Vision: Efficient Adaptive Coding

David Burr[1,2,*] and **Guido Marco Cicchini**[2]

[1]Department of Neuroscience, Psychology, Pharmacology and Child Health, University of Florence, Via S. Salvi 12, Florence, Italy

[2]Neuroscience Institute, National Research Council, Via Moruzzi 1, Pisa 56124, Italy

Abstract

Recent studies show that perception is driven not only by the stimuli currently impinging on our senses, but also by the immediate past history. The influence of recent perceptual history on the present reflects the action of efficient mechanisms that exploit temporal redundancies in natural scenes.

Did you notice how Harry Potter's T-shirt changes from a crewneck to a henley shirt in *The Order of the Phoenix*, or how Julia Roberts' croissant inexplicably morphs into a pancake in *Pretty Woman*? Do not worry if you did not: such failure to notice blatant continuity errors may reflect the operation of our highly efficient perceptual systems, which adapt to the redundancies of the world. Recent work is showing how our perceptual systems exploit the temporal redundancies of natural scenes, particularly the fact that objects tend to be constant, rarely changing abruptly from one moment.

Many perceptual studies show how much detail in the world escapes our awareness. The most well known are the stunning demonstrations of 'change blindness' [1,2]: when motion transients are masked, subjects fail to see huge changes in successive scenes, such as the disappearance of aeroplane engines (see examples in http://www.gocognitive.net/demo/change-blindness). Another example is 'motion silencing' [3]: failure to see large changes in form or colour of groups of moving shapes. A newly developed approach complements these studies to shed more light on why changes may go unnoticed: the system seems to exploit temporal redundancies of natural scenes by integrating information over tens of seconds, smoothing the unavoidable random fluctuations in their neural representation.

As they reported recently in *Current Biology*, Liberman *et al.* [4] asked subjects to report the identity of a series of sequentially presented faces. They showed that perceived identity (along a morphed continuum) does not just depend on the currently viewed face, but is strongly biased towards the identity of recently viewed faces. A series of clever and important control studies showed that the effects cannot be attributed to response hysteresis, and are definitely perceptual, not simply decision biases.

---

[*]dave@in.cnr.it.

This new study on faces — important because of their complexity and behavioural relevance — builds on previous work using more basic stimuli, such as patches of gratings [5] or clouds of dots [6]. The perception of orientation — a basic visual property probably mediated largely by the orientation-selective neural mechanisms of primary visual cortex [5] — is robustly biased towards the orientation of patches presented prior to the current stimulus, and even those two or three items back. Our own group [6] has recently reported significant serial dependencies in mapping number onto space. These dependencies can be strong enough to cause compression of the subjective spatial representation of number — the mental *number line* — previously thought to reflect logarithmic encoding of number [7]. Given that the non-linearities in the number line are predictors of poor mathematical performance in schoolchildren [8], understanding the mechanisms causing the non-linearities in the number line is of fundamental importance.

Serial effects in vision are not new. For example, *repetition priming* is a well-known phenomenon, where words, faces, and many objects are recognized more rapidly the second time around [9,10]. Priming is particularly robust in attention-based studies, such as of 'pop-out' effect [11], where repeated use of the same cueing colour greatly improves performance. Another example is *motion priming*, where ambiguous motion can be disambiguated by viewing motion with a clear direction. These priming effects can be quite subtle: viewing a sequence of motion events of alternating direction causes the ambiguous test to be seen to complete the pattern of alternation, suggesting that priming involves active prediction [12]. What is novel about the recent studies [4–6] is that they show that past experience not only speeds recognition and discrimination, but can warp perception itself. Most importantly, this warping of our perception is paradoxically beneficial, leading to more efficient perception: it may even be an optimal perceptual strategy.

The key to understanding how *misperception* of stimuli can be an efficient strategy is that the system is exploiting to advantage natural temporal redundancies. In the real world, objects tend to remain constant over time: croissants tend to remain croissants, not morph miraculously into pancakes. Thus, it is a good bet that transitory changes in stimuli may result from noisy neural fluctuations rather than actual changes in the external world.

One way to incorporate past history into the present would be to perform a rolling average, which would damp to some extent the neural fluctuations. But a more interesting idea — which has a good deal of support [12] — is that the past may play an active role in *predicting* the future [13,14], rather than simply merging with it. As the world tends to remain constant, a good prediction of the future is that it will be similar to the present. In these models, new data samples do not stand alone, but are used to update the prediction. A convenient tool to model predictive behaviour is the *Kalman filter*, a standard engineering algorithm that reduces noise by combining instantaneous estimates of a system state with a weighted prediction from past data. Although the model is simple, both conceptually and in its implementation, it predicts well many of the reported features of serial dependency [6].

