

This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Do not Forget to Attend to Uncertainty while Mitigating Catastrophic Forgetting

Vinod K Kurmi IIT Kanpur vinodkk@iitk.ac.in Badri N Patro* IIT Kanpur badri@iitk.ac.in Venkatesh K Subramanian IIT Kanpur venkats@iitk.ac.in Vinay P Namboodiri University of Bath vpn22@bath.ac.uk

Abstract

One of the major limitations of deep learning models is that they face catastrophic forgetting in an incremental learning scenario. There have been several approaches proposed to tackle the problem of incremental learning. Most of these methods are based on knowledge distillation and do not adequately utilize the information provided by older task models, such as uncertainty estimation in predictions. The predictive uncertainty provides the distributional information can be applied to mitigate catastrophic forgetting in a deep learning framework. In the proposed work, we consider a Bayesian formulation to obtain the data and model uncertainties. We also incorporate selfattention framework to address the incremental learning problem. We define distillation losses in terms of aleatoric uncertainty and self-attention. In the proposed work, we investigate different ablation analyses on these losses. Furthermore, we are able to obtain better results in terms of accuracy on standard benchmarks.

1. Introduction

Deep neural networks have shown impressive results on a variety of computer vision tasks. However, these models have been observed to be not well suited for generalization in terms of tasks. This has been particularly observed, for instance, in seminal work by Kirkpatrick et al. [16] that the deep learning models can face catastrophic forgetting. Learning a model on a task 'A' results in an optimal model being obtained for task 'A'. However, using this same model to solve for task 'B' results in the model converging to new optima that are significantly worse for solving task 'A'. This is in contrast to humans, where we are able to learn new tasks without forgetting. In order to obtain deep learning networks that can continually learn without forgetting, we need to solve this problem. There have been a number of methods proposed towards solving this problem, such as those based on regularization [16, 45] or those based on memory replay [41, 1]. However, the methods proposed are not focused on the uncertainty associated with the distillation. One approach that allows us to understand the working of the method is by considering the uncertainty distribution. Deep learning models usually consider the point estimate for predicting instead of adopting a probabilistic distribution-based approach. In the probabilistic approach, a fully Bayesian treatment is intractable in general. However, work by Gal and Ghahramani [8] and other related subsequent works [14, 15] suggest that it is possible to obtain the uncertainty of a model in its parameters or the data distribution during prediction. We adopt this as it is particularly relevant in the case of incremental learning. By incorporating uncertainty, we can ensure that the model will initially be uncertain about the new task, and as it is trained, its uncertainty reduces. Moreover, we consider the uncertainty based distillation loss between tasks that ensure the method distills the uncertainty along with the prediction.

Another approach that allows for interpretability is considering the attention and visual explanation based models. Attention-based methods have been very popular in machine translation [40], visual question answer [43], and generative adversarial networks (GANs) [46]. In our work, we consider self-attention based methods [46] for incremental learning. We particularly consider the distillation loss between tasks that ensure the method attends to similar discriminative regions. Moreover, obtaining the attentionbased distillation can result in improved incremental learning techniques being developed.

Our work is based on a baseline that is a variant of learning without forgetting (LwF) [25]. Particularly we propose a variant that incorporates a distillation loss between the previous task and the new task. However, all the samples used for training the previous task would not be available while training for a new task. Therefore, we consider a method that retains a small amount of data from the previous task and uses distillation loss to retain the knowledge while learning a new task [34].

In the data-based incremental learning setting, previous works have proposed different distillation based models to preserve the old task information while training to solve for

^{*}Currently working at Google

the current task. The main problem with these methods is that they distill the knowledge of previous data from the previous model's softmax output for a given input of current task data. Since the predictions of deep learning models are overconfident for any prediction, it means that any single class has probability near one, while other classes have almost zero probability: even when the prediction is wrong [27]. To solve this problem, it has been suggested that high temperatures can be used [12] for calculating the distillation. However, this does not truly reflect the uncertainty in prediction. In incremental learning, the model distills knowledge from the next task, for which it is not trained. Thus it should not be overconfident for one class; instead, the prediction should be distributed among all the classes. There will be uncertainty present in the prediction. It means this distillation can not truly reflect the older class information properly. So the softmax based distillation loss is not sufficient to preserve the previous task knowledge in the current model. Thus, in the proposed method, we use uncertainty-based distillation to reflect information about the older task data.

