

Encoder Based Lifelong Learning

Amal Rannen* Rahaf Aljundi* Matthew B. Blaschko Tinne Tuytelaars
 KU Leuven
 KU Leuven, ESAT-PSI, IMEC, Belgium
 firstname.lastname@esat.kuleuven.be

Abstract

This paper introduces a new lifelong learning solution where a single model is trained for a sequence of tasks. The main challenge that vision systems face in this context is catastrophic forgetting: as they tend to adapt to the most recently seen task, they lose performance on the tasks that were learned previously. Our method aims at preserving the knowledge of the previous tasks while learning a new one by using autoencoders. For each task, an under-complete autoencoder is learned, capturing the features that are crucial for its achievement. When a new task is presented to the system, we prevent the reconstructions of the features with these autoencoders from changing, which has the effect of preserving the information on which the previous tasks are mainly relying. At the same time, the features are given space to adjust to the most recent environment as only their projection into a low dimension submanifold is controlled. The proposed system is evaluated on image classification tasks and shows a reduction of forgetting over the state-of-the-art.

1. Introduction

Intelligent agents are able to perform remarkably well on individual tasks. However, when exposed to a new task or a new environment, such agents have to be retrained. In this process, they learn the specificity of the new task but tend to loose performance on the tasks they have learned before. For instance, imagine an agent that was trained to localize the defects on a set of factory products. Then, when new products are introduced and the agent has to learn to detect the anomalies in these new products, it faces the risk of forgetting the initial recognition task. This phenomenon is known as *catastrophic forgetting* [16, 21, 15, 7, 10]. It occurs when the datasets or the tasks are presented to the model separately and sequentially, as is the case in a lifelong learning setup [25, 26, 24].

The main challenge is to make the learned model adapt to new data from a similar or a different environment [19], without losing knowledge on the previously seen task(s). Most of the classical solutions for this challenge suffer from important drawbacks. Feature extraction (as in [5]), where the model / representation learned for the old task is re-used to extract features from the new data without adapting the model parameters, is highly conservative for the old task and suboptimal for the new one. Fine-tuning (as in [8]), adapts the model to the new task using the optimal parameters of the old task as initialization. As a result, the model is driven towards the newly seen data but forgets what was learned previously. Joint training (as in [4]) is a method where the model is trained jointly on previous and current tasks data. This solution converges to the best compromise between ta, but requires the presence of all the data at the same time. Such a requirement can be hard to meet, especially in the era of big data.

To overcome these drawbacks without the constraint of storing data from the previously seen tasks, two main approaches can be found in the literature. The first, presented in [14], proposes a way to train convolutional networks, where a shared model is used for the different tasks but with separate classification layers. When a new task is presented, a new classification layer is added. Then, the model is fine-tuned on the data of the new task, with an additional loss that incorporates the knowledge of the old tasks. This loss tries to keep the previous task predictions on the new data unchanged. Such a solution reduces the forgetting but is heavily relying on the new task data. As a consequence, it suffers from a build-up of errors when facing a sequence of tasks [1]. The work presented recently in [12] tackles the problem in a different way. Rather than having a data-oriented analysis, they consider the knowledge gained in the model itself, and transfer it from one task to another in a Bayesian update fashion. The method relies on approximating the weight distribution after training the model on the new task. A Gaussian distribution, for which the mean is given by the optimal weights for the first task, and the variance given by the diagonal of the Fisher information matrix

* Authors with equal contribution

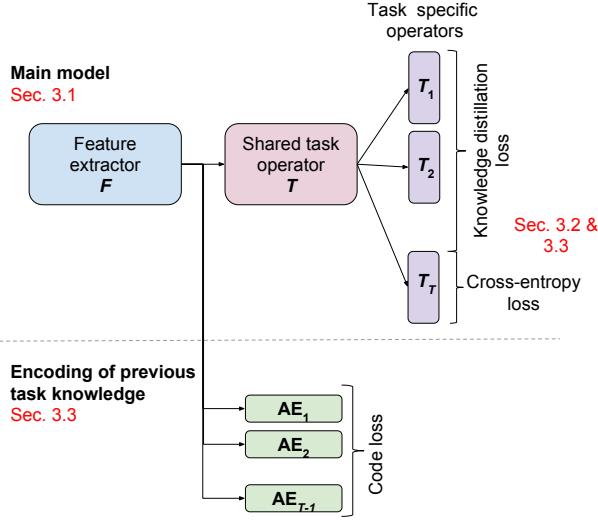


Figure 1. Diagram of the proposed model. Above the dotted line are the model components that are retained during test time, while below the dashed line are components necessary to our improved training scheme.

is used as an approximation. Such a solution is based on a strong principle and gives interesting results. However, it requires to store a number of parameters that is comparable to the size of the model itself.

In this work, we propose a compromise between these two methods. Rather than heavily relying on the new task data or requiring a huge amount of parameters to be stored, we introduce the use of autoencoders as a tool to preserve the knowledge from one task while learning another. For each task, an undercomplete autoencoder is trained after training the task model. It captures the most important features for the task objective. When facing a new task, this autoencoder is used to ensure the preservation of those important features. This is achieved by defining a loss on the reconstructions made by the autoencoder, as we will explain in the following sections. In this manner, we only restrict a subset of the features to be unchanged while we give the model the freedom to adapt itself to the new task using the remaining capacity. Figure 1 displays the model we propose to use.

Below, we first give a short description of the most relevant related work in Sec. 2. Then, in Sec 3, we describe how to use the autoencoders to avoid catastrophic forgetting and motivate our choice by a short analysis that relates the proposed objective to the joint training scheme. In Sec. 4, we describe the experiments that we conducted, report and discuss their results before concluding in Sec. 5.

