Repetition Suppression in High and Low Order Areas of Macaque Visual Cortex

A Thesis submitted to the Department of Biological Sciences

At Carnegie Mellon University

By

Nathaniel P. Williams

In partial fulfillment of the requirements for the degree of

Doctor of Philosophy

November, 2021

ACKNOWLEDGEMENTS

Realizing the work presented in this dissertation, and my journey as a scientist, would not have been possible without the guidance and support of a great many people over the years. I would like to acknowledge some of them here. First and foremost, I would like to thank Dr. Carl Olson, my research advisor and mentor throughout this process, for giving me this great opportunity to explore, and for always being willing to sit down and go over any problem no matter how seemingly monumental or trivial. I have learned a great deal simply by listening. I would also like to thank the members of my thesis advisory committee: Dr. Carol Colby, Dr. Sandy Kuhlman, and Dr. Tai Sing Lee. In particular I would like to thank Dr. Lee for giving me the great opportunity to step in and gain more experience in the surgery, as well as for insightful comments on my work, Dr. Kuhlman for excellent guidance and Dr. Colby for great discussions during meetings of the Brain Group journal club. I would also like to thank members of the labs that I shared space with, in particular Dr. Marlene Cohen, Dr. Matt Smith, Dr. Doug Ruff, Josh Alberts, Brittany Bowes and Lily Kramer and a special thanks to Karen McCracken for her care of the animals used in these studies.

I thank the staff in both the Biological Sciences Department and the Center for the Neural Basis of Cognition for keeping everything running behind the scenes, especially Ena Miceli and Melissa Stupka as well as the entire CNBC for being such a great community and source of inspiration and ideas.

I would like to thank my parents for always believing that I could achieve what I set out to do, and my extended family, especially my grandfather Mike Williams, who taught me at a very early age the value of experiment. I've been collecting data ever since. Finally but certainly not least of all I thank my partner for always being there through the ups and downs.

Table of Contents

List of Figures

ABSTRACT

In addition to recognizing objects, the visual system of our brains must adapt to different conditions, such as periods of rapid change, like a busy intersection in a large city, and to periods of stability, such as observing a painting in a museum. A key difference between these conditions is the recency with which visual features are observed. In the case of the painting, one might look at a particular feature for a long time and then return to it a short time later, whereas in the case of the busy intersection things are constantly moving and changing and the same thing is unlikely to appear at the same place again. One of the most robust ways that neurons in the visual system respond to stimulus recency is through repetition suppression. Repetition suppression is the phenomenon whereby a visual sensory neuron will reliably fire fewer spikes in response to the second presentation of an identical stimulus than to its initial presentation. Moreover, repetition suppression is associated with behavioral improvements. Judgments about object properties are faster for repeated versus non-repeated objects. Despite its prevalence and possible impact on behavior, we know little about how repetition suppression arises within the visual system. To shed light on the mechanisms and possible functional significance of repetition suppression we have performed several experiments in awake rhesus macaques recording from single neurons in multiple brain areas at different levels of the visual hierarchy while monkeys viewed sequential displays in which we controlled the repetition of different aspects of the stimulus. We found that surprisingly the degree of suppression depended not only on the properties of the images but also on the preferences of the neuron. We found that repetition suppression probably does not serve as a behaviorally relevant recency signal due its

homogeneous nature. Furthermore, we found no evidence to support that repetition suppression is driven by neuronal fatigue as a mechanism. We also found that surprisingly repetition suppression arises seemingly independently at multiple levels of the visual hierarchy without a clear bottom-up or top-down origin. We have also found evidence that content outside a neuron's classical receptive field can have an impact on repetition suppression in a context dependent manner, suggesting a possible role for lateral connections in the generation of repetition suppression.

CHAPTER I

GENERAL INTRODUCTION

1.1. Repetition suppression

1.1.1 Introduction

Our visual environment is populated with objects that are relatively stable: Objects rarely suddenly disappear, jump to new locations, or change identities. Extensive evidence suggests that our brains are sensitive to this stability. Specifically, regions of the brain that encode object identity – in particular area TE of inferotemporal cortex in the macaque - respond to objects most strongly when they are initially seen, and subsequently respond less as the same objects continue to be observed. This robust phenomenon, known as repetition suppression, is consistently observed in neural responses measured from both the hemodynamic signal using functional magnetic resonance imaging (fMRI) (*1*), in the local field potential (LFP) (*2*) and at the level of single neurons (*3*). Repetition related changes are also found in scalp electroencephalography (EEG) (*4*) and magnetoencephalography (MEG) (*5*). There is also evidence for repetition related adaptation effects in the auditory domain (*5*) as well as in somatosensory cortex and motor regions for repeated movements (*6*), although it's unclear whether these observations are generated by the same or similar underlying mechanisms as repetition suppression in visual cortex. Despite the robust and reproducible effects of stimulus repetition, and that sensitivity to repetition appears to be a fundamental aspect of neural processing, given its ubiquitous nature in multiple recording and sensory modalities, we know little about the mechanisms which give rise to repetition suppression within the visual system, or their functional consequences.

Although fMRI has been used to study repetition suppression, its poor spatial resolution, aggregating over thousands or more neurons, slow temporal resolution, aggregating over tens of

milliseconds, and the unclear relationship between blood oxygen level dependent (BOLD) signal change and spiking activity make it a poor tool for elucidating the neural mechanisms which give rise to repetition suppression. EEG/MEG and LFP have similar issues regarding spatial resolution and their relationship to spiking activity. In the following sections of this chapter, I will explore what is currently known about repetition suppression. I will focus on studies of single unit spiking activity in visual cortex, as this is the methodology used to study repetition suppression in the following chapters.

1.1.2 Repetition suppression was discovered in and is most commonly studied in inferotemporal cortex

Repetition suppression was first demonstrated in macaque inferotemporal cortex during the late 1980s (*3, 7*). Inferotemporal cortex, as the terminus of the ventral stream of visual areas, plays a crucial role in visual object recognition, a role thought to depend on the fact that its neurons respond selectively to complex images and on the fact that their responses can be modified by visual experience (*8*). The subdivision of inferotemporal cortex on which studies of repetition suppression have focused encompasses anterior area TE, spanning the ventral bank of the superior temporal sulcus and the adjacent inferior temporal gyrus at levels roughly 10-20 mm anterior to the ear canals (*9*).

1.1.3 Repetition suppression generalizes to images identical to the adapter except for size or location

Repetition suppression occurs even under conditions in which the location or size of an image changes between the first and second presentations. In an early study utilizing stimuli at the fovea or at 5° eccentricity, the strength of identity-specific suppression was statistically indistinguishable across conditions in which the sample and the test locations were the same or

different (*10*). Likewise, with foveal images presented at sizes of 2° or 4°, neurons exhibiting identity-specific suppression did so regardless of whether sample and test size were the same or different (*10*). These results would seem to indicate that repetition suppression is immune to changes in accidental viewpoint-dependent attributes, such as location and size. Later reports have qualified this conclusion. Suppression does generalize across location but the strength of suppression is reduced with a change in location (*2, 11*). Likewise, suppression does generalize across fourfold changes in size but with the strength of the effect falling off as the size difference increases (*12*). In these studies, suppression was measured as the decrement in firing rate from first to second presentation of an image. This measure cannot distinguish between identity-specific and identity-nonspecific suppression. Thus the results leave open the question: does identity-specific repetition suppression generalize completely across changes of size and location, as suggested by the earlier studies, or does it not? A recent study has resolved this issue by demonstrating that identity-specific suppression, while it indeed generalizes across locations, declines in strength as separation between the adapter and the test image increases (*13*).

1.1.4 Repetition suppression generalizes to new images resembling the adapter

How narrowly is image identity defined in identity-specific suppression? Investigators have addressed this issue by measuring neuronal responses to algorithmically generated images distributed at intervals along a parametric continuum (*2, 14*). The results indicate that neurons exhibit tuning in parametric space and that presentation of a given image as adapter creates a notch in the test-image tuning curve that is centered at the parametric location of the adapter and that spreads to parametrically nearby locations. Why should suppression spread locally over a particular distance in an arbitrary parametric space? One idea, based on the notion that suppression is a network phenomenon, is that cross-suppression should scale with the overlap

between the neuronal populations excited by the two images. To test this idea would require a full-fledged representational dissimilarity analysis based on responses of many neurons to many images (*15*). The only relevant analysis to date was based on single-neuron responses to a few images. The results show that images to which a neuron responds with comparable strength exhibit more cross-suppression than images to which the neuron responds with different strength (*16*). However, cross-suppression between two images matched by this measure is quite weak on average (*17*). It remains to be determined whether, if images were ranked for dissimilarity on the basis of activity in a large enough population, inter-image distance could predict the degree of cross-suppression.

1.1.5 The strength of repetition suppression depends on how recently and often an image has been seen

Repetition suppression is often described as reflecting the familiarity or recency of an image. On one hand, suppression tends to dissipate with the passage of time after an image has been presented, as if it were related to recency (*3, 7, 18-25*). On the other hand, if an image is presented multiple times during a session, then suppression tends to build up, as if it were related to familiarity (*7, 19, 20, 26, 27*). These phenomena are compatible with a mechanism based on fatigue and recovery at a not-as-yet-identified synaptic or somatic site.

1.1.6 Repetition suppression occurs regardless of task context

Repetition suppression appears to depend on the repetition of visual stimulus without regard to task context. It occurs in tasks requiring that visual stimuli be remembered (*3, 7, 12, 16, 18, 20, 22, 26, 28-34*), that they be actively processed without being remembered (*2, 3, 21, 34*) and that they be passively viewed (*2, 7, 11, 14, 17-19, 31, 33, 35-40*). No study characterizing repetition suppression in more than one task has revealed any systematic difference across tasks

(*2, 3, 7, 12, 16, 18, 31, 33, 34*). Repetition suppression may, however, be obscured by signals dependent on the task-relevance of stimuli. For example, some neurons respond with enhanced firing to a repeated image when it is a match to which the monkey must respond but not otherwise (*29*).

1.1.7 Repetition suppression can be disentangled from signals related to working memory

Repetition suppression has often been studied in the context of tasks requiring working memory. Standard tests of working memory require monkeys to hold a sample in mind and to respond to one or more subsequent test stimuli on the basis of whether they match or do not match the sample. Including repeated nonmatches in the test string can be used to force monkeys to engage in working memory, responding to each test on the basis of whether it matches the sample rather than on the basis of whether it is simply a repeat (*29, 41*). The engagement of working memory induces, in inferotemporal cortex, neuronal signals, including match enhancement, that reflect the relation between the test image and the sample (*16, 29, 41*). These signals allow decoding the match status of a test stimulus from population activity and thus could serve to guide task performance (*24, 41, 42*). Decoding accuracy is reduced on error trials, as would be expected if these signals did indeed guide behavior (*41*). Repetition suppression can be reliably measured in such tasks if the repetition status of an image is balanced against its task relevance, for example if the image is a nonmatch item during both its first and second presentations (*29*) or is a sample in both cases (*20*). However, comparing responses elicited by the image presented first as sample and then as match would confound the influence of working memory with the influence of the image's repetition status.

1.1.8 Repetition suppression can be disentangled from identity-nonspecific suppression

Another standard approach to the study of repetition suppression is to present two images, an adapter and a test, in rapid succession during a single trial. In this context, it is possible to distinguish repetition suppression (dependent on the second image's being the same as the first) from identity-nonspecific suppression (occurring for any image). Repetition suppression can be measured as the reduction in the strength of the response to the second image contingent on its matching the first image. However, comparing responses elicited by an image in leading and trailing positions would confound image-nonspecific with image-specific effects.

1.1.9 Repetition suppression is accompanied by an enhancement of oscillatory amplitude

Neurons in inferotemporal cortex are typically described as responding to stimuli with a transient or phasic burst followed by sustained or tonic firing at a reduced rate. However, the dynamics of the response are often more complex. In particular, the response can take the form of a damped 5 Hz oscillation with a peak at approximately 100 ms following stimulus onset, a trough at approximately 200 ms and a second peak at approximately 300 ms (*43*). The oscillatory tendency is accentuated by presenting an image against an already visible backdrop, as would be expected if oscillations arose from resonance in a reciprocally inhibitory circuit mediating divisive normalization (*44*). However resonance could also arise from slow adaptation in a recurrent excitatory network (*45*). Oscillatory responses are evident in population histograms of numerous studies based on presenting adapter and test in rapid succession so as to maximize suppression (*2, 11, 13, 14, 27, 34, 36-39*). As a general rule, although oscillatory activity varies across monkeys (*34, 36, 38*), oscillations are weakest in response to the adapter, stronger in response to a non-matching test and strongest in response to a matching test. Across these conditions, as the net response becomes progressively weaker the oscillatory component

becomes progressively stronger. This may help to explain why it is possible to decode the match status of an image from the temporal profile of the neuronal response (*42*). The fact that repetition enhances the oscillatory component of the response carries information about the mechanism of suppression. It tells us that repetition suppression is likely not the product of attenuated bottom-up input alone. Attenuation of bottom-up input would be expected to act like a reduction in image contrast, which brings down response strength without enhancing the oscillatory component of the response (*37, 46*). The neural trace giving rise to repetition suppression seems to reside at least in part in the resonant circuit giving rise to 5 Hz oscillations.

1.2. Relation of repetition suppression to familiarity, attention, memory and behavior

1.2.1 Repetition suppression may be an early stage of familiarity suppression

Inferotemporal neurons exhibit familiarity suppression, responding less strongly to highly familiar images than to novel images (*18, 25, 47-52*). Familiarity suppression may arise from long-term consolidation of short-term repetition suppression. If so, one might expect a floor effect whereby highly familiar images would be immune to further suppression when repeated. Contrary to this prediction, many studies have utilized familiar images to elicit repetition suppression. Moreover, most reports based on the use of both familiar and novel images note no difference between them with regard to repetition suppression (*7, 14, 18, 21, 31, 53*). One report describes repetition suppression as building up for novel images presented repeatedly during a session but not for familiar images (*20*). However, interpretation is clouded by the fact that data were collected in the context of a delayed-match-to-sample task with familiar images as nonmatch stimuli and novel images as sample and match stimuli, confounding familiarity status with task relevance.

1.2.2 Repetition suppression might reduce the salience of repeated images

If the strength of the neuronal response to an image determines its salience, which is to say its ability to capture attention automatically, then one might imagine that repetition suppression should result in a reduction of salience. There is some reason to believe that this is so. Humans and monkeys alike, when given a choice between viewing a previously seen image and a novel image, direct gaze preferentially to the novel image (*54, 55*). This effect is generally regarded as dependent on recognition of the repeated image, and not just on the reduced strength of the visual response, because it is eliminated by hippocampal lesions (*56-59*). However, it remains possible that low inferotemporal cortex response strength contributes to the effect.

1.2.3 Repetition suppression might support recognition memory

Monkeys and humans are able to recognize hundreds of distinctive images after prolonged sequential exposure, utilizing a form of recognition memory clearly distinct, in its exceedingly large capacity, from working memory (*60-62*). Early descriptions of repetition suppression considered the phenomenon to constitute a neural substrate for this form of memory. This idea is plausible and is supported by limited evidence, notably the finding that recognition memory and inferotemporal repetition suppression decline in parallel with an increase in the number of trials intervening between the first and second presentations of an image (*22*). However, no study published to date has posed the question whether, on trials in which the monkey fails to recognize an image, repetition suppression is reduced, as would be expected if it played a causal role in recognition. This approach is potentially practicable as demonstrated by its successful application to repetition signals in the hippocampus (*54*) and to match signals in inferotemporal cortex (*41*). Moreover, two studies using pharmacological and neurological interventions have suggested that recognition memory is dissociable from repetition suppression.

In one study, systemic administration of scopolamine, an anticholinergic agent acting at muscarinic receptors, impaired delayed-match-to-sample performance without affecting repetition suppression in inferotemporal cortex (*30*). This finding can be reconciled with the notion that repetition suppression mediates recognition on the assumption that scopolamine interfered with readout of the suppression signal downstream from inferotemporal cortex. The second study was carried out in monkeys following optic chiasm and posterior corpus callosum transection (*32*). Following this surgery, the right and left inferotemporal cortices continue to communicate via the intact anterior commissure whereas low-order ventral-stream visual areas are isolated from their counterparts in the opposite hemisphere (*63*). When a sample is presented to the right eye, activation propagates up the right ventral-stream hierarchical chain, presumably leaving a refractory trace, and crosses from the right to the left inferotemporal cortex. When a test is presented to the left eye, activity propagates up the left ventral-stream chain, bypassing the refractory trace left by the preceding sample, and activates left inferotemporal neurons without any sign of repetition suppression. Nevertheless, monkeys report the match with reasonable accuracy.

1.2.4 Repetition suppression probably does not underlie repetition priming

Humans making decisions about visual stimuli, for example naming a picture or indicating whether a string of letters is a word, exhibit repetition priming, responding more quickly the second time a given stimulus is presented (*64, 65*). Monkeys also exhibit repetition priming (*21*). The enhanced processing of the stimulus that underlies repetition priming might conceivably arise from repetition suppression. In the so-called "sharpening" model, repetition suppression is stronger among neurons weakly responsive to the stimulus than among neurons strongly responsive to it, so that, when it is repeated, the firing of the selective neurons stands out

especially strongly from background activity (*20, 66, 67*). Sharpening does indeed appear to occur when images are rendered familiar through repeated viewing over days and weeks (*48, 51, 68*). However, efforts to demonstrate sharpening in connection with short-term repetition have been fruitless (*2, 21*). Moreover, numerous reports explicitly note or implicitly support the existence of a "scaling" principle according to which suppression is proportional to response strength (*2, 16, 20, 21, 26, 27, 39*). The scaling principle is directly contrary to the prediction of the sharpening model that weak responses should be disproportionately reduced. Theoretical efforts to reconcile repetition suppression with repetition priming have fallen back on the principle that suppression may be accompanied by some secondary change favorable to priming such as increased gamma synchrony (*69*).