Another problem with incremental learning is that not all the information provided by the previous model is useful for preserving the older task data. The attention-based method can further improve the performance of incremental learning. In [5] authors propose a GradCAM [36] based attention distillation for class incremental learning. GradCAM based attention also depends upon the predicted class label of the previous model. In contrast to the predicted label based attention, we propose self-attention based distillation for the incremental learning. The idea of the self-attention module is to adapt the non-local model information using the residual connections.

In the next section, we discuss the related work. The proposed method is presented in detail in section 3. In section 4, we present thorough empirical analysis and results for the various evaluations on standard CIFAR-100 and ImageNet datasets. We finally conclude our paper with directions for future work. The major contributions of the proposed work are as follows:

- We propose an uncertainty and attention based distillation in incremental learning problem.
- We show that by transferring the uncertainty, student model captures more information about the teacher model.
- Attention framework further reduce the catastrophic forgetting.
- We thoroughly analyze the proposed approach through ablation analysis and visualization.

2. Related Work

Many methods have been proposed to solve catastrophic forgetting in the field of incremental or lifelong learning. One of the basic and intuitive methods among them is LwF [25]. For each task, this method uses the same set of feature extractors and task-specific layers for classification. In general, the incremental learning methods can be broadly divided into three categories: i) task-specific methods ii) regularization based methods and iii) methods based on memory replay. The task-specific methods deal with adding a separate model or layer for each new task [28, 25]. Regularization based methods use a regularization term to avoid changing those parameters that are sensitive to old tasks. Elastic weight consolidation (EWC) [16] is a regularization based model that quantifies the importance of weights for previous tasks, and selectively alters the learning rates of weights. Following EWC, the synaptic consolidation strength is measured by [45] in an online fashion and used it as regularization. In Incremental Moment Matching (IMM) [22], a trained network on each task is preserved, and the networks are later merged into one at the end of all the tasks. The above approaches focus on adjusting the network weights. In [31], the authors use the dropout based model selection for incremental learning.

The third category is data-based methods that use the manifold distance between the model trained on old tasks and model trained on the current task. This distance is minimized by suitable loss functions. The most commonly used manifold and loss is the distillation loss [12]. These methods need feature for obtaining the loss function; thus, a small amount of data of previous task is kept. Therefore, the amount of knowledge kept by knowledge distillation depends on the degree of similarity between the data distribution used to learn the previous tasks in the previous stages. Other data-generation based methods [23, 33, 38, 39, 42] use generative adversarial network [9] and by replying the total data, new model is trained. Gradient Episodic Memory (GEM) [26, 34] stores a subset of training data. The End-to-End incremental learning (E2E) framework [3] uses a task-specific distillation loss. In [13], a distillation and retrospection (D+R) used to preserve the previous task information. Recently in [21] the external data along with global distillation (GD) loss is used to improve the incremental learning performance. Another method iCaRL [34] also considers the distillation loss for incremental learning. Other data-based incremental methods are proposed in [24]. Similarly, [35] uses the two teacher models for model-based incremental learning. In [13] authors also use the distillation loss in task incremental learning setting. These methods consider the distillation loss without incorporating the uncertainty value.

Recently some attention based methods have been used in incremental learning [5, 37]. A GradCAM based atten-



Figure 1. Overview of the uncertainty distillation and self-attention in incremental learning. The upper part of the figure represents the model that is trained on the previous tasks and is set to be frozen. G_M is pretrained classifier trained on the previous tasks, G_P , and G_c are the classifiers of the current model for training the old tasks classes and new task classes, respectively.

tion used in [5], while authors in [37] use a attention module to avoid catastrophic forgetting. In [17], a Bayesian framework is discussed for the incremental learning problem. But these methods do not consider uncertainty for preserving the knowledge of previous tasks. Moreover, we differ by considering self-attention based methods for obtaining the attention [40, 46]. This allows us to learn the attention regions that are relevant for the task, unlike the work that uses GradCAM [36] that requires the true class label for obtaining a visualization.

Recently, work by [47] utilizes knowledge distillation to maintain discrimination within old classes to avoid catastrophic forgetting. The adversarial continual learning [7] method learns a disjoint representation for task-invariant and task-specific features. PODNet [6] approaches incremental learning as representation learning, with a distillation loss that constrains the evolution of the representation. iTaML [32] proposes a task agnostic meta-learning approach that seeks to maintain an equilibrium between all the encountered tasks. ScaIL [2] discusses an efficient scaling of past classifiers' weights to make them more comparable to those of new classes in incremental learning scenarios.

3. Method

We propose a class incremental learning framework based on the uncertainty distillation and multi-stage selfattention, where class data come in sequentially in a batch of classes. The model is illustrated in Fig. 1.

