

# AN EVOLUTIONARY APPROACH TO DYNAMIC INTRODUCTION OF TASKS IN LARGE-SCALE MULTITASK LEARNING SYSTEMS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Multitask learning assumes that models capable of learning from multiple tasks can achieve better quality and efficiency via knowledge transfer, a key feature of human learning. Though, state of the art ML models rely on high customization for each task and leverage size and data scale rather than scaling the number of tasks. Also, continual learning, that adds the temporal aspect to multitask, is often focused to the study of common pitfalls such as catastrophic forgetting instead of being studied at a large scale as a critical component to build the next generation artificial intelligence. We propose an evolutionary method capable of generating large scale multitask models that support the dynamic addition of new tasks. The generated multitask models are sparsely activated and integrates a task-based routing that guarantees bounded compute cost and fewer added parameters per task as the model expands. The proposed method relies on a knowledge compartmentalization technique to achieve immunity against catastrophic forgetting and other common pitfalls such as gradient interference and negative transfer. We demonstrate empirically that the proposed method can jointly solve and achieve competitive results on 69 public image classification tasks, for example improving the state of the art on a competitive benchmark such as cifar10 by achieving a 15% relative error reduction compared to the best model trained on public data.

## 1 INTRODUCTION

The success of machine learning continues to grow as it finds new applications in areas as diverse as language generation (Brown et al., 2020), visual art generation (Ramesh et al., 2021), chip design (Mirhoseini et al., 2020), protein folding (Senior et al., 2020) and competitive sports (Silver et al., 2016; Vinyals et al., 2019). The vast majority of machine learning models are designed and trained for a single task and specific data modality, and are often trained by starting with randomly initialized parameters, or with limited knowledge transfer from a pre-trained model. While this paradigm has shown great success, it uses a large amount of computational resources, and does not leverage knowledge transfer from many related tasks in order to achieve higher performance and efficiency.

The work presented in this paper is based on the intuition that significant advances can be enabled by dynamic, continual learning approaches capable of achieving knowledge transfer across a very large number of tasks. The method described in this paper can dynamically incorporate new tasks into a large running system, can leverage pieces of a sparse multitask ML model to achieve improved quality for new tasks, and can automatically share pieces of the model among related tasks. This method can enhance quality on each task, and also improve efficiency in terms of convergence time, amount of training examples, energy consumption and human engineering effort.

The ML problem framing proposed by this paper can be interpreted as a generalization and synthesis of the standard multitask and continual learning formalization, since an arbitrarily large set of tasks can be solved jointly. But also, over time, the set of tasks can be extended with a continuous stream of new tasks. Furthermore, it lifts the distinction between a pretraining task and a downstream task. As new tasks are incorporated, the system searches for how to combine the knowledge and representations already present in the system with new model capacity in order to achieve high quality

for each new task. Knowledge acquired and representations learned while solving a new task are available for use by any future task or continued learning for existing tasks.

We refer to the proposed method as “mutant multitask network” or  $\mu$ 2Net. This method generates a large scale multitask network that jointly solves multiple tasks to achieve increased quality and efficiency for each. It can continuously expand the model by allowing the dynamic addition of new tasks. The more accumulated knowledge that is embedded into the system via learning on previous tasks, the higher quality the solutions are for subsequent tasks. Furthermore, new tasks can be solved with increasing efficiency in terms of reducing the newly-added parameters per task. The generated multitask model is sparsely activated as it integrates a task-based routing mechanism that guarantees bounded compute cost per task as the model expands. The knowledge learned from each task is compartmentalized in components that can be reused by multiple tasks. As demonstrated through experiments, this compartmentalization technique avoids the common problems of multitask and continual learning models, such as catastrophic forgetting, gradient interference and negative transfer. The exploration of the space of task routes and identification of the subset of prior knowledge most relevant for each task is guided by an evolutionary algorithm designed to dynamically adjust the exploration/exploitation balance without need of manual tuning of meta-parameters. The same evolutionary logic is employed to dynamically tune the hyperparameters multitask model components.

## 2 RELATED WORK

The main novelty of the presented work is to propose and demonstrate a method that jointly provides all of the following properties: 1) ability to continually learn from an unbounded stream of tasks, 2) automate the selection and reuse of prior knowledge and representations learned for previous tasks in the solving of new tasks, 3) search the space of possible model architectures allowing the system to dynamically extend its capacity and structure without requiring random initialization, 4) automatically tune the hyperparameters of both the generated models and the evolutionary method, including the ability to learn schedules for each hyperparameter, rather than just constant values, 5) ability to optimize for any reward function, also including non-differentiable factors, 6) immunity from catastrophic forgetting, negative transfer and gradient interference, 7) ability to extend any pre-existing pre-trained model, including extending its architecture and adapting the domain on which such model have been trained to other domains automatically, 8) introduction of a flexible access control list mechanism that allows expression of a variety of privacy policies, including allowing the use or influence of task-specific data to be constrained to just a single task or to a subset of tasks for which data or higher-level representation use should be permitted. Different lines of research have focused on distinct subsets of the many topics addressed by the proposed method. In this section we highlight a few cornerstone publications. Refer to Appendix A for an extended survey.

Different methods have been proposed to achieve **dynamic architecture extensions** (Chen et al., 2016; Cai et al., 2018), some also focusing on an unbounded stream of tasks (Yoon et al., 2018), or achieving immunity from catastrophic forgetting (Rusu et al., 2016; Li & Hoiem, 2018; Rosenfeld & Tsotsos, 2020). Unlike our work, these techniques rely on static heuristics and patterns to define the the structural extensions, rather than a more open-ended learned search process. **Neural architecture search** (NAS) (Zoph & Le, 2017) methods aim to modularize the architectural components in search spaces whose exploration can be automated with reinforcement learning or evolutionary approaches (Real et al., 2019; Maziarz et al., 2018). More efficient (but structurally constrained) parameter sharing NAS techniques (Pham et al., 2018; Liu et al., 2019a) create a connection with routing methods (Fernando et al., 2017) and **sparse activation** techniques, that enable the decoupling of model size growth from compute cost growth (Shazeer et al., 2017; Du et al., 2021). Evolutionary methods have also been applied with success for **hyperparameter tuning** (Jaderberg et al., 2017).

Cross-task **knowledge transfer** has gained popularity, especially through transfer learning from a model pre-trained on a large amount of data for one or a few general tasks, and then fine-tuned on a small amount of data for a related downstream task. This approach has been shown to be very effective in a wide variety of problems and modalities (Devlin et al., 2019; Dosovitskiy et al., 2021). Large scale models have recently achieved novel transfer capabilities such as few/zero shot learning (Brown et al., 2020). More complex forms of knowledge transfer such as **multitask training** or **continual learning** often lead to interesting problems such as catastrophic forgetting (McCloskey & Cohen, 1989; French, 1999), negative transfer (Rosenstein, 2005; Wang et al., 2019) or gradient interference

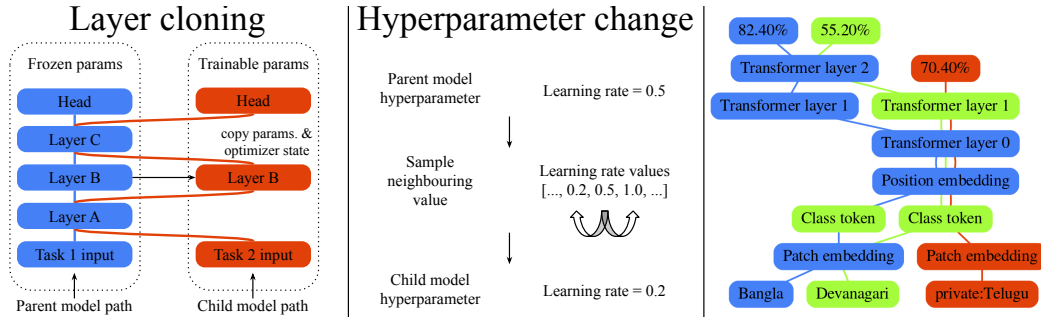


Figure 1: Graphical representation of the two mutation types used by the proposed method: layer cloning mutation (left) and hyperparameter change (center). The graph on the right represents the model generated by the preliminary experiment described in Section 4. The bottom nodes display the task names, the top nodes display the validation accuracy, and internal nodes are represented with the color of the task that has last updated the parameters of the corresponding layer.

(Chen et al., 2018; Yu et al., 2020). Research on these topics mostly focuses on approaches such as weighted combination methods (Liu et al., 2019b; Sun et al., 2020b) or gradient transformations (Sener & Koltun, 2018; Kendall et al., 2018), also methods automating knowledge selection at a layer level was proposed (Sun et al., 2020a).

