
Meta-Consolidation for Continual Learning

K J Joseph and Vineeth N Balasubramanian

Department of Computer Science and Engineering
Indian Institute of Technology Hyderabad, India
{cs17m18p100001,vineethnb}@iith.ac.in

Abstract

The ability to continuously learn and adapt itself to new tasks, without losing grasp of already acquired knowledge is a hallmark of biological learning systems, which current deep learning systems fall short of. In this work, we present a novel methodology for continual learning called MERLIN: Meta-Consolidation for Continual Learning. We assume that weights of a neural network ψ , for solving task t , come from a meta-distribution $p(\psi|t)$. This meta-distribution is learned and consolidated incrementally. We operate in the challenging online continual learning setting, where a data point is seen by the model only once. Our experiments with continual learning benchmarks of MNIST, CIFAR-10, CIFAR-100 and Mini-ImageNet datasets show consistent improvement over five baselines, including a recent state-of-the-art, corroborating the promise of MERLIN.

1 Introduction

The human brain is able to constantly, and incrementally, consolidate new information with existing information, allowing for quick recall when needed [5, 78]. In this natural setting, it is not common to see the same data sample multiple times, or even twice at times. Human memory capacity is also limited which forbids memorizing all examples that are seen during its lifetime [39]. Hence, the brain operates in an *online manner*, where it is able to adapt itself to continuously changing data distributions without losing grasp of its previously acquired knowledge [33]. Unfortunately, deep neural networks have been known to suffer from catastrophic forgetting [53, 25], where they fail to retain performance on older tasks, while learning new tasks.

Continual learning is a machine learning setting characterized by its requirement to have a learning model incrementally adapt to new tasks, while not losing its performance on previously learned tasks. Note that ‘task’ here can refer to a set of new classes, new domains (e.g. thermal, RGB) or even new tasks in general (e.g. colorization, segmentation) [63, 68, 89]. The last few years have seen many efforts to develop methods to address this setting from various perspectives. One line of work [88, 38, 47, 2, 16, 71, 56, 13, 79, 42] constrains the parameters of the deep network trained on Task A to not change much while learning a new Task B, while another - replay-based methods - store [63, 48, 14, 66, 9, 3, 4, 22] or generate [74, 81, 46, 45] examples of previous tasks to finetune the final model at evaluation time. Another kind of methods [51, 72, 69, 19, 68, 62, 87] attempt to expand the network to increase the capacity of the model, while learning new tasks. Broadly speaking, all these methods manipulate the data space or the weight space in different ways to achieve their objectives.

In this work, we propose a different perspective to addressing continual learning, based on the latent space of a weight-generating process, rather than the weights themselves. Studies of the human brain suggest that knowledge and skills to solve tasks are represented in a meta-space of concepts with a high-level semantic basis [28, 10, 50]. The codification from tasks to concepts, and the periodic consolidation of memory, are considered essential for a transferable and compact representation of knowledge that helps humans continually learn [11, 85, 5]. Current continual learning methods consolidate (assimilate knowledge on past tasks) either in the weight space [13,

38, 47, 16, 88, 56, 2, 71] or in the data space [14, 63, 74, 48, 3, 4, 45, 66]. Even meta-learning based continual learning methods that have been proposed in the recent past [34, 24, 6, 64], meta-learn an initialization amenable for quick adaptation across tasks, similar to MAML [23], and hence operate in the weight space. We propose MERLIN: Meta-Consolidation for Continual Learning, a new method for continual learning that is based on consolidation in a meta-space, viz. the latent space which generates model weights for solving downstream tasks.

We consider weights of a neural network ψ , which can solve a specific task, to come from a meta-distribution $p(\psi|\mathbf{t})$, where \mathbf{t} is a representation for the task. We propose a methodology to learn this distribution, as well as continually adapt it to be competent on new tasks by consolidating this meta-space of model parameters whenever a new task arrives. We refer to this process as “Meta-Consolidation”. We find that *continually learning in the parameter meta-space with consolidation* is an effective approach to continual learning. Learning such a meta-distribution $p(\psi|\mathbf{t})$ provides additional benefits: (i) at inference time, any number of models can be sampled from the distribution $\psi_{\mathbf{t}} \sim p(\psi|\mathbf{t})$, which can then be ensembled for prediction in each task (Sec 4.1.5); (ii) it is easily adapted to work in multiple settings such as class-incremental and domain-incremental continual learning (Sec 4); and (iii) it can work in both a task-aware setting (where the task is known at test time) and task-agnostic setting where the task is not known at test time (achieved by marginalizing over \mathbf{t} , Sec 5.2). Being able to take multiple passes through an entire dataset is an assumption that most existing continual learning methods make [2, 38, 47, 88, 63, 12]. Following [4, 3, 14, 48], we instead consider the more challenging (and more natural) online continual learning setting where ‘only a single pass through the data’ is allowed. We compare MERLIN against a recent state-of-the-art GSS [4], as well as well-known methods including GEM [48], iCaRL [63] and EWC [38] on Split MNIST [13], Permuted MNIST [88], Split CIFAR-10 [88], Split CIFAR-100 [63] and Split Mini-Imagenet [15] datasets. We observe consistent improvement across the datasets over baseline methods in Sec 4.

The key contributions of this work can be summarized as: (i) We introduce a new perspective to continual learning based on the meta-distribution of model parameters, and their consolidation over tasks arriving in time; (ii) We propose a methodology to learn this distribution using a Variational Auto-encoder (VAE) [37] with learned task-specific priors, which also allows us to ensemble models for each task at inference; (iii) We show that the proposed method outperforms well-known benchmark methods [48, 63, 38], as well as a recent state-of-the-art method [4], on five continual learning datasets; (iv) We perform comprehensive ablation studies to provide a deeper understanding of the proposed methodology and showcase its usefulness. To the best of our knowledge, MERLIN is the first effort to incrementally learn in the meta-space of model parameters.

2 Related Work

In this section, we review existing literature that relate to our proposed methodology from two perspectives: continual learning and meta-learning. We also discuss connections to neuroscience literature in the Appendix.

Continual learning methods: In continual learning, when a new task t_k comes by, the weights of the associated deep neural network, ψ_{t_k} , get adapted for the new task causing the network to perform badly on previously learned tasks. To overcome this issue (called catastrophic forgetting [53, 25]), one family of existing efforts [2, 38, 47, 88] force the newly learned weight configuration ψ_{t_k} , to be close to the previous weight configuration $\psi_{t_{k-1}}$, so that the performance on both tasks are acceptable. This approach, by design, can restrict the freedom of the model to learn new tasks. Another set of methods [14, 63, 48, 12] store a few samples (called exemplars) in a fixed-size, task-specific episodic memory, and use strategies like distillation [31] and finetuning to ensure that ψ_{t_k} performs well on all tasks seen so far. These methods work at the risk of memorization. Shin *et al.* [74] instead used a generative model to synthesize examples for all the previous tasks, which allows generating infinite data for the seen tasks. However, as more tasks are added, the capacity of the model reduces, and the generative model does not work well in practice. A recent group of methods [51, 72, 69] attempt to expand the network dynamically to accommodate learning new tasks. While this is an interesting approach, the model size in such methods can increase significantly, hampering scalability.