3.1. Problem Formulation

Let us assume dataset D consists of pairs of images xand its label y, i.e $D \in (x, y)$. In the proposed framework, we assume there are a total of T tasks i.e. the classification dataset is divided into T parts $D_1, D_2, ..., D_T$. where each $D_t = \{x_j^t, y_j^t\}_{j=1}^{n_t}$. Our goal is to learn a classifier C_T that can correctly classify all the images of all the tasks. Let's assume D_t^{old} contain a fixed number of representative samples from the old dataset, i.e. $D_t^{old} \subset D_1 \cup D_2 \cup ... \cup D_{t-1}$. In the task increment learning problem, at timestamp t, we can access the dataset of the current task and the $(t-1)^{th}$ representative data, i.e. $D_t^{train} = D_t \cup D_t^{old}$.

The classifiers of the model can be divided into three types: i) previous step's classifier (G_m) , parameters θ_m , trained on the previous task(s) data, ii) the current model classifier (G_p) , parameters θ_p , trained on the current task data of previous classes, and iii) current model classifier (G_c) , parameters θ_c , trained on new classes at t^{th} stage. At next $(t+1)^{th}$ stage the models parameters θ_m will be $\theta_p \cup \theta_c$. There are three-stage attention networks in the feature extractor, which has the parameters θ_f . The model is trained using the knowledge distillation from the previous model as well as the cross-entropy loss for the current task data.

3.2. Knowledge Distillation

Knowledge distillation [12] is obtained by minimizing the loss function between the distribution of class probabilities predicted by the teacher and student model. In the deep learning models, the correct class's predicted probability is very high, while other class probabilities are nearly equal to zero. If the ground truth label is provided to the student network, then the teacher model also predicts the same output. Thus the teacher model does not provide any useful information to students. This problem has been tackled using a softmax temperature by Hinton *et al.* [12].

$$p_{c,n} = \frac{\exp(q_{c,n}/\tau)}{\sum_{j=1}^{C} \exp(q_{j,n}/\tau)}$$
(1)

 $q_{c,n}$ is the logit value for $c^{th} \in C$ class for input $x_n \in D$ from the model M_t . τ is the temperature value. Then distillation loss is defined as follows:

$$\mathcal{L}_D(\theta_s, \mathcal{D}, \mathcal{C}) = \frac{1}{|\mathcal{D}|} \sum_{n=1}^N \sum_{c=1}^C -p_{c,n} \cdot \log\left(\hat{y}_{c,n}\right) \quad (2)$$

 $\hat{y}_{c,n}$ is the predicted c^{th} class logit for student network $(M_s(:, \theta_s))$ for input x_n . N is the total number of samples and C is the total number of classes.

3.3. Distillation in Incremental Learning

In the incremental learning setting, suppose the model M_{prev} is trained on the total $|C_{prev}|$ classes and at t^{th} task, there are new $|C_t|$ classes introduced. The objective is to train the current model M_t for all classes i.e $C_{cur} = C_t \cup C_{prev}$. The current model M_t is defined by the feature extractor (M_t^f, θ_f) , old class classifier (G_p, θ_p) and new class classifier (G_c, θ_c) .

For the old task model, distillation loss at t^{th} stage is given as:

$$\mathcal{L}_{\mathrm{D}}(\theta_f, \theta_p; \mathcal{C}_{prev}) = \frac{1}{|D_t^{train}|} \sum_{x \in D_t^{train}} \sum_{y \in \mathcal{C}_{prev}} \left[-p_c \log p(y|x)\right]$$
(3)

 p_c is the softmax output of the previous model, and p(y|x) is the output of the model G_p . Since the previous task model trained on the total $|C_{prev}|$, so distillation is calculated for the old task class logit value. For the current task data, the current model M_t is trained using the cross-entropy loss as given by the equation;

$$\mathcal{L}_{\mathcal{C}}(\theta_f, \theta_c; \mathcal{C}_t) = \frac{1}{|D_t^{train}|} \sum_{x \in D_t^{train}} \sum_{y_c \in \mathcal{C}_t} \left[-y_c \log(\hat{y}_c)\right]$$
(4)

where y_c is the ground truth label and \hat{y}_c is the prediction of G_c for the input x.

3.4. Limitation of Distillation loss

It has been shown that the deep learning models are overconfident for the wrong prediction [27]; thus, the prediction probabilities do not truly reflect the knowledge. In LwM [5], a GradCAM [36] based loss (a visual explanation method) is used for the distillation; it considers the class activation map for distillation. In contrast to the class activation map, the self-attention models learn the attention using the network itself, and an uncertainty estimation predicts the uncertainty associated with it, used to train the model. So the *overconfident* class activation map for given input does not affect the distillation loss.

