Flexible AI for Real-World Environments

T. Achler

Optimizing Mind Inc.
624 University Avenue, Palo Alto, CA 94301 USA, Achler@OptimizingMind.com

ABSTRACT

AI remains less capable compared to humans at quickly accumulating knowledge without forgetting what it previously learned. Everything that might happen in the work environment must be in the training set and rehearsed in cumbersome ways. AI based robots are limited where and when they can be applied: in rigidly controlled environments, with limited interactions with humans, autonomy and high costs of setup and retraining. Moreover even with severe limits, many hazards remain which safety organizations such as OSHA must carefully monitor.

This problem may be addressed by mimicking predominant feedback connections found in the brain as a primary form of computation; allowing for learning flexibility without cumbersome rehearsal.

Transfer learning is the current leading method for fast learning thus we Compare this approach with transfer learning demonstrating easier learning (faster learning with less data) and modularity (learning without forgetting old data, eliminating cumbersome rehearsal). The improvements are quite significant for the complex dataset tested. This solution is essential to enable computer vision and automation engineer customers to correct their applications in the field (without sending data back to retrain the whole network) reducing machine and customer downtime/disruption and increasing productivity. Overcoming this limitation is a key towards interacting with machines in a natural way as is depicted in science fiction.

Keywords: flexible ai, robotic rehearsal, learning on the fly, modular learning, easier learning, catastrophic forgetting

1 PROBLEM

Old movies, cartoons and science fiction depict robots interacting naturally in our environment, however this has not panned out in reality. Although very capable robots have been created that dance, perform acrobatics, and gallop, they still do not interact naturally. The learning algorithms that we have today may be powerful but they are deceivingly inflexible. This limits their application in real life environments, safety, and even sometimes misleads experts: when are truly self-driving cars due?

The limit is no longer how accurate the algorithms are but how flexible they are. Current algorithms require big data and rehearsing over and over again limiting their use in natural environments. Overcoming this limitation has challenged large and small companies, startups and large government grants in the Billions of dollars.

Training is always incomplete as issues arise in real natural environments. Currently everything that may be in the real environment must be predicted during learning. Predicting the real environment is tricky. Moreover if something is missing, new data must be captured and weights must be rerun and retrained from scratch on the training computer. Typically the machine in the recognition environment is smaller and less powerful because it requires being integrated with the rest of the robot, industrial applications, and consumes less space and power. Current training computers are much larger and more expensive. Thus training for even semi-realistic environments is difficult and costly with huge downtimes.

1.1 Unrealistic Rehearsal

These difficulties ultimately exist because current algorithms require rehearsal that is unnatural and unintuitive. This rehearsal is akin to a casino dealer shuffling training examples of everything we previously learned for anything new we want to learn. Imagine spending a summer in Hawaii, and by not rehearsing other environments (e.g. winter scenes, desert scenes, etc) forgetting how to recognize other scenes. Subsequently, current rehearsal solutions are very impractical and present severe limits to artificial intelligence solutions [1]. The restrictive rehearsal paradigm is an attempt to solve an underlying problem in the literature called catastrophic interference or forgetting [2-5] which it has not been solved satisfactorily [6]. Catastrophic forgetting describes how information fades away during learning. Rehearsal is a stopgap measure to manage this fading by periodically "reminding" the system with randomly selected patterns (because if the patterns are not random something else may be forgotten). This measure strongly limits flexibility. The core issue is the ability of the learning algorithm and network to be modular: to have the ability to add new information to one labeled class without affecting the information of other classes.

2 GUIDANCE BASED ON BRAIN STUDY

Guidance from brain studies of a system that is both powerful and flexible was used to develop an algorithm to overcome these limitations. The underlying method was developed based on the founder's Computational Neuroscience research in incorporating predominant top down feedback connections found throughout the brain as a primary form of computation. Current algorithms rarely use these types of connections because they do not allow for backpropagation, a key mechanism for current algorithms. Briefly: instead of learning distributions from the training set, that are determined by rehearsal, top-down feedback (from outputs modulating inputs) estimate distributions during recognition when the object(s) to be recognized are present. This process happens during recognition and allows faster, more flexible, and easier to comprehend learning during training without distributions. The Neural mechanisms and motivations are beyond the scope of this paper, additional information can be found in [7].