### 3 EVOLUTIONARY METHOD

This section defines the proposed method capable of generating a dynamic multitask ML system. The multitask system is initialized with one *root model*. This model can be either pretrained or randomly initialized. During the evolutionary process, the proposed method searches for the best model for a single task at a time, referred to as the *active task*. During the active phase of a task, a population of models trained on the active task is evolved: the *active population*. The first time a task becomes active, its active population is empty. For subsequent iterations, the active population is initialized with all the models trained on the active task that have been retained from previous iterations. The active population is iteratively extended by: 1) sampling a *parent model* (Section 3.1), 2) applying to the parent model a sampled set of mutations (Section 3.2) to produce a *child model*, 3) performing cycles of training and validation in order to train and score the child model. Each trained model is assigned a *score* (Section 3.3). Early population pruning is performed by discarding the models that did not achieve a better score than their parent. An active phase is composed of multiple *generations* in which multiple batches of child models are sampled and trained in parallel. At the end of a task active phase, only its best scoring model is retained as part of the multitask system. A task can become active multiple times. Details of the method are reported below (and in Algorithm 1).

#### 3.1 PARENT MODEL SAMPLING

The first attempt to sample a parent model for the active task is done over the active population of models for that task. The models in the active population are visited in decreasing order of score, starting with the highest scoring one. Each model,  $m$ , can be accepted as parent with probability:  $p_{parent}(m|t) = 0.5^{\#selections(m,t)}$ . Where  $\#selections(m,t)$  denotes the number of times the candidate model,  $m$ , has been previously selected as parent to generate a child models for task  $t$ . If the current candidate parent is not selected, then iteratively the model with the next best score is considered to be selected as parent with probability  $p_{parent}(\cdot|t)$ . If a full iteration on the active population is completed without a successful parent model selection, then the same method is applied to the randomly sorted list of all remaining models: all the models currently part of the multitask system that were trained on a task different from the current active task,  $t$ . This fallback list is randomly sorted since these models have not been scored for  $t$ . As a final fallback a parent is uniformly sampled among all the models currently in the system. This method prioritizes the exploitation of high scoring models that had few attempts at generating an improved model for the active task. But also, in combination with early pruning, it automatically transitions toward a more exploratory behavior in case the higher scoring models are unable to generate an improvement.

### 3.2 MUTATIONS

In this work we consider Deep Neural Networks (DNN) models. DNN are commonly defined by their architecture and hyperparameters. Architectures are composed of a sequence of neural network layers, each mapping an input vector into an output vector of variable dimensions. Hyperparameters specify the configuration details such as the optimizer or data preprocessing configurations. The presented method allows for two types of mutations (Figure 1):

**Hyperparameter mutations** can be applied to modify the configuration inherited from the parent. Each hyperparameter is associated with a sorted list of valid values (Table 1). If a hyperparameter is selected for mutation, then its new value is selected at random among the two neighbouring values in the sorted list (Figure 1). This constraints hyperparameter mutations to only incremental changes. Notice that, every hyperparameter of a child model is set to a single value. However, considering that a child model continues its ancestors training with mutated hyperparameters, then the method can be regarded as capable of defining a piece-wise constant schedule over time for each hyperparameter.

**Layer cloning mutations** create a copy of any parent model layer that can be trained by the child model. If a layer of the parent model is not selected for cloning, then it is shared with the child model in a frozen state to guarantee immutability of pre-existing models. Child models can train only the cloned copies of the parent layers. The cloned layers are trained with a possibly modified version of the parent optimizer. The configuration of the child optimizer is defined by the mutated hyperparameters. If such optimizer is of a type that stores a state (i.e. momentum), then the state is also cloned from the state saved by the ancestor that has last trained the cloned layer. Notice that, a trainable layer may be followed by frozen layers. In this case the gradients for the trainable layer are propagated through the frozen layers and applied only to the parameters of the trainable layers while frozen parameters are left unchanged. The head layer is always cloned since it always needs to be trainable. If a child model is trained on a task different from the parent’s task, then a new randomly initialized head layer is created with output shape matching the number of classes of the new task.

Each possible layer cloning or hyperparameter mutation is independently sampled to be applied with probability  $\mu$ .  $\mu$  is itself a hyperparameter that is mutated by the evolutionary process. Thus demonstrating that automatic tuning is not only applied to selecting the hyperparameters of the generated models, but can also be applied to self-tune the configuration of the evolutionary algorithm.

### 3.3 TRAINING AND SCORING

A newly sampled child model is trained on the active task for a given number of epochs. The model is evaluated on the validation set after each epoch. At each intermediate evaluation, the child model is assigned a score that the evolutionary algorithm aims to maximize. The score can be defined to optimize a mixture of factors such as validation quality, inference latency, training compute or model size, depending on the applications requirements. The presented experiments aim to compare against the state of the art for a large number of tasks without any size or compute constraint. Therefore, the validation accuracy is used directly as the score without additional factors. After training, only the parameters and optimizer state of the version of the child model achieving best score are retained.

### 3.4 DISCUSSION AND PROPERTIES

Notice that, none of the defined mutation actions or the evolutionary algorithm allow the creation of child models that can alter the parent model in any way. Once a model has been trained, the parameters storing its knowledge cannot be modified. This method guarantees immunity from **catastrophic forgetting**, since the knowledge of a trained model is always preserved. It also provides a solution to **negative transfer**, since it automates the selection of the knowledge that is most relevant for each new task. Furthermore, it also avoids **gradient interference**, that can arise when multiple gradients are synchronously applied to the same set of parameters. Nonetheless, models for new tasks can use knowledge and representations from prior tasks and even extend these to improve or specialize them.

The method compartmentalizes the knowledge of each task in a subset of components, allowing the implementation of different **dataset privacy control** policies. For example, we can introduce private tasks that can benefit from all the public knowledge embedded in the multitask system but are able to withhold the knowledge and representations derived from their private dataset from being used by other tasks. This is achieved by preventing other tasks from using or cloning components trained on

Table 1: Hyperparameters valid values. Bold vales are defaults. This search space consists of a parametrization of the configuration for the published ViT model definition library.

---

$\mu \in [0.10, 0.12, 0.14, 0.16, 0.18, \mathbf{0.20}, 0.22, 0.24, 0.26, 0.28, 0.30]$
Learning rate $\in [0.0001, 0.0002, 0.0005, 0.001, 0.002, 0.005, \mathbf{0.01}, 0.02, 0.05, 0.1, 0.2, 0.5]$
Cosine learning rate schedule warm up ratio $\in [0.01, 0.02, 0.05, \mathbf{0.1}, 0.2, 0.3, 0.4]$
Momentum $\in [0.5, 0.6, 0.7, 0.8, 0.85, \mathbf{0.9}, 0.95, 0.98, 0.99]$
Nesterov update $\in [\mathbf{False}, \text{True}]$
Crop input image $\in [\text{False}, \mathbf{True}]$
Cropped area range min $\in [0.01, 0.02, \mathbf{0.05}, 0.1, 0.2, 0.5, 1.0]$
Cropped aspect ratio range min $\in [0.25, 0.5, \mathbf{0.75}, 1.0]$
Flip left/right $\in [\text{False}, \mathbf{True}]$
Brightness delta $\in [\mathbf{0.0}, 0.01, 0.02, 0.05, 0.1, 0.2]$
Contrast delta $\in [\mathbf{0.0}, 0.01, 0.02, 0.05, 0.1, 0.2]$
Saturation delta $\in [\mathbf{0.0}, 0.01, 0.02, 0.05, 0.1, 0.2]$
Hue delta $\in [\mathbf{0.0}, 0.01, 0.02, 0.05, 0.1, 0.2]$

---

private data. This also allows to remove the knowledge learned from the private dataset at any future date by simply removing its components. This private/public distinction can be generalized into an access- control-list mechanism. For example, a set of private tasks can share representations that are withheld from any other task. Privacy control capabilities are empirically demonstrated in Section 4.

### 3.5 EXPERIMENTAL SET UP

This section details the instantiation of the proposed method employed in the experimental analysis. The task type for the presented set of experiment is image classification. This choice allows us to define a large benchmark of publicly available datasets with standardized framing. It also allows us to build on top of state-of-the-art models whose architecture definition and checkpoints are public: the Visual Transformer (ViT) is used as root model (Dosovitskiy et al., 2021).

