SLEEP PATTERNS IN THE FORWARD-FORWARD ALGORITHM

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper aims to explore the separation of the two forward passes in the Forward-Forward algorithm from a biological perspective in the context of sleep, following common patterns in the circadian rhythm of animals. We show the importance of the negative phase in the algorithm and how learning gradually collapses in its absence. Furthermore, we propose possible solutions to diminish the effects of sleep deprivation.

1 INTRODUCTION

The backpropagation algorithm (Rumelhart et al., 1986) is at the heart of deep learning and drives most of the research in the field, obtaining significant results in most practical applications. However, the main issue with backpropagation is its little resemblance to how the cerebral cortex learns. The network can only learn during the backward pass, which heavily restricts the possibility of separating the two parts[1] of the algorithm.

In contrast, the Forward-Forward algorithm (Hinton, 2022) proposes a new learning procedure that consists of two consecutive forward passes that aim to optimize opposite objectives. In the first forward pass, the network uses positive data and tries to update the weights to increase each layer's goodness over a certain threshold, making the algorithm greedy. During the second pass, the network is fed with negative data[2] and tries to minimize the goodness of each layer such that it ends below the same threshold. In a biological context, the negative pass represents REM sleep and is associated with a "reverse learning mechanism".

We analyze the properties and limitations of this approach by considering three different scenarios. Initially, we examine unihemispheric sleep (Manoach & Stickgold, 2016). Despite being uncommon for animal species to exhibit this sleeping pattern, the initial setup of the Forward-Forward algorithm is similar[3], providing a strong baseline for the rest of the experiments. Secondly, the more common monophasic sleep pattern is explored by alternating the two phases allocating equal time between them. Ultimately, we look at multiple scenarios that mimic sleep deprivation, trying to keep learning stable and efficient while sleeping for short periods.

2 SLEEP

No animal can survive without sleep. However, different types of sleep cycles emerged throughout evolution based on different environments and biological needs. In the Forward-Forward algorithm, we refer to sleep as a forward pass performed with negative data that has the same function as REM sleep (Crick & Mitchison, 1983). By performing a multitude of positive forward passes, the network gets overloaded with unwanted patterns. Thus, the model starts to "hallucinate", assigning the majority of data points a high goodness score. Once the model gets provided with negative data, the network suppresses the "parasitic" patterns of behavior, allowing it to learn a generalized design of the dataset.

[1] forward and backward pass

[2] explained in more depth in Appendix section A

[3] at each iteration, the two forward passes are calculated, and the gradients are updated with respect to both losses

SLEEP PATTERNS

As explained beforehand, there is a high coupling between the two passes in backpropagation, restricting the algorithm from performing efficient online learning. Due to the design of the FF algorithm, we achieve the necessary low-coupling in order to separate the awake-sleep phase while still learning, or unlearning, in either of the two forward passes.

The remaining question is how to minimize the sleep schedule of the algorithm without impairing its learning capabilities. To answer the question, we look at three common sleep patterns that emerged in nature.

3 EXPERIMENTS

To test the hypothesis presented in this paper, we perform the phase separation on two image classification datasets: MNIST and Fashion-MNIST (Xiao et al., 2017). For each dataset, we experiment with two scenarios: equal positive-negative phases and unbalanced with larger positive phases. We do not test the algorithm's behavior on the Unihemispheric sleep (Section B.1 of the Appendix) scenario as this was previously explored in Hinton's paper; this type of sleep allows for the best performance on the datasets.

We train seven different models as described in section Model Architecture of the Appendix. Each model has different awake and sleep phase durations. We start our experiments with 1 batch per phase and go up to 128. We train using two types of negative data, one generated by setting the wrong label on the train data point and the other by using a mask (Negative data of the Appendix).

Table 1: Showing the accuracy of models using equal awake and sleep phases. A clear performance advantage exists to using masks as negative data once the phases become larger.

