Sensory complexity and global gain in a DCNN codetermine optimal arousal state

Abstract

Arousal deeply impacts behaviour and sensory processing. Whereas for an easy task, perceptual performance linearly increases with higher arousal levels, for a more challenging task, an inverted-U-shaped relationship has been described. These findings are commonly referred to as the Yerkes-Dodson law (1908). Yet, it remains unclear why perceptual performance decays with high levels of arousal in difficult but not in simple tasks. Based on recent studies linking a global gain change in sensory processing to changes in arousal state, we augmented a deep convolutional neural network with a global gain mechanism to mimic the effects of cortical arousal. With this approach, we show that the Yerkes-Dodson law can be accounted for by this global gain mechanism, acting on a hierarchical sensory system that processes stimuli of varying sensory complexity. By leveraging the full observability of our model, we reconcile conflicting findings from previous studies on sensory processing, by showing that both linear as well inverted-U-shaped gain profiles emerge in the interaction of hierarchical sensory processing and global arousal changes.

1 Introduction

The effects of neuromodulation are pervasive in the brain. Recent estimates are that around 11% of the variance of spontaneous activity in visual cortex is related to arousal state [1]. For behaviour, already in 1908, Yerkes and Dodson [2] established that an animal's ability to learn a sensory discrimination task was deeply affected by its arousal state (which they manipulated by electrical stimulation). This effect took the shape of an inverted U for challenging perceptual discriminations, indicating that the animals learnt the discrimination best at intermediary levels of stimulation. Interestingly, this relationship became linear for perceptually easy tasks: The animals were fastest at learning when confronted with the strongest stimulation intensities. This interaction between arousal state and perceptual task difficulty and its effect on performance has been replicated many times (e.g. [3]). Yet, the origin of this interaction is much less clear. Some have proposed that an involvement of frontal cortices (e.g. prefrontal cortex) is the decisive factor for an inverted U-shaped performance profile [4], while other suggest that the saturation of activation levels in the locus coeruleus – the central source of noradrenaline – is at the core of this transition [5]. Yet, these broad proposals still do not specify *how* performance degrades with increasing levels of arousal for a perceptually difficult task.

Studies examining the effects of arousal (indexed by pupil dilation) on processing in sensory cortices and perceptual performance provide conflicting answers to this question [6]. For early visual cortex, a number of studies report a linear increase in firing rates with increases in arousal during sensory stimulation [7–10]. One of these studies also reported that the animals performed best at a go/no-go task with perceptual discrimination at the highest levels of arousal, identified through locomotion [8]. Yet, another study in auditory cortex that also used a go/no go-task discovered an inverted U-shaped

Submitted to 2nd Workshop on Shared Visual Representations in Human and Machine Intelligence (SVRHM).

Figure 1: A: Parametric perceptual discrimination task in which the network is tasked to distinguish a natural scene from an average image. The perceptual difficulty decreases with the number of images used for the average image. B: Different ResNet18 models with a single sigmoidal read-out neuron were tested on this task while varying global gain. While we see a clear inverted U-shaped performance profile for the difficult task, the link between global gain and performance becomes more monotonically increasing the easier the task gets.

link for both firing rates during sensory stimulation and performance levels [11]. In contrast to studies in visual cortex, this finding thus suggests an optimum for both a gain in firing rates as well performance at intermediate levels of arousal.

The differences across these studies have been mainly attributed to the difference in sensory modality [6], yet more recent studies suggest a high degree of similarity for the effects of arousal on both visual and auditory processing [12]. Beyond the sensory modality, another striking difference between these studies is the complexity of sensory stimuli that were presented to the animals. While the study with the inverted U-shaped relationship used quite complex sensory stimuli for their go/no go task (presence of pure tones embedded in sequences of structured noise, [11]), the studies suggesting a linear link used rather simple stimuli (presence of a large Gabor patch on a grey screen, [8, 10]). This prompted us to ask whether the different findings of arousal on evoked sensory processing could be due to a difference in experimental stimuli instead.

To answer this question, we implemented a global gain model to mimic the effects of cortical arousal on sensory processing in a deep convolutional neural network (DCNN) and tested it on a range of tasks with varying degrees of sensory complexity. With our approach, we show that long established effects of arousal on perceptual performance, as captured by the Yerkes-Dodson law, can be accounted for by a global gain mechanism acting on a hierarchical sensory system processing stimuli of varying complexity.

2 Global gain as a model for cortical arousal

Cortical arousal is linked to increased firing rates in response to sensory stimulation in sensory cortices, i.e. increased gain [6]. These increases in gain have been argued to be implemented via a global release of noradrenaline [13], making them spatially unspecific and temporally slow. For creating an analogous situation in a DCNN, we used a ResNet18 architecture [14] with sigmoidal-like activation functions [15] that was pretrained on ImageNet [16]. To simulate the effects of arousal, we targeted all activation functions in the network simultaneously with one global gain parameter. This parameter multiplicatively scaled the output of the activation functions, thus giving rise to response gain everywhere in the network.

