

Learning to Prune Deep Neural Networks via Reinforcement Learning

Manas Gupta

I2R, ASTAR Singapore

MANAS.GUPTA@I2R.A-STAR.EDU.SG

Siddharth Aravindan

National University of Singapore

SIDDHARTH.ARAVINDAN@COMP.NUS.EDU.SG

Aleksandra Kalisz[†]

I2R, ASTAR Singapore

KALISZ.OLA@GMAIL.COM

Vijay Chandrasekhar[†]

I2R, ASTAR Singapore

VIJAY@I2R.A-STAR.EDU.SG

Lin Jie

I2R, ASTAR Singapore

LIN-J@I2R.A-STAR.EDU.SG

Abstract

This paper proposes PuRL - a deep reinforcement learning (RL) based algorithm for pruning neural networks. Unlike current RL based model compression approaches where feedback is given only at the end of each episode to the agent, PuRL provides rewards at every pruning step. This enables PuRL to achieve sparsity and accuracy comparable to current state-of-the-art methods, while having a much shorter training cycle. PuRL achieves more than 80% sparsity on the ResNet-50 model while retaining a Top-1 accuracy of 75.37% on the ImageNet dataset. Through our experiments we show that PuRL is also able to sparsify already efficient architectures like MobileNet-V2. In addition to performance characterisation experiments, we also provide a discussion and analysis of the various RL design choices that went into the tuning of the Markov Decision Process underlying PuRL. Lastly, we point out that PuRL is simple to use and can be easily adapted for various architectures.

1. Introduction

Neural network efficiency is important for specific applications, e.g., deployment on edge devices and climate considerations Strubell et al. (2019). Weight pruning has emerged as a viable solution methodology for model compression Han et al. (2016), but pruning weights effectively remains a difficult task — the search space of pruning actions is large, and over-pruning weights (or pruning them the wrong way) leads to deficient models Frankle et al. (2019); Deng et al. (2009).

In this work, we approach the pruning problem from a decision-making perspective, and propose to automate the weight pruning process via reinforcement learning (RL). RL provides a principled and structured framework for network pruning, yet has been under-explored. There appears to be only one existing RL-based pruning method, namely AutoML for Model Compression (AMC) He et al. (2018). Here, we build upon AMC and contribute an improved framework: Pruning using Reinforcement Learning (PuRL).

Compared to AMC, PuRL rests on a different Markov Decision Process (MDP) for pruning. One key aspect of our model is the provision of “dense rewards” — rather than rely on the “sparse”

[†]For contributions made in 2019

rewards (given only at the end of each episode), we shape the reward function to provide reward feedback at each step of the pruning process. This results in a far shorter training cycle and decreases the number of training episodes required by as much as 85%. The remaining design changes are informed by ablation-style experiments; we discuss these changes in detail and elucidate the trade-offs of different MDP configurations.

2. Related Work

Various techniques have been proposed for compressing neural networks Cheng et al. (2017). Pruning comes out as a general approach not having restrictions in terms of the tasks it is applicable to. However, pruning too has a large search space size and hence, traditionally, human expertise has been relied upon to do pruning. But with the advent of new search techniques like deep reinforcement learning, we can now automate the process of pruning. In Runtime Neural Pruning, Lin et al. (2017) demonstrate one early approach for using RL to do pruning. They use RL to do a sub-network selection during inference. Thus, they actually don't really prune the network, but select a sub-network to do inference. He et al. (2018) demonstrate the first use of RL for pruning. However, they only reward the agent at the end of an episode (sparse rewards) and don't give it any reinforcement at each step in the episode. This slows down the learning process of the RL agent.

We improve upon this by creating a novel training procedure that rewards the agent at each step of the episode (dense rewards) and achieves faster convergence. Our approach is also general in nature and can be easily adapted for different architectures. We compare and report our performance with regards to AMC and other state-of-the-art pruning algorithms on the ImageNet dataset.

3. Pruning using Reinforcement Learning (PuRL)

This section details PuRL, our reinforcement learning method for network pruning. We formalize network pruning as an MDP; we specify the constituent elements, along with intuitions underlying their design.