2. Related work

Our end goal is to train a single model that can perform well on multiple tasks, with tasks learned sequentially. This problem is at the intersection of **Joint training** (or multi-task training) and **Lifelong learning**. Standard multi-task learning [4] aims to learn jointly from the data of the multiple tasks and uses inductive bias [17] in order to integrate the knowledge from the different domains in a single model. However, it requires the presence of data from all the tasks during training. In a lifelong learning scenario, on the other hand, tasks are treated in a sequential manner. The aim is then to exploit the knowledge from previous tasks while learning a new one. This knowledge is used to (i) preserve the performance on the previously seen data, (ii) improve this knowledge using inductive bias [17] from the new task data, and (iii) act as a regularizer for the new task, which can be beneficial for the performance. *In this work we aim at preserving the knowledge of the previous tasks and possibly benefiting from this knowledge while learning a new task, without storing data from previous tasks.*

Despite its potential benefits, this problem is under explored. **Learning without forgetting** (LwF), introduced in [14], proposes to preserve the previous performance through the knowledge distillation loss introduced in [11]. They consider a shared convolutional network between the different tasks in which only the last classification layer is task specific. When encountering a new task, the outputs of the existing classification layers given the new task data are recorded. During training, these outputs are preserved through a modified cross-entropy loss that softens the class probabilities in order to give a higher weight to the small outputs. More details about this loss and the method can be found in Sec. 3.2. This method reduces the forgetting, especially when the datasets come from related manifolds. Nevertheless, it has been shown by **iCaRL: Incremental Classifier and Representation Learning** [22] that LwF suffers from a build up of errors in a sequential scenario where the data comes from the same environment. Similarly, **Expert-gate** [1] shows that the LwF performance drops when the model is exposed to a sequence of tasks drawn from different distributions. [22] proposes to store a selection of the previous tasks data to overcome this issue – something we try to avoid. The goal in [1] is to obtain experts for different tasks (instead of a single joint model shared by all tasks). They suggest a model of lifelong learning where experts on individual tasks are added to a network of models sequentially. The challenge then is to decide which expert to launch based on the input. Interestingly, they also use undercomplete autoencoders – in their case to capture the context of each task based on which a decision is made on which task the test sample belongs to. *In this work, we build on top of LwF but reduce the cumulated errors using undercomplete autoencoders learned on the optimal representa-*

tions of the previous tasks.

Even more recently, another interesting solution to train shared models without access to the previous data, somewhat similar in spirit to our work, has been proposed in [12], in the context of reinforcement learning. The main idea of this method, called **Elastic weight consolidation**, is to constrain the weights W_i while training for a second task with an additional loss $\sum_i \frac{\lambda}{2} F_i(W_i - W_{1,i}^*)$ where $W_{1,i}^*$ are the optimal weights of the first task, and F_i the diagonal terms of their Fisher information matrix. The use of the Fisher matrix prevents the weights that are important for the first task to change much. In our point of view, this method, despite its success, has two drawbacks. First, the method keeps the weights in a neighborhood of one possible minimizer of the empirical risk of the first task. However, there could be another solution that can give a better compromise between the two tasks. Second, it needs to store a large number of parameters that grows with the total number of weights and the number of tasks. For these reasons, *rather than constraining the weights, we choose to constrain the resulting features, enforcing that those that are important for the previous tasks do not change much*. By constraining only a sub-manifold of the features, we allow the weights to adjust so as to optimize the features for the new task, while preserving those that ensure a good performance on the previous tasks.

3. Overcoming forgetting with autoencoders

In this work, we consider the problem of training a supervised deep model that can be useful for multiple tasks, in the situation where at each stage the data fed to the network come always from one single task, and the tasks enter in the training scenario successively. The best performance for all the tasks simultaneously is achieved when the network is trained on the data from all the considered tasks at the same time (as in joint training). This performance is of course limited by the capacity of the used model, and can be considered an upper bound to what can be achieved in a lifelong learning setting, where the data of previous tasks is no longer accessible when learning a new one.

3.1. Joint training

In the following, we will use the notations $\mathcal{X}^{(t)}$ (model input) and $\mathcal{Y}^{(t)}$ (target) for the random variables from which the dataset of the task t is sampled, and $X_i^{(t)}$ and $Y_i^{(t)}$ for the data samples. When we have access to the data from all \mathcal{T} tasks jointly, the network training aims to control the statistical risk:

$$\sum_{t=1}^{\mathcal{T}} \mathbb{E}_{(\mathcal{X}^{(t)}, \mathcal{Y}^{(t)})} [\ell(f_t(\mathcal{X}^{(t)}), \mathcal{Y}^{(t)})], \quad (1)$$

by minimizing the empirical risk:

$$\sum_{t=1}^{\mathcal{T}} \frac{1}{N_t} \sum_{i=1}^{N_t} \ell(f_t(X_i^{(t)}), Y_i^{(t)}), \quad (2)$$

where N_t is the number of samples and f_t the function implemented by the network for task t . For most of the commonly used models, we can decompose f_t as $T_t \circ T \circ F$ where:

- F is a feature extraction function (e.g. Convolutional layers in ConvNets)
- $T_t \circ T$ is a task operator. It can be for example a classifier or a segmentation operator. T is shared among all tasks, while T_t is task specific. (e.g. in ConvNets, T_t could be the last fully-connected layer, and T the remaining fully-connected layers.)

The upper part of Figure 1 gives a scheme of this general model. For simplicity, we will focus below on two-task training before generalizing to a multiple task scenario in section 3.4.

