Early recurrent feedback facilitates visual object recognition under challenging conditions

Dean Wyatte[1], David J. Jilk[2]* and Randall C. O'Reilly[1]

[1] Department of Psychology and Neuroscience, University of Colorado Boulder, Boulder, CO, USA
[2] eCortex, Inc., Boulder, CO, USA

Standard models of the visual object recognition pathway hold that a largely feedforward process from the retina through inferotemporal cortex leads to object identification. A subsequent feedback process originating in frontoparietal areas through reciprocal connections to striate cortex provides attentional support to salient or behaviorally-relevant features. Here, we review mounting evidence that feedback signals also originate within extrastriate regions and begin during the initial feedforward process. This feedback process is temporally dissociable from attention and provides important functions such as grouping, associational reinforcement, and filling-in of features. Local feedback signals operating concurrently with feedforward processing are important for object identification in noisy real-world situations, particularly when objects are partially occluded, unclear, or otherwise ambiguous. Altogether, the dissociation of early and late feedback processes presented here expands on current models of object identification, and suggests a dual role for descending feedback projections.

Keywords: object recognition, feedback, top–down attention, illusory contours, amodal completion

INTRODUCTION

Visual object recognition has traditionally been described as a largely feedforward process that operates independently of and prior to top–down signals that reflect strategic processing or attentional effects. This standard model of object recognition is supported by research that spans multiple levels of analysis, including single- and multi-unit recording, computational modeling, and behavioral experiments, all of which have been discussed in detail in recent reviews (e.g., Serre et al., 2007; DiCarlo et al., 2012). Feedback projections, nearly equal in density to feedforward neurons throughout the ventral visual stream (Felleman and Van Essen, 1991; Sporns and Zwi, 2004), are commonly thought to subserve slower, attention-mediated processing that happens after recognition processes are complete, but not the core object recognition processing itself (Hochstein and Ahissar, 2002).

The proposal advanced in this paper is that these local, recurrent feedback connections also provide an avenue for rapid top–down signals that influence object recognition-related processing *as it is being carried out*—well before the slower attention-mediated processes. The theory is inspired by the pioneering work of Dehaene et al. (2006) and Lamme (2003, 2006) in identifying the neural correlates of consciousness. Both of these researchers' theories dissociate between local recurrent processing and top–down signals from frontoparietal areas in terms of the effects that they have on awareness. The present work draws a similar distinction between top–down, attention-mediated processing, and local recurrent processing between hierarchically adjacent areas within the ventral stream.

We support this distinction by first providing evidence that there are two temporally dissociable processes operating on these feedback projections; and second by presenting results showing an important functional role for the earlier, local recurrent processing.

EVIDENCE FOR A TEMPORAL DISSOCIATION OF LOCAL RECURRENT AND TOP–DOWN PROCESSING

Top–down attention is known to be a consciously generated, executive signal originating in frontal and parietal areas (Thompson et al., 2005; Bressler et al., 2008). Signals reflecting these strategic processes do not manifest in early visual areas until 150–170 ms after stimulus onset at the earliest, with most reported effects occurring within the range of 200–300 ms (Mehta et al., 2000a,b; Martinez et al., 2001; Noesselt et al., 2002). The relatively long latency of attentional effects in early visual areas is thought to arise from top–down signals that target late stages of the ventral stream and then progress backward toward V1 (Buffalo et al., 2010). Local recurrent processing can also be thought of as a top–down process, except that the signal originates from within the ventral stream itself, as opposed to frontal or parietal areas. Local recurrent processing is completely involuntary and does not require conscious execution, evidenced by its observation in recordings from anesthetized animals (Roland et al., 2006; Roland, 2010) and is simply a consequence of signal propagation through recurrent corticocortical connectivity. Specifically, as soon as a given area responds, signals are routed both to higher-level and lower-level connected areas. Feedback to immediately lower levels occurs with very short latencies—as quickly as 10 ms after the initial feedforward responses (Hupé et al.,

2001; Pascual-Leone and Walsh, 2001)—and thus could plausibly be underway after initial feedforward IT neural responses (ca. 80–100 ms) but before the completion of the categorization process (ca. 150 ms).

Recent research using methods that temporarily interfere with cortical processing have revealed strong evidence that recurrent feedback circuits are engaged during the first 80–150 ms of visual processing. One line of evidence comes from experiments that use transcranial magnetic stimulation (TMS) to temporarily prevent a targeted brain area from responding. In a recent study, Koivisto et al. (2011) used fMRI-localized TMS to selectively inactivate V1/V2 while subjects categorized images according to whether they contained an animal. The authors found that applying TMS to V1/V2 with stimulus onset asynchronies (SOAs) of 90–210 ms impaired categorization performance and subjective perception of stimuli. Camprodon et al. (2010) observed similar results with TMS applied over V1 only, but found that there were actually two windows of impairment with SOAs of 100 and 220 ms. Earlier work from Corthout et al. also found an early window of activity with an SOA of around 100 ms during which applying TMS over V1 impairs letter recognition (Corthout et al., 1999a,b). Collectively, these experiments show that disruption of processing in early visual areas around 100 ms after stimulus presentation impairs visual recognition. Importantly, this time window occurs after the earliest contributions of IT neurons, opening up the possibility that the impairment is due to the disruption of feedback to lower-level areas in influencing the quality of object representations. Furthermore, several of these studies found a second, later time window around 200 ms during which TMS also impaired recognition. This later time window coincides with the latency of spatial attention-mediated processing (Mehta et al., 2000a,b; Martinez et al., 2001; Noesselt et al., 2002; Buffalo et al., 2010), providing a temporal dissociation from the rapid recurrent processing effects that are of interest here.

Visual backward masking experiments have also identified a similar time window for recurrent processing around 100 ms after stimulus onset (Fahrenfort et al., 2007, 2008; Boehler et al., 2008). In backward masking experiments, a first stimulus (the target) is followed by a second stimulus (the mask) at a particular latency. At very short latencies, backward masking can impair recognition of the target and in some cases, prevent it from reaching awareness (Macknik and Livingstone, 1998). While the effect of backward masking was initially accounted for with a feedforward explanation (Breitmeyer and Ganz, 1976), modern theories of backward masking emphasize recurrent processing between higher-level and lower-level areas (Enns and Di Lollo, 2000; Lamme and Roelfsema, 2000; Wyatte et al., 2012a). Specifically, if information about a target stimulus being processed in higher-level areas is fed back down to lower areas, but a masking stimulus is simultaneously being processed at that lower level, there will be a fundamental mismatch in the information being processed at each level (Lamme and Roelfsema, 2000). This mismatch causes a decoupling in the functional connectivity (i.e., co-activation) between the visual areas involved in processing the stimulus, which has the psychological effect of greatly reduced perceptual visibility (Dehaene et al., 2001; Haynes et al., 2005).

