EFFICIENT CONTINUAL LEARNING WITH MODULAR NETWORKS AND TASK-DRIVEN PRIORS

Anonymous authors
Paper under double-blind review

ABSTRACT

Existing literature in Continual Learning (CL) has focused on overcoming catastrophic forgetting, the inability of the learner to recall how to perform tasks observed in the past. There are however other desirable properties of a CL system, such as the ability to transfer knowledge from previous tasks and to scale memory and compute sub-linearly with the number of tasks. Since most current benchmarks focus only on forgetting using short streams of tasks, we first propose a new suite of benchmarks to probe CL algorithms across these new axes. Finally, we introduce a new modular architecture, whose modules represent atomic skills that can be composed to perform a certain task. Learning a task reduces to figuring out which past modules to re-use, and which new modules to instantiate to solve the current task. Our learning algorithm leverages a task-driven prior over the exponential search space of all possible ways to combine modules, enabling efficient learning on long streams of tasks. Our experiments show that this modular architecture and learning algorithm perform competitively on widely used CL benchmarks while yielding superior performance on the more challenging benchmarks we introduce in this work.

1 INTRODUCTION

Continual Learning (CL) is a learning framework whereby an agent learns through a sequence of tasks (Ring, 1994; Thrun, 1994; 1998), observing each task once and only once. Much of the focus of the CL literature has been on avoiding *catastrophic forgetting* (McClelland et al., 1995; McCloskey & Cohen, 1989; Goodfellow et al., 2013), the inability of the learner to recall how to perform a task learned in the past. In our view, remembering how to perform a previous task is particularly important because it promotes knowledge accrual and transfer. CL has then the potential to address one of the major limitations of modern machine learning: its reliance on large amounts of labeled data. An agent may learn well a new task even when provided with little labeled data if it can leverage the knowledge accrued while learning previous tasks.

Our first contribution is then to pinpoint general properties that a good CL learner should have, besides avoiding forgetting. In §3 we explain how the learner should be able to *transfer knowledge* from related tasks seen in the past. At the same time, the learner should be able to *scale sub-linearly* with the number of tasks, both in terms of memory and compute.

Our second contribution is to introduce a new benchmark suite, dubbed CTrL, to test the above properties, since current benchmarks only focus on forgetting. For the sake of simplicity and as a first step towards a more holistic evaluation of CL models, in this work we restrict our attention to supervised learning tasks. Our experiments show that while commonly used benchmarks do not discriminate well between different approaches, our newly introduced benchmark let us dissect performance across several new dimensions of transfer and scalability, helping machine learning developers better understand strengths and weaknesses of various approaches.

Our last contribution is a new model that is designed according to the above mentioned properties of CL methods. It is based on a modular neural network architecture (Eigen et al., 2014; Denoyer & Gallinari, 2015; Fernando et al., 2017; Li et al., 2019): Every task is solved by the composition of a handful of neural modules which can be either borrowed from past tasks or freshly trained on the new task. In principle, modularization takes care of all the fundamental properties we care about, as i) by design there is no forgetting as modules from past tasks are not updated when new tasks arrive,

ii) transfer is enabled via sharing modules across tasks, and iii) the overall model scales sublinearly with the number of tasks as long as similar tasks share modules. The key issue is how to efficiently select modules, as the search space grows exponentially in their number. In this work we overcome this problem by leveraging a data driven prior over the space of possible architectures, which allows only local perturbations around the architecture of the previous task whose features best solve the current task. We describe this in detail in §4.

Our experiments in §5 show that this model performs comparably well on standard benchmarks, and much better on our new and more challenging benchmark, exhibiting better transfer and ability to scale to streams with a hundred of tasks.

2 Related Work

CL methods can be categorized in three main families of approaches. **Regularization** based methods use a single shared predictor across all tasks with the only exception that there can be a task specific classification head depending on the setting. They rely on various regularization methods to prevent forgetting. Kirkpatrick et al. (2016); Schwarz et al. (2018) use an approximation of the Fisher Information matrix while Zenke et al. (2017) use the distance of each weight to its initialization as a measure of importance. These approaches work well in streams containing a limited number of tasks but will inevitably either forget or stop learning as streams grow in size and diversity (van de Ven & Tolias, 2019), due to their structural rigidity and fixed capacity.

Similarly, **rehearsal** based methods also share a single predictor across all tasks but attack forgetting by using rehearsal on samples from past tasks. For instance, Lopez-Paz & Ranzato (2017); Chaudhry et al. (2019b); Rolnick et al. (2019) store past samples in a replay buffer, while Shin et al. (2017) learn to generate new samples from the data distribution of previous tasks and Zhang et al. (2019) computes per-class prototypes. These methods share the same drawback of regularization methods: Their capacity is fixed and pre-determined which makes them ineffective at handling long streams of tasks.