1.3. Open questions regarding repetition suppression

1.3.1 It is not known whether neurons differ with regard to the strength of repetition suppression

The question whether some or all inferotemporal neurons exhibit repetition suppression is important because it has a direct bearing on the issue of whether suppression could be put to use in the service of recognition memory. In the absence of disambiguating information, no downstream area could differentiate between a weak population response arising from repetition suppression and a weak population response arising from inefficacy of the image due to some intrinsic feature such as low contrast (*46*). Desimone and colleagues proposed, as a solution to this problem, the existence of neurons that do not exhibit repetition suppression (*20, 24*). They suggested that strong firing by non-suppressing neurons, combined with weak firing by suppressing neurons, could signal unambiguously that an image has been repeated. However, there is as yet no convincing evidence for this scheme. When neurons are characterized

individually with an index such that greater suppression is indicated by greater positive values, the distribution of indices invariably has a positive mean and a flank that extends into the negative domain but without any obvious deviation from unimodality (*2, 14, 19, 22, 34, 36, 38, 39*). Whether the distribution reflects neuronal variability or measurement noise is unclear. Statistical tests applied to data from individual neurons invariably reveal some cases in which suppression does not achieve significance but this could be due to lack of power. A few neurons are commonly found to exhibit statistically significant repetition enhancement (*3, 10, 20, 22-24, 29*). However, interpretation of these cases is difficult for two reasons. First, some false positives are expected at the comparatively permissive significance thresholds typically employed. Second, in certain task contexts, enhancement could reflect factors correlated with but separate from image repetition, for example the match status of the image in delayed-match-to-sample context. At the present time, we simply do not know whether neurons differ systematically with respect to the strength or sign of repetition-related effects.

1.3.2 Evidence for repetition suppression in low-order visual areas is limited

Neurons in low-order visual areas exhibit adaptation when tested with displays such as drifting gratings drawn from the classical psychophysical toolbox (*70-75*). However the relation between adaptation (as based on fatigue of channels for orientation, spatial frequency and motion direction) and repetition suppression (as based on storage of the trace of a particular image) is poorly understood. Functional imaging in monkeys has revealed no sign of repetition suppression for natural images in lower-order areas, including V1, V2 and V4, while demonstrating its robust occurrence in inferotemporal cortex (*76*). The absence of repetition suppression at the level of the BOLD signal does not, however, rule out its presence at the level of neuronal activity. A few studies have indeed produced results supportive of the idea that

neurons in low-order areas exhibit effects resembling repetition suppression. Neurons in area V4 give habituating responses to rapid strings of letter-like stimuli (*77*) and respond more weakly to a simple geometric shape when its contour curvature matches than when it does not match the contour curvature of an immediately preceding image (*78*). Likewise, neurons in area V1 exhibit conjunction-specific adaptation to a plaid pattern formed by combining orthogonal gratings (*79*). In light of these observations, it seems possible that neurons in low-order areas would exhibit repetition suppression under standard conditions employed in studies of inferotemporal cortex including the use of complex natural images.

1.3.3 Evidence for repetition suppression outside the visual system is limited

Hippocampus. It is plausible that neurons of the hippocampal system might be sensitive to image repetition because hippocampal lesions are associated with a deficit in visual recognition memory (*80*). Accordingly, several groups have assessed the impact of image repetition on neuronal responses in the hippocampus and its extended system. Early reports described at most a small fraction of hippocampal neurons as exhibiting sensitivity to repetition, with the sign of the effect inconsistent across studies (*7, 18, 31, 81, 82*); repetition-sensitivity likewise was observed in only a small fraction of neurons in the amygdala and basal forebrain (*83, 84*). A more recent report has indicated that around two-thirds of hippocampal neurons respond to visual stimuli, with either an increase or a decrease in firing rate, and that around a third of these neurons fire differentially to the first or second presentation of an image (*54*). However, approximately as many neurons favor the second presentation as favor the first. The bulk of the evidence accordingly seems to indicate that hippocampal neurons, unlike inferotemporal neurons, do not exhibit consistent repetition suppression.

Prefrontal cortex. Two cortical regions outside the ventral stream visual system have exhibited repetition suppression in an fMRI-based monkey study: parietal area LIP and a region of dorsolateral prefrontal cortex lying between the principal sulcus and the ventral arm of the arcuate sulcus (*76*). No single-neuron study has yet characterized repetition suppression in LIP, whereas two studies have focused on the implicated region of dorsolateral prefrontal cortex. Both studies characterized neuronal activity in the context of delayed-match-to-sample tasks, with the attendant challenge of disentangling repetition suppression from neuronal signals related to working memory. Desimone and colleagues analyzed the impact on neuronal response strength of repeating a nonmatch item, thus holding task-relevance constant while manipulating recency (*85*). They found that a minority of neurons exhibit repetition suppression, with suppression a weak trend at the population level. Vogels and colleagues compared neuronal responses to test stimuli that matched or did not match the sample in shape or size (*12*). They observed weak match suppression at the population level and noted that suppression early in the response was greatest when the difference in size between adapter and test was least. It appears from these results that repetition suppression is weaker and more entangled with signals related to task-relevance in prefrontal than in inferotemporal cortex. If repetition suppression is propagated from inferotemporal to prefrontal cortex, then, due to transmission delay, prefrontal suppression should develop later. Comparison of population PSTHs reveals, however, no clear tendency toward later onset (*12*).

1.3.4 The mechanistic underpinnings of repetition suppression are unknown

Prediction suppression cannot explain repetition suppression. If two images have been presented in sequence repeatedly over the course of weeks, so that the first image dependably predicts the second, then inferotemporal neurons respond only weakly to the second image when it follows

the first (*49, 86-88*). This effect has been termed prediction suppression. It is possible, given the relative stability of the visual environment, that presenting an image once constitutes a prediction that it will appear again. If so, then repetition suppression could result from prediction suppression. Kaliukhovich and Vogels assessed this possibility by recording from inferotemporal neurons during blocks of trials differing with regard to the probability that an image would be repeated (*36*). They found that repetition suppression occurred with the same strength regardless of whether repetition probability was high or low. Insofar as the brain made predictions on the basis of within-block statistics, this outcome indicates that suppression was a product of repetition as such and not of prediction confirmation.

Neuronal fatigue cannot explain repetition suppression. Repetition suppression might arise from fatigue, due to spike-frequency adaptation of the very neurons in which the effect is measured. The idea that inferotemporal neurons are subject to fatigue is plausible because they typically respond to a visual stimulus with a phasic burst that declines to a lower tonic firing rate. However, mechanisms other than fatigue, for example recurrent inhibition, could underlie the tonic-phasic pattern. Vogels and colleagues directly tested the fatigue-based account of repetition suppression by probing whether two images equally effective in driving neuronal activity produced significant cross-suppression (*27*). They found that cross-suppression was minor, thus ruling out fatigue at the level of the recorded neuron as an explanation for repetition suppression. It is theoretically possible that repetition suppression arises from fatigue among neurons at an earlier processing stage, but this proposal is implausible because it rests on the dual assumptions that neurons at the earlier processing stage possess sufficient image selectivity to mediate identity-specific suppression and that they are subject to a form of fatigue not evident in inferotemporal cortex.

Repetition suppression presumably arises from synaptic plasticity. In the absence of an explanation based on neuronal fatigue, we must assume that repetition suppression is a product of stimulation-induced synaptic plasticity. However, we know nothing about the location or nature of the changes in synaptic strength underlying the effect. They might affect feedforward, feedback or horizontal connections. They might occur within inferotemporal cortex itself or at an earlier processing stage. They might depend, like fatigue, solely on the strength with which a synapse is activated, or, alternatively, on the combination of pre- and post-synaptic activation. They might occur at all activated synapses or act only on a subset of circuit components.

1.4 Experimental questions tested

In the following three chapters, several experimental questions regarding repetition suppression will be addressed. In Chapter II, the overarching question is:

1. For neurons in TE, does repetition suppression operate at the level of individual features, or at the level of the whole image?

Additional questions addressed in Chapter II include:

2. Are there separate populations of suppressing and non-suppressing neurons in TE?
3. Does firing rate fatigue play a role in repetition suppression in TE?
4. How specific to image repetition is the oscillatory response dynamic?

In Chapter III, the overarching question is:

1. Does repetition suppression occur in V2 under the same stimulation parameters used in IT?

Additional questions addressed in Chapter III include:

2. What is the degree of image specificity and location specificity of repetition suppression in V2, and do these correspond to signals seen in TE?

3. Do suppression signals in V2 precede those in TE (suggesting a bottom-up contribution) or lag those in TE (suggesting a top-down contribution)?

4. Can content outside the classical receptive field of V2 neurons contribute to repetition suppression in V2?

5. Does V2 exhibit an oscillatory response dynamic as seen in TE, and if so, how specific to image repetition is the oscillatory response dynamic in V2?

In Chapter IV, the overarching question is:

1. Does repetition suppression occur in V2 when content in the center and content in the surround of an image is repeated independently?

Additional questions addressed in Chapter IV include:

2. Do suppression signals induced by content in the surround occur at longer latency, suggesting a possible top-down signal?

The motivation for each of these questions and the paradigm for testing them is outlined at the beginning of each chapter, and the results of each experiment summarized.

1.5 REFERENCES

1. Z. Kourtzi, N. Kanwisher, Representation of perceived object shape by the human lateral occipital complex. *Science* **293**, 1506-1509 (2001).
2. W. De Baene, R. Vogels, Effects of Adaptation on the Stimulus Selectivity of Macaque Inferior Temporal Spiking Activity and Local Field Potentials. *Cerebral Cortex* **20**, 2145-2165 (2010).

3. G. C. Baylis, E. T. Rolls, Responses of neurons in the inferior temporal cortex in short term and serial recognition memory tasks. *Exp Brain Res* **65**, 614-622 (1987).
4. M. E. Nagy, M. D. Rugg, Modulation of Event-Related potentials by word repetition: The effects of Inter-Item lag. *Psychophysiology* **26**, 431-436 (1989).
5. A. Todorovic, F. van Ede, E. Maris, F. P. de Lange, Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *Journal of Neuroscience* **31**, 9118-9123 (2011).
6. A. F. d. C. Hamilton, S. T. Grafton, Repetition suppression for performed hand gestures revealed by fMRI. *Hum Brain Mapp* **30**, 2898-2906 (2009).
7. M. W. Brown, F. A. Wilson, I. P. Riches, Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res* **409**, 158-162 (1987).
8. L. G. Ungerleider, M. Mishkin, in *Analysis of Visual Behavior,* D. J. Ingle, M. A. Goodale, R. J. W. Mansfield, Eds. (MIT Press Cambridge, MA, 1982), pp. 549-586.
9. K. S. Saleem, N. K. Logothetis, *Atlas of the Rhesus Monkey Brain in Stereotaxic Coordinates*. (Academic Press, Amsterdam, 2007).
10. A. Lueschow, E. K. Miller, R. Desimone, Inferior temporal mechanisms for invariant object recognition. *Cerebral cortex (New York, N.Y. : 1991)* **4**, 523-531 (1994).
11. P. Kuravi, R. Vogels, Effect of adapter duration on repetition suppression in inferior temporal cortex. *Sci Rep* **7**, 3162 (2017).
12. B.-E. Verhoef, G. Kayaert, E. Franko, J. Vangeneugden, R. Vogels, Stimulus Similarity-Contingent Neural Adaptation Can Be Time and Cortical Area Dependent. *The Journal of Neuroscience* **28**, 10631-10640 (2008).
13. F. Fabbrini, R. Vogels, Within- and between-hemifield generalization of repetition suppression in inferior temporal cortex. *J Neurophysiol* **125**, 120-139 (2021).
14. Y. Liu, S. O. Murray, B. Jagadeesh, Time Course and Stimulus Dependence of Repetition-Induced Response Suppression in Inferotemporal Cortex. *Journal of Neurophysiology* **101**, 418-436 (2009).
15. N. Kriegeskorte *et al.*, Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* **60**, 1126-1141 (2008).
16. E. Miller, L. Li, R. Desimone, Activity of neurons in anterior inferior temporal cortex during a short- term memory task. *The Journal of Neuroscience* **13**, 1460-1478 (1993).
17. H. Sawamura, G. A. Orban, R. Vogels, Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. *Neuron* **49**, 307-318 (2006).
18. F. L. Fahy, I. P. Riches, M. W. Brown, Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. *Exp Brain Res* **96**, 457-472 (1993).
19. D. A. Kaliukhovich, R. Vogels, Neurons in macaque inferior temporal cortex show no surprise response to deviants in visual oddball sequences. *J Neurosci* **34**, 12801-12815 (2014).
20. L. Li, E. K. Miller, R. Desimone, The representation of stimulus familiarity in anterior inferior temporal cortex. *J Neurophysiol* **69**, 1918-1929 (1993).
21. D. B. T. McMahon, C. R. Olson, Repetition Suppression in Monkey Inferotemporal Cortex: Relation to Behavioral Priming. *Journal of Neurophysiology* **97**, 3532-3543 (2007).
22. T. Meyer, N. C. Rust, Single-exposure visual memory judgments are reflected in inferotemporal cortex. *Elife* **7**, (2018).
23. E. K. Miller, P. M. Gochin, C. G. Gross, Habituation-like decrease in the responses of neurons in inferior temporal cortex of the macaque. *Vis Neurosci* **7**, 357-362 (1991).

24. E. K. Miller, L. Li, R. Desimone, Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* **13**, 1460-1478 (1993).
25. S. Sobotka, J. L. Ringo, Investigation of long-term recognition and association memory in unit responses from inferotemporal cortex. *Exp Brain Res* **96**, 28-38 (1993).
26. E. Miller, L. L, R. Desimone, A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* **254**, 1377-1379 (1991).
27. H. Sawamura, G. A. Orban, R. Vogels, Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. *Neuron* **49**, 307-318 (2006).
28. A. Lueschow, E. K. Miller, R. Desimone, Inferior Temporal Mechanisms for Invariant Object Recognition. *Cerebral Cortex* **4**, 523-531 (1994).
29. E. K. Miller, R. Desimone, Parallel neuronal mechanisms for short-term memory. *Science* **263**, 520-522 (1994).
30. E. K. Miller, R. Desimone, Scopolamine affects short-term memory but not inferior temporal neurons. *Neuroreport* **4**, 81-84 (1993).
31. I. P. Riches, F. A. Wilson, M. W. Brown, The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. *J Neurosci* **11**, 1763-1779 (1991).
32. S. Sobotka, J. L. Ringo, Mnemonic responses of single units recorded from monkey inferotemporal cortex, accessed via transcommissural versus direct pathways: a dissociation between unit activity and behavior. *J Neurosci* **16**, 4222-4230 (1996).
33. R. Vogels, G. Sary, G. A. Orban, How task-related are the responses of inferior temporal neurons? *Visual Neuroscience* **12**, 207-214 (1995).
34. K. Vinken, H. P. Op de Beeck, R. Vogels, Face Repetition Probability Does Not Affect Repetition Suppression in Macaque Inferotemporal Cortex. *J Neurosci* **38**, 7492-7504 (2018).
35. T. Abel, R. Havekes, J. M. Saletin, M. P. Walker, Sleep, plasticity and memory from molecules to whole-brain networks. *Current biology : CB* **23**, R774-788 (2013).
36. D. A. Kaliukhovich, R. Vogels, Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral cortex (New York, N.Y. : 1991)* **21**, 1547-1558 (2011).
37. F. Fabbrini *et al.*, Probing the Mechanisms of Repetition Suppression in Inferior Temporal Cortex with Optogenetics. *Current biology : CB* **29**, 1988-1998.e1984 (2019).
38. D. A. Kaliukhovich, R. Vogels, Stimulus repetition affects both strength and synchrony of macaque inferior temporal cortical activity. *J Neurophysiol* **107**, 3509-3527 (2012).
39. D. A. Kaliukhovich, R. Vogels, Divisive Normalization Predicts Adaptation-Induced Response Changes in Macaque Inferior Temporal Cortex. *J Neurosci* **36**, 6116-6128 (2016).
40. P. Kuravi, R. Vogels, GABAergic and cholinergic modulation of repetition suppression in inferior temporal cortex. *Sci Rep* **8**, 13160 (2018).
41. M. Pagan, L. S. Urban, M. P. Wohl, N. C. Rust, Signals in inferotemporal and perirhinal cortex suggest an untangling of visual target information. *Nat Neurosci* **16**, 1132-1139 (2013).
42. E. N. Eskandar, B. J. Richmond, L. M. Optican, Role of inferior temporal neurons in visual memory. I. Temporal encoding of information about visual images, recalled images, and behavioral context. *J Neurophysiol* **68**, 1277-1295 (1992).
43. J. E. Rollenhagen, C. R. Olson, Low-Frequency Oscillations Arising From Competitive Interactions Between Visual Stimuli in Macaque Inferotemporal Cortex. *Journal of Neurophysiology* **94**, 3368-3387 (2005).
44. S. Moldakarimov, J. E. Rollenhagen, C. R. Olson, C. C. Chow, Competitive Dynamics in Cortical Responses to Visual Stimuli. *Journal of Neurophysiology* **94**, 3388-3396 (2005).

