Journal Club

Editor's Note: These short, critical reviews of recent papers in the *Journal*, written exclusively by graduate students or postdoctoral fellows, are intended to summarize the important findings of the paper and provide additional insight and commentary. For more information on the format and purpose of the Journal Club, please see http://www.jneurosci.org/misc/ifa_features.shtml.

Adaptation, the Coding Catastrophe and Disaster Management in Natural Vision

Koen V. Haak[1] and Juraj Mesik[2]

[1]Donders Institute for Brain Cognition and Behaviour, Centre for Cognitive Neuroimaging, Radboud University, Nijmegen 6525 EN, The Netherlands, and
[2]Department of Psychology, University of Minnesota, Minneapolis, Minnesota 55455

Review of Zavitz et al.

If we stare at a waterfall for a while and then shift our gaze away, our visual environment will appear to move upward. First reported by Aristotle ~330 BC, this phenomenon is known as the waterfall illusion or the motion aftereffect (Anstis et al., 1998). The motion aftereffect is a manifestation of neural adaptation: viewing a waterfall selectively reduces the responsiveness of neurons tuned to downward motion, leading to imbalanced responses to static stimuli with relatively stronger responses of neurons tuned to upward motion (Barlow and Hill, 1963). Neural adaptation is thought to be crucial for sensory processing, for instance, to cope with internal noise and/or to increase coding efficiency (Andrews, 1964; Ullman and Schechtman, 1982; Barlow and Földiák, 1989; Wainwright, 1999; Clifford et al., 2000; Kohn, 2007; Webster, 2011, Webster, 2015; Haak et al., 2014a). Yet under some conditions, these benefits come at the cost of aftereffects, such as the illusion of false movement, which is clearly not beneficial for survival.

The prevailing theory of why neural adaptation produces aftereffects concerns the idea that changing the neural code at one level of sensory processing can lead to misinterpretations downstream (Schwartz et al., 2007; Series et al., 2009; Dhruv and Carandini, 2014; Patterson et al., 2014; Webster, 2015). This notion is known as the coding catastrophe: if neurons at one level of sensory processing change their tuning properties, then downstream decoding neurons must be "aware" of these changes or they will erroneously attribute these changes to the environment.

As adaptation appears to involve both costs and benefits, it is important to understand their balance in natural vision. Intuitively, one might expect that the benefits outweigh the costs under natural viewing conditions and that the coding catastrophe is only evident under extreme laboratory conditions. In everyday life, we walk around, move our heads, and shift our gaze ~3 times per second. Therefore, we rarely have the same retinal input for more than a few hundred milliseconds. Studies of visual adaption, on the other hand, typically require subjects or animals to hold a steady gaze at the center of the screen for much longer periods. Indeed, while the decoding process in downstream areas falls short in accounting for the exaggerated tuning changes induced in the laboratory, it may be robust to relatively minor adaptive changes that occur in natural vision.

A recent publication by Zavitz et al. (2016) takes a step toward addressing this question by studying how the representation of visual motion in area MT changes during adaptation to a relatively naturalistic, continuous stream of motion. While recording spiking activity using multielectrode arrays, the authors presented three anesthetized marmoset monkeys with sequences of coherently moving random dot patterns that changed direction every 500 ms. Thus, akin to everyday experience, the paradigm kept the distribution of motion directions highly diverse over time. This way, the effects of adaptation could be studied on shorter time scales, where each stimulus could be treated both as an adaptor for subsequent stimuli and a test stimulus for the preceding adaptors.

Zavitz et al. (2016) showed that the relatively brief stimuli induced robust neuronal gain reductions that lasted over multiple intervening test stimuli and were strongest for neurons whose tuning matched the stimulus direction. However, whether the adaptation was profound enough to induce perceptual biases (i.e., the motion aftereffect) is unclear based on the observation of gain changes alone. To infer the induction of aftereffects, the authors used spiking patterns from a portion of the trials to train a multiclass linear model for motion direction decoding, and applied it to estimate motion direction from the remaining trials. Indeed, the di-

rection estimates were biased in a fashion consistent with the traditional psychophysical aftereffects: the estimates were repulsively shifted away from the adaptor such that they deviated from the veridical direction by as much 4–6 degrees, although the peak of this effect appeared at somewhat higher adaptor test offsets than in psychophysics (60–90 vs 30 degrees in e.g., Levinson and Sekuler, 1976).