3.2. Shortcomings of Learning without Forgetting

As a first step, we want to understand the limitations of LwF [14]. In that work, it is suggested to replace in Eq (1) $\ell(T_1 \circ T \circ F(\mathcal{X}^{(1)}), \mathcal{Y}^{(1)})$ with $\ell(T_1 \circ T \circ F(\mathcal{X}^{(2)}), T_1^* \circ T^* \circ F^*(\mathcal{X}^{(2)}))$, where $T_1^* \circ T^* \circ F^*$ is obtained from training the network on the first task. If we suppose that the model has enough capacity to integrate the knowledge of the first task with a small generalization error, then we can consider that

$$\mathbb{E}_{(\mathcal{X}^{(1)})} [\ell(T_1 \circ T \circ F(\mathcal{X}^{(1)}), T_1^* \circ T^* \circ F^*(\mathcal{X}^{(1)}))] \quad (3)$$

is a reasonable approximation of $\mathbb{E}_{(\mathcal{X}^{(1)}, \mathcal{Y}^{(1)})} [\ell(T_1 \circ T \circ F(\mathcal{X}^{(1)}), \mathcal{Y}^{(1)})]$. However, in order to be able to compute the measure (3) using samples from $\mathcal{X}^{(2)}$, further conditions need to be satisfied.

In other terms, if we consider that $T_1 \circ T \circ F$ tries to learn an encoding of the data in the target space $\mathcal{X}^{(1)}$, then one can say that the loss of information generated by the use of $\mathcal{X}^{(2)}$ instead of $\mathcal{X}^{(1)}$ is a function of the Kullback-Leibler divergence of the two related probability distributions, or equivalently of their cross-entropy. Thus, if the two data distributions are related, then LwF is likely to lead to high performance. If the condition of the relatedness of the data distributions fails, there is no direct guarantee that the use of $\ell(T_1 \circ T \circ F(\mathcal{X}^{(2)}), T_1^* \circ T^* \circ F^*(\mathcal{X}^{(2)}))$ will not result in an important loss of information for the first task. Indeed, it has been shown empirically in [1] that the use of significantly different data distributions may result in a significant decrease in performance for LwF.

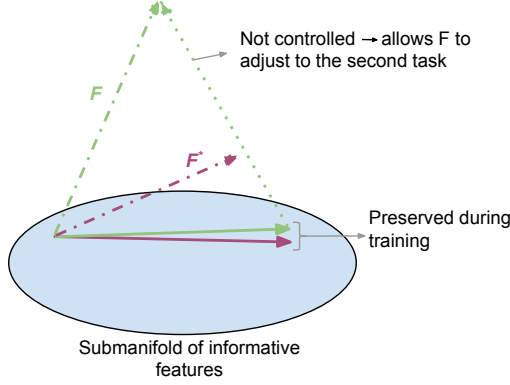


Figure 2. Preservation of the features that are important for task 1 while training on task 2. During training, we enforce the projection of F into the submanifold that captures these important features to stay close to the projection of F^* , the optimal features for the first task. The part of F that is not meaningful for the first task is allowed to adjust to the variations of the second task.

LwF is based on the knowledge distillation loss introduced in [11] to reduce the gap resulting from the use of different distributions. In this work, we build on top of the LwF method. In order to make the used approximation less sensitive to the data distributions, we see an opportunity in controlling $\|T_1 \circ T \circ F(\mathcal{X}^{(1)}) - T_1 \circ T \circ F(\mathcal{X}^{(2)})\|$. Under mild conditions about the model functions, namely Lipschitz continuity, this control allows us to use $T_1 \circ T \circ F(\mathcal{X}^{(2)})$ instead of $T_1 \circ T \circ F(\mathcal{X}^{(1)})$ to better approximate the first task loss in Eq. (1). Note that the condition of continuity on which this observation is based is not restrictive in practice. Indeed, most of the commonly used functions in deep models satisfy this condition (e.g. sigmoid, ReLU).

Our main idea is to learn a submanifold of the representation space $F(\mathcal{X}^{(1)})$ that contains the most informative features for the first task. Once this submanifold is identified, if the projection of the features $F(\mathcal{X}^{(2)})$ onto this submanifold do not change much during the training of a second task, then two consequences follow: (i) $F(\mathcal{X}^{(2)})$ will stay informative for the first task during the training, and (ii) at the same time there is room to adjust to the second task as only its projection in the learned submanifold is controlled. Figure 2 gives a simplified visualization of this mechanism. In the next paragraphs, we propose a method to learn the submanifold of informative features for a given task using autoencoders.

3.3. Informative feature preservation

When beginning to train the second task, the feature extractor F^* of the model is optimized for the first task. A feature extraction type of approach would keep this operator unchanged in order to preserve the performance on the previous task. This is, however, overly conservative, and usually suboptimal for the new task. Rather than preserving

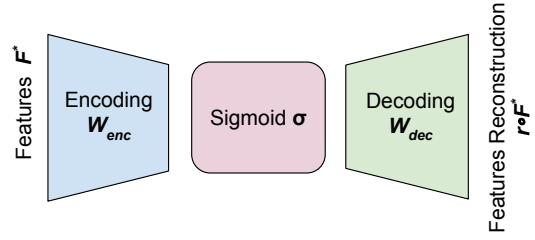


Figure 3. Scheme of an undercomplete autoencoder trained to capture the important features submanifold.

all the features during training, our main idea is to preserve only the features that are the most informative for the first task while giving more flexibility for the other features in order to improve the performance on the second task. An autoencoder [3] trained on the representation of the first task data obtained from an optimized model can be used to capture the most important features for this task.

