

Article

Transfer Incremental Learning Using Data Augmentation

Ghouthi Boukli Hacene ^{1,2,*} , Vincent Gripon ^{1,2} , Nicolas Farrugia ¹ and Matthieu Arzel ¹
and Michel Jezequel ¹

¹ Lab-STICC, IMT Atlantique, 29280 Plouzané, France; vincent.gripon@imt-atlantique.fr (V.G.); nicolas.farrugia@imt-atlantique.fr (N.F.); matthieu.arzel@imt-atlantique.fr (M.A.); michel.jezequel@imt-atlantique.fr (M.J.)

² Montreal Institute for Learning Algorithms (MILA), Montréal, QC H3C 3J7, Canada

* Correspondence: ghouthi.bouklihacene@imt-atlantique.fr

Received: 22 October 2018; Accepted: 28 November 2018; Published: 6 December 2018



Abstract: Deep learning-based methods have reached state of the art performances, relying on a large quantity of available data and computational power. Such methods still remain highly inappropriate when facing a major open machine learning problem, which consists of learning incrementally new classes and examples over time. Combining the outstanding performances of Deep Neural Networks (DNNs) with the flexibility of incremental learning techniques is a promising venue of research. In this contribution, we introduce Transfer Incremental Learning using Data Augmentation (TILDA). TILDA is based on pre-trained DNNs as feature extractors, robust selection of feature vectors in subspaces using a nearest-class-mean based technique, majority votes and data augmentation at both the training and the prediction stages. Experiments on challenging vision datasets demonstrate the ability of the proposed method for low complexity incremental learning, while achieving significantly better accuracy than existing incremental counterparts.

Keywords: transfer learning; incremental learning; computer vision

1. Introduction

Humans have the ability to incrementally learn new pieces of information through time, building over previously acquired knowledge. This process is, most often, nondestructive, and results in what is often referred to as “curriculum learning” in the literature [1]. On the contrary, it has been known for decades that neural networks’ learning procedures, despite the fact they originally were proposed as a simplifying model for brain mechanisms, suffer from “catastrophic forgetting” [2,3], or the fact that previously learned knowledge is destroyed when learning a new one.

In recent years, deep learning has become the golden standard in many supervised learning challenges, especially in the field of computer vision [4–6]. Deep Learning relies on the use of a large number of trainable parameters, that are carefully adjusted using stochastic gradient descent based algorithms. Learning novel data using the same set of parameters inevitably leads to the loss of the previously acquired knowledge. This is why many techniques have proposed learning distinct deep learning systems over the course of time, letting another algorithm decide which one to use at prediction stage [7,8]. Such methods can quickly result in very complex systems, that are likely to fail in adversarial conditions [9].

Formally, an incremental learning approach would satisfy the following criteria [10]:

1. An ability to learn data using one (or a few) example(s) at a time, in any order, without requiring to reconsider or store previous ones.

2. An ability to sustain a classification accuracy comparable to state-of-art methods while traversing successive incremental learning stages, thus avoiding catastrophic forgetting.
3. Low computation and memory footprints, during training and classifying phases, that should remain sublinear in both of number of examples and their dimension.

Satisfying these three criteria while maintaining a competitive accuracy of the proposed systems has remained a key open challenge so far.

A promising venue of research lies in “transfer learning” methods [7], which make use of very efficient pre-trained deep neural networks previously obtained using huge datasets of signals related to the tasks at hand. As a result, very high quality feature vectors can be used to feed the incremental learning techniques, which can then achieve reasonable performances despite using simplistic mechanisms [8].

In this paper, we introduce Transfer Increment Learning with Data Augmentation (TILDA), an incremental learning method that provides (a) a robust selection of feature vectors in subspaces, and (b) prediction procedures making use of data-augmentation. We stress the method using challenging vision datasets, namely CIFAR10, CIFAR100 and ImageNet ILSVRC 2012. As a result the proposed method allows us to:

- Perform incremental learning following the above-mentioned definition,
- Approach state-of-the-art performances on vision datasets,
- Reduce memory usage and computation time by several order of magnitude compared to other incremental approaches.

2. Related Work

There has been an interest in incremental learning for a long time [11–13]. For example, methods have been proposed [14–16] to address this problem with the aim at bounding memory footprint (see criterion 3). These approaches perform learning one subset at a time using Support Vector Machines (SVMs). More precisely, a new SVM is trained for each batch of new data, exploiting previous support vectors. Since the latter are not conveying the full extent of previous data, the newly trained SVM suffers from *catastrophic forgetting* [2,3], and thus violate criterion 2 defined in the introduction.

Another incremental learning algorithm, called “Learn++” was introduced [17,18]. This algorithm adds weak one-vs-all classifiers to accommodate new classes. Therefore, it may result in an excessive computational complexity and memory usage, disobeying criterion 3. It also needs training data for all classes to occur repeatedly, which contradicts criterion 1.