3.5. Uncertainty based Distillation

In the deep learning models, the class prediction based model is highly probable for one class, and so they cannot transfer the full distributional information to the student model. In the Bayesian neural networks, they can predict the uncertainty value along with their prediction, which can be applicable for transferring more detailed knowledge. Thus, in an incremental learning setting, uncertainty distillation combined with predictions can help the model to preserve useful information from the previous tasks.

3.5.1 Bayesian Model and Uncertainty Estimation

Bayesian modeling is beneficial for predicting uncertainty. In the deep learning models, obtaining the probabilistic inference is an intractable problem. However, [8, 14] proposed a dropout based method to approximate the posterior. We follow a similar method to obtain the Bayesian framework for our incremental learning model. In the proposed method, we work with two types of uncertainties, data uncertainty, also known as aleatoric uncertainty and model uncertainty or epistemic uncertainty. We obtain these uncertainties in a manner similar to [14, 20, 30, 19]. We train the classifier to predicts output class probabilities along with aleatoric uncertainty (data uncertainty). The predictive uncertainty includes epistemic uncertainty and data uncertainty. The epistemic uncertainty results from uncertainty in model parameters. It is obtained by sampling the model using MC-dropout [8, 14]. The estimation of the aleatoric uncertainty in the Bayesian neural network also makes the model more robust. Thus it is better for the incremental learning setting. For obtaining the variance value, we define a variance network G_{cv} for the current task model and G_{pv} for the older task model. The variance network takes the feature for input to predict the variance value. We divide the model M_t into feature extractor G_f and classifier networks G_c and G_p . Suppose x_i is the *i*th input images at current task t, the predicted class logits and variance value are obtain as follows:

$$\hat{y}_i^c = G_c(G_f(x_i^t, \theta_f), \theta_c) \qquad \hat{y}_i^p = G_p(G_f(x_i^t, \theta_f), \theta_p)$$
$$(\sigma_i^c)^2 = G_{cv}(G_f(x_i^t, \theta_f), \theta_{cv}) \tag{5}$$

$$(\sigma_i^p)^2 = G_{pv}(G_f(x_i^t, \theta_f), \theta_{pv})$$
(6)

The aleatoric loss is obtained as:

$$\hat{y}_{i,s}^c = \hat{y}_i^c + \sigma_i^c * \epsilon_s^c, \quad \epsilon_s^c \sim \mathcal{N}(0, I)$$
(7)

$$\hat{y}_{i,s}^p = \hat{y}_i^p + \sigma_i^p * \epsilon_s^p, \quad \epsilon_s^p \sim \mathcal{N}(0, I)$$
(8)

$$\mathcal{L}_{ca} = -\frac{1}{|D_t^{train}|} \sum_{x_i \in D_t^{train}} \log \frac{1}{\mathrm{T}_{\mathrm{a}}} \sum_{s=1}^{\mathrm{T}_{\mathrm{a}}} \mathcal{L}_{\mathcal{C}} \left(\hat{y}_{i,s}^c, y_i \right)$$
(9)

$$\mathcal{L}_{pa} = -\frac{1}{|D_t^{train}|} \sum_{x_i \in D_t^{train}} \log \frac{1}{\mathrm{T}_{\mathrm{a}}} \sum_{s=1}^{\mathrm{T}_{\mathrm{a}}} \mathcal{L}_{\mathcal{D}}\left(\hat{y}_{i,s}^p, p_c\right)$$
(10)

where σ_i^c and σ_i^p are the standard deviation obtained from Eq. 5 and Eq. 6. T_a is the number of Monte Carlo (MC) samples. The models are trained by jointly minimizing both the classification loss and aleatoric loss. The total aleatoric loss is given as:

$$\mathcal{L}_{ale} = \mathcal{L}_{ca} + \mathcal{L}_{pa} \tag{11}$$

3.5.2 Distillation using aleatoric uncertainty

The objective of incremental learning is to preserve the previous task information while training for the current tasks. Thus preserving uncertainty information provided by the Bayesian neural network on the previous task can also overcome the limitation of prediction-based distillation loss. In this work, we propose an aleatoric uncertainty distillation loss to transfer the previous task's information to the current model. The aleatoric uncertainty tells us how much the model is uncertain about the prediction due to the data. In the incremental learning setting, the previously trained model's data uncertainty should be similar for the current task models. This is because the data uncertainty indicates properties such as low illumination or blur that would be a factor of uncertainty for all methods. In incremental learning, at task t, the current model M_t is initialized with the previous task model M_{t-1} . The objective is to train the M_t model from the current task dataset i.e. D_t^{train} . We use the cross-entropy loss to train M_t from the current task data. For any input image, x of current task data $x \in D_t^{train}$ is forwarded to the current model and previous model to predict the aleatoric variance value. For transferring the uncertainty, we define the aleatoric distillation loss $\mathcal{L}_{\mathcal{A}}$ using the following equation:

$$\mathcal{L}_{\mathcal{A}} = \sum_{i} ||(\sigma_{i,t-1}^{p})^{2} - (\sigma_{i,t}^{p})^{2}||^{2}$$
(12)