3.1 Pruning as a Markov Decision Problem

We model the task of pruning a neural network as a Markov Decision Problem (MDP). We formulate and construct each element of the MDP tuple i.e. $\langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \gamma \rangle$ to enable us to use RL for pruning. In the below subsections, we elaborate on each of the tuple elements.

3.1.1 STATE REPRESENTATION

We represent the network state s through a tuple of features. We experiment with two different kind of representations. The first is a simple representation scheme consisting of three features, $s = \langle l, a, p \rangle$, where l is the index of the layer being pruned, a is the current accuracy achieved on the test set (after retraining) and p corresponds to the proportion of weights pruned thus far. The attributes serve as indicators of the network state. The second representation is a higher dimensional representation aimed at capturing more granular information on the state of the network. It is formulated as, $s = \langle a_1, p_1, a_2, p_2, \dots, a_n, p_n \rangle$, where a_i is the test accuracy after pruning layer i and doing retraining and p_i is the sparsification percentage of layer i . s is a tuple of zeros at the start of

every episode. Each tuple element is updated progressively as layer i is pruned. We report results from both state representations in the ablation experiments.

3.1.2 ACTION SPACE

The action space consists of actions where each action corresponds to an α value which decides the amount of pruning. We use a magnitude threshold derived from the standard deviation of the weights of a layer as our pruning criteria. We prune all weights smaller than this threshold in absolute magnitude. The set of weights that get pruned when an $\text{Action}_i(\alpha)$ is taken for layer i are given by Equation 1.

$$\text{Weights Pruned}_i(\alpha) = \{w \mid |w| < \alpha\sigma(w_i)\} \quad (1)$$

where $\sigma(w_i)$ is the standard deviation of weights in layer i . To further reduce the search complexity, we also experiment with increasing the step size of our actions from 0.1 to 0.2, to sample less actions but still achieve same target pruning rates. We report results in the ablation experiments. We use the same action space for all layers in the network i.e. $\alpha \in \{0.0, 0.1, 0.2, \dots, 2.2\}$. This is in contrast to current approaches like AMC and State of Sparsity which set a different pruning range for initial layers, in order to prune them by a lesser amount. Our approach is hence more general in this aspect.

3.1.3 REWARD FUNCTION

Since, the RL agent learns the optimal sparsification policy based on the objective of maximization of total reward per episode, reward shaping helps in faster convergence Ng et al. (1999). We formulate the total reward to be an accumulation of sub-rewards depending upon test accuracy and sparsification achieved. The reward function corresponding to a state s is given in Equation 2. Here, $A(s)$ and $P(s)$ denote the test accuracy and sparsity at state s , T_A and T_P denote the desired accuracy and sparsity target set by the user, and β corresponds to a fixed scaling factor of 5.

$$R(s) = -\beta(\max(1 - A(s)/T_A, 0) + \max(1 - P(s)/T_P, 0)) \quad (2)$$

The reward design ensures that the agent jointly optimises for the desired sparsity and accuracy.

3.2 The PuRL Algorithm

We design the PuRL algorithm, to be fast and efficient, when solving the above-mentioned MDP. The first aspect of this is the choice of a good RL agent. The second and more important consideration is the rewards scheme i.e. sparse vs. dense rewards. We elaborate on each of these in the below subsections.

3.2.1 CHOICE OF RL AGENT

To solve the MDP we choose amongst various available RL algorithms. Our primary focus is on sample efficiency and accuracy. Deep Q-Network (DQN) Mnih et al. (2013), a form of Q-learning, does a very fast exploration, however, it is not very stable. Through careful design of our reward structure, we make DQN stable and hence, utilise it for doing pruning.

