

Cognitively-Inspired Model for Incremental Learning Using a Few Examples

Ali Ayub, Alan R. Wagner
The Pennsylvania State University
{aja5755, alan.r.wagner}@psu.edu

Abstract

Incremental learning attempts to develop a classifier which learns continuously from a stream of data segregated into different classes. Deep learning approaches suffer from catastrophic forgetting when learning classes incrementally, while most incremental learning approaches require a large amount of training data per class. We examine the problem of incremental learning using only a few training examples, referred to as Few-Shot Incremental Learning (FSIL). To solve this problem, we propose a novel approach inspired by the concept learning model of the hippocampus and the neocortex that represents each image class as centroids and does not suffer from catastrophic forgetting. We evaluate our approach on three class-incremental learning benchmarks: Caltech-101, CUBS-200-2011 and CIFAR-100 for incremental and few-shot incremental learning and show that our approach achieves state-of-the-art results in terms of classification accuracy over all learned classes.

1. Introduction

Humans can continuously learn new concepts over their lifetime. In contrast, modern machine learning systems often must be trained on batches of data [15, 27]. Applied to the task of object recognition, *incremental learning* is an avenue of research that seeks to develop systems that are capable of continually updating the learned model as new data arrives [19]. Incremental learning gradually increases an object classifier’s breadth by training it to recognize new object classes [11].

This paper examines a sub-type of incremental learning known as *class-incremental learning*. Class-incremental learning attempts to first learn a small subset of classes and then incrementally expand that set with new classes. Importantly, class-incremental evaluation of a final model is tested on a single blended dataset, an evaluation known as single-headed evaluation [7]. To paraphrase Rebuffi et al. [38], a class-incremental learning algorithm must:

1. Be trainable from a stream of data that includes instances of different classes at different times;

2. Offer a competitively accurate multi-class classifier for any classes it has observed thus far;
3. Be bounded or only grow slowly with respect to memory and computational requirements as the number of training classes increase.

Creating a high accuracy classifier that incrementally learns, however, is a hard problem. One simple way to create an incremental learner is by tuning the model to the data of the new classes. This approach, however, causes the model to forget the previously learned classes and the overall classification accuracy decreases, a phenomenon known as *catastrophic forgetting* [13, 22]. To overcome this problem, most existing class-incremental learning methods avoid it altogether by storing a portion of the training data from the earlier learned classes and retrain the model (often a neural network) on a mixture of the stored data and new data containing new classes [38, 6, 48]. These approaches are, however, neither scalable nor biologically inspired i.e. when humans learn new visual objects they do not forget the visual objects they have previously learned, nor must humans relearn these previously known objects. Furthermore, current methods for incremental learning require a large amount of training data and are thus not suitable for training from a small set of examples.

We seek to develop a practical incremental learning system that would allow human users to incrementally teach a robot different classes of objects. In order to be practical for human users, an incremental learner should only require a few instances of labeled data per class. Hence, in this paper we explore the Few-Shot Incremental Learning (FSIL) problem, in which an agent/robot is required to learn new classes continually but with only a small set of examples per class.

With respect to class-incremental learning and FSIL, this paper contributes a novel cognitively-inspired method termed Centroid-Based Concept Learning (CBCL). CBCL is inspired by the concept learning model of the hippocampus and the neocortex [29, 40, 33]. CBCL treats each image as an episode and extracts its high-level features. CBCL uses a fixed data representation (ResNet [17] pre-trained on

ImageNet [43]) for feature extraction. After feature extraction, CBCL generates a set of concepts in the form of centroids for each class using a cognitively-inspired clustering approach (denoted as *Agg-Var* clustering) proposed in [3]. After generating the centroids, to predict the label of a test image, the distance of the feature vector of the test image to the n closest centroids is used. Since CBCL stores the centroids for each class independently of the other classes, the decrease in overall classification accuracy is not catastrophic when new classes are learned. CBCL is tested on three incremental learning benchmarks (Caltech-101 [12], CUBS-200-2011 [47], CIFAR-100 [23]) and it outperforms the state-of-the-art methods by a sizable margin. Evaluations for FSIL show that CBCL outperforms some class-incremental learning methods, even when CBCL uses only 5 or 10 training examples per class and other methods use the complete training set per class (500 images per class for CIFAR-100). For FSIL, CBCL even beats a few-shot learning baseline that learns from the training data of all classes (batch learning) on the three benchmark datasets. The main contributions of this paper are:

1. A cognitively-inspired class-incremental learning approach is proposed that outperforms the state-of-the-art methods on the three benchmark datasets listed above.
2. A novel centroid reduction method is proposed that bounds the memory footprint without a significant loss in classification accuracy.
3. A challenging incremental learning problem is examined (FSIL) and experimental evaluations show that our approach results in state-of-the-art accuracy when applied to this problem.

2. Related Work

The related work is divided into two categories: traditional approaches that use a fixed data representation and class-incremental approaches that use deep learning.