The aforementioned methods operate in an offline fashion, wherein once data for a specific task is available, it can be iterated over multiple times to learn the task. In contrast, online continual learning methods [4, 3, 14, 48] tackle a more challenging setting, closer to how humans operate, where all datapoints are seen only once during the lifetime of the model. [4, 3] proposes methods to choose




Figure 1: Overview of MERLIN. While learning a new task τ_k at time k , B models are trained and saved to Ψ_t . These weights are used to learn a task-specific parameter distribution with task-specific priors in Step 3. In Step 4, the model is consolidated by replaying parameters generated using earlier task-specific priors as in Algorithm 2. At test time, model parameters are sampled from $p_\theta(\phi|z, t)$ and ensembled.

optimal training datapoints for learning continually, which they found is critical in an online continual learning setting, while [14, 48] constrains the learning of the new task, such that the loss on the previously learned class should not increase. We refer the interested reader to [17, 59] for a detailed survey of current continual learning methods and practices. Contrary to these methods, we propose a new perspective to continual learning based on meta-consolidation in the parameter space which has its own benefits as mentioned earlier.

Meta-learning methods: Meta-learning encapsulates a wide variety of methods that can adapt to new tasks with very few training examples. Meta-learning algorithms - which mostly have focused on the standard learning setting (not continual learning) - can be broadly classified into black-box adaptation based methods [70, 55], optimization based [23, 57, 90] and non-parametric methods [83, 75, 77]. Hypernetworks [27] is one black-box adaptation method where a meta-neural network is used to predict the weights of another neural network. This opened up an interesting research direction, in exploiting the meta-manifold of model parameters for multi-task learning [67], neural architecture search [8], zero-shot learning [58, 82] and style transfer [73] in recent years. One can view our approach as having a similar view to online continual learning (which has not been done before), although our methodology of learning the distribution of weights of the network is different from [27] and allows more flexibility in sampling and generating multiple models at inference. Mandivarapu *et al.* [52] auto-encode the weights in a continual setting, but the size of their meta-model is many times the size of the task network and does-not operate in an online setting. The work closest to ours is by Johannes *et al.* [84], which recently tried using HyperNetworks in the context of continual learning. We differ from [84] in the following ways: (i) we operate in the more natural and challenging *online* continual learning setting, while they use multiple passes over the data (ii) they learn a deterministic function which when conditioned on a task embedding, generates target weights, while we model the weight distribution itself, which allows us to sample as many model parameters at inference time and ensemble them.

3 MERLIN: Methodology

We begin by summarizing the overall methodology of MERLIN, as in Algorithm 1 and Fig 1. We consider a sequence of tasks $\tau_1, \tau_2, \dots, \tau_{k-1}$ that have been seen by the learner, until now. A new task τ_k is introduced at time instance k . (Note that task can refer to a new class or a domain, both of which are studied in Sec 4.) Each task, $\tau_j, j \in \{1, \dots, k\}$ consists of $\tau_j^{tr}, \tau_j^{val}$ and τ_j^{test} corresponding to training, validation and test samples for that task. Each task τ_j is also represented by a corresponding vector t_j , as described in Sec 3.1. As the first step, we train a set of B base models on random subsets of τ_k^{tr} (sampled without replacement, note that this does not violate the online setting as this is done by choosing each streaming point with a certain probability in each base model) to obtain a collection of models $\Psi_k = \{\psi_k^1, \dots, \psi_k^B\}$ (as in first section of Fig 1, line 2 in Algorithm 1). Ψ_k is then used to learn a task-specific parameter distribution $p_\theta(\psi_k|t_k)$ using a VAE-like strategy further described in Sec 3.1. This is followed by a meta-consolidation phase, where we sample model parameters for all tasks seen so far, from the decoder of the VAE, each conditioned on a task-specific prior, and use them to refine the overall VAE, as described in Sec 3.2. At inference, we sample models from the parameter distributions for each task, and use them to evaluate on the test data following the methodology described in Sec 3.3. A set of ψ s are sampled from $p_\theta(\psi|z, t)$ for task t and ensembled to obtain the final result. We now describe each of the components below.

3.1 Modeling Task-specific Parameter Distributions: Each task τ_j is represented by a corresponding vector representation t_j . Our framework allows any fixed-length vector representation

Algorithm 1 MERLIN: Overall Methodology

Input: Tasks: $\mathcal{T} = \{\tau_1 \dots \tau_k\}$; Task training data: $\tau_j^{tr} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_j^{tr}}$; # of base models: B

- 1: **for** $j = 1$ to $|\mathcal{T}|$ **do**
- 2: $\Psi_j \leftarrow$ Train B models by randomly sub-sampling training data from τ_j^{tr} .
- 3: Learn task parameter distribution and task specific prior using methodology in Sec 3.1 and Ψ_j .
- 4: Consolidate using methodology in Sec 3.2, Algorithm 2.
- 5: Perform inference on \mathcal{T} , using methodology in Sec 3.3, Algorithm 3.

for this task descriptor, including semantic embeddings such as GloVe [61] or Word2Vec [54], or simply a one-hot encoding of the task. A richer embedding such as Task2Vec [1], which takes into account task correlations, may also be used. Details of our implementation are deferred to Sec 4. We consider ψ_t , the weights of a neural network which can solve a task t , as generated by a random process given by a continuous latent variable z , i.e. ψ_t is generated from $p_{\theta^*}(\psi|z, t)$, where z is sampled from a conditional prior distribution $p_{\theta^*}(z|t)$ (We denote the j^{th} task, t_j as t for simplicity of explaining the model in this subsection). θ^* refers to the unknown optimal parameters of the weight-generating distribution, which we seek to find. We achieve this objective using a VAE-like formulation [37], adapted for this problem. Computing the marginal likelihood of the weight distribution $p_{\theta}(\psi|t) = \int p_{\theta}(\psi|z, t)p_{\theta}(z|t)dz$ is intractable as its true posterior $p_{\theta}(z|\psi, t) = \frac{p_{\theta}(\psi|z, t)p_{\theta}(z|t)}{p_{\theta}(\psi|t)}$ is intractable to compute. We introduce the approximate posterior $q_{\phi}(z|\psi, t)$, parametrized by ϕ , as a variational distribution for the intractable true posterior $p_{\theta}(z|\psi, t)$. The marginal likelihood of the parameter distribution, $\log p_{\theta}(\psi|t)$, can then be written as (please refer the Appendix for the complete derivation):

$$\log p_{\theta}(\psi|t) = D_{KL}(q_{\phi}(z|\psi, t) || p_{\theta}(z|\psi, t)) + \underbrace{\int_z q_{\phi}(z|\psi, t) \log \frac{p_{\theta}(z, \psi|t)}{q_{\phi}(z|\psi, t)}}_{\mathcal{L}(\theta, \phi|\psi, t)} \quad (1)$$

