_i = \operatorname{argmax}_{j \leq c} F_{ij} \quad (5)$$

For efficiently using unlabeled instances, the Siamese network is first trained on labeled examples only, using triplet loss. Then the nearest-neighbor classifier is used to predict labels for unlabeled examples. Then, following that, labeled and unlabeled embeddings along with labels are passed to LLGC. After a certain number of iterations of LLGC, a fixed percentage p of unlabeled examples are chosen based on their LLGC score and added to the labeled examples for the next iteration. For more details see the pseudo-code in Listing 2.

Algorithm 2 Proposed approach based on LLGC self-training

```

1: Input: Labeled examples  $(x_L, y_L)$ , Unlabeled examples  $x_U$ , number of meta-
   iterations  $i$ , selection percentage  $p$ ,  $\alpha$  and  $\sigma$  parameters for LLGC.
2: for 1 to  $i$  do
3:    $train\_siamesenetwork(x_L, y_L)$ 
4:    $embed_U = siamesenetwork(x_U)$ 
5:    $embed_L = siamesenetwork(x_L)$ 
6:    $labels_U = KNN(embed_U, embed_L, y_L)$ 
7:    $LLGC\_labels, LLGC\_score = LLGC(embed_L, embed_U, [y_L, labels_U], \sigma, \alpha)$ 
8:    $labels_U = LLGC\_labels[len(x_L) :]$ 
9:    $x_{new}, y_{new} = select\_top(LLGC\_score, p, x_U, labels_U)$ 
10:   $x_L, y_L = concat((x_L, y_L), (x_{new}, y_{new}))$ 
11:   $x_U = delete\_from(x_U, x_{new})$ 
12: end for

```

5 Experiments

We consider four standard image classification problems for our evaluation. For all experiments, a small subset of labeled examples was chosen according to standard semi-supervised learning practice, with a balanced number of examples from each class, and the rest were considered as unlabeled. Final accuracy was calculated on the standard test split for each dataset. No data augmentation was applied to the training sets. Siamese networks were trained using triplet loss with margin $m = 0.3$ for all datasets.

A simple convolutional network architecture was chosen for each dataset to ensure performance achieved was due to the proposed method and not the network architecture. For more details about the network architectures, see Table 1. Layer descriptions use (feature-maps, kernel-size, stride, padding) for convolutional layers and (pool-size, stride) for pooling layers. The simple model is used for MNIST, Fashion MNIST, and SVHN, and produces 16-dimensional embeddings, while the CIFAR-10 model produces 64-dimensional embeddings. We trained the networks using mini-batch sizes 50, 100, and 200. We found that batch size 50 was insufficient and 200 did not yield significant improvements compared to batch size 100. Batch size = 100 is used for all experiments, with Adam [8] as the optimizer for updating network parameters for 200 epochs. Our proposed approaches Siamese self-training (Algorithm 1) and LLGC self-training (Algorithm 2) respectively were run for 25 meta-iterations. For LLGC, $\alpha = 0.99$ is used in all experiments, while σ is optimized for each dataset. The final test accuracy is computed using a k-NN classifier with $k = 1$ for simplicity. Our results were averaged over 3 random runs, using a different random initialization of the Siamese network parameters for each run and random selection of initially labeled examples except SVHN. We set a baseline by (a) training the network on the small number of the labeled instances only, and by (b) using all the labeled instances. These two baselines should provide good empirical lower and upper bounds for the semi-supervised error rates.

Table 1: Network Model

Simple(#parameters=163908)	CIFAR-10(#parameters=693792)
INPUT	INPUT
Conv-Relu(32,7,1,2)	Conv-Relu-BN(192,5,1,2)
Max-Pooling(2,2)	Conv-Relu-BN(160,1,1,2)
Conv-Relu(64,5,1,2)	Conv-Relu-BN(96,1,1,2)
Max-Pooling(2,2)	Max-Pooling(3,2)
Conv-Relu(128,3,1,2)	Conv-Relu-BN(96,5,1,2)
Max-Pooling(2,2)	Conv-Relu-BN(192,1,1,2)
Conv-Relu(256,1,1,2)	Conv-Relu-BN(192,1,1,2)
Max-Pooling(2,2)	Max-Pooling(3,2)
Conv(4,1,1,2)	Conv-Relu-BN(192,3,1,2)
Flatten()	Conv-Relu-BN(64,1,1,2)
	Avg-Pooling(8,1)

We now consider the datasets used in our experiments. The MNIST dataset consists of gray-scale 28 by 28 images of handwritten digits. We select only 100 instances (10 from each class) as labeled instances initially. We apply our algorithms with a selection percentage $p = 10\%$ and the LLGC-based method with $\sigma = 1.8$. Table 2 shows noticeable improvements over the supervised-only

approach when compared with the proposed semi-supervised approaches, when using the same number of labeled examples.

