

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.

Priority Maps Explain the Roles of Value, Attention, and Salience in Goal-Oriented Behavior

P. Christiaan Klink,^{1,2,3} Pia Jentgens,² and Jeannette A.M. Lorteije²

¹Department of Neuromodulation and Behaviour, and ²Department of Vision and Cognition, Netherlands Institute for Neuroscience, 1105 BA Amsterdam, The Netherlands, and ³Department of Psychiatry, Academic Medical Center, University of Amsterdam, 1070 AW, Amsterdam, The Netherlands
Review of Chelazzi et al.

What makes us tick? All behavior is essentially goal-oriented, and while these goals may vary from unconsciously generated urges to well conceived intentions, our choice of action generally involves a comparison of options with their associated costs and gains. This decision-making mechanism has been extensively studied in the visuomotor system, where the available visual information in a scene typically overflows the capacity of the visual system, urging a selection of scene elements for preferential processing.

Some elements may intrinsically attract attention because they stand out from their surround. The global conspicuity pattern of a scene is captured in a stimulus-salience map that provides a stimulus-driven prioritization of targets for the attraction of attention (Itti and Koch, 2001). Neural processing of visual information is additionally influenced by cognitive top-down factors like expectation, memory, or selective attention. Together, these cognitive factors create an internal motivational salience map that interacts with the stimulus-based salience

map to form a priority map that guides visuomotor behavior (Fecteau and Munoz, 2006). To illustrate this principle, consider a game of whack-a-mole, in which a player is challenged to, as fast as possible, hit a mole figure that pops up in random locations. While the popping-up is a salient event in itself, the player's response speed is hindered by having to spread attention across possible pop-up locations. If a player focuses more on one location, performance for this location may improve, but at a cost of performance for other locations.

A recent study by Chelazzi et al. (2014) addresses the formation of spatial priority maps by investigating the influence of learned space–reward associations on spatially confined discrimination performance. Observers identified one or two potentially presented characters or digits from among a briefly flashed array of eight stimuli, followed by a mask. Discrimination performance for each of the eight potential target locations was used as a baseline representing individual observers' spatial priority maps. Observers then trained with a visual search task in which they identified a single target among seven distractors as fast as possible. The same stimulus locations were used as in the discrimination task, but crucially, the probability of receiving a high instead of a low monetary reward for correct target discrimination varied with stimulus location. For two locations, the odds for receiving a high over a low reward were

80/20, for two others it was 20/80, while for the four remaining locations it was equally likely to get a high or low reward. The differently rewarded locations were placed so that the high reward locations were in a different hemifield than the low reward locations, while both hemifields also contained two 50/50 locations. Observers were not informed of these spatial biases in reward probability and while they did receive reward-related feedback after each trial, the task was deliberately made challenging to keep them from noticing the reward pattern.

When observers performed the discrimination task a second time, 4 d after training, their performance for target discrimination was not only better in general, it was also affected by the location–reward associations from the training phase. If two targets in the flashed stimulus array were presented at locations previously associated with different rewards, observers were more likely to only report the target at the high-reward location and miss the one at the low-reward location. This deviation from baseline was stronger if the difference in associated reward for the two locations was larger. It was furthermore specific for stimulus locations and did not generalize to complete hemifields, suggesting high spatial specificity. Performance in the one-target conditions showed a general improvement in performance, but no reward-related effects.

What does this mean? Since the same stimuli are used before and after training,

Received Aug. 6, 2014; accepted Sept. 1, 2014.

P.C.K. is supported by NWO-VENI Grant 451.13.023. J.A.M.L. is supported by Grant 823.02.010 of the NWO-ALW open program and by the European Union Seventh Framework Program (Project 269921 BrainScale5).

Correspondence should be addressed to P. Christiaan Klink, Netherlands Institute for Neuroscience, Royal Netherlands Academy of Arts and Sciences, Meibergdreef 47, 1105 BA Amsterdam, The Netherlands. E-mail: C.Klink@nin.knaw.nl.