45. S. Lim, Mechanisms underlying sharpening of visual response dynamics with familiarity. *Elife* **8**, (2019).
46. M. W. Oram, Contrast induced changes in response latency depend on stimulus specificity. *J Physiol Paris* **104**, 167-175 (2010).
47. B. Anderson, D. L. Sheinberg, Effects of temporal context and temporal expectancy on neural activity in inferior temporal cortex. *Neuropsychologia* **46**, 947-957 (2008).
48. D. J. Freedman, M. Riesenhuber, T. Poggio, E. K. Miller, Experience-Dependent Sharpening of Visual Shape Selectivity in Inferior Temporal Cortex. *Cerebral Cortex* **16**, 1631-1644 (2006).
49. T. Meyer, S. Ramachandran, C. R. Olson, Statistical Learning of Serial Visual Transitions by Neurons in Monkey Inferotemporal Cortex. *The Journal of Neuroscience* **34**, 9332-9337 (2014).
50. R. E. B. Mruczek, D. L. Sheinberg, Context Familiarity Enhances Target Processing by Inferior Temporal Cortex Neurons. *The Journal of Neuroscience* **27**, 8533-8545 (2007).
51. L. Woloszyn, D. L. Sheinberg, Effects of Long-Term Visual Experience on Responses of Distinct Classes of Single Units in Inferior Temporal Cortex. *Neuron* **74**, 193-205 (2012).
52. J. Z. Xiang, M. W. Brown, Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* **37**, 657-676 (1998).
53. E. K. Miller, L. Li, R. Desimone, A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* **254**, 1377-1379 (1991).
54. M. J. Jutras, E. A. Buffalo, Recognition memory signals in the macaque hippocampus. *Proceedings of the National Academy of Sciences of the United States of America* **107**, 401-406 (2010).
55. R. D. McKee, L. R. Squire, On the development of declarative memory. *J Exp Psychol Learn Mem Cogn* **19**, 397-404 (1993).
56. S. Nemanic, M. C. Alvarado, J. Bachevalier, The hippocampal/parahippocampal regions and recognition memory: insights from visual paired comparison versus object-delayed nonmatching in monkeys. *J Neurosci* **24**, 2013-2026 (2004).
57. O. Pascalis, J. Bachevalier, Neonatal aspiration lesions of the hippocampal formation impair visual recognition memory when assessed by paired-comparison task but not by delayed nonmatching-to-sample task. *Hippocampus* **9**, 609-616 (1999).
58. J. M. Reed, L. R. Squire, Impaired recognition memory in patients with lesions limited to the hippocampal formation. *Behav Neurosci* **111**, 667-675 (1997).
59. S. M. Zola *et al.*, Impaired recognition memory in monkeys after damage limited to the hippocampal region. *J Neurosci* **20**, 451-463 (2000).
60. A. A. Wright, Santiago, H.C., Sands, S.F., Kendrick, D.F., and Cook, R.G. , Memory Processing of serial lists by pigeons, monkeys, and people. *Science* **229**, 287-289 (1985).
61. S. F. Sands, A. A. Wright, Serial probe recognition performance by a rhesus monkey and a human with 10- and 20-item lists. *Journal of Experimental Psychology: Animal Behavior Processes* **6**, 386-396 (1980).
62. S. F. Sands, A. A. Wright, Primate Memory: Retention of Serial List Items by a Rhesus Monkey. *Science* **209**, 938-940 (1980).
63. J. H. Kaas, in *Epilepsy and the Corpus Callosum 2,* A. G. Reeves, D. W. Roberts, Eds. (Springer US, Boston, MA, 1995), pp. 15-27.
64. R. N. Henson, Neuroimaging studies of priming. *Prog Neurobiol* **70**, 53-81 (2003).
65. E. Tulving, D. L. Schacter, Priming and human memory systems. *Science* **247**, 301-306 (1990).
66. R. Desimone, Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences* **93**, 13494-13499 (1996).
67. C. L. Wiggs, A. Martin, Properties and mechanisms of perceptual priming. *Curr Opin Neurobiol* **8**, 227-233 (1998).

68. S. Lim *et al.*, Inferring learning rules from distributions of firing rates in cortical neurons. *Nat Neurosci* **18**, 1804-1810 (2015).
69. S. J. Gotts, C. C. Chow, A. Martin, Repetition Priming and Repetition Suppression: A Case for Enhanced Efficiency Through Neural Synchronization. *Cogn Neurosci* **3**, 227-237 (2012).
70. C. W. Clifford *et al.*, Visual adaptation: neural, psychological and computational aspects. *Vision research* **47**, 3125-3131 (2007).
71. A. Kohn, Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. *Journal of Neurophysiology* **97**, 3155-3164 (2007).
72. M. Snow, R. Coen-Cagli, O. Schwartz, Adaptation in the visual cortex: a case for probing neuronal populations with natural stimuli. *F1000Res* **6**, 1246 (2017).
73. S. G. Solomon, A. Kohn, Moving sensory adaptation beyond suppressive effects in single neurons. *Current biology : CB* **24**, R1012-1022 (2014).
74. M. A. Webster, Visual Adaptation. *Annual review of vision science* **1**, 547-567 (2015).
75. C. J. Whitmire, G. B. Stanley, Rapid Sensory Adaptation Redux: A Circuit Perspective. *Neuron* **92**, 298-315 (2016).
76. H. Sawamura, S. Georgieva, R. Vogels, W. Vanduffel, G. A. Orban, Using functional magnetic resonance imaging to assess adaptation and size invariance of shape processing by humans and monkeys. *J Neurosci* **25**, 4294-4306 (2005).
77. B. C. Motter, Modulation of Transient and Sustained Response Components of V4 Neurons by Temporal Crowding in Flashed Stimulus Sequences. *Journal of Neuroscience* **26**, 9683-9694 (2006).
78. K. M. Müller, M. Wilke, D. A. Leopold, Visual adaptation to convexity in macaque area V4. *Neuroscience* **161**, 655-662 (2009).
79. M. Carandini, H. B. Barlow, L. P. O'Keefe, A. B. Poirson, J. A. Movshon, Adaptation to contingencies in macaque primary visual cortex. *Philos Trans R Soc Lond B Biol Sci* **352**, 1149-1154 (1997).
80. D. M. Freed, S. Corkin, N. J. Cohen, Forgetting in H.M.: a second look. *Neuropsychologia* **25**, 461-471 (1987).
81. E. T. Rolls, P. M. Cahusac, J. D. Feigenbaum, Y. Miyashita, Responses of single neurons in the hippocampus of the macaque related to recognition memory. *Exp Brain Res* **93**, 299-306 (1993).
82. E. T. Rolls *et al.*, Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J Neurosci* **9**, 1835-1845 (1989).
83. F. A. Wilson, E. T. Rolls, Neuronal responses related to the novelty and familarity of visual stimuli in the substantia innominata, diagonal band of Broca and periventricular region of the primate basal forebrain. *Exp Brain Res* **80**, 104-120 (1990).
84. F. A. Wilson, E. T. Rolls, The effects of stimulus novelty and familiarity on neuronal activity in the amygdala of monkeys performing recognition memory tasks. *Exp Brain Res* **93**, 367-382 (1993).
85. E. K. Miller, C. A. Erickson, R. Desimone, Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* **16**, 5154-5167 (1996).
86. T. Meyer, C. R. Olson, Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences* **108**, 19401-19406 (2011).
87. S. Ramachandran, T. Meyer, C. R. Olson, Prediction suppression in monkey inferotemporal cortex depends on the conditional probability between images. *Journal of Neurophysiology* **115**, 355-362 (2016).
88. S. Ramachandran, T. Meyer, C. R. Olson, Prediction suppression and surprise enhancement in monkey inferotemporal cortex. *J Neurophysiol* **118**, 374-382 (2017).

CHAPTER II

CONTRIBUTION OF INDIVIDUAL FEATURES TO REPETITION SUPPRESSION IN MACAQUE INFEROTEMPORAL CORTEX

2.1 ABSTRACT

When an image is presented twice in succession, first as adapter and then as test, neurons in area TE of macaque inferotemporal cortex exhibit repetition suppression, responding less strongly to the second presentation than to the first. Suppression is known to occur if the adapter and the test image are subtly different from each other. However, it is not known whether cross-suppression occurs between images that are radically different from each other but that share a subset of features. To explore this issue, we measured repetition suppression using colored shapes as stimuli. On interleaved trials, the test image might be identical to the adapter, might share its shape alone, might share its color alone or might differ from it totally. At the level of the neuronal population as a whole, suppression was deepest when adapter and test were identical, intermediate when they shared only one attribute and minimal when they shared neither attribute. The depth of suppression when the images were identical was slightly greater than could be explained by linear summation of effects observed when shape alone or color alone was repeated. At the level of the individual neuron, the degree of suppression depended, surprisingly, not only on the properties of the two images but also on the preferences of the neuron. Suppression was deeper when the repeated color or shape was preferred by the neuron than when it was not. We show that this effect could arise either from feature-specific adaptation or from adapter-induced fatigue. We argue provisionally for a feature-based explanation and against a fatigue-based explanation on grounds of our own observations and previous findings. We report two additional fundamental observations. First, repetition suppression involves not

only weakening of the response to the test image but also induction of a trough-rebound response dynamic. Second, there is no obvious difference among TE neurons with regard to the strength of repetition suppression.

2.2 INTRODUCTION

Neurons in macaque inferotemporal cortex respond less strongly to a visual stimulus when it is repeated than on its first presentation. This phenomenon is termed repetition suppression. Repetition suppression was first described (*1, 2*) in anterior area TE (*3*), a division of inferior temporal cortex on which most subsequent studies of the phenomenon have focused. Repetition suppression appears to be the product of visual adaptation unmodulated by cognitive influences. It has been observed in tasks requiring that visual stimuli be remembered (*1, 2, 4-16*), that they be actively processed (*1, 16-18*) and that they simply be passively viewed (*2, 4, 11, 14, 17, 19-28*). No study characterizing the phenomenon in more than one task has revealed any systematic cross-task difference (*1, 2, 4, 8, 11, 13, 14, 16, 17*).

Repetition suppression does not depend on a precise match between the successive images. For example, it survives changes in location (*6, 17, 20, 25*) and size (*6, 13*). However, these alterations do not affect image identity. Parametric manipulations of image identity do reduce suppression, with the effect becoming progressively weaker as the difference between the adapter and the test image increases (*17, 27*). However, it is not known what general principle governs the transition from cases in which cross-suppression occurs to cases in which it does not. In particular, it is not known whether cross-suppression would occur under conditions in which some features of an image are drastically altered while others remain intact. To address this issue, we recorded neuronal responses to adapter and test images containing colors and shapes in

particular combinations. On any given trial, a test image might differ from the adapter in shape, in color, in neither attribute or in both.

At the level of the average activity of the neuronal population, suppression was deepest when test and adapter were identical, intermediate when they shared only one attribute and most shallow when they shared neither attribute. There was a weak but significant tendency for suppression to be deeper when both attributes were repeated than would have been predicted by linear summation of shape-specific and color-specific effects. These straightforward population-level effects represented the sum of more complex neuron-level processes. In particular, at the level of the individual neuron, the depth of suppression did not depend solely on the degree of similarity between the adapter and the test image but also was sensitive to whether the repeated shape or color provided strong drive to the neuron. The neuron-level results allow for two possible explanations. According to the first account, repetition suppression is feature-based: presenting a given shape or color in the adapter fractionally reduces drive from that shape or color when presented in the test image. According to the second account, image-specific repetition suppression is accompanied by a fatigue-based effect: suppression depends on the degree of similarity between the test-image and the adapter and also on the strength of the response to the adapter.

2.3 MATERIALS AND METHODS

2.3.1 Subjects

The experimental subjects were two adult rhesus macaques (*macaca mulatta*) here designated M1 (male; 12 kg; laboratory designation Ol) and M2 (male; 6 kg; laboratory designation Sc). All procedures were in accordance with guidelines set forth by the United States

Public Health Service Guide for the Care and Use of Laboratory Animals and were approved by the Carnegie Mellon University IACUC.

2.3.2 Task

Each trial began with presentation of a fixation spot. Once the monkey had attained fixation, the following displays appeared in sequence: fixation spot (320 ms), adapter stimulus (320 ms), fixation spot (320 ms), test stimulus (320 ms), fixation spot (320 ms). The monkey was required to maintain gaze, within a window subtending 1.4-2.1° throughout the trial. Gaze was monitored with an ISCAN video-based eye tracking system. Upon successful completion of a trial, a juice reward was delivered. Any fixation break terminated the trial and triggered onset of a checkerboard display which remained visible for two seconds or until the monkey had fixated it for a cumulative duration of 300 ms, whichever came first. Behavior was monitored and stimulus presentation and reward delivery were controlled by a PC running NIMH Cortex.

2.3.3 Stimuli

All stimuli were centered at the midpoint of an LCD monitor with a 60 Hz refresh rate at a viewing distance of 32 cm. The fixation spot was a white disk 0.3° in diameter. The adapter and test stimuli were Sloan letters, subtending 5.3° of visual angle, rendered in equiluminant saturated colors. The library of available stimuli included six letters selected for minimal confusability by human observers (D, F, G, H, W and X)) rendered in six colors spaced at equal intervals around the color wheel. Before beginning any recording session, we assessed neuronal shape and color selectivity by monitoring responses elicited by all six letters, rendered in white, and by disks of all six colors. We selected for use in a given session four shapes and four colors that elicited the strongest responses as judged by inspection of online raster and histogram displays. Arbitrarily numbering the shapes and colors from 1 to 4, we constructed from them two

tetrads of stimuli. Tetrad A consisted of all stimuli obtained by rendering shapes 1 and 2 in colors 1 and 2. Tetrad B consisted of all stimuli obtained by rendering shapes 3 and 4 in colors 3 and 4.

2.3.4 Session structure

From the four images in each tetrad, 16 adapter-test sequences could be created. Among these were four cases in which shape alone repeated, four in which color alone repeated, four in which both attributes repeated and four in which neither attribute repeated. Each block of 32 successful trials contained one instance of each adapter-test sequence for each tetrad. Sequencing was pseudo-random subject to the constraint that the adapter and test were drawn from different tetrads on alternate trials. We adopted this design to minimize the potential for cross-trial carryover of image-specific repetition suppression. Trials terminated due to fixation-break were repeated to ensure that one trial under each of the 32 conditions was completed successfully. A full run consisted of four successive blocks and thus of 128 successful trials encompassing four repetitions of each possible adapter-test sequence for each tetrad.

2.3.5 Neurophysiological data collection

A vertically oriented cylindrical Cilux recording chamber with an inner diameter of approximately 2 cm (Crist Instrument Co. Inc., Hagerstown, MD) was implanted over the left hemisphere of each monkey with its center approximately 16 mm anterior and 22 mm lateral to Horsley-Clarke zero. Chamber placement was guided by pre-surgical T1-weighted structural MRI scans in a 4.7-Tesla scanner and was confirmed by post-surgical scans. The chamber gave access to cortex both in the ventral bank of the superior temporal sulcus and on the adjacent inferior temporal gyrus. The recording sites, as judged by reference to standard maps of cytoarchitecturally defined areas, lay within anterior area TE (*29*) and did not encroach on

perirhinal cortex (*30*). On each recording day, a cylindrical plug containing guide holes arranged at 1 mm spacing in a square grid was inserted in the chamber (Crist Instrument Co. Inc., Hagerstown, MD). A dura-penetrating guide tube was inserted through one guide hole and a single varnish coated tungsten microelectrode (FHC Inc., Bowdoin, ME) with an impedance of 0.1-5 MΩ at 1kHz was advanced through the guide tube by use of a hydraulic micromanipulator (M0-10; Narishige International Inc., East Meadow, NY). Upon encounter with neurons giving phasic responses to visual stimuli, recording commenced. Threshold-crossing events, sampled at 40 KHz, were digitally recorded and stored for offline sorting by a Plexon MAP system (Plexon Inc., Dallas, TX).

2.3.6 Database

We sorted waveforms from each session using a PCA-based approach implemented by Plexon Offline Sorter (Plexon Inc., Dallas, TX). All further steps of analysis were carried out in Matlab on time-stamped action potential markers. We first analyzed data from each neuron to ensure that it met an arbitrary criterion for visual responsiveness. Only if a one-tailed t-test comparing mean firing rate 100-200 ms after adapter-stimulus onset to mean firing rate 100-200 ms before adapter-stimulus onset yielded an outcome of $p < 0.05$ did we include the neuron in the database for subsequent analyses.

