LEAP: LEARNING EMBEDDINGS FOR ADAPTIVE PACE

Anonymous authors
Paper under double-blind review

ABSTRACT

The parameterization of mini-batches for training Deep Neural Networks (DNN) is a non-trivial problem. In this paper, we propose a Self-Paced Learning (SPL)-fused Deep Metric Learning (DML) framework, which we call Learning Embeddings for Adaptive Pace (LEAP). Our method parameterizes mini-batches dynamically based on the *easiness* and *true diverseness* of the sample within a salient feature representation space. In LEAP, we train an *embedding* Convolutional Neural Network (CNN) to learn an expressive representation space by adaptive density discrimination using the Magnet Loss. The *student* CNN classifier dynamically selects samples to form a mini-batch based on the *easiness* from cross-entropy losses and *true diverseness* of examples from the representation space sculpted by the *embedding* CNN. We evaluate LEAP using deep CNN architectures for the task of supervised image classification on MNIST, FashionMNIST and CIFAR-10. We show that the LEAP framework can achieve a higher convergence w.r.t. the number of mini-batch updates required to achieve a comparable or better test performance on each of the datasets. Our framework is implemented in PyTorch and will be released as open-source on GitHub following review.

1 INTRODUCTION

Stochastic non-convex optimization algorithms, which are driven by gradient information on mini-batches of the dataset, are the most common way to train Deep Neural Networks (DNNs). Despite recent advances in DNNs making it possible to achieve high performance on several tasks across many application domains such as computer vision and natural language processing, their architectures have a large number of parameters, which raises the importance of data scheduling for training DNNs. The standard method to train DNNs is stochastic gradient descent (SGD) which employs backpropagation to compute gradients. It typically relies on fixed-size mini-batches of random samples drawn from a finite dataset. However, the contribution of each sample during model training varies across training iterations and configurations of the model's parameters (Lapedriza et al., 2013). Previous studies on Curriculum Learning (Bengio et al., 2009, CL) show that organizing training samples based on the ascending order of difficulty can favour model training. However, in CL, the curriculum remains fixed over the iterations and is determined without any knowledge or introspection of the model's learning. Self-Paced Learning (Kumar et al., 2010) presents a method for dynamically generating a curriculum by biasing samples based on their *easiness* under the current model parameters. This can lead to a highly imbalanced selection of samples, i.e. very few instances of some classes are chosen, which negatively affects the training process due to overfitting. Loshchilov & Hutter (2015) propose a simple batch selection strategy based on the loss values of training data for speeding up neural network training. However, their results are limited and the approach is time-consuming, as it achieves high performance on MNIST, but fails on CIFAR-10. Their work reveals that selecting the examples to present to a DNN is non-trivial yet the strategy of uniformly sampling the training data set is not necessarily the optimal choice.

Jiang et al. (2014b) show that partitioning the data into groups to use *diversity* and *easiness* in their Self-Paced Learning with Diversity (SPLD) framework, can have substantial effect on training. Rather than constraining the model to limited groups and areas, we can spread the sample selection as wide as possible to obtain diverse samples of similar easiness. However, their use of K-Means and Spectral Clustering to partition the data into groups can lead to sub-optimum clustering results when learning non-linear feature representations. Therefore, learning an appropriate metric by which to capture similarity among arbitrary groups of data is of great practical importance. Deep

Metric Learning (DML) approaches have recently attracted considerable attention and have been the focus of numerous studies (Bell & Bala (2015); Schroff et al. (2015)). The most common methods are supervised, in which a feature space in which distance corresponds to class similarity is obtained. The Magnet Loss (Rippel et al., 2015) presents state-of-the-art performance on fine-grained classification tasks. Song et al. (2016a) show that it achieves state-of-the-art on clustering and retrieval tasks.

To summarize, this paper makes a number of contributions toward improving mini-batch formation:

- We propose a general sample selection framework called Learning Embeddings for Adaptive Pace (LEAP) that is independent of model architecture or objective, and learns when to introduce certain samples to the DNN during training.
- We exploit a new type of knowledge —similar instance-level samples for knowledge transfer through DML from an *embedding* network for self-paced sample selection.