Research showed also the possibility for the sequential learning of data [19]. However, this requires to choose a correct ordering of the whole dataset, which does not fulfil criterion 1. In [20], the authors proposed use of a pre-trained and unchanged DNN as a feature extractor followed by the Nearest Class Mean classifier (NCM). NCM summarises each class using the average feature vector of all examples observed for the class so far. Classification processes by assigning the class of the most similar average vector using a metric that can be learned from data. Compared to other parametric classifiers [20–22], NCM showed better performances in incremental learning scenarios. However, NCM gives a lower accuracy than state-of-art methods even when it uses all the dataset, hence does not fulfil criterion 2.

In [23], a quite different incremental method called Budget Restricted Incremental Learning (BRIL) was proposed. BRIL combines “transfer learning” [7,8] with binary associative memories. A pre-trained DNN is used as feature extractor, as mentioned in [20], while binary associative memories act as a classifier. A product random sampling is performed as an intermediate between the pre-trained DNN and the classifier. Despite being compliant with criteria 1 and 3, the accuracy remains significantly lower than existing counterparts, which violates criterion 2.

Kuzborskij et al. [24] showed that new classes can be added to a multi-class classifier with limited impact on accuracy when the classifiers can be retrained from at least a small amount of data

belonging to all classes. Using this, in [10] the authors proposed an incremental learning method called “Incremental Classifier and Representation Learning” (iCaRL), based on a trainable DNN feature extractor, followed by a single classification layer. The classification process is inspired by NCM: it computes the mean of feature vectors for each class, and assigns the label of the nearest prototype. However, memory usage can easily increase, especially when the dataset is made of high resolution images such as ImageNet, which may violate criterion 3. Moreover, the iCaRL method, when trained on data streams containing only a few classes at a time, provides low accuracy as shown in [10], hence iCaRL does not respect criterion 2. To reach good performances and a comparable accuracy to state-of-art methods, iCaRL thus needs to be trained over batches of data containing a large part of the dataset, which does not correspond to an incremental learning scenario and infringes 1.

In this paper, we introduce TILDA that builds upon previously proposed work, attempting to cover all three criteria for efficient incremental learning. As in iCaRL and BRIL, TILDA uses a pre-trained DNN as feature extractor. TILDA also uses an NCM-inspired classifier over the feature vectors obtained from by the pre-trained DNN. Data augmentation is performed on both training and classification datasets, aiming to improve accuracy. Consequently, there is no need to retrain the system with previous data, nor to perform computationally intensive processing when new data comes in. In addition, learning new data does not damage previously learned information.

3. Proposed Method

In this section, we describe the TILDA method. We start by giving a high level overview of the process, and then we explain the details.

3.1. Overview of the Proposed Method

TILDA is built upon four main steps: (1) a pre-trained DNN to perform feature extraction, (2) a technique to project features into low dimensional subspaces, (3) an assembly of NCM-inspired Classifiers (NCMC) applied independently in each subspace (see Figure 1) and (4) a data augmentation inspired scheme to increase accuracy of the classifying process. We develop these steps in the following paragraphs.

The first step consists of using the internal layers of a pre-trained DNN [25] as a generic feature extractor on which subsequent learning is performed. This process has become increasingly popular in the past few years and is often referred to as “Transfer Learning” [26]. The aim is to transfer acquired knowledge on a dataset to another related problem [8].

In the following step we project feature vectors into multiple low dimensional subspaces. More precisely, we split feature vectors into P subvectors. For each class and each subspace, we produce k anchor vectors conveying robust statistical properties about corresponding feature subvectors.

Then, in each subspace, anchor vectors are exploited to perform weak classification of the input data. We use here a NCM inspired method. A majority vote is then performed to obtain an aggregate decision.

Finally, we perform data augmentation on the input signals, thus obtaining multiple decisions for each input data as well as more robust classifiers in each subspace. A second majority vote is performed using these decisions to generate a global prediction.

3.2. Details of the Proposed Method

3.2.1. Pre-Trained Deep Neural Networks

To obtain features from an input signal, TILDA relies on using DNNs that are pre-trained on a large number of examples. Consequently, using the pre-trained inner layers of the DNN acts as a generic feature extractor [8,26,27]. As a matter of fact, inner layers of a deep DNN offer a good generic description of an input image, even when it does not belong to the learning domain [26].

