This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation.

Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

Neural Architecture Search of Deep Priors: Towards Continual Learning without Catastrophic Interference

Martin Mundt^{*}, Iuliia Pliushch^{*}, Visvanathan Ramesh Goethe University, Frankfurt, Germany

{mmundt, pliushch, vramesh}@em.uni-frankfurt.de

Abstract

In this paper we analyze the classification performance of neural network structures without parametric inference. Making use of neural architecture search, we empirically demonstrate that it is possible to find random weight architectures, a deep prior, that enables a linear classification to perform on par with fully trained deep counterparts. Through ablation experiments, we exclude the possibility of winning a weight initialization lottery and confirm that suitable deep priors do not require additional inference. In an extension to continual learning, we investigate the possibility of catastrophic interference free incremental learning. Under the assumption of classes originating from the same data distribution, a deep prior found on only a subset of classes is shown to allow discrimination of further classes through training of a simple linear classifier.

1. Introduction

Prevalent research routinely inspects continual deep learning through the lens of parameter inference. As such, an essential desideratum is to overcome the threat of *catastrophic interference* [21, 27]. The latter describes the challenge to avoid accumulated knowledge from being continuously overwritten through updates on currently observed data instances. As outlined by recent reviews in this context [26, 22], specific mechanisms have primarily been proposed in incremental classification scenarios. Although precise techniques vary drastically across the literature, the common focus is a shared goal to maintain a deep encoder's representations, in an effort to protect performance from continuous degradation [19, 16, 6, 29, 24, 31, 25, 23, 1].

In this work, we embark on an alternate path, one that asks a fundamentally different question: What if we didn't need to infer a deep encoder's parameters and thus didn't have to worry about catastrophic interference altogether? This may initially strike the reader as implausible. However, successfully solving the discriminative classification task has but one intuitive requirement: data belonging to different classes must be easily separable.

Inspired by the possibility of recovering encodings generated by random projections in signal theory [4], selected prior works have thus investigated the utility of (neural) architectures with entirely random weights [30, 14], see [5] for a recent review. Their main premise is that if a deep network with random weights is sensitive to the low-level statistics, it can progressively enhance separability of the data and a simple subsequent linear classification can suffice [10]. Whereas promising initial results have been presented, a practical gap to fully trained systems seems to remain. We posit that this is primarily a consequence of the complexity involved in hierarchical architecture assembly.

In the spirit of these prior works on random projections, we frame the task of finding an adequate deep network with random weights, hence referred to as a *deep prior*, from a perspective of neural architecture search (NAS) [2, 28, 36]. We empirically demonstrate that there exist configurations that achieve rivalling accuracies to those of their fully trained counterparts. We then showcase the potential of deep priors for continual learning. We structure the remainder of the paper according to our contributions:

1. We formulate deep prior neural architecture search (DP-NAS) based on random weight projections. We empirically show that it is possible to find hierarchies of operations which enable a simple linear classification. As DP-NAS does not require to infer the encoders' parameters, it is magnitudes of order faster than conventional NAS methods.

2. Through ablation experiments, we observe that deep priors are not subject to a weight initialization lottery. That is, performance is consistent across several randomly drawn sets of weights. We then empirically demonstrate that the best deep priors capture the task through their structure. In particular, they do not profit from additional training.

3. We corroborate the use of deep priors in continual learning. Without parameter inference in the deep architecture, disjoint task settings with separate classifiers per task are trivially solved by definition. In scenarios with shared clas-

^{*}equal contribution

sifier output, we empirically demonstrate that catastrophic interference can easily be limited in a single prediction layer. Performances on incremental MNIST [18], Fashion-MNIST [34] and CIFAR-10 [17] are shown to compete with complex deep continual learning literature approaches.

We conclude with limitations and prospects.

2. Deep Prior Neural Architecture Search: Hierarchies of Functions with Random Weights as a Deep Prior

In the recent work of [33], the authors investigate the role of convolution neural networks' structure in image generation and restoration tasks. Their key finding is that a significant amount of image statistics is captured by the structure, even if a network is initialized randomly and subsequently trained on individual image instances, rather than data populations. Quoting the authors "the structure of the network needs to resonate with the structure of the data" [33], referred to as a *deep image prior*. We adopt this terminology.

To start with an intuitive picture behind such a deep prior, we first conduct a small experiment to showcase the effect in our classification context, before referring to previous theoretical findings. Specifically, in table 1, we take three popular image classification datasets: MNIST [18], FashionMNIST [34] and CIFAR-10 [17], and proceed to train a single linear classification layer to convergence. We report the accuracy on the test set in bold print in the first row and then repeat the experiment in two further variations.

