ANNs Dream of Augmented Sheep: An Artificial Dreaming Algorithm

Gustavo Assunção[1,2][a], Miguel Castelo-Branco[3][b] and Paulo Menezes[1,2][c]

[1]*University of Coimbra, Department of Electrical and Computer Engineering, Coimbra, Portugal*

[2]*Institute of Systems and Robotics, Coimbra, Portugal*

[3]*Coimbra Institute for Biomedical Imaging and Translational Research (CIBIT), Coimbra, Portugal*

Keywords: Machine Learning, Overfitting, Data Augmentation, ANN, Artificial Sleep.

Abstract: Sleep is a fundamental daily process of several species, during which the brain cycles through critical stages for both resting and learning. A phenomenon known as dreaming may occur during that cycle, whose purpose and functioning have yet to be agreed upon by the research community. Despite the controversy, some have hypothesized dreaming to be an overfitting prevention mechanism, which enables the brain to corrupt its small amount of statistically similar observations and experiences. This leads to better cognition through non-rigid consolidation of knowledge and memory without requiring external generalization. Although this may occur in numerous ways depending on the basis theory, some appear more adequate for homologous methodology in machine learning. Overfitting is a recurrent problem of artificial neural network (ANN) training, caused by data homogeneity/reduced size and which is often resolved by manual alteration of data. In this paper we propose an artificial dreaming algorithm, following the mentioned hypothesis, for tackling overfitting in ANNs using autonomous data augmentation and interpretation based on a network's current state of knowledge.

1 INTRODUCTION

When considering the average human spends roughly a third of their life sleeping (Colten and Altevogt, 2006), it is natural to want to understand the importance of this necessity. Not surprisingly, studies have linked sleep to several cognition-related factors which affect daily life. For instance, as a foundation for permanent establishment of recent experiences (Deak and Stickgold, 2010) and understanding of new knowledge (declarative or not). It is therefore only natural to expect certain learning-related processes to occur during sleep, such as data augmentation and scenario simulation, which eventually result in enhanced perception when we awake. To exemplify, these phenomena have been observed extensively by Matt Wilson on rodents (Foster and Wilson, 2006).

One of the most fascinating aspects of sleep is the occurrence of dreams: virtual concoctions of memory, emotion and knowledge, recent or consolidated, which result in multi-sensorial experiences of disputed significance. Some researchers have argued for an evolutionary take on dreaming, where brains have developed the ability to simulate threatening or unresolved situations and determine the course of action most likely to culminate with success and survival (Blackmore, 2012), (Adami, 2006). Others recorded examples of active problem solving being catalyzed on participants who, while dreaming, perceived apparent solutions they were not consciously aware of before (Barrett, 1993). This evaluation of internalized problems during sleep is even congruent with attempted task integration in dreams (Schoch et al., 2019). Yet, dreamless sleep has also been correlated with performance improvement and learning (Cao et al., 2020). In fact, neural pattern replay typical of non-dream stages is critical for abstracting core knowledge and consolidate memory, as evidence suggests (Lewis et al., 2018). However, this offers no concise explanation on the utility of idiosyncratic episodes during other sleep stages. Further, the odd and consistently scattered nature of dreams disfavors the objective usefulness of these unconscious experiences in daily life, as noted by (Hoel, 2021). Dream absurdity could still be associated with the emotional factor of subconscious simulations, contributing to a mental preparation by hyperbolization

[a] https://orcid.org/0000-0003-4015-4111
[b] https://orcid.org/0000-0003-4364-6373
[c] https://orcid.org/0000-0002-4903-3554

Assunção, G., Castelo-Branco, M. and Menezes, P.
ANNs Dream of Augmented Sheep: An Artificial Dreaming Algorithm.
DOI: 10.5220/0011055700003209
In *Proceedings of the 2nd International Conference on Image Processing and Vision Engineering (IMPROVE 2022)*, pages 135-141
ISBN: 978-989-758-563-0; ISSN: 2795-4943

of potential scenarios (Scarpelli et al., 2019). However, the fuzziness of emotionality discourages the notion that the same exact emotional state could continuously be experienced until a numbness is developed. Indeed there is an abundance of open-ended theories of sleep in the fields of Psychology and Neuroscience. Yet, traversability to current machine learning seems more plausible with the Overfitted Brain Hypothesis (OBH) (Hoel, 2021). This notion envisions the brain warping the statistically proximal instances observed throughout a day, in the form of dreams so as to prevent its overfitting. This begs the questions: Could dreams function as our data augmentation process for finer abstraction? If so, can ANNs benefit from a similar capability?