2.1. Traditional Methods

Early incremental learning approaches used SVMs [10]. For example, Ruping [42] creates an incremental learner by storing support vectors from previously learned classes and using a mix of old and new support vectors to classify new data. Most of the earliest approaches did not fulfill the criteria for class-incremental learning and many required old class data to be available when learning new classes: [24, 35, 36, 34].

Another set of early approaches use a fixed data representation with a *Nearest Class Mean* (NCM) classifier for incremental learning [31, 32, 41]. NCM classifier computes a single centroid for each class as the mean of all the feature

vectors of the images in the training set for each class. To predict the label for a test image, NCM assigns it the class label of the closest centroid. NCM avoids catastrophic forgetting by using centroids. Each class centroid is computed using only the training data of that class, hence even if the classes are learned in an incremental fashion the centroids for previous classes are not affected when new classes are learned. These early approaches, however, use SIFT features [28], hence their classification accuracy is not comparable to the current deep learning approaches, as shown in [38].

2.2. Deep Learning Methods

Deep learning methods have produced excellent results on many vision tasks because of their ability to jointly learn task-specific features and classifiers [5, 15, 27, 45]. However, deep learning approaches suffer from catastrophic forgetting on incremental learning tasks. Essentially, classification accuracy rapidly decreases when learning new classes [2, 13, 16, 22, 25, 30]. Various approaches have been proposed recently to deal with catastrophic forgetting for task-incremental and class-incremental learning [1, 38].

For task-incremental learning, a model is trained incrementally on different datasets and during evaluation it is tested on the different datasets separately [7]. Task-incremental learning utilizes multi-headed evaluation which is characterized by predicting the class when the task is known, which has been shown to be a much easier problem in [7] than the class-incremental learning considered in this paper.

2.2.1 Class-Incremental Learning Methods

Most of the recent class-incremental learning methods rely on storing a fraction of old class data when learning a new class [38, 19, 6, 48, 7]. iCaRL [38] combines knowledge distillation [18] and NCM for class-incremental learning. Knowledge distillation uses a distillation loss term that forces the labels of the training data of previously learned classes to remain the same when learning new classes. iCaRL uses the old class data while learning a representation for new classes and uses the NCM classifier for classification of the old and new classes. EEIL [6] improves iCaRL with an end-to-end learning approach. Hou et al. [19] uses cosine normalization, less-forget constraint and inter-class separation for reducing the data imbalance between old and new classes. The main issue with these approaches is the need to store old class data which is not practical when the memory budget is limited. To the best of our knowledge, there are only two approaches that do not use old class data and use a fixed memory budget: LWF-MC [38] and LWM [11]. LWF-MC is simply the implementation of LWF [26] for class-incremental learning. LWM uses at-

tention distillation loss and a teacher model trained on old class data for better performance than LWF-MC. Although both of these approaches meet the conditions for class-incremental learning proposed in [38], their performance is inferior to approaches that store old class data [38, 6, 48].

An alternative set of approaches increase the number of layers in the network for learning new classes [44, 46]. Another novel approach is presented in [49] which grows a tree structure to incorporate new classes incrementally. These approaches also have the drawback of rapid increase in memory usage as new classes are added.

Some researchers have also focused on using a deep network pre-trained on ImageNet as a fixed feature extractor for incremental learning. Belouadah et al. [4] uses a pre-trained network for feature extraction and then trains shallow networks for classification while incrementally learning classes. They also store a portion of old class data. The main issue with their approach is that they test their approach on the ImageNet dataset using the feature extractor that has already been trained on ImageNet which skews their results. FearNet [21] uses a ResNet-50 pre-trained on ImageNet for feature extraction and uses a brain-inspired dual memory system which requires storage of the feature vectors and co-variance matrices for the old class images. The feature vectors and co-variance matrices are further used for generating augmented data during learning. Our approach does not store any base class data or use any data augmentation, although it uses a ResNet pre-trained on ImageNet for feature extraction but we do not test our approach on ImageNet.

3. Methodology

Following the notation from [38], CBCL learns from a class-incremental data stream of sample sets X^1, X^2, \dots, X^N in which all samples from the set $X^y = \{x_1^y, \dots, x_{N_y}^y\}$ are from the class $y \in N$ with N_y samples.

The subsections below, first explain our method for class-incremental learning. Next, we explore how the memory footprint can be managed by restricting the total number of centroids. Finally, we demonstrate the use of our approach to incrementally learn using only a few examples per class.

3.1. Agg-Var Clustering

The complete architecture of our approach is depicted in Figure 1. Once the data for a new class becomes available, the first step in CBCL is the generation of feature vectors from the images of the new class using a fixed feature extractor. The proposed architecture can work with any type of image feature extractor or even for non-image datasets with appropriate feature extractors. In this paper, for the task of object-centric image classification, we use CNNs

(ResNet [17]) pre-trained on ImageNet [43] as feature extractors.