Similar to a VAE, we can maximize the log likelihood by maximizing the lower bound. $\mathcal{L}(\theta, \phi|\psi, t)$ can hence be rewritten as (complete derivation in the Appendix):

$$\mathcal{L}(\theta, \phi|\psi, t) = -D_{KL}(q_{\phi}(z|\psi, t) || p_{\theta}(z|t)) + \mathbb{E}_{q_{\phi}(z|\psi, t)}[\log p_{\theta}(\psi|z, t)] \quad (2)$$

where the second term is the expected negative reconstruction error. The KL divergence term on the RHS forces the approximate posterior of weights to be close to a *task-specific prior* $p_{\theta}(z|t)$. Note that this differentiates this formulation from a conditional VAE [36], where the latents are directly conditioned. In our model, the latents are conditioned indirectly through priors that are conditioned on task vectors. Assuming $q_{\phi}(\cdot)$ and $p_{\theta}(\cdot)$ to be Gaussian distributions, the KL divergence term can be computed in closed form. The second term requires sampling, and the model parameters, ϕ and θ , are trained using the reparameterization trick [37], backpropagation and Stochastic Gradient Descent.

Instead of choosing the prior to be an isotropic multivariate Gaussian $p_{\theta}(z) = \mathcal{N}(z|\mathbf{0}, \mathbf{I})$ as in a vanilla VAE, we use a learned task-specific prior, $p_{\theta}(z|t)$, given by:

$$p_{\theta}(z|t) = \mathcal{N}(z|\mu_t, \Sigma_t); \text{ where } \mu_t = \mathbf{W}_{\mu}^T t \text{ and } \Sigma_t = \mathbf{W}_{\Sigma}^T t \quad (3)$$

Here, \mathbf{W}_{μ} and \mathbf{W}_{Σ} are parameters which are learned alongside ϕ and θ , when maximizing the lower bound (Eqn 2) using backpropagation. We find that a simple linear model is able to learn \mathbf{W}_{μ} and \mathbf{W}_{Σ} effectively, as revealed in our experiments.

As already stated, the probabilistic encoder $q_{\phi}(z|\psi, t)$ and decoder $p_{\theta}(\psi|z, t)$ are materialized using two neural networks parametrized by ϕ and θ respectively. To train these, we consider $\Psi_j = \{\psi_j^i\}_{i=1}^B$, a set of B model parameters that are obtained by training on the j^{th} task by sampling the corresponding training data without replacement. These models are now used to learn the “meta-parameters” ϕ and θ , as well as μ_j and Σ_j of the Gaussian prior, specific to the task. At inference, one can sample as many ψ s (task-specific models) as needed from the decoder using different samples from the task-specific prior. Obtaining a model ψ for each task t from the parameter distribution $p_{\theta}(\psi|z, t)$, is a two-step process: (i) Sample z from task-specific prior distribution i.e. $z \sim \mathcal{N}(z|\mu_t, \Sigma_t)$; (ii) Sample ψ from the probabilistic decoder, using z i.e. $\psi \sim p_{\theta}(\psi|z, t)$. Note that we have only a single VAE model, which can generate network parameters for each task individually by conditioning the VAE on the corresponding task-specific prior.

3.2 Meta-consolidation: When the continual learner encounters the k^{th} task and receives the set of optimal model parameters for this task $\Psi_k = \{\psi_k^i\}_{i=1}^B$, directly updating the VAE on Ψ_k alone causes a distributional shift towards the k^{th} task. We address this by “meta-consolidating” the encoder and decoder, after accommodating the new task. Algorithm 2 summarizes the steps involved in this consolidation. We assume that all the learned task-specific priors are stored and available to us. This adds negligible storage complexity as it involves storing only means and covariances. For each task τ_1, \dots, τ_k seen so far, a task-specific z_t is sampled from the corresponding task-specific prior distribution. Next, P pseudo-models are sampled from the decoder $p_\theta(\psi|z_t, t)$. These generated pseudo-models are used to finetune the parameters of the encoder and the decoder by maximizing the lower bound, $\mathcal{L}(\theta, \phi|\psi, t)$ defined in Eqn 2. Note that we update only θ and ϕ , and not the parameters of the task-specific priors, which are fixed. One could consider this as a “replay” strategy, only using the meta-parameter space. This ensures that after learning each new task, the encoder and decoder are competent enough to generate model parameters for solving all tasks seen till then. Importantly, we do not learn separate encoder-decoders per task and do not store weight parameters (ψ_j) for any task, which makes our proposed method storage-efficient. The consolidation in this step takes place in the task parameter meta-space (unlike earlier methods), which we find to be effective for continual learning, as validated by our experimental results in Sec 4.

Algorithm 2 META-CONSOLIDATION IN MERLIN

Input: Encoder: $q_\phi(z|\psi, t)$; Decoder: $p_\theta(\psi|z, t)$; Last seen task: τ_k ; Task priors: $\{P_i\}_{i=1}^k, P_i = (\mu_i, \Sigma_i)$; # of psuedo-models: P

Output: Consolidated encoder-decoder parameters: ϕ and θ

- 1: **for** $j = 1$ to k **do**
- 2: $\mu_j, \Sigma_j \leftarrow P_j$
- 3: $z_j \sim \mathcal{N}(z|\mu_j, \Sigma_j)$ \triangleright Task specific z_j
- 4: $Loss \leftarrow \sum_{i=1}^P \mathcal{L}(\theta, \phi|\psi_i, t)$; where $\psi_i \sim p_\theta(\psi|z_t, t)$ $\triangleright \mathcal{L}(\cdot)$ is defined in Equation 2.
- 5: $g \leftarrow \nabla_{\theta, \phi} Loss$
- 6: $\phi, \theta \leftarrow$ Update parameters ϕ, θ , using gradient g .
- 7: **return** ϕ, θ

3.3 Inference: Having learned the distribution to generate model parameters for each task seen until now, the learned task-specific priors give us the flexibility to sample any number of models from this distribution at inference/test time. This allows for *ensembling* multiple models at test time (none of which need to be stored a priori), which is a unique characteristic of the proposed method.