Finally, approaches based on **modular architectures** directly tackle the issue of the limited capacity by enabling the architecture to evolve over time. Our approach belongs to this category. Rusu et al. (2016) introduced a new network on each task, with connection to all previous layers, resulting in a model that grows super-linearly with the number of tasks. This issue was later addressed by Schwarz et al. (2018) who proposed to distill the new network back to the original one after each task, henceforth yielding a fixed capacity predictor which is going to have severe limitations on long streams. Yoon et al. (2018); Hung et al. (2019) propose a heuristic algorithm to automatically add and prune weights. Li et al. (2019) propose to softly select between reusing, adapting and introducing a new module at every layer. This work is the most similar to ours, with the major difference that we restrict the search space, enabling better scaling to longer streams. Rajasegaran et al. (2019); Fernando et al. (2017) propose to learn a path in a graph of modules using random search and genetic algorithms which are rather inefficient as the number of modules grow.

There are of course other works that introduce new benchmarks for CL. Most recently, Wortsman et al. (2020) have proposed a stream with 2500 tasks all derived from MNIST permutations. Unfortunately, this may provide little insight in terms of how well models transfer knowledge across tasks. Other benchmarks like CORe50 (Lomonaco & Maltoni, 2017) and CUB-200 (Wah et al., 2011) are more realistic but once again do not enable precise assessment of how well models transfer and scale.

CL is also related to other learning paradigms, such as meta-learning (Finn et al., 2017; Nichol et al., 2018; Duan et al., 2016), but these only consider the problem of quickly adapting to a new task while in CL we are also interested in preventing forgetting and learning better over time. For instance, Alet et al. (2018) proposed a modular approach for robotic applications. However, only the performance on the last task was measured. There is also a body of literature on modular networks for multi-task and multi-domain learning (Ruder et al., 2019; Rebuffi et al., 2017; Zhao et al., 2020). The major differences are the static nature of the learning problem they consider and the lack of emphasis on scaling to a large number of tasks.

3 EVALUATING CONTINUAL LEARNING MODELS

Let us start with a general formalization of the CL framework. We assume that tasks arrive in sequence, and that each task is associated with an integer task descriptor $t = 1, 2, ...$ which corresponds to the order of the tasks. Each task is defined by a labeled dataset $\mathcal{D}^t$. We denote with $\mathcal{S}$ a sequence of such tasks. A predictor for a given task t is denoted by $f : \mathcal{X}^t \times \mathbb{Z} \to \mathcal{Y}^t$. The predictor has internally some trainable parameters whose values depend on the stream of tasks $\mathcal{S}$ seen in the past, therefore the prediction is: $f(x, t|\mathcal{S})$. Notice that in general, different streams lead to different predictors for the same task: $f(x, t|\mathcal{S}) \neq f(x, t|\mathcal{S}')$.

Desirable Properties of CL models And Metrics: Since we are focusing on supervised learning tasks, it is natural to evaluate models using accuracy. We denote the prediction accuracy of the predictor f as $\Delta(f(x, t|\mathcal{S}), y)$, where x is the input, t is the task descriptor of x, $\mathcal{S}$ is the stream of tasks seen by the learner and y is the ground truth label.

In this work, we consider four major properties of a CL algorithm. First, the algorithm has to yield predictors that are accurate by the end of the learning experience. This is measured by their **average accuracy** at the end of the learning experience:

$$\mathcal{A}(\mathcal{S}) = \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{(x,y)\sim\mathcal{D}^t}[\Delta(f(x, t|\mathcal{S} = 1, \ldots, T), y)]. \tag{1}$$

Second, the CL algorithm should yield predictors that do not forget, i.e. that are able to perform a task seen in the past without significant loss of accuracy. **Forgetting** is defined as:

$$\mathcal{F}(\mathcal{S}) = \frac{1}{T-1} \sum_{t=1}^{T-1} \mathbb{E}_{(x,y)\sim\mathcal{D}^t}[\Delta(f(x, t|\mathcal{S} = 1, \ldots, T), y) - \Delta(f(x, t|\mathcal{S}' = 1, \ldots, t), y)] \tag{2}$$

This measure of forgetting has been called backward transfer (Lopez-Paz & Ranzato, 2017), and it measures the average loss of accuracy on a task at the end of training compared to when the task was just learned. Negative values indicate the model has been forgetting. Positive value indicate that the model has been improving on past tasks by learning subsequent tasks.

Third, the CL algorithm should yield predictors that are capable of transferring knowledge from past tasks when solving a new task. **Transfer** can be measured by:

$$\mathcal{T}(\mathcal{S}) = \mathbb{E}_{(x,y)\sim\mathcal{D}^T}[\Delta(f(x, T|\mathcal{S} = 1, \ldots, T), y) - \Delta(f(x, T|\mathcal{S}' = T), y)] \tag{3}$$

which measures the difference of performance between a model that has learned through a whole sequence of tasks and a model that has learned the last task in isolation. We would expect this quantity to be positive if there exist previous tasks that are related to the last task. Negative values imply the model has suffered some form of interference or even lack of plasticity when the predictor has too little capacity left to learn the new task.