In the learning phase, for each new image class $1 \leq y \leq N$, *Agg-Var* clustering [3] is applied on the feature vectors of all the training images in the class $\{x_1^y, x_2^y, \dots, x_{N_y}^y\}$. In the hippocampal concept learning model [29, 40, 33], after the feature extraction step, the hippocampus calculates a term called the *memory-based prediction error*. This value represents the difference from the incoming episode to all of the previously experienced concepts. This step is replicated in *Agg-Var* clustering by finding the Euclidean distance between the incoming image to each centroid for a class. Initially there are no centroids for a new class y . Hence, this step begins by creating a centroid from the first image of class y . Next, for each image in training set of the class, feature vector x_i^y (for the i th image) is generated and compared using the Euclidean distance to all the centroids for the class y . If the distance of x_i^y to the closest centroid is below a pre-defined distance threshold D , the closest centroid is updated by calculating a weighted mean of the centroid and the feature vector x_i^y :

$$C_{new} = \frac{w_C \times C_{old} + x_i^y}{w_C + 1} \quad (1)$$

where, C_{new} is the updated centroid, C_{old} is the centroid before the update, w_C is the number of data points (images) already represented by the centroid. This step of *Agg-Var* clustering is meant to capture *memory integration* in the concept learning process of the hippocampus. Memory integration occurs when the *memory-based prediction error* of an episode to a previous concept is small. If, on the other hand, the *memory-based prediction error* of an episode to a previous concept is large, according to the concept learning process of the hippocampus, *pattern separation* occurs resulting in the creation of a new distinct concept based on the incoming episode. *Agg-Var* clustering captures this aspect of the process as: if the distance between the i th image and the closest centroid is higher than the distance threshold D , a new centroid is created for class y and equated to the feature vector x_i^y of the i th image. The result of this process is a collection containing a set of centroids for the class y , $C^y = \{c_1^y, \dots, c_{N_y^*}^y\}$, where N_y^* is the number of centroids for class y . This process is applied to the sample set X^y of each class incrementally once they become available to get a collection of centroids $C = C^1, C^2, \dots, C^N$ for all N classes in a dataset. It should be noted that using the same distance threshold for different classes can yield different number of centroids per class depending on the similarity among the images (intra-class variance) in each class. Hence, we only need to tune a single parameter (D) to get the optimal number of centroids that yield best validation accuracy in each class. Note that our approach calculates the centroids for each class separately. Thus, the perfor-

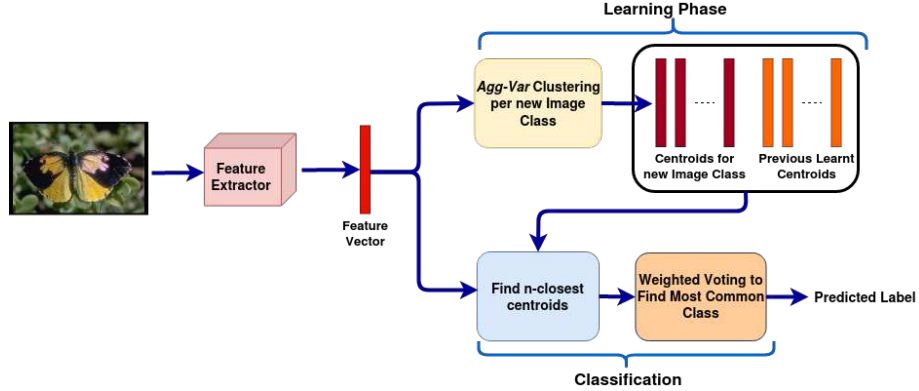


Figure 1: For each new image class in a dataset, the feature extractor generates the CNN features of all the training images in the image class and generates a set of centroids using *Agg-Var* clustering algorithm, concatenates them with the centroids of previously learned classes and uses the complete set of centroids for classifying unlabeled test images

mance of our approach is not strongly impacted when the classes are presented incrementally.

3.2. Weighted-Voting Scheme for Classification

To predict the label y^* of a test image we use the feature extractor to generate a feature vector x . Next, Euclidean distance is calculated between x and the centroids of all the classes observed so far. Based on the calculated distances, we select n closest centroids to the unlabeled image. The contribution of each of the n closest centroids to the determination of the test image’s class is a conditional summation:

$$Pred(y) = \sum_{j=1}^n \frac{1}{dist(x, c_j)} [y_j = y] \quad (2)$$

where $Pred(y)$ is the prediction weight of class y , y_j is category label of j th closest centroid c_j and $dist(x, c_j)$ is the euclidean distance between c_j and the feature vector x of the test image. The prediction weights for all the image classes observed so far are first initialized to zero. Then, for the n closest centroids the prediction weights are updated, using equation (2), for the classes that each of the n centroids belong to. The prediction weight for each class is further multiplied by the inverse of the total number of images in the training set of the class to manage class imbalance. Since classes with more training data most likely have more centroids than other classes, prediction weight can become biased towards such classes. The proposed weighting scheme avoid bias towards such classes during prediction:

$$Pred(y)' = \frac{1}{N_y} Pred(y) \quad (3)$$

where $Pred(y)'$ is the prediction weight of class y after multiplication with the inverse of total number of training

images of the class N_y with the previous prediction weight $Pred(y)$ of the class. The test image is assigned the class label with the highest prediction weight $Pred(y)'$.