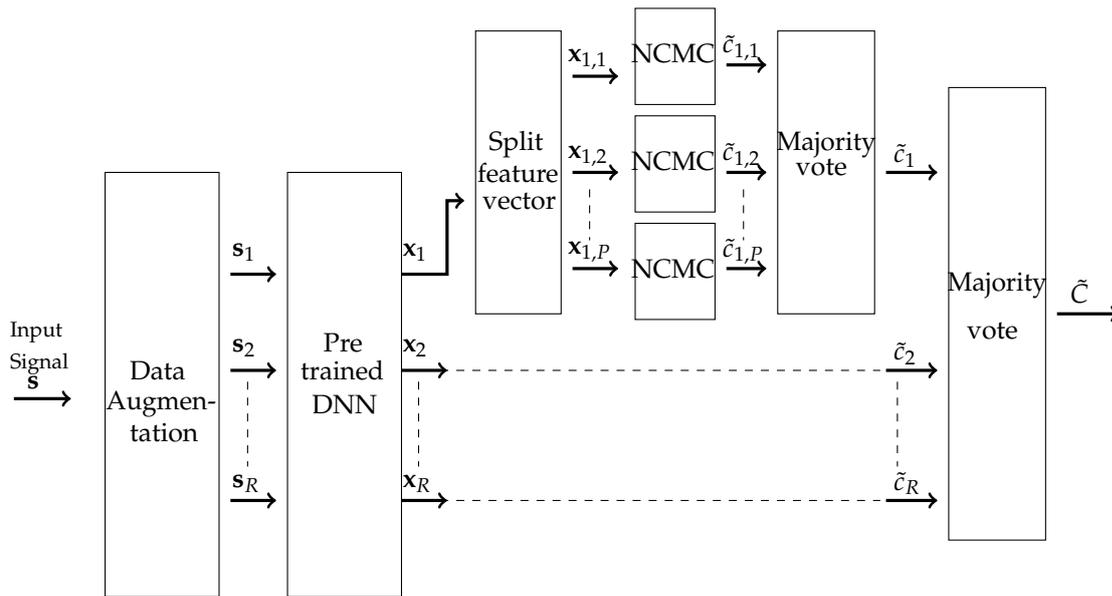


Figure 1. Overview of the proposed method. Given an input signal \mathbf{s} , we first use data augmentation to generate a multiple version of the input signal $(\mathbf{s}_r)_{1 \leq r \leq R}$. Then we use a pre-trained DNN for feature extraction and obtain the corresponding feature vectors $(\mathbf{x}_r)_{1 \leq r \leq R}$. Subsequently, we split each feature vector \mathbf{x}_r into P equal parts $(\mathbf{x}_{r,p})_{1 \leq p \leq P}$, and classify each part $\mathbf{x}_{r,p}$ using a NCM-inspired classifiers (NCMC) containing anchor vectors $(\mathbf{Y}_{c,p})_{1 \leq c \leq C}$. We obtain a class for each part $c_{r,p}$ and do a majority vote to get the class of \mathbf{x}_r . Finally, a second majority vote is done thanks to the obtained classes $(c_r)_{1 \leq r \leq R}$ of all generated signals to get assigned class \tilde{C} to the original input signal \mathbf{s} .

Using “Transfer Learning” ideas, we are not interested in this work in the network’s architecture details, as we simply use the appropriate layers to extract features from a given input.

In the remainder of this paper, we denote by \mathbf{s}^m the m -th input training signal and by \mathbf{x}^m its corresponding feature vector, where $1 \leq m \leq M$ and M is the total number of training signals.

3.2.2. Projection to Low Dimensional Subspaces

Feature extraction allows us to consider the feature vector \mathbf{x}^m instead of the input signal \mathbf{s}^m . Formally, let us denote \mathbf{x}_c^m the fact that feature vector \mathbf{x}^m belongs to class c . We split each \mathbf{x}_c^m into P parts, denoted $(\mathbf{x}_{c,p}^m)_{1 \leq p \leq P}$. For each class and each subspace, we create k anchor vectors initialised with 0, each of them associated with a counter, also initialised by 0. Considering the p -th subspace and the c -th class, we denote by $Y_{c,p} = [\mathbf{y}_{c,p,1}, \dots, \mathbf{y}_{c,p,k}]$ the corresponding anchor vectors and $N_{c,p} = [n_{c,p,1}, \dots, n_{c,p,k}]$ their associated counters.

For each c and p , we aim at using the corresponding anchor vectors as centroids of a clustering of $\{\mathbf{x}_{c,p}^m\}$. To this end, at each step of the training process, we ensure that each anchor vector is a centroid of a clustering of already processed input subvectors, and the associated counter accounts for the cardinality of the corresponding cluster.