Ideally, a continual learning algorithm should be able to solve all tasks encountered so far, without access to any task-specific information at inference time. While a handful of existing methods [29, 63] take this into consideration, many others [48, 38, 88] assume availability of task-identifying information during inference. This is also referred as single-head or multi-head evaluation in literature [13, 80, 21]. Our methodology works with both evaluation settings, which we refer to as *task-agnostic* and *task-aware* inference, as described below. Algorithm 3 outlines the inference steps. We use the

Algorithm 3 MERLIN INFERENCE

Input: Decoder: $p_\theta(\psi|z, t)$; Last seen task: τ_k ; Task priors: $P = \{P_i\}_{i=1}^k, P_i = (\mu_i, \Sigma_i)$; Exemplars: $\mathcal{E} = \{Ex_i\}_{i=1}^m, Ex_i = \{(x_i, y_i)\}$; Number of base models to ensemble from: E

- 1: **if** *Task-agnostic inference* **then** \triangleright Task-agnostic inference
- 2: $z \sim \mathcal{N}(z|\mu, \Sigma)$ where $\mu \leftarrow \frac{1}{k} \sum_{i=1}^k \mu_i$ and $\Sigma \leftarrow \frac{1}{k} \sum_{i=1}^k \Sigma_i$
- 3: $\Psi \leftarrow$ Sample E models from $p_\theta(\psi|z)$
- 4: $\Psi \leftarrow$ Fine-tune Ψ on \mathcal{E}
- 5: Ensemble results from Ψ to solve all tasks (τ_1, \dots, τ_k)
- 6: **if** *Task-aware inference* **then** \triangleright Task-aware inference
- 7: **for** $j = 1$ to k **do**
- 8: $z_j \sim \mathcal{N}(z|\mu_j, \Sigma_j)$ where $\mu_j, \Sigma_j \leftarrow P_j$
- 9: $\Psi_j \leftarrow$ Sample E models from $p_\theta(\psi|z_j, t_j)$
- 10: $\Psi_j \leftarrow$ Fine-tune Ψ_j on Ex_j
- 11: Ensemble results from Ψ_j to solve task τ_j .

consolidated decoder $p_\theta(\psi|z, t)$, and task-specific priors P at test time. In a *task-agnostic* setting, a single overall prior distribution is aggregated from all task-specific priors. We find that a simple averaging of parameters of the prior distributions works well in practice. The latent variable z is sampled from this aggregated prior. The model parameters sampled from the decoder using z has the

capability to solve all tasks seen till then. The generated ψ is amenable for quick adaptation to each task similar to MAML [23] to obtain task-dependent model parameters. To leverage this, we store a small set of randomly selected exemplars for each task, which is used to finetune ψ for each task. E such models are sampled from the decoder, each of which is finetuned on all tasks. An ensemble of these models is used to predict the final output. *Task-aware* inference works very similarly, with the only change that z_j is sampled from each task-specific prior. The task-specific model parameters ψ_j are then sampled from $p_\theta(\psi|z_j, t_j)$, finetuned on task exemplars Ex_j and finally ensembled. The ensembling step adds minimal inference overhead as explained in Sec 4.1.5.

4 Experiments and Results

We evaluate MERLIN against prominent methods for continual learning: GEM [48], iCaRL [63], EWC [38], a recent state-of-the-art GSS [4], and a crude baseline where a single model is trained across tasks (referred to as ‘Single’ in our results). GSS and GEM are inherently designed for online continual learning, while iCaRL and EWC are easily adapted to this setting following [48, 14]. These baselines consolidate in different spaces: EWC [38], a regularization-based method, consolidates in the weight space; GEM [48], GSS [4], and iCaRL [63] use exemplar memory and consolidate in the data space; while MERLIN consolidates in the meta-space of parameters. We used the official implementations of each of these baseline methods for fair comparison. Five standard continual learning benchmarks, viz. Split MNIST [13], Permuted MNIST [88], Split CIFAR-10 [88], Split CIFAR-100 [63] and Split Mini-Imagenet [15], are used in the experiments, following recent continual learning literature [12, 4, 63, 51, 13].

4.1 Experimental Setup: We describe the datasets, evaluation metrics and other implementation details below, before presenting our results. Our code¹ is implemented in PyTorch [60] and runs on a single NVIDIA V-100 GPU.

4.1.1 Datasets: We describe the datasets considered briefly below (a detailed description is presented in the Appendix):

Split MNIST and Permuted MNIST: Subsequent classes in the MNIST [43] dataset are paired together and presented as a task in *Split MNIST*, a well-known benchmark for continual learning. This results in 5 incremental tasks. In *Permuted MNIST*, each task is a unique spatial permutation of 28×28 images of MNIST. 10 such permutations are generated to create 10 tasks. We use 1000 images per task for training and the model is evaluated on the all test examples, following the protocol in [48].

Split CIFAR-10 and Split CIFAR-100: 5 and 10 tasks are created by grouping together 2 and 10 classes from CIFAR-10 [41] and CIFAR-100 [41] datasets respectively. Following [48], 2500 examples are used per task for training. Trained models are evaluated on the whole test set.

Split Mini-Imagenet: Mini-Imagenet [83] is a subset of ImageNet [18] with a total of 100 classes and 600 images per class. Each task consists of 10 disjoint subset of classes from these 100 classes. Similar to CIFAR variants, 2500 examples per task is used for training, as in [15].

We note a distinction between tasks in ‘‘Split’’ versions of the datasets and Permuted MNIST. In ‘‘Split’’ datasets, the label space expands with tasks, while for Permuted MNIST, the data space changes with tasks without changing the label space. The former is referred to as the *Class-Incremental* setting, while the latter as *Domain-Incremental* setting in literature [32, 80]. MERLIN works on both settings.

4.1.2 Base Classifier Architecture: For CIFAR and Mini-ImageNet datasets, a modified ResNet [30] architecture is used, which is 10 layers deep and has fewer number of feature maps in each of the four residual blocks (5, 10, 20, 40). This reduces the number of parameters from $0.27M$ to 34997. In spite of using a weaker base network (owing to computing constraints), our method outperforms baselines, as shown in our results. For the MNIST dataset, we use a two-layer fully connected neural network with 100 neurons each, with ReLU activation, following the experimental setting in GEM [48]. To train these base models (which are then used to train the VAE in MERLIN), batch size is set to 10 and Adam [35] is used as the optimizer, with an initial learning rate of 0.001 and weight decay of 0.001. To ensure the online setting, the model is trained only for a single epoch, similar to baseline methods [48, 4, 3]. For class-incremental experiments, we follow earlier methods [48, 4] to assume an upper-bound on the number of classes to expect, and modify the loss function to consider only classes that are seen so far. This is done by setting the final layer classification logits of the unseen class to a very low value (-10^{10}), as in [48, 14, 62].