The scarcity of definitive example cases and distinct data instances available during the day seems contradictory to the generalization capabilities of the human psyche. Moreover, the intuitiveness of the brain is what enables condensation of experiences with incorporation of scattered priors to generate new intelligence (Dubey et al., 2018). Therefore OBH becomes increasingly plausible, as the corrupted data of dreams could function as a source of creativity in machines. To exemplify, statistical corruption of data as that performed by Google's DeepDream (Mordvintsev et al., 2015) has been used for multi-candidate dreamed object classification, through correlation of real image feature vectors and decoded brain activity patterns obtained from sleeping subjects (Horikawa and Kamitani, 2017). Thus a logical next step would be to employ these data warping techniques in generating augmented data for ANNs to improve their training. Evidently not following the same procedure as the original training of the model, but implementing certain biological aspects outlined by OBH which prevent overfitting. To summarize, this paper advocates for potential overfitting prevention in ANNs through augmentation of input data based on a model's current knowledge, most similar to the imaginative characteristics of real dreaming.

2 BIOLOGICAL SLEEP

Sleep is commonly registered as alternating between two stages over the course of a single session, namely REM (rapid eye movement) and NREM (non-REM) sleep. Opposing wave activity makes for an easy distinction between the two stages (Martin et al., 2020). Whilst slow-wave activity (SWA) is characteristic of NREM, the accentuated excitatory-inhibitory oscillating behavior gradually stabilizes as the brain cycles to REM and shifts to higher frequency (gamma) activity. This shows a potential inverse correlation between SWA decrease, gamma increase and dream intensity/recallability (Siclari et al., 2018). Moreover, as replays of previous neural activity, NREM experiences are recalled as mundane and memory-related. Contrastingly, REM episodes display a surrealistic disposition (Manni, 2005). It is therefore presumable that NREM sleep accounts for the consolidation necessities of the brain, while REM deals with the cognitive and data augmenting aspects detailed by OBH. As such, our approach could be interpreted as simulating an artificial REM phase, whilst conventional training would correspond to NREM, and finally post-training usage be homologous to wakefulness.

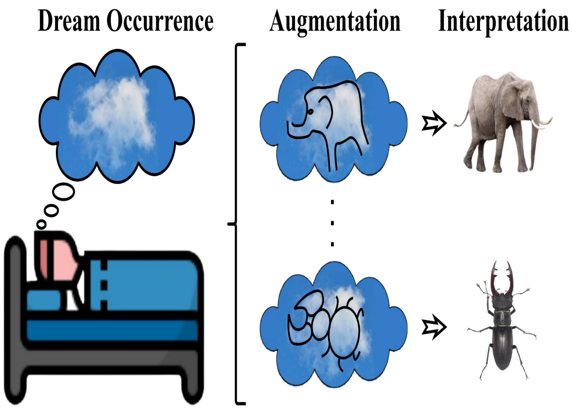

Figure 1: Dream data augmentation, from left to right. Dream starts either by random memory access or activation of brain structures. Based on current knowledge, personal interest, persisting thoughts and others, the dream is morphed to match the current state of the brain. Simultaneously it attempts to interpret this information in order to form a response to it. Consequently, it helps negate latent overfitting.

2.1 REM Augmentation

During REM sleep, the brain resembles its awake state with intense bursts of activity occurring in a seemingly random fashion. In addition to these desynchronized and fast brain waves, the body transitions to a new homeostatic balance (Ranasinghe et al., 2018) whose objectives differ from the usual optimized functioning. Thus the brain is capable of processing information in a manner unconventional with its waking operation, which suggests an exploration-exploitation dichotomy resonating with sleep-wakefulness. This would also be supported by the decreased activity of the prefrontal cortex, responsible for logic and planning, during REM dreaming which allows the dreamer to forgo risk aversion and consider novel associations (exploration). As such, considering OBH and how the small amount of examples available to a real brain does not hinder its recognition abilities in real life, it is possible that data augmentation occurs during REM sleep.