2.3.7 Models

We assessed the ability of four families of models to predict the strength of the response to the test image as a function of test shape, test color, shape match, color match, exact match and fatigue. Two sets of models were of direct interest. These embodied image-based and same-feature-based suppression. Cross-feature-based and null models were included for completeness. The full model in each family had the following form:

Image-based model:

$T = (I_{shape} + I_{color})\bullet(1 - M_{shape} - M_{color} - M_{exact}) - A\bullet M_{fatigue}$

Same-feature-based model:

$T = I_{shape}(1- M_{shape} - M_{exact}) + I_{color}(1 - M_{color} - M_{exact}) - A\bullet M_{fatigue}$

Cross-feature-based model:

$T = I_{shape}(1- M_{color} - M_{exact}) + I_{color}(1 - M_{shape} - M_{exact}) - A\bullet M_{fatigue}$

Null model

$T = I_{shape}(1 - M_{exact}) + I_{color}(1 - M_{exact}) - A\bullet M_{fatigue}$

where T represents the response to the test image, A represents the response to the antecedent adapter and the other terms were subject to the following constraints:

Shape and color inputs

$I_{shape} = S_1$ if test contains shape 1

S_2 if test contains shape 2

$I_{color} = C_1$ if test contains color 1

C_2 if test contains color 2

$0 \leq S_1, S_2, C_1, C_2$

Shape-match-based and color-match-based suppression:

$M_{shape} = W_{shape}$ if test shape matches adapter shape

0 if test shape does not match adapter shape

$M_{color} = W_{color}$ if test color matches adapter color

0 if test color does not match adapter color

$0 \leq W_{shape}, W_{color} \leq 1$

Exact-match-based suppression:

$M_{exact} = W_{exact}$ if test matches adapter in both shape and color

0 in any other case

$0 \leq W_{exact} \leq 1$

Fatigue-based suppression:

$M_{fatigue} = W_{fatigue}$

$0 \leq W_{fatigue} \leq 1$

We permitted W_{shape} and W_{color} to assume independent values because we could not rule out the possibility that shape-match and color-match would elicit different levels of suppression. We included an exact-match term because preliminary analysis suggested that an exact match elicited especially deep suppression. We included a fatigue term because preliminary analysis revealed an inverse relation between the strength of the response to the adapter and the strength of the response to the test image. The image-based, same-feature-based and cross-feature-based models, in their full form, possessed eight free parameters:

a. Strength of individual feature inputs (S_1, S_2, C_1, C_2).

b. Strength of suppression induced by shape-match and color-match (W_{shape}, W_{color})

c. Strength of suppression induced by exact match between adapter and test (W_{exact})

d. Strength of suppression induced by fatigue ($W_{fatigue}$)

The null model, in its full form, possessed only those six free parameters listed in categories a, c and d above. Each family of models comprised four nested variants: the full model, a version lacking exact-match-based suppression, a version lacking fatigue-based suppression and a version lacking both.

2.3.8 Model Fitting

Model fitting involved adjusting the values of the free parameters to minimize the residual sum of squared differences (RSS) between model output and observed firing rate across 16 conditions representing all possible combinations of test shape (2 levels), test color (2 levels), shape-match status (2 levels) and color-match status (2 levels). Fitting was implemented by use of the fmincon function in Matlab (Mathworks). In fitting any model either to population data or to individual-neuron data, we ran the function ten times with the free parameters assigned random initial values and selected the outcome associated with the lowest RSS. Increasing the number of iterations beyond ten yielded no improvement because, for any given model and dataset, repeated application of the procedure yielded only a few solutions at most.

2.3.9 Goodness of Fit (RSS)

Each family of models comprised four variants differentiated by the presence or absence of free parameters for exact-match-based suppression (M_{exact}) and fatigue-based suppression ($M_{fatigue}$). Within each family, we assessed, by means of an F-test (fcdf, Matlab, Mathworks), whether adding exact-match-based suppression and/or fatigue-based suppression to a model lacking it reduced RSS significantly more than would be expected from simply adding the extra degrees of freedom.

2.3.10 Efficiency of Fit (AICc)

We computed, for each model, a version of the Akaike Information Criterion (AICc) corrected for the case that the number of observations is fewer than forty times the number of free parameters:

$$AICc = n \bullet Ln(RSS/n) + 2 \bullet K \bullet (K+1)/(n-K-1)$$

where n is the number of observations (16 firing rates associated with 16 test conditions) and K is the number of free parameters (4-8 depending on the model). We judged the model for which AICc was lowest to afford the most efficient fit to the data.

2.3.11 Model Fitting Individual Cases

For each of 222 cases (111 neurons x 2 tetrads/neuron), we asked which model of the 16 yielded the most efficient fit to the data as indicated by lowest AICc. This analysis yielded 222 counts distributed across 16 models. The count for each model was the number of tetrads, out of the total of 222, for which that model provided the most efficient fit. When this approach was applied to random datasets, the counts, as expected, varied across models. To determine whether the pattern of counts obtained by applying the procedure to the actual data differed significantly from the pattern expected by chance, we carried out the procedure eight times on data randomly permuted so as to eliminate any systematic dependence of firing rate on the match status of the test image. Each of the 16 firing rates associated with the 16 test conditions had two labels, one marking the identity of the test image and one marking the match condition (shape-match, color-match, both-match or neither-match). On each of the eight iterations, for each of the 222 tetrads independently, we randomly shuffled the match-condition labels, leaving the image-identity labels intact. We then computed the average across the eight iterations, to the nearest integer, of each model's counts. Finally, for each model, we assessed the probability that the observed count was different from the mean shuffled count by use of Fisher's exact test (fishertest, Matlab, Mathworks) applied to the matrix [U, 222-U; S, 222-S] where S and U were counts obtained for shuffled and unshuffled data respectively.

2.3.12 Correlation across Trials between Adapter Response Strength and Test Response Strength

The aim of this analysis was to determine whether trial-to-trial fluctuations in the strength of the response to the test image were correlated to trial-to-trial fluctuations in the strength of the response to the adapter. The first step was to de-trend and normalize the data so as to remove the influence on firing rate of slow drift, image identity and match condition. For this purpose, we took advantage of the fact that each session consisted of four successive 32-trial blocks and that within each block there was one instance of each of the 32 possible conditions (2 shapes x 2 colors x 2 shape-match conditions x 2 color-match conditions for each of 2 tetrads). We carried out three steps designed to eliminate the influence of nuisance variables. Step 1, de-trending: Independently for each tetrad in each block, we normalized the 16 adapter-image firing rates by subtracting from each the mean of the 16. We did likewise for the 16 test-image firing rates. Step 2, removal of the influence of image identity: For each adapter image in each tetrad in each block, we normalized the four responses by subtracting from each the mean of the four. We did the same for each test image in each tetrad in each block. Step 3, removal of the influence of match status: For each match condition in each tetrad in each block, we normalized the four responses by subtracting from each the mean of the four. For each session, this procedure produced 128 de-trended and normalized adapter-image firing rates paired with 128 de-trended and normalized test-image firing rates, both still in units of spikes per second. In data combined across all 111 neurons, we proceeded to regress test-image response strength on adapter-image response strength.

2.3.13 Cluster-Based Permutation Test

To determine whether the post-image-onset time-varying population firing rates measured under two conditions were significantly different, we employed a nonparametric approach requiring neither arbitrary designation of a measurement window nor comparisons within multiple windows (*31*). The starting point for this analysis was a table of mean firing rates of each neuron under each condition in each 5 ms bin spanning a window from 0 ms to 560 ms following image onset. In each bin, we carried out a paired T-test on the two distributions of firing rates. If the test yielded a p-value < 0.05, we tagged the bin as positive or negative according to which condition was associated with the higher firing rate. If the test yielded a p-value ≥ 0.05, we tagged the bin as zero. The T-test was used only as a means for imposing an arbitrary threshold and not to assess statistical significance. For each cluster of bins of uniform sign, either positive or negative, we computed the sum across those bins of the associated T-statistics. The cluster with the greatest sum was classified as "best" for further steps of analysis. To test its statistical significance, we generated a permutation distribution, applying the above-described procedure 1,000 times to data in which the condition labels for each neuron had been randomly shuffled. We computed p as the fraction of iterations in which the best-cluster sum of T-statistics was greater than the best-cluster sum of T-statistics in the original data. We classified the original cluster as significant if $p < 0.05$. If this cluster was significant, we proceeded to assess the statistical significance of the observed cluster with the next highest sum of T-statistics. We computed p for this cluster as the fraction of cases in the previously generated permutation distribution for which the sum of T-statistics was greater than the observed value. We classified the cluster as significant if $p < 0.05$. We repeated this procedure until a non-significant result was obtained or no observed cluster larger than four bins remained.

2.4 RESULTS

We collected data from 111 visually responsive inferotemporal neurons (32 in M1 and 79 in M2) while the monkeys performed a task requiring them to maintain central fixation on each

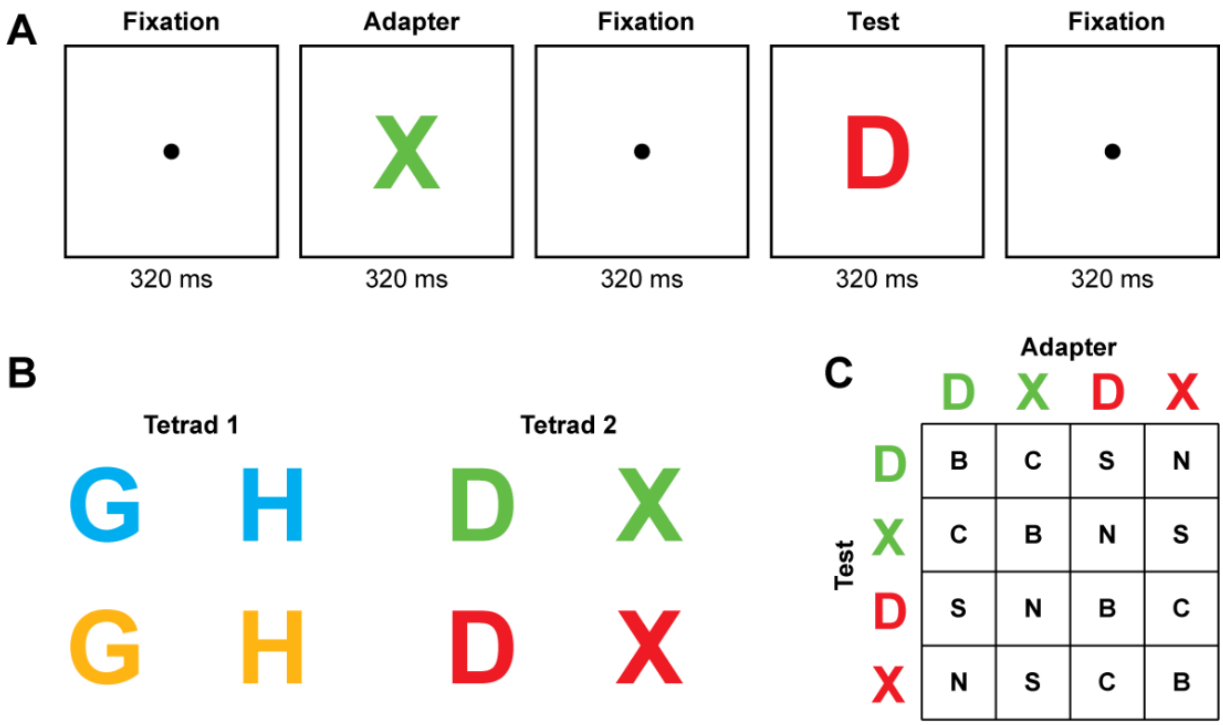

Figure 2.1. Experimental design. A. On each trial, a pair of colored shapes was presented in sequence. **B.** For testing each neuron, we employed two tetrads of images, each produced by crossing two shapes with two colors. **C.** On interleaved trials, images from a given tetrad were presented in all sixteen possible sequences. The test image might match the adapter in shape ("S"), color ("C"), neither ("N") or both ("B").

trial during successive central presentation of an adapter image and a test image (**Figure 2.1A**). For use with each neuron, we constructed two tetrads of images, with images in each tetrad representing all possible combinations of two shapes and two colors unique to that tetrad (**Figure 2.1B**). To minimize cross-trial adaptation, images from tetrad 1 and tetrad 2 were presented on alternate trials. Trials involving each tetrad conformed to 16 conditions representing all possible sequences of adapter and test. The test might match the adapter in shape, color, both attributes or

neither (**Figure 2.1C**). In describing the results, we will focus on data combined across the two monkeys. Key phenomena present in the combined data were evident in each monkey considered individually (**Figures S2.1-S2.12**).

2.4.1 Shape and color selectivity

To quantify the impact of adapter shape and color on firing rate following stimulus onset, we proceeded in the following steps. For each tetrad studied in each neuron, we identified the shape (and likewise color) eliciting stronger firing 75-375 ms following adapter onset. We then computed, for each neuron, the mean, across the two tetrads used in testing it, of responses measured under all four possible shape-color combinations (best-best, best-worst, worst-best and worst-worst). We then plotted the mean across neurons of the four firing rates as a function of time following adapter onset (**Figure 2.2A**). To analyze signal timing more finely, we computed, as a function of time following adapter onset, a shape index, a color index and an interaction index. The shape (or color) index was the difference in mean firing rate between trials in which the adapter had the best minus worst shape (or color) for a given neuron. The interaction index was the mean firing rate on trials when shape and color were both best or both worst minus the mean firing rate on trials when one was best and the other worst. This definition of the interaction index ensured that it would be positive for any interaction favoring conjunction selectivity and, in the event of perfect conjunction selectivity (manifest in a strong response to the image containing best shape and best color and weaker identical response to the other three images), would be equal in magnitude to the shape and color indices. During an analysis window 75-375 ms following stimulus onset, the shape, color and interaction indices had mean values, respectively, of 8.0, 11.9 and 1.0 spikes/second. The main effects of shape and color peaked rapidly following stimulus onset with a roughly equivalent time-course, whereas the interaction

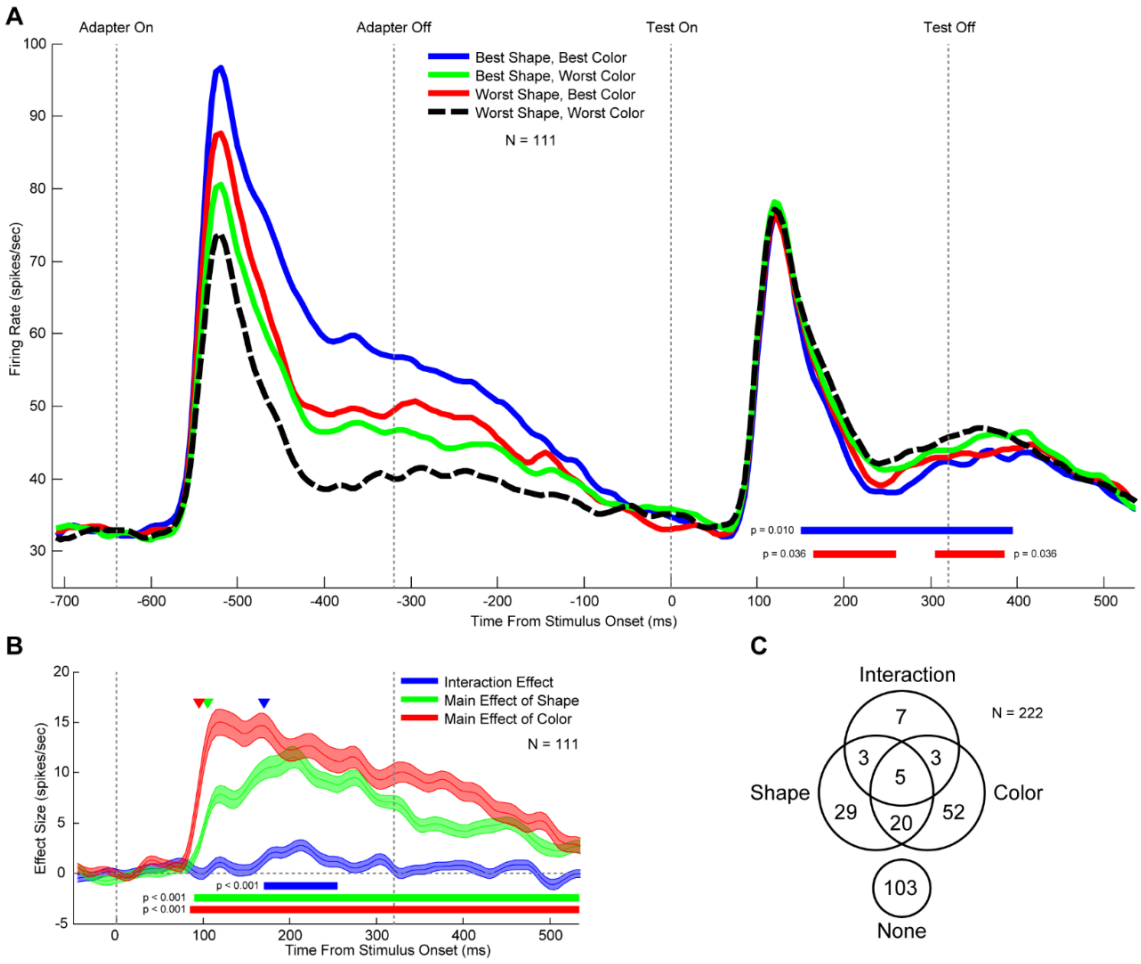

Figure 2.2. Responses to the adapter were selective for shape and color and moderately selective for their conjunction. A. Mean population firing rate as a function of time during trials sorted according to whether the adapter contained the best or worst shape and the best or worst color for the neuron in question (means in 5 ms bins smoothed with a 10 ms standard deviation Gaussian kernel). Although the main purpose of this plot is to demonstrate neuronal selectivity for the shape and color of the adapter, we also note that the strength of the response to the test image was inversely related to the strength of the response to the adapter. The underlying red and blue horizontal bars indicate periods during which the red and blue curves deviated significantly from the black dashed curve. The p-value juxtaposed to each bar indicates the statistical significance of the corresponding cluster. We consider possible explanations for this effect in the text. **B.** Strength of the shape-selective signal (green), color-selective signal (red) and interaction signal (blue) as a function of time following adapter onset. Ribbons represent ± standard error of the mean. Each colored triangle indicates the time to half-peak of the correspondingly colored curve. Horizontal bars indicate periods during which the correspondingly colored curves deviated significantly from zero as indicated by a cluster-based permutation test. The p-value juxtaposed to each bar indicates the statistical significance of the corresponding cluster. Time-to-half-peak markers are at 101 ms (shape), 90 ms (color) and 166 ms (interaction). **C.** Counts of cases (out of a total of 222 tetrads tested in 111 neurons) in which neuronal firing rate 75-375 ms following the onset of the adapter was significantly ($p < 0.05$) dependent on its shape, its color or their interaction.

effect, although positive on average, as expected in the case of conjunction selectivity, was relatively weak and poorly defined with respect to onset time (**Figure 2.2B**). All three effects achieved significance as indicated by a cluster-based permutation test (underlying bars in **Figure 2.2B**). To determine how often significant effects were present at the level of individual neurons, we applied to data collected using each tetrad in each neuron an ANOVA with shape (two levels) and color (two levels) as factors and with firing rate 75-375 ms following adapter onset as the dependent variable. This revealed a statistically significant ($p < 0.05$) main effect of shape, main effect of color and interaction effect in 57, 80 and 18 cases respectively (**Figure 2.2C**). The frequency of each effect was significantly greater than expected by chance (shape, color and interaction: $p < 0.00001$, $p < 0.00001$, and $p = 0.049$, chi-squared test, 1 df, with Yates correction). Among interaction effects, positive cases, in which the pattern of interaction favored conjunction selectivity, were significantly preponderant (14/18 cases, $p = 0.034$, chi-squared test, $df = 1$, with Yates correction).