2 Relevant Work

2.1 Learning Small and Easy

The perspective of "starting small and easy" for structuring the learning regimen of neural networks dates back decades to Elman (1993). Recent studies show that selecting a subset of *good* samples for training a classifier can lead to better results than using all the samples (Lee & Grauman, 2011; Lapedriza et al., 2013). Pioneering work in this direction is Curriculum Learning (Bengio et al., 2009), which introduced a heuristic measure of easiness to determine the selection of samples from the training data. By comparison, SPL (Kumar et al., 2010) quantifies the easiness by the current sample loss. The training instances with loss values larger than a threshold, λ, are neglected during training and λ dynamically increases in the training process to include more complex samples, until all training instances are considered. This theory has been widely applied to various problems, including dictionary learning for image classification (Tang et al., 2012), object detection (Sangineto et al., 2016), multimedia event detection (Jiang et al., 2014a), long-term tracking (III & Ramanan, 2013), visual tracking (Huang et al., 2017) and medical imaging analysis (Li et al., 2017). In SPLD Jiang et al. (2014b), training data are pre-clustered in order to balance the selection of the easiest samples with a sufficient inter-cluster diversity. However, the clusters and the feature space are fixed: they do not depend on the current self-paced training iteration. Adaptation of this method to a deep-learning scenario, where the feature space changes during learning, is non-trivial. Our self-paced sample selection framework aims at a similar goal but the diversity of samples is obtained with a DML approach to adaptively sculpt a representation space by autonomously identifying and respecting intra-class variation and inter-class similarity.

2.2 Learning Representations by Metric Embedding

Deep metric learning has gained much popularity in recent years, along with the success of deep learning. The objective of DML is to learn a distance metric consistent with a given set of constraints, usually minimizing the distances between pairs of data points from the same class and maximizing the distances between pairs of data points from different classes. DML approaches have shown promising results on various tasks, such as semantic segmentation (Fathi et al., 2017), visual product search (Kiapour et al., 2015), face recognition (Schroff et al., 2015), feature matching (Choy et al., 2016), fine-grained image classification (Zhang et al., 2015), zero-shot learning (Frome et al., 2013) and collaborative filtering (Hsieh et al., 2017). DML can also be used for challenging, extreme classification settings, where the number of classes is very large and the number of examples per class becomes scarce. Most of the current methods define the loss in terms of pairs (Song et al., 2016b; Ustinova & Lempitsky, 2016), triplets (Schroff et al., 2015; Wang et al., 2017) or n-pair tuples (Sohn, 2016) inside the training mini-batch. These methods require a separate data preparation stage which has very expensive time and space cost. Also, they do not take the global structure of the embedding space into consideration, which can result in reduced clustering. An alternative is the Magnet Loss (Rippel et al., 2015) and DML via Facility Location (Song et al., 2016a) which do not require the training data to be preprocessed in rigid paired format and are aware of the global structure of the embedding space. Our work employs the Magnet loss to learn a representation

space, where we compute centroids on the raw features and then update the learnt representation continuously.

To our knowledge, the concept of employing DML for SPL-based DNN training has not been investigated. Effectively, an end-to-end DML can be constructed to be a feature extractor using a deep CNN which can learn to sculpt an expressive representation space by metric embedding. We can use this feature representation space which maintains intra-class variations and inter-class similarity to select samples based on the *true diverseness* and *easiness* for the CNN we want to learn. Our architecture combines the strength of adaptive sampling with that of mini-batch online learning and adaptive representation learning to formulate a representative self-paced strategy in an end-to-end DNN training protocol.

3 Learning Embeddings for Adaptive Pace (LEAP)

The Learning Embeddings for Adaptive Pace (LEAP) framework consists of a dual DNN setup. An *embedding* DNN learns a salient representation space, then transfers its knowledge to the self-paced selection strategy to train the second DNN, called the *student*. Thus, an *embedding* CNN is trained alongside the *student* CNN we want to train (Figure 1). For now, but without loss of generality, we focus on training deep Convolutional Neural Networks (CNNs) for the task of supervised image classification. The embedding network is trained to learn a representation space using the Magnet Loss, because a more salient representation, in theory, enables improved classification performance. In this framework, we want to form mini-batches using the *easiness* and *true diverseness* as sample importance priors for the selection of training samples. Given that we are learning the representation space adaptively as training progresses, this has negligible computational cost compared to the actual training of the student CNN. We evaluate our framework on image classification datasets such as MNIST, FashionMNIST (Xiao et al., 2017) and CIFAR-10, allowing for a comparison of our results with other relevant sample selection strategies.