3.3. Centroid Reduction

The memory footprint is an important consideration for an incremental learning algorithm [38]. Real system implementations have limited memory available. We therefore propose a novel method that restricts the number of centroids while attempting to maintain classification accuracy.

If we assume that a system can store a maximum of K centroids and that currently the system has stored K_t centroids for t classes. For the next batch of classes the system needs to store K_{new} more centroids but the total number of centroids $K_t + K_{new} > K$. Hence, the system needs to reduce the total stored centroids to $K_r = K_t + K_{new} - K$ centroids. Rather than reducing the number of centroids for each class equally, CBCL reduces the centroids for each class based upon the previous number of centroids in the class. The reduction in the number of centroids N_y^* for each class y is calculated as (whole number):

$$N_y^*(new) = N_y^* \left(1 - \frac{K_r}{K_t}\right) \quad (4)$$

where $N_y^*(new)$ is the number of centroids for class y after reduction. Rather than simply removing the extra centroids from each class, we cluster the closest centroids in each class to get new centroids, keeping as much information as possible about the previous classes. This process is accomplished by applying k-means clustering [20] on the centroid set C^y of each class y to cluster them into a total of $N_y^*(new)$ centroids. Results on benchmark datasets show the effectiveness of our centroid reduction approach (Section 4).

Dataset	CIFAR-100	CUBS-200-2011	Caltech-101
# classes	100	100	100
# training images	500	80% of data	80% of data
# testing images	100	20% of data	20% of data
# classes/batch	2, 5, 10, 20	10	10

Table 1: Statistical details of the datasets in our experiments, same as in [11] for a fair comparison. Number of training and test images reported are for each class in the dataset.

3.4. Few-Shot Incremental Learning (FSIL)

For a traditional few-shot learning problem, an algorithm is evaluated on n -shots, k -way tasks. Hence, a model is given a total of n examples per class for k classes for training. After the training phase, the model is evaluated on a small number of test samples (usually 15 test samples for 1-shot, 5-shot and 10-shot learning) for each of the k classes. Some few-shot learning approaches have been proposed in which the model is tested on the new k classes and the base classes [14, 37, 39]. However, these approaches are not suitable for learning classes incrementally for a large number of increments using only a few samples per class. The few-shot incremental learning setting proposed here deals with this problem.

For an n -shot incremental learning setting, we propose to train a model on n examples per class for k classes in an increment. The training data for the l previously learned classes is not available to the model during the current increment. After training, the model is tested on the complete test set for all the classes learned so far ($k+l$). Although this problem becomes more difficult with each increment, we show that our approach performs well even for 5-shot and 10-shot incremental learning cases (Section 4) because even a limited number of instances per class generate centroids covering most of the class’s concept. FSIL is potentially important for applications where labeled data is difficult to obtain, perhaps such as a human incrementally teaching a robot. In such cases, the human is unlikely to be willing to provide more than few examples of a class.

4. Experiments

We evaluate CBCL on three standard class-incremental learning datasets: Caltech-101 [12], CUBS-200-2011 [47] and CIFAR-100 [23]. First, we present the datasets and the implementation details. CBCL is then compared to state-of-the-art methods for class-incremental learning and evaluated on 5-shot and 10-shot incremental learning. Finally, we perform an ablation study to analyze the contribution of each component of our approach.

4.1. Datasets

CBCL was evaluated on the three datasets used in [11]. LWM [11] was also tested on iLSVRC-small(ImageNet) dataset but since our feature extractor is pre-trained on ImageNet, comparing on this dataset would not be a fair comparison. Caltech-101 contains 8,677 images of 101 object categories with 40 to 800 images per category. CUBS-200-2011 contains 11,788 images of 200 categories of birds. CIFAR-100 consists of 60,000 32×32 images belonging to 100 object classes. There are 500 training images and 100 test images for each class. The number of classes, train/test split size and number of classes per batch used for training are described in Table 1. The classes that compose a batch were randomly selected. For the 5-shot and 10-shot incremental learning experiments, only the training images per class were changed to 5 and 10, respectively in Table 1 keeping the other statistics the same.

Similar to [11, 38], top-1 accuracy was used for evaluation. We also report the *average incremental accuracy*, which is the average of the classification accuracies achieved in all the increments [38].

Because CBCL’s learning time is much shorter than the time required to train a neural network, we are able to run all our experiments 10 times randomizing the order of the classes. We report the average classification accuracy and standard deviation over these ten runs.

4.2. Implementation Details

The Keras deep learning framework [9] was used to implement all of the neural network models. For Caltech-101 and CUBS-200-2011 datasets, the ResNet-18 [17] model pre-trained on the ImageNet [43] dataset was used and for CIFAR-100 the ResNet-34 [17] model pre-trained on ImageNet was used for feature extraction. These model architectures are consistent with [11] for a fair comparison. For the experiment with the CIFAR-100 dataset the model was allowed to store up to $K = 7500$ centroids requiring 3.87 MB versus 84 MB for an extra ResNet-34 teacher model as in [11]. Furthermore, compared to methods that store only 2000 images for previous classes [38, 6, 4], 7500 centroids for our approach require less memory (3.87 MB) than 2000 complete images (17.6 MB). For Caltech-101 $K = 1100$ centroids were stored (0.5676 MB versus 45 MB for a ResNet-18 teacher model as in [11]) and for CUBS-200-2011 $K = 800$ centroids were stored (0.4128 MB).