Table 2: MNIST Test error %.

# labels	100-Labeled	All (60000)
Supervised-only	9.73 ± 0.74	0.6 ± 0.04
Siamese self-training	3.24 ± 0.32	–
LLGC self-training	3.50 ± 0.14	–

The Fashion MNIST dataset consists of 28 by 28 gray-scale images showing fashion items. 100 instances are considered as labeled initially. Again, we use selection percentage $p = 10\%$ and $\sigma = 3.2$. Table 3 again shows noticeable improvement over the supervised-only approach when compared with the proposed semi-supervised approaches, when using the same amount of labeled data.

Table 3: Fashion MNIST Test error %.

# labels	100-Labeled	All (60000)
Supervised-only	26.72 ± 1.23	9.66 ± 0.10
Siamese self-training	23.33 ± 0.43	–
LLGC self-training	23.23 ± 0.67	–

SVHN comprises 32x32 RGB images of house numbers, taken from the Street View House Numbers dataset. Each image can have multiple digits, but only the digit in the center is considered for prediction. The proposed approaches are evaluated using 1000 labeled instances initially, with selection percentage $p = 5\%$, and $\sigma = 2.4$. Table 4 shows noticeable improvement over the supervised-only approach when compared to the proposed approaches when 1000 labeled examples are used. Interestingly, purely Siamese self-training again performs better than LLGC self-training in this case.

The CIFAR-10 dataset contains 32 by 32 RGB images of ten classes. The proposed semi-supervised approaches are evaluated using 4000 labeled instances initially, with selection percentage $p = 5\%$, and $\sigma = 2.4$. Table 5 shows little improvement over the supervised-only approach when compared to the proposed semi-supervised approaches. Siamese self-training performs better than LLGC self-training.

Figures 3, 4, 5 and 6 show a detailed comparison between Siamese self-training and LLGC self-training across three different runs of all four datasets; MNIST, Fashion MNIST, SVHN, and CIFAR-10. The accuracy curves show

Table 4: SVHN Test error %.

# labels	1000-Labeled	All (73275)
Supervised-only	30.33 ± 1.55	12.26 ± 0.52
Siamese self-training	20.09 ± 3.22	–
LLGC self-training	27.23 ± 0.99	–

Table 5: CIFAR-10 Test error %.

# labels	4000-Labeled	All (50000)
Supervised-only	40.87 ± 0.56	21.51 ± 0.88
Siamese self-training	36.56 ± 0.74	–
LLGC self-training	40.06 ± 0.62	–

definite improvement with respect to the supervised-only version on all datasets using Siamese self-training as well as LLGC self-training. However, CIFAR-10 and SVHN seem to get low or negligible additional improvement from LLGC self-training compared to Siamese self-training only.

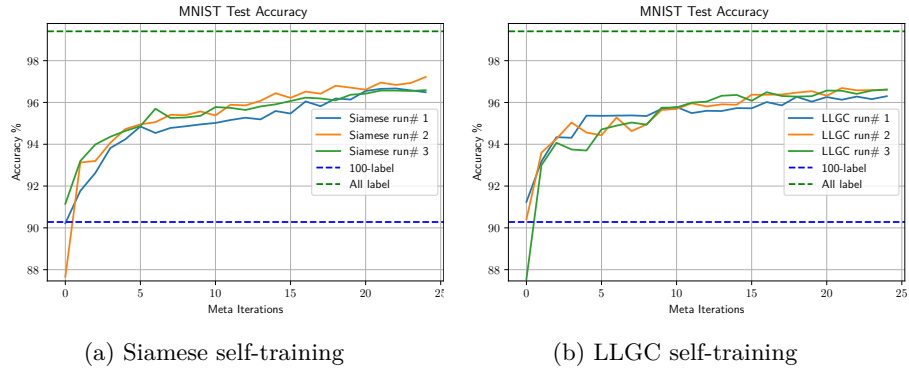


Fig. 3: MNIST-100 Comparison of Siamese self-training vs. LLGC self-training.