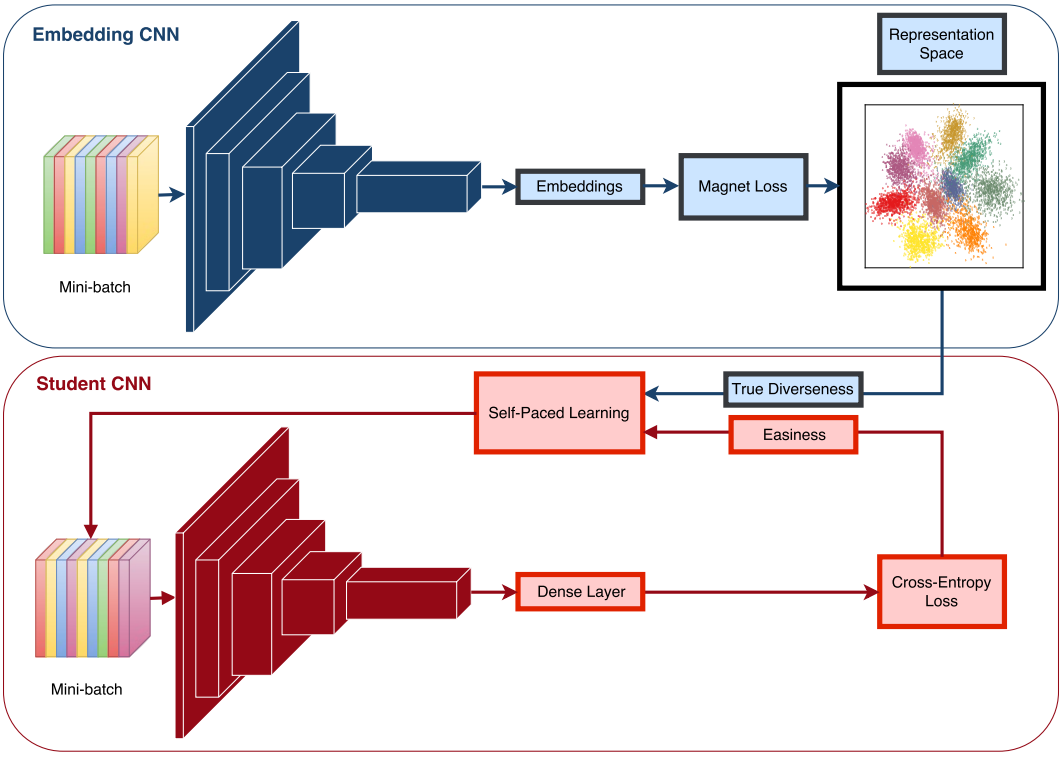

Figure 1: The LEAP framework, consisting of an embedding CNN that learns a representation for the student CNN to create a self-paced strategy using *easiness* and *true diverseness* as sample importance priors.

We adopt the Magnet Loss approach which is based on learning the distribution of distances for each sample, from k clusters assigned for each class. It then uses an intermediate K-Means clustering, to reposition the different assigned clusters. This proved to allow better accuracy than both margin-based Triplet Loss and Softmax regression. The samples are fed into the embedding CNN in mini-batches of random samples from the training data. The aim of LEAP can be formally described as follows. Assuming that the training samples $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_n) \in \mathbb{R}^{\mathrm{mxn}}$ are grouped into k clusters in the learned representation space through DML: $\mathbf{X}^{(1)}, \ldots, \mathbf{X}^{(k)}$, where columns of $\mathbf{X}^{(j)} \in \mathbb{R}^{\mathrm{mxn}_j}$ correspond to the samples in the jth cluster, n_j is the number of samples in each cluster and $\sum_{j=1}^{k} n_j = n$. The weight vector is denoted accordingly as $\mathbf{v} = [\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(k)}]$, where $\mathbf{v}^{(j)} = (\mathbf{v}_1^j, \ldots, \mathbf{v}_{n_j}^{(j)})^T \in [0,1]^{n_j}$. Similar to the conventional SPL, we need to assign non-zero weights v to easy samples, and increase the diversity by dispersing non-zero elements across possibly more groups' $v^{(j)}$. This leads to the optimization model similar to the one presented in SPLD, where λ, γ are the two pacing parameters for sampling based on *easiness* and the *true diverseness* as sample importance priors, respectively. The student CNN receives a diverse cluster of samples, the up-to-date model parameters $\boldsymbol{\theta}$, λ, γ and the outputs of v of $\min_v \mathbb{E}(\boldsymbol{\theta}, v; \lambda, \gamma)$ for extracting the global optimum of this optimization problem.

4 EXPERIMENTS

All experiments were conducted using the PyTorch framework, while leveraging containerized multi-GPU training on NVIDIA P100 Pascal GPUs through Docker. We compared our LEAP framework against the original SPLD algorithm and Random sampling on MNIST, FashionMNIST and CIFAR-10. The λ and γ pace parameters are kept consistent between the SPL strategy in LEAP and the original SPLD algorithm to ensure a fair evaluation. In our experiments, we mainly compare convergence in terms of number of mini-batches required to achieve a comparable state-of-the-art test performance. We also visualize the original high-dimensional representations using t-SNE (van der Maaten & Hinton, 2008), where the different colours correspond to different classes and the values to density estimates. The following sections discuss the experimental setups and results in more detail.