As mentioned in Section 1, none of the prior incremental learning techniques are suitable for FSIL because they require a large amount of training data per class. Hence, we compare CBCL against a few-shot learning baseline (FLB). FLB uses the features from the pre-trained ResNet neural network which are passed on to a linear layer which is trained with softmax loss (Figure 3). This procedure follows prior work on few-shot learning research [8], which in-




Figure 2: Average and standard deviation of classification accuracies (%) on CIFAR-100 dataset with (a) 2, (b) 5, (c) 10, (d) 20 classes per increment with 10 executions. Average incremental accuracies are shown in parenthesis. (For other methods, results are reported from the respective papers and different papers reported results on different increment settings. Best viewed in color)



Figure 3: Few-shot Learning Baseline (FLB) architecture

Methods	k-Shot	Classes per increment			
		2	5	10	20
FLB	5	41.1	39.9	41.3	44.4
CBCL	5	56.9	55.6	54.7	53.8
FLB	10	53.5	52.4	55.1	55.6
CBCL	10	61.9	61.4	61.3	60.7

Table 2: Comparison of CBCL with FLB on 5-shot and 10-shot incremental learning settings in terms of average incremental accuracy (%) on CIFAR-100 dataset with 2, 5, 10 and 20 classes per increment.

indicates that FLB is better than many other few-shot learning techniques that use a deeper backbone, such as ResNet-18 or ResNet-34. Since FLB is not suitable for few-shot incremental learning, we train the final linear layer of FLB with softmax loss using the complete training set of all the new and old class data in each increment. In other words, FLB does not learn incrementally. FLB was trained for 25 epochs in each increment using a fixed learning rate of 0.001 and cross-entropy loss with minibatches of size 8 optimized using stochastic gradient descent.

For CBCL, for each batch of new classes, the hyper-parameters D (distance threshold) and n (number of closest centroids used for classification) are tuned using cross-validation. We only use the previously learned centroids and the training data of the new classes for hyper-parameter tuning.

4.3. Results on CIFAR-100 Dataset

On CIFAR-100 dataset, our method is compared to seven different methods: *finetuning* (FT), LWM [11], LWF-MC [38], iCaRL [38], EEIL [6], BiC [48] and FearNet [21]¹. FT simply uses the network trained on previous classes and adapts it to the new incoming classes. LWM extends LWF [26] and uses attention distillation loss for class-incremental learning. LWF-MC uses distillation loss during the training phase. iCaRL also uses the distillation loss for representation learning but stores exemplars of previous classes and uses the NCM classifier for classification. EEIL improves iCaRL by offering an end-to-end learning approach which also uses the distillation loss and keeps exemplars from the old classes. BiC also uses the exemplars

from the old classes and adds a bias correction layer after the fully connected layer of the ResNet to correct for the bias towards the new classes. We also compare the classification accuracy after learning all the classes to an upper bound (68.6%) consisting of ResNet34 trained on the entire CIFAR-100 dataset in one batch.

Figure 2 compares CBCL to first six out of the seven methods mentioned above with 2, 5, 10 and 20 classes per increment. Even though a fair comparison of CBCL is only possible with FT, LWF-MC and LWM, since they are the only approaches that do not require storing the exemplars of the old classes, it outperforms all six methods on all increment settings. The difference in classification accuracy between CBCL and these other methods increases as the number of classes learned increases. Moreover, for smaller increments the difference in accuracy is larger. Unlike other methods, CBCL’s performance remains the same regardless of the number of classes in each increment (final accuracy after 100 classes for all increments is **60%**).

Table 2 compares CBCL with FLB for 5-shot and 10-shot incremental learning in terms of average incremental accuracy. CBCL beats FLB on both 5-shot and 10-shot incremental learning for all four incremental settings with significant margins. It should be noted that FLB uses the training set of all the old and new classes in each increment while CBCL uses the training set of new classes only. Fur-

¹Comparison with FearNet is in Supplementary File

# Classes	FT	LWM	CBCL (Ours)
10 (base)	97.78	97.78	97.75 ± 1.23
20	59.55	75.34	95.42 ± 2.00
30	52.65	71.78	92.58 ± 2.08
40	44.51	67.49	91.40 ± 1.77
50	35.52	59.79	90.23 ± 1.45
60	31.18	56.62	89.34 ± 1.57
70	32.99	54.62	88.37 ± 1.57
80	27.45	48.71	87.57 ± 1.07
90	28.55	46.21	86.99 ± 0.95
100	28.26	48.42	86.44 ± 0.70

Table 3: Comparison with FT and LWM [11] on Caltech-101 dataset in terms of classification accuracy (%) with 10 classes per increment. Average and standard deviation of classification accuracies per increment are reported

ther, the difference in accuracy between CBCL and FLB is higher when using 5 examples per class. Also, the difference is higher when using smaller number of classes per increment. This may suggest that CBCL is best suited for incremental learning situations when data is scarce.