Since the brain does draw from general knowledge, memory and episodic information to simulate scenarios of what it perceives (Foulkes, 2014), (Domhoff, 2003), a process through which biological dream data corruption can occur involves taking statistically similar observations idly flowing through the brain and forcing a set of patterns on them for interpretation according to what the brain knows and comprehends. Logically these must also correspond to the patterns most likely to fit on the initial data, and as the dream progresses they may be enhanced to form decipherable examples of either novel or recurring information. A diagram of this hypothesis is shown on Figure 1 using visual data, despite the process' applicability to other data types. In the figure an initial dream of clouds is forced with broad stroke animal shape patterns, likely according to the subject's interests and curiosity, which are afterwards interpreted as valid yet soft object recognition examples. Finally, this bio-based theory can be readily adjusted for machine learning as we detail in the following section.

3 ARTIFICIAL SLEEP

Much like the human brain functions according to OBH, with artificial sleep our primary objective is to come up with a mechanism for overfitting self-prevention in neural networks. This encompasses a conditional extra training stage during which models ”dream”, augmenting or inferring data. Such a process could be integrated in most ANN frameworks and architectures with little effort.

3.1 Hypothetical Formulation

The intense but incoherent neural activity approximating REM sleep to wakefulness suggests a continuous attempt of the brain to correlate and make sense of information flowing through it, which can only occur in accordance with the knowledge available. The source of this information can, from our perspective, be two-fold: noisy input, caused by random activation of certain neural regions or external stimuli to a dormant system; objective albeit out-of-context input from memory accesses induced by the shifted electrical connectivity of a brain in REM sleep. This rewiring could explain REM sleep's lack of correlation with neural replay, opposing NREM. In any case, those special inputs are adulterated by the network which tries to process them, inadvertently acquiring

Figure 2: Deepdream's maximization of layer activations used for augmentation of non-relevant data (rocky formations) on the left. The resulting image on the right, with relevant information (buildings resembling pagodas), using a network trained on places by MIT Computer Science and AI Laboratory as presented in (Mordvintsev et al., 2015). The new image could be useful for disrupting overfitting, as the network continues training on place identification, using soft labels to account for its augmented characteristics (i.e. the pagodas).

attributes specific to its current state of knowledge and structure. Homologously, ANNs can be made to integrate global patterns of learned data on new extrinsic data instances through excitation of the layers which detect them (Mordvintsev et al., 2015). Even though DeepDream's usage has exclusively targeted visual data, the premise still holds for any other input type, provided the goal is not solely art generation.

Generalizing DeepDream's maximization of layer activations to non-image processing networks can yield outputs with little meaning to human interpretation, even artistically speaking. Yet to the models inducing the patterns themselves, the resulting representations will feature non-relevant data augmented by relevant information (e.g. speaker X characteristics integrated into the sound features of an animal cry, by/for a speaker recognition model). An example of this augmentation, specifically on visual data in order to be more easily comprehensible, is shown in Figure 2. If employed correctly, these fabricated examples can be useful in terms of overfitting disruption. With that in mind, the question comes down to how the examples should be perceived post-dream, by the networks which generate them. Labels should reflect both the original and augmented characteristics of the corrupted data, so that its combination may enable new knowledge inference. Further, this labeling process must not influence the parameters of a network, as it would most likely lead to catastrophic interference. Minor forgetting of overfitted details should not be disregarded entirely as it could prove useful, and may be added to the presented technique in the future.

The ambiguity of artificial dream data characteristics is also congruent with real dream recalling (often fuzzy). To account for this level of imprecision while simultaneously assuring the usability of that data, the generated labels must necessarily be soft. This is imperative to the implementation, as it is more accurate to human memory and likely the reason for our occasional confusion. Labels can then be learnable as dream data is iteratively interpreted according to the current state of a model's knowledge. Hence label updates would minimize the impact of their corresponding instances on the training of that model. This process is agreed to have potential since Soft-Label Dataset Distillation (Sucholutsky and Schonlau, 2019) has used a similar notion successfully. In it, distilled labels are perfected to minimize the error of a model on real data, when trained with a single forward pass of distilled data.