3.3.1 Learning the informative submanifold with Autoencoders

An autoencoder is a neural network that is trained to reconstruct its input [9]. The network operates a projection r that can be decomposed in an encoding and a decoding function. The optimal weights are usually obtained by minimizing the mean ℓ_2 distance between the inputs and their reconstructions. If the dimension of the code is smaller than the dimension of the input (i.e. if the autoencoder is undercomplete), the autoencoder captures the submanifold that represents the best the structure of the input data. More precisely, we choose to use a two-layer network with a sigmoid activation in the hidden layer: $r(x) = W_{dec}\sigma(W_{enc}x)$. Figure 3 shows a general scheme of such an autoencoder.

Here, our aim is to obtain through the autoencoder a submanifold that captures the information that is not only important to reconstruct the features (output of the feature extraction operator F^*) of the first task, but also important for the task operator $(T_1^* \circ T^*)$. The objective is:

$$\arg \min_r \mathbb{E}_{(\mathcal{X}^{(1)}, \mathcal{Y}^{(1)})} [\lambda \|r(F^*(\mathcal{X}^{(1)})) - F^*(\mathcal{X}^{(1)})\|_2 + \ell(T_1^* \circ T^*(r(F^*(\mathcal{X}^{(1)}))), \mathcal{Y}^{(1)})], \quad (4)$$

where ℓ is the loss function used to train the model on the first task data. λ is a hyper-parameter that controls the compromise between the two terms in this loss. In this manner, the autoencoder represents the variations that are needed to reconstruct the input and at the same time contain the information that is required by the task operator.

3.3.2 Representation control with separate task operators

To explain how we use these autoencoders, we start with the simple case where there is no task operator shared among all tasks (i.e., $T = \emptyset$). The model is then composed of a common feature extractor F , and a task specific operator for each task T_t . Each time a new task is presented to the model, the corresponding task operator is then optimized. However, in order to achieve the aim of lifelong learning, we want to also update the feature extractor without damaging the performance of the model on the previously seen tasks.

In a two task scenario, after training the first task, we have T_1^* and F^* optimized for that task. Then, we train an undercomplete autoencoder using $F^*(X_i^{(1)})$ minimizing the empirical risk corresponding to (4). The optimal performance for the first task, knowing that the operator T_1 is kept equal to T_1^* , is obtained with F equal to F^* .

Nevertheless, preventing F from changing will lead to suboptimal performance on the second task. The idea here is to keep only the projection of F into the manifold represented by the autoencoder ($r \circ F$) unchanged. The second term of Eq. (4) explicitly enforces r to represent the submanifold needed for good performance on task 1. Thus, controlling the distance between $r \circ F$ and $r \circ F^*$ will preserve the necessary information for task 1. From the undercompleteness of the encoder, r projects the features into a lower dimensional manifold, and by controlling only the distance between the reconstructions, we give the features flexibility to adapt to the second task variations (cf. Figure 2).

3.3.3 Representation control with shared task operator

We now consider the model presented in Figure 1 where a part of the task operator is shared among the tasks as in the setting used in LwF [14]. This is clearly a preferable architecture in a lifelong learning setting, as the memory increase when adding a new task is much lower. Our main idea is to start from the loss used in the LwF method and add an additional term coming from the idea presented in Sec. 3.3.2. Thus, in a two task scenario, in addition to the loss used for the second task, we propose to use two constraints:

1. The first constraint is the knowledge distillation loss (ℓ_{dist}) used in [14]. If $\hat{\mathcal{Y}} := T_1 \circ T \circ F(\mathcal{X}^{(2)})$ and $\mathcal{Y}^* = T_1^* \circ T^* \circ F^*(\mathcal{X}^{(2)})$ then:

$$\ell_{dist}(\hat{\mathcal{Y}}, \mathcal{Y}^*) = -\langle \mathcal{Z}^*, \log \hat{\mathcal{Z}} \rangle \quad (5)$$

where log is operated entry-wise and

$$\mathcal{Z}_i^* = \frac{\mathcal{Y}_i^{*1/\theta}}{\sum_j \mathcal{Y}_j^{*1/\theta}} \text{ and } \hat{\mathcal{Z}}_i = \frac{\hat{\mathcal{Y}}_i^{1/\theta}}{\sum_j \hat{\mathcal{Y}}_j^{1/\theta}} \quad (6)$$

The application of a high temperature θ increases the small values of the output and reduces the weight of the high values. This mitigates the influence of the use of different data distributions.

2. The second constraint is related to the preservation of the reconstructions of the second task features ($r \circ F^*(\mathcal{X}^{(2)})$). The goal of this constraint is to keep $r \circ F$ close to $r \circ F^*$ as explained in Sec. 3.3.2.

For the second constraint, rather than controlling the distance between the reconstructions, we will here constrain the codes $\sigma(W_{enc})$. From sub-multiplicity of the Frobenius norm, we have:

$$\|r(x_1) - r(x_2)\|_2 \leq \|W_{dec}\|_F \|\sigma(W_{enc}x_1) - \sigma(W_{enc}x_2)\|_2.$$

The advantage of using the codes is their lower dimension. As the codes or reconstructions need to be recorded before beginning the training on the second task, using the codes will result in a better usage of the memory.

