

ENCODERS AND ENSEMBLES FOR TASK-FREE CONTINUAL LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

We present an architecture that is effective for continual learning in an especially demanding setting, where task boundaries do not exist or are unknown, and where classes have to be learned online (with each example presented only once). To obtain good performance under these constraints, while mitigating catastrophic forgetting, we exploit recent advances in contrastive, self-supervised learning, allowing us to use a pre-trained, general purpose image encoder whose weights can be frozen, which precludes forgetting. The pre-trained encoder also greatly simplifies the downstream task of classification, which we solve with an ensemble of very simple classifiers. Collectively, the ensemble exhibits much better performance than any individual classifier, an effect which is amplified through specialisation and competitive selection. We assess the performance of the encoders-and-ensembles architecture on standard continual learning benchmarks, where it outperforms prior state-of-the-art by a large margin on the hardest problems, as well as in less familiar settings where the data distribution changes gradually or the classes are presented one at a time.

1 INTRODUCTION

Supervised learning methods often rely on the assumption that data is identically and independently distributed (i.i.d.) and drawn from a fixed distribution. However, in many applications, as in ordinary human life, this assumption is unwarranted. Instead, data arrives a little at a time from an ever-changing world upon which the learning system has a limited window. The sub-field of *continual learning* studies algorithms and architectures that can cope with this more realistic scenario (Hadsell et al., 2020; Delange et al., 2021). The main difficulty to be overcome in devising neural network architectures for continual learning is that of *catastrophic forgetting*, wherein a model’s performance on earlier tasks in a sequence degrades as it learns new tasks (McCloskey & Cohen, 1989; French, 1999; Kirkpatrick et al., 2017). Neural networks are especially vulnerable to catastrophic forgetting because of the holistic nature of gradient-based weight updates (Hadsell et al., 2020). Each new example that a network encounters has the potential to perturb every weight in the network. If a whole class of examples is absent from the data for a while, weights tuned for that class soon drift. Researchers have addressed this problem in a number of ways, such as consolidating weights, retaining a memory of past experience, dividing the architecture into separate modules, and meta-learning (Hadsell et al., 2020; Delange et al., 2021).

Here we demonstrate that catastrophic forgetting can be mitigated, and current state-of-the-art results surpassed, in the most challenging continual learning setting, using a relatively straightforward combination of well-understood architectural components (Fig. 1). This is possible largely thanks to recent progress in self-supervised learning. Our architecture incorporates a powerful off-the-shelf encoder whose self-supervised pre-training has resulted in a general-purpose feature extractor (Grill et al., 2020; Mitrovic et al., 2021), which greatly simplifies the downstream classification task. (The encoder must be pre-trained on a dataset different from the one used to evaluate continual learning performance, of course.) By freezing the weights of the pre-trained encoder, we ensure that this part of the architecture is immune to forgetting. However, forgetting can still creep in downstream of the encoder. We address this by using an ensemble of very simple single-layer classifiers. Individually, these classifiers are prone to errors, but collectively they perform well. Their collective success is amplified through specialisation, which is achieved by associating each classifier with a random key drawn from the encoder’s latent space, and by selecting a subset of classifiers according to key-

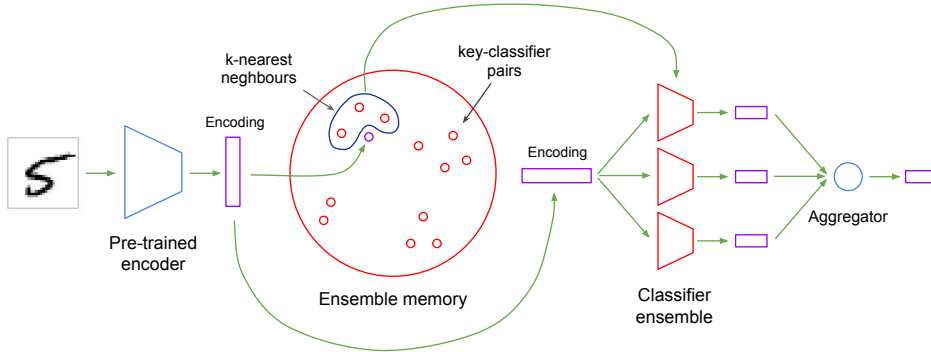


Figure 1: The Ensemble Memory Architecture

matching. Finally, a judicious choice of activation and loss functions further mitigates forgetting by confining updates to weights that affect the class of the current sample.

Much of the continual learning literature makes the unrealistic assumption that the learning problem is constituted by a set of sequentially presented tasks, each comprising data drawn from a different distribution. Indeed, many approaches rely on the existence of well defined task boundaries within the data. By contrast, the most realistic continual learning setting, as well as the most challenging, is one in which no clear task boundaries even exist (Aljundi et al., 2019b;a). This is the task-free continual learning setting we tackle in the present paper (although, for convenience, we still refer to tasks where appropriate). The architecture we present here does not require knowledge of task boundaries, nor does it try to infer them, and it can be applied to continual learning problems where the distribution changes gradually and no clear task boundaries exist. We evaluate the architecture on three standard image classification datasets, namely MNIST, CIFAR-10, and CIFAR-100. We look at three benchmark variations, each tailored for continual learning: 1) task-wise splits, where the classes are partitioned into subsets (tasks) and presented sequentially, 2) the fully incremental case, where the classes are presented one at a time, and 3) a variation in which the classes are drawn from a distribution that shifts gradually over the course of training, presenting an ever-changing blend of classes.

2 AN ENSEMBLE MEMORY ARCHITECTURE

Our architecture comprises a pre-trained image encoder with an ensemble of single-layer classifiers (Fig. 1). (For full details, including hyperparameters, see Appendix.) Each classifier has a fixed associated key, which is used for classifier look-up via k -nearest neighbours. Hence the classifiers can be thought of as residing in an ensemble memory. Crucially, the key space of the ensemble memory is the same as the latent space of the encoder. The architecture works as follows. An image is passed through the encoder, and the resulting encoding is used to look up the k classifiers with the closest keys by cosine similarity. The encoding is then passed through each of these classifiers, and the resulting vectors are aggregated to produce the model’s final output. The classifiers’ output vectors are aggregated by taking their weighted average, where the weighting for each classifier is given by the cosine similarity between the encoding and the classifier’s key. As we shall see, to mitigate catastrophic forgetting, the loss function and the classifiers’ activation functions must be of a specific form.

Let x be an input image and $y \in \{0, 1\}^m$ be the one-hot encoding of the class to which it belongs. Let f be an encoder with $z = f(x) \in \mathbb{R}^d$ the latent encoding. Let a t -classifier (tanh classifier) be a pair (W, b) where $W \in \mathbb{R}^{m \times d}$ is a matrix of trainable weights and $b \in \mathbb{R}^m$ is a vector of trainable biases. The output of a t -classifier is given by

$$v(W, b, z) = [\phi(\psi_1(z)), \dots, \phi(\psi_m(z))] \quad (1)$$

where $\psi_i(z) = w_i \cdot z^T + b_i$ and $\phi(x) = \tau \tanh(x/\tau)$ with w_i the i^{th} row of W . The scaling factor τ is a hyper-parameter. Now, let an ensemble memory of size n be a pair $M = (M_{\text{key}}, M_{\text{cfier}})$ of vectors of keys and t -classifiers, respectively, i.e. $M_{\text{key}} \in \mathbb{R}^{n \times d}$ and $M_{\text{cfier}} \in (\mathbb{R}^{n \times m \times d}, \mathbb{R}^{n \times m})$ representing the weights and biases of the t -classifiers. We denote the i^{th} key in M_{key} and the i^{th}

classifier in M_{cifer} by M_{key}^i and M_{cifer}^i , respectively. Let $\gamma(x, y)$ be the cosine similarity between two vectors x and y , and let $\mathcal{I}(i, z)$ denote the index of the i^{th} -ranked key in M_{key} according to cosine similarity to the encoding z . Then the output of the model is given by

$$V_M(z) = \frac{\sum_{i=1}^k \gamma(M_{\text{key}}^{\mathcal{I}(i,z)}, z) v(W^{\mathcal{I}(i,z)}, b^{\mathcal{I}(i,z)}, z)}{\sum_{i=1}^k \gamma(M_{\text{key}}^{\mathcal{I}(i,z)}, z)} \quad (2)$$

where $M_{\text{cifer}}^{\mathcal{I}(i,z)} = (W^{\mathcal{I}(i,z)}, b^{\mathcal{I}(i,z)})$. We seek to minimise the loss function

$$\mathcal{L}(y, \hat{y}) = -(y \cdot \hat{y}) \quad (3)$$

where $\hat{y} = V_M(f(x))$ (the predicted labels from the model).