Figure 2: A: Example images from the datasets used to probe the effects of global gain on object recognition. Specifically, we curated two datasets from the MS COCO database [17] with a homogeneous context and 8 possible target categories. B: Performance of the finetuned ResNet18 networks for both datasets. Both datasets also result in an inverted U-shaped profile.

3 Results

3.1 Yerkes–Dodson law – A global gain model accounts for behavioural effects of arousal

To evaluate the link between global gain and perceptual difficulty, we developed a task in which we could parametrically alter the visual complexity required to solve a binary discrimination task (Figure 1A). The network was tasked to distinguish a natural scene from an average image. The perceptual difficulty was determined by the number of images that were used to compute the average images, with average images that consisted of fewer images being more difficult to distinguish from the natural scenes. Next, we fine-tuned a single sigmoidal neuron connected to the fully-connected layer of the ResNet18. The remaining network weights were kept unchanged. For all degrees of visual complexity, a separate model was fine-tuned and tested with varying degrees of global gain (Figure 1B).

For the most difficult task, based on an average of two images, we observed an inverted U-shape relationship between global gain and performance. The best task performance was achieved close to a gain of 1, corresponding to the evaluation of the original model. Further in- or decreases in gain only hampered performance, thus resulting in an inverted U-shaped profile. Interestingly, the easier the task became the more this inverted U transformed into a monotonically increasing performance profile that had its optimum at very large gain values. This means that the model performed best in a state that was very different from training. Across all perceptual difficulties, performance decreased back to chance performance at a similar global gain value.

This pattern of findings directly replicates the Yerkes-Dodson law [2], assuming that very low and very high levels of global gain are not accessible experimentally or plausible to occur in a living organism. Moreover, we show that the transition from a linear to an inverted U-shape profile is a continuum and could arise directly from applying a global gain factor to a hierarchical sensory system such as a DCNN.

3.2 From binary to multi-class tasks - The Yerkes Dodson law also holds for object recognition

A commonality of many studies describing the Yerkes-Dodson law is the usage of sensory discrimination tasks with artificial stimuli that typically involve the detection of a single target class [2, 3, 8, 11]. Yet despite this commonality, these findings have been interpreted to apply to task difficulty more broadly [4, 5]. To test the generalization of our first experiment that the inverted U-shaped performance profile arises in tasks with complex visual stimuli, we therefore reasoned that it should also hold for other visual recognition tasks with more than one target class. To this end,

Figure 3: Mean activations of the activation layers after each addition layer in the ResNet18. The used stimuli were the same as in Figure 1.

we fine-tuned the same ResNet18 on a 8-way multi-label object detection task with a shared context (Figure 2A).

As in the binary task, we again only trained an output layer consisting of eight sigmoidal output nodes. The networks were first trained with multi-label images, but evaluated on single target images. We find that also here the same pattern of results holds: Across both contexts, we observe an inverted-U-shaped performance profile (Figure 2B). Comparing the datasets shows that networks performed less well on the food compared to the street scenes at baseline, a gain of 1. In the data by Yerkes and Dodson [2] more difficult tasks resulted in a narrower inverted U-shape than less difficult task. Similarly in our results, the food scenes featured a narrower profile compared to the street scenes. For high gain regimes, we observed some marked differences to the binary task for the street dataset. While performance completely decayed in the food scenes for global gain parameters exceeding 1.2, street scenes are being solved above chance far beyond that. This could be indicative of a higher chance of guessing based on an overly driven network for the street compared to the food scenes.

In sum, we have thus shown that the Yerkes-Dodson law also holds for more complex behavioural scenarios in our model, such as object recognition in natural scenes. This experiment also further supported our claim that the Yerkes-Dodson law in behaviour is likely moderated by sensory complexity rather than other task factors.

3.3 The inverted U arises from high-level features

In a final analysis, we sought to connect our results for performance to observations made experimentally in sensory cortices [6–11]. Specifically, arousal has been linked to both linear as well inverted-U-shaped neural gain profiles for evoked processing. To obtain an intuition on how activation levels were shaped by both stimulus complexity and global gain in our networks, we collected the mean activations of the activation functions after the addition in the ResNet18 across different gain parameters and stimulus complexities.

A general finding across all layers is that the two most perceptually difficult binary task datasets were producing the highest activation levels in almost all ResNet blocks (Figure 3). Furthermore, while there seems to be a linear effect of global gain for most early ResNet blocks (1 − 4), this link becomes

increasingly complex in the later blocks (5 – 8). Overall, this pattern is suggestive of increasingly specific features in later ResNet blocks, which cannot be activated to a similar extent by simpler visual stimuli such as those used for the perceptually easy task (cf. Figure 1A). Our analysis also makes it clear that both linear as well inverted-U-shaped activation profiles may occur in a hierarchical sensory system. The gain profiles observed experimentally [11, 10] that were first thought to be at odds with one another [6], are likely to co-occur when accounting for stimulus complexity and hierarchical processing.