Boehler et al. (2008) combined a backward masking paradigm with magnetoencephalography (MEG) recording to determine the time course of recurrent feedback to V1 during a recognition task. On trials where subjects correctly recognized the target stimulus (i.e., no impairment from the mask), there was modulation of the V1 MEG signal from 100 to 120 ms. This modulation occurred soon after (ca. 27 ms) the initial V1 signals and almost immediately after (ca. 11 ms) extrastriate generated signals, in strong accordance with being driven by rapid recurrent feedback from extrastriate areas to V1. Again, these rapid recurrent processing effects were dissociable from slower attentional modulation, which manifested 250–300 ms after stimulus presentation and only when subjects attended to the same region of the display that the target appeared in. In contrast, modulation from rapid recurrent processing occurred regardless of where subjects directed attention. Similar results have been demonstrated when combining backward masking with electroencephalography (EEG) recording with both rapid recurrent and slower attentional modulation, but with less emphasis on the specific neural generators of effects given the relatively poor spatial resolution of EEG (Fahrenfort et al., 2007, 2008).

Together, TMS and backward masking experiments provide strong support for the idea that recurrent visual processing engages striate and extrastriate areas around 100 ms after stimulus onset during visual recognition tasks. This local rapid recurrent processing is dissociable from attention-mediated or strategic processing both in terms of where the signals originate (within the ventral stream vs. frontal and parietal areas) and in terms of their relative time courses (ca. 100 ms vs. 150–170 ms at the earliest). Attention has long been known to modulate early EEG responses such as the P1 (first positive deflection, ca. 100 ms) and N1 (first negative deflection, ca. 150–200 ms) (Luck et al., 1990a; Hillyard and Anllo-Vento, 1998). Given the data reviewed here, it seems plausible that the P1 indexes recurrent feedback generated within the ventral stream while the N1 reflects the first influences of frontal and parietal attentional signals that progress backwards through visual areas toward V1 (Buffalo et al., 2010; see also Luck et al., 1990b).

While TMS impairment around 100 ms is consistent with the disruption of recurrent processing, it cannot rule out the possibility that the TMS is actually disrupting delayed feedforward responses. Specifically, low-level image properties such as local contrast can affect the temporal order of feedforward spikes, with lower contrast image regions exhibiting delay relative to more salient image regions (VanRullen and Thorpe, 2001, 2002). However, the information content of these regions is much lower than the salient regions that exhibit fast responses and thus it is unlikely that disrupting their contribution to object recognition processing will negatively impact recognition ability for relatively unambiguous images. More importantly, backward masking experiments that target impairment of recurrent processing provide additional constraints in interpreting TMS effects. Finally, 90–110 ms post stimulus onset is hypothesized to be the time at which peak feedback signals arrive at V1 from extrastriate areas (Roland et al., 2006; Roland, 2010).

Overall, it seems clear that recurrent processing operates within the established time course of object recognition, which

spans the first 150 ms of visual processing. The data reviewed in this section are summarized in **Table 1** with a rough sketch of overall feedforward and feedback events shown in **Figure 1**. Having established support for the idea that recurrent feedback occurs rapidly beginning around 100 after stimulus presentation, this paper now turns to discussion of its function.

EVIDENCE FOR A DISTINCT FUNCTIONAL ROLE FOR LOCAL RECURRENT SIGNALS

There is considerable evidence that local recurrent processing is important when stimuli are degraded, partial, or otherwise ambiguous, and we hypothesize that this is one important functional role for the dissociated process described in the previous section. The basic logic behind this proposal is that degrading a stimulus has been shown to weaken the initial responses in object-selective areas (Sclar et al., 1990; Kovacs et al., 1995; Nielsen et al., 2006; Williford and Maunsell, 2006), but recurrent processing over time can strengthen responses back to near undegraded levels and preserve selectivity via top–down reinforcement. Consistent with this idea, object-selective responses in IT cortex remain intact when stimuli are occluded, but take significantly more time to manifest than when stimuli are unoccluded (around an extra 50 ms on average, Kovacs et al., 1995; Nielsen et al., 2006).

Single-unit recordings that use reversible cooling to temporarily inactivate a particular brain area provide further support for our hypothesis. Hupé et al. (1998) applied cooling to area V5/MT, a visual area in the superior temporal sulcus of the monkey brain that sends feedback projections to areas V1, V2, and V3. Recordings from V1 through V3 indicated that

Table 1 | Summary of data that suggest a temporal dissociation of local recurrent and top–down attentional processing.

Reference	Milliseconds after Stimulus Onset	Comments
	10 20 30 40 50 60 70 80 90 100 110 120 130 140 150 160 170 180 190 200 210 220 230 240 250 260 270 280 290 300	
Koivisto et al., 2011		TMS over human V1. Scene categorization impairment at indicated times
Corthout et al., 1999a		TMS over human V1. Letter recognition impairment at indicated times
Corthout et al.,1999b		TMS over human V1. Letter recognition impairment at indicated times
Camprodon et al., 2010		TMS over human V1. Scene categorization impairment at indicated times
Boehler et al., 2008		MEG in human subjects. Effects of masking over V1 at indicated times
Fahrenfort et al., 2007		EEG in human subjects. Effect of masking over occipital channels at indicated times
Fahrenfort et al., 2008		EEG in human subjects. Effect of masking over occipital channels at indicated times
Martinez et al., 2001		EEG/fMRI in human subjects. V1 modulation with attention at indicated times
Noesselt et al., 2002		EEG/MEG/fMRI in human subjects. V1 modulation with attention at indicated times
Mehta et al., 2000a, 2000b		LFP in awake behaving monkey. V1 modulation with attention at indicated times

■ Putative feedforward   ■ Proposed local recurrent   ■ Putative attentional

FIGURE 1 | Proposed time course of feedforward and feedback events during early visual processing. Top row: Feedforward-dominant latencies, which are well-documented in the literature (e.g., Nowak and Bullier, 1997). Light pink shading refers to earliest reported latencies, likely corresponding to the depicted areas' first spikes, while darker pink shading corresponds to ongoing feedforward responses. **Bottom row:** Areas are shaded orange when they are known to be receiving recurrent feedback. Most reports of recurrent feedback to V1 center around an absolute latency of 100 ms after stimulus presentation, with some reports being slightly faster. Common methods used to detect feedback (coarse application of TMS, MEG, EEG) do not have the spatial resolution to distinguish between feedback to V1 and extrastriate areas, but the view taken here is that feedback originates in immediately adjacent areas, and thus those areas that fire earliest during the feedforward dominant phase will also be the first to receive feedback.

responses to moving bar stimuli were vastly weakened (fewer spikes observed per second) when V5/MT was inactive compared to control experiments in which it was active. This attenuation of lower-level responses was most dramatic in low salience conditions, such as when the bar had a very low contrast, a point that will be discussed in further detail later in this section. These results suggest that when higher-level visual areas are active, they provide additional excitatory input to lower levels. Similar effects have been shown for other recurrent circuits in other mammalian species such as those involving middle suprasylvian (MS) cortex and V1 (Galuske et al., 2002) as well as V2 and V1 (Sandell and Schiller, 1982; Mignard and Malpeli, 1991), suggesting that top–down amplification is a highly generic mechanism that occurs between any two recurrently connected areas.

