Vision Research 104 (2014) 68–79

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Probing the functions of contextual modulation by adapting images rather than observers

Michael A. Webster

Department of Psychology/296, University of Nevada, Reno, Reno, NV 89557, USA

ARTICLE INFO

Article history:
Received 30 April 2014
Received in revised form 21 August 2014
Available online 2 October 2014

Keywords:
Adaptation
Normalization
Contextual modulation
Color vision
Color constancy
Natural images

ABSTRACT

Countless visual aftereffects have illustrated how visual sensitivity and perception can be biased by adaptation to the recent temporal context. This contextual modulation has been proposed to serve a variety of functions, but the actual benefits of adaptation remain uncertain. We describe an approach we have recently developed for exploring these benefits by adapting images instead of observers, to simulate how images should appear under theoretically optimal states of adaptation. This allows the long-term consequences of adaptation to be evaluated in ways that are difficult to probe by adapting observers, and provides a common framework for understanding how visual coding changes when the environment or the observer changes, or for evaluating how the effects of temporal context depend on different models of visual coding or the adaptation processes. The approach is illustrated for the specific case of adaptation to color, for which the initial neural coding and adaptation processes are relatively well understood, but can in principle be applied to examine the consequences of adaptation for any stimulus dimension. A simple calibration that adjusts each neuron's sensitivity according to the stimulus level it is exposed to is sufficient to normalize visual coding and generate a host of benefits, from increased efficiency to perceptual constancy to enhanced discrimination. This temporal normalization may also provide an important precursor for the effective operation of contextual mechanisms operating across space or feature dimensions. To the extent that the effects of adaptation can be predicted, images from new environments could be "pre-adapted" to match them to the observer, eliminating the need for observers to adapt.

1. Introduction

While this feature issue is primarily devoted to the functions of spatial context, in the present article I approach the question of contextual modulation by focusing instead on temporal context – on processes that shape a neuron's response that are outside the classical receptive field in time rather than space. The effects of context show a number of parallels between space and time, and appear to modulate sensitivity and perception in similar ways, possibly because the statistics of the visual world are themselves similar in space and time (Schwartz, Hsu, & Dayan, 2007). It would therefore be surprising if spatial and temporal contextual processes were not designed to achieve similar goals. For example, both adaptation and spatial contrast interactions have long been argued to play an important role in color constancy (Brainard & Wandell, 1992; Hurlbert & Wolf, 2004; Land, 1986). A great deal of thinking has been invested in understanding why sensory systems adapt (for recent reviews see (Clifford et al., 2007; Kohn, 2007; Wark, Lundstrom, & Fairhall, 2007; Webster, 2011; Webster & MacLeod, 2011)). Reviewing the functional consequences of temporal context may thus shed light on some of the places to look for analogous roles of spatial context. Alternatively, this can also point to places where they might serve different goals. If temporal context already optimizes some aspects of visual coding, what roles are left for spatial context? And given that both are occurring, how might they interact? To examine these questions I begin by reviewing some of the purposes that have been suggested for adaptation, and then describe a recent approach we have developed for testing different hypothesized functions by "adapting images" to simulate theoretically optimal states of adaptation. This is followed by briefly considering the implications of these temporal processes for mechanisms that respond to spatial context. While the simulations are focused only on the consequences of adaptation for color vision, the principles are general and thus applicable to visual coding in general.

2. The functions of adaptation

Most reviews of visual adaptation include a litany of potential benefits. For example, a typical list might include the following

---

E-mail address: mwebster@unr.edu

http://dx.doi.org/10.1016/j.visres.2014.09.003

functions of adaptation, along with the ways in which these might be manifest in our perception:

1. *Sensitivity regulation:* Neurons have a very limited dynamic range of their response but must operate over a potentially enormous range of inputs. Adaptation adjusts sensitivity so that the responses are centered around the average stimulus level in the environment (Barlow, 1972). This allows the visual system to devote the full signaling capacity to registering small changes around the mean, where most of the information is concentrated. Behaviorally this shows up as a better ability to distinguish or discriminate stimuli within a scene when we are first adapted to the scene.
2. *Coding efficiency:* The limited capacity of neurons to carry information also exerts strong pressure to optimize efficient coding (Clifford et al., 2007; Wainwright, 1999; Wark, Lundstrom, & Fairhall, 2007). This can act at the level of an individual neuron to adjust its operating curve so that each response level is given equal weight (e.g. histogram matching) (Laughlin, 1981). It can also operate across neurons to remove redundancies in their responses. Adaptation has been proposed to play a role in both of these adjustments. For example, adaptation may adjust the gain of the neural response so that the average response occurs for the average level of the current stimulus gamut (Ohzawa, Sclar, & Freeman, 1982; Rieke & Rudd, 2009); and may potentially remove the correlations between neurons coding different attributes of the stimulus if those attributes are themselves correlated (Barlow, 1990b; Carandini et al., 1997). Signs of these benefits are typically searched for in the neural code rather than the performance of the observer, but should again reflect behavioral benefits in sensitivity and discrimination.
3. *Error correction:* Neurons might also be thought of as mismatched to the stimulus if they signal the "wrong" percept – if the world consistently looks tilted or too yellow. In this case adaptation might serve to recalibrate the neural response to remove the error (Andrews, 1967). For instance, if there are inherent differences in sensitivity to different orientations, then this might introduce biases in the population code for orientation. Adaptation can remove these biases by equating the responses and thus leveling the playing field. Perceptual signs of these corrections are seen in the stability of some percepts despite large sensitivity differences within or between observers.
4. *Perceptual constancy:* Often the visual system is trying to estimate invariant properties of the world from a retinal image in which multiple sources of stimulus variation are confounded. A classic example is color constancy, where the goal is to recover the reflectance of a surface from a spectrum that also varies with the lighting (Foster, 2011; Smithson, 2005). Adaptation can promote constancy by filtering out or discounting some of the sources of variation. Thus when the lighting becomes redder, adapting to become less sensitive to red can remove much of the color shift owing to the illuminant. Here we expect to see advantages of adaptation when we can recognize the same stimulus under different viewing contexts.
5. *Learning and predictive coding:* An efficient way to represent information is to build a prediction about the expected properties of the world and then signal only the deviations from this prediction (Srinivasan, Laughlin, & Dubs, 1982). In this way the generic state of the world need only be represented implicitly. Adaptation could adjust visual coding to null the responses to the expected level. Mechanistically, this should show up as stronger responses the more novel or unexpected the stimulus is (Ranganath & Rainer, 2003). In turn, perceptually this should be manifest so that novel stimuli are more salient, and so that what we notice most about the world are the very properties we are not adapted to.

This list is not exhaustive, but again gives a flavor of the variety of roles that adaptation has been hypothesized to play in sensory systems. Moreover, many of these putative roles are closely related, and may amount to different perspectives on the same problem. However, several issues continue to plague our understanding of the actual purpose of adaptation. First, the wide variety of proposed functions appears to run counter to the fact that, in many cases, it may only be a single type of adjustment that the visual system is making. How can one trick yield so many and so seemingly disparate advantages? This suggests that there is probably some more fundamental and general principle guiding why the visual system adapts, and what we now consider the functions of adaptation will turn out to be specific consequences of this principle. Here I explore the possibility that this principle is "normalization," and reflects the simple tendency of adaptation to balance the responses across the mechanisms coding a stimulus dimension so that they are either nulled or equated within the visual context we are currently adapted to. In the following I use normalization specifically to refer to a form of adaptation, though importantly the term has also taken on a broader meaning which includes both temporal and spatial calibrations (Carandini & Heeger, 2011). Norms are central to models of visual coding (Webster, 2011). In some cases the visual system uses an explicit "norm-based" code in which stimuli are represented relative to a reference or norm which itself is encoded by a null in the neural response and which has a special, "neutral" status in visual coding. Examples of such codes include color vision (where all hues are referenced to gray) and face perception (where all identities are referenced to the average or prototypical face) (Webster & MacLeod, 2011). In other cases, stimuli may be represented by a population code – by something like the peak response in a distribution of neurons tuned for different levels of the stimulus dimensions. Examples in this case include the encoding of orientation or spatial scale. Here the norm is implicit but corresponds to equal activity (no peak) across the set of mechanisms.

The function of adaptation, quite simply, may be to set these norms according to the current context, and can be accomplished simply by adjusting each neuron's gain so that the mean response occurs for the mean stimulus level each neuron is exposed to. There are likely to be many additional forms of adaptation, e.g. to match more complex characteristics of the input (Gollisch & Meister, 2010), or to decorrelate the responses of different neurons (Barlow, 1990b), but here I consider what a visual system can do with only a simple gain change. This modulation alone would assure in theory that the responses of all mechanisms would be the same on average for the current adapting context (i.e. the stimulus distribution we are adapted to). Moreover, this recalibration proceeds naturally whenever the context changes, and thus will always lead to the appropriate norms for the appropriate context. Thus norms are not distinct from adaptation, but rather are synonymous with the states of adaptation that the visual system is currently in. In this sense, the phrase "contextual modulation" is an understatement – context does not merely perturb neural responses, it defines them.

As we will see, it turns out that this normalization predicts each of the functional benefits highlighted above. Yet a second general problem in understanding the purpose of adaptation has been that the behavioral correlates of these functional improvements are often lacking. Strong visual aftereffects can be readily induced for most visual patterns, yet adapting to these patterns often fails to improve visual performance. Specifically, observers are not typically better at detecting or discriminating patterns after they have adapted to them (even though that adaptation leads to large changes in the appearance of the patterns) (Clifford et al., 2007). The clearest exception is light adaptation, where adjusting to the