Then, each time an input training vector is processed, we identify an anchor vector to be updated. The update simply consists of computing a new anchor vector obtained as a barycenter of the old one with weight given by its counter and the input subvector with weight 1, then incrementing the counter. This procedure is detailed in Algorithm 1. Namely, rather than simply associating the new subvector with the closest anchor vector, which would inevitably lead to unbalanced counters and thus poor performance in prediction, we prefer to take into account counters while performing this association. More precisely, we linearly penalize anchor vectors that are already made of the combination of many subvectors. Note that when two or more anchor vectors gives the same results (distances multiplied by counters), we choose uniformly, at random, one of these anchor vectors.

Note that the learning process is independent on the order of streaming data, and is performed one example at a time, thus enforcing criterion 1 of incremental learning methods described in the introduction.

Algorithm 1 Incremental Learning of Anchor Subvectors

Input: streaming feature vector \mathbf{x}_c^m

```

for  $p := 1$  to  $P$  do
  for  $i := 1$  to  $k$  do
     $d_i = \|\mathbf{x}_p - \mathbf{y}_{c,p,i}\|_2$ 
     $R_i = d_i n_{c,p,i}$ 
  end for
   $\tilde{k} = \arg \min_i R_i$ 
   $\mathbf{y}_{c,p,\tilde{k}} \leftarrow \mathbf{y}_{c,p,\tilde{k}} n_{c,p,\tilde{k}} + \mathbf{x}_{c,p}^m$ 
   $n_{c,p,\tilde{k}} \leftarrow n_{c,p,\tilde{k}} + 1$ 
   $\mathbf{y}_{c,p,\tilde{k}} \leftarrow \mathbf{y}_{c,p,\tilde{k}} / n_{c,p,\tilde{k}}$ 
end for

```

3.2.3. Aggregation of Subspaces Weak Classifiers

At the prediction stage, consider an input signal \mathbf{s} and the associated feature vector \mathbf{x} . We split \mathbf{x} into the corresponding P parts and obtain $(\mathbf{x}_p)_{1 \leq p \leq P}$. We compute Euclidean distances between each \mathbf{x}_p and all anchor subvectors $\mathbf{y}_{c,p,i}$ for which the counter is not 0. Note that there are at most kC such distances, where C is the number of classes seen so far. The class of the closest average anchor subvector is considered as the decision for the p -th subspace. Finally, we apply a majority vote over all subspaces to achieve an aggregate decision (see Algorithm 2). Note that more elaborate strategies can result in higher accuracy but may require more computation during the learning phase as well as memorisation of previously seen examples.

Algorithm 2 Predicting the Class of a Test Input Signal

Input: input signal \mathbf{s}

```

Compute the feature vector  $\mathbf{x}$  associated with  $\mathbf{s}$ 
Initialize the vote vector  $\mathbf{v}$  as the 0 vector with dimension  $C$ 
for  $p := 1$  to  $P$  do
   $v_p = \arg \min_c \left[ \min_i \|\mathbf{x}_p - \mathbf{y}_{c,p,i}\|_2 \right]$ 
   $\mathbf{v}_{v_p} = \mathbf{v}_{v_p} + 1$ 
end for
 $\tilde{c} = \arg \max_c (\mathbf{v}_c)$ 

```

Output: class \tilde{c} attributed to \mathbf{s}

3.2.4. Data Augmentation

We use two data augmentation methods to improve the accuracy and robustness during training and classification.

Data Augmentation during Training

To improve the accuracy without increasing memory usage, data augmentation is applied to the training dataset. We generate multiple version of each training input signal, and we consider the resulting dataset as an input to train the model.

Data Augmentation during Classification

In addition, we propose to obtain multiple predictions for each input signal \mathbf{s} using data augmentation [28]. The idea is to generate multiple versions of the input signal \mathbf{s} that we denote $(\mathbf{s}_r)_{1 \leq r \leq R}$. We perform a prediction of the class associated with each \mathbf{s}_r independently, and then perform a majority vote to obtain the final prediction.

3.2.5. Remarks

We point out multiple facts about the proposed method:

- (a) The learning procedure performs learning one example at a time,
- (b) The learning procedure is computationally light as it only requires performing of the order of d operations where d is the dimension of feature vectors,
- (c) The learning procedure has a small memory footprint, as it only stores averages of feature vectors,
- (d) The learning procedure is such that adding new examples can only increase robustness of the method, so that there is no catastrophic forgetting,
- (e) During prediction stage, memory usage is of the order of kCd and thus is independent on the number of examples and grows linearly with the number of classes,
- (f) During prediction stage, computations are of the order of $kCdR$ elementary operations.

From these facts we derive that TILDA is compliant with criteria 1 and 3 defined in the introduction. In the next section, we devise a set of experiments to evaluate the classification accuracy of the proposed method on challenging datasets (criterion 2).

4. Experiments

In this section we describe the protocol used to test the proposed method and compare its accuracy and memory usage with other incremental learning methods.