3.6. Attention based Distillation

The basic intuition is for attention-based distillation is that not all information needs to be preserved for the next task. In the traditional incremental learning, the method that has proposed the distillation loss apply it to preserve the full image information. It reduces a model's capability for learning a new task as it is learning unnecessary information. In [5], the authors use a GradCAM activation to distill the model information. These activations are obtained by the network's prediction, which is overconfident for even for the wrong prediction due to mismatch of train-test distribution [11], and thus do not capture the full prediction distribution. In contrast to the above global and activation based distillation, we are training the model to learn the attention map itself. Inspired by the self-attention methods [40, 46], we introduce an attention network to predict the attention map. These attention models are trained by the residual information to classification loss. Let the attention model $Att(:, \theta_{att})$ be parametrized by θ_{att} , then the attention obtained similar to [46] is given by:

$$y_{i,t}^{att} = Att(x_{i,t}^f, \theta_{att,t}) \tag{13}$$

 $(x_{i,t}^{f})$ is the feature representation at task t for input x_{i} .

3.6.1 Attention Distillation loss

In the incremental learning setting, if the current model M_t and previous model M_{t-1} have equivalent knowledge of base classes, they should have a similar response to these regions. Therefore the attention map obtained by the previous model and the current model should be similar. It transfers the old class data information to the current model. Since the previous model is trained on the old task data, the similarity between the previous model and the current model's attention map for new task data provides us with a signature of the old task data. So the current model objective is to preserve these signatures while learning the new task. We use the mean square loss for the attention distillation. L_2 loss or mean squared loss is obtained by squaring the difference between attention from the previous model and the current model. The loss function is defined as follows:

$$L_{att}^{t}(y_{i,t-1}^{att}, y_{i,t}^{att}) = \sum_{i} ||y_{i,t-1}^{att} - y_{i,t}^{att}||^{2}$$
(14)

3.6.2 Multi Stage Attention Distillation

In the convolution neural network, each stage learns a different feature for an input image. We use the attention module in each stage to predict the attention or essential part of images. These stage-wise features reveal the semantic structure of these images. In the proposed framework, we use these stage activations of the network to generate the attention map. For the incremental learning setting, we distill this attention map from the older model to the current model to better learns to preserve the older task information. In the proposed framework, we consider 3 stage self-attention net

Method	5 class		10 class		20 class	
Wiethou	ACC	FGT	ACC	FGT	ACC	FGT
Baseline	57.4	21.0	56.8	19.7	56.0	18.0
LwF-MC[25]	58.4	19.3	59.5	16.9	60.0	14.5
D+R[13]	59.1	19.6	60.8	17.1	61.8	14.3
E2E[3]	60.2	16.5	62.6	12.8	65.3	8.9
iCaRL[34]*	61.2	-	64.1	-	67.2	-
GD[21]	62.1	15.4	65.0	12.1	67.1	8.5
MD [47]	62.6	-	64.4	-	66.6	-
Ours-AU	61.8	15.2	65.4	11.4	67.6	8.2
Oracle	78.6	3.3	77.6	3.1	75.7	2.8

Table 1. Results on CIFAR-100. Baseline refers when distillation loss in not used and Oracle refers when we store full data samples. iCaRL uses ResNet-32 backbone.

in the incremental class setting. We provide the visualization of the 3rd stage attention map in Fig. 3. The multi-stage attention distillation loss is given as:

$$\mathcal{L}_m = \sum_{s=1}^{3} L_{att}^t(y_{i,t-1}^{att(s)}, y_{i,t}^{att(s)})$$
(15)

Here $y_{i,t-1}^{att(s)}$ corresponds to the attention map of i^{th} input at step t from the (t-1)th model at stage s.