First, we compute two convolutional neural architectures with random weights, drawn from Gaussian distributions according to [12], before again training a single linear classification layer to convergence on this embedding. One of these architectures is the popular LeNet [18] with 3 convolutional layers, intermittent average pooling operations and rectified linear unit (ReLU) activation functions. The other is a 2 layer convolutional layer architecture, again with Re-

	MNIST	Fashion	CIFAR10
Linear Classifier (LC)	91.48	85.91	41.12
Random LeNet + LC	88.76	<mark>80.33</mark>	43.40
Trained LeNet + LC	98.73	90.89	58.92
Random CNN-2L + LC	98.01	89.29	60.26
Trained CNN-2L + LC	98.86	92.13	70.86

Table 1. Example classification accuracies (in %) when computing the randomly projected embedding and training a subsequent linear classifier are compared to training the latter directly on the raw input data (bold reference value). Values are colored in blue if the random weight deep prior facilitates the task. Red values illustrate a disadvantage. Fully trained architecture accuracies are provided for completeness. Example architectures are the three convolutional layer LeNet [18] with average pooling and a similar two-layer convolutional architecture with max pooling. LUs, but with max pooling operations. For convenience we have color coded the results in red and blue, red if the obtained accuracy is worse than simply training the linear classifier on the raw input, blue if the result is improved. We can observe that the 2 layer max-pool architecture with random weights dramatically improves the classification accuracy, even though none of the encoder's weights were trained. In a second repetition of the experiment we also train the convolutional neural architectures' weights to full convergence. Expectedly, this improves the performance. However, we can also observe that the gap to the random encoder version of the experiment is less than perhaps expected.

Before proceeding to further build on these preliminary results, we highlight two seminal works, which detail the theoretical understanding behind the values presented in table 1. The work of Saxe et al. [30] has proven that the combination of random weight convolutions and pooling operations can have inherent frequency selectivity with wellknown local translation invariance. Correspondingly, they conclude that large portions of performance in classification stems from the choice of architecture, similar to observations of [33] for generation. The work by Girves et al. [10] further proves that deep neural networks with random i.i.d. Gaussian weights preserve the metric structure of the data throughout the layer propagation. For the specific case of piecewise linear activations, such as ReLU, they further show that a sensitivity to angles is amplified by modifying distances of points in successive random embedding. This mechanism is suspected to draw similar concepts closer together and push other classes away, promoting separation.

2.1. Deep Prior Neural Architecture Search

Leveraging previously outlined theoretical findings and encouraged by our initial ablation experiment, we formulate the first central hypothesis of this work:

Hypothesis 1 - deep prior neural architecture search

A hierarchical neural network encoder with random weights acts as a *deep prior*. We conjecture that there exist deep priors, which we can discover through modern architecture search techniques, that lead to a classification task's solution to the same degree of a fully trained architecture.

A crucial realization in the practical formulation of such a deep prior neural architecture search (DP-NAS) is that we are not actually required to infer the weights of our deep neural architectures. Whereas previous applications of NAS [2, 28, 36] have yielded impressive results, their practical application remains limited due to the excessive computation involved. This is because neural architecture search, in independence of its exact formulation, requires a reward signal to advance and improve. For instance, the works of



Figure 1. Illustration of the steps involved in Deep Prior Neural Architecture Search (DP-NAS). In DP-NAS, a reinforcement learning qagent is trained to learn to assemble suitable architectures with random weights across a diverse range of options (illustrated in the first box). In contrast to conventional NAS, these architectures do not have their weights trained. Instead, a randomly projected dataset embedding is calculated (second box) and only a linear classifier is trained (third box). The essential premise is that suitable deep priors contain functional forms of selectivity with respect to low level data statistics, such that classes are easily linearly separable in the transformed space. The q-agent learns to find the latter for a given task (fourth box with outer loop). For convenience, parts of the algorithm that do not involve any parameter inference are colored in blue, with red parts illustrating the only trained components.

Baker *et al.* [2] or Zoph *et al.* [36] require full training of each deep neural network to use the validation accuracy to train the agent that samples neural architectures. Reported consumed times for a search over thousands of architectures are thus regularly on the order of multiple weeks with tens to hundreds of GPUs used. Our proposed DP-NAS approach follows the general formulation of NAS, alas significantly profits from the architecture weights' random nature.

We visualize our procedure in figure 1. In essence, we adopt the steps of MetaQNN [2], without actual deep neural network training. It can be summarized in a few steps:

- 1. We sample a deep neural architecture and initialize it with random weights from Gaussian distributions.
- 2. We use this deep prior to calculate a single pass of the entire training dataset to compute its embedding.
- 3. The obtained transformed dataset is then used to evaluate the deep prior's suitability by training a simple linear classifier. Based on a small held-out set of training data, the latter's validation accuracy is stored jointly with the deep prior topology into a replay buffer.
- 4. The current architecture, together with random previous samples stored in the replay buffer, are then used to update the q-values of a reinforcement learner.