4.1. Benchmark Protocol

We propose an incremental learning scenario in which we have streaming data containing new classes or examples. We test and compare Budget Restricted Incremental Learning (BRIL), Nearest Neighbour search (NN), Nearest Class Mean classifier (NCM), Learn++, incremental Classifier and Representation Learning (iCaRL), and finally the proposed method (TILDA). Learn++ uses Classification And Regression Trees (CART) as weak classifiers.

We evaluate the different methods using CIFAR10, CIFAR100 and ImageNet ILSVRC 2012 [29]. We also use 50 ImageNet classes which have not been used to train the CNN (denoted ImageNet50), and which contains 900/100 training/test images per class. All methods take the same feature vectors extracted from Inception V3 [6] as input and use the whole dataset for training, unless explicitly mentioned. This requires to modify iCaRL method by replacing its CNN with a fully connected network. In the following, and for the iCaRL method, we use a MultiLayer Perceptron (MLP) with one hidden layer containing 1024 neurons, and output layer containing C neurons, where C is the number of classes.

The non-incremental learning methods (NI) used are denoted by TMLP and TSVM. TMLP uses transfer learning to compute feature extractors of input data through Inception V3, and then trains a MLP over feature vectors, using the hyperparameters previously described for iCaRL. TSVM method uses Inception V3 to get feature vectors as well, and uses them to train an SVM using Radial Basis Function kernel.

Data augmentation used in TILDA generates a horizontal flip of the original image, and shifts the pixels of the image by one pixel at a time (to the left, right, top, bottom, and on the four diagonals). Thus we generate $R = 10$ images (eight generated by shifting pixels on the image, one generated by horizontal flip and the original one).

4.2. Results

As a preliminary experiment, we aim to show that replacing the last layers of Inception V3 by the proposed method does not compromise the performances obtained on Imagenet ILSVRC 2012. The 5-top accuracy is 94.4% when we use TILDA with $p = 16$ and $k = 30$, and 96.5% when we use the last layers of Inception V3 to classify data. The accuracy obtained by TILDA approaches the one obtained by Inception V3, thus our method does not bring a considerable decrease in performances.

The second experiment is performed on CIFAR10/100, ImageNet50 and ImageNet ILSVRC 2012, in which we show the contribution of data augmentation, NCM-inspired classification, and subspace division on classification accuracy. Therefore, we define three methods: TILDA-DA does not use data augmentation and classifies only the original image, TILDA-NCM disregards NCM inspired classification and uses k feature vectors randomly chosen per class, and TILDA-P, which is TILDA method with no splitting of vectors. Table 1 summarises the accuracy of TILDA, TILDA-DA, TILDA-NCM and TILDA-P, when performing one-shot learning (learn one example at a time). We notice that TILDA-DA, TILDA-NCM and TILDA-P reach lower accuracy than TILDA, which confirms that the combination of data augmentation with NCM-inspired classification and subspace division can achieve good performances.

Table 1. Accuracy on CIFAR10/100, ImageNet50 and ImageNet ILSVRC 2012. TILDA uses the following parameter: $P = 16$ and $k = 30$. We learn incrementally one example at a time.

	TILDA	TILDA-DA	TILDA-NCM	TILDA-P
CIFAR100	69.6%	65.3%	60.7%	67%
CIFAR10	88.7%	86.6%	84.11%	87%
ImageNet50	76%	74.4%	69.2%	72%
ILSVRC 2012	94.4%	91%	89.6%	90%

In the third experiment, we study the effect of both quantization parameters P and k on the accuracy of TILDA (see Figure 2). This experiment demonstrates that TILDA reaches best performances for $P = 16$. In the following, we perform experiments using TILDA with $P = 16$ and $k = 30$. Note that in order to be fair in comparison with other techniques, we do not perform data-augmentation during training or prediction in TILDA in the upcoming experiments.

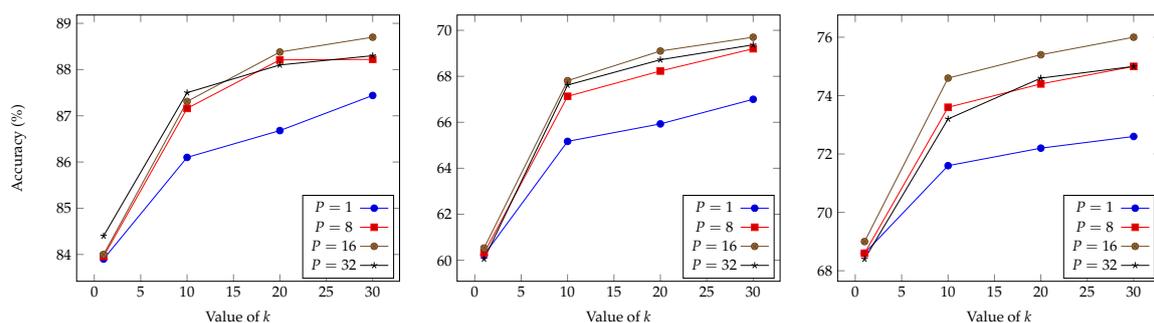


Figure 2. Evolution of the accuracy as a function of P and k for CIFAR10 (left), CIFAR100 (middle) and ImageNet50 (right).