DOI:10.1523/JNEUROSCI.3249-14.2014

Copyright © 2014 the authors 0270-6474/14/3413867-03\$15.00/0

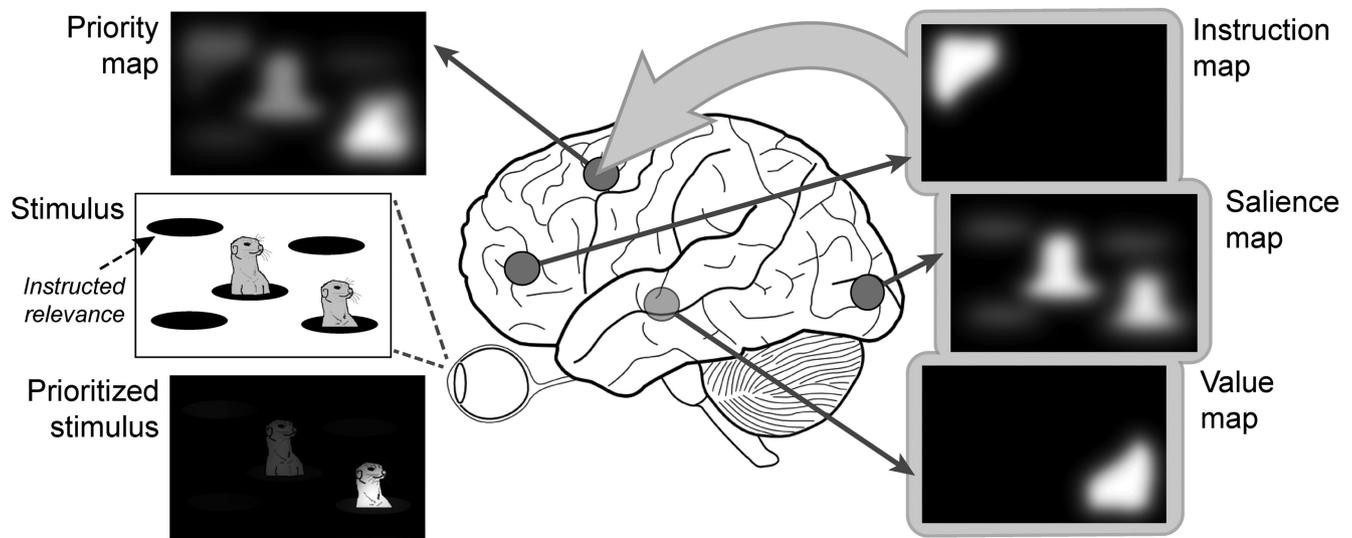


Figure 1. Priority maps in different parts of the brain can be combined to highlight elements of a visual stimulus and guide attention. Schematically depicted are a stimulus-driven saliency map in visual cortex, a learned value-based map in the midbrain (here favoring the bottom right corner), an instruction-based map in frontal cortex (“attend top left”), and a resulting combined priority map in the frontal eye fields.

the stimulus-driven saliency maps should not be different. The internal saliency map, however, appears to be altered by the newly learned space–reward associations, resulting in a value-adjusted overall spatial priority map. The fact that no differential effects were observed with single targets suggests that the stimulus saliency map still accurately attracts attention to each target location. However, when two targets compete for attention, the value-adjusted priority map biases this competition toward the highly rewarded location, and with such brief stimulus presentations, the second target goes unnoticed.

Since the stimuli in the discrimination and training tasks were completely different and the discrimination task lacked differential rewards, the effects of training should be purely driven by a lasting reevaluation of spatial locations. The type of task was, however, similar across the test and training phases. If the effect of training is truly a general spatial prioritization, it should also show up in completely different tasks or using other read-outs, e.g., in reduced reaction times while detecting targets at prioritized locations during visual search. It is furthermore striking that the conditioned priority map persists for at least 4 d after training without reinforcement of the new value distribution. It would be interesting to investigate the persistence of learned priority maps and their specificity to the contextual environment. Do these maps slowly fade over time, or is an explicit recalibration with new space–reward associations needed? Do the priorities generated during train-

ing in the lab generalize to other stimulus arrays and environments? Given our natural tendency to move our eyes, one may also wonder what kind of spatial reference frame would be best suited to encode learned priority maps. While a retinotopic reference frame does not immediately seem appropriate for such long-lasting biases, there is no way to dissociate retinotopic and spatiotopic contributions in the current study, because for stationary, fixating observers these reference frames coincide.

The effects of reward and attention on visuomotor behavior and neuronal activity have been notoriously difficult to disentangle (Maunsell, 2004; Stănișor et al., 2013). For visual priority maps, two different forms of attention can be discerned. First, the goal-oriented attraction of processing resources toward the visual elements that are selectively highlighted by the priority map is a direct behavioral consequence (effect) of this priority map. Selective attention, however, can be thought of as a top-down determining factor (cause) in the formation of the internal motivational saliency map. This distinction fits the recent suggestion of attention as a consequence of value-based decision-making mechanisms (Stănișor et al., 2013) rather than an independent cause for the filtering of sensory information (Krauzlis et al., 2014). Goal-directed behavior could thus be guided by a context-dependent priority map that combines acute internally generated value-based priorities (e.g., selective attention, task instructions), more slowly learned value-based

priorities (e.g., reinforcement learning), and stimulus salience.