Top–down amplification promotes visual awareness (Lamme, 2003, 2006; Dehaene et al., 2006), and some data indicate that amplification is a simple contrast gain operation, as some have suggested is implemented by attention (e.g., Reynolds et al., 2000; Reynolds and Heeger, 2009). However, there is mounting evidence that recurrent amplification also plays an important functional role in visual object recognition when stimuli are degraded or ambiguous, by promoting a complex grouping and "filling-in" process.

Wyatte et al. (2012a) degraded visual object stimuli using visual occlusion and contrast degradation and used backward masking to control whether recurrent processing mechanisms were available (Enns and Di Lollo, 2000; Lamme and Roelfsema, 2000). When relatively clear stimuli were masked using a relatively long latency 100 ms SOA pattern mask, there was little impairment in recognition performance. However, when heavily occluded or low contrast stimuli were masked, the mask had a much larger effect, suggesting that recurrent processing was crucial in resolving object identity in these conditions. Simulations using a computational model of object recognition that included recurrent feedback between hierarchically adjacent layers (O'Reilly et al., 2013) showed that responses in both lower layers (corresponding to striate/extrastriate regions) and upper layers (corresponding to IT cortex) strengthened over time when objects were occluded. Backward masking selectively interfered with this strengthening process, which was crucial when stimulus signals were weak due to degradation. Furthermore, the strengthening dynamic was found to be specifically due to recurrent feedback—purely feedforward versions of the model exhibited asymptotic response levels across areas.

One possibility for the mechanism underlying these recognition performance differences is a grouping and "filling-in" process similar to what is observed in the figure-ground literature in V1 (Figure 2A), but repeated between higher levels of the visual hierarchy. As an illustration, consider a population of IT neurons that respond to bicycle stimuli (Figures 2B,C). If a bicycle stimulus is occluded and only the wheels are visible, some members of this population will become active (specifically, those corresponding to wheel-like features), but the selective response across the full population will be unavailable. The partial responses, however, will be propagated back to earlier visual areas, which will drive neurons that are sensitive to visual features that are known to co-occur with bicycle wheels, such as a bicycle's frame, handlebars, and saddle. Importantly these responses occur in the absence of these features in the actual stimulus. These "illusory" responses in turn provide new driving potential to IT neurons, ultimately evoking the selective response corresponding to the unoccluded stimulus across the full IT population responsive to bicycles. The IT response is "object complete," meaning that there is little-to-no difference between the response to the partially occluded object and the complete object—the brain has filled in the missing information.

Computationally, recurrent processing's amplification effect is capable of supporting a grouping or surface-based encoding. The most convincing demonstrations of these computations are found in the figure-ground processing literature, where the term "contextual modulation" is used to describe them (Zipser et al., 1996; Lamme et al., 1998). In contextual modulation, neurons with non-overlapping receptive fields such as those found in V1 are capable of modulating and reinforcing each other by virtue shared connections through higher levels in the visual hierarchy where receptive fields do tend to overlap. This extra modulation has the effect of grouping together figural elements of a display and enhancing their activity relative to background elements effectively spreading activation throughout the figure interior and "filling" it in as a perceptually salient surface (Figure 2A). The models suggest that contextual modulation is driven by recurrent feedback, because lesions of feedback from extrastriate and dorsal structures to V1 obliterate the surface filling effect. They further illustrate that the timing of contextual modulation to area V1 would be on the order of 80–100 ms after stimulus presentation, coinciding with the known time course of feedback to striate areas during visual processing. Finally, contextual modulation is dissociable from slower top–down attentional effects, not just with respect to time course but also because its surface filling computations are retained even when attention is deployed away from the target stimulus (Poort et al., 2012).