To characterize shape and color selectivity of responses to the test images, we carried out analyses parallel to those applied to adapter responses. For each tetrad-neuron case, we sorted trials according to whether that image contained the best or worst shape and the best or worst color as identified by analysis of responses to the adapter. During an analysis window 75-375 ms following stimulus onset, the shape, color and interaction indices had mean values, respectively, of 5.5, 9.3 and 1.2 spikes/second. The dynamics of the population firing rate (**Figure 2.3A**) and of shape, color and interaction signals (**Figure 2.3C**) generally resembled the dynamics observed during the adapter response. The main effects of shape and color achieved significance as indicated by a cluster-based permutation test (underlying bars in **Figure 2.3C**). Before carrying out an ANOVA on data from each tetrad-neuron case, we factored out the influence of match-

status (shape alone, color alone, both or neither) by subtracting from each firing rate measured under a given match condition the mean of all firing rates measured under that condition. Across 222 neuron-tetrad cases, the ANOVA revealed a statistically significant ($p < 0.05$) main effect of shape, main effect of color and interaction effect in 68, 95 and 21 cases respectively (**Figure 2.3B**). The frequency of each of the three effects was significantly greater than expected by chance (shape, color and interaction: $p < 0.00001$, $p < 0.00001$, and $p = 0.0038$, chi-squared test, 1 df, with Yates correction). Among cases with a significant main effect of shape or color, the shape or color favored in the test image was nearly always the same as the shape or color favored in the adapter. Contrary outcomes occurred at a rate (6/68 cases for shape and 1/95 cases for color) within the expected false-positive range. We characterized an interaction effect as positive if the mean firing rate on trials when shape and color were both best or both worst was greater than the mean firing rate on trials when one was best and the other worst. Best and worst in this context were identified on the basis of the neuron's responses to the test images. Among interaction effects, positive cases, in which the pattern of interaction favored conjunction selectivity, were preponderant but the effect did not attain statistical significance (15/21 cases, $p = 0.081$, chi-squared test, $df = 1$, with Yates correction). We conclude, from the analysis of both adapter and test responses, that neurons were sensitive to both the color and shape of the adapter and exhibited a moderate tendency toward responding especially strongly to the image combining the favored shape with the favored color.

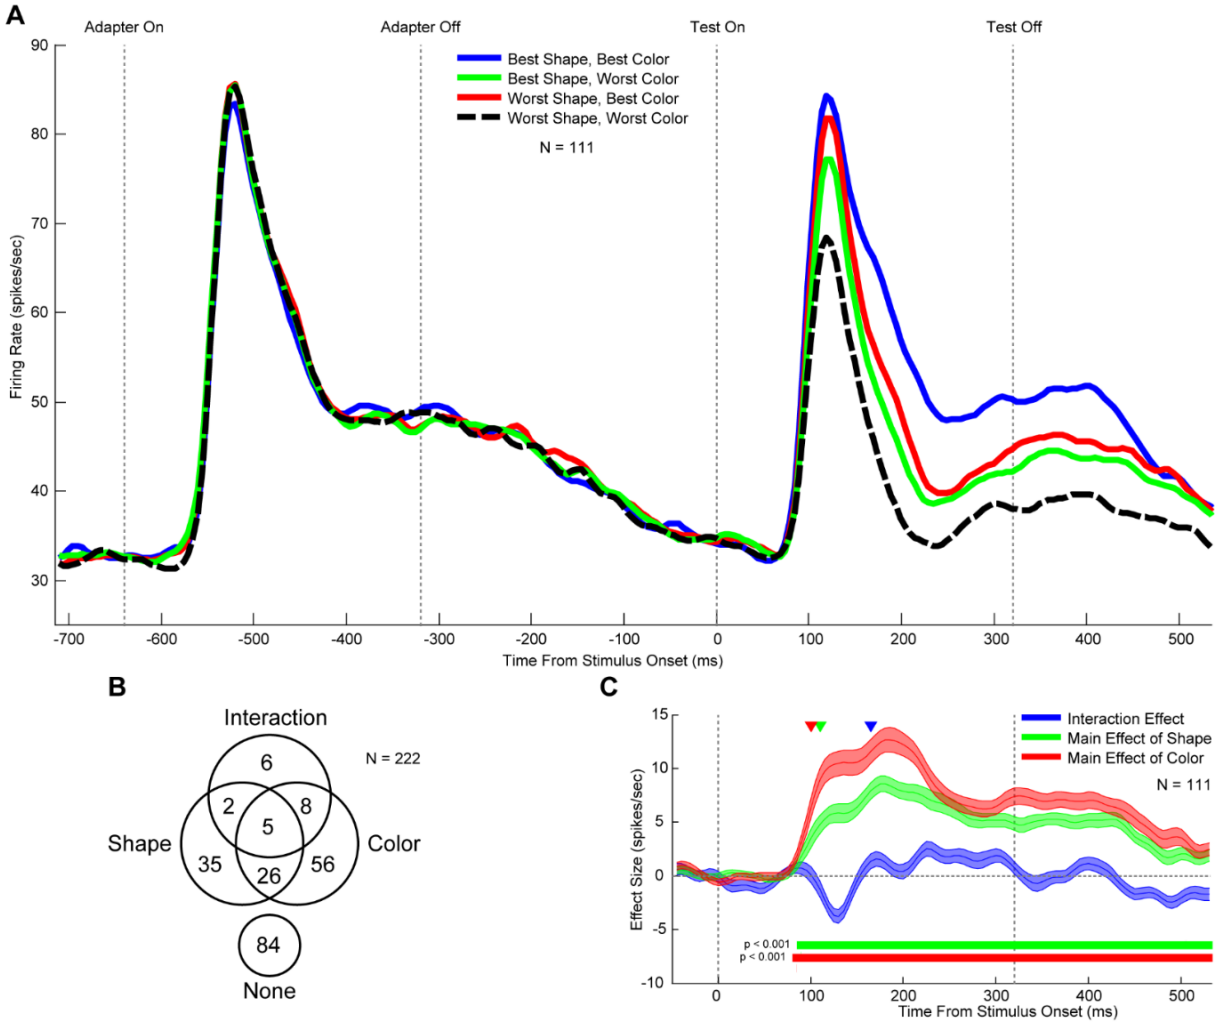

Figure 2.3. The pattern of selectivity for the test image paralleled the pattern of selectivity for the adapter. Conventions as in **Figure 2.2**. Time-to-half-peak markers are at 102 ms (shape), 96 ms (color) and 161 ms (interaction).

2.4.2 Repetition suppression

To characterize shape-dependent and color-dependent repetition suppression, we analyzed the population response as a function of whether the test image matched the adapter in shape, in color, in neither attribute or in both. At the level of the mean population firing rate, there was clear evidence for suppression under conditions of both shape-match and color-match, with the

effect of color-match moderately stronger than that of shape-match (**Figure 2.4A**). We quantified the impact of shape-match, color-match and their interaction as a function of time following stimulus onset by computing three time-varying suppression indices. The shape-match (or color-match) index was the mean firing rate on trials in which the test differed from the adapter in shape (or color) minus the mean firing rate on trials in which it had the same shape (or color). A positive value indicated repetition suppression. The interaction index was the mean firing rate on trials when the test image matched the adapter in one attribute or the other minus the mean firing rate when it matched in both or neither. This definition of the interaction index ensured that it would be positive for any interaction favoring conjunction selectivity and, in the event of perfect conjunction selectivity (manifest in deep suppression for the test image perfectly matching the adapter and weaker identical suppression for the other three cases), would be equal in magnitude to the shape and color indices. During an analysis window extending from 75-375 ms following stimulus onset, the shape-match, color-match and interaction indices had mean values, respectively, of 2.3, 3.4 and 0.88 spikes/second. The fact that the mean interaction index was positive indicates that suppression tended to be especially deep when the test image matched the adapter in both shape and color. All three effects achieved statistical significance as revealed by cluster-based permutation tests (underlying bars in **Figure 2.4C**). Overall, shape-match-dependent and color-match-dependent suppression (**Figure 2.4C**) developed more slowly and peaked later than shape selectivity and color selectivity (**Figures 2.2C** and **2.3C**).

We carried out a parallel analysis at the level of individual neurons, conducting, independently for each of two tetrads studied in each neuron, an ANOVA with firing rate as dependent variable and with shape-match status and color-match status as factors. Before carrying out the ANOVA, we factored out the influence of image identity by subtracting from

the firing rate on each trial the mean of all firing rates on all trials in which that trial's image was presented. Across 222 neuron-tetrad cases, the ANOVA revealed a statistically significant ($p < 0.05$) main effect of shape, main effect of color and interaction effect in 20, 40 and 12 cases respectively (**Figure 2.4B**). The frequency of the main effects was significantly greater than expected by chance whereas the frequency of interaction effects was not (shape: $p = 0.0097$, color: $p < 0.00001$, interaction: $p = 0.90$, chi-squared test, $df = 1$, with Yates correction). Main effects of shape-match and color-match were preponderantly positive (indicating repetition suppression), with the rate of positive cases significantly greater than expected by chance (shape-match: $n = 20/20$, $p = 0.000022$; color-match: $n = 35/40$, $p = 4.5E\text{-}6$; chi-squared test, $df = 1$, with Yates correction). Among interaction effects, positive cases (indicating conjunction selectivity) were preponderant although the effect did not achieve statistical significance ($n = 8/12$, $p = 0.39$, chi-squared test, $df = 1$, with Yates correction). Results obtained at the single-neuron level are in line with population-based analyses indicating that repetition suppression was sensitive to both color-match and shape-match, with especially deep suppression occurring when the test matched the adapter in both attributes.

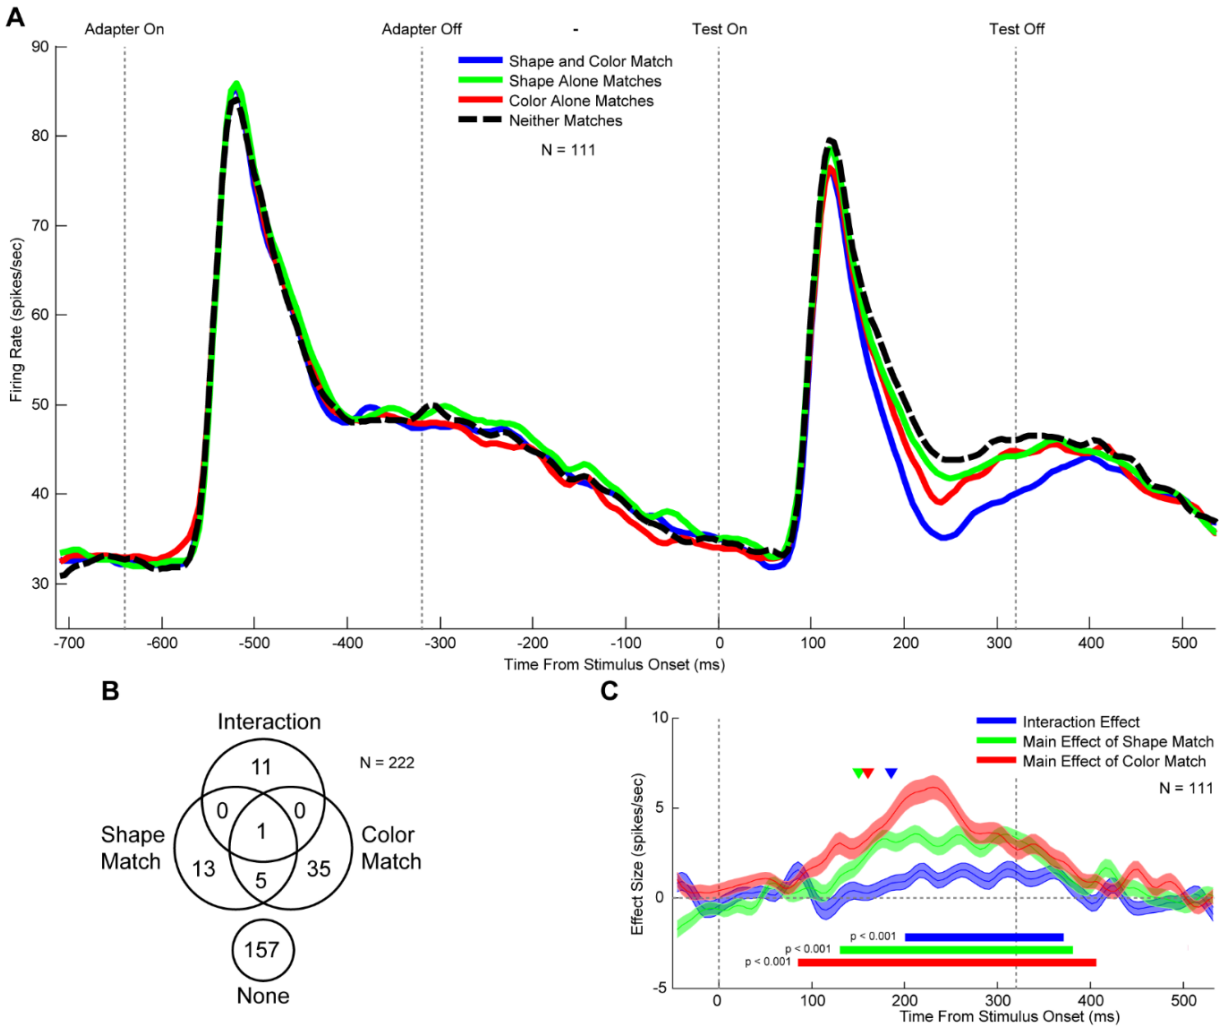

Figure 2.4. Suppression of the response to the test image depended on whether it resembled the adapter in shape and in color. A. For trials under each of four conditions (shape-match alone, color-match alone, both matching and neither matching), we computed the mean across all 111 neurons of firing rate as a function of time following test onset. The response to the test image was suppressed under all three match conditions (colored curves) as compared to the no-match condition (black dashed curve). **B.** Counts of cases (out of a total of 222 tetrads tested in 111 neurons) in which neuronal firing rate 75-375 ms following the onset of the test was significantly ($p < 0.05$) dependent on shape-match, color-match or their interaction. **C.** Strength of shape-match suppression (green), color-match suppression (red) and their interaction (blue) as a function of time following test onset. Time-to-half-peak markers are at 146 ms (shape-match), 154 ms (color-match) and 181 ms (interaction). Other conventions as in **Figure 2.2**.