Dataset	*Negative data*	1	2	4	8	16	32	64	128
MNIST	Wrong label	89%	88%	81%	63%	35%	11%	10%	9%
MNIST	Masks	88%	84%	85%	78%	74%	49%	23%	11%
Fashion-MNIST	Wrong label	73%	63%	59%	54%	20%	14%	10%	10%
Fashion-MNIST	Masks	56%	58%	53%	50%	45%	36%	30%	22%

Table 2: Showing the accuracy of models using unequal phases. The phase size of the positive data ranges from 1 to 16, and the negative phase is fixed at 1. The awake phase's learning rate is scaled down to have the same effect as the negative phase in the training procedure. Using masks as negative data offers a significant advantage on both datasets, allowing good inference quality even at large phase separations. The empty lines in the table represent models that experienced no learning.

Dataset	*Negative data*	1	2	4	8	16
MNIST	Wrong label	10%	9%	-	-	-
MNIST	Masks	88%	78%	75%	73%	10%
Fashion-MNIST	Wrong label	73%	56%	47%	14%	10%
Fashion-MNIST	Masks	56%	55%	54%	52%	49%

4 CONCLUSION

In this work, we present the relationship between different sleeping patterns and the Forward-Forward algorithm. We show the correlation between learning performance and choice of negative data, especially when the model is sleep deprived. Improving the quality of the negative data would allow the model to stay awake for longer periods, enhancing its practical usage. Future research in this area should be conducted in order to fully explore the sleep properties of the algorithm from both a practical and biological perspective.

URM STATEMENT

The authors acknowledge that at least one key author of this work meets the URM criteria of ICLR 2023 Tiny Papers Track.

REFERENCES

Francis Crick and Graeme Mitchison. The function of dream sleep. *Nature*, 304:111–114, 1983.

Geoffrey Hinton. The forward-forward algorithm: Some preliminary investigations. *arXiv preprint arXiv:2212.13345*, 2022.

Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020.

Dara S Manoach and Robert Stickgold. Sleep: keeping one eye open. *Current Biology*, 26(9): R360–R361, 2016.

David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

A NEGATIVE DATA

We employed the use of two different negative data generation strategies. Our first approach involved setting a wrong label on a train data point and using this as negative data. We later experimented with using a mixture of two data points combined using a predefined mask as presented in the Forward-Forward paper (Hinton, 2022). The latter approach allowed our model to perform significantly better than the former, allowing us to separate the awake and sleep phase more.

Figure 1: Negative data generation with masks on the Fashion-MNIST dataset. The image in the middle is the result of combining the lateral images with the generated mask and its inverse.

B SLEEP PATTERNS

B.1 UNIHEMISPHERIC SLEEP

This type of sleep pattern is predominant in aquatic mammals and birds. In order to avoid drowning or crashing during long migratory flights, the animals keep one eye open and the corresponding hemisphere active during sleep. In the FF algorithm, we represent this type of behavior by updating the model parameters after performing both forward passes and combining the two losses.

B.2 MONOPHASIC SLEEP

Monophasic sleep can be described as sleeping once during 24 hours, usually for more than 5-6 hours. Most common in mammals, the pattern represents the perfect example of how the awake and sleep phases can be successfully separated. Despite the usual difference between the time spent in the sleep and awake phase, our experiments consider the two phases to be of equal duration.

B.3 SLEEP DEPRIVATION

While the previous case allows for detachment between the positive and negative phases, there is no advantage as the model spends the same amount of time sleeping as it is awake. The negative phase of the training removes the parasitic patterns learned by the model; thus, a more efficient negative phase would decrease the time spent training. The two phases should be disproportionate, prioritizing awake time. We associate this with partial sleep deprivation, forcing a trade-off between learning capabilities and usage.

C MODEL ARCHITECTURE

All experiments are conducted with the same architecture consisting of three fully connected layers having 784, 500, and 500 neurons, respectively. Each layer is trained for 50 epochs with a threshold of 1.5. We make use of the Adam optimizer with a batch size of 512. The default learning rate is set at 0.001 for the equal phases and scaled accordingly for the disproportionate phase.