Once the search advances, we progressively decrease an epsilon value, a threshold value for a coin flip that determines whether a deep prior is sampled completely randomly or generated by the trained agent, from unity to zero.

To get a better overview, we have shaded the parts of figure 1 that require training in red and parts that do not in blue. As the sampling of a neural architecture is computationally negligible and a single computation of the deep prior embedding for the dataset on a single GPU is on the order of seconds, the majority of the calculation is now redirected to the training of a linear classifier and updating of our q-agent. Fortunately, the former is just a matrix multiplication, the latter is a mere computation of the Bellman equation in tabular q-learning. The training per classifier thus also resides in the seconds regime. Our DP-NAS code is available at: https://github.com/ccc-frankfurt/DP-NAS

2.2. Ablation study: DP-NAS on FashionMNIST

To empirically corroborate hypothesis 1, we conduct a DP-NAS over 2500 architectures on FashionMNIST. Here, the related theoretical works, introduced in the beginning of this section, serve as the main motivation behind our specific search space design. Correspondingly, we have presently limited the choice of activation function to ReLUs and the choice of pooling operation to max pooling. We search over the set of convolutions with {16, 32, 64, 128, 256, 512, 1024} random Gaussian filters, drawn according to [12], of size $\{1, 3, 5, 7, 9, 11\}$ with options for strides of one or two. Similarly, potential receptive field sizes and strides for pooling are both $\{2, 3, 4\}$. We allow for the sampling of skip connections in the style of residual networks [11, 35], where a parallel connection with convolution can be added to skip two layers. We presently search over architectures that contain at least one and a maximum of twelve convolutional layers. The subsequent linear classifier is trained for 30 epochs using an Adam optimizer [15] with a learning rate of 10^{-3} and a mini-batch size of 128. We have applied no additional pre-processing, data augmentation, dropout, weight decay regularization, or similar techniques. We start with a 1500 long exploration phase, $\epsilon = 1$, before starting to exploit the learned q-agent by reducing ϵ by 0.1 every subsequent 100 architectures.

Figure 2 shows the obtained DP-NAS results. The graph depicts the mean average reward, *i.e.* the rolling validation accuracy over time, as a function of the number of sampled architectures. Vertically dashed lines indicate the epsilon



Figure 2. DP-NAS on FashionMNIST. The y-axis shows the moving average reward, *i.e.* rolling validation accuracy, whereas the color of each architecture encodes the individual performance. Dashed vertical lines show the used epsilon greedy schedule.

greedy schedule. In addition, each plotted point is color coded as to represent the precisely obtained accuracy of the individual architecture. From the trajectory, we can observe that the agent successfully learns suitable deep priors over time. The best of these deep priors enable the linear classifier to reach accuracies around 92%, values typically reached by fully trained networks without augmentation.

At this point, we can empirically suspect these results to already support hypothesis 1. To avoid jumping to premature conclusions, we further corroborate our finding by investigating the role of the particularly drawn random weights from their overall distribution, as well as an experiment to confirm that the best found deep priors do in fact not improve significantly with additional parameter inference.

2.3. Did we get lucky on the initialization lottery?

We expand our experiment with a further investigation with respect to the role of the precisely sampled weights. In particular, we wish to empirically verify that our found deep priors are in fact deep priors. In other words, the observed success is actually due to the chosen hierarchy of functions projecting into spaces that significantly facilitate classification, rather than being subject to lucky draws of good sets of precise weight values when sampling from a Normal distribution. Our second hypothesis is thus:

Hypothesis 2 - deep priors and initialization lottery

A hierarchical neural network encoder with random weights acts as a *deep prior*, irrespectively of the precisely sampled weights. In particular, it is not subject to an initialization lottery.

This train of thought is motivated from recent findings on the lottery ticket hypothesis [9], where it is conjectured that there exists an initialization lottery in dense randomly initialized feed-forward deep neural networks. According to the original authors, winning this lottery is synonymous

FashionMNIST accuracy across 10 random intializations



Figure 3. Accuracy for 18 deep priors across 10 experimental repetitions with randomly sampled weights. Six deep priors for respective 3 performance segments of figure 2 (0-5 low, 6-11 median, 12-17 top) have been selected. Result stability suggests that the architecture composition is the primary factor in performance.

with finding an initial set of weights that enables training.