Finally, the objective for the training of the second task is the following:

$$\begin{aligned} \mathcal{R} = & \mathbb{E}[\ell(T_2 \circ T \circ F(\mathcal{X}^{(2)}), \mathcal{Y}^{(2)})] \\ & + \ell_{dist}(T_1 \circ T \circ F(\mathcal{X}^{(2)}), T_1^* \circ T^* \circ F^*(\mathcal{X}^{(2)})) \\ & + \frac{\alpha}{2} \|\sigma(W_{enc}F(\mathcal{X}^{(2)})) - \sigma(W_{enc}F^*(\mathcal{X}^{(2)}))\|_2^2. \end{aligned} \quad (7)$$

The choice of the parameter α will be done through model selection. An analysis of this objective is given in the appendix, giving a bound on the difference between $\mathbb{E}[\ell(T_2 \circ T \circ F(\mathcal{X}^{(2)}), \mathcal{Y}^{(2)})] + \ell(T_1 \circ T \circ F(\mathcal{X}^{(2)}), T_1^* \circ T^* \circ F^*(\mathcal{X}^{(2)}))$ and the statistical risk in a joint-training setting (1). It shows that (7) effectively controls this bound.

3.4. Training procedure

The proposed method in Sec. 3.3.3 generalizes easily to a sequence of tasks. An autoencoder is then trained after each task. Even if the needed memory will grow linearly with the number of tasks, the memory required by an autoencoder is a small fraction of that required by the global model. For example, in the case of AlexNet as a base model, an autoencoder comprises only around 1.5% of the memory.

In practice, the empirical risk is minimized:

$$\begin{aligned} R_N = & \frac{1}{N} \sum_{i=1}^N \left(\ell(T_T \circ T \circ F(X_i^{(\mathcal{T})}), Y_i^{(\mathcal{T})}) \right. \\ & + \sum_{t=1}^{\mathcal{T}-1} \ell_{dist}(T_t \circ T \circ F(X_i^{(\mathcal{T})}), T_t^* \circ T^* \circ F^*(X_i^{(\mathcal{T})})) \\ & \left. + \sum_{t=1}^{\mathcal{T}-1} \frac{\alpha_t}{2} \|\sigma(W_{enc,t}F(X_i^{(\mathcal{T})})) - \sigma(W_{enc,t}F^*(X_i^{(\mathcal{T})}))\|_2^2 \right). \end{aligned} \quad (8)$$

Algorithm 1 Encoder based Lifelong Learning

Input :

F^* shared feature extractor; T^* : shared task operator;
 $\{T_t\}_{t=1..T-1}$ previous task operators;
 $\{W_{enc,t}\}_{t=1..T-1}$ previous task encoders;
 $(X^{(\mathcal{T})}, Y^{(\mathcal{T})})$ training data and ground truth of the new task \mathcal{T} ;
 α_t and λ //hyper parameters

Initialization :

- 1: $Y_t^* = T_t^* \circ T^* \circ F^*(X^{(\mathcal{T})})$ //record task targets
- 2: $C_t^* = \sigma(W_{enc,t} F^*(X^{(\mathcal{T})}))$ //record new data codes
- 3: $T_{\mathcal{T}} \leftarrow \text{Init}(Y^{(\mathcal{T})})$ // initialize new task operator

Training :

- 4: $\hat{Y}_t = T_t \circ T \circ F(X^{(\mathcal{T})})$ // task outputs
 - 5: $\hat{C}_t = \sigma(W_{enc,t} F(X^{(\mathcal{T})}))$ //current codes
 - 6: $T_t^*, T^*, F^* \leftarrow \arg \min_{T_t, T, F} [\ell(\hat{Y}_t, Y^{(\mathcal{T})}) + \sum_{t=1}^{T-1} \ell_{dist}(\hat{Y}_t, Y_t^*) + \sum_{t=1}^{T-1} \frac{\alpha_t}{2} \|\hat{C}_t - C_t^*\|^2]$
 - 7: $(W_{enc,\mathcal{T}}, W_{dec,\mathcal{T}}) \leftarrow \text{autoencoder}(T_{\mathcal{T}}^*, T^*, F^*, X^{(\mathcal{T})}, Y^{(\mathcal{T})}; \lambda)$ // minimizes Eq. (4)
-

The training is done using stochastic gradient descent (SGD) [2]. The autoencoder training is also done by SGD but with an adaptive gradient method, AdaDelta [28] which alleviates the need for setting the learning rates and has nice optimization properties. Algorithm 1 shows the main steps of the proposed method.

4. Experiments

We compare our method against the state-of-the-art and several baselines on image classification tasks. We consider sets of 2, 3 and 5 tasks learned sequentially, in two settings: 1) when the first task is a large dataset, and 2) when the first task is a small dataset.

Architecture We experiment with AlexNet [13] as our network architecture due to its widespread use and similarity to other popular architectures. The feature extraction block F corresponds to the convolutional layers. By default, the shared task operator T corresponds to all but the last fully connected layers (i.e., $fc6$ and $fc7$), while the task-specific part T_i contains the last classification layer ($fc8$). Other choices for F and T are possible. More details can be found in the supplementary materials. During the training, we used an α in the order of 10^{-3} for ImageNet and 10^{-2} for the rest of the tasks. Note that this parameter sets the trade off between the allowed forgetting on the previous task and the performance on the new task.

For the autoencoders, we use a very shallow architecture, to keep their memory footprint low. Both the encoding as well as the decoding consist of a single fully connected layer, with a sigmoid as non-linearity in between. The dimensionality of the codes is 100 for all datasets, except for

ImageNet where we use a code size of 300. The size of the autoencoder is 3MB, compared to 250MB for the size of the network model. The training of the autoencoders is done using AdaDelta as explained in Sec. 3.4. During training of the autoencoders, we use a hyperparameter λ (cf. Eq. (4)) to find a compromise between the reconstruction error and the classification error. This parameter is tuned manually and is set to 10^{-6} in all cases.