The results of Zavitz et al. (2016) thus suggest that the coding catastrophe also occurs under relatively naturalistic conditions. Why then does it not seem to affect our everyday perceptual experience? As mentioned above, one possibility is that the downstream decoders are robust to relatively minor coding changes induced by adaptation in everyday life, having learned from past experience how to correct for adaptation effects under natural conditions. Such "learning" might be an ongoing process, and accordingly, Zavitz et al. (2016) hypothesize that prolonged exposure to a visual environment with biased feature statistics could reduce the strength of perceptual aftereffects. Indeed, recent work in humans reported a reduction in the magnitude of perceptual aftereffects after many hours of exposure to altered image statistics (Haak et al., 2014b), suggesting a form of perceptual learning whereby high-level decoding neurons learn to reinterpret the new neural code.

Another possibility, not considered by Zavitz et al. (2016), is that the brain may also account for the coding catastrophe by taking advantage of contextual information (Powell et al., 2012). Zavitz et al. (2016) measured adaptation under relatively naturalistic presentation durations, but the stimuli themselves were not very naturalistic because the motion direction could only be inferred from a single source of information. In contrast, natural visual stimulation is generally much richer, with many cues from which the direction of motion can potentially be inferred. For instance, when we watch a car driving on a distant road, we can infer its direction from both the motion itself and the direction of the road because the road constrains the car's driving direction. Thus, in the presence of such contextual information, representations of individual cues may not need to be very accurate because they can be disambiguated by their context.

Given that adaptation is known to occur throughout the cortex, it is also interesting to consider that the decoders themselves may adapt as well. If such

"meta-adaptation" involves response gain changes whereby prolonged neuronal stimulation from upstream neuronal populations leads to gain reductions at the level of decoders, then this could allow for the down-weighing of inaccurate decoder outputs feeding into yet higher-level neurons. "Meta-adaptation" could therefore serve to correct the coding catastrophe by reducing the strength of the decoder outputs by just the right amount to balance out the inaccuracies produced by the strength of adaptation upstream. This way, high-level neurons that integrate the decoder's output with the contextual information would not have to be "aware" of any coding changes upstream.

What benefits could outweigh going through all this trouble? One proposal is that adaptation decorrelates neuronal population activity (e.g., Barlow and Földiák, 1989; Wainwright, 1999; Clifford et al., 2000; Kohn, 2007; Webster, 2011; Haak et al., 2014a). In line with this, Zavitz et al. (2016) show that stronger adaptation is associated with greater decorrelation, a feature that has also been observed for orientation adaptation in primary visual cortex (Benucci et al., 2013). The main benefit of decorrelation is likely that it improves coding efficiency, thereby leading to fewer active neurons and hence metabolic savings (e.g., Kohn, 2007). Indeed, the brain consumes much of our energy, so having a more energy-efficient brain could mean the difference between life and death when food is scarce.

Decorrelation is further thought to be beneficial for decoding performance. A stimulus drives not only those neurons whose tuning curves are centered on it, but also all other neurons with partially overlapping tuning. Without adaptation (and hence decorrelation), prolonged presentation of the same stimulus would increase correlations between the responses of these differently tuned neurons. This is thought to impair decoding—another coding catastrophe that would occur in the absence of adaptation (see, e.g., Benucci et al., 2013). However, it seems silly to prevent one coding catastrophe by a mechanism that produces another. Furthermore, there might not be a correlation-driven coding catastrophe in the first place. If correlations impair decoding, then removing them should improve it. However, when Zavitz et al. (2016) explicitly removed all correlations before training their decoder, this

did not improve decoding. Thus, it appears unlikely that adaptation serves to improve decoding performance.