**Architecture** Layer cloning mutations can create a copy of any of ViT’s layers: 1) *Patch embedding*: the first layer of the model maps the input image into a sequence of embedded tokens, each corresponding to a patch of the input image. 2) *Class token*: a classification token is prepended to the sequence. The final hidden state corresponding to this token is used as the aggregate sequence representation for classification tasks (Devlin et al., 2019). 3) *Position embedding*: the sequence representation is then augmented with an embedding that carries each patch positional information. 4) *Transformer layers*: the sequence representation generated by the input layers is iteratively transformed by a stack of transformer layers (Vaswani et al., 2017). 5) *Model head*: a final fully connected layer mapping the representation produced by the top-most transformer layer into the logits.

**Parameters** The parameters of the root model can be either randomly initialized or loaded from a checkpoint. The preliminary experiment demonstrates the evolution from random initialization (see Section 4), while the large scale experiment starts from a pretrained large ViT model (see Section 5).

**Hyperparameters** As default hyperparameters we use those resulting from the extensive study conducted by Steiner et al. (2021): SGD momentum optimizer, cosine decay schedule, no weight decay, gradient clipping at global norm 1 and  $386 \times 386$  image resolution. The evolutionary method can change the hyperparameters of optimizer, image preprocessing and architecture (see Table 1).

## 4 PRELIMINARY EXPERIMENT

This section describes a small scale preliminary experiment that introduces details of the method application and illustrates the privacy control and initialization capabilities. This experiment demonstrates the ability to generate a multitask model from random initialization and a minimal architecture rather than evolving a pretrained state-of-the-art model. Therefore, a randomly initialized ViT Ti/16 architecture (Steiner et al., 2021) stripped of transformer layers is used as root model. To allow the method to build a capable architecture, we add an extra mutation action that can insert a new randomly initialized transformer layer just before the head layer.

Furthermore, the dataset privacy control technique (see Section 3) is demonstrated by adding a private task. Three numerals recognition tasks are used as benchmark: {[bangla](#), [devanagari](#), [telugu](#)}. Telugu is introduced as a private task, so that no other task can access the knowledge introduced into the system by its dataset. However, its models can leverage the knowledge provided by the public tasks.

This short experiment is configured to perform 2 active task iterations for each task. During each active task iteration 4 model generations are produced. In each generation 8 child models are sampled and trained in parallel on each of the 8 TPUv3 cores. The choice of small datasets, architecture and training budget is intended to facilitate reproducibility and fast experimental iterations. The experiment can be reproduced by using the published artifacts and completes in less than 5 minutes.

Figure 1 (right) displays the resulting multitask model solving jointly the 3 tasks. We observe a high degree of cross-task knowledge and components sharing throughout the [evolutionary process](#). Even though the root model has no transformer layers, multiple randomly initialized transformer layers are inserted and trained improving the score of each task. Note that, at any point during the evolution, the components trained on the private task (red) are only used by the private task.

## 5 LARGE SCALE CONTINUAL LEARNING EXPERIMENT

This section reports a single [large scale continual learning experiment](#) producing a multitask system jointly solving 69 visual tasks. A pretrained ViT L/16 is used as root model, which has been selected for its pretraining validation accuracy on the imagenet-21k dataset following Steiner et al. (2021).

### 5.1 ViT BENCHMARK

The first tasks introduced to the system are 3 tasks on which ViT was evaluated in Dosovitskiy et al. (2021). This experiment is configured to perform 5 active iterations for each task, and 4 model generations for each iteration. During each generation, 8 child models are trained in parallel on each of the 8 TPUv3 cores. Each model training performs 4 validation cycles. The number of train samples between validation cycles is set to  $\min(1 \text{ epoch}, 128116 \text{ samples})$  to smooth the distribution of compute across datasets of different size. 128116 is equivalent to  $1/10^{\text{th}}$  of an epoch of the imagenet2012 training set. The same configuration is applied in following experiments.

This configuration results in 8 epochs for imagenet, and 80 epochs for cifar. This is roughly equivalent to the fine-tuning setup of the baseline model we compare against (Dosovitskiy et al., 2021): 8 epochs for imagenet and 102.4 for cifar. The proposed method can be considered cheaper since: 1) ViT fine-tuning has been repeated multiple times for the hyperparameters tuning process and 2) setting  $\mu = 0.2$  results in cheaper training steps, since parameters updates can be skipped for frozen layers and gradient propagation can be skipped for frozen layers at the base of the model (preliminary experiments have shown a 2.5-4% training speed-up attributable to this). In order to provide a fair comparison, as a root model is used the same ViT L/16 architecture, same checkpoint pretrained on the i12k dataset, same  $384 \times 384$  image resolution, and optimizer and preprocessing configuration.

Table 2 reports the top 1 test accuracy achieved.  $\mu$ 2Net outperforms fine-tuning with comparable training steps per task. Extending the training with 5 additional tasks iterations leads to moderate gains on imagenet2012 and cifar10. Notice that, for cifar100 the accuracy decreases. This can happen since the best models are selected according the validation accuracy and, as the model gets close to convergence, a small validation accuracy gain may lead to a noisy perturbation of the test accuracy.

To quantify knowledge transfer, we consider the model produced for each task and examine the dataset on which each layer’s ancestors were trained. On average, the layers composing the model generated

Table 2: Test accuracy achieved by  $\mu$ 2Net and by fine-tuning a comparable pretrained ViT model.

Model	imagenet2012	cifar100	cifar10
ViT L/16 fine-tuning (Dosovitskiy et al., 2021)	85.30	93.25	99.15
$\mu$ 2Net after 5 task iterations	86.38	94.75	99.35
$\mu$ 2Net after 10 task iterations	86.66	94.67	99.38
$\mu$ 2Net cont. after adding VTAB-full tasks	<b>86.74</b>	94.67	99.41
$\mu$ 2Net cont. after adding VDD tasks	<b>86.74</b>	94.74	99.43
$\mu$ 2Net cont. after adding all 69 tasks	<b>86.74</b>	<b>94.95</b>	<b>99.49</b>

Table 3: Test accuracy achieved on the VTAB-full benchmark by: 1) fine-tuning with matching architecture and checkpoint (ViT L/16 i21k) reported by Steiner et al. (2021), 2) the Sup-rotation method (Gidaris et al., 2018) that achieved the best result in the VTAB-full leaderboard (Zhai et al., 2019b), 3-4)  $\mu$ 2Net results after 2 task iterations, 5) and after an additional iteration performed after the VDD benchmark introduction. Underlined models transfer knowledge from VDD tasks.

Model	caltech101	cifar100	dtd	flowers102	pets	sun397	svhn	cameleon	eurosat	resisc45	retinopathy	clevr-count	clevr-dist	dmlab	dspr-loc	dspr-orient	kitti-dist	snorb-azim	snorb-elev
Steiner et al. (2021)	<b>95.5</b>	94.1	80.3	<b>99.6</b>	95.0	83.4	97.4	86.4	99.0	96.6	83.3	99.8	91.7	75.6	100	90.4	<b>84.7</b>	27.5	76.5
Zhai et al. (2019b)	94.6	84.8	75.9	94.7	91.5	70.2	97.0	85.9	98.8	94.9	79.5	99.8	92.5	76.5	100	<b>96.5</b>	82.3	<b>100</b>	<b>98.4</b>
$\mu$ 2Net cont. 1st iter.	92.6	94.6	79.8	<b>99.6</b>	<b>95.3</b>	84.5	97.3	87.5	99.1	96.3	83.7	99.8	93.2	<b>76.9</b>	100	96.2	83.0	32.3	94.5
$\mu$ 2Net cont. 2nd iter.	92.6	94.6	80.5	<b>99.6</b>	<b>95.3</b>	<b>84.8</b>	<b>97.8</b>	88.4	<b>99.2</b>	<b>97.0</b>	<b>84.0</b>	99.8	<b>94.0</b>	<b>76.9</b>	100	96.4	83.0	33.3	95.1
$\mu$ 2Net cont. after VDD	<u>93.0</u>	<u>94.7</u>	<b>81.0</b>	<b>99.6</b>	<b>95.3</b>	<b>84.8</b>	<b>97.8</b>	<b>91.1</b>	99.1	<b>97.0</b>	<b>84.0</b>	99.8	<b>94.0</b>	<b>76.9</b>	100	96.4	<u>82.3</u>	33.3	95.1

for the imagenet2012 task have performed only 60.6% of the training steps on the imagenet2012 dataset, and have received 31.5% of the gradient updates from cifar100 and 7.9% from cifar10. The layers comprising the cifar100 model have performed 42.3% of their training on imagenet2012 and 20.6% on cifar10. And layers comprising the cifar10 model performed 46.1% of training on imagenet2012 and 35.9% on cifar100. The tasks heterogeneity improves the representations produced by the different layers, and results in generally higher performance, as shown in Table 2.