As hinted at earlier, several unconventional features of the model are essential to its success. First, the loss function is simply the dot product of the model’s output with a one-hot vector, with no softmax or other form of normalisation applied. Second, the activation function for a t-classifier is tanh with a scaling factor τ applied, which allows the output of a neuron to grow ever closer to τ without ever reaching it. A third unconventional feature of our approach is the choice of optimiser. We found training to be most effective when, for each batch, the magnitude of a parameter’s gradient was discarded having determined its sign, so that each parameter is raised or lowered by a fixed step size (the learning rate).

The pre-trained encoder f is an essential component of the architecture’s design. Here we employ two kinds of encoder. For the easier MNIST dataset, a simple variational autoencoder (VAE) is sufficient (Kingma & Welling, 2014; Rezende et al., 2014). The encoder half of our VAE comprises two convolutional layers followed by two linear layers, while the mirror-image decoder comprises two linear layers followed by two transpose convolutional layers. Crucially, for the architecture to count as performing continual learning, the encoder must be pre-trained on a *different dataset* to the one used during the continual learning phase itself. In our case, for evaluation on MNIST, we pre-trained the VAE to reconstruct Omniglot characters (Lake et al., 2015), then discarded the decoder half. The encoder half, with its weights frozen, was retained for the full model.

For more challenging datasets than MNIST, a more capable pre-trained encoder is needed. Recently there has been considerable progress in the use of self-supervised contrastive learning to pre-train image encoders (Grill et al., 2020; Chen et al., 2020; Mitrovic et al., 2021). We obtained results using both ReLIC (Mitrovic et al., 2021) and BYOL (Grill et al., 2020), in each case with a ResNet-50 encoder (He et al., 2016) pre-trained on the ImageNet dataset (Deng et al., 2009). Best performances were obtained using ReLIC, where we closely followed the setup in Mitrovic et al. (2021), using the same hyperparameters and optimisation settings. The results obtained with BYOL were also good, although not quite as impressive as with ReLIC (see Table A5). By leveraging a comparison to both similar and dissimilar points, ReLIC learns a latent space in which representations are more tightly clustered according to the latent classes, whereas BYOL only focuses on a comparison to similar points and thus learns a latent space which is less tightly clustered.

3 EVALUATION

Benchmarks. We evaluated our architecture on three standard image classification datasets: MNIST, CIFAR-10, and CIFAR-100. For MNIST and CIFAR-10, we trained models on three continual learning benchmarks: 5-way split, 10-way split (one class at a time), and a Gaussian schedule (Fig. A6). (For full details of each protocol, see Appendix.) In the conventional (i.i.d.) setting, a model is presented with samples drawn uniformly from every class in the training set throughout training, but in these continual learning benchmarks, samples are drawn from a subset of classes that changes as training proceeds. In the 5-way split, the class labels are partitioned into five subsets of two labels each (five tasks) and the model is presented with one subset at a time. Confusingly, the term “split MNIST” has been used by different authors to designate different protocols. As Aljundi, *et al.* point out, “comparing reported results in continual learning requires great diligence because of the plethora of experimental settings” (Aljundi et al., 2019a). For example, a number of authors (eg: Zenke et al. (2017); Farajtabar et al. (2020)) use the term “split MNIST” to designate a protocol in which models are tested on five binary classification tasks, where the first task is to distinguish between digits 0 and 1, the second task is to distinguish between digits 2 and 3, and so on. By

Table 1: Accuracy (%) (higher is better). Results for GEN-MIR and ER-MIR are taken from Aljundi et al. (2019a). Results for CN-DPM are taken from Lee et al. (2020). Results for GDumb are taken from Prabhu et al. (2020). For our models, means and standard deviations for 20 runs are shown.

	MNIST 5-way split (low data)	MNIST 5-way split	MNIST 10-way split	MNIST Gaussian schedule	CIFAR-10 5-way split	CIFAR-10 10-way split	CIFAR-10 Gaussian schedule	CIFAR-100 20-way split	CIFAR-100 100-way split	CIFAR-100 Gaussian schedule
GEN-MIR	82.1 ± 0.3	—	—	—	—	—	—	—	—	—
ER-MIR	87.6 ± 0.7	—	—	—	47.6 ± 1.1	—	—	—	—	—
CN-DPM	—	93.23	—	—	45.21	—	—	20.10	—	—
GDumb	91.9 ± 0.5	—	—	—	61.3 ± 1.7	—	—	—	—	—
Vanilla classifier	58.2 ± 3.5	61.7 ± 3.8	71.7 ± 2.5	56.8 ± 8.1	23.2 ± 4.8	10.6 ± 2.2	11.4 ± 2.4	3.8 ± 0.5	1.0 ± 0.2	6.9 ± 2.9
Tanh classifier	73.8 ± 2.4	76.7 ± 1.9	67.9 ± 3.8	64.9 ± 6.7	11.8 ± 2.7	9.3 ± 2.3	9.6 ± 1.3	1.9 ± 0.3	1.0 ± 0.2	3.7 ± 1.7
Ensemble	83.5 ± 1.3	91.0 ± 0.4	92.0 ± 0.5	84.2 ± 4.6	79.0 ± 0.4	78.3 ± 0.4	50.1 ± 9.5	55.3 ± 0.4	54.1 ± 0.5	39.0 ± 1.4

contrast, following Aljundi et al. (2019a) and Lee et al. (2020), for our 5-way split benchmark, we calculate accuracies for 10-way classification, since there are 10 classes in the dataset. Hsu et al. (2018) provide a clear account of these different evaluation protocols. (See also van de Ven & Tolia (2019); Prabhu et al. (2020).) According to their taxonomy, we perform “incremental class learning” in the 5-way split benchmark, also known as the “single head” or “shared classifier” setting (van de Ven & Tolia, 2019), which, as they show, is the most difficult.

The other two benchmarks we use are of a less common type. In the *fully incremental* 10-way split, the model receives samples from one class label at a time. To be successful at this benchmark, the model has to learn the first class having never seen examples from any other class, and then to correctly re-identify members of that class throughout training. The 5-way and 10-way split benchmarks both feature distinct tasks and clear task boundaries. However, our model makes no use of this fact, and works in the more natural task-free setting. The *Gaussian schedule* benchmarks test for this capability. A *schedule*, in this context, defines how a data distribution evolves over training. In the Gaussian schedule, there are no task boundaries. Rather, each label appears in the data with a probability that follows a bell curve, and each label’s probability peaks at a different time (see Fig. A6). In our case, the probability peaks corresponding to the ten classes are evenly spread out over training, and the respective curves overlap significantly, so that one class blends smoothly into the next. For the CIFAR-100 dataset, we trained the model on three benchmarks: 20-way split, 100-way split, and a Gaussian schedule. In the 20-way split benchmark, the dataset’s 100 classes are partitioned into 20 subsets of five classes each, according to their coarse class labels. In the fully incremental 100-way split, the classes are presented and learned one at a time, analogously to 10-way split MNIST and 10-way split CIFAR-10. For both benchmarks, accuracies are calculated for 100-way classification.