Finally, the conjectural nature of dream elements requires an escape in case no interpretation is found plausible by the network. In the case of humans, when dealing with content we are unable to make sense of, a "nonsense" classification is frequently employed and that specific content later disregarded. Respectively, the same can be implemented in ANN dreaming with an additional absurdity class which accounts for meaningless dream aspects. This class can later be ignored for post-dream model purposes such as recognition applications.

3.2 Workflow

A fundamental requirement of this algorithm is its applicability to existing neural network frameworks. As such we attempted to use typical nomenclature as much as possible. Additionally, from here on we employ the same basic notation as (Sucholutsky and Schonlau, 2019), which presumes a K-layered neural network f parameterized by θ, with typical backpropagation based on a twice-differentiable loss function $l_1(x_i^r, y_i^r, \theta)$. The goal of such models is usually to find an optimal set of parameters θ^*, using a training dataset $\mathbf{r} = \{x_i^r, y_i^r\}_{i=1}^N$, according to (1).

$$\begin{aligned} \theta_{new} &= \underset{\theta}{arg\,min} \frac{1}{N} \sum_{i=1}^{N} l_1(x_i^r, y_i^r, \theta) \\ &\triangleq \underset{\theta}{arg\,min}\, l_1(x^r, y^r, \theta) \end{aligned} \tag{1}$$

This is achieved iteratively with small batches of training data and stochastic gradient descent (SGD), according to (2), where η denotes a preset learning rate. Even though additional parameters may be present, such as a momentum α, we disregarded these for the sake of simplicity as our technique is independent from them.

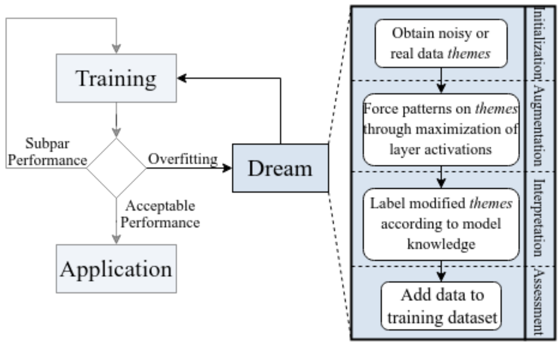

Figure 3: Overview of the neural network's workflow, with the dream stage being used as a mechanism to tackle overfitting.

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta_t} l_1(x_{batch}^r, y_{batch}^r, \theta_t) \tag{2}$$

Assuming that at some moment $t > 1$ the network is deemed overfitted, training is halted and a dream process starts to disrupt the overfitting. Following the hypothesis described in the previous section, this process can be divided into 4 major phases: initialization, augmentation, interpretation and assessment. These are shown orderly in the flow diagram of Figure 3. The first is meant to parallel the random flow of information occurring in REM sleep. Here we consider a dream dataset $\mathbf{d} = \{x_i^d, y_i^d\}_{i=1}^L$ whose instances, named augmentation *themes*, each make up the general setting of a single dream and which will be augmented using a model's current knowledge. The main requirement for $\mathbf{d}$, however, is that it must not include data existent in the regular dataset $\mathbf{r}$ being used for training. There are several ways to achieve this, depending on data availability, application goals and other factors. The most simple would be noise initialization, in which case each *theme* would have no inherent meaning. Real data may also be used, in which case it may or may not be related with the training dataset. In the former case, data can be split into train, test and dream data to accommodate this extra set, whereas in the latter, data may be sourced from another completely unrelated dataset with the imposition of it being available in the same modality as the training dataset (e.g. using face images as *themes* when training with CIFAR10 (Krizhevsky et al., 2009) datasets).