3.7. Total Loss

We train the model using the cross entropy loss (Eq. 4), distillation loss (Eq. 3), aleatoric uncertainty (Eq. 11), uncertainty distillation (Eq. 12) and attention distillation loss (Eq. 15). The total loss is given as follows:

$$L_{total}^{t} = \lambda (\mathcal{L}_{m} + \mathcal{L}_{\mathcal{A}} + \mathcal{L}_{ale}) + \mathcal{L}_{C} + \mathcal{L}_{D}$$
(16)

Here λ is the weighting parameter for the uncertainty and attention based distillation.

4. Experiments and Analysis

We evaluate our proposed method by comparing it with state-of-the-art methods for dataset CIFAR-100 and Imagenet dataset. The ablation analysis with uncertainty and attention-based distillation is also shown in subsequent sections. Also, we have provided attention visualization for various methods, as shown in Fig. 3. For our evaluation purpose, we use CIFAR-100 and Imagenet dataset, as mentioned in section 4.1.

4.1. Dataset

CIFAR-100: In the CIFAR-100 [18] dataset, there are 100 classes containing 60,000 RGB images of size 32 x 32. Each class contains 500 training images and 100 test images. We follow the benchmark protocol given in [34] for our experiments. The classes are shuffled similar to [34, 3]. We train all 100 classes of CIFAR-100 in incremental

batches of 5, 10, and 20 classes, so respectively, there are 20,10 and 5 tasks present.

ImageNet: For the Imagenet dataset [4], we follow a similar approach provide in [21] by sample 500 images per 100 randomly chosen classes for each trial, and then split the classes. These images also have a dimension of 32 x 32. We use the same split of [21] in our paper. In ImageNet, we also used incremental batches of 5, 10, and 20 classes, so respectively, there are 20, 10, and 5 tasks present.

4.2. Implementation Details

We follow similar protocols, as discussed in [22]. We use Wide-ResNet [44] architecture as our base model. We also obtain a comparable performance with ResNet [10]. The model is implemented in the PyTorch [29] framework. The aleatoric variance layer is a linear layer followed by the soft plus layer. The attention modules are consist of convolution layers. All the models are initialized with random weights. During training, we use a batch size of 128 in all the experiments. We standardize input images by subtracting them with the mean of the training set and dividing them by the standard deviation. In all methods, we use the total buffer size of the representative image number of 2000 [34, 21]. These representative images are randomly selected. Other details are provided in the project page¹.

4.3. Comparison with state-of-the-art methods

We provide comparison with other state-of-the-art methods such as Learning without forgetting (LwF) [25], distillation and retrospection (D+R) [13], End-to-End incremental learning (E2E) [3], Global Distillation (GD) [21], iCaRL [34] and maintaining discrimination (MD) [47]. The upper bound results correspond when we consider all the task data. Baseline results, when we are not using any distillation and uncertainty estimation. The CIFAR-100 results are provided in Table 1 with increment tasks size 5, 10, and 20. The state-of-the-art method's values in the table have been evaluated from [21]. In Fig. 2, we have plotted the accuracy and forgetting of these increment tasks. Similarly, ImageNet dataset's results are also provided in Table 2 with increment tasks size are 5, 10, and 20. In both of the tables, the baseline refers to the model where distillation loss is not used; oracle refers to when we store full data samples. ACC refers to average accuracy (higher is better), and FGT refers to the amount of catastrophic forgetting, by averaging the performance decay (lower is better). We followed a similar approach of [21] to obtain the ACC and FGT. Our Attention Uncertainty (AU) methods improves 4% from the baseline method [25] in 5 class increment steps, around 6% improvement as in 10 class increment steps, and around 7% improvement in terms of accuracy in 20 class increment steps.

¹https://github.com/DelTA-Lab-IITK/

Incremental-learning-AU



Figure 2. This figure shows accuracy and forgetting plot for 5 class , 10 class and 20 class increments for CIFAR-100 dataset

Also, we have 1.6% improvement in accuracy in comparison with the state of the art method [3] in 5 steps, 2.8%in 10 steps, and 1.5% in 20 steps. Note that this additional improvement in accuracy is accompanied by the ability to interpret the model through the obtained by uncertainty estimates and the attention map.