Findings. Our main findings are summarised in Table 1, alongside previous state-of-the-art performances, where applicable. To reiterate, all final accuracies are for n -way classification, where n is the total number of class labels. So $n = 10$ for MNIST and CIFAR-10, while $n = 100$ for CIFAR-100. The models we compare with are the Maximal Interfered Retrieval (MIR) technique of Aljundi et al. (2019a), the Continual Neural Dirichlet Process Mixture (CN-DPM) method of Lee et al. (2020), and the Greedy Sampler and Dumb Learner (GDumb) method of Prabhu et al. (2020). MIR is a replay-based method that improved on prior art, such as Lopez-Paz & Ranzato (2017), by selectively, rather than randomly, sampling from memory. It comes in two variants: GEN-MIR, which samples from encodings produced by a generative model, and ER-MIR, which directly samples stored images. CN-DPM is a mixture-of-experts style approach, where the set of “experts” (Dirichlet processes) expands dynamically to accommodate shifts in data distribution. GDumb is another a replay-based method, but unlike methods that mitigate forgetting by fine-tuning on stored samples, in GDumb the whole model is re-trained from scratch on all the stored samples at inference time. The precludes forgetting by design, but at the cost of inference-time computation.

Following Aljundi, *et al.* and Lee, *et al.*, our training regime is “online” in the sense that it does not involve multiple passes through the data (multiple epochs). This contrasts with, for example, iCaRL, whose reported results were obtained with multiple epochs (Rebuffi *et al.*, 2017; van de Ven & Tolias, 2019). Where Aljundi, *et al.*, Lee, *et al.*, and Prabhu *et al.* use a common protocol (5-way split CIFAR-10), we followed suit to obtain our results. Where they use different protocols (5-way split MNIST), we produced results for the protocols used by each set of authors (ie: for both a low data regime and for the full training set). In the case of CIFAR-100, only Lee, *et al.* supply a previous result. It should be born in mind that, as far as methods are concerned, we are not comparing like-with-like. For example, our method exploits the availability of a pre-trained encoder, unlike the other methods, while GDumb has a large computational overhead at inference time, unlike other methods (including ours). Claims of state-of-the art should be taken in this context. The ultimate value of the comparison is to contribute to a map of design possibilities for architectures for continual learning.

For 5-way split MNIST, our ensemble model’s performance is slightly below the previous best accuracy with the full training set (CN-DPM), and falls between GEN-MIR and ER-MIR for the low data regime. The best performance in for MNIST in this setting is obtained by GDumb, albeit with a heavy price tag in terms of computation at inference time. We note that MNIST is an “easy” dataset for which state-of-the-art accuracies are already high, and the pre-trained encoder we used to obtain MNIST results was a simple VAE. Our real focus here is the sort of realistic colour images exemplified by the CIFAR datasets, for which current state-of-the-art accuracies remain low. This is where powerful encoders pre-trained using contemporary self-supervised methods come into their own. For 5-way split CIFAR-10, the ensemble model achieves a final accuracy of 79.0%, which improves on ER-MIR (state-of-the-art in 2019) by over 30% and on GDumb by 18% (previous best reported accuracy). Finally, for 20-way split CIFAR-100, the most demanding dataset, our model achieves a final accuracy of 55.3%, which also represents an improvement on prior state-of-the-art (CN-DPM) by over 30%. We additionally report accuracies for the fully incremental (10-way or 100-way) split for each of MNIST, CIFAR-10, and CIFAR-100, as well as accuracies for the Gaussian schedule for each dataset. The performance of the ensemble model on the 10-way split is within 1% of its performance on the 5-way split for both MNIST and CIFAR-10. For 100-way CIFAR-100, the ensemble model attains 54.1%, which is close to its performance on the 20-way split. The Gaussian schedule presents more of a challenge for each dataset, although in every case the ensemble method is the best performer.

Overall, our findings show that the encoders-and-ensembles approach is effective for task-free continual learning, and that it outperforms other models on harder datasets. We note that, in addition to the differences between methods mentioned above, we have not compared like-with-like in terms of memory requirements; our model incorporates an encoder with a large number of parameters, and includes a memory-hungry ensemble. That said, even a small ensemble is sufficient to achieve state-of-the-art results (see ablations below), and replay-based methods also have extra memory requirements, thanks to their need for a replay buffer. Moreover, in other branches of machine learning, progress is often made by building bigger models, reflecting the ever-increasing memory capacity of commodity hardware. Although optimising memory usage is clearly a worthwhile aim, the continual learning sub-field also needs to explore ways to make good use of more memory when it is available.

Forgetting. In addition to classification accuracy, we assessed each model’s propensity to forget what it has learned during the course of training. To do this, we generalise the measure defined in Chaudhry *et al.* (2018) to cover the task-free setting. We define the *generalised forgetting* of a model over a training period with n time steps (batches) as

$$\frac{1}{m} \sum_{i=1}^m \max_t (a_t^i - a_n^i) \quad (4)$$

where m is the number of classes and a_t^i is the accuracy of the model for class i at time t . The forgetting metrics for our experiments are reported in Table 2. As with accuracy, we show results for previous state-of-the-art where applicable. (We note that generalised forgetting, which we report for our model, is a strict upper bound on task-wise forgetting, as reported in Aljundi *et al.* (2019c). Lower forgetting is better, of course.) We found that the ensemble model outperforms the state-of-the-art where previous figures were reported. On 5-way split MNIST, our model yields 6.2% (generalised) forgetting (albeit with high variance), compared to 7.0% forgetting for ER-MIR, while on 5-way split CIFAR-10, our model yields 7.5% compared to 17.4% for ER-MIR.

Table 2: Forgetting (lower is better). Results for GEN-MIR and ER-MIR are taken from Aljundi et al. (2019a). Means and standard deviations for 20 runs are shown.

	MNIST 5-way split (low data)	MNIST 5-way split	MNIST 10-way split	MNIST Gaussian schedule	CIFAR-10 5-way split	CIFAR-10 10-way split	CIFAR-10 Gaussian schedule	CIFAR-100 20-way split	CIFAR-100 100-way split	CIFAR-100 Gaussian schedule
GEN-MIR	17.0 ± 0.4	—	—	—	—	—	—	—	—	—
ER-MIR	7.0 ± 0.9	—	—	—	17.4 ± 2.1	—	—	—	—	—
Vanilla classifier	35.1 ± 4.4	31.7 ± 4.1	20.5 ± 2.6	34.1 ± 6.9	57.3 ± 4.8	20.9 ± 5.1	42.5 ± 12.9	7.6 ± 1.3	4.2 ± 0.5	22.6 ± 3.8
Tanh classifier	8.9 ± 2.9	7.1 ± 2.3	11.8 ± 3.9	12.9 ± 5.7	14.9 ± 4.8	20.7 ± 3.7	31.9 ± 7.9	3.5 ± 0.9	3.6 ± 0.6	10.0 ± 2.6
Ensemble	6.2 ± 1.7	3.2 ± 0.6	3.3 ± 0.5	6.2 ± 3.4	7.5 ± 0.9	8.7 ± 1.0	20.8 ± 9.0	6.6 ± 0.5	12.4 ± 0.7	14.6 ± 1.8