Figure 1 shows an example of how predictive serial dependency can reduce noise. It assumes that a particular dimension — orientation, face identity and so on — tends to be constant, but can change abruptly (thin black trace). The signal will be corrupted by noise, which

varies depending on conditions (such as lighting). The red trace shows the output of the recursive Kalman filter, which is closer to the veridical signal than is the raw (grey) trace (see the root-mean square errors in Figure 1B). Importantly, the Kalman filter is adaptive, changing on-line the weighting given to the past, to maximize efficiency. For example, the predictive (serial-dependent) component is automatically down weighted by the squared difference between current and previous trials, so large variations in stimulus dimensions are attributed to reality rather than noise: prior history has little influence at these points. The lower (blue) trace of Figure 1A shows the action of a simple low-pass filter. Although this smoothes out some noise, it weakens the response to real changes in the stimulus, and overall fails to reduce noise.

Although very successful, predictive coding cannot account for all serial perceptual phenomena. For example, after viewing downward motion, such as a waterfall, for some time, stationary objects seem to move upwards, the *opposite* direction to the prior stimulus [15]. *Negative aftereffects* like these go in the opposite direction to the serial dependencies discussed above. Empirically, the two are easily dissociated: negative aftereffects occur after relatively prolonged exposure to very salient stimuli, are well localized spatially, often in retinotopic coordinates, and are largely independent of attention; serial dependencies are strongest for weak stimuli, have weak spatial localization (largely spatiotopic), and are highly dependent on attention [5,16].

Again we should ask, what are the functional roles of these opposing forms of serial dependency? Most likely they trade off competing perceptual requirements. In space perception, there exists a well-known trade-off between the need to integrate to reduce noise, and to segregate to perceive small objects. Similar requirements occur over time, and are dealt with at many levels. For example, all sensory systems tend to have two classes of receptors, slow-adapting or tonic receptors that integrate over time, and fast-adapting or phasic receptors that respond to change. Serial dependencies and negative aftereffects may continue this dichotomy at higher levels of analysis: positive dependencies cause an adaptive form of integration, reducing noise by exploiting temporal consistencies; negative aftereffects aid segregation, amplifying small but potentially important changes over time.

Interestingly, like serial dependencies, adaptation effects also exploit temporal correlations, allocating resources to the representation of the parameter values in the vicinity of the adaptor at that point in time [17]. Even complex contingency aftereffects are thought to work by reducing temporal correlations in multi-dimensional cortical signals [18], maximizing the efficiency of low bandwidth neural systems. Negative adaptation is clearly also important for calibrating perceptual systems to the prevailing conditions [19]. However, it remains very much an open question exactly how these two opposing effects interact.

The studies discussed here provide further evidence for how neural perceptual systems maximize efficiency by exploiting temporal properties of natural scene statistics. As objects do not normally morph spontaneously from croissants to pancakes, averaging recent perceptual history with the current percept reduces noise and helps stabilize what we see over time, leading to a more robust, predictable and stable sensation of the world. This is clearly a fundamental perceptive mechanism. It would be important to examine whether

malfunction of these mechanisms could be a contributing factor to phenomena such as the 'sensory overload' experienced by some people, including those with autism [20].

References

1. Rensink RA, O'Regan JK, Clark JJ. To see or not to see: the need for attention to perceive changes in scenes. Psychol Sci. 1997; 8:368–373.

2. Simons DJ, Levin DT. Change blindness. Trends Cogn Sci. 1997; 1:261–267. [PubMed: 21223921]

3. Suchow JW, Alvarez GA. Motion silences awareness of visual change. Curr Biol. 2011; 21:140–143. [PubMed: 21215632]

4. Liberman A, Fischer A, Whitney D. Serial dependence in the perception of faces. Curr Biol. 2014; 24:2569–2574. [PubMed: 25283781]

5. Fischer J, Whitney D. Serial dependence in visual perception. Nat Neurosci. 2014; 17:738–743. [PubMed: 24686785]

6. Cicchini GM, Anobile G, Burr DC. Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. Proc Natl Acad Sci USA. 2014; 111:7867–7872. [PubMed: 24821771]

7. Dehaene S. The neural basis of the Weber-Fechner law: a logarithmic mental number line. Trends Cogn Sci. 2003; 7:145–147. [PubMed: 12691758]

8. Siegler RS, Booth JL. Development of numerical estimation in young children. Child Dev. 2004; 75:428–444. [PubMed: 15056197]

9. Kristjansson A, Campana G. Where perception meets memory: a review of repetition priming in visual search tasks. Atten Percept Psychophys. 2010; 72:5–18. [PubMed: 20045875]

10. Ellis AW, Young AW, Flude BM, Hay DC. Repetition priming of face recognition. Q J Exp Psychol A. 1987; 39:193–210. [PubMed: 3615943]

11. Maljkovic V, Nakayama K. Priming of pop-out: I. Role of features. Mem Cognit. 1994; 22:657–672.

12. Maloney LT, Dal Martello MF, Sahm C, Spillmann L. Past trials influence perception of ambiguous motion quartets through pattern completion. Proc Natl Acad Sci USA. 2005; 102:3164–3169. [PubMed: 15710897]

13. Rao RP, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci. 1999; 2:79–87. [PubMed: 10195184]

14. Friston K, Kiebel S. Predictive coding under the free-energy principle. Philos Trans R Soc Lond B. 2009; 364:1211–1221. [PubMed: 19528002]