¹<https://github.com/JosephKJ/merlin>

| Datasets → | Split MNIST | | Permuted MNIST | | Split CIFAR-10 | | Split CIFAR-100 | | Split Mini-ImageNet | |
|------------|-------------------|------------------|-------------------|------------------|-------------------|-------------------|-------------------|------------------|---------------------|------------------|
| | A (↑) | F (↓) | A (↑) | F (↓) | A (↑) | F (↓) | A (↑) | F (↓) | A (↑) | F (↓) |
| Single | 44.8 ± 0.3 | 98.3 ± 0.5 | 73.1 ± 2.2 | 15.7 ± 1.9 | 73.2 ± 3.1 | 12.6 ± 4.4 | 30.8 ± 3.5 | 20.5 ± 2.6 | 27.5 ± 2.6 | 17.1 ± 2.7 |
| EWC [38] | 45.1 ± 0.1 | 98.4 ± 0.2 | 74.9 ± 2.1 | 12.4 ± 2.5 | 74.2 ± 2.2 | 14.5 ± 3.4 | 29.2 ± 3.3 | 22.5 ± 4.1 | 28.1 ± 2.5 | 18.0 ± 4.6 |
| GEM [48] | 86.7 ± 1.5 | 23.4 ± 1.8 | 82.5 ± 4.9 | 0.8 ± 0.4 | 79.1 ± 1.6 | 5.9 ± 1.7 | 40.6 ± 1.9 | 1.3 ± 1.8 | 34.1 ± 1.2 | 4.7 ± 0.9 |
| iCaRL [63] | 89.9 ± 0.9 | 1.7 ± 1.3 | - | - | 72.6 ± 1.3 | 4.1 ± 1.5 | 27.1 ± 2.9 | 1.2 ± 1.3 | 38.8 ± 1.6 | 3.5 ± 0.6 |
| GSS [4] | 88.3 ± 0.8 | 33.3 ± 2.4 | 81.4 ± 1.2 | 8.8 ± 1.1 | 57.9 ± 2.6 | 49.2 ± 7.6 | 19.1 ± 0.7 | 42.7 ± 1.5 | 14.8 ± 0.9 | 31.3 ± 3.2 |
| MERLIN | 90.7 ± 0.8 | 6.4 ± 1.2 | 85.5 ± 0.5 | 0.4 ± 0.4 | 82.9 ± 1.2 | -0.9 ± 1.9 | 43.5 ± 0.6 | 2.9 ± 3.7 | 40.1 ± 0.9 | 2.8 ± 3.2 |

Table 1: Average accuracy (A) and average forgetting measure (F) of five baseline methods and MERLIN, across five datasets. MERLIN consistently outperforms the baselines across datasets.

4.1.3 Task Descriptors: The inputs to train models for each task are training data points $\tau_k^{tr} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_{tr}^k}$, corresponding to task τ_k . As in Sec 3.1, in order to condition the prior distribution $p_\theta(\mathbf{z}|t)$ on a task, each task, τ_k , is represented by a corresponding fixed-length task descriptor $\mathbf{t}_k \in \mathbb{R}^D$. We use the simplest approach in our experiments: one-hot encoding of the task sequence number. Using semantic task descriptors such as GloVe or Word2Vec can allow our framework to be extended to few-shot/zero-shot continual learning. The label embedding of a zero/few-shot task can be used to condition the prior, which can subsequently generate help decoder models for the zero/few-shot task. We leave this as a direction of future work at this time.

4.1.4 Training the VAE: The first step of MERLIN is to train task-specific classification models from training data τ_j^{tr} using the abovementioned architectures. The weights of these models are used to train the VAE. 10 models are learned for each task by random sampling of subsets (with replacement) from τ_k^{tr} . Considering that base classification models can be large, in order to not make the VAE too large, we use a chunking trick proposed by Johannes *et al.* in [84]. The weights of the base classification models are flattened into a single vector and split into equal sized chunks (last chunk zero-padded appropriately). We use a chunk size of 300 for all experiments, and show a sensitivity analysis on the chunk size in Sec 5.6. The VAE is trained on the chunks (instead of the full models) for scalability, conditioned additionally on the chunk index. At inference, the classifier weights are assembled back by concatenating the chunks generated by the decoder, conditioned on the chunk index. We observed that this strategy worked rather seamlessly, as shown in our results. The approximate posterior $q_\phi(\mathbf{z}|\psi, t)$, is assumed to be a 2-D isotropic Gaussian, whose parameters are predicted using an encoder network with 1 fully connected layer of 50 neurons, followed by two layers each predicting the mean and the covariance vectors. The decoder $p_\theta(\psi|\mathbf{z}, t)$ mirrors the encoder’s architecture. The network that generates the mean and diagonal covariance vectors of the learned prior, as in Eqn 3, is modeled as a linear network. AdaGrad [20] is used as the optimizer with an initial learning rate of 0.001. Batch size is set to 1 and the VAE network is trained for 25 epochs.

4.1.5 Inference: At test time, we sample 30 models from the trained decoder $p_\theta(\psi|\mathbf{z}, t)$ to solve each task. Algorithm 3 shows how these are obtained for task-aware and task-agnostic settings. These models are ensembled using majority voting. An ablation study varying the number of sampled models at inference is presented in Sec 5.4. We ensure that this adds minimal inference overhead by loading sampled weights into the model sequentially and saving only the final logits for ensembling. Hence, only one model is stored at a given time in memory, allowing our framework to scale up effectively. Any additional steps in our methodology is offset by the choice of very small models to train, and our chunking strategy, leading to minimal overhead in training time over the baseline methods. At test time, our method is real-time, and has no difference from baseline methods.

4.1.6 Evaluation Metrics: For all experiments, we consider the *average accuracy* across tasks and *average forgetting measure* as the evaluation criteria, following previous works [13, 48]. Average accuracy ($A \in [0, 100]$) after learning the k^{th} task (τ_k), is defined as $A_k = \frac{1}{k} \sum_{j=1}^k a_{k,j}$; where $a_{k,j}$ is the performance of the model on the test set of task j , after the model is trained on task k . Forgetting is defined as the difference in performance from the peak accuracy of a model learned exclusively for a task and the accuracy of the continual learner on the same task. Average forgetting measure ($F \in [-100, 100]$) after learning the k^{th} task can be defined as $F_k = \frac{1}{k-1} \sum_{j=1}^{k-1} \left(\max_{l \in \{1, \dots, k-1\}} a_{l,j} - a_{k,j} \right)$. We ran each experiment with five different seed values and report the mean and standard deviation.