There are two phenomena in the experimental literature that support the grouping and filling-in roles of recurrent processing during object recognition. The first is the perception of illusory contours, such as in displays containing Kanizsa shapes (Figure 3). V1 neurons have been shown to respond to the illusory contours that compose Kanizsa shapes, such as the edges of the illusory square in Figure 3. Multi-unit recordings have indicated that these responses occur beginning around 100 ms after stimulus presentation, which is shortly after the V1 responses to a physical contour with the same orientation and location, suggesting a role for feedback in their encoding (Lee and Nguyen, 2001; Seghier and Vuilleumier, 2006). Specifically, recurrent feedback from extrastriate areas could support the perception of illusory contours in the Kanizsa illusion by grouping similarly oriented contours at the V1 level that fall within the shape's receptive field; this would cause the shape to be perceived as perceptually salient surface similar to the way texture-defined shapes are perceived (Figure 3B). As such, a recent experiment has indicated that global contour information emerges in V1 responses shortly after the first V4 responses, implicating recurrent feedback in this grouping process (Chen et al., 2014).
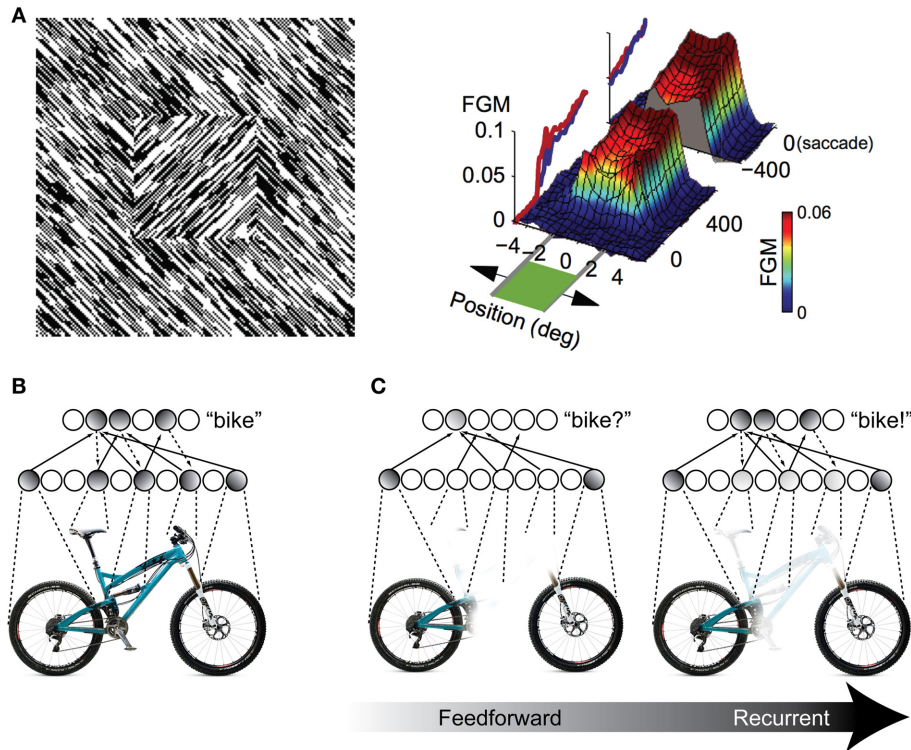
FIGURE 2 | Illustration of recurrent processing's filling-in computations during figure-ground processing and object recognition. (A) Processing of an orientation-defined square stimulus results in enhancement of the figural elements compared to the background elements. This enhancement comes in the form of recurrent feedback that groups together common image elements and spreads activation throughout the interior of the square, effectively "filling" it in as a perceptually salient surface. FGM, Figure Ground Modulation, i.e., difference between figure and background responses. Adapted from Lamme et al. (1998) and Poort et al. (2012). **(B,C)** The same feedback-based "filling-in" principle can be applied to object recognition processing when stimuli are occluded. When object features are occluded, only a partial representation is elicited by the first feedforward responses. However, recurrent feedback (e.g., between IT and extrastriate areas) propagates these partial responses back to early visual areas, driving neurons that respond to co-occurring features that might be occluded in the physical stimulus. This recurrent processing between hierarchically adjacent visual areas can effectively "fill in" the occluded features in the object representation.
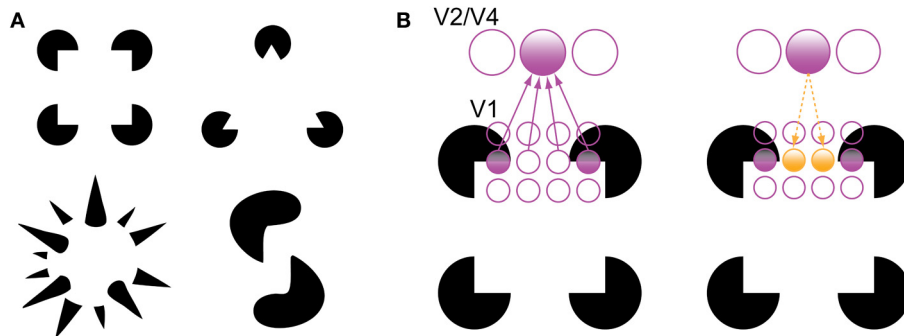
FIGURE 3 | Illusory contour perception in Kanizsa shapes (Kanizsa, 1979). (A) Traditional and more complex Kanizsa shapes that evoke strong illusory contour percepts. Complex shapes courtesy of Steven Lehar (http://cns-alumni.bu.edu/~slehar/Lehar.html). **(B)** Perception of illusory contours has been suggested to arise by virtue of recurrent feedback from extrastriate areas to V1. Specifically, feedback from an extrastriate neuron drives neurons that code similarly oriented contours that fall within its receptive field, spreading activation across the gap in the Kanizsa shape. Adapted from Lee (2003).

The second supportive phenomenon is an actual object completion effect, which has gained support from fMRI studies that show little-to-no difference in the activation levels of occluded and unoccluded stimuli in object-selective regions of cortex (Lerner et al., 2002). Intact activation, however, could simply reflect increased gain of the encoded object fragments without a more complex completion process. To differentiate between these two possibilities, one can use an fMRI adaptation paradigm,

which depends on neural mechanisms that decrease response levels for repeated stimuli that are perceived as the same. This method gives an experimenter an index of how perceptually similar two experimental conditions are. For example, Kourtzi and Kanwisher (2001) presented observers with images that contained occluding bars either in front of or behind target objects (in which case, the targets were effectively unoccluded). The experiment measured the hemodynamic response in the lateral occipital cortex (LOC), which has been strongly suggested as the human homolog of IT cortex in monkey (Grill-Spector et al., 2001; Orban et al., 2004). The results indicated that there was no significant change in hemodynamic response when subjects were presented with two identical objects in sequence, as well as when subjects were presented with occluded and unoccluded versions of an object in sequence. Thus, at the level of LOC, there is little difference in the way that unoccluded and occluded versions of the same object are represented. More recent techniques such as representational similarity (Kriegeskorte et al., 2008a,b) or decoding analyses (Tang et al., in press) might further illuminate how occluded objects are represented in various regions of cortex.

While the perception of illusory contours has been linked to recurrent feedback (Lee, 2003), this explanation has not been has explored as extensively in the object completion literature, likely due to most studies using relatively coarse measures like fMRI (e.g., the aforementioned studies that rely on fMRI adaptation). Computationally, illusory contour perception and object completion could be implemented by the same mechanism, whereby higher-level neurons with overlapping receptive fields feed responses back to lower-level neurons in the absence of the visual information itself and produce the perception of illusory object features. According to this view, when operating between extrastriate levels and V1, the mechanism produces illusory contours; when operating between IT cortex and extrastriate areas, it produces more complex illusory object features. There is some support for this idea in the literature. For example, Rauschenberger et al. (2006) demonstrated object completion effects in LOC as well as in extrastriate areas when stimuli were presented for longer durations, suggesting that there is a "temporal unfolding" of object completion from higher levels of the ventral stream to lower-level areas.

However, illusory contour stimuli evoke a perceptually salient completion phenomenon, whereas the filling-in of objects does not. These processes have been distinguished in the literature as "modal" and "amodal" completion, respectively (Johnson and Olshausen, 2005; Seghier and Vuilleumier, 2006). Modal completion has been shown to elicit illusory responses in V1 (Lee and Nguyen, 2001), supporting the idea that whatever representation is present in V1 is what we "perceive" (Bullier, 2001). It is unclear whether amodal completion processing also reaches back to the level of V1. Some studies indicate that V1 represents completed shapes (Rauschenberger et al., 2006), whereas others show that the complete representation is only present in extrastriate and higher-level areas (Weigelt et al., 2007). More recently, Emmanouil and Ro (2014) showed that object completion can occur rapidly and without visual awareness, further supporting the dissociation of object completion from top–down attention.

If our proposal is correct, the time course of object completion effects should agree with the time course of recurrent processing as described above. Some studies show object completion effects beginning to manifest over temporal and parietal sites (as indexed by EEG scalp recordings) around 130 ms at the earliest and continuing to evolve until around 200 ms into processing (Johnson and Olshausen, 2005; Chen et al., 2009). These data are consistent with the explanation of object completion rapidly engaging recurrent processing with striate and extrastriate areas, assuming the 50 ms delay typically observed when the brain is processing occluded object stimuli (Kovacs et al., 1995; Nielsen et al., 2006).