Another important proposal regarding the functional benefits of adaptation involves the idea that it acts as a "graphic equalizer" (e.g., Andrews, 1964; Ullman and Schechtman, 1982; Anstis et al., 1998; Kohn, 2007; Webster, 2011; Haak et al., 2014a). This classic idea derives from luminance adaptation in the retina, where it reflects a form of dynamic range adjustment that allows us to see optimally in very dark or bright environments. Benucci et al. (2013) have recently shown that equalization also occurs in the context of orientation adaptation in primary visual cortex. It is unfortunate that Zavitz et al. (2016) have not explored this possible benefit because it would provide important additional insights into the balance between the costs and benefits of adaptation in natural vision.

In conclusion, the study of Zavitz et al. (2016) provides a nice demonstration of some of the possible costs and benefits of sensory adaptation under relatively natural conditions. An important direction of future research is to shed light on how the brain deals with the costs of adaptation, if at all.

References

Andrews DP (1964) Error-correcting perceptual mechanisms. Q J Exp Psychol 16:104–115. CrossRef

Anstis S, Verstraten FA, Mather G (1998) The motion aftereffect. Trends Cogn Sci 2:111–117. CrossRef Medline

Barlow HB, Földiák P (1989) Adaptation and decorrelation in the cortex. In: The computing neuron (Durbin R, Miall C, Mitchinson G, eds), pp 52–79. New York: Addison-Wesley.

Barlow HB, Hill RM (1963) Evidence for a physiological explanation for the waterfall phenomenon and figural aftereffects. Nature 200:1345–1347. CrossRef Medline

Benucci A, Saleem AB, Carandini M (2013) Adaptation maintains population homeostasis in primary visual cortex. Nat Neurosci 16:724–729. CrossRef Medline

Clifford CW, Wenderoth P, Spehar B (2000) A functional angle on some after-effects in cortical vision. Proc Biol Sci 267:1705–1710. CrossRef Medline

Dhruv NT, Carandini M (2014) Cascaded effects of spatial adaptation in the early visual system. Neuron 81:529–535. CrossRef Medline

Haak KV, Fast E, Baek Y, Mesik J (2014a) Equalization and decorrelation in primary visual cortex. J Neurophysiol 112:501–503. CrossRef Medline

Haak KV, Fast E, Bao M, Lee M, Engel SA (2014b) Four days of visual contrast deprivation reveals limits of neuronal adaptation. Curr Biol 24:2575–2579. CrossRef Medline

Kohn A (2007) Visual adaptation: physiology, mechanisms, and functional benefits. J Neurophysiol 97:3155–3164. CrossRef Medline

Levinson E, Sekuler R (1976) Adaptation alters perceived direction of motion. Vision Res 16: 779–781. CrossRef Medline

Patterson CA, Duijnhouwer J, Wissig SC, Krekelberg B, Kohn A (2014) Similar adaptation effects in primary visual cortex and area MT of the macaque monkey under matched stimulus conditions. J Neurophysiol 111:1203–1213. CrossRef Medline

Powell G, Bompas A, Sumner P (2012) Making the incredible credible: afterimages are modulated by contextual edges more than real stimuli. J Vis 12(10):17. CrossRef Medline

Schwartz O, Hsu A, Dayan P (2007) Space and time in visual context. Nat Rev Neurosci 8:522–535. CrossRef Medline

Seriès P, Stocker AA, Simoncelli EP (2009) Is the homunculus "aware" of sensory adaptation? Neural Comput 21:3271–3304. CrossRef Medline

Ullman S, Schechtman G (1982) Adaptation and gain normalisation. Proc R Soc Lond B Biol Sci 216:299–313. CrossRef Medline

Wainwright MJ (1999) Visual adaptation as optimal information transmission. Vision Res 39:3960–3974. CrossRef Medline

Webster MA (2011) Adaptation and visual coding. J Vis 11(5):3. CrossRef Medline

Webster MA (2015) Visual adaptation. Annu Rev Vis Sci 1:547–567. CrossRef Medline

Zavitz E, Yu HH, Rowe EG, Rosa MG, Price NS (2016) Rapid adaptation induces persistent biases in population codes for visual motion. J Neurosci 36:4579–4590. CrossRef Medline