The following sections 5.2 and 5.3 describe the extensions of the system performed by introducing additional benchmarks. After the introduction of each benchmark, we perform an additional iteration on imagenet and cifar tasks to analyze effects of further knowledge enrichment. As a representative example: the VDD benchmark (Section 5.3) includes a low resolution version of cifar100. The model that will be generated for vdd/cifar100 will be a mutation of the current full resolution cifar100 model. Afterward, the additional active task iteration on cifar100 will be performed, and the resulting improved cifar100 model will be a mutation of the low resolution vdd/cifar100 model.

After a final iteration, we note that 99.49 is the best cifar10 accuracy reported for a model trained only on public data: to the best of our knowledge, this constitutes a 15% relative error reduction compared to the 99.40 state of the art achieved by Touvron et al. (2021). Dosovitskiy et al. (2021) achieves 99.50 with a double size ViT-Huge model trained on proprietary data. The achieved cifar100 accuracy is currently outperformed only by Ridnik et al. (2021) (95.10) and Foret et al. (2021) (96.08).

## 5.2 VTAB-FULL BENCHMARK

Next, we introduce to the system the 19 VTAB-full tasks (Zhai et al., 2019a), plus 5 additional task variants that are not included in the standard evaluation set (Table 8). From this experiment onward, the infrastructure is scaled from 8 to 32 cores, as detailed in Appendix B. The number of task iterations is reduced from 10 to 2. These changes lead to a roughly similar exploratory and training budget per task. However, the increased parallelism results in faster task iterations.

Table 3 reports the achieved results along with reference models that use limited knowledge transfer capabilities. Steiner et al. (2021) reports the quality achieved by fine-tuning a model equivalent to our root model. This outperforms  $\mu$ 2Net on only 2 tasks, even if it has been trained multiple times to perform hyperparameter tuning. Zhai et al. (2019b) reports the results of the best model identified with a large scale study. This state of the art model outperforms  $\mu$ 2Net on 4 tasks. Again, increasing number of task iterations and additional knowledge (VDD) in the system, seem to yield better quality.

Table 4: Test accuracy mean and std.dev. achieved on the VDD benchmark by 3 system replicas.

Model:	imagenet	svhn	cifar100	gtsrb	daimler	omniglot	ucf101	aircraft	dtd	flowers
$\mu$ 2Net cont.										
1st iter.	89.1±0.19	98.2±0.23	97.2±0.34	99.9±0.04	99.9±0.03	84.5±1.18	83.6±1.79	64.8±3.57	75.2±1.20	99.0±0.11
2nd iter.	89.2±0.12	98.4±0.23	97.2±0.40	99.9±0.05	99.9±0.04	84.3±0.27	85.7±1.36	65.4±3.03	76.0±0.49	99.2±0.11



Table 5: Test accuracy achieved on the Multitask Character Classification Benchmark by  $\mu$ 2Net continued extension with 2 active task iterations and by the model that has set the state of the art using comparable data splits: Jeevan & Sethi (2022) for digits, Kabir et al. (2020) for letters, Ajayan & James (2021) for kmnist, An et al. (2020) for mnist, Hazra et al. (2021) for cmaterdb/bangla. Underlined models reuse knowledge introduced by other character classification tasks. Datasets are listed in decreasing size from the biggest emnist/digits (240k samples) to the smallest telugu (2.5k).

Model	emnist/ digits	emnist/ letters	kmnist	mnist	omniglot	cmaterdb/ bangla	cmaterdb/ devanagari	cmaterdb/ telugu
State of the art	99.77	<b>95.88</b>	95.40	<b>99.91</b>	—	99.00	—	—
$\mu$ 2Net cont. 1st iteration	<b>99.82</b>	93.60	<b>98.68</b>	99.75	98.72	98.60	<u>96.60</u>	<u>97.80</u>
$\mu$ 2Net cont. 2nd iteration	<b>99.82</b>	93.68	98.60	99.69	<b>99.84</b>	<b>99.10</b>	<b>98.00</b>	<b>99.40</b>

### 5.3 VISUAL DOMAIN DECATHLON (VDD) BENCHMARK

The VDD benchmark (Bilen et al., 2017) is introduced next. The ML methodology proposed in this paper, can achieve higher efficiency by focusing the available compute on the training of a single multitask system. Though, the standard approach to measure variance relies on experiment repetitions. This section demonstrates how variance can be measured for any chosen segment of the training. In practice, 2 task iterations are performed to introduce the VDD tasks starting from the state achieved after the introduction of the last benchmark as usual. But, this experiment is run on 3 parallel copies of the system, allowing us to compute variance of the metrics for this set of task iterations.

The VDD benchmark is composed of 10 diverse tasks. This is also reflected in the diverse variance ranges measured (see Table 4). Variance is low for most of the tasks. However, for ucf101 and aircraft is significantly higher. The metrics that have highest correlation with standard deviation are error rate (linear proportionality in log scale) and number of training samples per class (inverse proportionality in log scale) (Figure 4). These can be considered metrics indicative of the complexity of the task. Furthermore, variance decreases with the second iteration: average standard deviation of 0.87 after 1 iteration and 0.61 after the second. These findings can support the intuitive hypothesis that tasks with higher complexity may benefit from more iterations to decrease variance and approach convergence. The next system extension continues from the state of one randomly selected replica.

### 5.4 MULTITASK CHARACTER CLASSIFICATION BENCHMARK

We continue extending the system by adding a set of 8 character classification tasks. Thus offering the opportunity to study knowledge transfer across tasks with high domain correlation.

Table 5 reports the test accuracy achieved with 2 active tasks iterations. We observe that tasks with more training data (left) achieve convergence in the first iteration, this hypothesis is supported by lack of significant accuracy gains with the second iteration. While tasks with less training data (right) show a significant gain from a second training iteration. Smaller tasks use transferred in domain knowledge: bangla top model reuses components that embed knowledge introduced by emnist/letters,

Table 6: Test accuracy achieved on the VTAB-1k benchmark by: 1) fine-tuning ViT L/16 i21k (matching root model) Dosovitskiy et al. (2021), 2) the Sup-rotation method (Gidaris et al., 2018) that achieved the best result in the VTAB-1k leaderboard (Zhai et al., 2019b). Underlined models have at least one ancestor trained on the corresponding full form task. Doubly underlined model inherit directly from the current best model for the matching full form task.

Model	caltech101	cifar100	dtd	flowers102	pets	sun397	svhn	camelyon	eurosat	resisc45	retinopathy	clevr-count	clevr-dist	dmlab	dspr-loc	dspr-orient	kitti-dist	snorb-azim	snorb-elev
Dosovitskiy et al. (2021)	90.8	84.1	74.1	99.3	92.7	<b>61.0</b>	80.9	82.5	95.6	85.2	75.3	70.3	56.1	41.9	74.7	64.9	79.9	30.5	41.7
Zhai et al. (2019b)	<b>91.7</b>	53.7	69.5	90.8	88.1	32.8	88.5	83.4	96.0	82.0	71.1	47.3	57.2	36.6	88.3	52.1	77.1	51.6	33.7
$\mu$ 2Net cont. 1st iter.	87.1	<u>89.4</u>	<u>77.6</u>	<u>99.2</u>	<u>94.5</u>	57.6	<u>97.5</u>	<b>86.0</b>	<u>98.6</u>	<u>93.4</u>	78.0	<u>91.2</u>	59.9	47.6	<u>58.4</u>	<b>96.2</b>	<b>81.9</b>	<u>32.1</u>	<b>92.5</b>
$\mu$ 2Net cont. 2nd iter.	89.9	<b>90.6</b>	<b>78.1</b>	<b>99.7</b>	<b>94.5</b>	57.6	<b>97.5</b>	<b>86.0</b>	<u>98.3</u>	<b>93.4</b>	<b>83.5</b>	<b>99.8</b>	<b>90.6</b>	<b>76.3</b>	<b>100</b>	<b>96.2</b>	81.7	<b>33.7</b>	<b>92.5</b>






Figure 2: Activated and added parameters per task as percentage with respect to the total number of parameters of the multitask system along the duration the large scale experiment (see Section 5. Vertical lines highlight the start of the introduction for each of the considered benchmark.

while devanagari transfers from omniglot and bangla, and telugu transfers from bangla and devanagari. Furthermore, the achieved quality is comparable to the state of the art published for each task.

### 5.5 VTAB-1K BENCHMARK

The system is further extended by adding the 1k-samples version of the VTAB tasks. Since the system contains already the knowledge learned from the full version of each task, this set allows to study how effective is the proposed method at retrieving knowledge that is already embedded in the system.