However, other studies have suggested a much later time course for object completion effects, beginning around 200 ms and completing around 400 ms (Doniger et al., 2000; Sehatpour et al., 2006, 2008). One consistent characteristic of these latter studies is that they use fragmented line drawings of objects, whereas studies that associate an early time course with object completion have used photorealistic images of objects. It is unclear whether this late temporal correlate of object completion is due to relatively slow, attention-mediated processing, or due to a fundamentally different type of processing. For example, photorealistic occlusion might recruit the surface-coded computations associated with recurrent processing since there are explicitly depicted depth planes (an occluder and an object) whereas resolving contour fragmentation might rely on a completely different computation since depth planes are less well-defined in line drawings. Furthermore, the studies that associate the later time course with object completion have not used a paradigm such as response adaptation that crucially allows inference about whether an unoccluded and occluded object are represented similarly.

In summary, it seems clear that recurrent processing promotes signal amplification between reciprocally connected brain regions. There is substantial evidence that this is not a simple multiplicative gain operation, but a considerably more complex grouping or surface-based computation that spreads activation between related object features. This idea has been well-studied in the literature on illusory contour perception and the data support the explanation that illusory contour perception is due to V1 neurons receiving recurrent feedback from extrastriate regions. The same idea can be applied to object completion effects in IT cortex, predicting that they are due to feedback-rectified signals from extrastriate regions. This recurrent processing-based explanation has received little attention in the literature, but is generally supported by the timing of object completion effects.

SUMMARY AND FUTURE RESEARCH

Over the last 5–10 years, evidence has accumulated that local recurrent signals are an integral part of early visual processing. TMS studies have indicated that recurrent processing engages striate and extrastriate areas during visual recognition tasks in as little as 100 ms (Camprodon et al., 2010; Koivisto et al., 2011) and theories of backward masking have provided additional accordant timing data as well as suggested a general theory of how corticocortical interactions support visual perception (Fahrenfort et al., 2007, 2008; Boehler et al., 2008). Surprisingly though, relatively little work has focused on synthesizing these ideas with theories of visual object recognition, which is commonly held

to be primarily a feedforward process (DiCarlo et al., 2012). Instead, theories of recurrent processing have focused on the role of interactions between brain areas in promoting visual awareness (Lamme, 2003, 2006; Dehaene et al., 2006). Object perception has long been known to benefit from top–down signals that reflect attention or strategic processing, but its time course has been considered to be too slow to support the initial rapid recognition processes (Hochstein and Ahissar, 2002; VanRullen, 2007).

This paper has attempted to map out the time course of feedforward- and feedback-based events during the first 150 ms of visual processing and establish the function that rapid recurrent processing between brain areas plays within this time frame. Specifically, we propose the following overall process: A feedforward-dominant wave of activation flows up to IT in the first 80–100 ms after stimulus presentation, quickly evoking object-selective responses, while, simultaneously, activation is also feeding backward through this pathway. In the following 20 ms (an absolute latency of 100–120 ms), prefrontal areas that support the actual object categorization decision receive their first feedforward responses from IT neurons, while simultaneously, recurrent feedback from extrastriate areas has had sufficient time to more fully engage V1 populations. Recurrent feedback to V1 amplifies neurons' initial responses by grouping the responses to similar object features and enhancing them relative to other responses (Zipser et al., 1996; Lamme et al., 1998; Poort et al., 2012). In some cases, these grouping computations can cause the perception of illusory contours and surfaces (Lee and Nguyen, 2001; Seghier and Vuilleumier, 2006), but they also seem to be important when objects are degraded in order to rectify signals (Hupé et al., 1998). At an absolute latency of 120–140 ms after the initial stimulus presentation, the now extensive recurrent processing between IT and extrastriate areas can cause the representation of more complex illusory features that support object completion, by propagating these illusory responses back toward IT populations. We have recently developed a biologically-based computational model that exhibits just these dynamics (O'Reilly et al., 2013), and can provide a platform for integrating the various data cited here, while generating further testable predictions.

It is unlikely that object completion in IT cortex is a sole function of rectified responses from extrastriate areas being propagated forward in the range of 120–140 ms (or 170–190 ms, assuming the 50 ms delay observed when the brain processes occluded object stimuli; see Kovacs et al., 1995; Nielsen et al., 2006). Object completion likely also benefits from the first recurrent responses from prefrontal areas that arrive shortly after this time frame. This feedback from prefrontal areas could reflect top–down predictions that constrain the space of potential object representations in IT cortex (Bar et al., 2006; Kveraga et al., 2007), which might also have the effect of filling in visual information when it is missing from the physical stimulus. It is also plausible that lateral interactions within IT cortex itself could support object completion by enforcing statistical co-occurrences and mutual exclusions between object features (Akrami et al., 2009; Daelli and Treves, 2010). It would not be surprising if a combination of rectified feedforward responses, feedback from prefrontal areas, and lateral interactions within IT cortex itself support object completion by bringing the brain as a whole into

an attractor that combines bottom–up sensory information with top–down task demands and appropriate local constraints (e.g., Spivey, 2008). Future research that uses sophisticated techniques to rapidly and systematically disable feedforward, recurrent and lateral connectivity (e.g., optogenetics, Deisseroth, 2011) might be necessary to disentangle the relative contributions of each of these influences. Nevertheless, any contribution to object completion from local recurrent processes is supportive of the distinct functional role in resolving degraded or ambiguous stimuli proposed here.

One remaining question concerns whether recurrent processing is necessary for recognizing relatively unambiguous stimuli. "Core object recognition" (DiCarlo et al., 2012) of stimuli that vary in terms of their spatial position, scale, pose, and illumination can be rapidly decoded from the first IT responses (Hung et al., 2005). Early IT responses are also known to exhibit invariance to limited clutter (Missal et al., 1997; Zoccolan et al., 2005), suggesting that the bulk of object recognition is solved by a largely feedforward process. Importantly, these data are not fundamentally incompatible with the theory proposed here. Feedback acts on immediately lower areas with latencies as short as 10 ms (Hupé et al., 2001; Pascual-Leone and Walsh, 2001) and might be important for the Winner-Take-All (WTA) or "max" computations (Riesenhuber and Poggio, 1999; Wyatte et al., 2012b; O'Reilly et al., 2013) that have been suggested to contribute to core object recognition. Our theory has focused on recurrent processing under challenging object recognition conditions such as when stimuli are occluded or otherwise degraded. However, more substantial variability in the spatial properties of inputs might also benefit from recurrent processing. A variant of the "animal/no animal" recognition task used in many studies has shown that increasing target viewing distance in the stimulus causes backward masking to have a greater effect (Serre et al., 2007, supporting information), implicating recurrent processing for robust recognition under these conditions (Wyatte et al., 2012a). Further research with stimuli whose spatial properties can be manipulated parametrically (DiCarlo et al., 2012; Cadieu et al., 2013) combined with methods like TMS and backward masking will be necessary to determine the exact conditions under which recurrent processing is necessary.