We also tried to visualize the quality of embeddings learned using the proposed method. We trained an additional model by slightly modifying the simple model 1. In order to get a 2-dimensional embedding, two feature-maps are used instead of 4 in the last convolutional layer, followed by average-pooling(2,2) before the final flattening layer. For this purpose, we considered MNIST. Figure 7 (a) depicts the embeddings for test instances marked in color according to their true class after random initialization of the network. Figure 7 (b) depicts the em-

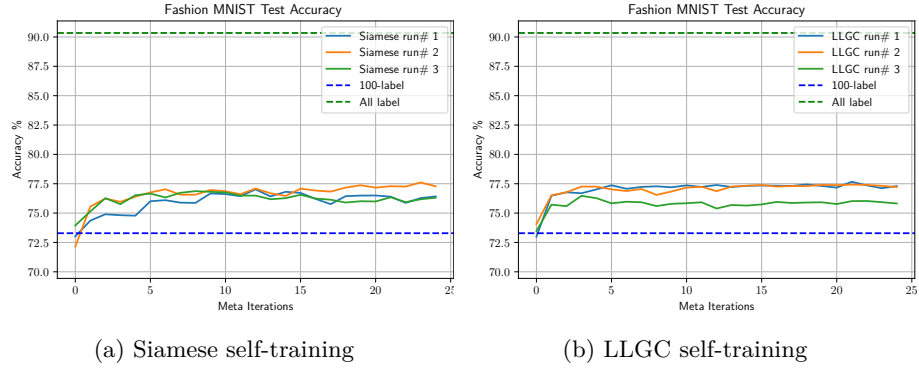


Fig. 4: Fashion MNIST-100 Comparison of Siamese self-training vs. LLGC self-training.

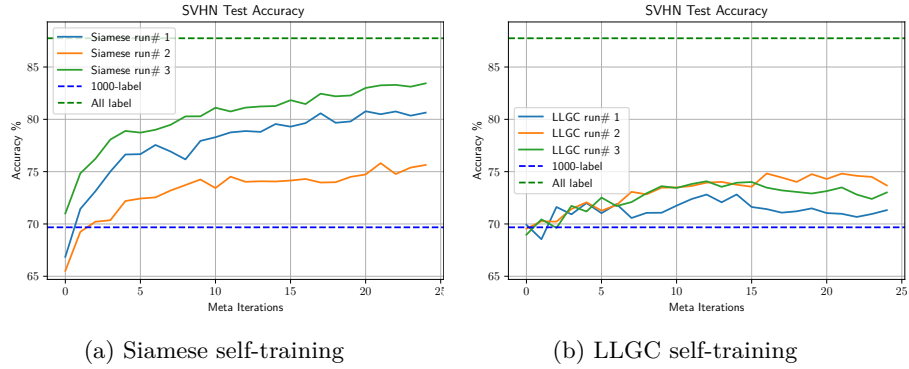


Fig. 5: SVHN-1000 Comparison of Siamese self-training vs. LLGC self-training.

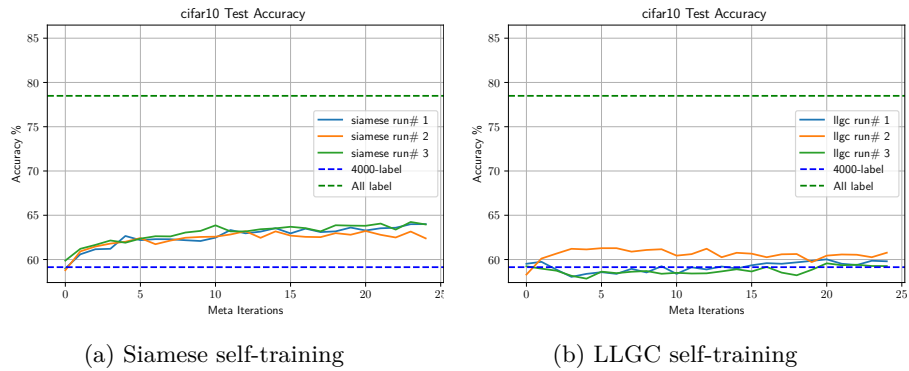


Fig. 6: CIFAR10-4000 Comparison of Siamese self-training vs. LLGC self-training.

beddings for test instances after training the Siamese network with only the 100 labeled MNIST instances. It can be seen that the 10000 test examples' embeddings form clusters in Euclidean space after training of the network according to the class labels; test examples' embeddings are largely scattered randomly throughout the 2D space before the network is trained.