Finally, the CL algorithm has to yield predictors that **scale sublinearly** with the number of tasks both in terms of memory and compute. In order to quantify this, we simply report the total memory usage and compute by the end of the learning experience. Since in the CL setting we are more interested in the memory consumption of the learner as opposed to the one of the final predictor, we include in the memory consumption everything a learner has to keep around to be able to continually learn (e.g., regularization parameters of EWC or the replay buffer for experience replay).

Streams The metrics introduced above can be applied to any stream of tasks. While current benchmarks are constructed to assess forgetting, they fall short at enabling a comprehensive evaluation of *transfer* and *scalability* because they do not control for task relatedness and they are composed of too few tasks. Therefore, we propose a new suite of streams. If t is a task in the stream, we denote with t^- and t^+ a task whose data is sampled from the same distribution as t, but with a much smaller or larger labeled dataset, respectively. Finally, t' and t'' are tasks that are similar to task t, while we assume no relation between t_i and t_j, for all $i \neq j$.

We consider five streams to assess different kinds of transfer. First, in the "ideal" transfer case $\mathcal{S}^- = (t_1^+, t_2, t_3, t_4, t_5, t_1^-)$, the last task is the same as the first one but with much less data to learn

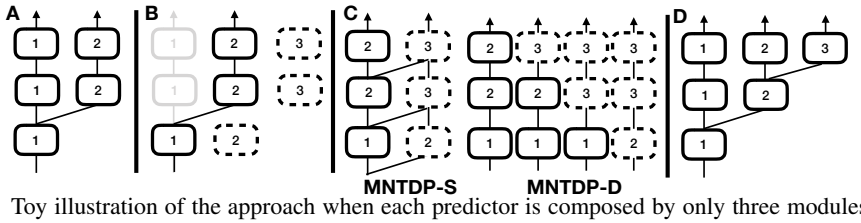

Figure 1: Toy illustration of the approach when each predictor is composed by only three modules and only two tasks have already been observed. **A)**: The predictor of the first task uses modules (1,1,1) (listing modules by increasing depth in the network) while the predictor of the second task uses modules (1,2,2); the first layer module is shared between the two predictors. **B)**: When a new task arrives, first we add one new randomly initialized module at each layer (the dashed modules). Second, we search for the most similar past task and retain only the corresponding architecture. In this case, the second task is most similar and therefore we remove (gray out) the modules used only by the predictor of the first task. **C)**: We train on the current task by learning both the best way to combine modules and their parameters. However, we restrict the search space. In this case, we only consider four possible compositions, all derived by perturbing the predictor of the second task. In the stochastic version (MNTDP-S), for every input a path (sequence of modules) is selected stochastically. Notice that the same module may contribute to multiple paths (e.g., the top-most layer with id 3). In the deterministic version instead (MNTDP-D), we train in parallel all paths and then select the best. Note that only the parameters of the newly added (dashed) modules are subject to learning. **D)**: Assuming that the best architecture found at the previous step is (1,2,3), module 3 at the top layer is added to the current library of modules.

from. This is useful to assess whether the learner can transfer knowledge from the relevant task. Second, we consider $\mathcal{S}^+ = (t_1^-, t_2, t_3, t_4, t_5, t_1^+)$ where the last task has much more data than the first task with intermediate tasks that are unrelated. In this case, there should not be much need to transfer anything from previous tasks. Next, we consider the case where the last task is similar to the first task but the input distribution changes $\mathcal{S}^{\text{in}} = (t_1, t_2, t_3, t_4, t_5, t_1')$, or the output distribution changes $\mathcal{S}^{\text{out}} = (t_1, t_2, t_3, t_4, t_5, t_1'')$. Finally, we consider a stream where all tasks are unrelated, $\mathcal{S}^{\text{pl}} = (t_1, t_2, t_3, t_4, t_5)$, which is useful to measure plasticity (ability to still learn) and potential interference (erroneous transfer from unrelated tasks) when learning the last task. All these tasks are evaluated using $\mathcal{T}$ in eq. 3.

In order to evaluate scalability, we simply consider $\mathcal{S}^{\text{long}}$, a stream with 100 tasks of various degree of relatedness and with varying amount of training data. See §5.1 for more details.

4 Modular Networks with Task Driven Prior (MNTDP)

In this section we describe an approach, dubbed MNTDP, designed according to the properties discussed in §3. The class of predictor functions $f(x, t|\mathcal{S})$ we consider in this work is *modular*, in the sense that predictors are composed by modules which are potentially shared across different (but related) tasks. A module can be any parametric function, for instance a ResNet block. The only restriction we use in this work is that all modules in a layer should differ only in terms of their actual parameter values, while modules across layers can implement different classes of functions. For instance, in the toy illustration of Fig. 1-A (see caption for details), there are two predictors, each composed by three modules (all ResNet blocks), the first one being shared.