The second phase is meant to incept knowledge the network already possesses onto the dream *themes*. This is achieved through a loss objective $\mathcal{L}_2$, which depends on layer activations $a = f_\theta(x^d)$ (i.e. chosen layers' outputs given a forward pass of a *theme*). The augmentation stems from the excitation of a few randomly chosen layers $k \in K \setminus \{input, output\}$, which will force the patterns they generally identify into the *themes* they receive. Considering network layers deal with different levels of abstraction and feature complexity, as described in the functioning of Deepdream (Mordvintsev et al., 2015), depending on factors such as depth and purpose each layer will impose different characteristics on the data being augmented. This is achieved using a differentiable loss function $l_2(xd_{batch}, a_k)$ based on layer activation, calculated individually for each *theme*.

$$\mathcal{L}_2(\tilde{x}^d, \theta_t) := l_2(\tilde{x}^d, a_k) \tag{3}$$

$$\begin{aligned} \tilde{x}^d_{new} &= \underset{\tilde{x}^d}{arg\,max}\, \mathcal{L}_2(\tilde{x}^d, \theta_t) \\ &= \underset{\tilde{x}^d}{arg\,max}\, l_2(\tilde{x}^d, a_k) \end{aligned} \tag{4}$$

With new and possibly meaningful information now present on the *themes* post-augmentation, the third phase is carried out based on the one-step loss objective $\mathcal{L}_1$, here minimized by learning those *themes* and their corresponding labels, for $\theta_1 = \theta_0 - \eta \nabla_{\theta_0} l_1(\tilde{x}^d, \tilde{y}^d, \theta_0)$.

$$\mathcal{L}_1(\tilde{x}^d, \tilde{y}^d; \theta_0) := l_1(x^r, y^r, \theta_1) \tag{5}$$

However, this minimization of the $\mathcal{L}_1$ objective is carried out over $\tilde{y}^d$ only, and it does not consider $\tilde{x}^d$ as dataset distillation (Sucholutsky and Schonlau, 2019) would. That is because we intend for the model to merely interpret dream data rather than optimize it according to its current knowledge.

$$\begin{aligned} \tilde{y}^d_{new} &= \underset{\tilde{y}^d}{arg\,min}\, \mathcal{L}_1(\tilde{x}^d, \tilde{y}^d; \theta_0) \\ &= \underset{\tilde{y}^d}{arg\,min}\, l_1(x^r, y^r, \theta_0 - \eta \nabla_{\theta_0} l_1(\tilde{x}^d, \tilde{y}^d, \theta_0)) \end{aligned} \tag{6}$$

Evidently, the greater number of times the second and third phases are carried out, the more prevalent the forced patterns will be on each dream *theme* as well as the network's interpretation of them. Additionally the number of recognized classes, whose vector is updated by (6), is also extended to include an absurdity class which can be dropped once the algorithm concludes, as noted in the previous section. Algorithm 1 realizes the described procedure for implementation. Once the algorithm finishes its execution, the resulting pairs of augmented dream *themes* and soft labels are integrated into the regular dataset **r** to complete the fourth and last phase of the dreaming process. Finally overfitting disruption occurs once regular training is resumed and goes through these new data samples.

3.3 Advantages and Limitations

Despite its capability to modify existing data with information meaningful to the training being carried out, this technique is not similar to data augmentation performed by generative models. Some advantages include its architecture agnosticism, as the algorithm is applicable regardless of network architecture. This is unlike generative data augmentation, where typically a full model or extra generative section is added to the main architecture, and trained to augment data, incurring additional memory and computational power requirements. Additionally, generative data augmentation does not have an interpretation phase following the network's current state of knowledge, with labels being retained from the original data source or inferred from latent space distribution.

Algorithm 1: Overfitting Disruption.

```
Input: M: Number of themes to occur during dream; α: step size; n: batch size; T: Dream depth in steps; ỹ₀^d: initial value for ỹ^d.
Output: Augmented dream data (x̃^d, ỹ^d).
  Dream Data of Theme Initialisation :
  ỹ^d = {ỹ_i^d}_{i=1}^M ← ỹ_0^d
  x̃^d = {x̃_i^d}_{i=1}^M randomly OR sample batch from dream dataset d
  for each training step t = 1 to T do
      for each layer k ∈ {randomly chosen layers} do
          for dream theme x̃_i^d do
              Forward pass the theme
              Evaluate objective function on activations
                  L_2^(k,i) = l_2(x̃_i^d, a_k)
          end for
          Compute updated model parameters with SGD:
          θ_1 = θ_0 − η∇_θ0 l_1(x̃^d, ỹ^d, θ_0)
          Evaluate objective function on real training data:
              L_1^(k) = l_1(x^r_batch, y^r_batch, θ_1^(k))
      end for
      Update dream data:
      ỹ^d ← ỹ^d − α∇_ỹd Σ_k L_1^(k), and
      x̃_i^d ← x̃_i^d + α∇_x̃i^d Σ_k L_2^(k,i)
  end for
```