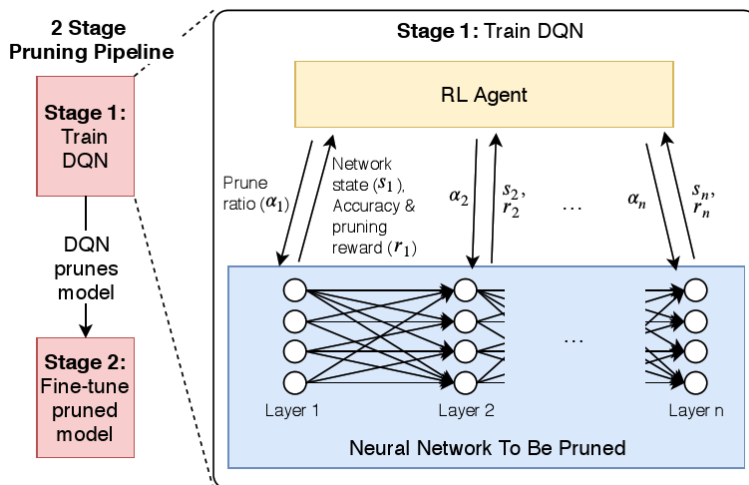


Figure 1: A high level view of the PuRL algorithm with dense rewards. PuRL assigns a unique compression ratio α to each layer. It then gets feedback on the test accuracy and sparsity achieved, after pruning that layer. This is in contrast to current approaches which only give feedback at the end of pruning the whole network. As a result, PuRL learns the optimal sparsity policy 85% faster than current approaches

3.2.2 MAKING RL FAST: DENSE REWARDS

The pruning procedure consists of pruning the weights in a layer based on their magnitude first. The remaining weights are then retrained to get back the accuracy. Retraining is an important aspect of this process.

By setting a different α for each layer, we try to prune away the maximum redundancy specific to each layer. As mentioned in section 2, one way this has been done is to 1) assign alphas to each layer 2) prune each layer and 3) retrain the pruned network in the end. While this method works, as shown by He et al. (2018), it may not be the fastest since it does not directly ascribe accuracy to the α of each layer. In other words, since retraining is only conducted after pruning all the layers and not after each layer (sparse rewards), the network cannot directly infer how accuracy is linked to α of each layer. This might elongate the training period since more samples are required to deduce effect of α of each layer to the network’s final accuracy.

We try to mitigate this by giving rewards after pruning each layer as opposed to giving them at the end of the episode (dense rewards). We retrain the network after pruning each layer to get a test accuracy value. We do retraining by using only a small training set of 1000 images in the case of ImageNet experiments, so as not to add training overhead. We measure accuracy after each layer is pruned and pass it to the agent through the reward and state embedding. As mentioned in section 4, this method of giving dense rewards helps achieve convergence much faster compared to giving sparse rewards. A high level view of the PuRL algorithm is presented in Figure 1.

4. Experiments & Analysis

In this section, we describe computational experiments comparing PuRL to ablated variants, as well as baseline and state-of-the-art methods. Our primary goal was to clarify the effect of different design choices (described in section 4.1) to the pruning performance. Secondly, we demonstrate

that PuRL achieves comparable results to state-of-the-art while using a 85% shorter RL training cycle by testing it on CIFAR-100 and ImageNet datasets and different architectures like ResNet-50, MobileNet-V2 and WideResNet-28-10 (refer to section 4.2 and B.3). Lastly, we showcase the generalization ability of PuRL by using the exact same settings to prune all architectures on ImageNet.

4.1 Understanding the RL Design Space

We conduct a series of ablation experiments to understand what components of the RL design space help make a good RL agent. The choices that give superior result over the baseline are then eventually used. Due to space constraints, we elaborate on some choices in Appendix B.1. We experiment on the ResNet-50 architecture trained on the CIFAR-10 dataset. We set a target sparsity of 60% and target accuracy of 95% for our agent (via the reward function), in all experiments.

Experiment	State Space	Action Size	Reward Space				Accuracy%	Sparsity%
			Prune Penalty	Acc. Penalty	Acc. Upside	Cubic Upside		
Sparse Rewards	3	0.1	✓	✓			68.0±13.9	66.1±12.7
Dense Rewards (Baseline)	3	0.1	✓	✓			91.8±2.0	70.1±2.3
Baseline + Magnitude Target	3	0.1	✓	✓			86.6±3.9	60.3±0.8
Baseline + Reward 2	3	0.1	✓	✓	✓		91.4±0.7	68.3±2.2
Baseline + Reward 3	3	0.1	✓	✓	✓	✓	90.7±0.2	69.4±2.9
Baseline + Action 2	3	0.2	✓	✓			92.2±0.5	70.6±0.6
Baseline + State 2	108	0.1	✓	✓			78.9±8.6	72.2±4.5