Fig. 7: MNIST-100: visualisation of 2-dimensional embeddings

6 Conclusion

In this work, we have shown how neural networks can be used to learn in a semi-supervised setting using small sets of labeled data by replacing the classification objective with an objective for learning a similarity function. This objective is compliant with standard techniques of training the deep neural network and requires no modification of the embedding model. For improving the intermediate prediction of unlabeled instances, we evaluated LLGC, but this yielded little additional benefit compared to k-NN classification alone. Using the method in this work, we were able to achieve significant improvement compared to supervised learning only on MNIST, Fashion MNIST and SVHN, when training on a small subset of labeled examples, but obtained little improvement on CIFAR-10. We speculate that instead of a fixed selection of unlabeled instances from LLGC's predictions, a threshold-based selection based on the LLGC score will be more beneficial for subsequent iterations of our meta-algorithm. Also, a more robust convolutional model may help the network in learning distinctive embeddings and achieving state-of-the-art results for the semi-supervised setting.

References

1. Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: Proceedings of the eleventh annual conference on Computational learning theory. pp. 92–100. ACM (1998)
2. Brefeld, U., Scheffer, T.: Co-em support vector learning. In: Proceedings of the twenty-first international conference on Machine learning. p. 16. ACM (2004)

3. Bromley, J., Bentz, J., Bottou, L., Guyon, I., LeCun, Y., Moore, C., Sackinger, E., Shah, R.: Signature verification using a “siamese” time delay neural network. *Int. J. Pattern Recognit. Artif Intell* **7** (1993)
4. Chapelle, O., Schölkopf, B., Zien, A.: *Semi-supervised learning*, ser. Adaptive computation and machine learning. Cambridge, MA: The MIT Press (2006)
5. Chopra, S., Hadsell, R., LeCun, Y., et al.: Learning a similarity metric discriminatively, with application to face verification. In: *CVPR* (1). pp. 539–546 (2005)
6. Dai, Z., Yang, Z., Yang, F., Cohen, W.W., Salakhutdinov, R.R.: Good semi-supervised learning that requires a bad gan. In: *Advances in Neural Information Processing Systems*. pp. 6513–6523 (2017)
7. Hoffer, E., Ailon, N.: Semi-supervised deep learning by metric embedding. *arXiv preprint arXiv:1611.01449* (2016)
8. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
9. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242* (2016)
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. Maaløe, L., Sønderby, C.K., Sønderby, S.K., Winther, O.: Auxiliary deep generative models. *arXiv preprint arXiv:1602.05473* (2016)
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. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *arXiv preprint arXiv:1704.03976* (2017)
14. Nigam, K., McCallum, A., Mitchell, T.: Semi-supervised text classification using em. *Semi-Supervised Learning* pp. 33–56 (2006)
15. Peters, J., Janzing, D., Schölkopf, B.: *Elements of causal inference: foundations and learning algorithms*. MIT Press (2017)
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. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In: *Advances in Neural Information Processing Systems*. pp. 1163–1171 (2016)
18. 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)
19. 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)
20. Wei, X., Gong, B., Liu, Z., Lu, W., Wang, L.: Improving the improved training of wasserstein gans: A consistency term and its dual effect. *arXiv preprint arXiv:1803.01541* (2018)
21. Weston, J., Ratle, F., Mobahi, H., Collobert, R.: Deep learning via semi-supervised embedding. In: *Neural Networks: Tricks of the Trade*, pp. 639–655. Springer (2012)
22. Zhou, D., Bousquet, O., Lal, T.N., Weston, J., Schölkopf, B.: Learning with local and global consistency. In: *Advances in neural information processing systems*. pp. 321–328 (2004)