

DEEP NEAREST CLASS MEAN CLASSIFIERS

Samantha Guerriero
VANDAL Laboratory
Sapienza Rome University

Barbara Caputo
VANDAL Laboratory
Italian Institute of Technology

Thomas Mensink*
Computer Vision Group
University of Amsterdam

ABSTRACT

In this paper we introduce DeepNCM, a Nearest Class Mean classification method enhanced to directly learn highly non-linear deep (visual) representations of the data. To overcome the computational expensive process of recomputing the class means after every update of the representation, we opt for approximating the class means with an online estimate. Moreover, to allow the class means to follow closely the drifting representation we introduce per epoch mean condensation. Using online class means with condensation, DeepNCM can train efficiently on large datasets. Our (preliminary) experimental results indicate that DeepNCM performs on par with SoftMax optimised networks.

1 INTRODUCTION

We are interested in Nearest Class Mean (NCM) classifiers, a well-known model for multi-class classification, where a test image is classified to the class with the closest class data mean. This in contrast to the SoftMax classifier, where for each class a weight vector is learned, in NCM the class representation is based solely on the mean representation of the images belonging to that class. While NCM has gained popularity and has widely been used, *e.g.* in the Fisher Discriminant Analysis and in the Rocchio classifier [Webb (2002); Manning et al. (2008)], it has often been stated to be weak because it is a linear classifier, and recently due to its inability to learn deep representations.

In this paper, we introduce our ongoing work on DeepNCM, which learns directly a highly non-linear deep representation while enforcing the desired property that images will be mapped closer to their class mean than any other class means. DeepNCM can be considered as the extension of the Metric Learning approach for NCM (Mensink et al., 2013), where a *shared visual representation* has been learned on top of hand designed fixed features, which generalises well over novel classes and therefore could be successfully used for class incremental learning.

We believe that DeepNCM, to date, has not yet been fully explored because computing class means over a changing representation would be overly expensive. Here, we circumvent expensive updates by approximating class means with an online estimate. This approximates the real mean, while guaranteeing high adaptability to the changing deep representation as new samples are processed.

Relation to other approaches Most notably, our proposed DeepNCMs relate to:

- (i) The incremental classifier learning approach of Rebuffi et al. (2017), where NCM classifiers are applied on SoftMax learned visual representations. While NCM classifiers are used, they do use a representation learned using a (rather) standard SoftMax classifiers, on a fixed-size training size to meet memory constraints, and with a distillation loss to overcome catastrophic forgetting. Our proposed DeepNCM would allow to directly learn the visual representation.
- (ii) The few-shot learning approach of Prototypical Networks (Snell et al., 2017), where a deep representation is learned based on NCM classifiers, however they consider the scenario of episode-few-shot classification (Vinyals et al., 2016), where each episode is a classification task using just a handful classes and data points. Computing exact class mean over all images in an episode is in this scenario trivial. In contrast to DeepNCM, their method would not extend to a larger number of images or classes, *e.g.* when not all classes are present in a (mini-)batch.

*thomas.mensink@uva.nl

Contributions We propose DeepNCM, a NCM-based method for Deep Learning. Our key contributions are using an online approximation for the class means, together with a procedure which we coin mean condensation, which aim to summarise the mean of previous epochs. We present some exciting results, which show that DeepNCM performs on par with its (standard) SoftMax counterpart. We are currently working to extend the DeepNCM method to class Incremental Learning and Open-Set Classification, where it is expected that the proposed NCM-based learned representation will have an edge over SoftMax-based representations.

2 DEEP NEAREST MEAN CLASSIFICATION

The nearest class mean (NCM) classifier is a distance-based classifier, which assigns an image to the class with the closest mean:

$$y^* = \operatorname{argmin}_{y \in \{1, \dots, Y\}} d(\mathbf{x}, \boldsymbol{\mu}_y), \quad (1)$$

where Y denotes the number of classes, class mean $\boldsymbol{\mu}_y = \frac{1}{N_y} \sum_{i: y_i=y} \mathbf{x}_i$, and N_y the number of examples in class y . The success of NCM classifiers critically depends on the distance metric used.

