CellPress

Visual attention mitigates information loss in small- and large-scale neural codes

Thomas C. Sprague[1], Sameer Saproo[2], and John T. Serences[1,3]

[1] Neurosciences Graduate Program, University of California San Diego, La Jolla, CA 92093-0109, USA
[2] Department of Biomedical Engineering, Columbia University, New York, NY, USA
[3] Department of Psychology, University of California San Diego, La Jolla, CA 92093-0109, USA

The visual system transforms complex inputs into robust and parsimonious neural codes that efficiently guide behavior. Because neural communication is stochastic, the amount of encoded visual information necessarily decreases with each synapse. This constraint requires that sensory signals are processed in a manner that protects information about relevant stimuli from degradation. Such selective processing – or selective attention – is implemented via several mechanisms, including neural gain and changes in tuning properties. However, examining each of these effects in isolation obscures their joint impact on the fidelity of stimulus feature representations by large-scale population codes. Instead, large-scale activity patterns can be used to reconstruct representations of relevant and irrelevant stimuli, thereby providing a holistic understanding about how neuron-level modulations collectively impact stimulus encoding.

Visual attention and information processing in visual cortex

Complex visual scenes contain a massive amount of information. To support fast and accurate processing, behaviorally-relevant information should be prioritized over behaviorally-irrelevant information (Figure 1). For example, when approaching a busy intersection while driving it is crucial to detect changes in your lane's traffic-light rather than one nearby to prevent a dangerous collision. This capacity for selective information processing, or selective visual attention, is supported by enhancing the amount of information that is encoded about relevant visual stimuli relative to the amount of information that is encoded about irrelevant stimuli. Importantly, understanding how relevant visual stimuli are represented with higher fidelity requires considering more than only the impact of attention on the response properties of individual neurons. Instead, examining activity patterns across large neural populations can provide insights into how different unit-level attentional modulations synergistically improve the quality of stimulus representations in visual cortex.

In the scenario above, neurons can undergo several types of modulation in response to the relevant light compared to one that is irrelevant: response amplitudes can increase (response gain), responses can become more

Corresponding authors: Sprague, T.C. (tsprague@ucsd.edu);
Serences, J.T. (jserences@ucsd.edu).

Keywords: vision; visual attention; stimulus reconstruction; information theory; neural coding..

1364-6613/

 http://dx.doi.org/10.1016/j.tics.2015.02.005

Glossary

Bit: unit of entropy (base 2).
Decoder: algorithm whereby a feature or features about a stimulus (orientation, spatial position, stimulus identity, etc.) is/are inferred from an observed signal (spike rate, BOLD signal). Typically, the signal is multivariate across many neurons/voxels, but in principle a decoder can use a univariate signal.
Dynamic range: the set of response values a measurement unit can take. An increase in the response gain of a unit will increase the range of possible response values, and this will increase its entropy.
Encoding model: a description of how a neuron (or voxel) responds across a set of stimuli (e.g., a spatial receptive field can be a good encoding model for many visual neurons and voxels, see Box 2).
Entropy: a measure of uncertainty in a random process, such as a coin flip or observation of a neuron's spike count. A variable with a single known value will have 0 entropy, whereas a fair coin would have >0 entropy (1 bit).
Feature space: after reconstruction using the IEM technique, data exist in feature space, with each datapoint being defined by a vector of values corresponding to the activation of a single feature-selective population response (e.g., orientation, spatial position); common across all participants and visual areas.
Inverted encoding model (IEM): when encoding models are estimated across many measurement units, it may be possible to use all encoding models to compute a mapping from signal space into feature space which allows reconstruction of stimulus representations from multivariate patterns of neural activity across the modeled measurement units (Box 2).
Multivariate: when analyses are multivariate, signals from more than one measured unit are analyzed; utilizing information about the pattern of responses across units rather than simplifying the data pattern by taking a statistic over the units (e.g., mean).
Mutual information: the amount of uncertainty about a variable (e.g., state of the environment) that can be reduced by observation of the state of another random variable (e.g., the voxel or the neuron's response).
Noise entropy: variability in one signal that is unrelated to changes in another signal.
Receptive field (RF): region of the visual field which, when visually stimulated, results in a response in a measured neuron or voxel (population RF, or pRF).
Tuning function (TF): the response of a neuron or voxel to each of several values of a feature, such as orientation or motion direction.
Signal entropy: variability in one signal that is related to changes in another signal.
Signal space: data as measured exist in signal space, with a dimension for each measurement unit (fMRI voxel, EEG scalp electrode, electrocorticography subdural surface electrode, animal single cell firing rate, or calcium signal); cannot be directly compared across individual subjects without a potentially suboptimal coregistration transformation.