Table 6 reports results along with reference models that use limited knowledge transfer capabilities. During the first iteration, the models generated for the short form tasks can retrieve the knowledge of the corresponding full form task also without directly mutating its model, but rather mutating a model having at least one ancestor trained on the full task. For example, the model generated for `flowers1021k` mutates the `ddd1k` model, that has 28 ancestors, of which only the 21<sup>st</sup> was trained on `oxford_flowers102full`. However, that is enough for `flowers1021k` to achieve 99.2% test accuracy.

After only one task iteration, 5 tasks achieve better test accuracy than the reference models without reusing any knowledge introduced by the corresponding full form task. Particularly interesting is the case of `kitty-dist1k` (a.k.a `kitty/closest_vehicle_distance1k`), that achieves a strong performance without transferring from the matching full form task but composing the knowledge of related tasks: `kitty/closest_object_distancefull` (8 ancestors), `kitty/count_vehiclesfull` (3 ancestors) and `kitty/count_vehicles1k` (3 ancestors). Thus learning to estimate distance of the closest vehicle by combining the knowledge of recognizing vehicles and estimating the distance of the closest object. Also, `clevr-dist1k` achieves a strong performance by inheriting from the semantically equivalent task `kitti/closest_object_distancefull` without reusing the knowledge introduced by `clevr-distfull`.

## 6 CONCLUSION

We introduced the  $\mu$ 2Net method, aimed at achieving state-of-the-art quality on a large task set, with the ability to dynamically introduce new tasks into the running system. The more tasks are learned the more knowledge is embedded in the system. A ViT-L architecture (307M parameters) was evolved into a multitask system with 13'087M parameters jointly solving 69 tasks. However, as the system grows, the sparsity in parameter activation keeps the amount of compute and the memory usage per task constant. The average added parameters per task decreases by 38% through the experiment, and the resulting multitask system activates only 2.3% of the total parameters per task (see Figure 2 and Table 7). The proposed method allows decoupling the growth of useful representations for solving new tasks from the growth of parameters/compute per task. Furthermore, experimenting with a large number of tasks allowed us to identify different patterns of positive knowledge transfer and composition, achieving higher efficacy on small datasets and across related tasks. The proposed approach to mutations allows to achieve immunity against common pitfalls of multitask systems such as catastrophic forgetting, negative transfer and gradient interference, and demonstrates the key data privacy properties we want to achieve in a continual learning system. Future work can continue to build toward systems that can acquire further capabilities and knowledge across multiple modalities.

All the experiments reported in this paper can be **reproduced** by using the public datasets and published  $\mu$ 2Net code, and can be extended by using the published checkpoints (see Appendix B).

## REFERENCES

- Arun Ajayan and A P James. Edge to quantum: hybrid quantum-spiking neural network image classifier. *Neuromorphic Computing and Engineering*, 1, 2021.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning. *ArXiv*, abs/2204.14198, 2022.
- Sanghyeon An, Min Jun Lee, Sanglee Park, He Sarina Yang, and Jungmin So. An ensemble of simple convolutional neural network models for mnist digit recognition. *ArXiv*, abs/2008.10400, 2020.
- Paul Barham, Aakanksha Chowdhery, Jeffrey Dean, Sanjay Ghemawat, Steven Hand, Daniel Hurt, Michael Isard, Hyeontaek Lim, Ruoming Pang, Sudip Roy, Brennan Saeta, Parker Schuh, Ryan Sepassi, Laurent El Shafey, Chandramohan A. Thekkath, and Yonghui Wu. Pathways: Asynchronous distributed dataflow for ml. *ArXiv*, abs/2203.12533, 2022.
- James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *NIPS*, 2011.
- Hakan Bilen, Sylvestre Rebuffi, and Tomas Jakab. Visual domain decathlon. 2017.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020.
- Han Cai, Tianyao Chen, Weinan Zhang, Yong Yu, and Jun Wang. Efficient architecture search by network transformation. In *AAAI*, 2018.
- Tianqi Chen, Ian J. Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. *CoRR*, abs/1511.05641, 2016.
- Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *ICML*, 2018.
- Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105:1865–1883, 2017.
- Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3606–3613, 2014.
- Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *ArXiv*, abs/1812.01718, 2018.
- Gregory Cohen, Saeed Afshar, Jonathan C. Tapson, and André van Schaik. Emnist: Extending mnist to handwritten letters. *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 2921–2926, 2017.
- N. Das, Jagan Mohan Reddy, Ram Sarkar, Subhadip Basu, Mahantapas Kundu, Mita Nasipuri, and Dipak Kumar Basu. A statistical-topological feature combination for recognition of handwritten numerals. *Appl. Soft Comput.*, 12:2486–2495, 2012a.
- N. Das, Ram Sarkar, Subhadip Basu, Mahantapas Kundu, Mita Nasipuri, and Dipak Kumar Basu. A genetic algorithm based region sampling for selection of local features in handwritten digit recognition application. *Appl. Soft Comput.*, 12:1592–1606, 2012b.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929, 2021.
- Nan Du, Yanping Huang, Andrew M. Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, Barret Zoph, Liam Fedus, Maarten Bosma, Zongwei Zhou, Tao Wang, Yu Emma Wang, Kellie Webster, Marie Pellat, Kevin Robinson, Kathleen S. Meier-Hellstern, Toju Duke, Lucas Dixon, Kun Zhang, Quoc V. Le, Yonghui Wu, Z. Chen, and Claire Cui. Glam: Efficient scaling of language models with mixture-of-experts. *ArXiv*, abs/2112.06905, 2021.
- Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. 2004.
- Chrisantha Fernando, Dylan S. Banarse, Charles Blundell, Yori Zwols, David R Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *ArXiv*, abs/1701.08734, 2017.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *ArXiv*, abs/2010.01412, 2021.
- Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3:128–135, 1999.
- Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3354–3361, 2012.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *ArXiv*, abs/1803.07728, 2018.
- Abhishek Hazra, Prakash Choudhary, Sanasam Chanu Inunganbi, and Mainak Adhikari. Bangla-meitei mayek scripts handwritten character recognition using convolutional neural network. *Applied Intelligence*, 51:2291–2311, 2021.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- Patrick Helber, Benjamin Bischke, Andreas R. Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12:2217–2226, 2019.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *ICML*, 2019.
- Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M. Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, Chrisantha Fernando, and Koray Kavukcuoglu. Population based training of neural networks. *ArXiv*, abs/1711.09846, 2017.
- Pranav Jeevan and Amit Sethi. Wavemix: Resource-efficient token mixing for images. *ArXiv*, abs/2203.03689, 2022.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1988–1997, 2017.

- Norman P. Jouppi, Cliff Young, Nishant Patil, David A. Patterson, Gaurav Agrawal, Raminder Singh Bajwa, Sarah Bates, Suresh Bhatia, Nanette J. Boden, Al Borchers, Rick Boyle, Pierre luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert B. Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Daniel Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle A. Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. In-datacenter performance analysis of a tensor processing unit. *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*, pp. 1–12, 2017.
- H. M. D. Kabir, Moloud Abdar, Seyed Mohammad Jafar Jalali, Abbas Khosravi, Amir F. Atiya, Saeid Nahavandi, and Dipti Srinivasan. Spinalnet: Deep neural network with gradual input. *ArXiv*, abs/2007.03347, 2020.
- Kaggle and EyePacs. Kaggle diabetic retinopathy detection. <https://www.kaggle.com/c/diabetic-retinopathy-detection/data>, 2015.
- Jared Kaplan, Sam McCandlish, T. J. Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeff Wu, and Dario Amodei. Scaling laws for neural language models. *ArXiv*, abs/2001.08361, 2020.
- Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7482–7491, 2018.
- David A. Klindt, Lukas Schott, Yash Sharma, Ivan Ustyuzhaninov, Wieland Brendel, Matthias Bethge, and Dylan M. Paiton. Towards nonlinear disentanglement in natural data with temporal sparse coding. *ArXiv*, abs/2007.10930, 2021.
- Efi Kokkioopoulou, Anja Hauth, Luciano Sbaiz, Andrea Gesmundo, Gábor Bartók, and Jesse Berent. Fast task-aware architecture inference. *ArXiv*, abs/1902.05781, 2019.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350:1332 – 1338, 2015.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86:2278–2324, 1998.
- Yann LeCun, Fu Jie Huang, and Léon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, 2:II–104 Vol.2, 2004.
- Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *ICML*, 2019.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:2935–2947, 2018.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *ArXiv*, abs/1806.09055, 2019a.
- Shengchao Liu, Yingyu Liang, and Anthony Gitter. Loss-balanced task weighting to reduce negative transfer in multi-task learning. In *AAAI*, 2019b.
- Krzysztof Maziarz, Andrey Khorlin, Quentin de Laroussilhe, and Andrea Gesmundo. Evolutionary-neural hybrid agents for architecture search. *ArXiv*, abs/1811.09828, 2018.