4.2 Results: Table 1 shows how MERLIN compares against baseline methods on all the considered datasets. iCaRL [63] results are not reported for Permuted MNIST as it is not a Domain-Incremental methodology. We see that the average accuracy (higher is better) obtained by MERLIN shows significant improvement over baseline methods across datasets and settings. The difference is larger

with increasing complexity of the dataset. MERLIN also shows strong improvement in the forgetting metric on most datasets, including a negative value (showing improved performance on earlier tasks, known as positive backward transfer [14]) on Split CIFAR10. While learning a new task, iCaRL uses distillation loss [31] to enforce that logits of exemplars from the previous tasks, do not alter much with the current task learning. This, we believe, helps iCaRL achieve lower forgetting, when compared to other methods. MERLIN is able to outperform iCaRL too, in two out of four datasets, with reduced forgetting. GSS [4] failed on CIFAR-100 and Mini-ImageNet, even after we tried different hyperparameters. The original paper does not report results on these datasets. We did not include comparison with A-GEM [14] as it reports only marginal improvement over GEM [48]. The evolution of test accuracy with addition of tasks, which showed improvement in accuracy over the baselines from the very first task, is presented in the Appendix due to space constraints.

5 Discussions and Analysis

We analyze the effectiveness of MERLIN using more studies, as described below.

5.1 Is the Probabilistic Decoder Learning the Parameter Distribution? The key component of MERLIN is the decoder that models the parameter distribution $p_{\theta}(\psi|z_t, t)$. Should the decoder fail to learn, then each model $\psi_i \sim p_{\theta}(\psi|z_t, t)$ would be an ineffective sample from an untrained decoder. In such a scenario, the only reason why inference (Algo 3) would work would be due to finetuning of these sampled models on the exemplars (L4, L10; Algo 3). In order to study this, we run an experiment where we intentionally skip training the VAE and the consolidation step that follows. With all other components the same, this corresponds to only finetuning with exemplars. The ‘w/o training VAE’ row in Tab 2 shows the results. Consistently, on all datasets, the performance drops significantly while removing this key component. This suggests that the proposed meta-consolidation through the VAE is actually responsible for the performance in Tab 1.

5.2 Task Agnostic vs Task Aware Inference: All results of MERLIN in Tab 1 do not assume task information during evaluation, and hence operate in the task-agnostic setting. MERLIN can work in both task-agnostic and task-aware settings, as shown in the last two rows of Tab 2. Expectedly, access to task information during inference boosts performance for simple datasets like Split MNIST. For other datasets, average accuracy and forgetting measure is almost similar for both. This supports our choice of aggregation strategy of priors across tasks, which captures the task-agnostic setting quite well.

| Datasets → | Split MNIST | | Permuted MNIST | | Split CIFAR-10 | | Split CIFAR-100 | | Split Mini-ImageNet | |
|------------------|-------------|------------|----------------|-----------|----------------|------------|-----------------|-----------|---------------------|-----------|
| | A (↑) | F (↓) | A (↑) | F (↓) | A (↑) | F (↓) | A (↑) | F (↓) | A (↑) | F (↓) |
| w/o training VAE | 44.8 ± 5.7 | 19.9 ± 2.3 | 29.9 ± 4.7 | 5.8 ± 0.2 | 53.5 ± 6.6 | 3.8 ± 1.3 | 15.7 ± 1.63 | 3.8 ± 0.6 | 16.1 ± 1.9 | 3.4 ± 3.1 |
| Task Agnostic | 90.7 ± 0.8 | 6.4 ± 1.2 | 85.5 ± 0.5 | 0.4 ± 0.4 | 82.9 ± 1.2 | -0.9 ± 1.9 | 43.5 ± 0.6 | 2.9 ± 3.7 | 40.1 ± 0.9 | 2.8 ± 3.2 |
| Task Aware | 97.4 ± 0.3 | 0.1 ± 0.3 | 85.1 ± 0.4 | 1.1 ± 0.8 | 82.3 ± 1.1 | -0.4 ± 1.2 | 44.4 ± 2.8 | 1.7 ± 0.7 | 41.8 ± 1.5 | 1.1 ± 0.8 |

Table 2: MERLIN variants: (i) without training VAE and subsequent consolidation (ii) Task-agnostic (iii) Task-aware MERLIN. Task-agnostic MERLIN performs almost well as Task-aware MERLIN on several datasets.

5.3 Visualization of Task-specific Priors and Sampled Models: In Fig 2, we visualize the latent variables sampled from the learned prior distribution $z_t \sim p_{\theta}(z|t)$ (left) - which are two-dimensional themselves, and the 2D t-SNE embeddings [49] of their corresponding model parameters generated using z_t : $\psi_i \sim p_{\theta}(\psi|z_t, t)$ (right). The model is trained on a 10 task split CIFAR-100 dataset. The separation of the latent samples and the task parameters of different models in the learned meta-space supports the usefulness of our method in overcoming catastrophic forgetting in continual learning.




Figure 2: (Left) Plot of 20 samples drawn from each task-specific prior $z_t \sim p_{\theta}(z|t)$. (Right) t-SNE plot of weights obtained from decoder using corresponding z_t ; $\psi_i \sim p_{\theta}(\psi|z_t, t)$. Each color corresponds to 10 diff tasks from split CIFAR100.

5.4 Varying Number of Models for Ensemble: In Sec 4, we sampled 30 models for each task from the decoder used in MERLIN. We

| # of models → | 1 | 30 | 50 | 100 |
|---------------------|------------|------------|------------|------------|
| Split MNIST | 86.6 ± 1.4 | 90.6 ± 0.8 | 91.4 ± 0.3 | 91.6 ± 0.3 |
| Permuted MNIST | 79.7 ± 3.9 | 85.4 ± 0.4 | 85.6 ± 0.3 | 85.7 ± 0.2 |
| Split Mini-ImageNet | 30.1 ± 1.8 | 39.1 ± 2.1 | 40.5 ± 2.9 | 40.9 ± 1.3 |

Table 3: Test accuracy with varying number of models in our inference ensemble.

now vary the number of models in Tab 3. We observe a 3.98%, 6.34% and 8.99% improvement on Split-MNIST, Permuted-MNIST and Split Mini-Imagenet, when using 30 models for ensembling, over using only 1. The performance improvement however is not very significant on further increase beyond 30 models. This suggests that while ensembling is required, the number of models need not be very large. Improving our ensembling strategy (beyond majority voting) is a direction of our future work.

5.5 Storage Requirements: MERLIN requires only parameters of task-specific prior $p_\theta(z|t)$ and decoder of the VAE, at inference time. The encoder is required only if we need to continually learn further tasks. None of the components in the architecture grows with number of tasks, making MERLIN scalable. Quantitative analysis on the size of classifier and VAE (meta-model) is presented in Tab 4. We note that meta-model size is always smaller than classifier, and is $8\times$ smaller than the storage requirement of GSS, GEM, EWC and iCaRL which use ResNet-18 with 272,260 parameters. Similar to GEM, we maintain an exemplar buffer of 200 and 400 for MNIST and other datasets respectively.

| # of params in \rightarrow | Clf | VAE |
|------------------------------|-------|-------|
| Split MNIST | 89610 | 31426 |
| Permuted MNIST | 89610 | 31446 |
| Split CIFAR-10 | 31307 | 17910 |
| Split CIFAR-100 | 34997 | 33810 |
| Mini-ImageNet | 34997 | 33810 |

Table 4: Num of params in classifier and VAE.

