Is Active Learning Always Beneficial? (Student Abstract)

Anna Kravchenko, Rhodri Cusack
Trinity College Dublin
College Green, Dublin, Ireland
kravchea@tcd.ie, cusackrh@tcd.ie

Abstract

This study highlights the limitations of automated curriculum learning, which may not be a viable strategy for tasks in which the benefits of the chosen curriculum are not apparent until much later. Using a simple convolutional network and a two-task training regime, we show that in some cases a network is not able to derive an optimal learning strategy using only the data available during a single training run.

In some cases a network is best trained for a task by following a curriculum, in which simpler concepts are learned before more complex ones. This curriculum could be handcrafted by the engineer or optimised like other hyperparameters. However, an attractive alternative is that the network by monitoring its learning, could choose its own curriculum. In machine learning, this is called automated curriculum learning (Bengio et al. 2009); in human learning and developmental robotics it is active or curiosity-driven learning. Metrics, including the performance on the current task, the gradient of learning, and information theory measures, have been used to help the learner to determine when to switch between task stages. Active learning is appealing, as it makes design simpler and could allow networks to adapt to new datasets. It is also viewed as one of the methods to prevent catastrophic forgetting during continual learning (Parisi et al. 2019). However, is it always possible? We hypothesized that in some cases learners may not be able to choose the optimal curriculum if the benefits will only be apparent with hindsight.

To test this we created a network and task for which curriculum learning was important. The curriculum of our network consisted of two tasks, a simple supporting task of recognizing single-digit numbers (Task 1) and a more complex main task (Task 2), requiring the network to discriminate whether the sum of two digits in the picture was odd or even. The training datasets were based on the MNIST dataset of handwritten digits. In both tasks the network was trained in a supervised way.

We evaluated a pre-training curriculum, where Task 1 always preceded Task 2, with varying degrees of training on each. Task 2 required the network to recognize two digits simultaneously and analyse higher-level abstract information.

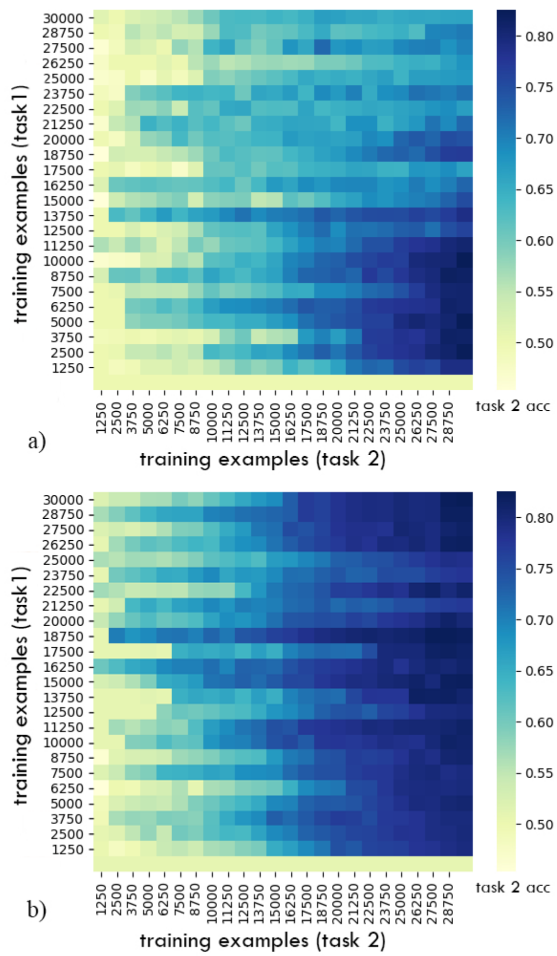

Figure 1: Performance on Task 2 with respect to training on each of the individual tasks. Both for sequential continual (a) and regular learning (b) performance on Task 2 was strongly dependent on the amount of pre-training on Task 1. In both cases without pre-training on Task 1 the network never learned Task 2. For continual learning too much Task 1 pre-training could also be detrimental (a). We attribute this effect to catastrophic forgetting as it does not appear during regular learning.