Mensink et al. (2013) proposed a probabilistic interpretation of the class mean distances, to learn a (squared) low-rank Mahalanobis distance, by maximising the log-likelihood:

$$\mathcal{L} = \frac{1}{N} \sum_i \ln p(y_i | \mathbf{x}_i), \quad p(y | \mathbf{x}) \propto \exp -\frac{1}{2} d_{xy}^W, \quad d_{xy}^W = (\mathbf{x} - \boldsymbol{\mu}_y)^\top W^\top W (\mathbf{x} - \boldsymbol{\mu}_y) \quad (2)$$

where $\mathbf{x} \in \mathbb{R}^D$ and $W \in \mathbb{R}^{m \times D}$, with $m \leq D$ as intrinsic dimension of the metric space. The final NCM classifier can be identified as a multi-class SoftMax classifier, with constrained bias terms and weight vectors, *e.g.* the weight vector is constraint to: $\mathbf{w}_y = W^\top W \boldsymbol{\mu}_y$, see (Mensink et al., 2013).

NCM for Learned Representations Instead of relying on a fixed representation and learning a Mahalanobis metric W , the current standard is to learn the parameters of a deep representation $\phi(\cdot)$. We postulate that the highly non-linear nature of deep representations eliminates the need of a linear metric W and allows to use the Euclidean distance between the deep representations:

$$d_{xy}^\phi = (\phi(x) - \boldsymbol{\mu}_y^\phi)^\top (\phi(x) - \boldsymbol{\mu}_y^\phi), \quad \boldsymbol{\mu}_y^\phi = \frac{1}{N_y} \sum_{i: y_i=y} \phi(x_i), \quad (3)$$

where $\phi(x)$ is the deep representation of image x , which evolves during training. For example $\phi(x)$ could be the penultimate layer of a ConvNet, such as AlexNet or ResNet.

Online Mean Updates It is – in theory – straightforward to learn $\phi(\cdot)$ via the gradients of the log-loss by stochastic gradient descent. However, due to the drift caused by learning the representation, after each mini batch of images, (re-)computing class means would cost a full epoch over the train data. This makes training prohibitively (computationally) expensive for large datasets.

Approximating the class means with the per-batch class means is not a sound solution and has two main disadvantages. First, when using small batch sizes (*e.g.* due to a large network structure) the batch class mean is likely to be a poor approximation of the real mean. Second, when not all classes are present in a batch, ad hoc modifications to the classification problem are required, *i.e.* removing classes not present in a given batch will change the classification problem and will change the normaliser of the probability distribution per batch. This latter will become more severe when small batch sizes are used, especially in combination with a large number of classes in the dataset.

For DeepNCM, we therefore propose to use an online estimate of the class means:

$$\boldsymbol{\mu}_{y_i}^\phi \leftarrow \frac{n_{y_i}}{n_{y_i} + 1} \boldsymbol{\mu}_{y_i}^\phi + \frac{1}{n_{y_i} + 1} \phi(x_i), \quad (4)$$

where n_{y_i} denotes the current number of samples in class y_i . The class means are updated with the relevant samples every batch, following smoothly the drifting learned representation $\phi(\cdot)$.

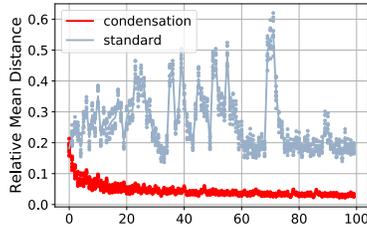


Figure 1: **Mean condensation experiment:** we evaluate the relative distance, $R_y = \|\mu_y^\phi - \hat{\mu}_y^\phi\|_2^2 / \|\hat{\mu}_y^\phi\|_2^2$, after each epoch between the *online-mean* and the *true-training-mean* for each class y . We study the influence of mean condensation, using the CIFAR-10 (Krizhevsky, 2009) dataset, for 100 epochs. Based on the decreasing value of R_y with mean condensation, we conclude that restarting the mean every epoch is beneficial for learning.