Mathod	5 class		10 class		20 class	
Wiethou	ACC	FGT	ACC	FGT	ACC	FGT
Baseline	44.2	23.6	44.1	21.5	44.7	18.4
LwF-MC	45.6	21.5	47.3	18.5	48.6	15.3
D+R	46.5	22.0	48.7	18.8	50.7	15.1
E2E	47.7	17.9	50.8	13.4	53.9	8.8
GD	50.0	16.8	53.7	12.8	56.5	8.4
Ours-AU	51.2	15.4	54.3	11.9	56.9	8.1
Oracle	68.0	3.3	66.9	3.1	65.1	2.7

Table 2. Results on ImageNet-100. Baseline refers when distillation loss in not used and Oracle refers when we store full data samples

4.4. Class Incremental Accuracy

For showing the class increment effect in incremental class learning, we have plotted each incremental step average accuracy as well as forgetting in Fig. 2 for CIFAR-100 dataset. The Fig. 2(a) shows for 5 class increments, (b) for 10 class increments, and (c) for the 20 class increments. Similarly, for the same CIFAR-100 dataset, we have plotted the forgetting value for the different models on (d) 5 class increments, (e) for 10 class increments, and (f) for the 20 class increment task in the Fig. 2. From the figure, it is clear the proposed uncertainty, and attention distillation based method outperforms as compare to other state-of-the-art methods Learning without forgetting (LwF) [25], Progressive distillation and retrospection (D+R) [13], iCaRL [34], End-to-End incremental learning (E2E) [3] and Global distillation (GD) [22].



Figure 3. Visualization of attention map in CIFAR-100 in different task learning stage of 10 class increment setting. Each row represents the task stage,

Methods	ACC	FGT
AU	63.81	12.22
AD	64.37	12.03
AU +UD	64.54	11.84
AU + AD	64.52	11.92
AU +UD + AD	65.42	11.43

Table 3. Ablation analysis for different loss functions on CIFAR-100 dataset for 10 class incremental setup.

λ	Model	ACC	FGT
0.1	AU + AD	64.32	11.97
0.1	AU +UD + AD	64.63	11.76
0.5	AU + AD	64.54	11.84
0.5	AU +UD + AD	65.42	11.43
1	AU + AD	64.13	12.11
1	AU +UD + AD	64.25	12.08

Table 4. Ablation analysis for different weights of loss functions on CIFAR-100 dataset for 10 class incremental setup.

4.5. Effect of Attention and Uncertainty Distillation

To analyze the effect of uncertainty and attention in incremental learning, we provide the ablation analysis of different loss functions in Table 3 for the CIFAR-100 dataset for 10 class increment setup. Here, for all, we use the $\lambda = 0.5$ for uncertainty and attention distillation loss. In Table, **AU:** Aleatoric Uncertainty, **AD:** Attention Distillation and **UD**: Uncertainty Distillation. In this experiment, we can see that the uncertainty distillation and attention distillation improve the model performance.

4.6. Effect of weights of Attention and Uncertainty Distillation loss

To further analyze the effect of attention and uncertainty on the incremental learning setup, we experiment with different weight values for the loss, and the results are reported in Table 4. In the table, λ refers to the weight value of uncertainty and attention distillation. This experiment is performed for the CIFAR-100 dataset and 10 incremental class setup. We experiment with weight value for 0.1, 0.5, and 1. We can see that when the distillation loss weighted by 0.5 gives better performance in terms of better accuracy and less forgetting. This is expected because as we increment the distillation loss's contribution, it does not perform well on the current task; similarly, if the weight value is less, it does not keep the previous task information very well.

4.7. Attention Visualization

In Fig. 3, we provide the attention map visualization for a given input of different class CIFAR-100 dataset in 10 class increment setting. In each column, we have shown different instances of an example, while along the column, we have shown the attention visualization in different step sizes. For example, in the first row, the attention map is generated when the model generated an attention map for a model trained on the first 10 classes. Similarly, the second row shows the attention after the model is trained on till 20 classes in increment setting. Note that visualization is shown here from the last stage of the attention module. From the figure, we observe that the attention map is located in the object region for each task. The model distills the attention information to the next task; this helps to avoid the catastrophic forgetting in incremental learning.

5. Conclusion

In this paper, we consider the task of incremental learning, where we aim to have models that are able to preserve their performance across tasks. Several methods have been proposed for this task; these have mostly considered point estimates during prediction; for instance, those obtained using softmax. We show that these methods are susceptible to be incorrectly certain about performance, and by incorporating uncertainty, we are able to obtain a method that incrementally learns on tasks while also accurately providing the uncertainty of the model. Further, by incorporating self-attention, we can also visualize the regions on which the method attends across tasks. Further, our results show consistent improvements over the state of the art methods on standard datasets.