4.1 Training

Once the new task t arrives, the algorithm follows three steps. First, it temporarily adds new randomly initialized modules at every layer (these are denoted by dashed boxes in Fig. 1-B) and it then defines a search space over all possible ways to combine modules. Second, it minimizes a loss function, which in our case is the cross-entropy loss, over both ways to combine modules and module parameters, see Fig. 1-C. Note that only the parameters of the newly added modules are subject to training, de facto preventing forgetting of previous tasks by construction. Finally, it takes the resulting predictor for the current task and it adds the parameters of the newly added modules (if any) back to the existing library of module parameters, see Fig. 1-D.

Since predictors are uniquely identified by which modules compose them, they can also be described by the *path* in the grid of module parameters. We denote the j-th path in the graph by π_j. The

parameters of the modules in path π_j are denoted by $\theta(\pi_j)$. Note that in general $\theta(\pi_j) \cap \theta(\pi_i) \neq \emptyset$, for $i \neq j$ since some modules may be shared across two different paths.

Let Π be the set of all possible paths in the graph. This has size equal to the product of the number of modules at every layer, after adding the new randomly initialized modules. If Γ is a distribution over Π which is subject to learning (and initialized uniformly), then the loss function is:

$$\Gamma^*, \theta^* = \arg\min_{\theta,\Gamma} \mathbb{E}_{j\sim\Gamma,(x,y)\sim\mathcal{D}^t} \mathcal{L}(f(x,t|\mathcal{S},\theta(\pi_j)),y) \tag{4}$$

where $f(x,t|\mathcal{S},\theta(\pi_j)$ is the predictor using parameters $\theta(\pi_j)$, $\mathcal{L}$ is the loss, and the minimization over the parameters is limited to only the newly introduced modules. The resulting distribution Γ^* is a delta distribution, assuming no ties between paths. Once the best path has been found and its parameters have been learned, the corresponding parameters of the new modules in the optimal path are added to the existing set of modules while the other ones are disregarded (Fig. 1-D). In this work we consider two instances of the learning problem in eq. 4, which differ in a) how they optimize over paths and b) how they share parameters across modules.

Stochastic version: This algorithm alternates between one step of gradient descent over the paths via REINFORCE (Williams, 1992) as in (Veniat & Denoyer, 2018), and one step of gradient descent over the parameters for a given path. The distribution Γ is modelled by a product of multinomial distributions, one for each layer of the model. These select one module at each layer, ultimately determining a particular path. Newly added modules may be shared across several paths which yields a model that can support several predictors while retaining a very parsimonious memory footprint thanks to parameter sharing. This version of the model, dubbed MNTDP-S, is outlined in the left part of Fig. 1-C and in Alg. 1 in the Appendix. In order to encourage the model to explore multiple paths, we use an entropy regularization on Γ during training.

Deterministic version: This algorithm minimizes the objective over paths in eq. 4 via exhaustive search, see Alg. 2 in Appendix and the right part of Fig. 1-C. Here, paths do not share any newly added module and we train one independent network per path, and then select the path yielding the lowest loss on the validation set. While this requires much more memory, it may also lead to better overall performance because each new module is cloned and trained just for a single path. Moreover, training predictors on each path can be trivially parallelized.

4.2 Data Driven Prior

Unfortunately, the algorithms as described above do not scale to a large number of tasks (and henceforth modules) because the search space grows exponentially. If there were N modules per layer and L layers, the search space would have size N^L. In order to restrict the search space we only allow paths that branch to the right: A newly added module at layer l can only connect to another newly added module at layer $l+1$, but it cannot connect to an already trained module at layer $l+1$. The underlying assumption is that for most tasks we expect changes in the output distribution as opposed to the input distribution, and therefore, if tasks are related, the base trunk is a good candidate for being shared. We will see in §5.3 what happens when this assumption is not satisfied, e.g., when applying this to $\mathcal{S}^{\text{in}}$.

To further restrict the search space we employ a *data driven prior*. The intuition is to limit the search space to perturbations of the path corresponding to the past task (or to the top-k paths) that is the most similar to the current task. There are several methods to assess which task is the closest to the current task without accessing data from the past tasks and also different ways to perturb a path. We propose a very simple approach, but others could have been used. We take the predictors from all the past tasks and select the path that yields the best nearest neighbor classification accuracy when feeding data from the current task using the features just before the classification head. This process is shown in Fig. 1-B. The search space is reduced from T^L to only L, and Γ of eq. 4 is allowed to have non-zero support only in this restricted search space.