3 GOALS AND METHODS

The goal of this paper is to evaluate the performance of this algorithm and flexibility within deep learning vision with a perspective towards practical commercial applications.

3.1 Transfer Learning

To demonstrate flexibility we chose the best known paradigm for those who implement deep learning: transfer learning method.

We used the Tensorflow (TF) tutorial notebook of transfer learning [8], a very popular and accessible example that was specifically created to demonstrate easier learning by learning on top of a partially trained network. It is a state-of-the-art example for practical learning, so we perform the Optimizing Mind (OM) benchmarks and comparisons against it.

3.2 Defining Flexibility

There many several ways to demonstrate flexibility and we will show two:

- Easier learning:
  - Less time needed to learn
  - Less examples needed to achieve the same accuracy (reducing big data)
- Modularity:
  - Learn or update without having to rehearse old data

3.3 Evaluating Easier learning

To demonstrate easier learning we trained using the exact same pictures of cats and dogs as the transfer learning example and used the exact same evaluations. We had to however adjust the original notebook batch sizes (number of data examples before running validation tests) because the original obtained 32 examples per training instance and our solution completed learning with less than 32 examples (making our solution look instantaneous). Thus we reduced the batch size to one sample each. The validation set(s) was run after every batch in order to get the best granularity

3.4 Evaluating Modularity

To demonstrate modularity, the ability to add new information without affecting the old information or requiring rehearsal, we added a set of pictures of birds which we obtained from the animals 151 database [9]. To provide diverse samples we included 34 species of birds into our category of birds: (e.g. penguins, ostriches, hawks, hummingbirds, chickens, parrots, etc.).

Next we test flexibility based on modularity: the ability to separately train new classes without degradation of old information. This is essential in order to avoid the huge rehearsal costs that exist today.

As part of the modularity paradigm a new class node was added representing birds to the transfer learning layer of both OM and TF models, and all of this node's weights were set to zero. All of the weights of the previous class nodes from the example above (cats and dogs) were retained and remained unchanged.

To assure modularity (ability to separately train without rehearsal) only birds were subsequently trained without the other classes: cats or dogs. This is reflective of real world scenarios where all things previously learned would not be available for training.

3.5 Validation sets to show Modularity

To demonstrate how well old information was retained we tested the networks with two validation data sets. The validation set labeled OLD contained the exact same examples of cats and dogs used in the Easier Learning demonstration example.

To demonstrate how well new information was retained, a validation set labeled NEW was created that contained validation examples of all three categories: cats, dogs and birds. The validation set had the same number of each.

4 RESULTS

Tests were run 7 times with different batch sizes (1 & 32). The results were similar regardless of batchsize however batchsize=1 is displayed for better granularity

4.1 Flexibility through Ease of Learning:

The results show that using less than 10 examples of cats and dogs, the OM solution was able to achieve a similar accuracy which the original achieved after 1500 examples. This is a huge improvement in ease from the perspectives of both speed (more than 100x faster) and less data (requiring less than 1% of the data). These improvements are essential for learning on the machine that is in the field; on the smaller computers within the real recognition environment they are exposed to.