References

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings* of the European Conference on Computer Vision (ECCV), pages 139–154, 2018.
- [2] Eden Belouadah and Adrian Popescu. Scail: Classifier weights scaling for class incremental learning. In *The IEEE Winter Conference on Applications of Computer Vi*sion, pages 1266–1275, 2020.
- [3] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European Conference* on Computer Vision (ECCV), pages 233–248, 2018.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009.
- [5] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyan Wu, and Rama Chellappa. Learning without memorizing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5138–5146, 2019.
- [6] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning: Supplementary material. ECCV, 2020.
- [7] Sayna Ebrahimi, Franziska Meier, Roberto Calandra, Trevor Darrell, and Marcus Rohrbach. Adversarial continual learning. *ECCV*, 2020.
- [8] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Advances in neural information processing systems, pages 2672–2680, 2014.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [11] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations*, 2019.
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531, 2015.
- [13] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Lifelong learning via progressive distillation and retrospection. In *Proceedings of the European Conference* on Computer Vision (ECCV), pages 437–452, 2018.
- [14] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances*

in neural information processing systems, pages 5574–5584, 2017.

- [15] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.
- [16] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [17] Max Kochurov, Timur Garipov, Dmitry Podoprikhin, Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Bayesian incremental learning for deep neural networks. *Workshop ICLR 2018*, 2018.
- [18] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [19] Vinod Kumar Kurmi, Shanu Kumar, and Vinay P. Namboodiri. Attending to discriminative certainty for domain adaptation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2019.
- [20] kyle dorman. *Bayesian Uncertainty*, 2017 (accessed Sep 3, 2018).
- [21] Kibok Lee, Kimin Lee, Jinwoo Shin, and Honglak Lee. Overcoming catastrophic forgetting with unlabeled data in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 312–321, 2019.
- [22] Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. In *Advances in neural information processing systems*, pages 4652–4662, 2017.
- [23] Timothée Lesort, Hugo Caselles-Dupré, Michael Garcia-Ortiz, Andrei Stoian, and David Filliat. Generative models from the perspective of continual learning. In 2019 International Joint Conference on Neural Networks (IJCNN), pages 1–8. IEEE, 2019.
- [24] Yu Li, Zhongxiao Li, Lizhong Ding, Yijie Pan, Chao Huang, Yuhui Hu, Wei Chen, and Xin Gao. Supportnet: solving catastrophic forgetting in class incremental learning with support data. *ICLR*, 2018.
- [25] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelli*gence, 40(12):2935–2947, 2017.
- [26] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In Advances in Neural Information Processing Systems, pages 6467–6476, 2017.
- [27] Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. In Advances in Neural Information Processing Systems, pages 7047–7058, 2018.
- [28] Nicolas Y Masse, Gregory D Grant, and David J Freedman. Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization. *Proceedings of the National Academy of Sciences*, 115(44):E10467–E10475, 2018.

- [29] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [30] Badri N. Patro, Mayank Lunayach, Shivansh Patel, and Vinay P. Namboodiri. U-cam: Visual explanation using uncertainty based class activation maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (*ICCV*), October 2019.
- [31] Jathushan Rajasegaran, Munawar Hayat, Salman H Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for continual learning. In Advances in Neural Information Processing Systems, pages 12669–12679, 2019.
- [32] Jathushan Rajasegaran, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Mubarak Shah. itaml: An incremental task-agnostic meta-learning approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13588–13597, 2020.
- [33] Amal Rannen, Rahaf Aljundi, Matthew B Blaschko, and Tinne Tuytelaars. Encoder based lifelong learning. In Proceedings of the IEEE International Conference on Computer Vision, pages 1320–1328, 2017.
- [34] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017.
- [35] Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. In *International Conference on Machine Learning*, pages 4535–4544, 2018.
- [36] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–626, 2017.
- [37] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4555–4564, 2018.
- [38] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In Advances in Neural Information Processing Systems, pages 2990–2999, 2017.
- [39] Gido M van de Ven and Andreas S Tolias. Generative replay with feedback connections as a general strategy for continual learning. arXiv preprint arXiv:1809.10635, 2018.
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. 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 5998–6008. 2017.
- [41] Chenshen Wu, Luis Herranz, Xialei Liu, Joost van de Weijer, Bogdan Raducanu, et al. Memory replay gans: Learning to

generate new categories without forgetting. In *Advances In Neural Information Processing Systems*, pages 5962–5972, 2018.

- [42] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, Zhengyou Zhang, and Yun Fu. Incremental classifier learning with generative adversarial networks. arXiv preprint arXiv:1802.00853, 2018.
- [43] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29, 2016.
- [44] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. NIN, 8:35–67.
- [45] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings* of the 34th International Conference on Machine Learning-Volume 70, pages 3987–3995. JMLR. org, 2017.
- [46] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International Conference on Machine Learning*, pages 7354–7363, 2019.
- [47] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13208–13217, 2020.