By construction, the model does not forget, because we do not update modules of previous tasks. The model can transfer well because it can re-use modules from related tasks encountered in the past while not being constrained in terms of its capacity. And finally, the model scales sub-linearly in the number of tasks because modules can be shared across similar tasks. We will validate empirically in §5.3 whether the choice of the restricted search space works well in practice.

5 EXPERIMENTS

In this section we first introduce our benchmark in §5.1 and the modeling details §5.2, and then report results both on standard benchmarks as well as ours in §5.3.[1]

5.1 THE CTRL BENCHMARK

The CTrL (Continual Transfer Learning) benchmark is a collection of streams of tasks built over seven popular computer vision datasets, namely: CIFAR10 and CIFAR100 (Krizhevsky, 2009), DTD (Cimpoi et al., 2014), SVHN (Netzer et al., 2011), MNIST (LeCun et al., 1998), Rainbow-MNIST (Finn et al., 2019) and Fashion MNIST (Xiao et al., 2017); see Table 3 in Appendix for basic statistics. These datasets are desirable because they are diverse (hence tasks derived from some of these datasets can be considered unrelated), they have a fairly large number of training examples to simulate tasks that do not need to transfer, and they have low spatial resolution enabling fast evaluation. CTrL is designed according to the methodology described in §3, to enable evaluation of various transfer learning properties and ability of models to scale to a large number of tasks. Each task consists of a training, validation and test datasets corresponding to a 5-way and 10-way classification problem for the transfer streams and the long stream, respectively. The last task of $\mathcal{S}^{\text{out}}$ consists of a shuffling of the output labels of the first task. The last task of $\mathcal{S}^{\text{in}}$ is the same as its first task except that MNIST images have a different background color. $\mathcal{S}^{\text{long}}$ has 100 tasks, and it is constructed by first sampling a dataset, then 5 classes at random, and finally the amount of training data from a distribution that favors small tasks by the end of the learning experience. Therefore, $\mathcal{S}^{\text{long}}$ tests not only the ability of a model to scale to a relatively large number of tasks but also to transfer more efficiently with age.

All images have been rescaled to 32x32 pixels in RGB color format, and per-channel normalized using statistics computed on the training set of each task. During training we perform data augmentation by using random crops (4 pixels padding and 32x32 crops) and random horizontal reflection. Please, refer to Appendix A for further details.

5.2 METHODOLOGY AND MODELING DETAILS

Models learn over each task in sequence; data from each task can be replayed several times but each stream is observed only once. Since each task has a validation dataset, hyper-parameters (e.g., learning rate and number of weight updates) are task specific and they are cross-validated on the validation set of each task. Once the learning experience ends, we test the resulting predictor on the test sets of all the tasks. Notice that this is a stricter paradigm than what usually employed in the literature (Kirkpatrick et al., 2016), where hyper-parameters are set at the stream level (by replaying the stream several times). Our model selection criterion is more realistic because it does not assume that the learner has access to future tasks when cross-validating on the current task, and this is more consistent with the CL's assumptions of operating on a stream of data.

All models use the same backbone, a small variant of the ResNet-18 architecture which is divided into 7 modules; please, refer to Table 8 for details of how many layers and parameters each block contains. Each predictor is trained by minimizing the cross-entropy loss with a small L2 weight decay on the parameters. In our experiments, MNTDP adds 7 new randomly initialized modules, one for every block. The search space does not allow connecting old blocks from new blocks, and it considers two scenarios: considering old blocks from the past task that is deemed most similar ($k = 1$, default setting) or considering the whole set of old blocks ($k =$ all) resulting in a larger search space.

We compare to several baselines: **Independent Models** which instantiates a randomly initialized predictor for every task (as many paths as tasks without any module overlap), 2) **Finetuning** which trains a single model to solve all the tasks without any regularization and initializes from the model of the previous task (a single path shared across all tasks), 3) **New-Head** which also shares the trunk of the network across all tasks but not the classification head which is task specific, 4) **New-Leg** which shares all layers across tasks except for the very first input layer which is task specific,

[1]The scripts allowing to rebuild the streams as well as to train models will be provided upon publication.

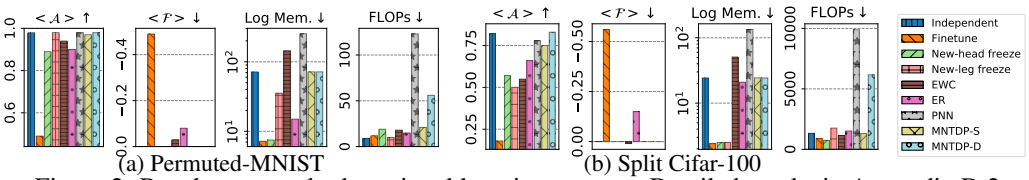

(a) Permuted-MNIST (b) Split Cifar-100

Figure 2: Results on standard continual learning streams. Detailed results in Appendix D.2.