As any other technique however, the presented algorithm also has its drawbacks. Specifically the corruption of dream *themes* with training data patterns, which relies on gradient ascent, imposes a computational cost proportional with dream depth (i.e. the depth of the layers whose activations are maximized over the *themes*, through gradient ascent). Nonetheless, this increase is coherent with the processing intensity of a brain in REM sleep which, according to energy expenditure, matches and often exceeds that of wakefulness. Another depth-related potential issue is related with the escape absurdity class included in the soft label vectors of the dream data. The shallower the dream process is made to be, the more likely it is that the network will be unable to interpret the

Table 1: Exemplary run of two CIFAR10 images as themes (top - airplane, bottom - ship) over a single dream iteration, using a double-layered CNN trained exclusively for MNIST handwritten digit recognition. 'Original' rows show the untouched CIFAR10 images, while 'Conv1' and 'Conv2' each refer to a 800-step run of the Deepdream technique over the CIFAR10 images activating the first and second convolutional layers, respectively. Probabilities, shown as percentages, refer to the evaluation of their respective images by the MNIST-trained CNN (i.e. the soft-labels they would attributed after this initial iteration).

											
Original		13.7	4.8	15.9	6.5	14.8	5.2	10.2	2.3	23.6	3.0
Conv1		0.0	0.0	10.1	0.0	0.0	0.0	0.0	0.0	89.9	0.0
Conv2		0.0	0.0	83.1	1.3	0.1	12.1	0.0	0.7	2.7	0.0
Original		41.4	0.3	16.8	5.0	3.7	0.6	16.2	0.0	15.8	0.2
Conv1		0.0	0.0	98.9	0.9	0.0	0.0	0.0	0.0	0.2	0.0
Conv2		0.0	0.0	66.0	7.8	0.0	26.1	0.0	0.1	0.0	0.0

theme which it augments. As a consequence, the absurdity class can overwhelm the others in the label vector, making the augmented instance insignificant for later overfitting prevention purposes. It should be noted however that either of these issues can be mitigated with enough dream depth. As such, they can be negligible if enough computational resources are available or no critical time constraints are imposed for training.

4 DEMO POTENTIAL AND FUTURE WORK

The presented procedure is a working idea. However, as demonstrated in Table 1, it clearly shows potential in terms of autonomous data augmentation. Despite the fact that CIFAR10 images are completely unrelated with handwritten digit recognition, the forcing of patterns over those images by a CNN trained with the MNIST (Deng, 2012) dataset yields data with potentially meaningful information to that same network, which may disrupt the eventual monotony of training data and consequential overfitting. The interpretation phase assures this by allowing the network to label the results according to its current state of knowledge, so they may be later added to the training set. For instance, in Table 1 the image of a ship enhanced by the second convolutional layer of the MNIST-trained CNN produces something closely resembling the digits 2, 3 and 5. Since the CNN applied here was not overfitted, its low loss implies that it is more certain of its predictions than an overfitted counterpart would be. Thus, more digits could possibly also be interpreted by the network, were it overfitted or also if more iterations were carried out.

Evidently, specific experimental scenarios are required for validation. With that in mind, in addition to the ongoing implementation of the described algorithm, we plan to devise adequate experiments for its validation. This will include but is not limited to:

1. Comparison of integration in shallow and deep ANNs;
2. Performance assessment with different data types;
3. Exploration of layer suitability for excitation;

Ultimately we intend to build on this first autonomous step against overfitting, until obtaining a technique capable of dealing with the greatest amount of scenarios possible, while also being applicable to common ANN architectures.

5 CONCLUSION

Research into the intricate cognitive aspects of sleep is ongoing and contributing to a growing understanding of dream phenomena. Despite the potential usefulness of such advancements in addressing common issues of machine learning, such as overfitting, there is a continuous disregard for neuroscientific findings. Ultimately the issues remain latent in machine learning models, consistently hindering their performance until being resolved through manual intervention.