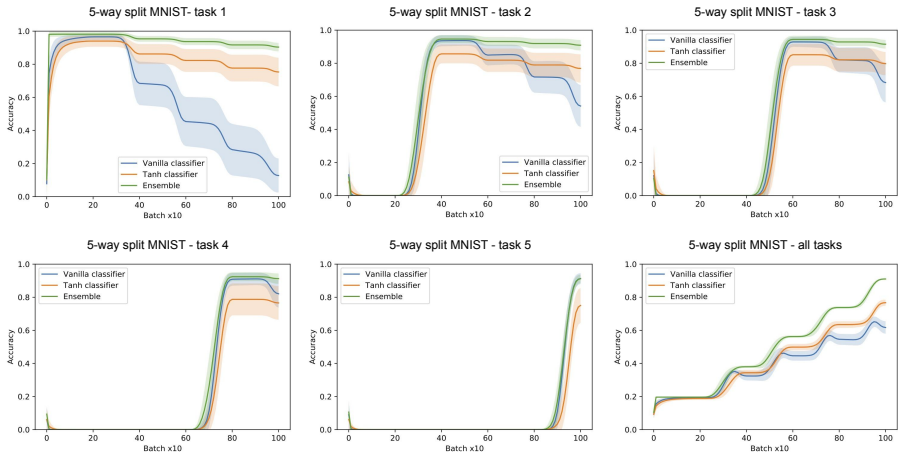


Figure 2: Task-wise accuracies over training for 5-way split MNIST

3.1 BASELINES AND ABLATIONS

For a more qualitative appreciation of the ensemble model in action, see Fig. 2, bottom right, and Fig. 3, bottom right, which show the evolution of overall accuracy on all ten classes on the test set for 5-way split MNIST and 5-way split CIFAR-10, respectively. The model’s progress is steadily upwards, and where there are discrete changes in distribution (new tasks) – in the 5-way and 10-way splits – their effects are clearly visible. These figures also show the performance of two non-ensemble baseline models, one in which the output of the pre-trained encoder is delivered to a single vanilla classifier with a conventional softmax applied, and a similar model with a single t-classifier

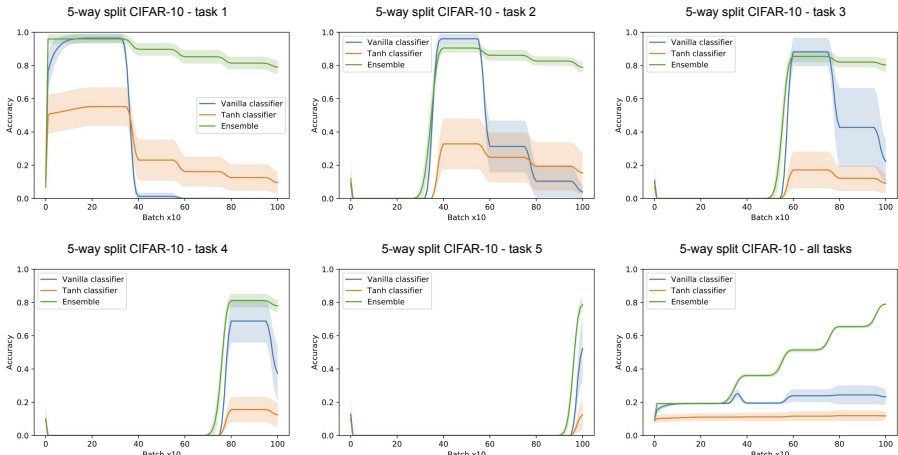


Figure 3: Task-wise accuracies over training for 5-way split CIFAR-10

(with a tanh activation function and no softmax). (See Table 1 for numerical final accuracies.) For the easier MNIST dataset (Fig. 2), forgetting is non-catastrophic even with the vanilla classifier, and is further reduced with the stand-alone tanh classifier, although neither model matches the performance of the ensemble. This suggests that each of the three elements of the architecture – the pre-trained encoder, the activation / loss function combination, and the ensemble – can help to reduce catastrophic forgetting for a simple enough dataset. However, with the harder CIFAR-10 dataset (Fig.3), catastrophic forgetting is mitigated only with the full ensemble model, incorporating all three of these features.

The remaining plots in Figs. 2 and 3 depict task-wise accuracies for 5-way split MNIST and 5-way split CIFAR-10, respectively. The propensity for forgetting in the vanilla classifier is clear, although it is more rapid for CIFAR-10 than MNIST. This contrasts with the ensemble model, which exhibits much less forgetting with both datasets. The stand-alone tanh classifier exhibits markedly less forgetting than the vanilla classifier for MNIST, but it attains a worse initial performance on each task, putting its final accuracy between that of the vanilla classifier and the ensemble model. A similar trade-off is apparent for CIFAR-10 with the stand-alone classifier, where its initial performance on each task is especially poor.

We carried out a number of further experiments to assess the extent to which ensemble size and top- k classifier selection contribute to the model’s success. We found that the ensemble model is still effective even with a significantly reduced ensemble size, though its performance degrades (Table A4, left). Even with a “small” ensemble (128 classifiers, $k = 8$), the model still beats ER-MIR on 5-way split CIFAR-10 in terms of final accuracy by 17%, and is comparable to GDumb, which has large overheads in terms of both memory and inference-time computation. This is encouraging, as an increase in memory requirements is one of the costs of our approach. The best performance, however, is still achieved with the largest ensemble (1024 classifiers), as shown in Table 1. To study the impact of top- k classifier selection, we looked at the effect of larger values for k . We found that performance slowly degrades for $k > 32$ (Table A4, right), which suggests that the competition and consequent specialisation induced by top- k selection is working as expected.

4 HOW DOES THE MODEL WORK?

Here we offer some intuition for how the model works (visualised in Fig. A4). Given a pre-trained encoder with fixed weights, the success of the model depends on two further features that operate in tandem: the dot-product loss function / tanh activation function combination, and the ensemble. Let’s first consider a single t-classifier. How, in the context of a fixed encoder, does the combination of a dot-product loss function and a tanh activation function mitigate catastrophic forgetting? Let f be the frozen, pre-trained encoder, and let v_i be the i^{th} neuron in the t-classifier v (representing class i). (So $v_i(x) = \phi_i(\psi_i(x))$) (Eqn. 1).) Let’s consider $E(i, j) = \mathbb{E}(v_i(f(x)) | x \in C_j)$, the expected value output by neuron i given an input image x from class C_j . How does this expectation evolve during training? Because of the form of the loss and activation functions, given an image from class $j \neq i$, the gradients of the weights on incoming connections to v_i are always zero. So incoming images from any class $j \neq i$ will not affect those weights, and therefore will not affect $E(i, j)$, and we can safely ignore this case. But what about incoming images from class i itself? Given an image from class i , the weights on incoming connections to v_i will be adjusted to push $v_i(f(x))$ up, which will drive up $E(i, i)$ over time. Let’s call these *targeted* increases in E . Targeted increases in $E(i, i)$ are good, as they steer the model towards correct classifications. However, as a side-effect of increasing $E(i, i)$, $E(i, j)$ where $j \neq i$ can also increase. Let’s call these *collateral* increases in E . Collateral increases in $E(i, j)$ where $j \neq i$ are bad, as they draw the model away from correct classifications.

Crucially, over time, as long as the data is roughly class-balanced, any collateral increases in $E(i, j)$ where $j \neq i$ will tend to be outpaced by the targeted increases in $E(i, i)$. Consequently, $E(i, i) - E(j, i)$ will tend to increase for all i and j , which is what we really care about since this is what determines the predicted class after taking $\arg \max_i v_i(f(x))$. Now, the tendency for targeted increases in E to outpace collateral increases in E is sometimes enough for a single t-classifier to mitigate catastrophic forgetting, as we see with the MNIST benchmarks in Fig. 2. However, as this desirable upward trend is rather weak, it isn’t sufficient to prevent poorly performing t-classifiers from arising, as they do most of the time for the more challenging CIFAR-10 and CIFAR-100 bench-

marks. This is where the ensemble comes in. The weak statistical trend we’re looking for can be amplified by training many t-classifiers. Averaged over a population of t-classifiers, the tendency for targeted increases to outpace collateral increases will be greater than for a single t-classifier.