Datasets We use multiple datasets of moderate size: MIT *Scenes* [20] for indoor scene classification (5,360 samples), Caltech-UCSD *Birds* [27] for fine-grained bird classification (5,994 samples), Oxford *Flowers* [18] for fine-grained flower classification (2,040 samples) and VOC *Actions* [6], the human action classification subset of VOC challenge 2012. These datasets were also used in both [1] and [14].

For the scenario based on a large initial dataset, we start from ImageNet (LSVRC 2012 subset) [23], which has more than 1 million training images. For the small dataset scenario, we start from Oxford *Flowers*, which has only 2,040 training and validation samples.

The reported results are obtained with respect to the test sets of Scenes, Birds, Flowers and Actions, and on the validation set of ImageNet. As in LwF [14], we need to record the targets corresponding to the old tasks before starting the training procedure for a new task. Here, we perform an offline augmentation with 10 variants of each sample (different crops and flips). This setting differs slightly from what has been done in [14], which explains the higher performance on the individual tasks in our experiments. We therefore compare against a stronger baseline than the accuracies reported in [14].

Compared Methods We compare our method (**Ours**) with Learning without Forgetting (**LwF**) [14], which represents the current state-of-the-art. Additionally, we consider two baselines: **Finetuning**, where each model (incl. F and T) is learned for the new task using the previous task model as initialization, and **Feature extraction**, where the weights of the previous task model (F and T) are fixed and only the classification layer (T_t) is learned for each new task. Further, we also report results for a variant of our method, **Ours - separate FCs** where we only share the representation layers (F) while each task has its own fully connected layers (i.e., $T = \emptyset$ and $T_i = \{fc6 - fc7 - fc8\}$). This variant aims at finding a universal representation for the current sequence of tasks while allowing each task to have its own fully connected layers. With less sharing, the risk of forgetting is reduced, at the cost of a higher memory consumption and less regularization for new tasks. Note that in this case the fully connected layers of the previous tasks are not retrained, and thus there is no need to use the knowledge distillation loss. Moreover, task autoencoders can be used at test time to activate only the fully connected layers of the

	ImageNet → Scenes		ImageNet → Birds		ImageNet → Flowers		Average loss	
	Acc. on Task1	Acc. on Task2	Acc. on Task1	Acc. on Task2	Acc. on Task1	Acc. on Task2	Task 1	Task 2
Finetuning	48.0% (-9%)	65.0% (ref)	41.3% (-15.7%)	59.0% (ref)	50.8% (-6.2%)	86.4% (ref)	-10.3% (ref)	
Feature extraction	57.0% (ref)	60.6% (-4.4%)	57.0% (ref)	51.6% (-7.4%)	57.0% (ref)	84.6% (-1.8%)	(ref)	-4.5%
LwF	55.4% (-1.6%)	65.0% (-0%)	54.4% (-2.6%)	58.9% (-0.1%)	55.6% (-1.4%)	85.9% (-0.5%)	-1.9%	-0.2%
Ours	56.3% (-0.7%)	64.9% (-0.1%)	55.3% (-1.7%)	58.2% (-0.8%)	56.5% (-0.5%)	86.2% (-0.2%)	-1.0%	-0.4%
Ours - separate FCs	57.0% (-0%)	65.9% (+0.9%)	57.0% (-0%)	57.7% (-1.3%)	56.5% (-0.5%)	86.4% (-0%)	-0.2%	-0.1%

Table 1. Classification accuracy for the Two Task scenario starting from ImageNet. For the first task, the reference performance is given by Feature extraction. For the second task, we consider Finetuning as the reference as it is the best that can be achieved by one task alone.

task that a test sample belongs to, in a similar manner to what was done in [1].

Setups We consider sequences of 2, 3 and 5 tasks. In the **Two Tasks** setup, we are given a model trained on one previously seen task and then add a second task to learn. This follows the experimental setup of LwF [14]. In their work, all the tested scenarios start from a large dataset, ImageNet. Here we also study the effect of starting from a small dataset, Flowers.¹ Further, we also consider a setup involving **Three Tasks**. First, we use a sequence of tasks starting from ImageNet, i.e. ImageNet → Scenes → Birds. Additionally, we consider Flowers as a first task in the sequence Flowers → Scenes → Birds. Note that this is different from what was conducted in [14] where the sequences were only composed of splits of one dataset i.e. one task overall. Finally, for a stronger validation of our method, we also test on a longer sequence: ImageNet → Scenes → Birds → Flowers → Actions.

Results Table 1 shows, for the different compared methods, the achieved performance on the Two Tasks scenario with ImageNet as the first task. While **Finetuning** is optimal for the second task, it shows the most forgetting of the first task. The performance on the second task is on average comparable for all methods except for **Feature extraction**. Since the **Feature extraction** baseline doesn't allow the weights of the model to change and only optimizes the last fully connected layers, its performance on the second task is suboptimal and significantly lower than the other methods. Naturally, the performance of the previous task is kept unchanged in this case. **Ours - separate FCs** shows high performance on both tasks, with the performance of the second task being comparable or better to the methods with shared FCs. This variant of our method has a higher capacity as it allocates separate fully connected layers for each task, yet its memory consumption increases more rapidly as tasks are added, a severe drawback. Our method with a complete shared model **Ours** systematically outperforms the **LwF** method on the previous task and on average achieves a similar performance on the second task.