If the theory proposed here is true, the standard description of object recognition as a feedforward process is somewhat misleading. Simply put, there is always ongoing brain activity that must be combined with new incoming sensory information, so that the notion of a strictly "feedforward sweep" is fundamentally ill-conceived (Arieli et al., 1996; Tsodyks et al., 1999). Ongoing activity could be used to establish moment-to-moment constraints that effectively guide coherent perception via recurrent processing mechanisms. While the seminal research on object recognition often focused on simple spike counts of anesthetized animals to map out the receptive field characteristics of neurons throughout the ventral stream in a well-controlled manner, future research should emphasize more complex corticocortical interactions in the awake, behaving brain to determine how neural interactions involving feedforward, lateral, and recurrent processing mechanisms combine to give rise to the visual system's robust perceptual abilities even in difficult stimulus conditions.

REFERENCES

Akrami, A., Liu, Y., Treves, A., and Jagadeesh, B. (2009). Converging neuronal activity in inferior temporal cortex during the classification of morphed stimuli. *Cereb. Cortex* 19, 760–776. doi: 10.1093/cercor/bhn125

Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273, 1868–1871. doi: 10.1126/science.273.5283.1868

Bar, M., Kassam, K., Ghuman, A., Boshyan, J., and Schmidt, A. (2006). Top-down facilitation of visual recognition. *Proc. Natl. Acad. Sci. U.S.A.* 103, 449–454. doi: 10.1073/pnas.0507062103

Boehler, C. N., Schoenfeld, M. A., Heinze, H. J., and Hopf, J. M. (2008). Rapid recurrent processing gates awareness in primary visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 105, 8742–8747. doi: 10.1073/pnas.0801999105

Breitmeyer, B. G., and Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing. *Psychol. Rev.* 83, 1–36. doi: 10.1037/0033-295X.83.1.1

Bressler, S. L., Tang, W., Sylvester, C. M., Shulman, G. L., and Corbetta, M. (2008). Top-down control of human visual cortex by frontal and parietal cortex in anticipatory visual spatial attention. *J. Neurosci.* 28, 10056–10061. doi: 10.1523/JNEUROSCI.1776-08.2008

Buffalo, E. A., Fries, P., Landman, R., Liang, H., and Desimone, R. (2010). A backward progression of attentional effects in the ventral stream. *Proc. Natl. Acad. Sci. U.S.A.* 107, 361–365. doi: 10.1073/pnas.0907658106

Bullier, J. (2001). Feedback connections and conscious vision. *Trends Cogn. Sci.* 5, 369–370. doi: 10.1016/S1364-6613(00)01730-7

Camprodon, J. A., Zohary, E., Brodbeck, V., and Pascual-Leone, A. (2010). Two phases of v1 activity for visual recognition of natural images. *J. Cogn. Neurosci.* 22, 1262–1269. doi: 10.1162/jocn.2009.21253

Cadieu, C. F., Hong, H., Yamins, D., Pinto, N., Majaj, N. J., and DiCarlo, J. J. (2013). "The neural representation benchmark and its evaluation on brain and machine," in *International Conference on Learning Representations (ICLR)* (Scottsdale, AZ).

Chen, J., Liu, B., Chen, B., and Fang, F. (2009). Time course of amodal completion in face perception. *Vision Res.* 49, 752–758. doi: 10.1016/j.visres.2009.02.005

Chen, M., Yan, Y., Gong, X., Gilbert, C., Liang, H., and Li, W. (2014). Incremental integration of global contours through interplay between visual cortical areas. *Neuron* 82, 682–694. doi: 10.1016/j.neuron.2014.03.023

Corthout, E., Uttl, B., Walsh, V., Hallett, M., and Cowey, A. (1999a). Timing of activity in early visual cortex as revealed by transcranial magnetic stimulation. *Neuroreport* 10, 2631–2634. doi: 10.1097/00001756-199908200-00035

Corthout, E., Uttl, B., Ziemann, U., Cowey, A., and Hallett, M. (1999b). Two periods of processing in the (circum)striate visual cortex as revealed by transcranial magnetic stimulation. *Neuropsychologia* 37, 137–145. doi: 10.1016/S0028-3932(98)00088-8

Daelli, V., and Treves, A. (2010). Neural attractor dynamics in object recognition. *Exp. Brain Res.* 203, 241–248. doi: 10.1007/s00221-010-2243-1

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., and Sergent, C. (2006). Conscious preconscious and subliminal processing: a testable taxonomy. *Trends Cogn. Sci.* 10, 204–211. doi: 10.1016/j.tics.2006.03.007

Dehaene, S., Naccache, L., Cohen, L., Le Bihan, D., Mangin, J. F., Poline, J. B., et al. (2001). Cerebral mechanisms of work masking and unconscious repetition priming. *Nat. Neurosci.* 4, 752–758. doi: 10.1038/89551

Deisseroth, K. (2011). Optogenetics. *Nat. Methods* 8, 26–29. doi: 10.1038/nmeth.f.324

DiCarlo, J. J., Zoccolan, D., and Rust, N. C. (2012). How does the brain solve visual object recognition? *Neuron* 73, 415–434. doi: 10.1016/j.neuron.2012.01.010

Doniger, G. M., Foxe, J. J., Murray, M. M., Higgins, B. A., Snodgrass, J. G., Schroeder, C. E., et al. (2000). Activation timecourse of ventral visual stream object-recognition areas: high density electrical mapping of perceptual closure processes. *J. Cogn. Neurosci.* 12, 615–621. doi: 10.1162/089892900562372

Emmanouil, T. A., and Ro, T. (2014). Amodal completion of unconsciously presented objects. *Psychon. Bull. Rev.* doi: 10.3758/s13423-014-0590-9. [Epub ahead of print].