Comparing the average incremental accuracies of the other six methods (Figure 2), which use the complete training set per class, even with only 5 or 10 training examples per class CBCL outperforms other methods that do not store class data (FT, LWF-MC, LWM) and is only slightly inferior to methods that do store old class data (2, 5, 10 classes per increment). For 20 classes per increment, CBCL is slightly inferior to the other methods when using 5 and 10 examples per class. For further comparison, when ResNet-34 was trained on 5-shot and 10-shot settings in a single batch it yielded only 8.22% and 12.15% accuracies, respectively. These results clearly show that CBCL offers excellent performance on few-shot incremental learning for object classification.

4.4. Results on Caltech-101 Dataset

For the Caltech-101 dataset CBCL was compared to *fine-tuning* (FT) and LWM [11] with learning increments of 10 classes per batch (Table 3). FT and LWM were introduced in Subsection 4.3. CBCL outperforms FT and LWM by a significant margin. The difference between CBCL and LWM and FT continues to increase as more classes are learned. FT performs the worst with a classification accuracy after 100 classes that is about one fourth of the base accuracy (decreases by 69.52%). LWM is an improvement compared to FT. Nevertheless, the accuracy after 100 classes is almost the half of the base accuracy (decreases by 49.36%). For CBCL the decrease in accuracy is only **11.31%** after incrementally learning 100 classes. The average incremental accuracies for FT, LWM, CBCL, CBCL on 5-shot incremental learning and CBCL on 10-shot incremental

# Classes	FT	LWM	CBCL (Ours)
10 (base)	99.17	99.17	92.83 ± 1.78
20	57.92	78.75	84.22 ± 1.62
30	41.11	70.83	74.55 ± 3.06
40	35.42	58.54	70.73 ± 3.26
50	32.33	53.67	66.06 ± 3.32
60	29.03	47.92	62.26 ± 2.79
70	22.14	43.79	59.58 ± 2.54
80	22.27	43.83	57.44 ± 2.58
90	20.52	39.85	56.37 ± 2.28
100	17.4	34.52	54.75 ± 2.66

Table 4: Comparison with FT and LWM [11] on CUBS-200-2011 dataset in terms of classification accuracy (%) with 10 classes per increment. Average and standard deviation of classification accuracies per increment are reported

cremental learning are 43.84%, 62.67%, **90.61%**, **87.70%** and **89.92%**, respectively. Hence, CBCL improves accuracy over the current best method (LWM) by a margin of **27.94%** in terms of average incremental accuracy when the complete training set is used. Even for the 5-shot and 10-shot incremental learning, CBCL outperforms LWM by margins of **25.03%** and **27.25%**, respectively.

We also compare CBCL for 5-shot and 10-shot incremental learning against FLB trained on all the classes data in each increment (batch learning). FLB achieves 72.48% and 83.81% average incremental accuracies for 5-shot and 10-shot incremental learning settings, respectively, which are significantly inferior (**15.22%** and **6.11%**) to CBCL. These results are in accordance with CIFAR-100 results.

4.5. Results on CUBS-200-2011 Dataset

For the CUBS-200-2011 dataset we again compare our approach to FT and LWM with learning increments of 10 classes per batch (Table 4). The classification accuracy of CBCL is greater than FT and LWM after the 10 classes (base). As the learning increments increase the performance margin also increases. The accuracy of FT decreases by 81.77% after 10 increments and LWM’s accuracy decreases by 64.65%. The decrease in classification accuracy of CBCL after 10 increments is **38.0%** lower than both of these approaches. The average incremental accuracies for FT, LWM, CBCL, CBCL for 5-Shot incremental learning and CBCL for 10-shot incremental learning are 37.7%, 57.0%, **67.8%**, 56.2% and **63.8%** respectively. CBCL is an improvement over LWM by a **10.7%** margin in terms of average incremental accuracy. Furthermore, even for 10-shot incremental learning setting CBCL improves over LWM and it is slightly below LWM for 5-shot incremental learning.

Similar to CIFAR-100 and Caltech-101, we compare CBCL on 5-shot and 10-shot incremental learning against

FLB trained on all the classes data in each increment (batch learning). FLB achieves 37.48% and 55.00% average incremental accuracies for 5-shot and 10-shot incremental learning settings, respectively, which are inferior (**18.68%** and **8.8%**) to CBCL. These results are in accordance with CIFAR-100 and Caltech-101 FSIL results.

4.6. Ablation Study

We performed an ablation study to examine the contribution of each component in our approach to the overall system’s accuracy. This set of experiments was performed on CIFAR-100 dataset with increments of 10 classes and memory budget of $K = 7500$ centroids using all the training data per class. We report average incremental accuracy for these experiments.