To empirically confirm that the structure is the imperative element, we randomly select 18 deep priors from our previous FashionMNIST search. 6 of these are sampled from the lowest performing architectures, 6 are picked around the median, and 6 are chosen to represent the top deep priors. For each of these deep priors, we repeat the linear classifier training for 10 independently sampled sets of weights. The respective figure 3 shows the median, upper and lower quartiles, and deviations of the measured test accuracy. We observe that the fluctuation is limited, the results reproducible and the ordering of the deep priors is thus preserved. Whereas minor fluctuations for particular weight samples seem to exist, the results suggest that the architecture itself is the main contributor to obtained performance.

2.4. Are Deep Priors predictive of performance with parameter inference?

To finalize our initial set of ablation experiments, we empirically confirm that the best deep priors do "fully resonate" with underlying dataset statistics. If we conjecture the best deep priors to extract all necessary low level image statistics to allow for linear decision boundaries to effectively provide a solution, then we would expect no additional parameter inference to yield significant improvement.

Hypothesis 3 - deep priors and parameter inference

A deep image prior performs at least equivalently, if not significantly better, when parameters are additionally inferred from the data population. However, we posit that the best deep priors are already close to the achievable performance.



Figure 4. Accuracy comparison when only a linear classifier is trained on top of the deep prior and when the entire architecture is trained. The same 18 deep priors as picked in figure 3, representing respective 3 performance segments of figure 2 (0-5 low, 6-11 median, 12-17 top), are shown. Results demonstrate that the best deep priors enable a solution without parametric inference and performance improvement is negligible with additional training.

To investigate the above hypothesis, we pick the same 18 architectures as analyzed in the last subsection to again represent the different deep prior performance segments. In contrast to the random weight deep prior experiments, we now also fully train the deep architecture jointly with the classifier. We show the obtained test accuracies in comparison with the deep prior with exclusive training of the linear classifier in figure 4. We first observe that all trained architectures perform at least as well as their random counterparts. As expected, the full training did not make it worse. Many of the trained architectures with non-ideal deep priors still perform worse with respect to the best untrained deep priors. The latter all improve only very marginally, suggesting that adjustments to the parameters provide negligible value. The best deep priors seem to solve the task to the same degree that a fully trained deep network does. We thus see our third hypothesis to be empirically confirmed, and in turn the initial first hypothesis validated.

Although not crucial to our main hypotheses, interestingly, we can also observe that the ordering for the worst to intermediate deep priors in terms of achieved accuracy is not retained. We speculate that this is a consequence of heavy over-parametrization of many constructed deep architectures. The latter entails a high fitting capacity when trained and thus a higher ability to compensate misalignment.

3. Continual Learning with Deep Priors

With the foregoing section focusing on the general feasibility of deep priors, we now extend our investigation towards implications for continual learning. In neural networks the latter is particularly difficult, given that training overwrites parameters towards the presently observed data sub-population: the phenomenon of *catastrophic interference* [21, 27]. As highlighted in recent reviews [26, 22], the challenge is already significant when considering simple *class incremental* tasks, such as the prevalent scenario of splitting datasets into sequentially presented disjoint sets of classes. In contrast, we formulate our central hypothesis with respect to continual learning with deep priors:

Hypothesis 4 - deep priors and continual learning

If classes in a dataset originate from the same data distribution, finding a deep prior on a subset of dataset classes can be sufficient to span prospective application to the remainder of the unseen classes.

The above hypothesis is motivated from the expectation that a deep prior primarily captures the structure of the data through low-level image statistics. The latter can then safely be considered to be shared within the same data distribution. In particular, the hypothesis seems reasonable with hindsight knowledge of the findings in prior theoretical works, which we re-iterate to have proven that Gaussian random projections enhance seprability through inherent frequency and angular sensitivity [4, 30, 10]. A deep prior, found to respond well to a subset of classes, can then also be assumed to act as a good prior for other labels of instances drawn from the same data distribution. To give a practical example, we posit that a deep prior found for the t-shirt and coat classes transfers to images of pullovers and shirts under shared sensor and acquisition characteristics.

From this hypothesis, an intriguing consequence arises for deep continual learning. For instance, learning multiple disjoint tasks in sequence, that is sharing the neural network feature extractor backbone but training a separate task classifier, struggles with catastrophic interference in conventional training because the end-to-end functional is subject to constant change. In contrast, if the deep prior holds across these tasks, the solution is of trivial nature. Due to the absence of tuning the randomly initialized weights, by definition, the deep prior encoder is not subject to catastrophic interference in such continual data scenarios. We start with an empirical investigation of the practical validity of the above hypothesis in this scenario. With empirical confirmation in place, we then proceed to extend this investigation to a more realistic scenario, where a single task of continuously growing complexity is considered. We posit that a deep prior significantly facilitates this more generic formulation, as we only need to regulate inference at the prediction level.