2.4.3 Dependence of repetition suppression on neuronal shape and color preferences

Analyses up to this point have been based on data collapsed across neurons without regard to their individual shape and color preferences. This is a standard approach in studies of

repetition suppression. It would be justified if the depth of suppression depended, for all neurons, solely on the degree of similarity between the adapter and the test image. However, it is possible that suppression induced by a shape-match or a color-match depended on whether the shape or the color was preferred by the neuron. To investigate this possibility, we computed the population mean firing rate independently for all 16 conditions determined by test shape (best or worst for a given neuron), test color (best or worst for a given neuron), shape-match status (repeated or not) and color-match status (repeated or not). We found that shape-match-induced suppression predominated when the test image contained the best shape and worst color and that color-match-induced suppression predominated when the test image contained the best color and worst shape (**Figure 2.5A**). This pattern is incompatible with a mechanism in which the depth of suppression depends solely on the degree of similarity between the test-image and the adapter

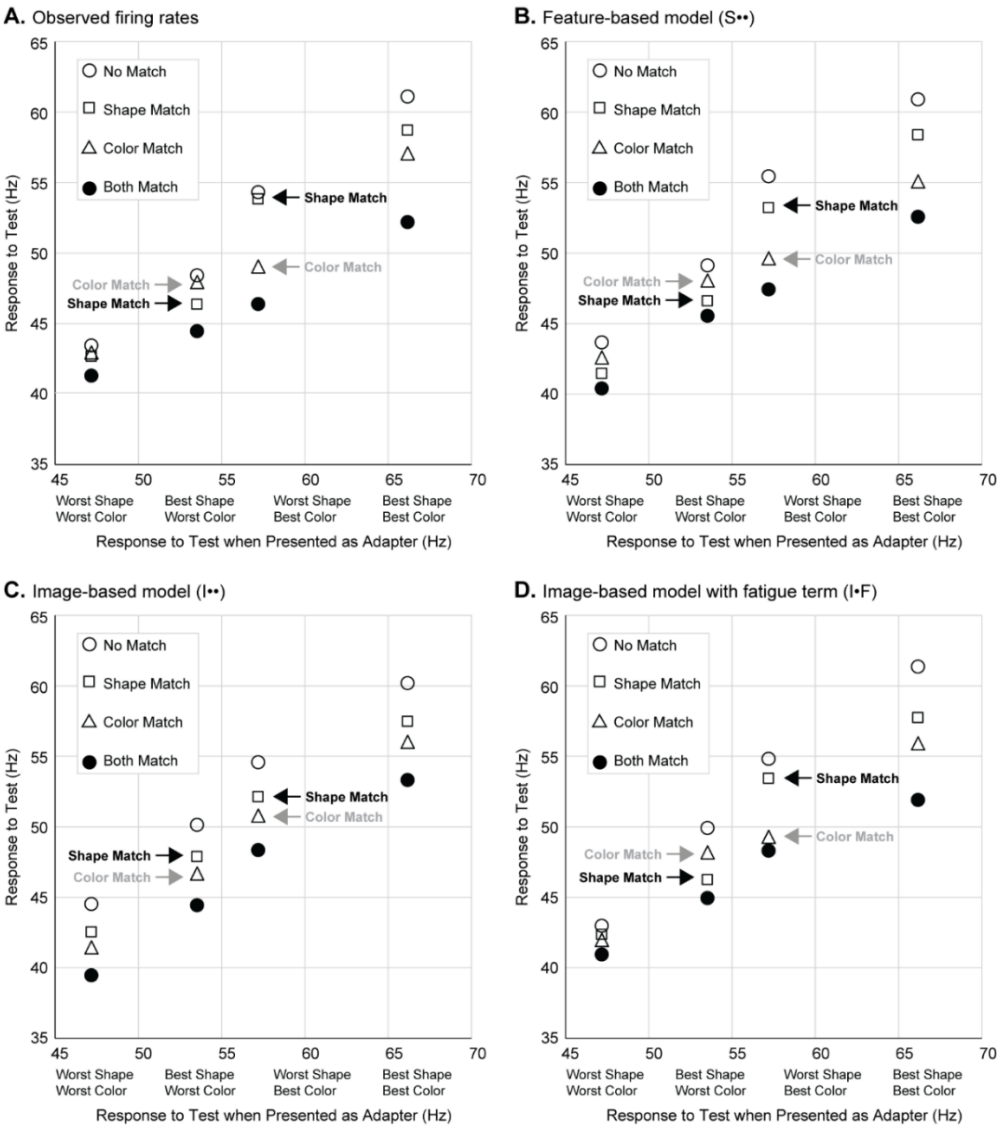

Figure 2.5. Dependence of repetition suppression on neuronal shape and color preference is captured by models incorporating feature-based suppression alone or image-based suppression in combination with fatigue. A. Images were sorted, for each neuron, according to whether they contained the preferred (best) or nonpreferred (worst) shape or color. The mean strength of the population response to the test image is plotted as a function of the mean strength of the population response elicited by the same image when presented as adapter. When the test image contained the best shape and worst color, shape-match suppression predominated. When it contained the best color and worst shape, color-match suppression predominated. **B.** Firing rates produced by best-fit model incorporating feature-based suppression. **C.** Firing rates produced by best-fit model incorporating image-based suppression. **D.** Firing rates produced by best-fit model incorporating both image-based suppression and fatigue.

To explore the possible origins of this phenomenon, we fitted multiple models to the experimental data. These embodied several mechanisms by which stimulus identity and repetition status might determine firing rate. In image-based suppression, repetition of shape or color equally suppressed shape and color input. In same-feature-based suppression, repetition of shape suppressed shape input and likewise for color. In cross-feature-based suppression, repetition of shape suppressed color input and vice versa. In conjunction-based suppression, there was a fractional reduction of input strength unique to the case in which both shape and color were repeated. In fatigue-based suppression, there was a decrement in response strength proportional to the mean firing rate elicited by the antecedent adapter. See Methods for further details. Note that, in the context of the modeling analysis, we use the term "same-feature-based suppression" to denote the phenomenon referred to outside this context simply as "feature-based suppression". We will refer to models by means of acronyms obeying the following conventions: image-based = I, same-feature-based = S, cross-feature-based = C, conjunction-based = X, fatigue-based = F and absent = •. For example, SX• denotes the model embodying same-feature-based and conjunction-based suppression but not fatigue-based suppression. The models formed four nested families comprising four members each. The models in different families obeyed different principles with regard to how repetition of shape or color affected the strength of shape or color input. The principles were, respectively, image-based, same-feature-based, cross-feature-based and none of the above.

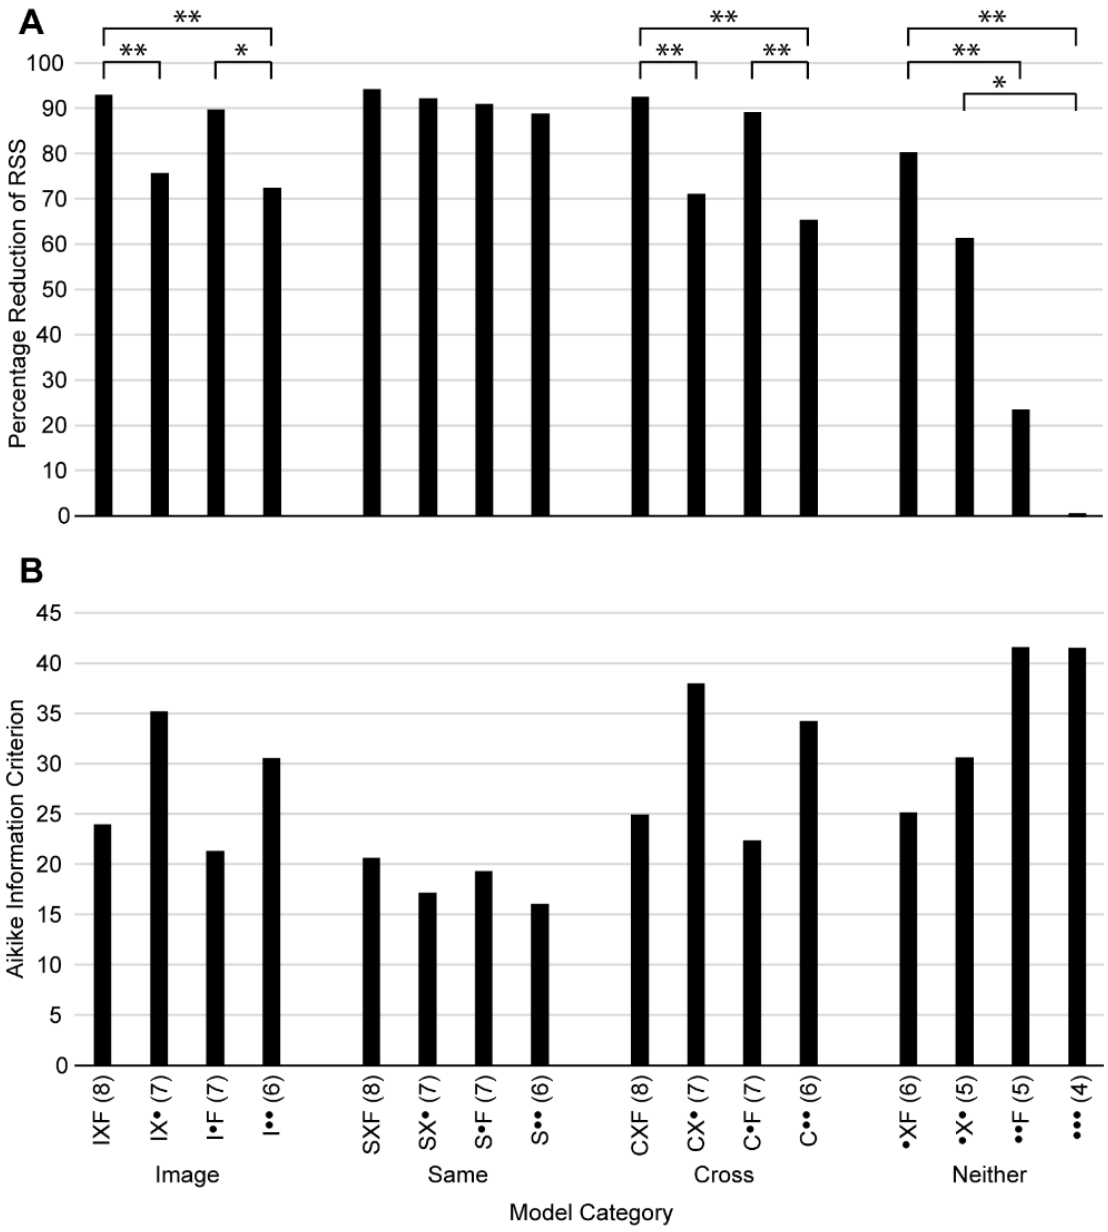

Figure 2.6. Performance measures for 16 models parametrically adjusted to provide the best fit to population data shown in Figure 2.5. A. Percentage reduction of residual sum of squared differences afforded by each model. The height of each bar represents 100*(B-M)/B where B = 103.4 which was the residual sum of squared differences for the model that contained free parameters only for shape and color drives (•••) and M was the residual sum of squared differences for the model in question. Asterisks indicate cases in which an F-test applied to two models with different degrees of freedom revealed a greater improvement in the full model than could be explained by additional degrees of freedom alone (* = $p < 0.05$; ** = $p < 0.01$). The F-test was based on raw RSS and not on the percentage reduction of RSS depicted in the plot. **B.** The Akaike Information Criterion for each model. The parenthetic number appended to each model's acronym indicates its number of degrees of freedom.

Upon fitting the 16 models to the population mean firing rates shown in **Figure 2.5A**, we obtained results summarized in **Figure 2.6**. Model SXF yielded the best fit to the mean population firing rates, as reflected in the greatest reduction of the residual sum of squared differences (**Figure 2.6A**). Model S•• yielded the most efficient fit as reflected in the lowest Akaike Information Criterion (**Figure 2.6B**). On these grounds, we might conclude that suppression arises from a feature-based rather than an image-based mechanism. However, it should be noted that a model combining image-based suppression with fatigue-based suppression (I•F) performed as well as the feature-based model (S••). The difference between feature-based and image-based suppression thus concerns their ability, when operating in isolation, to reproduce the experimentally observed pattern of results. The feature-based model performed well in isolation (**Figure 2.5B**). The image-based model performed poorly in isolation (**Figure 2.5C**) and showed a marked improvement in performance due to the addition of a fatigue-based term (**Figure 2.5D**). The improvement was significantly greater than expected simply from the presence of an additional free parameter (as indicated by asterisks in **Figure 2.6A**). The interaction term (X) in contrast did not produce a significant improvement when added either to the image-based or the feature-based model (**Figure 2.6A**). This observation, because it is based solely on ranking models with regard to their ability to fit the population mean firing rates, does not contradict the preceding conclusion, based on an analysis of variance, that repetition suppression exhibited a statistically significant conjunctive interaction effect.

We proceeded to apply the same approach to data from individual neuron-tetrad cases. For each of 222 cases obtained by testing 111 neurons with two tetrads each, we determined which model afforded the best fit to the data, as reflected in the lowest RSS, and which model afforded the most efficient fit, as reflected in the lowest AICc. We then counted, for each model, the

number of cases in which it prevailed. Data obtained from a single neuron with only four trials per condition are extremely noisy. To control for differences among models with regard to their ability to fit noise, we did not compare the counts for different models to each other but, instead, compared the count obtained for a given model upon processing the raw data (signal plus noise) to the average count obtained for the same model when processing data after random shuffling of the match-status labels (noise alone). See Methods for details. This approach revealed three significant effects. Cases in which the CXF model produced the best fit, as reflected in the lowest RSS, were significantly less numerous in the raw than in the shuffled data ($78 < 101$; $p = 0.033$; Fisher's exact test). Cases in which the S•• model provided the most efficient fit, as reflected in lowest AICc, were significantly more numerous in the raw than in the shuffled data ($24 > 3$; $p = 3.0$ E-5). Cases in which the ••• model provided the most efficient fit, as reflected in the lowest AICc, were significantly less numerous in the raw than in the shuffled data ($138 < 181$; $p = 8.1$ E-6). These effects, although small, as expected due to limits imposed by noise, nevertheless provide statistical support for the population-level observations. The underperformance of the CXF model argues against the occurrence of cross-feature suppression (and thus against an image-based mechanism); the overperformance of the S•• model argues for the occurrence of same-feature suppression (and thus in favor of a feature-based mechanism); and the underperformance of the ••• model indicates that free parameters absent from this model and present in the other models captured significant variance.

2.4.4 Arbitrating between feature-based suppression and image-based suppression with fatigue

Results presented up to this point are compatible with two interpretations of repetition suppression based respectively on feature-based suppression and image-based suppression

conjoined to fatigue. It might be thought that the existence of an inverse correlation between the strength of the response to the adapter and the strength of the response to the test-image (**Figure 2.2A**) favors an interpretation based on fatigue. However, feature-based suppression might also produce such a pattern under the conditions of our experiment. To determine whether this was so, we measured, for the best-fit S•• model, the correlation, across all 16 experimental conditions, between adapter response strength and test-image response strength. The strength of the response to the test image was indeed inversely correlated with the strength of the response to the adapter. The slope was -0.15 and the Pearson's correlation coefficient was -0.17. These values are nearly identical to the values yielded by the same analysis applied to the raw data: a slope of -0.18 and a correlation coefficient of -0.21. The best-fit S•• model, although it lacks any term for fatigue, thus can account for the fatigue-like pattern of results.

By use of data internal to this study, the only additional way in which to arbitrate between accounts based on feature-based suppression without fatigue and image-based suppression with fatigue was to ascertain whether, with all experimenter-controlled factors held constant, random trial-to-trial fluctuations in the strength of the test response were negatively correlated with random trial-to-trial fluctuations in the strength of the adapter response, as expected from a fatigue-based mechanism. To carry out this analysis required first factoring out the effects of trial condition and of slow drift across the session (see Methods for details). Following this step, we observed not a negative but a positive correlation ($p = 5.6$ E-25). The value of the measured Pearson's correlation coefficient ($r = 0.086$) is commensurate with the values of pairwise spike-count correlations thought to arise from slow uncontrolled fluctuations in hidden state variables (*32*). Any fatigue-based effect, if present, was so small as to be washed out by this positive correlation.

2.4.5 Impact of repetition on the dynamics of the visual response

It is evident in population histograms that the response to the test image differed from the response to the adapter not only by virtue of reduced amplitude but also because of an alteration in dynamic whereby the firing rate fell off and then rebounded over the course of a few hundred milliseconds (**Figure 2.4A**). It appears in the same figure that this effect was greatest when suppression was deepest owing to the repetition of both shape and color. To determine how consistent this effect was, we computed, for each neuron, independently for cases in which the image was an adapter, a non-matching test and a fully matching test, an index designed to assume a positive value in the event of rebound. This was the firing rate during a 50 ms window centered 400 ms following stimulus onset minus the firing rate in a 50 ms window centered 250 ms following stimulus onset. The rebound index for a test image matching the adapter in both shape and color (mean = 6.7 spikes/sec) was significantly greater than the rebound index for the adapter (mean = -2.1 spikes/sec, $p = 2.0 \text{ E-}7$, two-sided paired-sample rank sum test) as shown in **Figure 2.7A**. It was also significantly greater than the rebound index for the test image matching the adapter in neither shape nor color (mean = -0.71 spikes/sec, $p = 1.3 \text{ E-}5$) as show in **Figure 2.7B**. We conclude that image-specific repetition suppression was consistently accompanied by a shift in the temporal pattern of the response toward a trough-rebound dynamic. This result, taken in conjunction with the fact that match-status-specific suppression (**Figure 2.4B**) peaked at the time of the trough, later than shape and color signals peaked (**Figure 2.3B**), implies a close relation between the trough-rebound dynamic and repetition suppression.

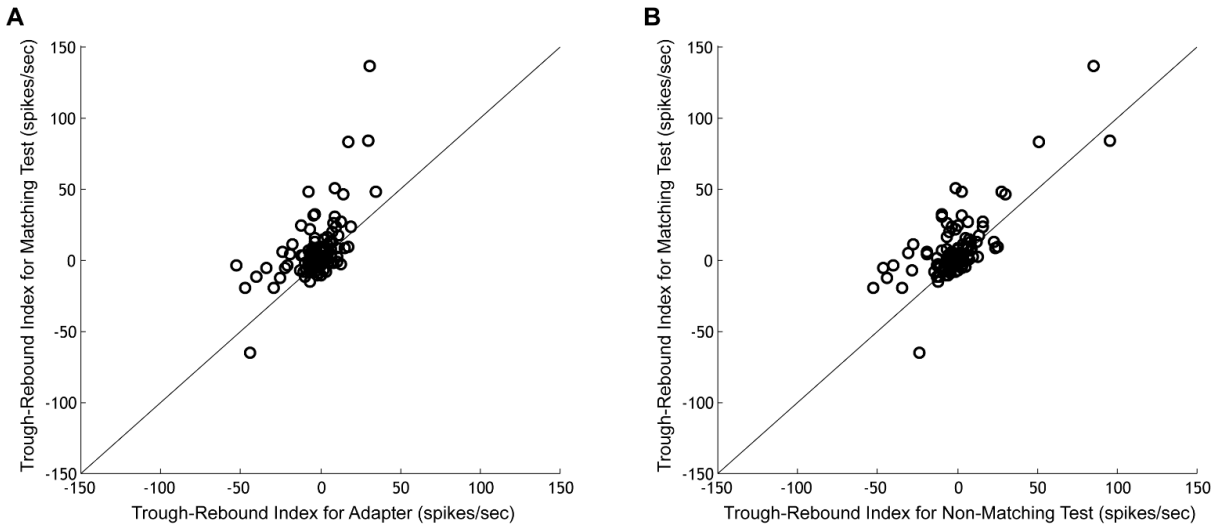

Figure 2.7. The trough-rebound response dynamic was most pronounced for a test image matching the adapter in both shape and color. A. The trough-rebound index for a fully matching test image is plotted against the trough-rebound index for the adapter. **B.** The trough-rebound index for a fully matching test image is plotted against the trough-rebound index for a test image matching the adapter in neither shape nor color.