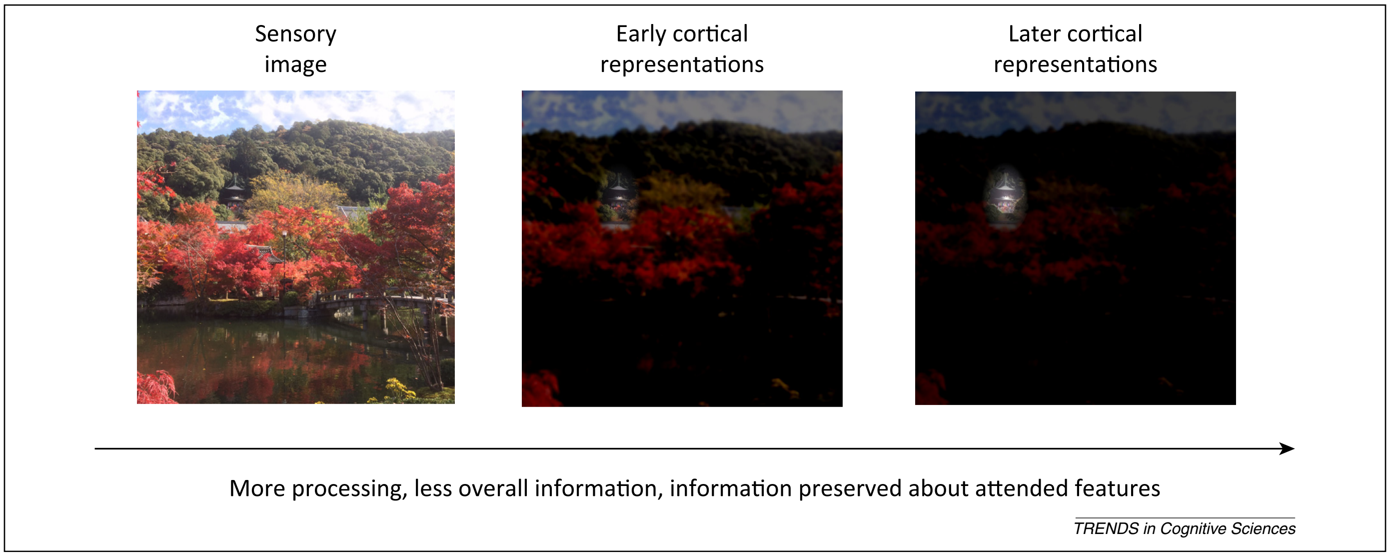

Figure 1. Attention filters behaviorally-relevant information. When viewing a complex natural scene (left), visual processing by a noisy neural system will necessarily result in an overall loss of information. If your eyes were fixated on the center of the image, but you were directing attention to the temple nestled among the trees near the top, information about the attended temple would be selectively preserved from degradation by noisy neural processing – such that, even at successively later stages of computation, information about the attended location and/or features of the image is still maintained, despite substantial loss of information about unattended components of the image (right panel).

reliable, and receptive field properties can shift (e.g., some neurons will shift their spatial receptive field to encompass the attended light). Thus, neural responses associated with attended stimuli generally have a higher signal-to-noise ratio and are more robust compared to responses evoked by unattended stimuli. Accordingly, the behavioral effects associated with visual attention are thought to reflect these relative changes in neural activity: when stimuli are attended, participants exhibit decreased response times, increased discrimination accuracy, and improved spatial acuity ([1–3] for reviews).

This selective prioritization of relevant over irrelevant stimuli follows from two related principles of information theory [4–7] (Box 1). First, the data-processing inequality [7] states that information is inevitably lost when sent via noisy communication channels, and that lost information cannot be recaptured via any amount of further processing. Second, the channel capacity of a communication system is determined by the amount of information that can be transmitted and received, and by the degree to which that information is corrupted during the process of transmission. In the brain, channel capacity is finite because there is a fixed (albeit large) number of neurons and because synaptic connections are stochastic such that information cannot be transmitted with perfect fidelity. Given this framework, different types of attention-related neural modulations can be viewed as a concerted effort to attenuate the unavoidable decay of behaviorally-relevant information as it is passed through subsequent stages of visual processing [8,9]. This framing also highlights the importance of understanding how attention differentially impacts responses across neurons, and, more importantly, how these modulations at the single-unit level interact to support population codes that are more robust to the information-processing limits intrinsic to the architecture of the visual system.

Box 1. Information content of a neural code

Information is related to a reduction in uncertainty [4,5,7]. A code is informative insofar as measurement of one variable (e.g., the firing rate of a single neuron) reduces uncertainty about another variable (e.g., feature of a stimulus). The amount of uncertainty in a random variable (e.g., the outcome of a coin toss or the spiking output of a cell) can be quantified by its entropy, which increases with increasing randomness. Mutual information (MI) is a measure of the reduction in uncertainty of one variable after knowing the state of another variable. MI would be zero for independent variables (e.g., two different coins), whereas MI would be high for two variables that strongly co-vary.

If a neuron noisily responds at the same level to each feature value, then the MI between the state of the stimulus and the state of the neuron's response is low because signal entropy (variability associated with changes in the stimulus) is low and noise entropy (variability unrelated to changes in the stimulus) is high (Figure IA). Instead, if the neuron exhibits a Gaussian-like orientation tuning function (TF; Figure IB,C), then MI is higher because more of the variability in the neuron's response is related directly to changes in the state of the stimulus. In this latter case, if the amplitude of the neuronal TF increases while noise remains approximately constant, then the ratio of signal entropy to noise entropy increases, resulting in greater MI between the neuron's response and the stimulus orientation. However, if the tuning width of the orientation TF changes, this could result in either an increase or decrease in the information about the stimulus, and would be contingent upon several factors such as the original tuning width, noise structure, dimensionality of the stimulus, and the responses of other neurons (Figure IF–H) [37,126–128]. For a widely-tuned neuron, a decrease in tuning width would result in an increase in signal entropy relative to noise entropy, increasing the information content of the neuron about orientation. At the other extreme, for a neuron perfectly tuned for a single stimulus value, with noisy baseline responses to other values, a broadening in tuning would result in greater variability associated with stimulus features, and consequently greater information (Figure IF–H). Thus, an increase in the amplitude of a neural response (under simple noise models) will increase the dynamic range and entropy, whereas a change in tuning width can either increase or decrease the information content of a neural code.