This ablation study investigates the effect of the following components: feature extractor, clustering approach, number of centroids used for classification, and the impact of centroid reduction. Hybrid versions of CBCL are created to ablate each of these different components. Hybrid-1 termed *VGG-16* uses a VGG-16 pre-trained on ImageNet as a feature extractor. Hybrid-2 termed *Trad-Agg* uses traditional agglomerative clustering and hybrid-3 termed *K-means* uses k-means clustering to generate centroids for all the image classes. Hybrid-4 termed *Single-Centroid-Pred* uses only a single closest centroid for classification (same as NCM classifier). Hybrid-5 termed *Remove-Centroids* simply removes the extra centroids when the memory limit is reached rather than using the proposed centroid reduction technique. Lastly, hybrid-6 termed *NCM* uses an NCM classifier with the ImageNet pre-trained feature extractor. Except for the changed component, all the other components in the hybrid approaches are the same as CBCL.

Table 5 shows the results for the ablation study. All of the hybrid methods are less accurate than the complete CBCL algorithm. *VGG-16* hybrid achieve slightly lower accuracy than CBCL with ResNet-34, depicting the robustness of our method against the choice of the feature extractor. *Trad-Agg* and *K-means* achieve similar average incremental accuracy but is significantly inferior when compared to CBCL. This difference in accuracy reflects the effectiveness of the *Agg-Var* clustering algorithm for object-centric image classification. *Single-Centroid-Pred* achieves slightly lower accuracy than CBCL, illustrating that the accuracy gain resulting from using multiple centroids for classification is about 1.15%. Finally, the *Remove-Centroid* hybrid’s accuracy is the closest to CBCL’s. This small difference may reflect the fact that the memory budget is large enough such that the algorithm does not need to reduce centroids until the last increment. Hence, only in the last increment is there a slight change, which does not effect the average incremental accuracy for all 10 increments by a significant margin. The effectiveness of our centroid reduction technique is more

Methods	Accuracy (%)
VGG-16	68.5
Trad-Agg	59.2
K-means	60.0
Single-Centroid-Pred	68.7
Remove-Centroids	69.0
NCM	58.25
CBCL	69.85

Table 5: Effect on *average incremental accuracy* by switching off each component separately in CBCL. All of the hybrids show lower performance than CBCL demonstrating each of their contribution to get the best results using CBCL

apparent when using smaller memory budgets. For example, for $K=3000$ centroids limit the average incremental accuracies for CBCL and *Remove-Centroids* are **67.5%** and 64.0%, respectively, depicting the effectiveness of our proposed centroid reduction technique. Lastly, for NCM, we again see a drastic decrease in accuracy because this hybrid uses a single centroid to represent each class. This ablation study clearly indicates that the most important component of CBCL is the cognitively-inspired *Agg-Var* clustering approach, based upon drastic decrease in performance for *Trad-Agg*, *K-means* and *NCM* hybrids. Note that the average incremental accuracy for all the other hybrids (*ResNet-18*, *VGG-16*, *Single-Centroid-Pred* and *Remove-Centroids*) is also higher than the state-of-the-art methods.

5. Conclusion

In this paper we have proposed a novel cognitively-inspired approach (CBCL) for class-incremental learning which does not store previous class data. The centroid-based representation of different classes not only produces the state-of-the-art results but also opens up novel avenues of future research, like few-shot incremental learning. Although CBCL offers superior accuracy to other incremental learners, its accuracy is still lower than single batch learning on the entire training set. Future versions of CBCL will seek to match the accuracy of single batch learning. Although, for FSIL CBCL beats the batch learning baseline.

CBCL contributes methods that may one day allow for real-world incremental learning from a human to an artificial system. Few-shot incremental learning, in particular, holds promise as a method by which humans could conceivably teach robots about important task-related objects. Our upcoming work will focus on this problem.

Acknowledgments

This work was supported by Air Force Office of Scientific Research contract FA9550-17-1-0017.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *The European Conference on Computer Vision (ECCV)*, September 2018. 2
- [2] Bernard Ans, Stéphane Rousset, Robert M. French, and Serban Musca. Self-refreshing memory in artificial neural networks: learning temporal sequences without catastrophic forgetting. *Connection Science*, 16(2):71–99, 2004. 2
- [3] Ali Ayub and Alan Wagner. CBCL: Brain inspired model for RGB-D indoor scene classification. *arXiv:1911.00155*, 2019. 2, 3
- [4] Eden Belouadah and Adrian Popescu. Deesil: Deep-shallow incremental learning. In *The European Conference on Computer Vision (ECCV) Workshops*, September 2018. 3, 5
- [5] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, Aug 2013. 2
- [6] Francisco M. Castro, Manuel J. Marin-Jimenez, Nicolas Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *The European Conference on Computer Vision (ECCV)*, September 2018. 1, 2, 3, 5, 6
- [7] Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *The European Conference on Computer Vision (ECCV)*, September 2018. 1, 2
- [8] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019. 5
- [9] François Chollet et al. Keras. <https://github.com/fchollet/keras>, 2015. 5
- [10] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Mach. Learn.*, 20(3):273–297, Sept. 1995. 2
- [11] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2, 5, 6, 7
- [12] Li Fei-Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):594–611, April 2006. 2, 5
- [13] Robert M. French. Dynamically constraining connectionist networks to produce distributed, orthogonal representations to reduce catastrophic interference. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pages 335–340, 2019. 1, 2
- [14] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 5
- [15] Ross Girshick. Fast r-cnn. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015. 1, 2
- [16] Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv:1312.6211*, 2013. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 1, 3, 5
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. 2
- [19] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2
- [20] Anil K Jain, M Narasimha Murty, and Patrick J Flynn. Data clustering: a review. *ACM computing surveys (CSUR)*, 31(3):264–323, 1999. 4
- [21] Ronald Kemker and Christopher Kanan. Fearnert: Brain-inspired model for incremental learning. In *International Conference on Learning Representations*, 2018. 3, 6
- [22] James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3521–3526, 2017. 1, 2
- [23] Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009. Technical report, University of Toronto. 2, 5
- [24] I. Kuzborskij, F. Orabona, and B. Caputo. From n to n+1: Multiclass transfer incremental learning. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3358–3365, June 2013. 2
- [25] Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4652–4662. Curran Associates, Inc., 2017. 2
- [26] Z. Li and D. Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, Dec 2018. 2, 6
- [27] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 1, 2
- [28] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, Nov. 2004. 2
- [29] Michael L. Mack, Bradley C. Love, and Alison R. Preston. Building concepts one episode at a time: The hippocampus and concept formation. *Neuroscience Letters*, 680:31–38, 2018. 1, 3
- [30] Michael McCloskey and Neil J. Cohen. Catastrophic interference in connectionist networks: The sequential learn-