Table 1: Ablation results on perturbing State, Action and Reward spaces for the PuRL algorithm on the CIFAR-10 dataset. Error denotes standard error as measured on 3 trials. Dense rewards outperform sparse rewards by a huge margin on accuracy (rows 1 & 2). Stepping the action space by 0.2 (row 6) leads to a Pareto dominant solution over the baseline (row 2)

4.1.1 ARE DENSE REWARDS BETTER THAN SPARSE REWARDS?

We compare sparse rewards i.e. rewards given to the agent only at the end of the episode and dense rewards i.e. rewards at each step of the episode, and try to answer which is better. Referring to Table 1, we compare sparse rewards (row 1) to dense rewards (row 2). Our dense rewards approach outperforms the sparse rewards by a huge margin, 4% on sparsity and 24% on accuracy. Dense rewards help the agent learn much faster by guiding the agent at each step instead of only at the end of the episode. We then use dense rewards as our baseline to conduct all further ablations.

4.1.2 ARE FEWER ACTIONS BETTER?

In the experiment using Action 2 (row 6), we modify the action space to cover the same breadth of actions but have lesser number of actions. So the range remains the same but the step size between the actions increases. So instead of the actions being (0.0, 0.1, .. , 2.2), we now have (0.0, 0.2, .. , 2.2). We see that this experiment Pareto dominates the baseline i.e. it exceeds the baseline in

Method	ResNet-50				
	Sparsity	Starting Acc.	Pruned Acc.	RL Episodes	Fine-tuning Epochs
State of Sparsity	80%	76.69%	76.52%	NA	153
AMC	80%	76.13%	76.11%	400	120
PuRL	80.27%	76.13%	75.37%	55	120

Table 2: We compare PuRL against the global state-of-the-art pruning results, not just for RL but for all pruning algorithms, and report the Top-1 accuracy performance on ImageNet. PuRL uses 85% less RL episodes than AMC.

both sparsity and accuracy. This is likely because with less number of actions to try, the agent is able to sample each action more and gain better knowledge of each action vis-a-vis the resultant performance metrics. Hence, it picks out better actions i.e. learns a better pruning policy given a particular layer in the network.

4.2 Generalization across ImageNet

To evaluate the performance of our agent on large scale tasks we experiment with the ImageNet dataset. We prune a pretrained ResNet-50 model using an iterative pruning scheme as mentioned in Han et al. (2015) to preserve accuracy by providing gradual pruning targets for the network. We compare our performance with the state-of-the-art pruning algorithms AMC: AutoML for Model Compression He et al. (2018) and State of Sparsity Gale et al. (2019). We prune more than 80% and achieve comparable accuracy to state-of-the-art methods (see Table 2 for full results).

Furthermore, PuRL finishes each RL training cycle in just 55 episodes, compared to 400 episodes required by AMC, due to the dense reward training procedure. We also conduct experiments on other state-of-the-art efficient architectures like MobileNet-V2 Sandler et al. (2018) and EfficientNet-B2 Tan and Le (2019). Referring to supplementary document, PuRL achieves more than 1.5x sparsity compared to AMC without much loss in accuracy. At the same time, PuRL achieves this performance on MobileNet-V2 without any changes in the underlying hyper-parameters compared to ResNet-50. Thus, PuRL can be easily used across architectures without the requirement of modifying the underlying MDP.

5. Conclusion

In this paper, we present PuRL - a fully autonomous RL algorithm for doing large scale compression of neural networks. By improving the rewards structure compared to current approaches, we shorten the training cycle of the RL agent from 400 to 55 episodes. We further do a detailed set of ablation experiments to determine the impact of each MDP component to the final sparsity and accuracy achieved by the agent. We achieve results comparable to current state-of-the-art pruning algorithms on the ImageNet dataset, sparsifying a ResNet-50 model by more than 80% and achieving a Top-1 accuracy of 75.37%. We also benchmark PuRL on other architectures like WideResNet-28-10 including already efficient architectures like MobileNet-V2 and EfficientNet-B2. Lastly, our algorithm is simple to adapt to different neural network architectures and can be used for pruning without a search for each MDP component.