When we start from a smaller dataset, Flowers, the same trends can be observed, but with larger differences in accuracy (Table 2). The performance on the second task is lower than that achieved with ImageNet as a starting point for all the compared methods. This is explained by the fact that the

¹Due to the small size of the Flowers dataset, we use a network pre-trained on ImageNet as initialization for training the first task model. The main difference hence lies in the fact that in this case we do not care about forgetting ImageNet.

representation obtained from ImageNet is more meaningful for the different tasks than what has been finetuned for Flowers. Differently from the ImageNet starting case, **Ours - separate FCs** achieves a considerably better performance on the second task than **Ours** and **LwF** while preserving the previous task performance. **Finetuning** shows the best performance on the second task while suffering from severe forgetting on the previous task. The pair of tasks here is of a different distribution and finding a compromise between the two tasks is a challenging problem. As in the previous case, **Ours** reduces the forgetting of **LwF** while achieving a similar average performance on the second task.

Overall, the **Ours-separate FCs** achieves the best performance on the different pairs of tasks. However, it requires allocating separate fully connected layers for each task which requires a lot of memory. Thus, for the sequential experiments we focus on the shared model scenario.

In Table 3 we report the performance achieved by **Ours**, **LwF** and **Finetuning** for the sequence of ImageNet → Scenes → Birds. As expected, the **Finetuning** baseline suffers from severe forgetting on the previous tasks. The performance on ImageNet (the first task) drops from 57% to 37.9% after finetuning on the third task. As this baseline does not consider the previous tasks in its training procedure, it has the advantage of achieving the best performance on the last task in the sequence.

Ours continually reduces forgetting compared to **LwF** on the previous tasks while showing a comparable performance on the new task in the sequence. For example, **Ours** achieves 54.9% on ImageNet compared to 53.3% by LwF. Similar conclusions can be drawn regarding the sequential scenario starting from Flowers as reported in Table 4.

Effect of the code length In order to examine the effect of varying the code size, we apply our method to the two-task experiment ImageNet → Scenes using AlexNet and autoencoders with code size varying from 20 to 9216 (full feature dimension) trained on the output of *conv5*. Figure 4 shows the classification accuracies for both of the datasets for different code sizes. The results indicate that the smaller code sizes favor the performance on the new task, while the larger code sizes lean towards a better preservation of the old task knowledge but a lower performance on the new task. This experiment confirms the motivation behind using autoencoders to encode the knowledge of previous tasks: a larger code means then a better preservation (see Figure 2 and Section 3.3.1), but a very large code preserves a noisy information, which explains the drop in the first task performance for the largest codes.

	Flowers → Scenes		Flowers → Birds		Average loss	
	Acc. on Task1	Acc. on Task2	Acc. on Task1	Acc. on Task2	Task 1	Task 2
Finetuning	61.6% (-24.8%)	63.9% (ref)	66.6% (-19.8%)	57.5% (ref)	-22.3% (ref)	
Feature extraction	86.4% (ref)	59.6% (-4.3%)	86.4% (ref)	48.6% (-8.9%)	(ref)	-6.6%
LwF	83.7% (-2.7%)	62.2% (-1.7%)	82.0 % (-4.4%)	52.2 % (-5.3%)	-3.6%	-3.5%
Ours	84.9% (-1.5%)	62.3% (-1.6%)	83.0% (-3.4%)	52.0 % (-5.5%)	-2.4%	-3.5%
Ours - separate FCs	86.4% (-0%)	63.0% (-0.9%)	85.4% (-1.0%)	55.1% (-2.4%)	-0.5%	-1.6%

Table 2. Classification accuracy for the Two Task scenario starting from Flowers. For the first task, the reference performance is given by Feature extraction. For the second task, we consider Finetuning as reference as it is the best that can be achieved by one task alone.

	ImageNet	Scenes	Birds	Average Acc.
Finetuning	37.5%	45.6%	58.1%	47.2%
LwF	53.3%	63.5%	57.2 %	58.0%
Ours	54.9%	64.7%	56.9%	58.8%

Table 3. Classification accuracy for the Three Task scenario starting from ImageNet. **Ours** achieves the best trade off between the tasks in the sequence with less forgetting to the previous tasks.

	Flowers	Scenes	Birds	Average Acc.
Finetuning	51.2%	48.1%	58.5%	51.6%
LwF	81.1%	59.1%	52.3%	64.1%
Ours	82.8%	61.2 %	51.2%	65.0%

Table 4. Classification accuracy for the Three Task scenario starting from Flowers. **Ours** achieves the best trade off between the tasks in the sequence with less forgetting to the previous tasks.

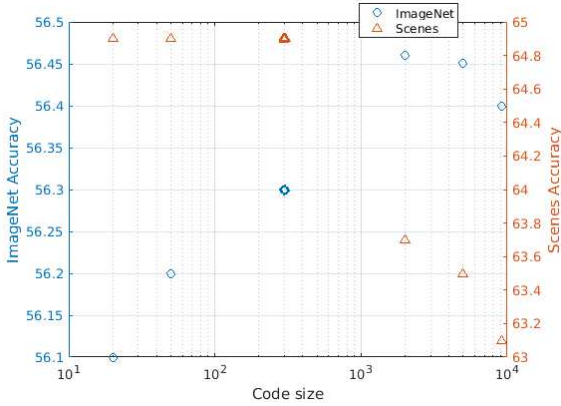


Figure 4. Classification accuracy for the Two Task scenario ImageNet → Scenes with different code sizes. The points in bold mark the code size corresponding to the setting of the experiments reported in Table 1.