The fourth experiment stresses the effect of class-incremental learning. We adopt a class-incremental scenario (CI), in which methods are trained over streaming data providing all examples from one class simultaneously, one class at a time. We test and compare TILDA-DA, NCM, Learn++ and iCaRL on CIFAR10/100 and ImageNet50 (c.f. Figure 3). Learn++ adds one weak classifier each time a novel class is introduced, and iCaRL stores 30 feature vectors per class. We can see that TILDA-DA outperforms the other methods in this setting.

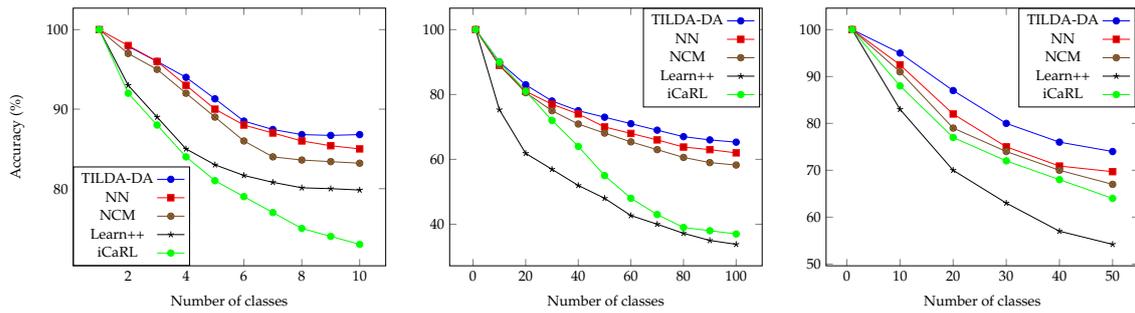


Figure 3. Evolution of the accuracy as a function of number of classes for CIFAR10 (left), CIFAR100 (middle) and ImageNet50 (right).

The next experiment illustrates the behaviour of the accuracy when trying to obtain incremental information from new examples of the same class. We adopt an example-incremental scenario (EI), in which we train the method over streaming data providing new examples without introducing new classes. We test and compare TILDA-DA, NCM, NN, Learn++ and BRIL on CIFAR10/100 and ImageNet50. We divide these datasets into 10 equal size parts, each part containing 5000 examples (500/50 example per class) for CIFAR10/100 and 4500 examples (90 per class) for ImageNet50, and learn one part at a time. Learn++ adds one weak classifier each time a new part is learned. Figure 4 shows that all methods handle example-incremental learning and improve their accuracy each time they learn new information provided by new examples. TILDA-DA consistently obtains higher accuracy than Learn++, NCM, NN and BRIL regardless of the quantity of provided data. Note that Learn++ needs large number of examples to perform, and obtains a low accuracy when only few examples are provided.

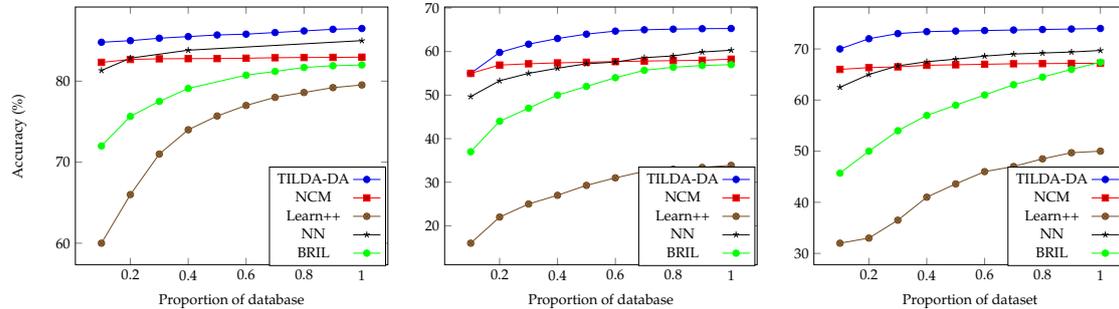


Figure 4. Evolution of the accuracy as a function of number of learning examples for CIFAR10 (left), CIFAR100 (middle) and ImageNet50 (right).

