Going Beyond the Point Neuron:

Active Dendrites and Sparse Representations for Continual Learning

Karan Grewal, Jeremy Forest, Subutai Ahmad

Numenta, Inc.

Numenta

Correspondence to
`kgrewal@numenta.com`

Introduction

- The point neuron –used by most artificial neural networks (ANNs)– abstracts away many details of real biological neurons.

Point Neuron

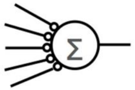

Real Pyramidal Neuron

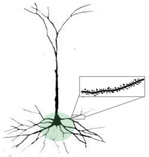

- ANNs suffer from **catastrophic forgetting** when learning tasks in sequence.
- We hypothesize that the physiology of real neurons is integral to how the brain learns continually without forgetting.
- Real neurons have **active dendrites** that act as independent pattern recognizers and can depolarize the cell, making it more likely to generate an action potential.

Model

- We augment ANNs with (1) **active dendrites** and (2) **sparse representations.** We measure resilience to catastrophic forgetting.

1. Active Dendrites

- Each artificial neuron has a set of dendritic segments (i.e., additional weights) which all receive a **context signal** and modulate the activity of the neuron.
- Each dendritic segment computes an activation value via a linear weighted sum of the context signal, and a single 'winning' activation is chosen to modulate the neuron.
- In our implementation, the absolute maximum activation d acts as a non-linear modifier by multiplying the neuron's output by $\sigma(d)$.

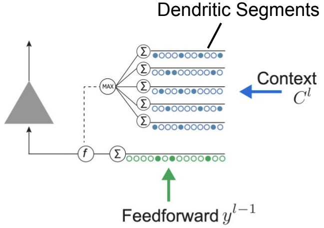

2. Sparse Representations

- We apply a k-**Winner Take All** non-linearity to all hidden layers, where only the top k-ranked activations keep their value.
- In effect, k winners sparsifies the activations of a layer, thus allowing for minimally-overlapping representations in each layer.
- Dendrite activity modulates the probability of a neuron winning.

Context Signal

- Dendrites on real neurons identify patterns that impact neural activity; our active dendrites receive a context signal that aims to serve as a task identifier.
- We use a **task prototype** by averaging all training examples in data space, and this serves as the context signal.

Training samples & prototype for task A

Training samples & prototype for task B

- At inference time, the closest prototype to a given test example is chosen as the context signal.

Results

On permutedMNIST, our dendrite networks achieve >80% accuracy when learning 50 tasks in sequence, far better than ANNs. This is competitive with a similar model, **XdG** [1], however task ID is given to XdG at test time. Our networks automatically infer the task information.

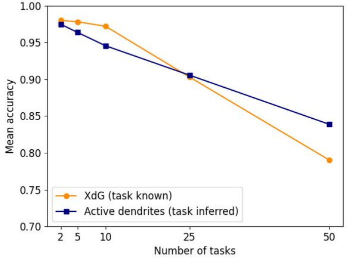

Sparse representations are important since accuracy on 50 tasks decreases as representations become more dense (*below*).

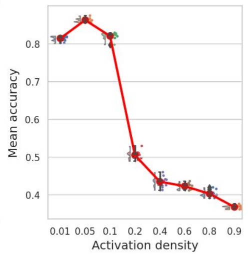

Active dendrites select subnetworks of neurons to learn each task, minimizing interference between tasks. We show activation frequency of 64 hidden units in a single layer across 10 tasks (*right*).

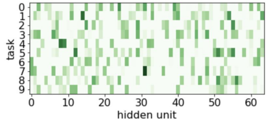

References

[1] Masse NY, Grant GD, Freedman DJ. Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization. *Proceedings of the National Academy of Sciences*. 2018, 115(44).