5) **EWC** (Kirkpatrick et al., 2016) which is like "finetuning" but with a regularizer to alleviate forgetting, 6) **Experience Replay** (Chaudhry et al., 2019b) which is like finetuning except that the model has access to some samples from the past tasks to rehearse and alleviate forgetting (we use 15 samples per class to obtain a memory consumption similar to other baselines), 7) **Progressive Neural Networks** (PNNs) (Rusu et al., 2016) which adds both a new module at every layer as well as lateral connections once a new task arrives. For new-head and new-leg we have two variants. In the "freeze" variant the modules shared across tasks are not updated after the first task, while in the "finetune" version they are updated throughout the whole stream. See Appendix C for more details.

We report performance across different axes as discussed in §3: Average accuracy as in eq. 1, forgetting as in eq. 2, transfer as in eq. 3 and applicable only to the transfer streams, "Mem." [MB] which refers to the average memory consumed at training time, and "Flops" [T] which corresponds to the average amount of computation needed to reach the best performance.

5.3 Results

Results over existing benchmarks: In Fig. 2 we compare MNTDP against several baselines on two standard streams with 10 tasks, Permuted MNIST and Split CIFAR100. We observe that all models do well, with the only exception of finetuning which exhibits strong forgetting. PNNs has good average accuracy but requires more memory. MNTDP-D yields the best average accuracy but requires more computation; notice however that its wall-clock training time can be easily reduced to that of "independent models", since all candidate paths (seven in our case) can be trained in parallel on modern GPU devices. More importantly, in these standard streams MNTDP trivially reduces to "independent models" without any module sharing. This is optimal because each task is fairly distinct and has a relatively large amount of data. It is therefore not possible to assess how well models can transfer knowledge across tasks, nor it is possible to assess how well models scale. Fortunately, we can leverage the CTrL benchmark to better study these properties, as described next.

Results on CTrL: We first evaluate models in terms of their ability to transfer by evaluating them on the streams $\mathcal{S}^-, \mathcal{S}^+, \mathcal{S}^{\text{in}}, \mathcal{S}^{\text{out}}$ and $\mathcal{S}^{\text{pl}}$ introduced in §3. Tab. 1 shows that "independent models" is again a strong baseline, because even on the first four streams, all tasks except the last one are unrelated and therefore instantiating an independent model is optimal. However, MNTDP yields the best average accuracy overall. MNTDP-D achieves the best transfer on streams $\mathcal{S}^-, \mathcal{S}^+, \mathcal{S}^{\text{out}}$ and $\mathcal{S}^{\text{pl}}$, and indeed it discovers the correct path in each of these cases (e.g., it discovers to reuse the path of the first task when learning on $\mathcal{S}^-$ and to just swap the top modules when learning the last task on $\mathcal{S}^+$). Examples of dicovered path are presented in appendix D.4. MNTDP underperforms on $\mathcal{S}^{\text{in}}$ because its prior does not match the data distribution, since in this case, it is the input distribution that has changed but swapping the first module is out of MNTDP search space. This highlights the importance of the choice of prior for this algorithm. In general, MNTDP offers a clear trade-off between accuracy, i.e. how broad the prior is which determines how many paths can be evaluated, and memory/compute budget. Computationally MNTDP-D is the most demanding, but in practice its wall clock time is comparable to "independent" because GPU devices can store in memory all the paths (in our case, seven) and efficiently train them all in parallel. We observe also that MNTDP-S has a clear advantage in terms of compute at the expense of a lower overall average accuracy, as sharing modules across paths during training can lead to sub-optimal convergence.

Overall, MNTDP has much higher average accuracy that methods with a fixed capacity. It also beats PNNs, which seems to struggle with interference when learning on $\mathcal{S}^{\text{pl}}$, as all new modules connect to all old modules which are irrelevant for the last task. Moreover, PNNs uses much more memory.

	$< \mathcal{A} >$	$< \mathcal{F} >$	Mem.	FLOPs	$\mathcal{T}(\mathcal{S}^-)$	$\mathcal{T}(\mathcal{S}^+)$	$\mathcal{T}(\mathcal{S}^{\text{in}})$	$\mathcal{T}(\mathcal{S}^{\text{out}})$	$\mathcal{T}(\mathcal{S}^{\text{pl}})$
Independent	0.58	0.0	14.1	308	-0.4	0.4	-0.4	-0.3	0.0
Finetune	0.19	-0.3	2.4	284	-0.3	0.2	-0.4	-0.3	-0.1
New-head freeze	0.48	0.0	2.4	307	-0.0	0.1	-0.6	0.0	-0.4
New-leg freeze	0.41	0.0	2.5	366	-0.1	0.0	-0.0	-0.4	-0.4
Online EWC	0.43	-0.1	7.3	310	-0.0	0.0	-0.1	-0.0	-0.4
ER	0.44	-0.1	13.1	604	-0.4	0.1	-0.4	-0.2	-0.2
PNN	0.57	0.0	48.2	1459	-0.1	0.2	-0.3	-0.1	-0.1
MNTDP-S	0.59	0.0	11.7	363	0.0	0.2	-0.3	-0.0	-0.1
MNTDP-D	0.64	0.0	11.6	1512	0.0	0.4	-0.4	0.0	-0.0

Table 1: Aggregated results on the transfer streams over multiple relevant baselines (complete table with more baselines provided in Appendix D).