Table 2 presents the different incremental learning methods, obtained accuracies and memory footprints. Learn++ uses either class-incremental scenario (CI) or example-incremental scenario (EI). iCaRL performs learning process using CI. TILDA, NN, NCM, and BRIL use one-shot learning to process one example at a time providing a novel class or additional information, thus they handle both class-incremental and example-incremental at the same time. TILDA outperforms all other incremental learning methods on both accuracy and memory usage.

The last evaluation we perform aims to compare TILDA with a non incremental learning method such as TMLP and TSVM. To do this, we store and train these methods on the whole dataset. The parameters used for TILDA are $P = 16$ and $k = 30$ for CIFAR10/100 and ImageNet50. Table 3 shows that TILDA reaches an accuracy comparable to state-of-art methods.

As shown by the different evaluation, the TILDA method can at any instant classify data with a good accuracy (see Figures 3 and 4), outperforms other incremental learning methods (see Table 2), and approaches state-of-art accuracy (see Table 3). Consequently, TILDA fulfils criterion 2.

Table 2. Comparison of accuracy (Acc) and memory usage (M) relative to full dataset (corresponding to 100%) for the different methods. Note that memory usage of Learn++ method represents the size of weak classifiers, and for iCaRL represents the stored feature vectors and the size of the trainable neural network.

	Only CI		Both CI and EI					Only EI
	Learn++	iCaRL	TILDA	TILDA-DA	NN	NCM	BRIL	Learn++
Acc (CIFAR100)	34%	30%	69.6%	65.3%	60.2%	58.25%	57%	34%
M (CIFAR100)	10.5%	8%	6%	6%	100%	0.2%	6%	6.8%
Acc (CIFAR10)	79.8%	41%	88.7%	86.6%	85%	83%	82%	79.5%
M (CIFAR10)	0.65%	2.7%	0.6%	0.6%	100%	0.02%	0.6%	0.65%
Acc (ImageNet50)	54.2%	64%	76%	74.4%	69.7%	67.2%	67.4%	50%
M (ImageNet50)	4.7%	5.6%	3.3%	3.3%	100%	0.11%	3.3%	3%

Table 3. Comparison of TILDA with non-incremental learning methods.

	TILDA	TILDA-DA	TMLP	TSVM
Acc (CIFAR100)	69.6%	65.16%	68.6%	67.6%
Acc (CIFAR10)	88.7%	86.6%	90%	89.2%
Acc (ImageNet50)	76%	74.4%	75.2%	75%

5. Conclusions

In this paper, we have introduced TILDA, a new incremental learning approach inspired by recently proposed methods. TILDA relies on a pre-trained DNN to process data, a projection technique that defines low-dimensional subspaces, NCM inspired classifiers, and data augmentation at both the training and prediction phases. This addresses previous concerns from previous methods, specifically: (a) iCaRL as TILDA reaches a good accuracy when the streamed data contains one class at a time, (b) BRIL as TILDA provides a good accuracy comparable to state-of-art method, (c) NCM as TILDA uses k anchor vectors instead of one and other methods to increase the accuracy, and (d) Learn++ as TILDA still performs well even if steam data does not contains examples of all classes each time. Experiments on challenging datasets show that: (a) TILDA does not suffer from catastrophic forgetting, (b) TILDA approaches state-of-art accuracy, (c) TILDA uses much less memory usage and gets the same accuracy as nearest neighbour search or even better, (d) TILDA still gives a good accuracy even in the case where we have only one class each time, and (e) finally, to our knowledge, TILDA is the incremental method that reaches the best accuracy. This method is also promising for embedded devices, since it is not necessary to train a DNN or compute extensive operations for learning.

In future work, we plan to explore further the methods for splitting feature vectors, data augmentation strategies and a weighted majority vote to improve the accuracy. We also plan to propose a hardware architecture of TILDA for incremental learning on chip.

Author Contributions: G.B.H. and V.G. conceived and designed the experiments; G.B.H. performed the experiments; G.B.H. and V.G. analyzed the data; G.B.H. wrote the original draft, V.G., N.F. and M.A. edited and reviewed the paper, M.J. supervised this work.

Funding: The research for this paper was financially supported by Pôle de Recherche Avancée en Communications (Pracom).