- Michael McCloskey and Neal J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, 24:109–165, 1989.
- Thomas Mensink, Jasper R. R. Uijlings, Alina Kuznetsova, Michael Gygli, and Vittorio Ferrari. Factors of influence for transfer learning across diverse appearance domains and task types. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2021.
- Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim M. Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Sungmin Bae, Azade Nazi, Jiwoo Pak, Andy Tong, Kavya Srinivasa, Will Hang, Emre Tuncer, Ananda Babu, Quoc V. Le, James Laudon, Richard Ho, Roger Carpenter, and Jeff Dean. Chip placement with deep reinforcement learning. *ArXiv*, abs/2004.10746, 2020.
- Yuval Netzer, Tao Wang, Adam Coates, A. Bissacco, Bo Wu, and A. Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pp. 722–729, 2008.
- Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3498–3505, 2012.
- Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *ICML*, 2018.
- Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *ArXiv*, abs/1910.10683, 2020.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *ArXiv*, abs/2102.12092, 2021.
- Rahul Ramesh and Pratik Chaudhari. Model zoo: A growing brain that learns continually. In *ICLR*, 2022.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. Regularized evolution for image classifier architecture search. In *AAAI*, 2019.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. In *NIPS*, 2017.
- Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley D. Edwards, Nicolas Manfred Otto Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A generalist agent. *ArXiv*, abs/2205.06175, 2022.
- T. Ridnik, Gilad Sharir, Avi Ben-Cohen, Emanuel Ben-Baruch, and Asaf Noy. MI-decoder: Scalable and versatile classification head. 2021.
- Amir Rosenfeld and John K. Tsotsos. Incremental learning through deep adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:651–663, 2020.
- Michael T. Rosenstein. To transfer or not to transfer. In *NIPS 2005*, 2005.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115: 211–252, 2015.
- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *ArXiv*, abs/1606.04671, 2016.

- Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *NeurIPS*, 2018.
- Andrew W. Senior, Richard Evans, John M. Jumper, James Kirkpatrick, L. Sifre, Tim Green, Chongli Qin, Augustin Zidek, Alexander W. R. Nelson, Alex Bridgland, Hugo Penedones, Stig Petersen, Karen Simonyan, Steve Crossan, Pushmeet Kohli, David T. Jones, David Silver, Koray Kavukcuoglu, and Demis Hassabis. Improved protein structure prediction using potentials from deep learning. *Nature*, 577:706–710, 2020.
- Noam M. Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc V. Le, Geoffrey E. Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *ArXiv*, abs/1701.06538, 2017.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, L. Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–489, 2016.
- Jasper Snoek, H. Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *NIPS*, 2012.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, 2010.
- Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *ArXiv*, abs/2106.10270, 2021.
- Ximeng Sun, Rameswar Panda, and Rogério Schmidt Feris. Adashare: Learning what to share for efficient deep multi-task learning. *ArXiv*, abs/1911.12423, 2020a.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. Ernie 2.0: A continual pre-training framework for language understanding. *ArXiv*, abs/1907.12412, 2020b.
- Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2815–2823, 2019.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam M. Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, Yaguang Li, Hongrae Lee, Huaixiu Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Yanqi Zhou, Chung-Ching Chang, I. A. Krivokon, Willard James Rusch, Marc Pickett, Kathleen S. Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Hartz Søraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravindran Rajakumar, Alena Butryna, Matthew Lamm, V. O. Kuzmina, Joseph Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. Lamda: Language models for dialog applications. *ArXiv*, abs/2201.08239, 2022.
- Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 32–42, 2021.
- Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *ArXiv*, abs/1706.03762, 2017.
- Bastiaan S. Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant cnns for digital pathology. *ArXiv*, abs/1806.03962, 2018.

- Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, L. Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander Sasha Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom Le Paine, Caglar Gulcehre, Ziyun Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy P. Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, pp. 1–5, 2019.
- Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime G. Carbonell. Characterizing and avoiding negative transfer. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11285–11294, 2019.
- Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 3485–3492, 2010.
- Huaxiu Yao, Yingbo Zhou, Mehrdad Mahdavi, Zhenhui Jessie Li, Richard Socher, and Caiming Xiong. Online structured meta-learning. *ArXiv*, abs/2010.11545, 2020.
- Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *ArXiv*, abs/1708.01547, 2018.
- Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *ArXiv*, abs/2001.06782, 2020.
- Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruysen, Carlos Riquelme, Mario Lucic, Josip Djolonga, André Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, Lucas Beyer, Olivier Bachem, Michael Tschannen, Marcin Michalski, Olivier Bousquet, Sylvain Gelly, and Neil Houlsby. The visual task adaptation benchmark. *ArXiv*, abs/1910.04867, 2019a.
- Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruysen, Carlos Riquelme, Mario Lucic, Josip Djolonga, André Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, Lucas Beyer, Olivier Bachem, Michael Tschannen, Marcin Michalski, Olivier Bousquet, Sylvain Gelly, and Neil Houlsby. A large-scale study of representation learning with the visual task adaptation benchmark. *arXiv: Computer Vision and Pattern Recognition*, 2019b.
- Jun Zhang, Zhi hui Zhan, Ying Lin, Ni Chen, Yue jiao Gong, Jinghui Zhong, Henry Shu hung Chung, Yun Li, and Yu hui Shi. Evolutionary computation meets machine learning: A survey. *IEEE Computational Intelligence Magazine*, 6:68–75, 2011.
- Wen Zhang, Lingfei Deng, Lei Zhang, and Dongrui Wu. A survey on negative transfer. 2020.
- Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. *ArXiv*, abs/1611.01578, 2017.



## A EXTENDED RELATED WORK SURVEY

The proposed method is designed to learn a unbounded number of tasks in a continual learning fashion. In such contexts it aims to learn each task with higher quality and efficiency by automating and optimizing the knowledge transfer among any subset of tasks that can provide useful knowledge to one another. The proposed model is designed to be immune from common multitask learning pitfalls: catastrophic forgetting, gradients interference, negative transfer. Cross-task **transfer-learning** has gained popularity, especially through transfer learning from a model pre-trained on a large amount of data for one or a few general tasks, and then fine-tuned on a small amount of data for a related downstream task. This approach has been shown to be very effective in a wide variety of problems across many modalities, including language (Devlin et al., 2019; Raffel et al., 2020) and vision (Dosovitskiy et al., 2021; He et al., 2016).. The success of transfer-learning applications hinges on adequate prior knowledge selection to avoid typical **negative transfer** pitfalls (Rosenstein, 2005; Wang et al., 2019). Common solutions rely on data or model selection techniques, often putting emphasis on the efficiency of the exploration (Zhang et al., 2020; Mensink et al., 2021), also method aiming to automate knowledge selection at a layer level have been proposed Sun et al. (2020a). Transfer learning capabilities are critical for **multitask models**. ML models trained jointly on multiple tasks can be affected by **gradients interference** if any subset of parameters receive gradients jointly from multiple sources (Chen et al., 2018; Yu et al., 2020), and by **catastrophic forgetting** of prior knowledge as new tasks are learned (McCloskey & Cohen, 1989; French, 1999). These knowledge loss problems can be alleviated with weighted combination of tasks (Liu et al., 2019b; Sun et al., 2020b) and gradient transformation methods (Chen et al., 2018; Sener & Koltun, 2018; Kendall et al., 2018). Stronger guarantees are provided by methods that compartmentalize task specific knowledge in dedicated parameter subsets (Rebuffi et al., 2017; Houlsby et al., 2019; Rusu et al., 2016; Rosenfeld & Tsotsos, 2020). Addressing catastrophic forgetting and identifying what subset of parameters/knowledge that is beneficial to share with each task is also critical for **continual learning** or life long learning methods (McCloskey & Cohen, 1989; French, 1999; Ramesh & Chaudhari, 2022).

The proposed method relies on an evolutionary approach to jointly search the spaces of models architectures, hyperparameters, and prior knowledge selection while optimizing for an possibly multi-factor non-differentiable reward function. The automation of **hyperparameter tuning** has been commonly addressed with Bayesian optimization (Srinivas et al., 2010; Bergstra et al., 2011; Snoek et al., 2012), evolutionary methods have also been explored for this purpose (Jaderberg et al., 2017; Zhang et al., 2011). Hyperparameters tuning can be considered related to the **neural architecture search** (NAS), as architectures can be defined by the selection of a sequence of architectural hyperparameters. Initially, NAS methods have been based on reinforcement learning techniques (Zoph & Le, 2017) but also sample efficient evolutionary approaches have also proposed (Real et al., 2019; Maziarz et al., 2018). Alternative NAS methods focusing on more efficient parameter-sharing (Pham et al., 2018; Liu et al., 2019a; Kokiopoulou et al., 2019) or optimization for multi-factor quality/cost trade-offs (Tan et al., 2019) have also been explored.