5.6 Varying Chunk Size: As in Sec 4.1.4, we split the weight vector into equal-sized chunks before encoding it using VAE. We vary chunk size and report the average accuracy in Tab 5. As chunk size increases, the accuracy drops - we hypothesize this is to do with the capacity of VAE. We use simple architectures in our encoder-decoder, and modeling larger weight chunks (or whole models) may require more complex VAE.

| Chunk Size \rightarrow | 100 | 300 | 1000 | 2000 |
|--------------------------|----------------|----------------|----------------|----------------|
| Split MNIST | 90.8 \pm 0.4 | 90.6 \pm 0.8 | 87.5 \pm 3.7 | 87.4 \pm 2.2 |
| Permuted MNIST | 85.6 \pm 0.4 | 85.5 \pm 0.5 | 83.7 \pm 2.8 | 82.5 \pm 3.6 |
| Split Mini-ImageNet | 40.9 \pm 1.2 | 40.5 \pm 2.9 | 38.8 \pm 2.8 | 37.8 \pm 2.1 |

Table 5: Avg test accuracy when varying chunk size.

5.7 Varying Exemplar Memory Size We vary the memory size of the exemplar buffer for MERLIN, and compare the performance against our best competitors in Table 1 of the main paper: GEM [48] and iCaRL [63] in CIFAR-10 experiments. The results in Table 6 show that the increase in performance for MERLIN is significantly better when compared to GEM and iCaRL. We hypothesize this is because the weights sampled from our parameter distribution use these exemplars better, similar to how other meta-learning approaches like MAML [23] finetune on just a few samples. Using better exemplar selection methods other than random sampling (used in this work) can further enhance MERLIN, and is a direction of our future work.

| Memory Size \rightarrow | 100 | 500 | 1000 | 2000 |
|---------------------------|----------------|----------------|----------------|----------------|
| GEM | 77.4 \pm 2.6 | 79.9 \pm 1.9 | 80.9 \pm 1.9 | 80.5 \pm 1.5 |
| iCaRL | 72.5 \pm 2.6 | 73.6 \pm 1.3 | 73.7 \pm 2.6 | 74.8 \pm 1.3 |
| MERLIN | 77.9 \pm 1.3 | 81.9 \pm 0.3 | 86.9 \pm 0.6 | 88.4 \pm 0.8 |

Table 6: Results of varying memory size of exemplar buffer for MERLIN, GEM [48] and iCaRL [63].

6 Conclusion

We introduce MERLIN, a novel approach for online continual learning, based on consolidation in a meta-space of model parameters. Our method is modeled based on a VAE, however adapted to this problem through components such as task-specific learned priors. Our experimental evaluation on five standard continual learning benchmark datasets against five baseline methods brings out the efficacy of our approach. MERLIN can handle both class-incremental and domain-incremental settings, and can work with or without task information at test time. Understanding the dynamics of the latent space of the parameter distribution to further enhance memory retention and extending the methodology to a few-shot setting can be immediate follow-up efforts to MERLIN.

Acknowledgements

We thank TCS for funding Joseph through its PhD fellowship, and DST, Govt of India, for partly supporting this work through the IMPRINT program (IMP/2019/000250). We also thank the members of Lab1055, IIT Hyderabad and the ContinualAI community for the engaging and fruitful discussions, which helped to shape MERLIN. Last but not the least, we thank all our anonymous reviewers for their insightful comments and suggestions, which helped to improve the quality and presentation of the paper.

Broader Impact

Continual learning is a key desiderata for Artificial General Intelligence (AGI). Hence, this line of research has the benefits as well as the pitfalls of any other research effort geared in this direction. In particular, our work can help deliver impact on making smarter AI products and services, which can learn and update themselves on-the-fly when newer tasks and domains are encountered, without forgetting previously acquired knowledge. This is a necessity in any large-scale deployments of machine learning and computer vision, including in social media, e-commerce, surveillance, e-governance, etc - each of which have newer settings, tasks or domains added continually over time. Any negative effect of our work, such as legal and ethical concerns, are not unique to this work - to the best of our knowledge, but are shared with any other new development in machine learning, in general.

References

- [1] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhansu Maji, Charles Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. *arXiv preprint arXiv:1902.03545*, 2019.
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018.
- [3] 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*, pages 11849–11860, 2019.
- [4] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, pages 11816–11825, 2019.
- [5] Pablo Alvarez and Larry R Squire. Memory consolidation and the medial temporal lobe: a simple network model. *Proceedings of the national academy of sciences*, 91(15):7041–7045, 1994.
- [6] Shawn Beaulieu, Lapo Frati, Thomas Miconi, Joel Lehman, Kenneth O Stanley, Jeff Clune, and Nick Cheney. Learning to continually learn. *arXiv preprint arXiv:2002.09571*, 2020.
- [7] Sören Becker, Marcel Ackermann, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech Samek. Interpreting and explaining deep neural networks for classification of audio signals. *CoRR*, abs/1807.03418, 2018.
- [8] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Smash: one-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344*, 2017.
- [9] Lucas Caccia, Eugene Belilovsky, Massimo Caccia, and Joelle Pineau. Online learned continual compression with adaptative quantization module. *arXiv preprint arXiv:1911.08019*, 2019.
- [10] Alfonso Caramazza and Bradford Z Mahon. The organization of conceptual knowledge: the evidence from category-specific semantic deficits. *Trends in cognitive sciences*, 7(8):354–361, 2003.
- [11] Alfonso Caramazza and Jennifer R Shelton. Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of cognitive neuroscience*, 10(1):1–34, 1998.
- [12] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 233–248, 2018.
- [13] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 532–547, 2018.
- [14] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. In *ICLR*, 2019.
- [15] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. Continual learning with tiny episodic memories. *arXiv preprint arXiv:1902.10486*, 2019.
- [16] Yu Chen, Tom Diethe, and Neil Lawrence. Facilitating bayesian continual learning by natural gradients and stein gradients. *arXiv preprint arXiv:1904.10644*, 2019.
- [17] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. Continual learning: A comparative study on how to defy forgetting in classification tasks. *arXiv preprint arXiv:1909.08383*, 2019.
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern*