Acknowledgments

This research is supported by the Agency for Science, Technology and Research (A*STAR) under its AME Programmatic Funds (Project No.A1892b0026 and No.A19E3b0099).

References

- Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks, 2017.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Stabilizing the lottery ticket hypothesis, 2019.
- Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*, 2019.
- Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pages 1135–1143, 2015.
- Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *International Conference on Learning Representations*, 2016.
- Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784–800, 2018.
- Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-100 (canadian institute for advanced research). URL <http://www.cs.toronto.edu/~kriz/cifar.html>.
- Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2181–2191. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6813-runtime-neural-pruning.pdf>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *NIPS Deep Learning Workshop*, 2013.
- Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, pages 278–287, 1999.
- Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018.

Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in nlp. *57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.

Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning*, 2019.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Edwin R. Hancock Richard C. Wilson and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, September 2016. ISBN 1-901725-59-6. doi: 10.5244/C.30.87. URL <https://dx.doi.org/10.5244/C.30.87>.

Appendix A. PuRL Algorithm

Algorithm 1 describes a DQN procedure which learns to select the sparsity threshold for each layer in the model. At the start of each episode, the original model with pre-trained weights is loaded. The agent then takes an action and the layer is pruned by that amount. The model is then retrained for one epoch on a small subset of training data (1000 images out of the 1.2million images in ImageNet). After that, the validation accuracy is calculated, to effect the state transition and the reward is then calculated using the validation accuracy and the pruning percentage. The training is done for $max_episodes$ episodes. Once, training is completed, the model is pruned using the trained agent and then fine-tuned on the full ImageNet dataset.

Algorithm 1 The PuRL Algorithm

```

1: Stage 1: Train DQN agent
2:  $episodes \leftarrow 0$ 
3: while  $episodes \leq max\_episodes$  do
4:    $model \leftarrow$  load original model
5:   for each layer  $t$  in the  $model$  do
6:      $a_t \leftarrow$  Sample action from DQN agent
7:     Prune layer  $t$  using  $a_t$ 
8:     Retrain  $model$  on small subset of data for 1 epoch
9:     Calculate reward  $r_t$  and new state  $s_t$  based on resultant sparsity and accuracy achieved
10:    Return  $r_t$  and  $s_t$  to agent
11:    $episodes \leftarrow episodes + 1$ 
12:
13: Stage 2: Prune and fine-tune model
14:  $model \leftarrow$  load original model
15: Prune  $model$  using the trained DQN agent (averaged over 5 episodes)
16: Fine-tune  $model$ 
17: return  $model$ 

```

Appendix B. Experimental Results

B.1 Understanding the RL Design Space

B.1.1 IS ABSOLUTE MAGNITUDE BETTER THAN STANDARD DEVIATION?

Referring to the Magnitude Target based ablation Table 1 (row 3), we compare absolute magnitude based pruning to standard deviation based pruning (Section 3). For absolute magnitude, we set a sparsity target for a layer and then remove all small weights until we hit the desired sparsity level. As we observe, both the sparsity and accuracy results are lower in this case as compared to our baseline experiment (dense rewards). Thus, standard deviation based pruning is better.

B.1.2 DOES REWARD SHAPING HELP?

In the experiments using Reward 2 and Reward 3 (Table 1), we investigate if reward shaping can help the agent achieve higher accuracy. For Reward 2, we allow the agent to receive positive rewards if it

Experiment		State Space	Action Size	Reward Space			Accuracy	Sparsity
				Prune Penalty	Acc. Penalty	Acc. Up-side Cubic Up-side		
Low State	Dimensional	3	0.2	✓	✓		47.6 ± 2.3	87.4 ± 1.3
Higher State	Dimensional	108	0.2	✓	✓		51.0 ± 0.3	82.0 ± 2.0

Table 3: Follow-up experiment on perturbing the State space on the ImageNet dataset. Error denotes standard error as measured on 3 trials. The higher dimensional state space (108 dimensions) performs better than the simple low dimensional state space (3 dimensions)

surpasses the given accuracy target (Equation 3). This is in contrast to the baseline reward function in which there is a cap on the maximum reward that the agent can achieve i.e. zero.