3.1. Preliminaries: scenarios and methods

Before delving into specific experiments, we provide a short primer on predominantly considered assumptions and result interpretation, in order to place our work in context. We do not at this point provide a review on the vast continual learning literature and defer to the referenced surveys.

Continual learning scenarios

For our purposes of continual learning with deep priors, we investigate two configurations:

- 1. *Multi-head incremental learning:* in this commonly considered simplified scenario, sets of disjoint classes arrive in sequence. Here, each disjoint set presents its own task. Assuming the knowledge of a task id, separate prediction layers are trained while attempting to preserve the joint encoder representations for all tasks. Although often considered unrealistic, complex strategies are already required to address this scenario.
- 2. *Single-head incremental learning:* the scenario mirrors the above, alas lifts the assumption on the presence of a task identifier. Instead of inferring separate predictors per task, a single task is extended. The prediction layer's amount of classes is expanded with every arrival of a new disjoint class set. In addition to catastrophic interference in a deep encoder, interference between tasks is now a further nuisance. For instance, a softmax based linear prediction will now also tamper with output confidences of former tasks.

We do not add to ongoing discussions on when specific scenario assumptions have practical leverage, see *e.g.* [8] for the latter. In the context of this work, the multi-head scenario is compelling because a separate classifier per task allows to directly investigate hypothesis four. Specifically, we can gauge whether deep priors found on the first task are suitable for prospective tasks from the same distribution. For the single-head scenario, we need to continuously train a linear prediction layer on top of the deep prior. As such, we will also need to apply measures to alleviate catastrophic interference on this single learned operation. However, in contrast to maintaining an entire deep encoder, we would expect this to work much more efficiently.

A brief primer on interpreting reported results

Independently of the considered scenario, techniques to address catastrophic interference typically follow one of three principles: explicit parameter regularization [16, 19, 6], retention and rehearsal of a subset of real data (a core set) [29, 24], or the training of additional generative models to replay auxiliary data [1, 23, 25, 31]. For the latter two families of methods, the stored or generated instances of older tasks get interleaved with new task real data during continued training. Once more, we defer to the survey works for detailed explanations of specific techniques [26, 22]. For our purposes of demonstrating an alternative to researching continual deep learning from a perspective of catastrophic interference in deep encoders, it suffices to know that all

	Multi-head Accuracy [%]		
Method	MNIST	FashionMNIST	
EWC [16]	99.3 [<mark>6</mark>]	95.3 [8]	
RWalk [6]	99.3 [<mark>6</mark>]	-	
VCL + Core [24]	98.6 [<mark>8</mark>]	97.1 [<mark>8</mark>]	
VCL [24]	97.0 [<mark>8</mark>]	80.6 [8]	
VGR [8]	99.3 [<mark>8</mark>]	99.2 [8]	
DP	99.79	99.37	

Table 2. Average accuracy across 5 disjoint tasks with 2 classes for FashionMNIST and MNIST. Our deep prior approach provides competitive performance, even though the corresponding DP-NAS has only been conducted on the initial task.

these methods train encoders and construct complex mechanisms to preserve its learned representations.

We point out that contrasting performances between these techniques in trained deep neural networks is fairly similar to a comparison of apples to oranges. Whereas the essential goal is shared, the amount of used computation, storage or accessibility of data varies dramatically. Correspondingly, in our result tables and figures we provide a citation to each respective technique and an additional citation next to a particular accuracy value to the work that has reported the technique in the specific investigated scenario. Whereas we report a variety of these literature values, we emphasize that our deep prior architectures do not actually undergo any training. Our upcoming deep prior results should thus be seen from a perspective of providing an alternate way of thinking about the currently scrutinized continual learning challenges. In fact, we will observe that our deep prior experiments yield remarkably competitive performances to sophisticated algorithms.

3.2. Disjoint tasks: deep priors as a trivial solution

We start by investigating hypothesis four within the multi-head incremental classification framework. For this purpose, we consider the disjoint FashionMNIST and MNIST scenarios, where each subsequent independent task is concerned with classification of two consecutive digits. We repeat our DP-NAS procedure once for each dataset *on only the initial task* and thus only the first two classes. We do not show the precise search trajectories as they look remarkably similar to figure 2, with the main difference being a shift in accuracy towards a final reached 99.9% as a result of narrowing down the problem to two classes. Thereafter, we use the top deep prior and proceed to continuously learn linear classification predictions for the remaining classes.

We report the final averaged accuracy across all five tasks in table 2 and compare them with achieved accuracies in prominent literature. We can observe that for these investigated datasets our deep prior hypothesis four holds. With the average final accuracy surpassing 99 % on both datasets,



the originally obtained deep prior for the first task seems to transfer seamlessly to tasks two to five. By definition, as predictions of disjoint tasks do not interfere with each other, a trivial solution to this multi-head continual learning scenario has thus been obtained. This is in stark contrast to the referenced complexly tailored literature approaches.