Acknowledgments: We would like to thank Nvidia which gave us GPUs we used to compute experiments.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bengio, Y.; Louradour, J.; Collobert, R.; Weston, J. Curriculum learning. In Proceedings of the 26th Annual International Conference on Machine Learning, Montreal, QC, Canada, 14–18 June 2009; ACM: New York, NY, USA, 2009; pp. 41–48.
2. Kasabov, N. *Evolving Connectionist Systems: Methods and Applications in Bioinformatics, Brain Study and Intelligent Machines*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2013.
3. French, R.M. Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* **1999**, *3*, 128–135. [[CrossRef](#)]
4. Iandola, F.N.; Han, S.; Moskewicz, M.W.; Ashraf, K.; Dally, W.J.; Keutzer, K. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5 MB model size. *arXiv* **2016**, arXiv:1602.07360.
5. Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv* **2014**, arXiv: 1409.1556.
6. Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; Wojna, Z. Rethinking the inception architecture for computer vision. *arXiv* **2015**, arXiv:1512.00567.
7. Girshick, R.; Donahue, J.; Darrell, T.; Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 580–587.
8. Pan, S.J.; Yang, Q. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **2010**, *22*, 1345–1359. [[CrossRef](#)]
9. Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; Fergus, R. Intriguing properties of neural networks. *arXiv* **2013**, arXiv:1312.6199.
10. Rebuffi, S.A.; Kolesnikov, A.; Sperl, G.; Lampert, C.H. iCaRL: Incremental Classifier and Representation Learning. In Proceedings of the 2017 Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017.
11. Schlimmer, J.C.; Fisher, D. A case study of incremental concept induction. In Proceedings of the 5th National Conference on Artificial Intelligence, AAAI 1986, Philadelphia, PA, USA, 14–15 August 1986; pp. 496–501.
12. Thrun, S. Is learning the n-th thing any easier than learning the first? In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 1996; pp. 640–646.
13. Zhou, Z.H.; Chen, Z.Q. Hybrid decision tree. *Knowl.-Based Syst.* **2002**, *15*, 515–528. [[CrossRef](#)]
14. Syed, N.A.; Huan, S.; Kah, L.; Sung, K. *Incremental Learning with Support Vector Machines*; CiteSeer^X: University Park, PA, USA, 1999.
15. Poggio, T.; Cauwenberghs, G. Incremental and decremental support vector machine learning. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2001; Volume 13, p. 409.
16. Zheng, J.; Shen, F.; Fan, H.; Zhao, J. An online incremental learning support vector machine for large-scale data. *Neural Comput. Appl.* **2013**, *22*, 1023–1035. [[CrossRef](#)]
17. Polikar, R.; Upda, L.; Upda, S.S.; Honavar, V. Learn++: An incremental learning algorithm for supervised neural networks. *IEEE Trans. Syst. Man Cybern. Part C (Appl. Rev.)* **2001**, *31*, 497–508. [[CrossRef](#)]
18. Muhlbaier, M.D.; Topalis, A.; Polikar, R. Learn++. NC: Combining Ensemble of Classifiers with Dynamically Weighted Consult-and-Vote for Efficient Incremental Learning of New Classes. *IEEE Trans. Neural Netw.* **2009**, *20*, 152–168. [[CrossRef](#)] [[PubMed](#)]
19. Pentina, A.; Sharmanska, V.; Lampert, C.H. Curriculum learning of multiple tasks. In Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition, Boston, MA, USA, 7–12 June 2015; pp. 5492–5500.
20. Mensink, T.; Verbeek, J.; Perronnin, F.; Csurka, G. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE Trans. Pattern Anal. Mach. Intell.* **2013**, *35*, 2624–2637. [[CrossRef](#)] [[PubMed](#)]
21. Mensink, T.; Verbeek, J.; Perronnin, F.; Csurka, G. Metric learning for large scale image classification: Generalizing to new classes at near-zero cost. In Proceedings of the Computer Vision–ECCV 2012, Florence, Italy, 7–13 October 2012; pp. 488–501.
22. Ristin, M.; Guillaumin, M.; Gall, J.; Van Gool, L. Incremental learning of NCM forests for large-scale image classification. In Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 3654–3661.

23. Hacene, G.B.; Gripon, V.; Farrugia, N.; Arzel, M.; Jezequel, M. Budget restricted incremental learning with pre-trained convolutional neural networks and binary associative memories. In Proceedings of the 2017 IEEE International Workshop on Signal Processing Systems (SiPS), Lorient, France, 3–5 October 2017; IEEE: Piscataway, NJ, USA, 2017; pp. 1–6.
24. Kuzborskij, I.; Orabona, F.; Caputo, B. From n to $n + 1$: Multiclass transfer incremental learning. In Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition, Portland, OR, USA, 23–28 June 2013; pp. 3358–3365.
25. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2012; pp. 1097–1105.
26. Oquab, M.; Bottou, L.; Laptev, I.; Sivic, J. Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks. In Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Columbus, OH, USA, 23–28 June 2014.
27. Hong, S.; You, T.; Kwak, S.; Han, B. Online Tracking by Learning Discriminative Saliency Map with Convolutional Neural Network. In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 597–606.
28. Ciresan, D.; Meier, U.; Gambardella, L.; Schmidhuber, J. Deep Big Simple Neural Nets Excel on Handwritten Digit Recognition. *arXiv* **2010**, arXiv 1003.0358.
29. Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* **2015**, *115*, 211–252. [[CrossRef](#)]



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).