The proposed method is capable to dynamically extend the system, adding capacity or novel structures in an unconstrained fashion. A few methods have been proposed to achieve **dynamic architecture extensions** (Chen et al., 2016; Cai et al., 2018), some also focusing on an unbounded stream of tasks (Yoon et al., 2018; Yao et al., 2020), or achieving immunity from catastrophic forgetting (Rusu et al., 2016; Li & Hoiem, 2018; Li et al., 2019; Rosenfeld & Tsotsos, 2020).

The proposed method is sparsely activated, thus the unbounded growth of knowledge and parameters is decoupled from the growth of computational cost. The growth in capabilities of state of the art models often requires growth in terms of trainable parameters (Kaplan et al., 2020). **Sparse activation** techniques at sub-layer level (Shazeer et al., 2017; Du et al., 2021) or network route level (Fernando et al., 2017) allow to decouple model size growth from compute cost. This is achieved by integrating a **routing technique** that selects the appropriate subset of parameters storing the most relevant knowledge for each task, sample or token/patch.

The ability of jointly solve a **large amount of tasks** is commonly associated with progress toward Artificial General Intelligence (AGI). Advancements in scaling language models (Brown et al., 2020; Thoppilan et al., 2022) allowed to achieve novel discourse, reasoning and zero/few shot learning capabilities that can be applied to new tasks without/minimal additional training. Recent work aims

to extend these achievements beyond text modality by defining static architectures for an extended subset of modalities (Alayrac et al., 2022; Reed et al., 2022). These are few examples of the ML models contributing to the line of research achieving incremental milestone toward AGI. Though, each models is trained from scratch with considerable resources consumption. The introduction of abstractions allowing to modularize, dynamically extend and reuse these large models may contribute to accelerate the rate of innovation.

## B EXPERIMENTS REPRODUCIBILITY AND DETAILS

All the experiments reported in this paper can be reproduced by using the following public resources:

- The ViT model definition and checkpoints published by Steiner et al. (2021). These resources are available at [https://github.com/google-research/vision\\_transformer](https://github.com/google-research/vision_transformer) and distributed under the [Apache License 2.0](#).
- Published code of the proposed method: <https://...> (Anonymized, refer to supplementary material).
- All the used datasets are publicly available via the Tensorflow Datasets image classification, catalog. Refer to <https://www.tensorflow.org/datasets/catalog/overview> for detailed information regarding each dataset licence and other metadata. Table 8 reports exact dataset splits and reference for each task.

We also publish the  $\mu$ 2Net checkpoint resulting from the large-scale multitask experiment reported in Section 5. This checkpoint can be used for inference on any of the 69 learned image classification tasks, or for further analysis, or even to be extended with additional tasks or methods. For information about the checkpoint and its license refer to: <https://...> (Anonymized)

References for mentioned state of the art public model for different tasks are sourced from <https://paperswithcode.com> as of September 2022.

The VTAB-1k results are reported for reference and are not directly comparable to the state of the art, as the benchmark definition specifies that VTAB-full tasks cannot be used for pre-training.

The initial experiments reported in Sections 4 and 5.1 have been executed on a TPUv3 (Jouppi et al., 2017) machine with 8 cores. While, all the following experiments have been executed on a larger scale infrastructure using 32 TPUv4 chips in MegaCore mode, by using the Pathways orchestration layer (Barham et al., 2022).

Table 7 reports more details for each training segment.

Table 7: Details for the different training segments of the large scale continual learning experiment described in Section 5.

Training segment	TPU			#Tasks	#Params (M)	Activated params per task (%)
	core-hours	#cores	type			
ViT tasks 10 iters	4949	8	TPUv3	3	659	46.6%
VTAB-full 2 iters	5927	32	TPUv4	26	4400	7.0%
ViT tasks +1 iter	541	32	TPUv4	26	4577	5.2%
VDD 2 iters	1507	32	TPUv4	36	7790	3.9%
ViT tasks +1 iter	564	32	TPUv4	36	7854	3.9%
VTAB-full +1 iter	2266	32	TPUv4	36	7596	4.0%
Char. class. 2 iters	1785	32	TPUv4	44	9368	3.3%
VTAB-1k 2 iters	271	32	TPUv4	69	13087	2.3%
ViT tasks +1 iter	568	32	TPUv4	69	13090	2.3%




Figure 3: Graph representing the architecture of the multitask system solving jointly 69 image classification tasks generated by the large scale continual learning experiment described in Section 5. Each task is identified with a unique color. Bottom triangular nodes represent the data input of each task. Top rectangular nodes represent the head layer of each task. Each edges sequence of the same color connecting a task input to its head, a *path*, defines the layers sequence composing the model for each task. Each path traverses 27 round nodes representing ViT L/16 internal layers (see Section 3.5) in the following order from bottom to top: patch embedding, class token, position embedding and 24 transformer layers. Internal nodes are represented with the color of the task on which the parameters of the corresponding layer were trained last. Except for the gray nodes that have not received gradient updates from any of the 69 tasks and still carry the parameters of the root model that were loaded from a checkpoint of a ViT L/16 pretrained on the imagenet-21k dataset (see Section 5) (Video: [youtu.be/Hf88Ge0eiQ8](https://youtu.be/Hf88Ge0eiQ8)).




Figure 4: Correlation in log scale with the standard deviation measured during the variance analysis conducted on Visual Domain Decathlon benchmark by running the training on 3 parallel replicas of the system. We display the 2 metrics that are most correlated with the standard deviation: error rate computed on the test set (left) and training samples per class (right). The red line is fitted to minimize the squared distance to the set of points.




Figure 5: Distributions of the hyperparameter values used by the best models of the 69 image classification tasks at the end of the large scale continual learning experiment described in Section 5.

Table 8: Datasets splits and reference (part 1 of 2). For each dataset used in the experiments, this table reports: 1) dataset name indicative of the Tensorflow Datasets Catalogs identification string and linking to the corresponding catalog page ("visual\_domain\_decathlon" has been abbreviated as "vdd", and "diabetic\_retinopathy\_detection" as "drd"), 2) train, validation and test data splits, represented with the [standard Tensorflow Datasets format](#) ("validation" has been abbreviated as "val"). 3) corresponding scientific publication reference. Datasets are listed in the order of introduction into the system.

Notes:

[1] The test split of the [imagenet\\_v2](#) dataset is used as validation set for [imagenet2012](#).

[2] The test split of the [cifar10\\_1](#) dataset is used as validation set for [cifar10](#).

[3] The VTAB-full benchmark also includes the cifar100 task. Cifar100 has been introduced to the  $\mu$ 2Net system as part of the initial benchmark. In the VTAB-full results tables we refer to the top 1 test accuracy achieved in the latest cifar100 training iteration without retraining it as part of the VTAB-full active training iteration.

[4] The definition for the VTAB standard and additional tasks has been sourced from [https://github.com/google-research/task\\_adaptation/tree/master/task\\_adaptation/data](https://github.com/google-research/task_adaptation/tree/master/task_adaptation/data).

[5] VTAB additional task, not included in the standard scoring set. These tasks were added to further scale the system and analyze transfer across related tasks.