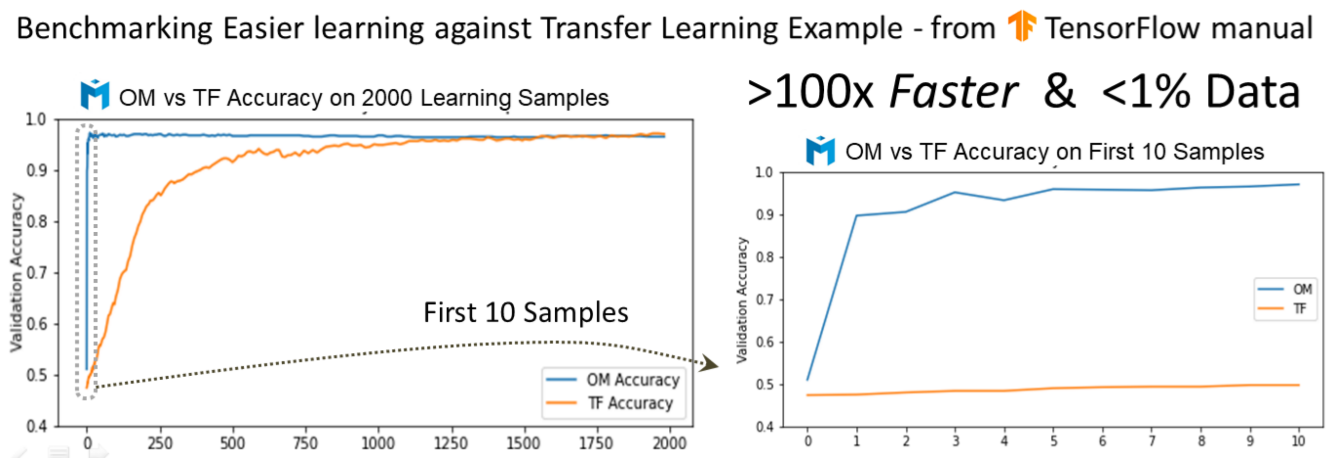

Figure 1: Comparing the validation accuracy performance of transfer learning layer in the tensorflow notebook (orange trace labeled TF) with the Optimizing Mind transfer learning layer (blue trace labeled OM). The batch size is one (one sample per batch) and the accuracy increased for TF gradually throughout the 2000 training samples. OM accuracy increased too quickly (within the first 8 samples) to display so a graph depicting the first 10 learning samples is shown on the right. The equivalent OM accuracy at batch 10 was achieved approximately in batch 1500 of TF. This represents a speed of over 100x and decrease to less than 1% in data required to achieve the same accuracy.

4.2 Flexibility through Modularity

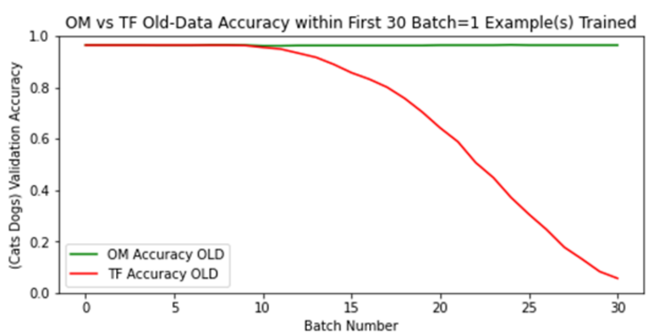

Figure 2: Performance on the OLD validation set from example 4.1, figure 1 (using cats dogs) was evaluated while the third class bird was trained. TF performance (red trace) decreased as more examples of birds were presented without any dogs or cats. OM performance (green trace) remained unchanged throughout the training of birds, demonstrating modular learning without forgetting.

Initial accuracy of both transfer learning Tensorflow (TF) and Optimizing Mind (OM) layers were high using the OLD validation set because old training information was retained (~97% for both) and birds were not learned yet. As training of birds progressed, OM accuracy remained the same using the OLD validation set because previously trained information (of Cats and Dogs) did not degrade with newly trained information (of birds), demonstrating modularity. The same was not true with TF: accuracy decreased to zero as old information of cats and dogs faded. The OM trace appears to address the problem of catastrophic forgetting the TF trace displayed.

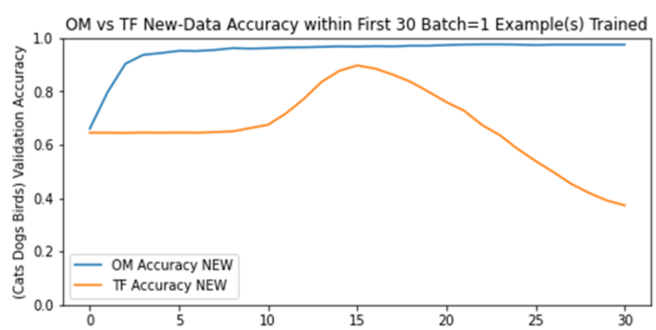

Figure 3: Performance on the NEW validation set including cats, dogs and birds. OM performance (blue trace) increased as birds were learned and by 7-10 examples reached peak performance on all three classes even though it was continually trained on only one class. TF performance initially increased as the network learned birds, however at the same time lost the ability to recognize cats and dogs and degraded until it could only recognize birds.