3.3. Alleviating catastrophic interference in a single prediction layer on the basis of a deep prior

To further empirically corroborate our conjecture of hypothesis four, we conduct additional class incremental continual learning experiments in the single-head scenario. This scenario softens the requirement of task labels by continuing to train a single joint prediction layer that learns to accommodate new classes as they arrive over time. Whereas our deep prior (again only obtained on the initial two classes) does not require any training, we thus need to limit catastrophic interference in our linear classifier. Note that we can use any and all of the existing continual learning techniques for this purpose. However, we have decided to use one of the easiest conceivable techniques in the spirit of variational continual learning (VCL) [24] and iCarl [29]. The authors suggest to store a core set, a small sub-set of original data instances, and continuously interleave this set in the training procedure. Although they suggest to use involved extraction techniques, such as k-center or herding, we sample uniformly in favor of maximum simplicity. In contrast to prior works, as the deep prior remains untrained, we have the option to conserve memory by storing randomly projected embeddings instead of raw data inputs.

3.3.1 FashionMNIST revisited: single-headed

Our first set of single-head incremental classification experiments uses the deep prior found for FashionMNIST on the initial two classes, similar to our previous multi-head experimental section. To protect the single linear prediction, we store a random core set of 40 examples, in correspondence



Figure 6. Single-head FashionMNIST for varying core set sizes.

with prior experiments in the literature [8]. Figure 5 provides the respective result comparison. Among the reported techniques, storing 40 exemplars in a deep prior seems to significantly outperform storing 40 exemplars in VCL. Naturally, this is because VCL requires protection of all representations in the entire neural network, whereas the deep prior is agnostic to catastrophic interference and only the single layer classifier needs to be maintained. The popular parameter regularization technique Elastic Weight Consolidation (EWC) [16] is similarly outperformed.

We nevertheless observe a significant amount of performance degradation. To also corroborate our hypothesis four in this single-head scenario and show that this forgetting can be attributed exclusively to the linear classifier, we further repeat this experiment with increased amount of instances stored in the core set. For ease of readability, these additional results are shown in figure 6 for a stored amount of 40, 120, 240, 300 instances, corresponding to a respective 0.33%, 1%, 2% and 2.5% of the overall data. The black solid curve shows the achieved accuracy when all real data is simply accumulated progressively. Our first observation is that the final accuracy of the latter curve is very close to the final performance values reported in the full DP-NAS of figure 2, even though we have only found a deep prior on the first two classes. Once more, we find additional evidence in support of hypothesis four. This is further substantiated when observing the curves for the individual core set sizes. We can see that in the experiments with 2 and 2.5 % stored data instances, our deep prior beats very complex generative replay techniques. All three reported techniques: deep generative replay (DGR) [31], variational generative replay (VGR) [8] and open-set denoising variational auto-encoder (OCDVAE) [23] employ additional deep generative adversarial networks or variational auto-encoders to generate a full sized dataset to train on. Although this is an intriguing topic to explore in the context of generative modelling, our experiments indicate that for classification purposes our simple deep prior approach seems to have the upper hand.

Accuracy [%]	MNIST	CIFAR-10	
Method	A10, D2	A5, D1	A10, D5
EWC [16]	55.80 [<mark>6</mark>]	-	37.75 [13]
IMM [19]	67.25 [13]	32.36 [13]	62.98 [<mark>13</mark>]
DGR [31]	75.47 [13]	31.09 [13]	65.11 [<mark>13</mark>]
PGMA [13]	81.70 [13]	40.47 [13]	69.51 [<mark>13</mark>]
RWalk [6]	82.50 [6]	-	-
iCarl [29]	55.80 [<mark>6</mark>]	57.30 [1]	-
DGM [25]	-	64.94 [1]	-
EEC [1]	-	85.12 [1]	-
DP + core	76.31	58.13	65.15

Table 3. Average final single-head accuracy. MNIST is reported on 10 classes, after 5 increments containing two classes (A10, D2). For CIFAR-10, accuracies on 5 classes with class increments of size 1 (A5, D1), and on all 10 classes after two increments of 5 classes (A10, D5) are shown. Used core set sizes are 10 instances per task for MNIST, following the experiments of [6], and a total memory of 2000 instances for CIFAR-10, according to iCarl [29].