Moreover, an ensemble approach allows for specialisation, which further improves performance. In our architecture, this is achieved through top- k , key-based classifier lookup, where they keys reside in the same space as the image encodings. Any degree of class-relevant clustering in the encoder’s latent space will confer an advantage here, further magnifying the statistical trend described above. To see this, recall that we are interested in increases in $E(i, i) - E(j, i)$ for all i and j . We have already seen that, for a standalone t-classifier, the targeted increases in $E(i, i)$ will tend to increase faster than the collateral increases in $E(i, j)$. Now consider a t-classifier from the ensemble. Assuming the latent space exhibits a degree of class-relevant clustering, the chances of the same classifier appearing in the top- k set for two given images will be higher for images belonging to the same class than for images from different classes. This entails that the classifier will tend to be updated more often through images of some classes than others. Now suppose the classifier is more likely to be updated by images from class i than class j . This entails that not only will the targeted increases in $E(i, i)$ tend to be larger per training step than the collateral increases in $E(i, j)$, they will also be more frequent. We further enhance this effect by taking the weighted average of the output of the top- k classifiers, giving more weight to classifiers whose keys are close to an image’s encoding, in other words to those that are more likely to have received more training on similar images.

5 RELATED WORK

Ensembles. The literature on continual learning is extensive (see Hadsell et al. (2020) and Delange et al. (2021) for recent reviews), so we focus here on related work that either a) uses techniques similar to ours (especially ensemble methods), or b) addresses the task-free scenario. First, we consider ensemble methods, which have been studied widely within machine learning Sagi & Rokach (2018). Ensembles of one sort or another feature in several prior approaches to continual learning. For example, a number of authors have applied mixture of experts models Jacobs et al. (1991) to continual learning, which explicitly incorporate ensembles Aljundi et al. (2017); Kruszewski et al. (2021); Lee et al. (2020). Others use ensembles in a more implicit way, in the form of sub-networks, to improve continual learning. For example, in Fernando et al. (2017), a genetic algorithm is used to discover an effective sub-network for a task, which is then frozen and can be reused when training on the next task. In other work, dropout Srivastava et al. (2014), which creates implicit sub-networks by silencing neurons at random during training, has been shown to mitigate catastrophic forgetting in a continual learning setting Goodfellow et al. (2013); Farajtabar et al. (2020). In Wen et al. (2020), a parameter-efficient ensemble method is described where a rank-one weight matrix is learned per task and used to scale a weight matrix learned on the first task which is shared among ensemble members.

Replay. Replay-based methods are a major category in the taxonomy of approaches to continual learning, and certain replay-based methods are naturally applicable to the task-free setting. For example, in Isele & Cosgun (2018); Rolnick et al. (2019), reservoir sampling is used to maintain a uniform sample of all previously seen data in the replay buffer and does not require any knowledge of task boundaries. Other methods mitigate forgetting by curating the contents of the buffer Aljundi et al. (2019c) or by being selective about which examples to replay Aljundi et al. (2019a) in a task-free manner. Certain “pseudo-rehearsal” methods that train generative models to mimic past data are also suited to the task-free setting Shin et al. (2017); Kamra et al. (2017). Memory-based parameter adaptation Sprechmann et al. (2018) stores a record of past examples like a replay method, but at test-time uses k -nearest neighbour look-up to retrieve similar examples to the current input, and uses them to locally adapt network weights. The authors show that a few locally targeted updates restore the network’s ability to deal with the retrieved exemplars, and hence with the current input. Our architecture does not employ a replay buffer. However, replay methods, such as those cited, are typically compatible with our approach, since our architecture is indifferent to where incoming data comes from. An outer loop that injected examples drawn from a replay buffer, either before or after the pre-trained encoder, would be a straightforward addition.

Meta-learning. Some meta-learning approaches to continual learning do not require the incoming data to be split up into discrete tasks. In Javed & White (2019), a representation-learning network is

meta-trained to minimise catastrophic forgetting on sequences of tasks (e.g. split Omniglot) when combined with a prediction network trained in the inner loop. In Beaulieu et al. (2020), a neuromodulatory network is meta-trained instead, and used to gate the prediction network. In He et al. (2019), a network is meta-trained to infer task representations from incoming data which are then used to condition another network that performs the task at hand. Like our method, these meta-learning approaches involve a kind of pre-training, in the form of the meta-training phase. Other methods exist that are agnostic to the *location* of task boundaries, but nevertheless rely on the fact that the data is piecewise stationary and try to detect task boundaries Kirkpatrick et al. (2017); Aljundi et al. (2019b).

Gradual distribution change. A number of other authors have presented continual learning methods that can handle gradual distribution change or soft task boundaries, and have proposed benchmarks akin to our Gaussian schedule to test this Rao et al. (2019); Zeno et al. (2018); Lee et al. (2020). The CURL approach (Continual Unsupervised Representation Learning) Rao et al. (2019) is a generative replay method similar to GEN-MIR Aljundi et al. (2019a) that can carry out unsupervised clustering in a continual learning setting without knowledge of task boundaries, and can be tailored to supervised learning in a task-free incremental class learning setting as used to evaluate our model. The authors also show that CURL is effective for unsupervised clustering of MNIST digits in the context of a gradually shifting distribution similar to our Gaussian schedule, although they don't report accuracies for supervised learning in this setting, nor do they apply their model in an online (single epoch) setting, as we do.

6 DISCUSSION

We have presented an architecture that combines three straightforward ideas – a pre-trained encoder, an ensemble of simple classifiers, and a particular activation / loss function pairing – and shown the effectiveness of this combination in the particularly demanding continual learning setting where task boundaries are either unknown or absent. We have shown that, for a simple dataset like MNIST, each of these features independently helps to mitigate catastrophic forgetting in this setting, but they are most potent when working together. Used together, this combination of features proves dramatically more effective on harder datasets, such as CIFAR-10 and CIFAR-100, than other methods.

Ensemble methods belong to a family of architectures with a long pedigree in artificial intelligence (Selfidge, 1959; Minsky, 1988; Nii, 1986; Jacobs et al., 1991; Shanahan, 2006). All of these architectures feature sets of parallel, independent modules or processes that compete and / or co-operate with each other to collectively determine a system's behaviour. The profound benefits of such architectures – which are manifest in the biological brain as well as in engineered systems – can be summarised in three maxims. 1) Competition: selection from a pool of processes or modules encourages specialisation for different contexts, and the right specialist in the right context will outperform a generalist (“jack of all trades, master of none”). 2) Co-operation: when independent processes or modules specialise in different aspects of a situation, their expertise can be combined in a compositional fashion, which aids generalisation (“divide and rule”). 3) Collectivity: the aggregated contributions of a diverse set of separate processes or modules yields better performance than any single monolithic process can (“the wisdom of the crowd”). In this respect, the choice of an ensemble method to deal with catastrophic forgetting is not *ad hoc*, but is part of a larger picture wherein modular architectures are used to address some of the deepest problems in building AI that approaches human-level intelligence.

Of course, our method is not without limitations. Chief among these is its reliance on a pre-trained encoder. But recent advances in self-supervised learning have helped to make the case for general-purpose, off-the-shelf encoders that can be applied to datasets different from the one they were trained on. Our work further strengthens this case. For new types of data (for audio, say, or language), a new encoder is needed to make use of our method. Yet if we consider the human capacity for lifelong learning – something today's AI systems can only aspire to – this idea seems natural. If a child is shown a picture book and learns the names of a series of animals it has never seen before, it starts with a perceptual system that has been thoroughly pre-trained on several years of previous visual experience. Similar remarks apply to sound, to touch, to the whole panoply of human experience. Indeed, the ability to build on what has been learned before is the essence of continual learning.