Enns, J. T., and Di Lollo, V. (2000). What's new in visual masking? *Trends Cogn. Sci.* 4, 345–352. doi: 10.1016/S1364-6613(00)01520-5

Fahrenfort, J. J., Scholte, H. S., and Lamme, V. A. F. (2007). Masking disrupts reentrant processing in human visual cortex. *J. Cogn. Neurosci.* 19, 1488–1497. doi: 10.1162/jocn.2007.19.9.1488

Fahrenfort, J. J., Scholte, H. S., and Lamme, V. A. F. (2008). The spatiotemporal profile of cortical processing leading up to visual perception. *J. Vision* 8, 12.1–12.12. doi: 10.1167/8.1.12

Felleman, D. J., and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* 1, 1–47. doi: 10.1093/cercor/1.1.1

Galuske, R. A. W., Schmidt, K. E., Goebel, R., Lomber, S. G., and Payne, B. R. (2002). The role of feedback in shaping neural representations in cat visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 99, 17083–17088. doi: 10.1073/pnas.242399199

Grill-Spector, K., Kourtzi, Z., and Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Res.* 41, 1409–1422. doi: 10.1016/S0042-6989(01)00073-6

Haynes, J. D., Driver, J., and Rees, G. (2005). Visibility reflects dynamic changes of effective connectivity between v1 and fusiform cortex. *Neuron* 46, 811–821. doi: 10.1016/j.neuron.2005.05.012

Hillyard, S. A., and Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proc. Natl. Acad. Sci.U.S.A.* 95, 781–787. doi: 10.1073/pnas.95.3.781

Hochstein, S., and Ahissar, M. (2002). View from the top: hierarchies and reverse hierarchies in the visual system. *Neuron* 36, 791–804. doi: 10.1016/S0896-6273(02)01091-7

Hung, C. P., Kreiman, G., Poggio, T., and DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferotemporal cortex. *Science* 310, 863–866. doi: 10.1126/science.1117593

Hupé, J. M., James, A. C., Girard, P., and Bullier, J. (2001). Response modulations by static texture surround in area v1 of the macaque monkey do not depend on feedback connections from v2. *J. Neurophysiol.* 85, 146–163.

Hupé, J. M., James, A. C., Payne, B. R., Lomber, S. G., Girard, P., and Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by v1 v2 and v3 neurons. *Nature* 394, 784–787. doi: 10.1038/29537

Johnson, J. S., and Olshausen, B. A. (2005). The recognition of partially visible natural objects in the presence and absence of their occluders. *Vision Res.* 45, 3262–3276. doi: 10.1016/j.visres.2005.06.007

Kanizsa, G. (1979). *Organization in Vision: Essays on Gestalt Perception*. Westport, CT: Praeger Publishers.

Koivisto, M., Railo, H., Revonsuo, A., Vanni, S., and Salminen-Vaparanta, N. (2011). Recurrent processing in v1/v2 contributes to categorization of natural scenes. *J. Neurosci.* 31, 2488–2492. doi: 10.1523/JNEUROSCI.3074-10.2011

Kourtzi, Z., and Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science* 293, 1506–1509. doi: 10.1126/science.1061133

Kovacs, G., Vogels, R., and Orban, G. A. (1995). Selectivity of macaque inferior temporal neurons for partially occluded shapes. *J. Neurosci.* 15, 1984–1997.

Kriegeskorte, N., Mur, M., and Bandettini, P. (2008a). Representational similarity analysis – connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* 2:4. doi: 10.3389/neuro.06.004.2008

Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., et al. (2008b). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60, 1126–1141. doi: 10.1016/j.neuron.2008.10.043

Kveraga, K., Ghuman, A., and Bar, M. (2007). Top-down predictions in the cognitive brain. *Brain Cogn.* 65, 145–168. doi: 10.1016/j.bandc.2007.06.007

Lamme, V., and Roelfsema, P. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci.* 23, 571–579. doi: 10.1016/S0166-2236(00)01657-X

Lamme, V., Super, H., and Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Curr. Opin. Neurobiol.* 8, 529–535. doi: 10.1016/S0959-4388(98)80042-1

Lamme, V. A. (2003). Why visual attention and awareness are different. *Trends Cogn. Sci.* 7, 12–18. doi: 10.1016/S1364-6613(02)00013-X

Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends Cogn. Sci.* 10, 494–501. doi: 10.1016/j.tics.2006.09.001

Lee, T. S. (2003). Computations in the early visual cortex. *J. Physiol.* 97, 121–139. doi: 10.1016/j.jphysparis.2003.09.015

Lee, T. S., and Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 98, 1907–1911. doi: 10.1073/pnas.98.4.1907

Lerner, Y., Hendler, T., and Malach, R. (2002). Object-completion effects in the human lateral occipital cortex. *Cereb. Cortex* 12, 163–177. doi: 10.1093/cercor/12.2.163

Luck, S. J., Heinze, H. J., Mangun, G. R., and Hillyard, S. A. (1990a). Visual event-related potentials index focused attention within bilateral stimulus arrays. I. Evidence for early selection. *Electroencephalogr. Clin. Neurophysiol.* 75, 511–527. doi: 10.1016/0013-4694(90)90139-B

Luck, S. J., Heinze, H. J., Mangun, G. R., and Hillyard, S. A. (1990b). Visual event-related potentials index focused attention within bilateral stimulus arrays. II. Functional dissociation of P1 and N1 components. *Electroencephalogr. Clin. Neurophysiol.* 75, 528–542. doi: 10.1016/0013-4694(90)90139-B

Macknik, S. L., and Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nat. Neurosci.* 1, 144–149. doi: 10.1038/393

Martinez, A., Di Russo, F., Anllo-Vento, L., Sereno, M. I., Buxton, R. B., and Hillyard, S. A. (2001). Putting spatial attention on the map: timing and localization of stimulus selection processes in striate and extrastriate visual areas. *Vision Res.* 41, 1437–1457. doi: 10.1016/S0042-6989(00)00267-4

Mehta, A. D., Ulbert, I., and Schroeder, C. E. (2000a). Intermodal selective attention in monkeys. I: distribution and timing of effects across visual areas. *Cereb. Cortex* 10, 343–358. doi: 10.1093/cercor/10.4.343

Mehta, A. D., Ulbert, I., and Schroeder, C. E. (2000b). Intermodal selective attention in monkeys. II: physiological mechanisms of modulation. *Cereb. Cortex* 10, 359–370. doi: 10.1093/cercor/10.4.359

Mignard, M., and Malpeli, J. G. (1991). Paths of information flow through visual cortex. *Science* 251, 1249–1251. doi: 10.1126/science.1848727

Missal, M., Vogels, R., and Orban, G. A. (1997). Responses of macaque inferior temporal neurons to overlapping shapes. *Cereb. Cortex* 7, 758–767. doi: 10.1093/cercor/7.8.758

Nielsen, K. J., Logothetis, N. K., and Rainer, G. (2006). Dissociation between local field potentials and spiking activity in macaque inferior temporal cortex reveals diagnosticity-based encoding of complex objects. *J. Neurosci.* 26, 9639–9645. doi: 10.1523/JNEUROSCI.2273-06.2006

Noesselt, T., Hillyard, S. A., Woldorff, M. G., Schoenfeld, A., Hagner, T., Jncke, L., et al. (2002). Delayed striate cortical activation during spatial attention. *Neuron* 35, 575–587. doi: 10.1016/S0896-6273(02)00781-X

Nowak, L., and Bullier, J. (1997). "The timing of information transfer in the visual system," in *Cerebral Cortex* Vol. 12: *Extrastriate Cortex in Primates*, eds K. S. Rockland, J. H. Kaas, and A. Peters (New York, NY: Plenum), 205–241.