	MNIST	CIFAR-10	CIFAR-100
Pixel representation			
SoftMax	92.0		
NCM	91.4		
Simple Deep			
SoftMax	99.3	72.9	
DeepNCM	98.8	71.6	
ResNet-32 He et al. (2015)			
SoftMax		91.3	66.6
DeepNCM		91.4	65.4

(a) **Comparison of DeepNCM with SoftMax**

Method	Initialisation	NCM using novel means	Fine tuning	
			1 ep	5 ep
SoftMax	scratch	n/a	25.4	52.8
	SFT-C50	42.3	58.9	66.6
NCM	NCM-C50	40.7	57.2	65.1

(b) **Generalisation experiment:** CIFAR-100 is split into two disjoint subsets, 50 classes are used for training (C50), the remaining for generalisation and fine-tuning (G50). Using online-means, 1 epoch fine-tuning has similar computation cost to computing just the new class means, while improving accuracy by over 15%.

Table 1: Experiments on MNIST and CIFAR-10/100. DeepNCM performs on par with SoftMax.

Mean Condensation When computing the online mean, the most recent example counts as much as the first example, while the representation has drifted since. To assign more weight to recent examples, we introduce per epoch mean condensation. At the start of each epoch, we set $n_y = 1 \forall y \in Y$, i.e. the *historical* class means count just as a single data point in the current epoch. Based on preliminary results, we conclude that mean condensation is beneficial for training, see Fig. 1.

3 EXPERIMENTS

In this section we describe some of our initial experiments using the proposed DeepNCM classifier on the MNIST, CIFAR-10 and CIFAR-100 datasets, reporting Top-1 accuracy. All models are trained using SGD with Nesterov momentum, with dropout and explicit weight regularisation.

Comparison to SoftMax First, we compare the NCM classifiers to SoftMax classifiers, the current standard for image classification, on multi-class image classification. We evaluate a (fixed) pixel based representation, a shallow ConvNet, and a ResNet-32He et al. (2015) architecture. From the ResNet-32 network point of view, the NCM classifiers uses the penultimate layer as $\phi(x)$ and the class means to learn the network from scratch. The results in Tab. 1a show that DeepNCM is able to learn strong visual representations, and performs on par with the SoftMax classifier.

Generalisation to new classes A selling point of the NCM over fixed representations is that it generalises well to new classes. Here we perform a preliminary generalisation experiment using the CIFAR-100 dataset, split in two subsets, see Tab. 1b. Starting from the learned representation on the C50 subset, we either evaluate NCM using the means of the novel classes or we do fine-tuning for 1 (or 5) epochs on the G50 subset. In the latter case, we either use the SoftMax classifier or the DeepNCM classifier on online means. The results indicate that DeepNCM and SoftMax both learn representations which generalise well and comparably to novel classes.

4 CONCLUSION

We have introduced DeepNCM, an NCM based deep learning framework, which is computationally feasible due to the use of online class mean approximations, training is improved by using mean condensations. Experimentally we have shown that DeepNCM offers a compelling alternative to training (standard) SoftMax-based ConvNets. Our current work is in extending DeepNCM in the context of class incremental learning and open-set classification, where class means seem a natural way to describe categories, to incorporate new classes, and to define class boundaries.

REFERENCES

- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2015.
- A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- C. Manning, P. Raghavan, and H. Schütze. *Introduction to information retrieval*. Cambridge University Press Cambridge, 2008.
- T. Mensink, J. Verbeek, F. Perronnin, and G. Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE Trans. PAMI*, 2013.
- S. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert. icarl: Incremental classier and representation learning. In *CVPR*, 2017.
- J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017.
- O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *NIPS*, 2016.
- A. Webb. *Statistical pattern recognition*. Wiley, New-York, NY, USA, 2002.