3.3.2 The easier MNIST and more difficult CIFAR-10

We finalize our experiments with an additional investigation of the MNIST and CIFAR-10 datasets in the singlehead scenario. Once more, we follow previous literature and store 10 randomly sampled data instances per task for MNIST, as suggested in Riemannian walk [6], and a maximum memory buffer of size 2000 for CIFAR-10, as suggested in iCarl [29]. Following our hypothesis, the DP-NAS is again conducted exclusively on the initial classes.

In table 3 we report the final obtained accuracy and compare it to literature approaches. Similar to VCL with core sets, we observe that storing a core set for our linear classifier on top of the deep prior outperforms the core set approach in iCarl. In comparison with the remaining methods we can see that the simple deep prior approach surpasses multiple methods and only falls behind a few select works. The latter can be attributed to additional storage and auxiliary model assumption. For instance, in EEC [1] an additional generative model learns to replay the full dataset embeddings. This procedure could find straightforward transfer to our randomly projected deep prior embeddings and is left for future work. For now, we conclude our work by highlighting that there exist alternate methods without parameter inference as potential solutions to explore for both deep neural network classification and continual learning.

4. Limitations and prospects

Domain incremental scenarios: Above continual learning experiments are limited in that they do not consider sequential data stream scenarios where arriving data x no longer is drawn from the same data distribution, *i.e.* the domain

 $p_t(x) \neq p_{t+1}(x)$ [7]. Previously postulated hypothesis four can no longer be expected to hold due to potential changes in image statistics. However, one could define a progressive version of DP-NAS, such that the random architecture found for the initial distribution is extended with functions and connections to accommodate further distributions.

Search space transformations: In similar spirit to the aforementioned point, other datasets may require more than the current angular and frequency selectivity of convolutional ReLU blocks. An intriguing future direction will be to explore random deep neural networks in a wider sense, with flexible activation functions or even interpretable transformations added to the search space. This could in turn provide a chance for a more grounded understanding of the relationship between the data distribution and the necessary transformations to accomplish a certain task.

Fully catastrophic interference free classification: For our current single-head continual learning experiments we have optimized a growing linear softmax classifier. Naturally this single prediction layer still suffers from catastrophic interference that needs to be alleviated. It will be interesting to lift this by examining generative alternatives, *e.g.* distance based decisions or mixture models.

Deterministic vs. stochastic deep prior: Our current deep priors do not fully leverage the weight distributions. After weights are sampled, the deep prior is treated as a deterministic processing block. As we do not train the deep prior, we conjecture that full sampling of weights with propagation of uncertainties, in the sense of Bayesian neural networks [3], can provide additional crucial information.

Autoencoding, inversion and compression: The present focus has been on classification. However, there also exists prior work on weight-tied random autoencoders [20]. [32] state that random autoencoders work surprisingly well, due to the symmetric and invertible structure, discarding only information on color, but preserving the one on geometry. In a similar vein to the experiments conducted in this paper, an appropriate deep prior could thus also be searched for.

5. Conclusion

In this paper we have analyzed the classification performance of the neural network structure independently of parametric inference. Using the proposed deep prior neural architecture search, we have shown that it is possible to find random weight architectures that rival their trained counterparts. Further experiments in continual learning lay open promising novel research directions that pursue an entirely different path from the present focus on catastrophic interference in deep encoder architectures.

References

- Ali Ayub and Alan R. Wagner. EEC: Learning to Encode and Regenerate Images for Continual Learning. *International Conference on Learning Representations (ICLR)*, 2021. 1, 6, 8
- [2] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing Neural Network Architectures using Reinforcement Learning. *International Conference on Learning Representations (ICLR)*, 2017. 1, 2, 3
- [3] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. 32nd International Conference on Machine Learning (ICML), 2:1613–1622, 2015. 8
- [4] Emmanuel J. Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, 2006. 1, 5
- [5] Weipeng Cao, Xizhao Wang, Zhong Ming, and Jinzhu Gao. Neurocomputing A review on neural networks with random weights. *Neurocomputing*, 275:278–287, 2018. 1
- [6] Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian Walk for Incremental Learning: Understanding Forgetting and Intransigence. *European Conference on Computer Vision (ECCV)*, 2018. 1, 6, 8
- [7] Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 8
- [8] Sebastian Farquhar and Yarin Gal. Towards Robust Evaluations of Continual Learning. International Conference on Machine Learning (ICML), Lifelong Learning: A Reinforcement Learning Approach Workshop, 2018. 6, 7, 8
- [9] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *International Conference on Learning Representations (ICLR)*, 2019. 4
- [10] Raja Giryes, Guillermo Sapiro, and Alex M. Bronstein. Deep Neural Networks with Random Gaussian Weights: A Universal Classification Strategy? *IEEE Transactions on Signal Processing*, 64(13):3444–3457, 2016. 1, 2, 5
- [11] Kun He, Yan Wang, and John Hopcroft. A Powerful Generative Model Using Random Weights for the Deep Image Representation. *Neural Information Processing Systems* (*NeurIPS*), 2016. 3
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *International Conference* on Computer Vision (ICCV), 2015. 2, 3
- [13] Wenpeng Hu, Zhou Lin, Bing Liu, Chongyang Tao, Zhengwei Tao, Dongyan Zhao, Jinwen Ma, and Rui Yan. Overcoming Catastrophic Forgetting for Continual Learning via Model Adaptation. *International Conference on Learning Representations (ICLR)*, 2019. 8