4.1 LEAP WITH LENET + LENET ON MNIST

In the experiments for MNIST, we extract the feature embeddings from a LeNet (LeCun et al., 1998), as the embedding CNN, and learn a representation space using the Magnet Loss. The fully-connected layer of the LeNet is replaced with an embedding layer for compacting the distribution of the learned features for feature similarity comparison using the Magnet Loss. The student CNN (classifier) was also a LeNet which we then trained with our LEAP framework. The embedding CNN was trained with randomly sampled mini-batches of size 64 and optimized using Adam with a learning rate of 0.0001. The results in Figure 2 show that the test performance is comparable to that of Random sampling and shows better convergence on the standard SPLD algorithm. The learned representation space of the MNIST dataset using LeNet is presented in Figure 3b and the convergence in Figure 3a. The MNIST experiments were primarily carried out to show that the LEAP framework can be deployed as an end-to-end DNN training protocol.

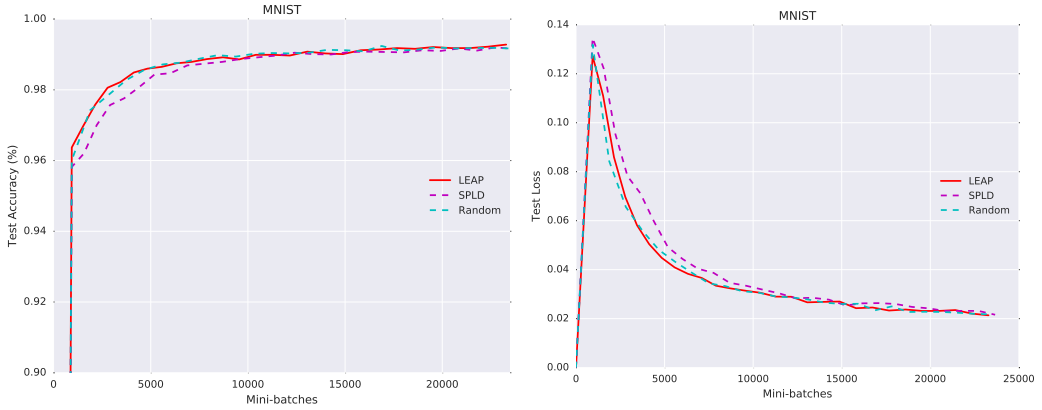

Figure 2: Averaged results of 5 independent runs on MNIST. The solid line represents our LEAP framework.

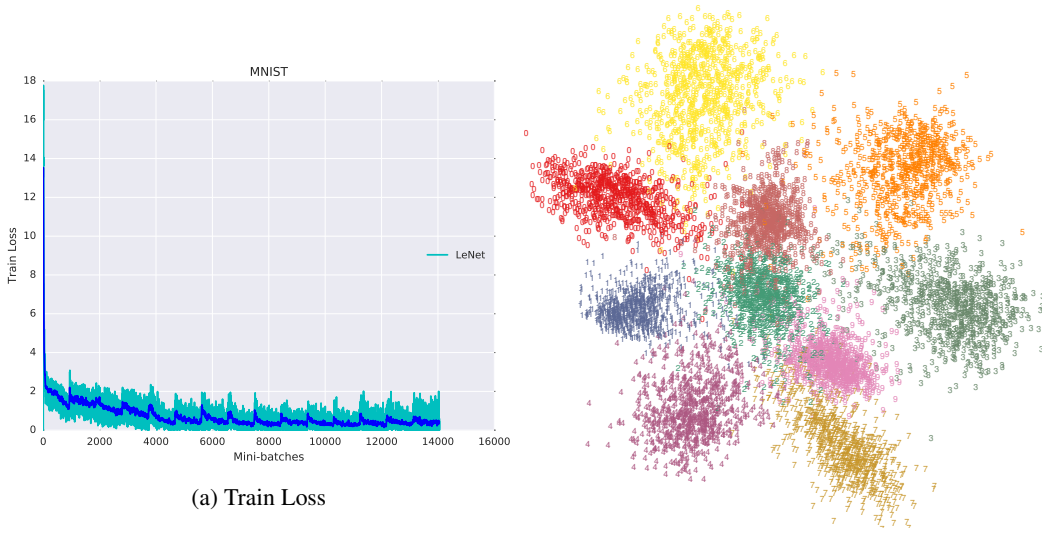

(a) Train Loss

(b) MNIST feature representation space

Figure 3: (a) The train loss with the moving average and (b) the feature representation space learned by LeNet (embedding CNN) using the LEAP framework on MNIST.