Name	Splits			Reference
	Train	Val.	Test	
<a href="#">imagenet2012</a>	train	<a href="#">imagenet_v2:test</a> <sup>[1]</sup>	val	(Russakovsky et al., 2015)
<a href="#">cifar100</a>	train[:98%]	train[98%:]	test	(Krizhevsky, 2009)
<a href="#">cifar10</a>	train	<a href="#">cifar10_1:test</a> <sup>[2]</sup>	test	(Krizhevsky, 2009)
<b>VTAB-full benchmark</b> <sup>[3][4]</sup>				
<a href="#">caltech101</a>	train[:2754]	train[2754:]	test	(Fei-Fei et al., 2004)
<a href="#">dtd</a>	train	val	test	(Cimpoi et al., 2014)
<a href="#">oxford_flowers102</a>	train	val	test	(Nilsback & Zisserman, 2008)
<a href="#">oxford_iiit_pet</a>	train[:2944]	train[2944:]	test	(Parkhi et al., 2012)
<a href="#">sun397</a>	train	val	test	(Xiao et al., 2010)
<a href="#">svhn_cropped</a>	train[:65931]	train[65931:]	test	(Netzer et al., 2011)
<a href="#">patch_camelyon</a>	train	val	test	(Veeling et al., 2018)
<a href="#">eurosat/rgb</a>	train[:16200]	train[16200:21600]	train[21600:]	(Helber et al., 2019)
<a href="#">resisc45</a>	train[:18900]	train[18900:25200]	train[25200:]	(Cheng et al., 2017)
<a href="#">drd/btgraham-300</a>	train	val	test	(Kaggle & EyePacs, 2015)
<a href="#">clevr/count_cylinders</a> <sup>[5]</sup>	train[:63000]	train[63000:]	val	(Johnson et al., 2017)
<a href="#">clevr/count_all</a>	train[:63000]	train[63000:]	val	(Johnson et al., 2017)
<a href="#">clevr/closest_object_distance</a>	train[:63000]	train[63000:]	val	(Johnson et al., 2017)
<a href="#">dmlab</a>	train	val	test	(Zhai et al., 2019a)
<a href="#">dsprites/label_x_position</a>	train[:589824]	train[589824:663552]	train[663552:]	(Klindt et al., 2021)
<a href="#">dsprites/label_orientation</a>	train[:589824]	train[589824:663552]	train[663552:]	(Klindt et al., 2021)
<a href="#">kitti/closest_object_distance</a> <sup>[5]</sup>	train	val	test	(Geiger et al., 2012)
<a href="#">kitti/count_vehicles</a> <sup>[5]</sup>	train	val	test	(Geiger et al., 2012)
<a href="#">kitti/closest_vehicle_distance</a>	train	val	test	(Geiger et al., 2012)
<a href="#">smallnorb/label_category</a> <sup>[5]</sup>	train	test[:50%]	test[50%:]	(LeCun et al., 2004)
<a href="#">smallnorb/label_lighting</a> <sup>[5]</sup>	train	test[:50%]	test[50%:]	(LeCun et al., 2004)
<a href="#">smallnorb/label_azimuth</a>	train	test[:50%]	test[50%:]	(LeCun et al., 2004)
<a href="#">smallnorb/label_elevation</a>	train	test[:50%]	test[50%:]	(LeCun et al., 2004)
<b>Visual domain decathlon benchmark</b>				
<a href="#">vdd/imagenet12</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/svhn</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/cifar100</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/gtsrb</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/daimlerpedcls</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/omniglot</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/ucf101</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/aircraft</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/dtd</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)
<a href="#">vdd/vgg-flowers</a>	train	val[:50%]	val[50%:]	(Bilen et al., 2017)

Continues in Table 9 ...

Table 9: Datasets splits and reference (part 2 of 2).

Name	Splits			Reference
	Train	Val.	Test	
... Continues from Table 8				
<b>Multitask Character Classification Benchmark</b>				
emnist/digits	train[5%:]	train[:5%]	test	(Cohen et al., 2017)
emnist/letters	train[5%:]	train[:5%]	test	(Cohen et al., 2017)
kmnist	train[5%:]	train[:5%]	test	(Clanuwat et al., 2018)
mnist	train[5%:]	train[:5%]	test	(LeCun et al., 1998)
omniglot	train	small1	small2	(Lake et al., 2015)
cmaterdb/bangla	train[20%:]	train[:20%]	test	(Das et al., 2012b;a)
cmaterdb/devanagari	train[20%:]	train[:20%]	test	(Das et al., 2012b;a)
cmaterdb/telugu	train[20%:]	train[:20%]	test	(Das et al., 2012b;a)
<b>VTAB 1k benchmark<sup>[4]</sup></b>				
caltech101	train[:800]	train[2754:2954]	test	(Fei-Fei et al., 2004)
cifar100	train[:800]	train[45000:45200]	test	(Krizhevsky, 2009)
cifar10	train[:800]	train[45000:45200]	test	(Krizhevsky, 2009)
dtd	train[:800]	val[:200]	test	(Cimpoi et al., 2014)
oxford_flowers102	train[:800]	val[:200]	test	(Nilsback & Zisserman, 2008)
oxford_iiit_pet	train[:800]	train[2944:3144]	test	(Parkhi et al., 2012)
sun397	train[:800]	val[:200]	test	(Xiao et al., 2010)
svhn_cropped	train[:800]	train[65931:66131]	test	(Netzer et al., 2011)
patch_camelyon	train[:800]	val[:200]	test	(Veeling et al., 2018)
eurosat/rgb	train[:800]	train[16200:16400]	train[21600:]	(Helber et al., 2019)
resisc45	train[:800]	train[18900:19100]	train[25200:]	(Cheng et al., 2017)
drd/btgraham-300	train[:800]	val[:200]	test	(Kaggle & EyePacs, 2015)
clevr/count_cylinders <sup>[5]</sup>	train[:800]	train[63000:63200]	val	(Johnson et al., 2017)
clevr/count_all	train[:800]	train[63000:63200]	val	(Johnson et al., 2017)
clevr/closest_object_distance	train[:800]	train[63000:63200]	val	(Johnson et al., 2017)
dmlab	train[:800]	val[:200]	test	(Zhai et al., 2019a)
dsprites/label_x_position	train[:800]	train[589824:590024]	train[663552:]	(Klindt et al., 2021)
dsprites/label_orientation	train[:800]	train[589824:590024]	train[663552:]	(Klindt et al., 2021)
kitti/closest_object_distance <sup>[5]</sup>	train[:800]	val[:200]	test	(Geiger et al., 2012)
kitti/count_vehicles <sup>[5]</sup>	train[:800]	val[:200]	test	(Geiger et al., 2012)
kitti/closest_vehicle_distance	train[:800]	val[:200]	test	(Geiger et al., 2012)
smallnorb/label_category <sup>[5]</sup>	train[:800]	test[:200]	test[50%:]	(LeCun et al., 2004)
smallnorb/label_lighting <sup>[5]</sup>	train[:800]	test[:200]	test[50%:]	(LeCun et al., 2004)
smallnorb/label_azimuth	train[:800]	test[:200]	test[50%:]	(LeCun et al., 2004)
smallnorb/label_elevation	train[:800]	test[:200]	test[50%:]	(LeCun et al., 2004)

**Algorithm 1** Pseudocode for one active task iteration

---

```

1: Active task:  $t$ 
2: Set of all the models currently in the multitask system:  $\mathcal{M}$ 
3: Active population:  $\mathcal{A} \leftarrow \{m \mid m \in \mathcal{M} \wedge m \text{ trained on } t\}$ 
4: for #generations do
5:   for #child-models do
6:     ▷ Sample parent model
7:     Parent model:  $p \leftarrow \text{none}$ 
8:     for Candidate parent model:  $\hat{p} \in [\text{sorted}_{score}(\mathcal{A}), \text{sorted}_{random}(\mathcal{M} \setminus \mathcal{A})]$  do
9:       if  $0.5^{\#selections(\hat{p}, t)} > x \sim \text{Uniform}([0, 1])$  then
10:          $p \leftarrow \hat{p}$ 
11:         break
12:       end if
13:     end for
14:     if  $p = \text{none}$  then
15:        $p \sim \text{Uniform}(\mathcal{A} \cup \mathcal{M})$ 
16:     end if
17:     ▷ Sample child model
18:     Set of mutations:  $\Delta \leftarrow \{make\text{-trainable-head}\}$ 
19:     for Candidate mutation:  $\hat{\delta} \in possible\text{-mutations}(p)$  do
20:       if  $\mu > x \sim \text{Uniform}([0, 1])$  then
21:          $\Delta \leftarrow \Delta \cup \{\hat{\delta}\}$ 
22:       end if
23:     end for
24:     Untrained child model:  $c_0 \leftarrow apply\text{-mutations}(p, \Delta)$ 
25:     ▷ Train child model
26:     Retained child model:  $c \leftarrow \text{none}$ 
27:     for  $i \in [1, \dots, \#train\text{-cycles}]$  do
28:        $c_i \leftarrow train(c_{i-1}, \min(1 \text{ epoch}, \#samples\text{-cap}))$ 
29:       if  $score(c_i) \geq \max(\{score(c) \mid c \neq \text{none}\} \cup \{score(p) \mid p \text{ trained on } t\} \cup \{-\infty\})$  then
30:          $c \leftarrow c_i$ 
31:       end if
32:     end for
33:     if  $c \neq \text{none}$  then
34:        $\mathcal{A} \leftarrow \mathcal{A} \cup \{c\}$ 
35:     end if
36:   end for
37: end for
38:   ▷ Keep only the best model for  $t$ 
39:  $\mathcal{M} \leftarrow \{\operatorname{argmax}_{m \in \mathcal{A}} score(m)\} \cup \{m \mid m \in \mathcal{M} \wedge m \text{ not trained on } t\}$ 

```

---