REFERENCES

- Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3366–3375, 2017.
- Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *Advances in Neural Information Processing Systems*, volume 32, pp. 11849–11860, 2019a.
- Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019b.
- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019c.
- Shawn Beaulieu, Lapo Frati, Thomas Miconi, Joel Lehman, Kenneth O. Stanley, Jeff Clune, and Nick Cheney. Learning to continually learn. In *Proceedings 24th European Conference on Artificial Intelligence (ECAI)*, pp. 992–1001, 2020.
- Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings European Conference on Computer Vision (ECCV)*, pp. 556–572, 2018.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)*, pp. 1597–1607, 2020.
- Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 248–255, 2009.
- Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *ArXiv preprint arXiv:1701.08734*, 2017.
- Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999.
- Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *ArXiv preprint arXiv:1312.6211*, 2013.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 21271–21284, 2020.
- Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in Cognitive Sciences*, 24(12):1028–1040, 2020.
- 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 (CVPR)*, pp. 770–778, 2016.

- Xu He, Jakub Sygnowski, Alexandre Galashov, Andrei A Rusu, Yee Whye Teh, and Razvan Pascanu. Task agnostic continual learning via meta learning. *ArXiv preprint arXiv:1906.05201*, 2019.
- Yen-Chang Hsu, Yen-Cheng Liu, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *ArXiv preprint arXiv:1810.12488*, 2018.
- David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991.
- Khurram Javed and Martha White. Meta-learning representations for continual learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Nitin Kamra, Umang Gupta, and Yan Liu. Deep generative dual memory network for continual learning. *ArXiv preprint arXiv:1710.10368*, 2017.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ArXiv preprint arXiv:1312.6114*, 2014.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Canadian Institute for Advanced Research, 2009.
- Germán Kruszewski, Ionut-Teodor Sorodoc, and Tomas Mikolov. Evaluating online continual learning with calm. *ArXiv preprint arXiv:2004.03340*, 2021.
- Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(1):2278–2324, 1998.
- Soochan Lee, Junsoo Ha, Dongsu Zhang, and Gunhee Kim. A neural dirichlet process mixture model for task-free continual learning. In *International Conference on Learning Representations (ICLR)*, 2020.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of Learning and Motivation*, volume 24, pp. 109–165. Elsevier, 1989.
- Marvin L. Minsky. *The Society of Mind*. Simon & Schuster, 1988.
- Jovana Mitrovic, Brian McWilliams, Jacob C Walker, Lars Holger Buesing, and Charles Blundell. Representation learning via invariant causal mechanisms. In *International Conference on Learning Representations (ICLR)*, 2021.
- H. Penny Nii. The blackboard model of problem solving and the evolution of blackboard architectures. *AI Magazine*, 7(2):38–53, 1986.
- Ameya Prabhu, Puneet Kumar Dokania, and Philip H.S. Torr. GDumb: A simple approach that questions our progress in continual learning. In *European Conference on Computer Vision (ECCV)*, 2020.

- Dushyant Rao, Francesco Visin, Andrei Rusu, Razvan Pascanu, Yee Whye Teh, and Raia Hadsell. Continual unsupervised representation learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- 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 (CVPR)*, pp. 2001–2010, 2017.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning (ICML)*, pp. 1278–1286, 2014.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Omer Sagi and Lior Rokach. Ensemble learning: A survey. *WIREs Data Mining and Knowledge Discovery*, 8(4):e1249, 2018.
- Oliver Selfidge. Pandemonium: A paradigm for learning. In *Proceedings of the Symposium on Mechanisation of Thought Processes*, pp. 511–529. Her Majesty’s Stationery Office, 1959.
- Murray Shanahan. A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition*, 15(2):433–449, 2006.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Pablo Sprechmann, Siddhant Jayakumar, Jack Rae, Alexander Pritzel, Adria Puigdomenech Badia, Benigno Uria, Oriol Vinyals, Demis Hassabis, Razvan Pascanu, and Charles Blundell. Memory-based parameter adaptation. In *International Conference on Learning Representations (ICLR)*, 2018.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning. *ArXiv preprint arXiv: 1904.07734*, 2019.
- Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. In *International Conference on Learning Representations (ICLR)*, 2020.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning (ICML)*, pp. 3987–3995, 2017.
- Chen Zeno, Itay Golan, Elad Hoffer, and Daniel Soudry. Task agnostic continual learning using online variational bayes. *ArXiv preprint arXiv:1803.10123*, 2018.

A1 APPENDIX

A1.1 CODE

Code to accompany this paper is available from `xxxx`. This is a Colab notebook, using the JAX library, that works for split MNIST using a pre-trained autoencoder, as described above. The notebook is best run with a GPU to obtain reasonable speed. Running the notebook should reproduce the bottom, right-hand plot of Fig. 2. To get good results for other datasets (eg: CIFAR-10), use an off-the-shelf pre-trained encoder such as BYOL or ReLIC. Code for BYOL is available from <https://github.com/deepmind/deepmind-research/tree/master/byol>.

A1.2 VISUALISING THE WORKING OF THE ARCHITECTURE

Fig. A4 depicts a representative training run for 5-way split MNIST (low data regime). Final layer activations are shown for two sample digits at different points in training. Magenta denotes the correct digit. Red denotes the maximal activation where different from the correct digit. See the discussion in Section 4.

A1.3 ENCODER ARCHITECTURE AND PRE-TRAINING

MNIST. For the MNIST experiments, a variational autoencoder (VAE) was pre-trained on the Omniglot training set. The *encoder* half of the autoencoder comprised two convolutional layers (kernel size = 4, no. of channels = 16, activation function = ReLU) followed by two heads, one for the encoding means, and the other for the encoding standard deviations. Each head comprised two linear layers. The first linear layer had an output size of 128 with a ReLU activation function, and the second produced the latent encoding with size 512. The final activation function for the means was tanh, and for the standard deviations was ReLU. The *decoder* half of the autoencoder was the mirror

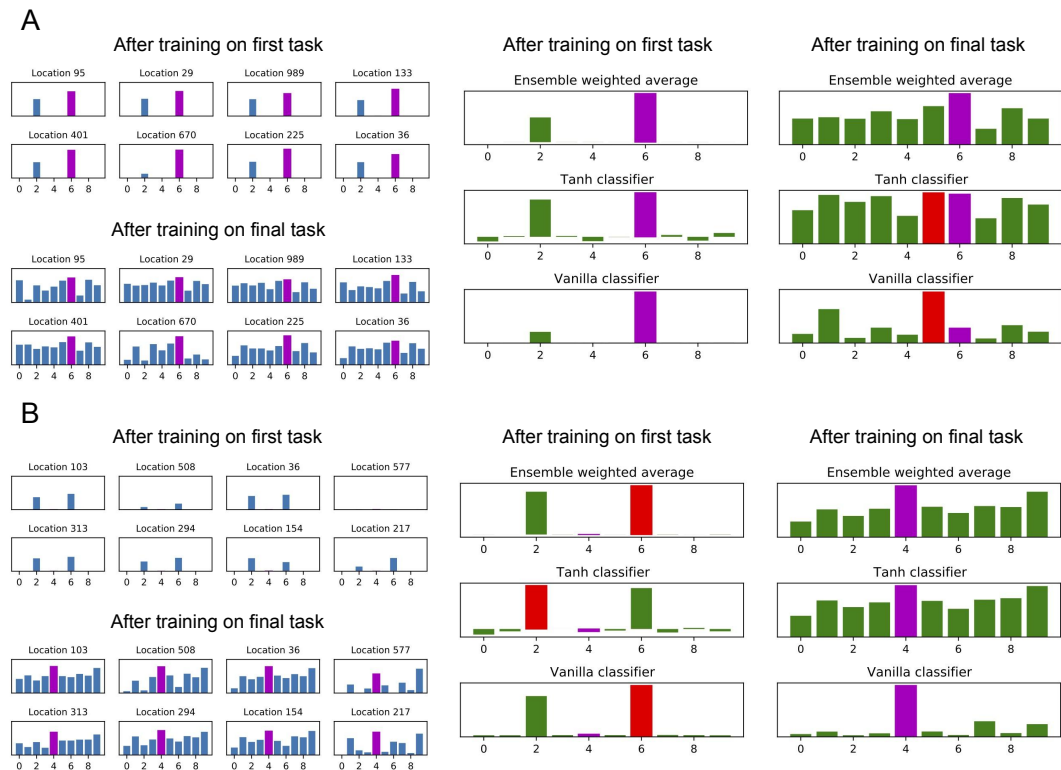


Figure A4: Visualisation of a representative training run for 5-way split MNIST (low data regime). (A): The sample digit (8) belongs to the first task. (B): The sample digit (3) belongs to the final task.