- recognition*, pages 248–255. Ieee, 2009.
- [19] Tom Diethe, Tom Borchert, Eno Thereska, Borja de Balle Pigem, and Neil Lawrence. Continual learning in practice. In *NeurIPS Continual Learning Workshop*, 2018.
 - [20] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
 - [21] Sayna Ebrahimi, Mohamed Elhoseiny, Trevor Darrell, and Marcus Rohrbach. Uncertainty-guided continual learning with bayesian neural networks. In *International Conference on Learning Representations*, 2020.
 - [22] Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning, 2019.
 - [23] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
 - [24] Chelsea Finn, Aravind Rajeswaran, Sham Kakade, and Sergey Levine. Online meta-learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1920–1930, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
 - [25] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.
 - [26] Oscar C Gonzalez, Yury Sokolov, Giri Krishnan, and Maxim Bazhenov. Can sleep protect memories from catastrophic forgetting? *BioRxiv*, page 569038, 2019.
 - [27] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
 - [28] Giacomo Handjaras, Emiliano Ricciardi, Andrea Leo, Alessandro Lenci, Luca Cecchetti, Mirco Cosottini, Giovanna Marotta, and Pietro Pietrini. How concepts are encoded in the human brain: a modality independent, category-based cortical organization of semantic knowledge. *Neuroimage*, 135:232–242, 2016.
 - [29] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. *arXiv preprint arXiv:1910.02509*, 2019.
 - [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
 - [31] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
 - [32] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018.
 - [33] Almut Hupbach, Rebecca Gomez, Oliver Hardt, and Lynn Nadel. Reconsolidation of episodic memories: A subtle reminder triggers integration of new information. *Learning & memory*, 14(1-2):47–53, 2007.
 - [34] Khurram Javed and Martha White. Meta-learning representations for continual learning. In H. Wallach, H. Larochelle, A. Beygelzimer, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 1818–1828. Curran Associates, Inc., 2019.
 - [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
 - [36] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pages 3581–3589, 2014.
 - [37] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
 - [38] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
 - [39] Carolin Konrad, Nora D Dirks, Annegret Warmuth, Jane S Herbert, Silvia Schneider, and Sabine Seehagen. Sleep-dependent selective imitation in infants. *Journal of sleep research*, 28(1):e12777, 2019.
 - [40] Mark A Kramer. Autoassociative neural networks. *Computers & chemical engineering*, 16(4):313–328, 1992.
 - [41] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *CiteSeer*, 2009.

- [42] Richard Kurl, Botond Cseke, Alexej Klushyn, Patrick van der Smagt, and Stephan Günnemann. Continual learning with bayesian neural networks for non-stationary data. In *International Conference on Learning Representations*, 2020.
- [43] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [44] 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*, 2020.
- [45] Timothée Lesort, Hugo Caselles-Dupré, Michael Garcia-Ortiz, Jean-François Goudou, and David Filliat. Generative models from the perspective of continual learning. In *IJCNN - International Joint Conference on Neural Networks*, Budapest, Hungary, Jul 2019.
- [46] Timothée Lesort, Alexander Gepperth, Andrei Stoian, and David Filliat. Marginal replay vs conditional replay for continual learning. In *International Conference on Artificial Neural Networks*, pages 466–480. Springer, 2019.
- [47] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018.
- [48] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pages 6467–6476, 2017.
- [49] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [50] Bradford Z Mahon, Stefano Anzellotti, Jens Schwarzbach, Massimiliano Zampini, and Alfonso Caramazza. Category-specific organization in the human brain does not require visual experience. *Neuron*, 63(3):397–405, 2009.
- [51] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018.
- [52] Jaya Krishna Mandivarapu, Blake Camp, and Rolando Estrada. Self-net: Lifelong learning via continual self-modeling. *Frontiers in Artificial Intelligence*, 3:19, 2020.
- [53] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [54] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [55] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.
- [56] Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. *arXiv preprint arXiv:1710.10628*, 2017.
- [57] Alex Nichol and John Schulman. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2, 2018.
- [58] Arghya Pal and Vineeth N Balasubramanian. Zero-shot task transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2189–2198, 2019.
- [59] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019.
- [60] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- [61] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [62] Jathushan Rajasegaran, Munawar Hayat, Salman H. Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for incremental learning. *CoRR*, abs/1906.01120, 2019.
- [63] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5533–5542. IEEE, 2017.
- [64] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018.
- [65] Anthony Robins. Consolidation in neural networks and in the sleeping brain. *Connection Science*, 8(2):259–276, 1996.

- [66] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, pages 348–358, 2019.
- [67] Clemens Rosenbaum, Tim Klinger, and Matthew Riemer. Routing networks: Adaptive selection of non-linear functions for multi-task learning. *arXiv preprint arXiv:1711.01239*, 2017.
- [68] Amir Rosenfeld and John K. Tsotsos. Incremental learning through deep adaptation, 2018.
- [69] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- [70] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, pages 1842–1850, 2016.
- [71] Jonathan Schwarz, Jelena Luketina, Wojciech M Czarnecki, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. *arXiv preprint arXiv:1805.06370*, 2018.
- [72] Joan Serra, Dídac Surís, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. *arXiv preprint arXiv:1801.01423*, 2018.
- [73] Falong Shen, Shuicheng Yan, and Gang Zeng. Neural style transfer via meta networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8061–8069, 2018.
- [74] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, pages 2990–2999, 2017.
- [75] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- [76] Larry R Squire, Lisa Genzel, John T Wixted, and Richard G Morris. Memory consolidation. *Cold Spring Harbor perspectives in biology*, 7(8):a021766, 2015.
- [77] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018.
- [78] Jakke Tamminen, Jessica D Payne, Robert Stickgold, Erin J Wamsley, and M Gareth Gaskell. Sleep spindle activity is associated with the integration of new memories and existing knowledge. *Journal of Neuroscience*, 30(43):14356–14360, 2010.
- [79] Michalis K Titsias, Jonathan Schwarz, Alexander G de G Matthews, Razvan Pascanu, and Yee Whye Teh. Functional regularisation for continual learning with gaussian processes. In *International Conference on Learning Representations*, 2019.
- [80] Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019.
- [81] Michiel van der Ven and Andreas S. Tolias. Generative replay with feedback connections as a general strategy for continual learning. *ArXiv*, abs/1809.10635, 2018.
- [82] Vinay Kumar Verma, Dhanajit Brahma, and Piyush Rai. A meta-learning framework for generalized zero-shot learning. *arXiv preprint arXiv:1909.04344*, 2019.
- [83] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- [84] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F. Grewe. Continual learning with hypernetworks. In *International Conference on Learning Representations*, 2020.
- [85] Matthew A Wilson and Bruce L McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, 1994.
- [86] John T Wixted. The psychology and neuroscience of forgetting. *Annu. Rev. Psychol.*, 55:235–269, 2004.
- [87] Jaehong Yoon, Saehoon Kim, Eunho Yang, and Sung Ju Hwang. Oracle: Order robust adaptive continual learning. *arXiv preprint arXiv:1902.09432*, 2019.
- [88] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3987–3995. JMLR. org, 2017.
- [89] Jeffrey O Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. Side-tuning: Network adaptation via additive side networks. *arXiv preprint arXiv:1912.13503*, 2019.
- [90] Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. In *International Conference on Machine Learning*, pages 7693–7702, 2019.