- ing problem. *The Psychology of Learning and Motivation*, 24:104–169, 1989. 2
- [31] T. Mensink, J. Verbeek, F. Perronnin, and G. Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2624–2637, Nov 2013. 2
- [32] Thomas Mensink, Jakob J. Verbeek, Florent Perronnin, and Gabriela Csurka. Metric learning for large scale image classification: Generalizing to new classes at near-zero cost. In *ECCV*, 2012. 2
- [33] Morris Moscovitch, Roberto Cabeza, Gordon Winocur, and Lynn Nadel. Episodic memory and beyond: The hippocampus and neocortex in transformation. *Annual Review of Psychology*, 67(1):105–134, Apr 2016. 1, 3
- [34] M. D. Muhlbaier, A. Topalis, and R. Polikar. Learn⁺⁺.nc: Combining ensemble of classifiers with dynamically weighted consult-and-vote for efficient incremental learning of new classes. *IEEE Transactions on Neural Networks*, 20(1):152–168, Jan 2009. 2
- [35] Anastasia Pentina, Viktoriia Sharmanska, and Christoph H. Lampert. Curriculum learning of multiple tasks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 2
- [36] R. Polikar, L. Upda, S. S. Upda, and V. Honavar. Learn⁺⁺: an incremental learning algorithm for supervised neural networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 31(4):497–508, Nov 2001. 2
- [37] Hang Qi, Matthew Brown, and David G. Lowe. Low-shot learning with imprinted weights. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 5
- [38] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. iCaRL: Incremental classifier and representation learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1, 2, 3, 4, 5, 6
- [39] Mengye Ren, Renjie Liao, Ethan Fetaya, and Richard Zemel. Incremental few-shot learning with attention attractor networks. In *Advances in Neural Information Processing Systems 32*, pages 5275–5285, 2019. 5
- [40] Louis Renoult, Patrick S. R. Davidson, Erika Schmitz, Lillian Park, Kenneth Campbell, Morris Moscovitch, and Brian Levine. Autobiographically significant concepts: More episodic than semantic in nature? an electrophysiological investigation of overlapping types of memory. *Journal of Cognitive Neuroscience*, 27(1):57–72, 2015. 1, 3
- [41] Marko Ristin, Matthieu Guillaumin, Juergen Gall, and Luc Van Gool. Incremental learning of ncm forests for large-scale image classification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 2
- [42] S. Ruping. Incremental learning with support vector machines. In *Proceedings 2001 IEEE International Conference on Data Mining*, pages 641–642, Nov 2001. 2
- [43] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vision*, 115(3):211–252, Dec. 2015. 2, 3, 5
- [44] Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv:1606.04671*, 2016. 3
- [45] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1, NIPS’14*, pages 568–576, Cambridge, MA, USA, 2014. MIT Press. 2
- [46] Alexander V. Terekhov, Guglielmo Montone, and J. Kevin O’Regan. Knowledge transfer in deep block-modular neural networks. In *Proceedings of the 4th International Conference on Biomimetic and Biohybrid Systems - Volume 9222, Living Machines 2015*, pages 268–279, New York, NY, USA, 2015. Springer-Verlag New York, Inc. 3
- [47] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge J. Belongie. The caltech-ucsd birds-200-2011 dataset, 2011. Technical Report CNS-TR-2011-001, California Institute of Technology. 2, 5
- [48] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2, 3, 6
- [49] Tianjun Xiao, Jiaying Zhang, Kuiyuan Yang, Yuxin Peng, and Zheng Zhang. Error-driven incremental learning in deep convolutional neural network for large-scale image classification. In *Proceedings of the 22nd ACM International Conference on Multimedia, MM ’14*, pages 177–186, New York, NY, USA, 2014. ACM. 3