Performance on longer sequences We compare the performance of our method to LwF on a longer sequence of five tasks: ImageNet → Scenes → Birds → Flowers → Actions. Figure 5 shows the obtained accuracies for the five datasets with both methods. It shows that our method outperforms LwF by 1.32% on average over all the tasks. The performance of our method on ImageNet after five tasks (53.6%) is higher than the performance of LwF (51.2%) . It worths noticing that our performance after 5 tasks is better than LwF performance after only 3 tasks (53.3%, Table 3).

5. Conclusions and future work

Strategies for efficient lifelong learning is still an open research problem. In this work, we tackled the problem of learning a sequence of tasks using only the data from the most recent environment, aiming at obtaining a reasonable

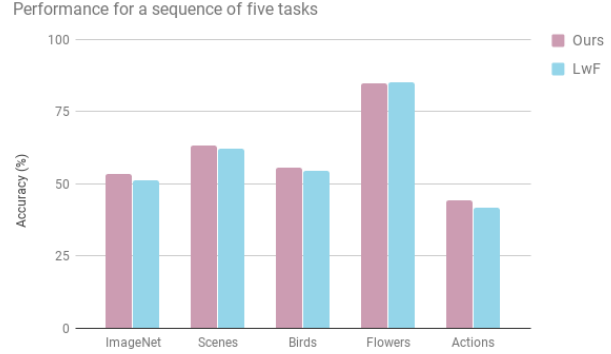


Figure 5. Classification accuracy for a Five Task scenario. The tasks (horizontal axis) are presented sequentially from left to right to AlexNet. The training is done using LwF and our method with shared fully connected layers for comparison.

performance on the whole sequence. Existing works consider solutions to preserve the knowledge of the previous tasks either by keeping the corresponding system predictions unchanged during training of the new task, or by keeping the model parameters in a neighborhood of the sequence of the previous optimal weights. While the first suffers from the difference in the task distributions, the second needs to store a large number of parameters.

The solution presented here reduces forgetting of earlier tasks by controlling the distance between the representations of the different tasks. Rather than preserving the optimal weights of the previous tasks, we propose an alternative that preserves the features that are crucial for the performance in the corresponding environments. Undercomplete autoencoders are used to learn the submanifold that represents these important features. The method is tested on image classification problems, in sequences of two, three or five tasks, starting either from a small or a large dataset. An improvement in performance over the state-of-the-art is achieved in all the tested scenarios. Especially, we showed a better preservation of the old tasks.

Despite the demonstrated improvements, this work also identifies possible further developments. A direction that is worth exploring is to use the autoencoders as data generators rather than relying on the new data. This would give a stronger solution in the situation where the new data does not represent previous distributions well.

Acknowledgment: Rahaf Aljundi's PhD is funded by an FWO scholarship. This work is partially funded by Internal Funds KU Leuven, FP7-MC-CIG 334380, GOA-project CAMETRON and an Amazon Research Award.

References

- [1] R. Aljundi, P. Chakravarty, and T. Tuytelaars. Expert gate: Lifelong learning with a network of experts. *IEEE conference on Computer Vision and Pattern Recognition*, 2017.
- [2] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of International Conference on Computational Statistics*, pages 177–186. Springer, 2010.
- [3] H. Bourlard and Y. Kamp. Auto-association by multilayer perceptrons and singular value decomposition. *Biological cybernetics*, pages 291–294, 1988.
- [4] R. Caruana. Multitask learning. In *Learning to learn*, pages 95–133. Springer, 1998.
- [5] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International Conference on Machine Learning*, volume 32, pages 647–655, 2014.
- [6] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>.
- [7] R. M. French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, pages 128–135, 1999.
- [8] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 580–587, 2014.
- [9] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [10] I. J. Goodfellow, M. Mirza, A. Courville, and Y. Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *stat*, page 6, 2014.
- [11] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *stat*, 1050:9, 2015.
- [12] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, page 201611835, 2017.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [14] Z. Li and D. Hoiem. Learning without forgetting. In *European Conference on Computer Vision*, pages 614–629. Springer, 2016.
- [15] J. L. McClelland, B. L. McNaughton, and R. C. O’reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, page 419, 1995.
- [16] M. McCloskey and N. J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of learning and motivation*, pages 109–165, 1989.
- [17] T. M. Mitchell. *The need for biases in learning generalizations*. Department of Computer Science, Laboratory for Computer Science Research, Rutgers Univ. New Jersey, 1980.
- [18] M.-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*, 2008.
- [19] A. Pentina and C. H. Lampert. Lifelong learning with non-iid tasks. In *Advances in Neural Information Processing Systems*, pages 1540–1548, 2015.
- [20] A. Quattoni and A. Torralba. Recognizing indoor scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 413–420. IEEE, 2009.
- [21] R. Ratcliff. Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological review*, pages 285–308, 1990.
- [22] S.-A. Rebuffi, A. Kolesnikov, and C. H. Lampert. icarl: Incremental classifier and representation learning. *IEEE conference on Computer Vision and Pattern Recognition*, 2017.
- [23] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, pages 211–252, 2015.
- [24] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell. Progressive neural networks. *arXiv:1606.04671*, 2016.
- [25] D. L. Silver and R. E. Mercer. The task rehearsal method of life-long learning: Overcoming impoverished data. In *Conference of the Canadian Society for Computational Studies of Intelligence*, pages 90–101. Springer, 2002.
- [26] D. L. Silver, Q. Yang, and L. Li. Lifelong machine learning systems: Beyond learning algorithms. In *AAAI Spring Symposium: Lifelong Machine Learning*, pages 49–55. Citeseer, 2013.
- [27] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.
- [28] M. D. Zeiler. Adadelata: an adaptive learning rate method. *arXiv:1212.5701*, 2012.