Our position reflects support for the introduction of a dream stage in the training process of machine learning models, to function as a self-sufficient mechanism for preventing overfitting. To this end, we proposed an artificial ”dreaming” algorithm to achieve this goal in ANNs. The process works by augmenting data through excitation of layer activations and interpretation of that same data according to current model knowledge. We hope to obtain successful results in the future with the implementation of the outlined ANN ”dreaming” procedure, further supporting general autonomy in artificial intelligence.

ACKNOWLEDGEMENTS

This work was supported in part by scholarhip 2020.05620.BD of the Fundação para a Ciência e a Tecnologia (FCT) of Portugal and OE - National funds of FCT/MCTES under project number UIDB/00048/2020.

REFERENCES

Adami, C. (2006). What do robots dream of? *Science*, 314(5802):1093–1094.

Barrett, D. (1993). The ”committee of sleep”: A study of dream incubation for problem solving. *Dreaming*, 3(2):115–122.

Blackmore, S. (2012). *Consciousness: an introduction*. Oxford University Press, New York.

Cao, J., Herman, A. B., West, G. B., Poe, G., and Savage, V. M. (2020). Unraveling why we sleep: Quantitative analysis reveals abrupt transition from neural reorganization to repair in early development. *Science Advances*, 6(38).

Colten, H. R. and Altevogt, B. M. (2006). *Sleep Disorders and Sleep Deprivation: An Unmet Public Health Problem*. The National Academies Press, Washington, DC.

Deak, M. C. and Stickgold, R. (2010). Sleep and cognition. *WIREs Cognitive Science*, 1(4):491–500.

Deng, L. (2012). The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142.

Domhoff, G. W. (2003). *The scientific study of dreams: Neural networks, cognitive development, and content analysis.* American Psychological Association.

Dubey, R., Agrawal, P., Pathak, D., Griffiths, T., and Efros, A. (2018). Investigating human priors for playing video games. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1349–1357, Stockholmsmässan, Stockholm Sweden. PMLR.

Foster, D. J. and Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680–683.

Foulkes, D. (2014). *Dreaming: A cognitive-psychological analysis*. Routledge.

Hoel, E. (2021). The overfitted brain: Dreams evolved to assist generalization. *Patterns*, 2(5):100244.

Horikawa, T. and Kamitani, Y. (2017). Hierarchical neural representation of dreamed objects revealed by brain decoding with deep neural network features. *Frontiers in Computational Neuroscience*, 11.

Krizhevsky, A. et al. (2009). Learning multiple layers of features from tiny images. *Drug Invention Today*.

Lewis, P. A., Knoblich, G., and Poe, G. (2018). How memory replay in sleep boosts creative problem-solving. *Trends in Cognitive Sciences*, 22(6):491–503.

Manni, R. (2005). Rapid eye movement sleep, non-rapid eye movement sleep, dreams, and hallucinations. *Curr Psychiatry Rep*, 7(3):196–200.

Martin, J., Andriano, D., Mota, N., Mota-Rolim, S., Araújo, J., Solms, M., and Ribeiro, S. (2020). Structural differences between rem and non-rem dream reports assessed by graph analysis. *PloS one*, 15(7):1806–1809.

Mordvintsev, A., Olah, C., and Tyka, M. (2015). Deepdream - a code example for visualizing neural networks.

Ranasinghe, A., Gayathri, R., and Priya, V. (2018). Awareness of effects of sleep deprivation among college students. *Drug Invention Today*, pages 1806–1809.

Scarpelli, S., Bartolacci, C., D'Atri, A., Gorgoni, M., and Gennaro, L. D. (2019). The functional role of dreaming in emotional processes. *Front. Psychol.*, 10.

Schoch, S. F., Cordi, M. J., Schredl, M., and Rasch, B. (2019). The effect of dream report collection and dream incorporation on memory consolidation during sleep. *J Sleep Res*, 28(1).

Siclari, F., Bernardi, G., Cataldi, J., and Tononi, G. (2018). Dreaming in NREM sleep: A high-density EEG study of slow waves and spindles. *J. Neurosci.*, 38(43):9175–9185.

Sucholutsky, I. and Schonlau, M. (2019). Soft-label dataset distillation and text dataset distillation. arXiv preprint.