Initially accuracy of both transfer learning Tensorflow (TF) and Optimizing Mind (OM) layers began at ~66% using the NEW validation set because both models were initially recognizing the cats and dogs but not birds.

OM quickly learned birds while retaining cats and dogs. This is reflected by the NEW validation set score increasing from 66% to 98% total, within the first 5 samples.

TF started increasing accuracy, as birds were learned. However as birds were trained, the weights of cats and dogs degraded. This caused performance to decrease and with additional training, this degradation continued until only birds were recognized and the NEW validation set score settled on 33% (correctly recognizing only birds). Because the OM model retained knowledge of cats and dogs, and only changing its representation of birds, it displays modular learning.

5 SUMMARY & CONCLUSIONS

We showed increased flexibility in the OM transfer learning layer through increased speed (>100x) less data required (<1%) and modularity – classes that were not trained did not change or lose information. These properties are essential to allow AI to interact with the real world. These properties will help end users address their learning needs quicker and more efficiently.

The increased flexibility of this approach can help reduce costs, increase safety, and enable greater automation in more variable environments bringing more AI to more customers.

Our continuing development at Optimizing Mind will bring more brain-like flexibility capabilities. The brain allows interleaving logic and expectation within recognition. For example when looking for a specific bird e.g. hawk, birds not relevant to the task can be ignored (e.g. hummingbird etc.). Such control of recognition logic within the mechanisms of recognition is difficult to achieve with current algorithms. With modular mechanisms such manipulations are much easier and this work is in process at Optimizing Mind.

5.1 Connecting with Customers

This solution allows you to train and update your networks based on a transfer-learning paradigm even in your real world test environment. This helps relieve the difficult development cycle associated with AI, whenever an issue arises.

Our tools and libraries are designed to seamlessly integrate with TensorFlow and produce a TensorFlow layer that is indistinguishable from other TensorFlow layers. We also have integrations with other standard Machine Learning libraries. Thus we reduce any setup costs, NREs or adoption barriers for customers already using their own tools; making it simple and easy to boost applications with OM's technologies.

Our overarching business strategy is to first start in applications such as robots where Deep Learning AI is well used such as industrial automation. This includes applications where customers have been able to control the environment enough to adopt computer vision and machine learning methods into their successful business practices. Our goal is to reduce costs, downtimes, increase safety and expand the AI use into less controlled environments, to provide reduced setup and update issues and increased efficiency.

We are looking for partners who are actively using vision and would like to increase their flexibility, speed and modularity. Please contact us if this solution will improve your bottom line. Overcoming this limitation is a key to propel towards the ability to interact with machines in a natural way as is depicted in science fiction.

AKNOWLEDGEMENTS

This work is funded by NSF SBIR grant 2127085. I would also like to thank Efi Dror and Jingyi Zha for their invaluable help and support.

REFERENCES

[1] T Achler, 2021 Updating Your AI Capabilities for Life-Like AI; and Demonstrating difficulty with rehearsal (video) https://youtu.be/xYw0sg1OaAk

[2] JL McClelland, BL McNaughton, RC O'Reilly. "Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning & memory" Psychological Review 102:419-57 (1995)

[3] M McCloskey, NJ Cohen "Catastrophic interference in connectionist networks: The sequential learning problem" The Psychology of Learning and Motivation, 24, 109-165 (1989).

[4] RM French. Catastrophic Forgetting in Connectionist Networks. In Encyclopedia of Cognitive Science, volume 1, pages 431–435, London, 2003. Nature Publishing Group.

[5] R Coop, I Arel, Mitigation of catastrophic interference in neural networks using a fixed expansion layer. Midwest Symposium on Circuits and Systems. 726-729. (2012).

[6] Ole-Marius Moe-Helgesen and Håvard Stranden. "Catastrophic forgetting in neural networks", 2005 https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.622.7385&rep=rep1&type=pdf

[7] T Achler, 2016 Neural Phenomena Focus (video) https://youtu.be/9gTJorBeLi8

[8] TensorFlow Manual Transfer Learning Example https://www.tensorflow.org/tutorials/images/transfer_learning

[9] https://www.kaggle.com/datasets/sharansmenon/animals141