This last result indicates that the characteristic inverted U-shape observed in perceptually difficult tasks might stem from an interaction between complex sensory features that are tuned to a specific activation range and global gain increases. We speculate that increases in arousal lead to activation patterns that do not match these complex features anymore, thus leading to a decay of these features and hence to lower mean activations at these stages of processing. For behaviour, this interaction between global gain and stimulus complexity can have direct consequences for perceptual performance: If the read-out layer learnt to capitalize on these features during training, this can either result in a linear or an inverted U-shaped performance profile.

4   Conclusions

In this paper, we implemented a global gain mechanism in a DCNN to mimic the effects of arousal on sensory processing. Combining performance-optimized computer vision models with recent evidence from neuroscience allowed us to address the long-standing question of why perceptual performance decays with high levels of arousal in perceptually difficult but not in simple tasks. Specifically, we found that with our straightforward global gain model we could robustly replicate the Yerkes-Dodson law for both binary as well as more challenging multi-label object recognition tasks. Beyond that, we have provided an intuition of how the switch from linear to inverted-U shaped performance profiles can arise, namely from an interaction between high-level sensory features and global gain increases. Importantly, we hereby put forward that this often reported change in performance profile for perceptual discriminations is at least in part due to an interaction with sensory processing, and not just the involvement of frontal cortices [4, 5]. Based on our results, we predict that an inverted-U-shaped gain profile with complex sensory stimuli should transition to a more linear gain profile with simpler sensory stimulation, both for perceptual performance as well as sensory activity.

Neuromodulation and arousal in particular have been suggested to belong to different behavioural repertoires in the context of reinforcement learning [5, 18]: While medium levels of arousal are thought to be particularly suitable for exploitative regimes, high levels of arousal were proposed to be particularly suitable for exploratory regimes, thus facilitating the switch to a new task. Our results provide a starting point for addressing these types of questions more directly in a performance-oriented framework with complex sensory processing.

References

[1] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):255, April 2019.

[2] Robert M Yerkes and John D Dodson. The relation of strength of stimulus to rapidity of habit-formation. *Punishment: Issues and experiments*, pages 27–41, 1908.

[3] P L Broadhurst. Emotionality and the Yerkes-Dodson law. *J. Exp. Psychol.*, 54(5):345–352, November 1957.

[4] David M Diamond, Adam M Campbell, Collin R Park, Joshua Halonen, and Phillip R Zoladz. The temporal dynamics model of emotional memory processing: a synthesis on the neurobiological basis of stress-induced amnesia, flashbulb and traumatic memories, and the Yerkes-Dodson law. *Neural Plast.*, 2007:60803, 2007.

[5] Gary Aston-Jones and Jonathan D Cohen. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.*, 28:403–450, 2005.

[6] Matthew J McGinley, Martin Vinck, Jacob Reimer, Renata Batista-Brito, Edward Zagha, Cathryn R Cadwell, Andreas S Tolias, Jessica A Cardin, and David A McCormick. Waking state: Rapid variations modulate neural and behavioral responses. *Neuron*, 87(6):1143–1161, September 2015.

[7] Cristopher M Niell and Michael P Stryker. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron*, 65(4):472–479, February 2010.

[8] Corbett Bennett, Sergio Arroyo, and Shaul Hestrin. Subthreshold mechanisms underlying state-dependent modulation of visual responses. *Neuron*, 80(2):350–357, October 2013.

[9] Jacob Reimer, Emmanouil Froudarakis, Cathryn R Cadwell, Dimitri Yatsenko, George H Denfield, and Andreas S Tolias. Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron*, 84(2):355–362, October 2014.

[10] Martin Vinck, Renata Batista-Brito, Ulf Knoblich, and Jessica A Cardin. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron*, 86(3):740–754, May 2015.

[11] Matthew J McGinley, Stephen V David, and David A McCormick. Cortical membrane potential signature of optimal states for sensory signal detection. *Neuron*, 87(1):179–192, July 2015.

[12] Daisuke Shimaoka, Kenneth D Harris, and Matteo Carandini. Effects of arousal on mouse sensory cortex depend on modality. *Cell Rep.*, 22(12):3160–3167, March 2018.

[13] Jacob Reimer, Matthew J McGinley, Yang Liu, Charles Rodenkirch, Qi Wang, David A McCormick, and Andreas S Tolias. Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nat. Commun.*, 7:13289, November 2016.

[14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint*, 1512.03385, December 2015.

[15] Davide Zambrano, Roeland Nusselder, H Steven Scholte, and Sander M Bohté. Sparse computation in adaptive spiking neural networks. *Front. Neurosci.*, 12:987, 2018.

[16] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, December 2015.

[17] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C Lawrence Zitnick, and Piotr Dollár. Microsoft COCO: Common objects in context. *arXiv preprint*, 1405.0312, May 2014.

[18] T Pfeffer, A Ponce-Alvarez, T Meindertsma, C Gahnström, R L van den Brink, G Nolte, K Tsetsos, A K Engel, G Deco, and T H Donner. Dissociation between catecholamines and acetylcholine in the human cerebral cortex. *bioRxiv preprint*, September 2020.