Orban, G. A., Van Essen, D., and Vanduffel, W. (2004). Comparative mapping of higher visual areas in monkeys and humans. *Trends Cogn. Sci.* 8, 315–324. doi: 10.1016/j.tics.2004.05.009

O'Reilly, R., Wyatte, D., Herd, S., Mingus, B., and Jilk, D. (2013). Recurrent processing during object recognition. *Front. Psychol.* 4:124. doi: 10.3389/fpsyg.2013.00124

Pascual-Leone, A., and Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science* 292, 510–512. doi: 10.1126/science.1057099

Poort, J., Raudies, F., Wannig, A., Lamme, V. A. F., Neumann, H., and Roelfsema, P. R. (2012). The role of attention in figure-ground segregation in areas v1 and v4 of the visual cortex. *Neuron* 75, 143–156. doi: 10.1016/j.neuron.2012.04.032

Rauschenberger, R., Liu, T., Slotnick, S. D., and Yantis, S. (2006). Temporally unfolding neural representation of pictorial occlusion. *Psychol. Sci.* 17, 358–364. doi: 10.1111/j.1467-9280.2006.01711.x

Reynolds, J., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of v4 neurons. *Neuron* 26, 703–714. doi: 10.1016/S0896-6273(00)81206-4

Reynolds, J. H., and Heeger, D. J. (2009). The normalization model of attention. *Neuron* 61, 168–185. doi: 10.1016/j.neuron.2009.01.002

Riesenhuber, M., and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat. Neurosci.* 2, 1019–1025. doi: 10.1038/14819

Roland, P. (2010). Six principles of visual cortical dynamics. *Front. Syst. Neurosci.* 4:28. doi: 10.3389/fnsys.2010.00028

Roland, P. E., Hanazawa, A., Undeman, C., Eriksson, D., Tompa, T., Nakamura, H., et al. (2006). Cortical feedback depolarization waves: a mechanism of top-down influence on early visual areas. *Proc. Natl. Acad. Sci. U.S.A.* 103, 12586–12591. doi: 10.1073/pnas.0604925103

Sandell, J. H., and Schiller, P. H. (1982). Effect of cooling area 18 on striate cortex cells in the squirrel monkey. *J. Neurophysiol.* 48, 38–48.

Sclar, G., Maunsell, J. H., and Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res.* 30, 1–10. doi: 10.1016/0042-6989(90)90123-3

Seghier, M. L., and Vuilleumier, P. (2006). Functional neuroimaging findings on the human perception of illusory contours. *Neurosci. Biobehav. Rev.* 30, 595–612. doi: 10.1016/j.neubiorev.2005.11.002

Sehatpour, P., Molholm, S., Javitt, D. C., and Foxe, J. J. (2006). Spatiotemporal dynamics of human object recognition processing: an integrated high-density electrical mapping and functional imaging study of "closure" processes. *Neuroimage* 29, 605–618. doi: 10.1016/j.neuroimage.2005.07.049

Sehatpour, P., Molholm, S., Schwartz, T. H., Mahoney, J. R., Mehta, A. D., Javitt, D. C., et al. (2008). A human intracranial study of long-range oscillatory coherence across a frontal-occipital-hippocampal brain network during visual object processing. *Proc. Natl. Acad. Sci. U.S.A.* 105, 4399–4404. doi: 10.1073/pnas.0708418105

Serre, T., Kreiman, G., Kouh, M., Cadieu, C., Knoblich, U., and Poggio, T. (2007). A quantitative theory of immediate visual recognition. *Prog. Brain Res.* 165, 33–56. doi: 10.1016/S0079-6123(06)65004-8

Spivey, M. J. (2008). *The Continuity of Mind.* New York, NY: Oxford University Press.

Sporns, O., and Zwi, J. D. (2004). The small world of the cerebral cortex. *Neuroinformatics* 2, 145–162. doi: 10.1385/NI:2:2:145

Tang, H., Buia, C., Madhavan, R., Madsen, J., Anderson, W., Crone, N., et al. (in press). Spatiotemporal dynamics underlying object completion in human ventral visual cortex. *Neuron.*

Thompson, K. G., Biscoe, K. L., and Sato, T. R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *J. Neurosci.* 25, 9479–9487. doi: 10.1523/JNEUROSCI.0741-05.2005

Tsodyks, M., Kenet, T., Grinvald, A., and Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286, 1943–1946. doi: 10.1126/science.286.5446.1943

VanRullen, R. (2007). The power of the feed-forward sweep. *Adv. Cogn. Psychol.* 3, 167–176. doi: 10.2478/v10053-008-0022-3

VanRullen, R., and Thorpe, S. J. (2001). Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Comput.* 13, 1255–1284. doi: 10.1162/08997660152002852

VanRullen, R., and Thorpe, S. J. (2002). Surfing a spike wave down the ventral stream. *Vision Res.* 42, 2593–2615. doi: 10.1016/S0042-6989(02)00298-5

Weigelt, S., Singer, W., and Muckli, L. (2007). Separate cortical stages in amodal completion revealed by functional magnetic resonance adaptation. *BMC Neurosci.* 8:70. doi: 10.1186/1471-2202-8-70

Williford, T., and Maunsell, J. H. R. (2006). Effects of spatial attention on contrast response functions in macaque area v4. *J. Neurophysiol.* 96, 40–54. doi: 10.1152/jn.01207.2005

Wyatte, D., Curran, T., and O'Reilly, R. (2012a). The limits of feedforward vision: recurrent processing promotes robust object recognition when objects are degraded. *J. Cogn. Neurosci.* 24, 2248–2261. doi: 10.1162/jocn_a_00282

Wyatte, D., Herd, S., Mingus, B., and O'Reilly, R. (2012b). Competitive inhibition and top-down feedback in binding during object recognition. *Front. Psychol.* 3:182. doi: 10.3389/fpsyg.2012.00182

Zipser, K., Lamme, V. A. F., and Schiller, P. H. (1996). Contextual modulation in primary visual cortex. *J. Neurosci.* 16, 7376–7373.

Zoccolan, D., Cox, D. D., and DiCarlo, J. J. (2005). Multiple object response normalization in monkey inferotemporal cortex. *J. Neurosci.* 25, 8150–8164. doi: 10.1523/JNEUROSCI.2058-05.2005