$$R_2 = -S((A/T_A - 1) + \max(1 - PP/T_{PP}, 0)) \quad (3)$$

For Reward 3, we build on Reward 2 and give a cubic reward to the agent. The agent now sees cubic growth in positive reinforcement as it approaches and surpasses the accuracy target (Equation 4). Hence, by taking the same step size towards accuracy improvement as compared to Reward 2, the agent now gets rewarded more for it.

$$R_3 = -S(((A/T_A)^3 - 1) + \max(1 - PP/T_{PP}, 0)) \quad (4)$$

The performance of both these functions is close to the baseline (dense rewards), but the baseline still outperforms them. The added complexity of these functions might require the agent to sample more steps to learn them well. Hence, given a tight training budget, the baseline reward function performs well.

B.1.3 IS MORE INFORMATION BETTER FOR THE AGENT?

In the last experiment with State 2 (Table 1), we vary the state space and make it 108 dimensional instead of 3 dimensional. The idea here is to give the agent more information on the state representation (See Section 3 for details). We see that in this experiment, the agent achieves less accuracy than the baseline however, prunes more than it. Hence, none of the experiments Pareto dominate each other and its inconclusive to determine which one is better. To get more evidence on this, we carry out a further ablation on the ImageNet dataset. Referring to Table 3, we see that the 108 dimensional state outperforms the 3 dimensional state. Hence, more information is indeed better and we use this feature in the final configuration.

B.2 Scaling PuRL to CIFAR100

We first experiment with PuRL on the WideResNet-28-10 architecture Zagoruyko and Komodakis (2016) on the CIFAR-100 Krizhevsky et al. dataset. We compare it to a uniform pruning baseline where every layer is pruned by the same amount to achieve a target sparsity of 93.5%. PuRL outperforms the baseline in Table 4 on both the sparsity and final accuracy.

Method	WideResNet-28-10			
	Sparsity	Top-1 Pre-Pruning	Acc.	Top-1 Post-Pruning
Baseline	93.50%	82.63%		72.42%
PuRL	93.90%	82.63%		80.63%

Table 4: Comparison of the PuRL algorithm to a uniform pruning baseline on the WideResNet-28-10 architecture on CIFAR-100 dataset. PuRL beats the baseline by a huge margin

Method	MobileNet-V2				
	Sparsity	Flops reduction	Top-1 Pre Pruning	Acc.	Top-1 Post Pruning
AMC	Not reported	30%	71.8%		70.8%
PuRL	43.3%	47.9%	71.9%		69.8%

Table 5: Comparison of PuRL to AMC for the MobileNet-V2 architecture.

Method	EfficientNet-B2			
	Sparsity	Top-1 Pre Pruning	Acc.	Top-1 Post Pruning
Baseline	59.0%	79.8%		68.9%
PuRL	59.5%	79.8%		74.5%

Table 6: Comparison of PuRL to uniform pruning baseline on the state-of-the-art EfficientNet-B2 architecture on the ImageNet dataset. PuRL outperforms the baseline on both the sparsity and accuracy

B.3 Generalization across ImageNet

We also conduct experiments on other state-of-the-art efficient architectures on ImageNet to see whether our pruning algorithm can make these architectures even more sparse. We experiment on MobileNet-V2 Sandler et al. (2018) and EfficientNet-B2 Tan and Le (2019). Referring to Table 5, PuRL achieves more than 1.5x sparsity compared to AMC without much loss in accuracy.

At the same time, PuRL achieves this performance on MobileNet-V2 without any changes in the underlying hyper-parameters compared to ResNet-50. Thus, PuRL can be easily used across architectures without the requirement of modifying the underlying MDP. For EfficientNet-B2, Table 6, we compare PuRL to a uniform pruning baseline. PuRL outperforms the baseline on both sparsity and final accuracy, achieving an accuracy improvement of more than 5%. In this case as well, we set the exact same hyper-parameters and MDP setting as that of ResNet-50 and MobileNet-V2. However, since Efficient-B2 is very deep, having 116 layers compared to 54 in ResNet-50, we do early-stopping of the RL episode, to make the training even faster. We stop the episode if the test accuracy drops to less than 0.1% and move on to the next episode.