	$< \mathcal{A} >$	$< \mathcal{F} >$	Mem.	FLOPs
Independent	0.57	0.00	243.7	16
Finetune	0.20	-0.34	2.4	21
New-head	0.43	0.00	2.6	24
Onl. EWC	0.27	-0.25	7.4	17
MNTDP-S	0.68	0.00	158.6	24
MNTDP-D	0.75	0.00	102.0	114

Table 2: Results on the long evaluation stream composed of 100 tasks.

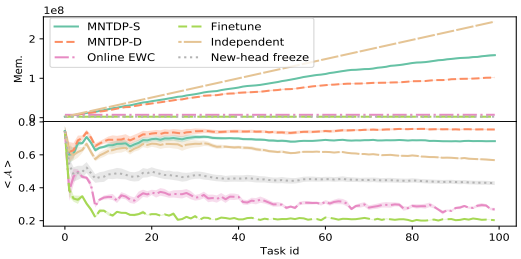

Figure 3: Evolution of $< \mathcal{A} >$ and Mem. over tasks.

"New-leg" and "new-head" models perform very well only when the tasks in the stream match their assumptions, showing the advantange of the adaptive prior of MNTDP. Finally, EWC shows great transfer on $\mathcal{S}^-, \mathcal{S}^{\text{in}}, \mathcal{S}^{\text{out}}$, which probe the ability of the model to retain information. However, it fails at $\mathcal{S}^+, \mathcal{S}^{\text{pl}}$ that require additional capacity allocation to learn a new task.

We conclude by reporting results on $\mathcal{S}^{\text{long}}$ composed of 100 tasks. Tab.2 report the results of all the approaches we could learn without running into out-of-memory issues. We notice that MNTDP-D yields an absolute 18% improvement over the baseline *independent model* while using less than half of its memory, thanks to the discovery many of paths with shared modules. Fig. 3 shows the average accuracy up to the current task over time. MNTDP-D attains the best performance while growing sublinearly in terms of memory usage. Online EWC's performance deteriorates with age instead.

Ablation: We first study the importance of the prior. Instead of selecting the nearest neighbor path, we pick one path corresponding to one of the previous task at random. In this case, $\mathcal{T}(\mathcal{S}^-)$ decreases from 0 to -0.2 and $\mathcal{T}(\mathcal{S}^{out})$ goes from 0 to -0.3. With a random path, MNTDP learns not to share any module, demonstrating that it is indeed important to form a good prior over the search space. Finally, we explore the robustness to the number of modules by splitting each module in two, yielding a total of 10 modules per path, and by merging adjacent modules yielding a total of 3 modules for the same overall number of parameters. We find that $\mathcal{T}(\mathcal{S}^{\text{out}})$ decreases from 0 to -0.1, with a 9% decrease on t_1^-, when the number of modules decreases but stays the same when the number of modules increases, suggesting that the algorithm has to have sufficient number of modules to flexibly grow.

6 Conclusions

We introduced a new benchmark to enable a more comprehensive evaluation of CL learning algorithms, not only in terms of average accuracy and forgetting but also knowledge transfer and scaling. We have also proposed a modular network that can grow its capacity and share components across tasks as learning proceeds. Our experiments show that our approach yields a very desirable trade-off between accuracy and compute/memory usage, suggesting that modularization is a promising avenue of investigation for continual learning and for knowledge transfer.

REFERENCES

Ferran Alet, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Modular meta-learning. In 2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings, volume 87 of Proceedings of Machine Learning Research, pp. 856–868. PMLR, 2018. URL http://proceedings.mlr.press/v87/alet18a.html.

Arslan Chaudhry, Marc'Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with A-GEM. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019a. URL https://openreview.net/forum?id=Hkf2_sC5FX.

Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet Kumar Dokania, Philip H. S. Torr, and Marc'Aurelio Ranzato. Continual learning with tiny episodic memories. CoRR, abs/1902.10486, 2019b. URL http://arxiv.org/abs/1902.10486.

M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2014.

L. Denoyer and P. Gallinari. Deep sequential neural networks. EWRL, 2015.

Yan Duan, John Schulman, Xi Chen, Peter L. Bartlett, Ilya Sutskever, and Pieter Abbeel. Rl^2: Fast reinforcement learning via slow reinforcement learning. CoRR, abs/1611.02779, 2016. URL http://arxiv.org/abs/1611.02779.