Where and how are priority maps represented in the brain? Chelazzi et al. (2014) give an overview of brain regions that either reflect a priority map, or are involved in selective attention and reward processing. From this overview, it seems that priority maps emerge from a distributed network involving the midbrain, hippocampus, and the frontal and parietal cortices. Priority signals of different origins may be initiated in different parts of the brain (Fig. 1). The stimulus-driven saliency, for instance, could come from early visual cortex and subcortical structures like the superior colliculus (Itti and Koch, 2001), while subconscious value-sets can be provided by evolutionary old structures like the hippocampus, amygdala, and striatum (Pennartz et al., 2011). Top-down instruction-related value-sets are more likely initiated in regions like the anterior cingulate and orbitofrontal cortex (Kennerley et al., 2011). The fact that different priority signals originate in different brain areas does not mean that their consequences cannot be measured elsewhere. It is more likely that interaction among different priority signals broadly influence neural activity. Furthermore, some value-sets may require more time to compute than others, causing the overall priority map to evolve as contributions of new value-sets become available to the system. The distributed priority network hypothesis also matches the idea of attention as an inescapable consequence of value-based prioritization (Krauzlis et al.,

2014). In that context, the priority maps that have been revealed in the frontal eye fields (Thompson and Bichot, 2005) and intraparietal cortex (Bisley and Goldberg, 2010) can be seen as behavioral planning maps that result from distributed value signals and guide covert attention shifts and eye-movements.

The general concept of priority maps does not restrict itself to the spatial dimension or the visuomotor system. A similar framework has recently been proposed as a novel way of looking at social interactions and dysfunctional behavior in psychiatric disorders (de Haan et al., 2013). The idea is that the choice for particular behaviors is driven by a combination of external possibilities for such behaviors (landscape of affordances) and a map of internal values that modulates the relative priorities of potential behaviors. Dysfunctional goal-oriented or social behavior as seen in psychiatric patients then arises from pathologically biased internal value maps that could give rise to obsessive compulsions, impulsivity, or social avoidance.

Chelazzi et al. (2014) provide support for a flexible, subconscious, value-based priority map that guides goal-oriented behavior, a concept that may extend well beyond the spatial dimension. Given the distributed underlying neural mechanisms, it seems likely that for each behavioral context, there will be different brain areas and networks contributing to the eventual priority map. It is up to future research to find out how these different priorities coexist and interact.

References

- Bisley JW, Goldberg ME (2010) Attention, intention, and priority in the parietal lobe. *Annu Rev Neurosci* 33:1–21. [CrossRef Medline](#)
- Chelazzi L, Eštočinová J, Calletti R, Lo Gerfo E, Sani I, Della Libera C, Santandrea E (2014) Altering spatial priority maps via reward-based learning. *J Neurosci* 34:8594–8604. [CrossRef Medline](#)
- de Haan S, Rietveld E, Stokhof M, Denys D (2013) The phenomenology of deep brain stimulation-induced changes in OCD: an enactive affordance-based model. *Front Hum Neurosci* 7:653. [CrossRef Medline](#)
- Fecteau JH, Munoz DP (2006) Saliency, relevance, and firing: a priority map for target selection. *Trends Cogn Sci* 10:382–390. [CrossRef Medline](#)
- Itti L, Koch C (2001) Computational modelling of visual attention. *Nat Rev Neurosci* 2:194–203. [CrossRef Medline](#)
- Kennerley SW, Behrens TE, Wallis JD (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci* 14:1581–1589. [CrossRef Medline](#)
- Krauzlis RJ, Bollimunta A, Arcizet F, Wang L (2014) Attention as an effect not a cause. *Trends Cogn Sci* 18:457–464. [CrossRef Medline](#)
- Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci* 8:261–265. [CrossRef Medline](#)
- Pennartz CM, Ito R, Verschure PF, Battaglia FP, Robbins TW (2011) The hippocampal–striatal axis in learning, prediction and goal-directed behavior. *Trends Neurosci* 34:548–559. [CrossRef Medline](#)
- Stănișor L, van der Togt C, Pennartz CM, Roelfsema PR (2013) A unified selection signal for attention and reward in primary visual cortex. *Proc Natl Acad Sci U S A* 110:9136–9141. [CrossRef Medline](#)
- Thompson KG, Bichot NP (2005) A visual salience map in the primate frontal eye field. *Prog Brain Res* 147:251–262. [CrossRef Medline](#)