2.4.6 Consistency of repetition suppression across neurons

Having assessed each neuron's responses to eight test images allowed us to ask whether some neurons exhibited repetition suppression consistently across multiple images while others did not. To answer this question, we counted, for each neuron, the number of images for which the response to the test was lower when it perfectly matched the adapter than when it differed in both shape and color. We used the sign of the difference as a criterion rather than its significance because the number of trials for a single image was too low to allow establishing statistical significance. We compared the observed distribution of counts to the distribution expected by chance if the probability of a negative value for any particular neuron and any particular test image depended only on the overall frequency of negative values across all neurons and all test images (544/888). The two distributions appeared closely similar (**Figure 2.8A**) and were not statistically distinguishable ($p = 0.39$, chi-squared test, $df = 8$, $n = 111$).

Although the tendency to exhibit repetition suppression might be consistent across neurons, the strength of the effect might differ. To investigate this possibility, we computed, for each tetrad studied in each neuron, the mean difference in firing rate between conditions in which the test perfectly matched the adapter and those in which it differed in both shape and color. The measures of suppression obtained for tetrad 1 and 2 were highly correlated across neurons ($r = 0.44$, $p < 0.00001$, $n = 111$). However, the correlation could have arisen from cross-

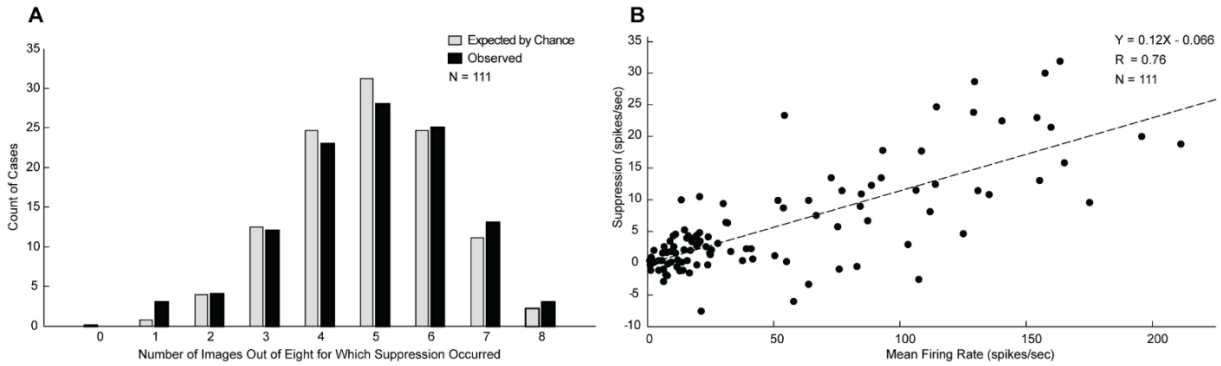

Figure 2.8. TE neurons are homogeneous with respect to repetition suppression. A. For each neuron, for each of eight test images, we determined whether the mean firing rate 75-375 ms following test-image onset was lower when the test image matched the adapter in shape and color than when it matched in neither attribute. If it was lower, we considered it to be a case of suppression. For each neuron, the number of cases of suppression could range from a minimum of zero to a maximum of eight. The height of each black bar indicates the number of neurons with the corresponding count. We compared this distribution to the distribution of counts expected by chance on the principle that the probability of suppression for any particular neuron and any particular test image depended only on the overall frequency of cases of suppression across all neurons and all test images (gray bars). The black and gray distributions are not significantly different. **B.** Collapsing data across all images studied in each neuron, we computed two measures: the mean firing rate and the mean difference in firing rate between trials when the test image differed from the adapter in both shape and color and trials when it was the same with regard to both attributes. The two measures were positively and significantly correlated, in accordance with the idea that suppression removes a fixed fraction of the response.

neuronal differences in mean firing rate if the magnitude of suppression were proportional to firing rate. To investigate this possibility, we analyzed the dependence of suppression strength on firing rate (**Figure 2.8B**). The two factors were significantly positively correlated ($r = 0.76$, $p <$

0.00001, n = 111). Suppression scaled as a function of firing rate with a proportionality constant of 0.12. When we divided the mean depth of suppression for each neuron by the mean firing rate for that neuron and recomputed the correlation across neurons between tetrad-1-based and tetrad-2-based measures of suppression, the correlation became vanishingly small and insignificant (r = 0.079, p = 0.41, n = 111). We conclude that neurons do not obviously differ in their tendency to exhibit repetition suppression.

It has been noted that the existence of two subpopulations of neurons, one exhibiting repetition suppression and the other not, would allow downstream areas to distinguish a repeated stimulus (eliciting a weak response only in the suppressing subpopulation) from a stimulus ineffective due to some visual feature such as low contrast (eliciting weak responses in both subpopulations) (*5, 33*). The results just described imply that such an approach is unlikely to work. To be sure of this conclusion, we carried out an additional test. For each neuron, we identified the "worst" image - the one combining the least effective color and shape - and the "best" image - the one combining the most effective color and shape. These served as proxies for "ineffective" and "effective" images. We then computed, for each neuron, a suppression index for tetrad 1. This was the mean "both-match" firing rate minus the mean "no-match" firing rate divided by the sum of the means. We ranked neurons according to this index. We then made a cut separating a group with low measured suppression (mean suppression index = 0.00, n = 80) from the remaining group with high measured suppression (mean suppression index = -0.41, n = 31). When we represented images as points in an activation space with principal axes representing the firing rates of the putative low-suppression and high-suppression subpopulations, it was trivially possible, for tetrad 1, to draw a decision boundary between points representing trials with and without repetition, thus avoiding the confound between the

effectiveness of an image and its repetition status (**Figure 2.9A**). The non-trivial question was whether, using the responses of subpopulations defined on the basis of tetrad 1 data, we could apply the same approach successfully to data from trials involving tetrad 2. We could not (**Figure 2.9B**). We obtained an equivalent result from a symmetric procedure using the firing of neurons classified on the basis of tetrad 2 data to distinguish repeated from ineffective images on tetrad 1 trials. These observations undercut the notion that differences among neurons with

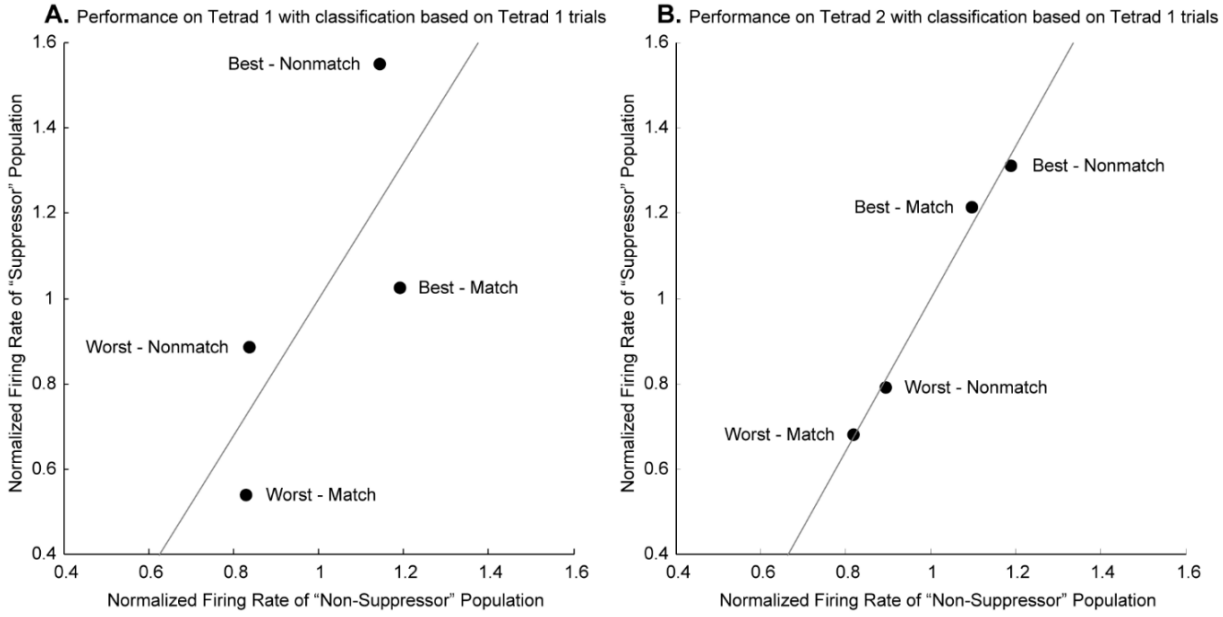

Figure 2.9. For a downstream area to distinguish a repeated image from an ineffective image on the basis of TE firing rate would be difficult because TE neurons do not form distinct low-suppression and high-suppression populations. A. When neurons were classified into putative low-suppression and high-suppression groups on the basis of responses to the test image during tetrad 1 trials, it was trivially possible, in data from tetrad 1 trials, to place a decision boundary in activation space that separated repeated (full match) from unrepeated (full nonmatch) conditions without regard to whether the test image was relatively ineffective (worst) or relatively effective (best). **B.** This approach failed when applied to data from tetrad 2 trials because subgroups yielding low suppression measures and high suppression measures for tetrad 1 did not exhibit low and high suppression for tetrad 2. For each tetrad studied in each neuron, firing rate was normalized by dividing the firing rate for each condition (best-match, best-nonmatch, worst-match and worst-nonmatch) by the mean of the four firing rates. For each of the two populations under each of the four conditions, the population mean firing rate was computed as the average across neurons of the normalized firing rates.

respect to the occurrence or strength of repetition suppression could be leveraged downstream from TE to distinguish between ineffective and repeated stimuli.

2.5 DISCUSSION

The first and most basic question posed by this study was whether repetition suppression at the level of population activity in TE would even occur when a test image shared only half its features with the adapter. We have answered this question in the affirmative. A test image matching the adapter only in shape or only in color elicits a response intermediate in strength between responses elicited by a fully matching and a fully nonmatching image. This outcome would be expected of a system in which shape-match-induced suppression and color-match-induced suppression sum linearly in modulating neuronal response strength. However, linear summation cannot entirely explain the results. When the test image matched the adapter in both shape and color, suppression was deeper than could be explained by linear summation. This suggests the operation of a conjunction-specific mechanism. Linear summation predominated over conjunction specificity, as indicated by the fact that the ratio of interaction-effect magnitude to the average of the shape main-effect and color main-effect magnitudes was 0.31, a value closer to 0.0, reflective of pure linear summation, than to 1.0, reflective of pure conjunction specificity. However, it is worth note that the tendency for suppression to be especially deep when both shape and color matched was markedly greater than the tendency for neurons to respond especially strongly to an image combining the preferred shape with the preferred color. The corresponding ratio for conjunction selectivity as reflected in visual response strength was 0.050 for the adapter and 0.068 for the test image.

Given that cross-suppression occurs even for images sharing only one feature, we were able to investigate whether the governing principle is image-based or feature-based. According to the

image-based account, repetition of an image as a whole fractionally reduces the strength of the response to that image when repeated. According to the feature-based account, repetition of a shape or a color fractionally reduces the strength of input to the neuron from that shape or that color when repeated. These accounts are dissociable in a neuron preferring one shape over the other or one color over the other because, in such a neuron, only the feature-based model predicts that repetition of a preferred shape or color will produce deeper suppression than repetition of a non-preferred shape or color. Our results conformed to the predictions of the feature-based model. We note, however, one potentially important qualification, namely that the same pattern of results could have arisen from image-based suppression accompanied by neuronal fatigue. By fatigue, we mean a reduction in the response to the test image proportional to the strength of the response to the antecedent adapter. There are reasonably compelling arguments against interpreting the present results in terms of fatigue. We found no evidence for adapter-induced fatigue in our analysis of trial-to-trial fluctuations of response strength. Moreover, intense optogenetic activation of a TE neuron does not significantly reduce the strength of its response to an immediately ensuing visual probe (*19*). It is true that an adapter to which a neuron is strongly responsive produces greater suppression than an adapter to which the neuron does not respond when the following test is a third image to which the neuron is strongly responsive (*28*). However, this phenomenon, although it could be interpreted in terms of fatigue, could also arise from feature-specific suppression if the two effective images shared a feature for which the neuron was selective.

We have shown that the neuronal response to a repeated image has a trough-rebound dynamic. This pattern is evident retrospectively in post-stimulus-time histograms from numerous previous studies involving presentation of adapter and test in rapid succession (*16, 17, 19, 22-25,*

27, 28) and may underlie an early report to the effect that the match status of an image can be decoded from the temporal profile of the neuronal response (*34*). A seemingly identical trough-rebound dynamic is evident in neuronal responses to images rendered familiar by hundreds or thousands of exposures (*35, 36*). This effect appears to arise from accentuation of a tendency, evident even in studies not involving repetition or familiarization, for inferotemporal visual responses to take the form of a damped 5 Hz oscillation (*37*). The phenomenon could arise from a simple mechanism involving resonance in an adapting recurrent inhibitory (*38*) or excitatory (*35*) circuit. The fact that repetition of an image and its prolonged familiarization both cause accentuation of the trough-rebound pattern suggests that they share a site of action and that this site is within the resonant circuit mediating 5 Hz damped oscillations.

TE neurons in this study were homogeneous with regard to the occurrence and strength of repetition suppression. This is of interest because it has a direct bearing on the question of whether repetition suppression could mediate recognition memory. A weak population response does not constitute an unambiguous signal that an image has been repeated because unrepeated low-contrast images also elicit a weak population response (*39*). Downstream areas might achieve disambiguation by comparing the activity of a suppressing population to the activity of a non-suppressing population: both would fire weakly in response to a low-contrast image but only the suppressing population would fire weakly in response to a repeated image (*5, 33*). However, the idea that TE neurons fall into such categories or even cover a range with regard to repetition suppression has never been put to a decisive test. When neurons are characterized individually with an index such that greater suppression is indicated by greater positive values, the distribution of indices invariably has a positive mean and a flank that extends into the negative domain without any obvious deviation from unimodality (*15-17, 21-24, 27*). This spread could

reflect genuine differences among the neurons or could be due just to measurement error. Statistical tests applied to data from individual neurons invariably reveal a mix of cases in which suppression achieves and does not achieve statistical significance. This could arise simply from lack of statistical power. A few cases have been noted in which repetition induces significantly enhanced firing (*1, 5, 6, 9, 15, 33, 40*). These cases must, however, be interpreted with caution for two reasons. First, they might represent rare false positives. Second, they may have been related to match detection in studies requiring monkeys to perform delayed-match-to-sample tasks. The present study is the first to have addressed this issue systematically by assessing the consistency with which each neuron exhibits repetition suppression across multiple image sets. Our findings rule out the idea that TE neurons differ with regard to repetition suppression and thereby call into question the idea that a weak population response in TE is a signal useful for recognition of a repeated image.

2.6 REFERENCES

1. G. C. Baylis, E. T. Rolls, Responses of neurons in the inferior temporal cortex in short term and serial recognition memory tasks. *Exp Brain Res* **65**, 614-622 (1987).
2. M. W. Brown, F. A. Wilson, I. P. Riches, Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res* **409**, 158-162 (1987).
3. K. S. Saleem, N. K. Logothetis, *Atlas of the Rhesus Monkey Brain in Stereotaxic Coordinates*. (Academic Press, Amsterdam, 2007).
4. F. L. Fahy, I. P. Riches, M. W. Brown, Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. *Exp Brain Res* **96**, 457-472 (1993).
5. L. Li, E. K. Miller, R. Desimone, The representation of stimulus familiarity in anterior inferior temporal cortex. *J Neurophysiol* **69**, 1918-1929 (1993).
6. A. Lueschow, E. K. Miller, R. Desimone, Inferior Temporal Mechanisms for Invariant Object Recognition. *Cerebral Cortex* **4**, 523-531 (1994).
7. E. Miller, L. L, R. Desimone, A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* **254**, 1377-1379 (1991).
8. E. Miller, L. Li, R. Desimone, Activity of neurons in anterior inferior temporal cortex during a short- term memory task. *The Journal of Neuroscience* **13**, 1460-1478 (1993).
9. E. K. Miller, R. Desimone, Parallel neuronal mechanisms for short-term memory. *Science* **263**, 520-522 (1994).