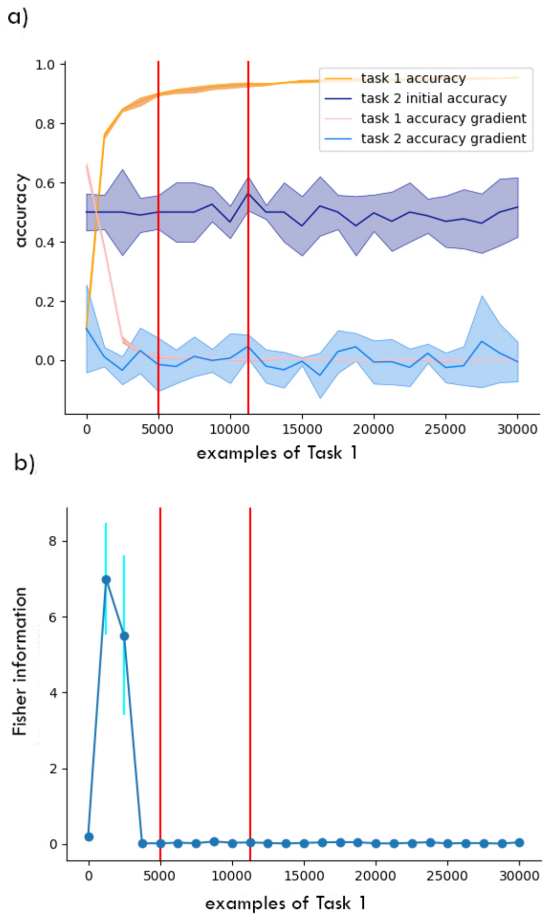

Figure 2: a) Mean of 8 runs and standard deviation bands for four potential switching metrics during continual learning. Neither auxiliary nor primary task, nor their gradients provided a signal that could guide active learning. b) The mean of the Fisher Information of weights with respect to the loss function for Task 1 across 8 runs. The bars show one standard deviation. On both figures red vertical lines indicate the optimal switching window.

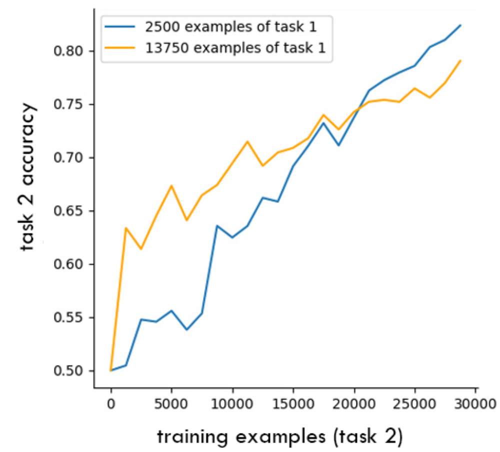

Figure 3: a) During continual learning too much pre-training on Task 1 was detrimental. Although in the early stages of Task 2 learning it did improve performance, after 20,000 training examples it limited performance

Because the error signal from Task 2 was uninformative with regards to the presence of any single digit, we predicted that pre-training on Task 1 would be a valuable stage towards optimizing performance on Task 2.

Task 2 turned out to be unlearnable without initial pre-training on task 1 (see Fig. 1). By exhaustively evaluating many training regimes, we also found that for continual learning of Task 1 an intermediate quantity of supporting task training was optimal. We then investigated under what conditions the network itself could determine the optimal moment for switching between tasks using only the data available during a single training run. As can be seen from Fig. 2 we found that neither main nor supporting task performance or their gradients, nor Fisher information of weights provided a signal that could guide active learning. These results provide no evidence that the optimal learning strategy indeed can be derived from information available to the network.

This highlights the limitations of active learning, which may not be a viable strategy for some tasks in which the benefits of the chosen curriculum are not apparent until much later (Fig. 3).

This is supported by the effectiveness of scaffolding approach in pedagogy, in which teachers support learners to achieve the next level by challenging them to the optimal degree, keeping them in their ”zone of proximal development” (Wood, Bruner, and Ross 1976). Building upon these lessons from human learning, it seems likely that ANNs may too benefit from a combination of fixed curriculum and active learning.

Acknowledgments

This work was supported by the European Research Council (ERC Advanced Grant ERC-2017-ADG, FOUNDCOG, 787981).

References

Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML ’09, 41–48. ISBN 9781605585161. doi:10.1145/1553374.1553380.

Parisi, G. I.; Kemker, R.; Part, J. L.; Kanan, C.; and Wermter, S. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks* 113: 54–71. doi:10.1016/j.neunet.2019.01.012. URL https://doi.org/10.1016/j.neunet.2019.01.012.

Wood, D.; Bruner, J. S.; and Ross, G. 1976. The Role of Tutoring in Problem Solving*. *Journal of Child Psychology and Psychiatry* 17(2): 89–100. ISSN 1469-7610. doi:10.1111/j.1469-7610.1976.tb00381.x.