4.2 LEAP WITH LENET + RESNET-18 ON FASHIONMNIST

The experiments for FashionMNIST had a similar setup to MNIST, however we use a ResNet-18 (He et al., 2016) for the classifier. The embedding CNN remains the same LeNet for feature extraction, and is trained using an identical setup to the MNIST experiments. The classifier is trained using SGD with momentum of 0.9, weight decay of 0.0005 and a learning rate of 0.001. The FashionMNIST dataset is considered a direct drop-in replacement for the original MNIST dataset, with a training set of 60000 examples and a test set of 10000 examples. Each example is a 28×28 grayscale image, associated with a label from 10 classes. We performed data augmentation on the training set with normalization, random horizontal flip, random vertical flip, random translation, random crop and random rotation. In comparison to SPLD and Random sampling, the results in Figure 4 reveal that LEAP converges to a higher test accuracy with a fewer number of mini-batch updates before saturating. Figure 5b depicts a learned representation space of the FashionMNIST dataset and the convergence of the embedding CNN in Figure 5a.

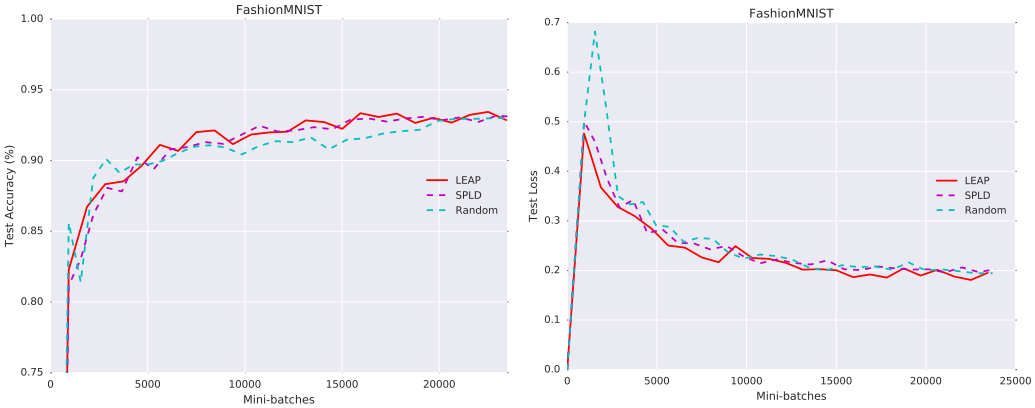

Figure 4: Averaged results of 5 independent runs on FashionMNIST. The solid line represents our LEAP framework.

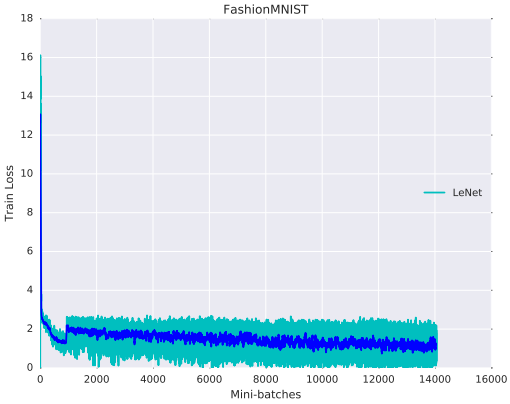

(a) Train Loss

(b) FashionMNIST feature representation space

Figure 5: (a) The train loss with the moving average and (b) the feature representation space learned by LeNet (embedding CNN) using the LEAP framework on FashionMNIST.

4.3 LEAP WITH VGG-16 + RESNET-18 ON CIFAR-10

We ran two sets of experiments on CIFAR-10, one with a fixed learning rate and the other with a learning rate scheduler identical to that of the Wide ResNet (Zagoruyko & Komodakis, 2016) training scheme. The CIFAR-10 training set was augmented with normalization, random horizontal flip and random crop. We used VGG-16 as our embedding CNN and ResNet-18 as our classifier in the CIFAR-10 experiments. We chose ResNet-18 over other architectures because it is faster to train and achieves good performance on CIFAR-10. Our experiments revealed that VGG-16 (Simonyan & Zisserman, 2014) learned strong and rich feature representations which yielded the best convergence on the Magnet Loss. Therefore, we treat the VGG-16 model as a feature extraction engine and use it for feature extraction from CIFAR-10 images, without any fine-tuning. In the first experiment, the classifier is trained using SGD with a momentum of 0.9, weight decay of 0.0005, a fixed learning rate of 0.001 and batch size of 128. In the second experiment, the classifier is trained with batch sizes of 128 as well, using SGD with a momentum of 0.9, weight decay of 0.0005 and a starting learning rate of 0.1 which is dropped by a factor of 0.1 at 60, 120 and 160

epochs. In our experiments, this would translate to 23400, 46800 and 62400 mini-batch updates. In both experiments, VGG-16 is trained with randomly sampled mini-batches of size 64 and optimized using Adam at a learning rate of 0.0001.