- [14] Guang Bin Huang, Dian Hui Wang, and Yuan Lan. Extreme learning machines: A survey. *International Journal of Machine Learning and Cybernetics*, 2(2):107–122, 2011.
- [15] Diederik P. Kingma and Jimmy Lei Ba. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations (ICLR)*, 2015. 3
- [16] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences (PNAS)*, 114(13):3521–3526, 2017. 1, 6, 7, 8
- [17] Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Technical report, Toronto, 2009. 2
- [18] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323, 1998.
 2
- [19] Sang Woo Lee, Jin Hwa Kim, Jaehyun Jun, Jung Woo Ha, and Byoung Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. *Neural Information Processing Systems (NeurIPS)*, pages 4653–4663, 2017. 1, 6, 8
- [20] Ping Li and Phan-Minh Nguyen. On random deep weighttied autoencoders: Exact asymptotic analysis, phase transitions, and implications to training. *International Conference* on Learning Representations (ICLR), 2019. 8
- [21] Michael McCloskey and Neal J. Cohen. Catastrophic Interference in Connectionist Networks : The Sequential Learning Problem. *Psychology of Learning and Motivation - Advances in Research and Theory*, 24(C):109–165, 1989. 1, 5
- [22] Martin Mundt, Yong Won Hong, Iuliia Pliushch, and Visvanathan Ramesh. A Wholistic View of Continual Learning with Deep Neural Networks: Forgotten Lessons and the Bridge to Active and Open World Learning. arXiv preprint arXiv:2009.01797, 2020. 1, 5, 6
- [23] Martin Mundt, Sagnik Majumder, Iuliia Pliushch, Yong Won Hong, and Visvanathan Ramesh. Unified Probabilistic Deep Continual Learning through Generative Replay and Open Set Recognition. arXiv preprint arXiv:1905.12019, 2019. 1, 6, 8
- [24] Cuong V. Nguyen, Yingzhen Li, Thang D. Bui, and Richard E. Turner. Variational Continual Learning. International Conference on Learning Representations (ICLR), 2018. 1, 6, 7
- [25] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jahnichen, and Moin Nabi. Learning to remember: A synaptic plasticity driven framework for continual learning. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019-June:11313–11321, 2019. 1, 6, 8
- [26] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual Lifelong Learning with Neural Networks: A Review. *Neural Networks*, 113:54– 71, 2019. 1, 5, 6

- [27] Roger Ratcliff. Connectionist Models of Recognition Memory: Constraints Imposed by Learning and Forgetting Functions. *Psychological Review*, 97(2):285–308, 1990. 1, 5
- [28] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Quoc Le, and Alex Kurakin. Large-Scale Evolution of Image Classifiers. *International Conference on Machine Learning (ICML)*, 2017. 1, 2
- [29] Sylvestre A. Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 6, 7, 8
- [30] Andrew M. Saxe, Pang Wei Koh, Zhenghao Chen, Maneesh Bhand, Bipin Suresh, and Andrew Y. Ng. On Random Weights and Unsupervised Feature Learning. *International Conference on Machine Learning (ICML)*, 2011. 1, 2, 5
- [31] Hanul Shin, Jung K. Lee, Jaehong J. Kim, and Jiwon Kim. Continual Learning with Deep Generative Replay. *Neural Information Processing Systems (NeurIPS)*, 2017. 1, 6, 8
- [32] Yao Shu, Man Zhu, Kun He, John Hopcroft, and Pan Zhou. Understanding Deep Representations through Random Weights. arXiv preprint arXiv:1704.00330, 2017. 8
- [33] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep Image Prior. *International Journal of Computer Vision* (*IJCV*), 128(7):1867–1888, 2020. 2
- [34] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv preprint arXiv: 1708.07747*, 2017. 2
- [35] Sergey Zagoruyko and Nikos Komodakis. Wide Residual Networks. British Machine Vision Conference (BMVC), 2016. 3
- [36] Barret Zoph and Quoc V. Le. Neural Architecture Search with Reinforcement Learning. *International Conference on Learning Representations (ICLR)*, 2017. 1, 2, 3