15. Thompson P, Burr D. Visual aftereffects. Curr Biol. 2009; 19:R11–R14. [PubMed: 19138580]

16. Yoshimoto S, Uchida-Ota M, Takeuchi T. The reference frame of visual motion priming depends on underlying motion mechanisms. J Vis. 2014; 14(1):10. [PubMed: 24413393]

17. Gepshtein S, Lesmes LA, Albright TD. Sensory adaptation as optimal resource allocation. Proc Natl Acad Sci USA. 2013; 110:4368–4373. [PubMed: 23431202]

18. Barlow, HB.; Földiák, P. Adaptation and decorrelation in the cortex. The Computing Neuron. Miall, C.; Durbin, RM.; Mitchinson, GJ., editors. Wokingham: Addison-Wesley; 1989.

19. Chopin A, Mamassian P. Predictive properties of visual adaptation. Curr Biol. 2012; 22:622–626. [PubMed: 22386314]

20. Pellicano E, Burr D. When the world becomes 'too real': a Bayesian explanation of autistic perception. Trends Cogn Sci. 2012; 16:504–510. [PubMed: 22959875]
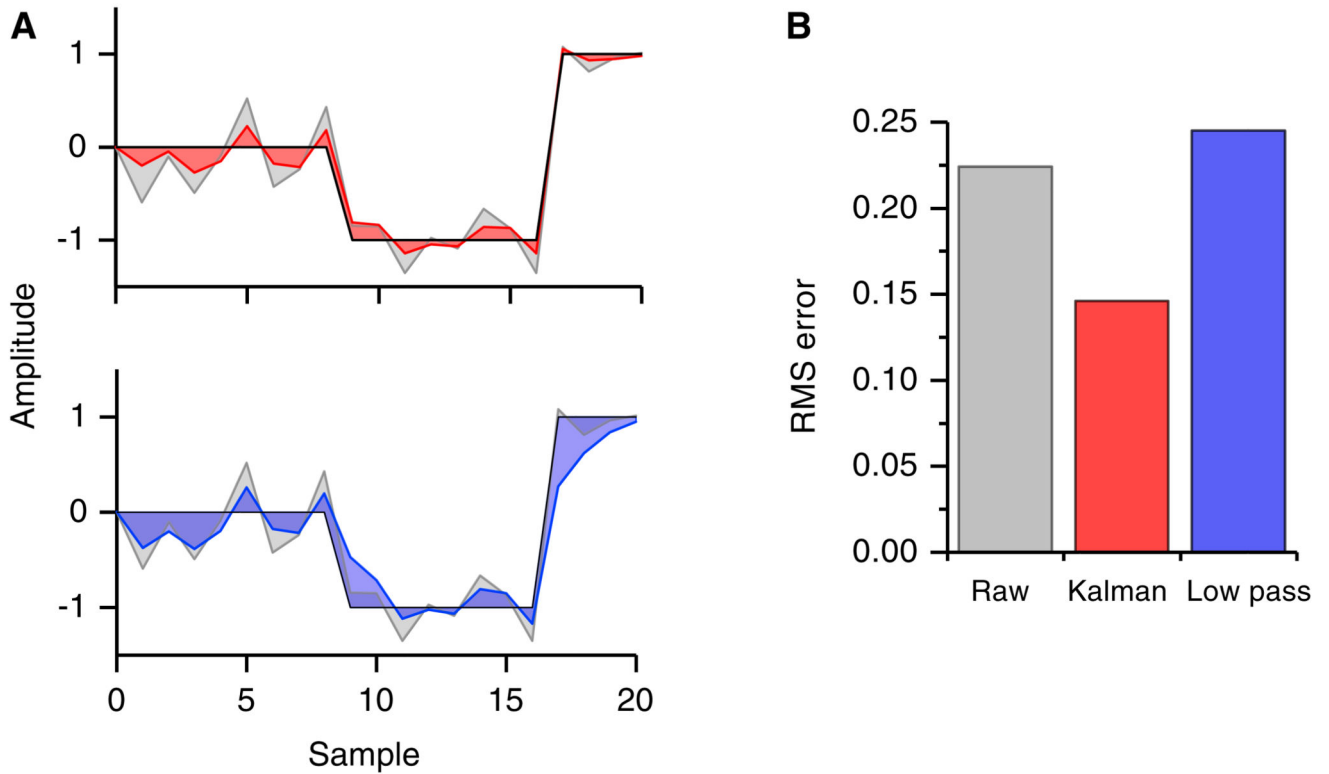
Figure 1. Illustration of how serial dependencies, modelled by the Kalman filter, can reduce noise and improve accuracy [6].

(A) Black lines show the changes in the physical dimension, grey the signal with added gaussian noise of space constant 0.3 in the high section, 0.2 in the low section. The red trace in the upper plot shows the output of the Kalman, clearly closer to the veridical stimulus than the raw noisy signal. The blue, lower trace shows a simple low-pass unadaptive filter, performing poorly, particularly near the edge. (B) Root mean squared (RMS) error from the real signal for the raw trace, Kalman filter and low-pass smoothing.