At a fixed learning rate of 0.001, training ResNet-18 with the LEAP framework results in a faster convergence to achieve higher test accuracy than either Random sampling or SPLD (Figure 6). Interestingly, SPLD and Random sampling show comparable results on an average of 5 runs. This is because SPLD uses K-Means to partition its data into k clusters at the start of training, which would lead to sub-optimal clustering results containing samples that are not of similar-instance level. As a result, the SPLD would not be selecting diverse samples of similar-instance level when traversing through the different clusters. On the other hand, our LEAP framework ensures that optimal clustering is achieved using the Magnet Loss and the classifier is able to use the learned representation space adaptively as training progresses. This ensures that during early stages of training, the mini-batches being fed into the student CNN are parameterized with truly diverse and easy samples. As the model matures, the mini-batches maintain diversity, as well as a mix of easy and hard samples.

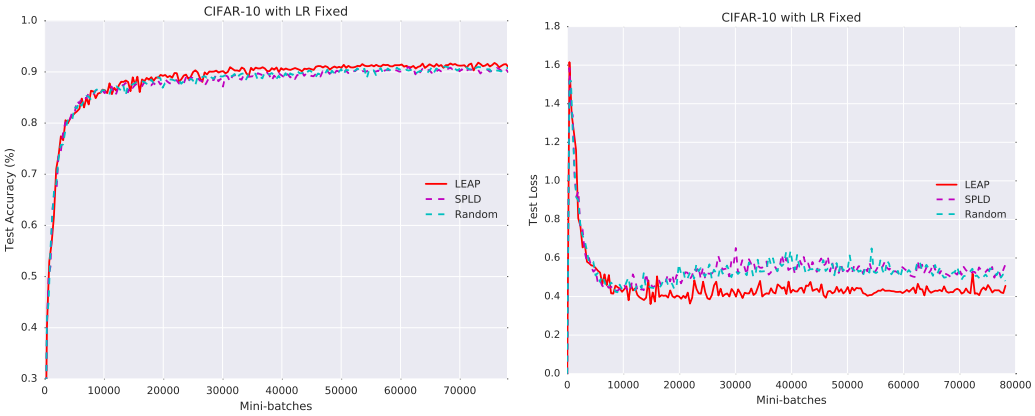

Figure 6: Averaged results of 5 independent runs on CIFAR-10 with a fixed learning rate of 0.001. The solid line represents our LEAP framework.

The second set of experiments we ran on CIFAR-10 with the learning rate scheduler shows that LEAP converges faster earlier on in training compared to SPLD and Random sampling (Figure 7). As the learning rate drops by a factor of 0.1 at 23400, 46800 and 62400 mini-batch updates, the classifier under the LEAP training protocol eventually achieves a higher test accuracy than SPLD and Random sampling. Also, the test loss of LEAP is lower than that of SPLD and Random. This is expected because LEAP is designed to sample for heterogeneity, thus maximizing the diversity relevant to the learning stage of the student CNN, which also helps prevent overfitting. The learned feature representation space of the augmented CIFAR-10 data is shown in Figure 8b and the convergence of the embedding CNN is presented in Figure 8a.

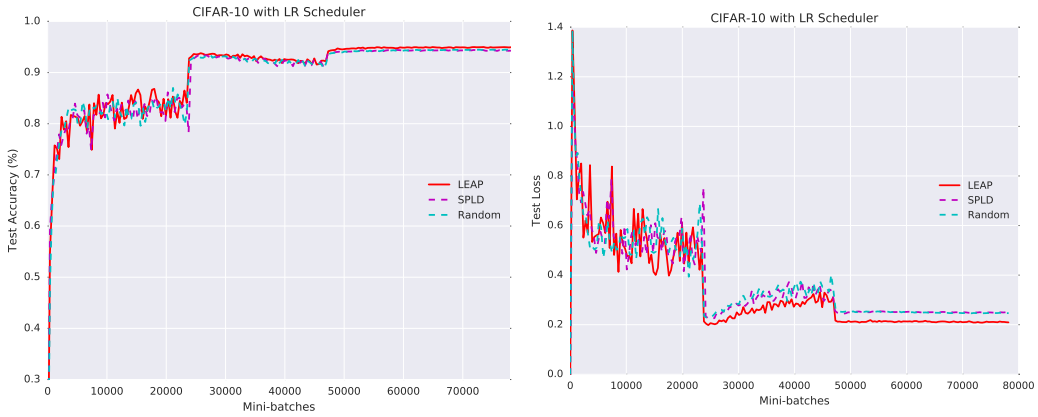

Figure 7: Averaged results of 5 independent runs on CIFAR-10 with a learning rate scheduler identical to the Wide ResNet training scheme. The solid line represents our LEAP framework.