D. Eigen, I. Sutskever, and M. Ranzato. Learning factored representations in a deep mixture of experts. ICLR, 2014.

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. CoRR, abs/1701.08734, 2017. URL http://arxiv.org/abs/1701.08734.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In Doina Precup and Yee Whye Teh (eds.), Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, volume 70 of Proceedings of Machine Learning Research, pp. 1126–1135. PMLR, 2017. URL http://proceedings.mlr.press/v70/finn17a.html.

Chelsea Finn, Aravind Rajeswaran, Sham Kakade, and Sergey Levine. Online meta-learning. In International Conference on Machine Learning, 2019.

I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio. An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks. arXiv, 2013.

Steven C. Y. Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada, pp. 13647–13657, 2019. URL http://papers.nips.cc/paper/9518-compacting-picking-and-growing-for-unforgetting-continual-learning.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. URL http://arxiv.org/abs/1412.6980.

James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. CoRR, abs/1612.00796, 2016. URL http://arxiv.org/abs/1612.00796.

Alex Krizhevsky. Learning multiple layers of features from tiny images. University of Toronto, technical report, 2009.

Y. LeCun, L. Bottou, and and P. Haffner Y. Bengio. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.

Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, volume 97 of Proceedings of Machine Learning Research, pp. 3925–3934. PMLR, 2019. URL http://proceedings.mlr.press/v97/li19m.html.

Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. arXiv:1705.03550, 2017.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA, pp. 6467–6476, 2017. URL http://papers.nips.cc/paper/7225-gradient-episodic-memory-for-continual-learning.

James L McClelland, Bruce L McNaughton, and Randall C O'reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychological review, 1995.

Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. Psychology of learning and motivation, 1989.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In NIPS Workshop on Deep Learning and Unsupervised Feature Learning, 2011.

Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. CoRR, abs/1803.02999, 2018. URL http://arxiv.org/abs/1803.02999.

Jathushan Rajasegaran, Munawar Hayat, Salman H. Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for incremental learning. CoRR, abs/1906.01120, 2019. URL http://arxiv.org/abs/1906.01120.

Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA, pp. 506–516, 2017. URL http://papers.nips.cc/paper/6654-learning-multiple-visual-domains-with-residual-adapters.

Mark B. Ring. Continual Learning in Reinforcement Environments. PhD thesis, University of Texas at Austin, Austin, Texas 78712, 1994.

David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy P. Lillicrap, and Gregory Wayne. Experience replay for continual learning. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada, pp. 348–358, 2019. URL http://papers.nips.cc/paper/8327-experience-replay-for-continual-learning.

Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. Latent multi-task architecture learning. In The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019,

The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019, pp. 4822–4829. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33014822. URL https://doi.org/10.1609/aaai.v33i01.33014822.

Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. CoRR, abs/1606.04671, 2016. URL http://arxiv.org/abs/1606.04671.

Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. In Jennifer G. Dy and Andreas Krause (eds.), Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018, volume 80 of Proceedings of Machine Learning Research, pp. 4535–4544. PMLR, 2018. URL http://proceedings.mlr.press/v80/schwarz18a.html.

Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA, pp. 2990–2999, 2017. URL http://papers.nips.cc/paper/6892-continual-learning-with-deep-generative-replay.

S. Thrun. A lifelong learning perspective for mobile robot control. Proceedings of the IEEE/RSJ/GI Conference on Intelligent Robots and Systems, 1994.

Sebastian Thrun. Lifelong learning algorithms. In Learning to learn. Springer, 1998.

Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning. CoRR, abs/1904.07734, 2019. URL http://arxiv.org/abs/1904.07734.

Tom Veniat and Ludovic Denoyer. Learning time/memory-efficient deep architectures with budgeted super networks. In Intl. Conference on Compute Vision and Pattern Recognition, 2018.

C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Caltech-ucsd birds-200-2011 dataset. Tech Report: CNS-TR-2011-001, 2011.

Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine Learning, 8:229–256, 1992.

Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. CoRR, abs/2006.14769, 2020. URL https://arxiv.org/abs/2006.14769.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.

Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net, 2018. URL https://openreview.net/forum?id=Sk7KsfW0-.

Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In Doina Precup and Yee Whye Teh (eds.), Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, volume 70 of Proceedings of Machine Learning Research, pp. 3987–3995. PMLR, 2017. URL http://proceedings.mlr.press/v70/zenke17a.html.

Mengmi Zhang, Tao Wang, Joo Hwee Lim, and Jiashi Feng. Prototype reminding for continual learning. CoRR, abs/1905.09447, 2019. URL http://arxiv.org/abs/1905.09447.

Hanbin Zhao, Hao Zeng, Xin Qin, Yongjian Fu, Hui Wang, Bourahla Omar, and Xi Li. What and where: Learn to plug adapters via NAS for multi-domain learning. CoRR, abs/2007.12415, 2020. URL https://arxiv.org/abs/2007.12415.