image of the encoder, comprising two linear layers (with 128 and $28 \times 28 \times 16 = 12844$ output channels, respectively) followed by two transpose convolutional layers (kernel size = 4, no of channels = 16 and 1, respectively, activation function = ReLU and sigmoid, respectively). The autoencoder was trained on 10,000 batches of Omniglot images (batch size = 48), using Adam (learning rate = 0.001), minimising the usual VAE loss function

$$\|x - \hat{x}\|^2 + \beta \text{KL}(\mathcal{N}(\mu_x, \sigma_x), \mathcal{N}(0, 1)) \quad (\text{A1})$$

where x is the input image, \hat{x} is the output of the decoder given input $z \sim \mathcal{N}(\mu_x, \sigma_x)$, and (μ_x, σ_x) is the output of the encoder given image x . We set $\beta = 0.001$. We noted, by visual inspection alone, that the resulting model produced satisfactory reconstructions of both Omniglot and MNIST characters, and didn't attempt to tune the model further. Once trained, the decoder half was discarded, along with the standard deviation head of the encoder. The weights of the remaining portion of the encoder were frozen for subsequent use in the main ensemble model. Prior to each training run of the main model, we pre-trained the autoencoder from scratch (on Omniglot, of course, not MNIST), to ensure statistical variation. Autoencoders were re-trained if the reconstruction loss $\|x - \hat{x}\|^2$ exceeded a certain threshold (0.025).

CIFAR-10 and CIFAR-100. For the CIFAR-10 and CIFAR-100 experiments, we used a ResNet-50 encoder He et al. (2016) pre-trained on the ImageNet dataset Deng et al. (2009). After the pre-training, the weights of the encoder are frozen. The results reported in the main paper were obtained with the ReLIC training method Mitrovic et al. (2021). We also obtained results using the BYOL method Grill et al. (2020) (A5). In both cases, we used the settings and hyperparameters described in the respective original publications cited. In both cases, we upsampled the CIFAR-10 and CIFAR-100 images to a size of 128×128 using cubic interpolation. The size of the encoder output is 2048.

The success of the model depends on the ability of the pre-trained encoder to map images into a latent space where unseen classes are to a degree linearly separable. Fig. A5 shows tSNE plots for batches of 1000 images from the MNIST and CIFAR-10 datasets after being passed through the relevant pre-trained encoders. In each case, we can see that the encodings exhibit a degree of class-relevant clustering.

A1.4 ENSEMBLE MODEL, BASELINES, AND TRAINING

Ensemble model. Encoded images are passed on to the ensemble for classification. Each classifier in the ensemble comprises a single linear layer followed by a scaled tanh activation function, as described in the main paper. The weights and biases of the classifiers in the ensemble are the model's only trainable parameters. Weights were initialised using fan-in variance scaling with a scaling factor of 1.0 drawn from a truncated normal distribution. Biases were initialised to 0. Each classifier is paired with a unique key, which is used for k -nearest neighbour lookup, using cosine similarity, as described in the main paper. The keys, which are fixed throughout training, were drawn from a standard normal distribution. The ensemble model was trained using a "naive" optimiser with weight decay, with learning rate 0.0001 and decay factor 0.0001. The naive optimiser works by discarding the magnitude of a gradient having computed its sign, then updating the associated parameter upwards or downwards (according to the sign) by the learning rate. With the exception of

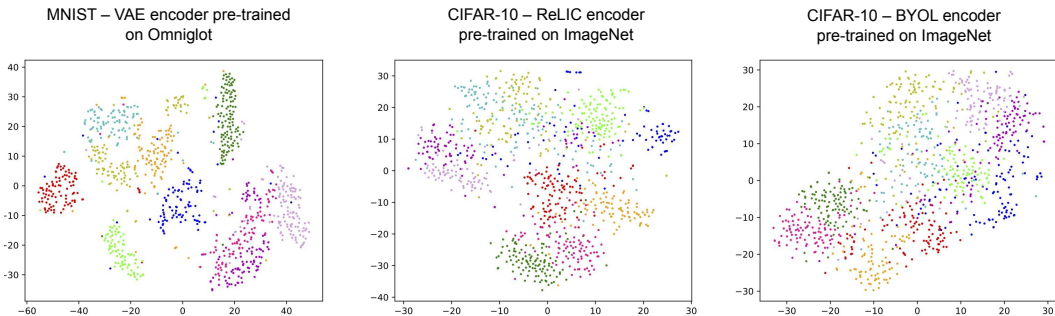


Figure A5: tSNE Plots for image encodings

Table A3: Hyperparameters

Parameter	Value
Learning rate	0.0001
Weight decay	0.0001
Ensemble size	1024
k (top- k selection)	32
τ (tanh scaling factor)	250
Runs per experiment	20

ablations, these and all other hyperparameters, were used for all experiments described in the paper (Table A3). We found performance to be robust to hyperparameter variation, and we did not have to tune them for different datasets or experimental settings.

Baselines. Results for two baseline models are reported throughout the paper: a tanh classifier and a vanilla classifier. These baseline models take the same input as the ensemble mode, namely the output of the pre-trained encoder, but they both consist of a single, stand-alone classifier comprising one linear layer. The only difference between the tanh classifier and the vanilla classifier is their activation functions. The tanh classifier uses a scaled tanh, with the same scaling factor τ as the classifiers in the ensemble (Table A3, while the vanilla classifier uses a conventional log-softmax function. Weights were initialised using fan-in variance scaling and a truncated normal distribution, as for the ensemble, and biases were initialised to 0. We found that a larger variance scaling factor of 10.0 improved performance for the baselines, so we adopted this. As for the ensemble model, the baselines were trained using a naive optimiser with weight decay, with learning rate 0.0001 and decay factor 0.0001.

A1.5 EXPERIMENTAL SETUP AND PROTOCOLS

We use standard MNIST, CIFAR-10, and CIFAR-100 datasets LeCun et al. (1998); Krizhevsky (2009). All accuracies were calculated over the full set of class labels (ie: 10 for MNIST and CIFAR-10 and 100 for CIFAR-100). CIFAR-100 accuracies are for the top class label predicted by the model (not top-5, as reported in some papers; our accuracies would be higher if they were top-5). All reported accuracies and all plots in the paper are for the relevant held-out test set. 20 training runs of every experiment were performed, and mean accuracies are reported in tables along with standard deviations, or shown in plots along with error bands of width one standard deviation. For all three datasets, and for every benchmark, we trained the ensemble for exactly 1000 batches. For MNIST, the batch size was five (low data regime) or 60 (otherwise). For CIFAR-10 and CIFAR-100, the batch size was 48. Only one pass through the dataset (one epoch) was performed, corresponding to the “online” setting. (To allow for schedules with gradual distribution shift, every batch was drawn at random from the dataset, after shuffling. So the model is likely to have seen a number of images more than once. However, the total number of images seen in all cases was less than or equal to the size of the training set, in other words a single epoch.)