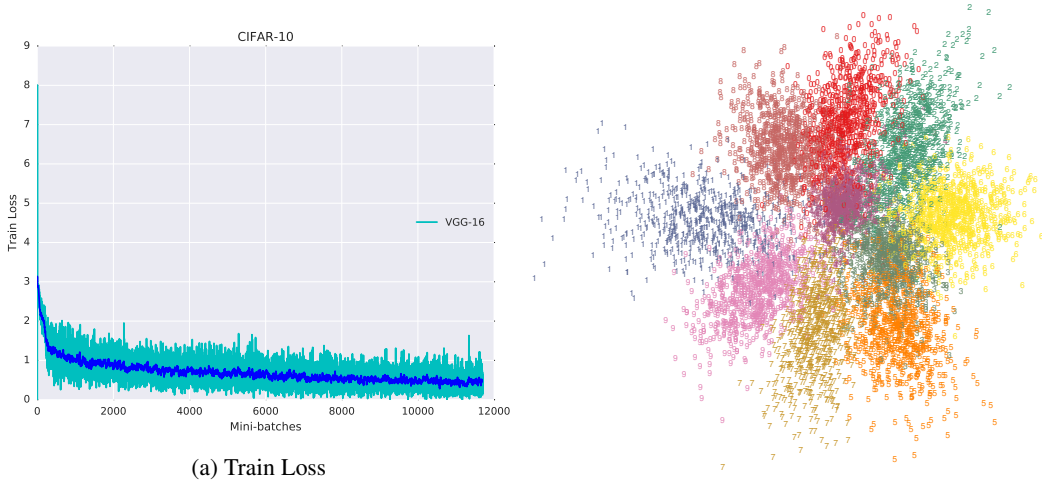

(a) Train Loss

(b) CIFAR-10 feature representation space

Figure 8: (a) The train loss with the moving average and (b) the feature representation space learned by VGG-16 (embedding CNN) using the LEAP framework on CIFAR-10.

5 Discussion and Future Work

An important finding is that fusing a salient non-linear representation space with a dynamic learning strategy can help a DNN converge towards an optimal solution. A random curriculum or a dynamic learning strategy without a good representation space were consistently found to converge more slowly. Biasing samples based on the *easiness* and *true diverseness* to select mini-batches shows improvement in convergence to achieve classification performance comparable or better than state-of-the-art. In cases where the classification dataset is balanced and the classes are clearly identifiable, we showed that our end-to-end LEAP training protocol is practical. However, an interesting line of work would be to apply LEAP on more complex classification datasets where there is room for improvement on classification results. Another interesting area of application would be learning representations using DML for different computer vision tasks (i.e. human pose estimation, human activity recognition, semantic segmentation, etc.) and fusing a representative SPL strategy to train the student CNN.

6 CONCLUSION

We introduced LEAP, an end-to-end representation learning SPL strategy for mini-batch formation. Our method uses an embedding CNN for learning an expressive representation space through a DML technique called the Magnet Loss. The student CNN is a classifier which can exploit this new knowledge from the representation space to place the *true diverseness* and *easiness* as sample importance priors during online mini-batch selection. LEAP training achieves faster convergence and higher test performance on MNIST, FashionMNIST and CIFAR-10 using a combination of deep CNN architectures. We hope this will help foster progress of end-to-end SPL fused DML strategies for DNN training, where a number of potentially interesting directions can be considered for further exploration. Code reproducing these results will be available on GitHub post-review.

ACKNOWLEDGMENTS

Acknowledgements were removed for double blind review.

REFERENCES

Sean Bell and Kavita Bala. Learning visual similarity for product design with convolutional neural networks. *ACM Trans. on Graphics (SIGGRAPH)*, 34(4), 2015.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pp. 41–48, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-516-1. doi: 10.1145/1553374.1553380. URL http://doi.acm.org/10.1145/1553374.1553380.

Christopher B Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. In *Advances in Neural Information Processing Systems 30*. 2016.

Jeffrey L. Elman. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99, 1993.

Alireza Fathi, Zbigniew Wojna, Vivek Rathod, Peng Wang, Hyun Oh Song, Sergio Guadarrama, and Kevin P. Murphy. Semantic instance segmentation via deep metric learning. *CoRR*, abs/1703.10277, 2017. URL http://arxiv.org/abs/1703.10277.

Andrea Frome, Greg Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, MarcAurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Neural Information Processing Systems (NIPS)*, 2013.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90. URL https://doi.org/10.1109/CVPR.2016.90.

Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. Collaborative metric learning. In *Proceedings of the 25th International Conference on World Wide Web (WWW)*, 2017.

Wenhui Huang, Jason Gu, Xin Ma, and Yibin Li. Self-paced model learning for robust visual tracking. *Journal of Electronic Imaging*, 26(1):013016–013016, 2017.

James Steven Supancic III and Deva Ramanan. Self-paced learning for long-term tracking. In *2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, June 23-28, 2013*, pp. 2379–2386, 2013. doi: 10.1109/CVPR.2013.308. URL https://doi.org/10.1109/CVPR.2013.308.

Lu Jiang, Deyu Meng, Teruko Mitamura, and Alexander G. Hauptmann. Easy samples first: Self-paced reranking for zero-example multimedia search. In Kien A. Hua, Yong Rui, Ralf Steinmetz, Alan Hanjalic, Apostol Natsev, and Wenwu Zhu (eds.), *ACM Multimedia*, pp. 547–556. ACM, 2014a. ISBN 978-1-4503-3063-3. URL http://dblp.uni-trier.de/db/conf/mm/mm2014.html#JiangMMH14.

Lu Jiang, Deyu Meng, Shoou-I Yu, Zhenzhong Lan, Shiguang Shan, and Alexander Hauptmann. Self-paced learning with diversity. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2078–2086. Curran Associates, Inc., 2014b. URL http://papers.nips.cc/paper/5568-self-paced-learning-with-diversity.pdf.

M. Hadi Kiapour, Xufeng Han, Svetlana Lazebnik, Alexander C. Berg, and Tamara L. Berg. Where to buy it: Matching street clothing photos in online shops. In *ICCV*, pp. 3343–3351. IEEE Computer Society, 2015. ISBN 978-1-4673-8391-2. URL http://dblp.uni-trier.de/db/conf/iccv/iccv2015.html#KiapourHLBB15.

M. Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 1*, NIPS'10, pp. 1189–1197, USA, 2010. Curran Associates Inc. URL http://dl.acm.org/citation.cfm?id=2997189.2997322.

Agata Lapedriza, Hamed Pirsiavash, Zoya Bylinskii, and Antonio Torralba. Are all training examples equally valuable? *arXiv preprint arXiv:1311.6510*, 2013.

Yann LeCun, Lon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, volume 86, pp. 2278–2324, 1998. URL http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.42.7665.

Yong Jae Lee and K. Grauman. Learning the easy things first: Self-paced visual category discovery. In *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition*, CVPR '11, pp. 1721–1728, Washington, DC, USA, 2011. IEEE Computer Society. ISBN 978-1-4577-0394-2. doi: 10.1109/CVPR.2011.5995523. URL http://dx.doi.org/10.1109/CVPR.2011.5995523.

Xiang Li, Aoxiao Zhong, Ming Lin, Ning Guo, Mu Sun, Arkadiusz Sitek, Jieping Ye, James Thrall, and Quanzheng Li. Self-paced convolutional neural network for computer aided detection in medical imaging analysis. In *International Workshop on Machine Learning in Medical Imaging*, pp. 212–219. Springer, 2017.

Ilya Loshchilov and Frank Hutter. Online batch selection for faster training of neural networks. *CoRR*, abs/1511.06343, 2015. URL http://arxiv.org/abs/1511.06343.

Oren Rippel, Manohar Paluri, Piotr Dollár, and Lubomir D. Bourdev. Metric learning with adaptive density discrimination. *CoRR*, abs/1511.05939, 2015. URL http://arxiv.org/abs/1511.05939.

Enver Sangineto, Moin Nabi, Dubravko Culibrk, and Nicu Sebe. Self paced deep learning for weakly supervised object detection, 2016.

Florian Schroff, Dmitry Kalenichenko, and James Philbin. 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.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. URL http://arxiv.org/abs/1409.1556.

Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 29*, pp. 1857–1865. Curran Associates, Inc., 2016. URL http://papers.nips.cc/paper/6200-improved-deep-metric-learning-with-multi-class-n-pair-loss-objective.pdf.

Hyun Oh Song, Stefanie Jegelka, Vivek Rathod, and Kevin Murphy. Deep metric learning via facility location, 2016a.

Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 4004–4012, 2016b. doi: 10.1109/CVPR.2016.434. URL https://doi.org/10.1109/CVPR.2016.434.

Ye Tang, Yu-Bin Yang, and Yang Gao. Self-paced dictionary learning for image classification. In *Proceedings of the 20th ACM international conference on Multimedia*, pp. 833–836. ACM, 2012.

Evgeniya Ustinova and Victor S. Lempitsky. Learning deep embeddings with histogram loss. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 4170–4178, 2016. URL `http://papers.nips.cc/paper/6464-learning-deep-embeddings-with-histogram-loss`.

L.J.P. van der Maaten and G.E. Hinton. Visualizing high-dimensional data using t-sne. 2008.

Jian Wang, Feng Zhou, Shilei Wen, Xiao Liu, and Yuanqing Lin. Deep metric learning with angular loss. *CoRR*, abs/1708.01682, 2017. URL `http://arxiv.org/abs/1708.01682`.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.

Xiaofan Zhang, Feng Zhou, Yuanqing Lin, and Shaoting Zhang. Embedding label structures for fine-grained feature representation, 2015.