5-way split MNIST. The ten class labels were partitioned at random into a sequence of five subsets (or tasks). Each training run was carried out on a new random partitioning of labels. (We note that some authors report results for where the partitioning is the same in every run, usually [0, 1, 2, 3, ...].) In each training run, the five tasks were presented sequentially, each task comprising 200 batches randomly drawn from the relevant subset of the training set. The total number of images seen is $5 \times 200 \times 60 = 60000$, which is equivalent to a single pass through the data (one epoch).

5-way split MNIST (low data). The protocol for this benchmark is identical to that for 5-way split MNIST, but with a smaller batch size (5). The total number of images seen is $5 \times 200 \times 5 = 5000$, which is equivalent to a 1/12 of a pass through the data (1/12 of an epoch).

10-way split MNIST. This is the “fully incremental” benchmark, where the classes are learned one at a time. The ten class labels were partitioned into a sequence of ten singleton subsets (or tasks), each containing just one label. Each training run was carried out on a new random partitioning of labels (ie: a new random ordering of the digits). In each training run, the ten tasks were presented sequentially, each task comprising 100 batches randomly drawn from the relevant subset of the

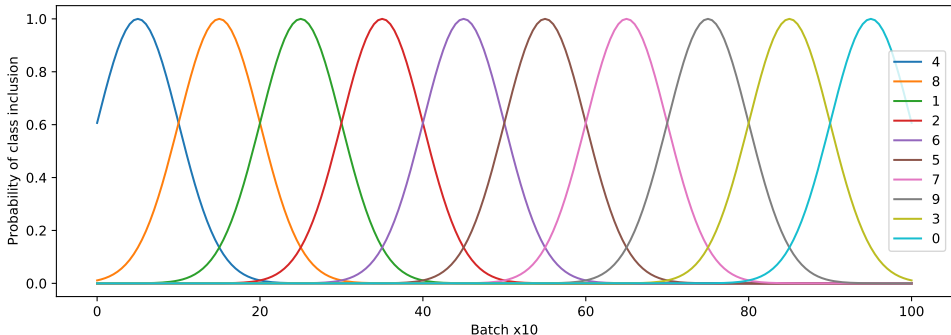


Figure A6: Gaussian schedule

training set. The total number of images seen is $10 \times 100 \times 60 = 60000$, which is equivalent to a single pass through the data (one epoch).

Gaussian schedule MNIST. In this benchmark, the data distribution shifts gradually as training proceeds. A series of 200 “micro-tasks” T_1 to T_{200} were presented to the model, each comprising 5 batches (of 60 images each). Each micro-task T_i was associated with a random set of class labels $L(T_i)$, and all images in that micro-task were drawn at random from the corresponding subset of the training set. The composition of $L(T_i)$ was determined by the following algorithm, where $m = 10$ classes (see Fig. A6). The m class labels are partitioned into a sequence of m singleton subsets C_1 to C_{10} , each containing just one label. (Fig. A6 illustrates one such partitioning.) Each of the ten classes C_i to C_{10} is then associated with a conditional probability distribution $P(C_i|B)$, where B is a batch number. $P(C_i|B)$ is given by a Gaussian function of height 1 centred at $10i$ with width 50. In other words,

$$P(C_i|B) = h.e^{-\frac{(B-10i)}{2w^2}} \quad (\text{A2})$$

where $h = 1$ and $w = 50$. The total number of images seen is $100 \times 10 \times 60 = 60000$, which is equivalent to a single pass through the data (one epoch).

5-way split CIFAR-10. The protocol for this benchmark is identical to that of 5-way split MNIST. The total number of images seen is $5 \times 200 \times 48 = 48000$, which is equivalent to a slightly less than a single pass through the data (less than one epoch). (We chose a batch size of 48 rather than 50 for a fair comparison with Aljundi et al. (2019a), who train on 48750 images for split CIFAR-10.)

10-way split CIFAR-10. The protocol for this benchmark is identical to that of 10-way split MNIST. The total number of images seen is $10 \times 100 \times 48 = 48000$, which is equivalent to a slightly less than a single pass through the data (less than one epoch).

Gaussian schedule CIFAR-10. The protocol for this benchmark is identical to that of Gaussian schedule MNIST. The total number of images seen is $100 \times 10 \times 48 = 48000$, which is equivalent to a slightly less than a single pass through the data (less than one epoch).

20-way split CIFAR-100. In the CIFAR-100 dataset each image has an associated fine class label (one of 100) and super-class label (one of 20). If two images belong to the same fine class, then they must belong to the same super-class. For the 20-way split CIFAR-100 benchmark, the 100 fine class labels were partitioned into a randomly ordered sequence of 20 subsets (tasks), with each subset corresponding to one of the super-classes. In each training run, the 20 tasks were presented sequentially, with each task comprising 50 batches (batch size 48) randomly drawn from the relevant subset of the training set.

100-way split CIFAR-100. This is the “fully incremental” benchmark, where the classes are learned one at a time. The 100 fine class labels were partitioned into a sequence of 100 singleton subsets (or tasks), each containing just one label. Each training run was carried out on a new random partitioning of labels (ie: a new random ordering of the digits). In each training run, the 100 tasks were presented sequentially, each task comprising 10 batches (batch size 48) randomly drawn from the relevant subset of the training set.

Table A4: Ablations. Left: decreasing the ensemble size results in worse performance. Right: increasing k in top- k classifier selection results in worse performance (ensemble size is 1024). All results are final accuracy (%) over 20 runs for all ten classes.

	CIFAR-10 5-way split	CIFAR-10 10-way split	CIFAR-10 Gaussian schedule	MNIST 5-way split (low data)	
Small ensemble (128 classifiers)	64.6 ± 1.3	60.3 ± 1.9	35.9 ± 8.6	$k=64$	84.0 ± 1.0
Tiny ensemble (16 classifiers)	36.4 ± 7.9	23.9 ± 4.5	16.0 ± 4.7	$k=128$	83.3 ± 1.1
				$k=192$	82.7 ± 1.0

Table A5: BYOL encoder results.

	CIFAR-10 i.i.d.	CIFAR-10 5-way split	CIFAR-10 10-way split	CIFAR-10 Gaussian schedule
Vanilla classifier	84.1 ± 0.1	22.0 ± 5.2	11.0 ± 2.2	11.8 ± 2.6
Tanh classifier	57.7 ± 4.0	12.0 ± 1.7	9.8 ± 1.7	10.1 ± 1.3
Ensemble	76.3 ± 0.7	77.7 ± 0.5	76.8 ± 0.6	46.0 ± 9.9

Gaussian schedule CIFAR-100. The protocol for this benchmark is the same as for Gaussian MNIST and Gaussian CIFAR-10, and uses the same algorithm to generate a schedule of micro-tasks but with $m = 100$.

We also evaluated the model (and baselines) for each dataset in the i.i.d. setting (Table A6). The setup here is the same as for all the other benchmarks: the images are fed through the pre-trained encoder and the resulting encodings are passed on the the ensemble (or baseline classifier). However, each batch is drawn uniformly from the set of all classes.

A1.6 ABLATIONS AND OTHER EXPERIMENTS

We carried out a number of ablations, as described in the main text, relating to ensemble size and k in top- k lookup. The results are shown in Table A4. We also carried out the CIFAR-10 and CIFAR-100 experiments using the BYOL encoder rather than ReLIC, yielding the results in Table A5. Table A6 shows accuracies for all benchmarks using the i.i.d. protocol. Finally, Fig. A5 shows tSNE plots illustrating the clustering of the encoded images with each type of pre-trained encoder we used.

Table A6: i.i.d. accuracies.

	MNIST i.i.d.	CIFAR-10 i.i.d.	CIFAR-100 i.i.d.
Vanilla classifier	88.2 ± 0.4	83.4 ± 0.1	46.4 ± 1.1
Tanh classifier	80.8 ± 0.9	61.9 ± 2.9	27.1 ± 3.2
Ensemble	86.7 ± 0.5	77.5 ± 0.6	48.0 ± 0.8