10. E. K. Miller, R. Desimone, Scopolamine affects short-term memory but not inferior temporal neurons. *Neuroreport* **4**, 81-84 (1993).
11. I. P. Riches, F. A. Wilson, M. W. Brown, The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. *J Neurosci* **11**, 1763-1779 (1991).
12. S. Sobotka, J. L. Ringo, Mnemonic responses of single units recorded from monkey inferotemporal cortex, accessed via transcommissural versus direct pathways: a dissociation between unit activity and behavior. *J Neurosci* **16**, 4222-4230 (1996).
13. B.-E. Verhoef, G. Kayaert, E. Franko, J. Vangeneugden, R. Vogels, Stimulus Similarity-Contingent Neural Adaptation Can Be Time and Cortical Area Dependent. *The Journal of Neuroscience* **28**, 10631-10640 (2008).
14. R. Vogels, G. Sary, G. A. Orban, How task-related are the responses of inferior temporal neurons? *Visual Neuroscience* **12**, 207-214 (1995).
15. T. Meyer, N. C. Rust, Single-exposure visual memory judgments are reflected in inferotemporal cortex. *Elife* **7**, (2018).
16. K. Vinken, H. P. Op de Beeck, R. Vogels, Face Repetition Probability Does Not Affect Repetition Suppression in Macaque Inferotemporal Cortex. *J Neurosci* **38**, 7492-7504 (2018).
17. W. De Baene, R. Vogels, Effects of Adaptation on the Stimulus Selectivity of Macaque Inferior Temporal Spiking Activity and Local Field Potentials. *Cerebral Cortex* **20**, 2145-2165 (2010).
18. D. B. T. McMahon, C. R. Olson, Repetition Suppression in Monkey Inferotemporal Cortex: Relation to Behavioral Priming. *Journal of Neurophysiology* **97**, 3532-3543 (2007).
19. F. Fabbrini *et al.*, Probing the Mechanisms of Repetition Suppression in Inferior Temporal Cortex with Optogenetics. *Current biology : CB* **29**, 1988-1998.e1984 (2019).
20. F. Fabbrini, R. Vogels, Within- and between-hemifield generalization of repetition suppression in inferior temporal cortex. *J Neurophysiol* **125**, 120-139 (2021).
21. D. A. Kaliukhovich, R. Vogels, Neurons in macaque inferior temporal cortex show no surprise response to deviants in visual oddball sequences. *J Neurosci* **34**, 12801-12815 (2014).
22. D. A. Kaliukhovich, R. Vogels, Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral cortex (New York, N.Y. : 1991)* **21**, 1547-1558 (2011).
23. D. A. Kaliukhovich, R. Vogels, Stimulus repetition affects both strength and synchrony of macaque inferior temporal cortical activity. *J Neurophysiol* **107**, 3509-3527 (2012).
24. D. A. Kaliukhovich, R. Vogels, Divisive Normalization Predicts Adaptation-Induced Response Changes in Macaque Inferior Temporal Cortex. *J Neurosci* **36**, 6116-6128 (2016).
25. P. Kuravi, R. Vogels, Effect of adapter duration on repetition suppression in inferior temporal cortex. *Sci Rep* **7**, 3162 (2017).
26. P. Kuravi, R. Vogels, GABAergic and cholinergic modulation of repetition suppression in inferior temporal cortex. *Sci Rep* **8**, 13160 (2018).
27. Y. Liu, S. O. Murray, B. Jagadeesh, Time Course and Stimulus Dependence of Repetition-Induced Response Suppression in Inferotemporal Cortex. *Journal of Neurophysiology* **101**, 418-436 (2009).
28. H. Sawamura, G. A. Orban, R. Vogels, Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. *Neuron* **49**, 307-318 (2006).
29. v. G. Bonin, P. Bailey, *The Neocortex of Macaca Mulatta*. (University of Illinois Press, Urbana, 1947).
30. W. A. Suzuki, Y. Naya, The Perirhinal Cortex. *Annual Review of Neuroscience* **37**, 39-53 (2014).

31. E. Maris, R. Oostenveld, Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* **164**, 177-190 (2007).
32. M. R. Cohen, A. Kohn, Measuring and interpreting neuronal correlations. *Nat Neurosci* **14**, 811-819 (2011).
33. E. K. Miller, L. Li, R. Desimone, Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* **13**, 1460-1478 (1993).
34. E. N. Eskandar, B. J. Richmond, L. M. Optican, Role of inferior temporal neurons in visual memory. I. Temporal encoding of information about visual images, recalled images, and behavioral context. *J Neurophysiol* **68**, 1277-1295 (1992).
35. S. Lim, Mechanisms underlying sharpening of visual response dynamics with familiarity. *Elife* **8**, (2019).
36. T. Meyer, C. Walker, R. Y. Cho, C. R. Olson, Image familiarization sharpens response dynamics of neurons in inferotemporal cortex. *Nat Neurosci* **17**, 1388-1394 (2014).
37. J. E. Rollenhagen, C. R. Olson, Low-Frequency Oscillations Arising From Competitive Interactions Between Visual Stimuli in Macaque Inferotemporal Cortex. *Journal of Neurophysiology* **94**, 3368-3387 (2005).
38. S. Moldakarimov, J. E. Rollenhagen, C. R. Olson, C. C. Chow, Competitive Dynamics in Cortical Responses to Visual Stimuli. *Journal of Neurophysiology* **94**, 3388-3396 (2005).
39. M. W. Oram, Contrast induced changes in response latency depend on stimulus specificity. *J Physiol Paris* **104**, 167-175 (2010).
40. E. K. Miller, P. M. Gochin, C. G. Gross, Habituation-like decrease in the responses of neurons in inferior temporal cortex of the macaque. *Vis Neurosci* **7**, 357-362 (1991).

2.7 Individual monkey figures

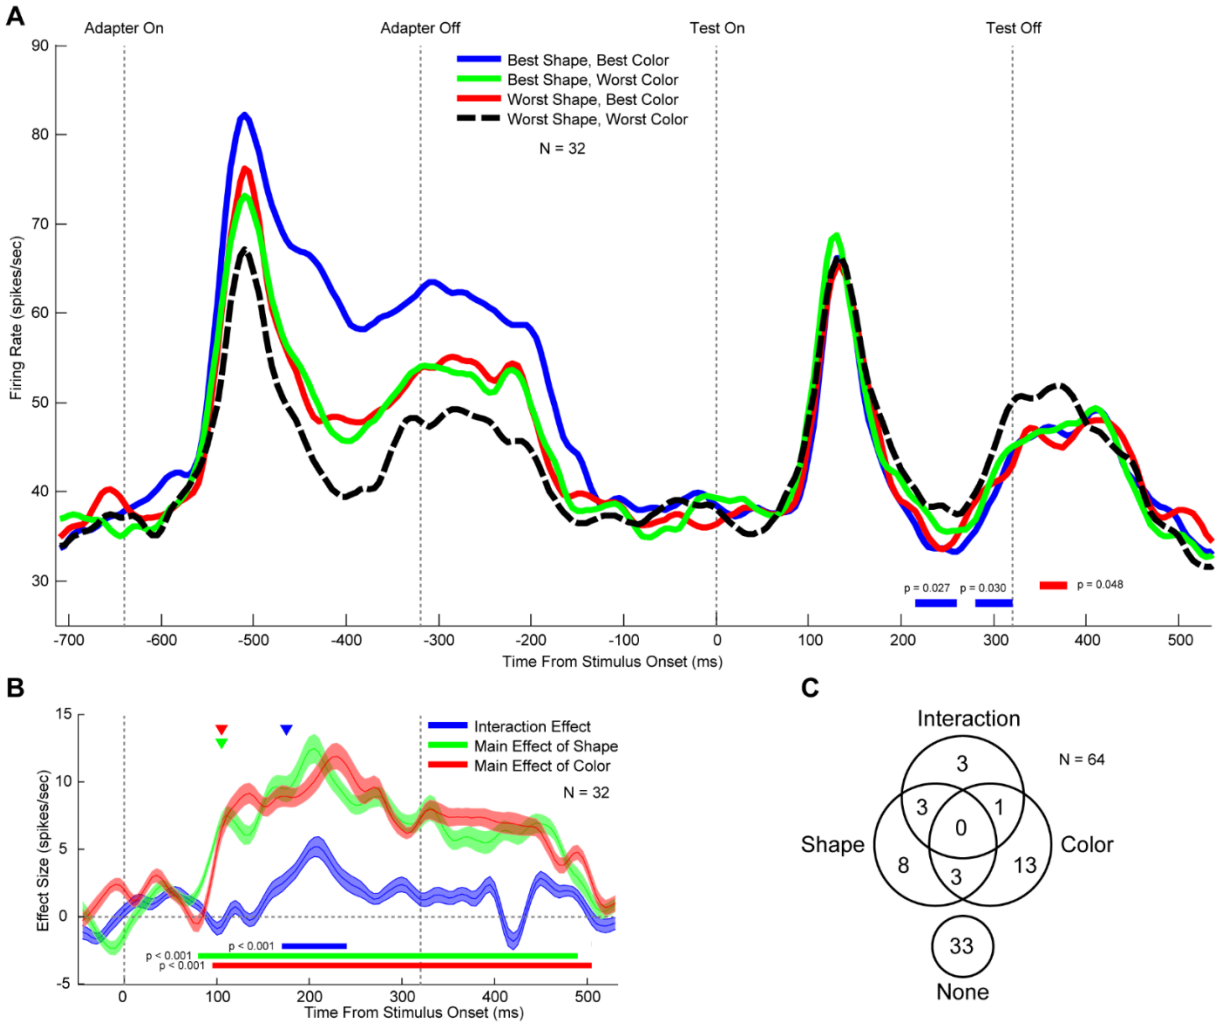

Figure S2.1. Responses to the adapter selective for shape and color in Monkey 1. Five out of seven interaction effects were positive and thus indicative of conjunction selectivity.

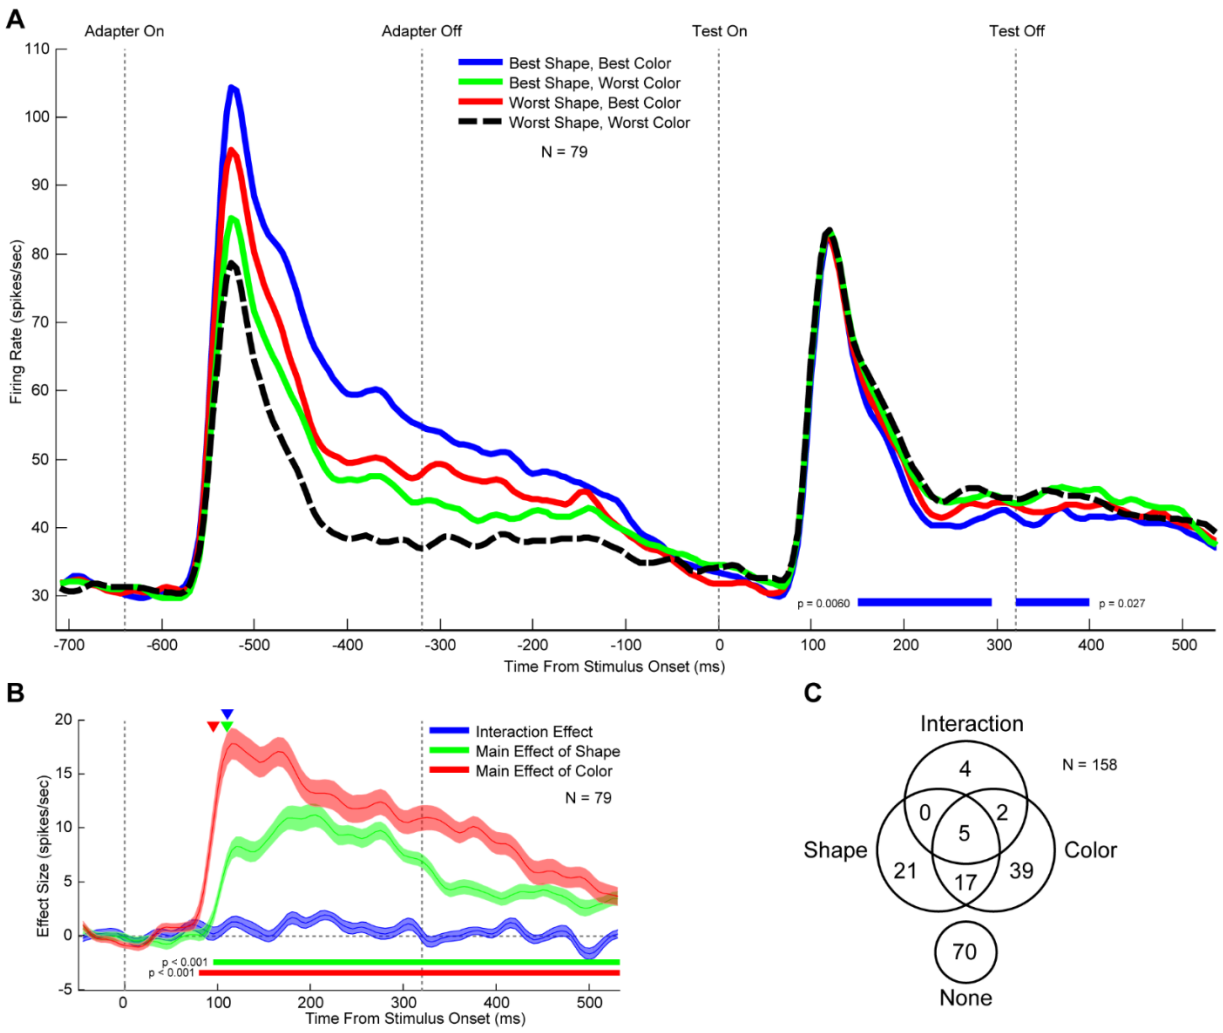

Figure S2.2. Responses to the adapter selective for shape and color in Monkey 2. Nine out of eleven interaction effects were positive and thus indicative of conjunction selectivity.

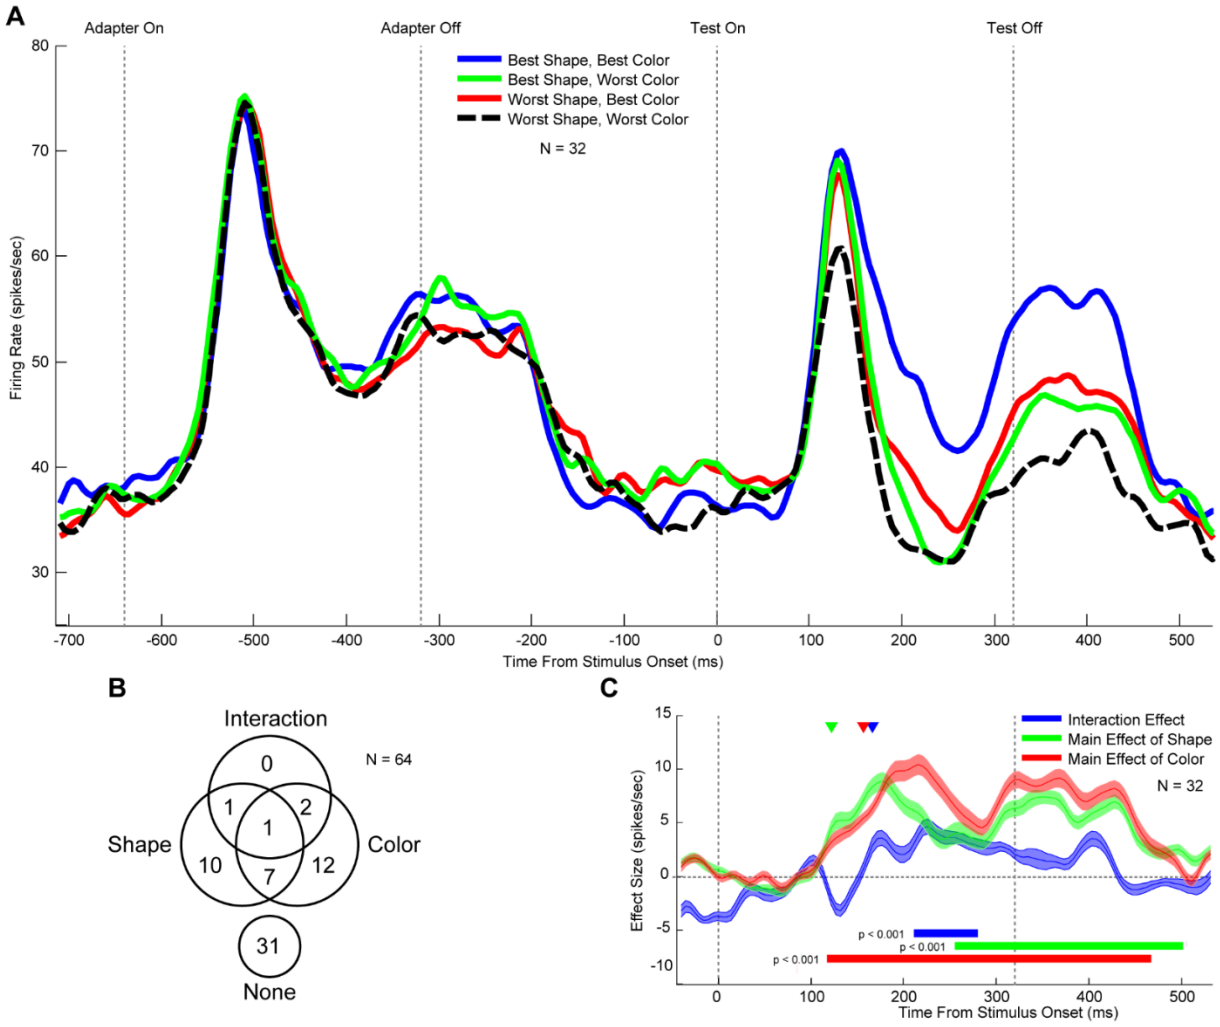

Figure S2.3. Responses to the test image selective for shape and color in Monkey 1. Significant main effects of shape were positive in 17/19 cases, indicating congruence with shape selectivity in responses to the adapter. Significant main effects of color were positive in 22/22 cases, indicating congruence with color selectivity in responses to the adapter. Interaction effects in 2/4 cases were positive and thus indicative of conjunction selectivity.

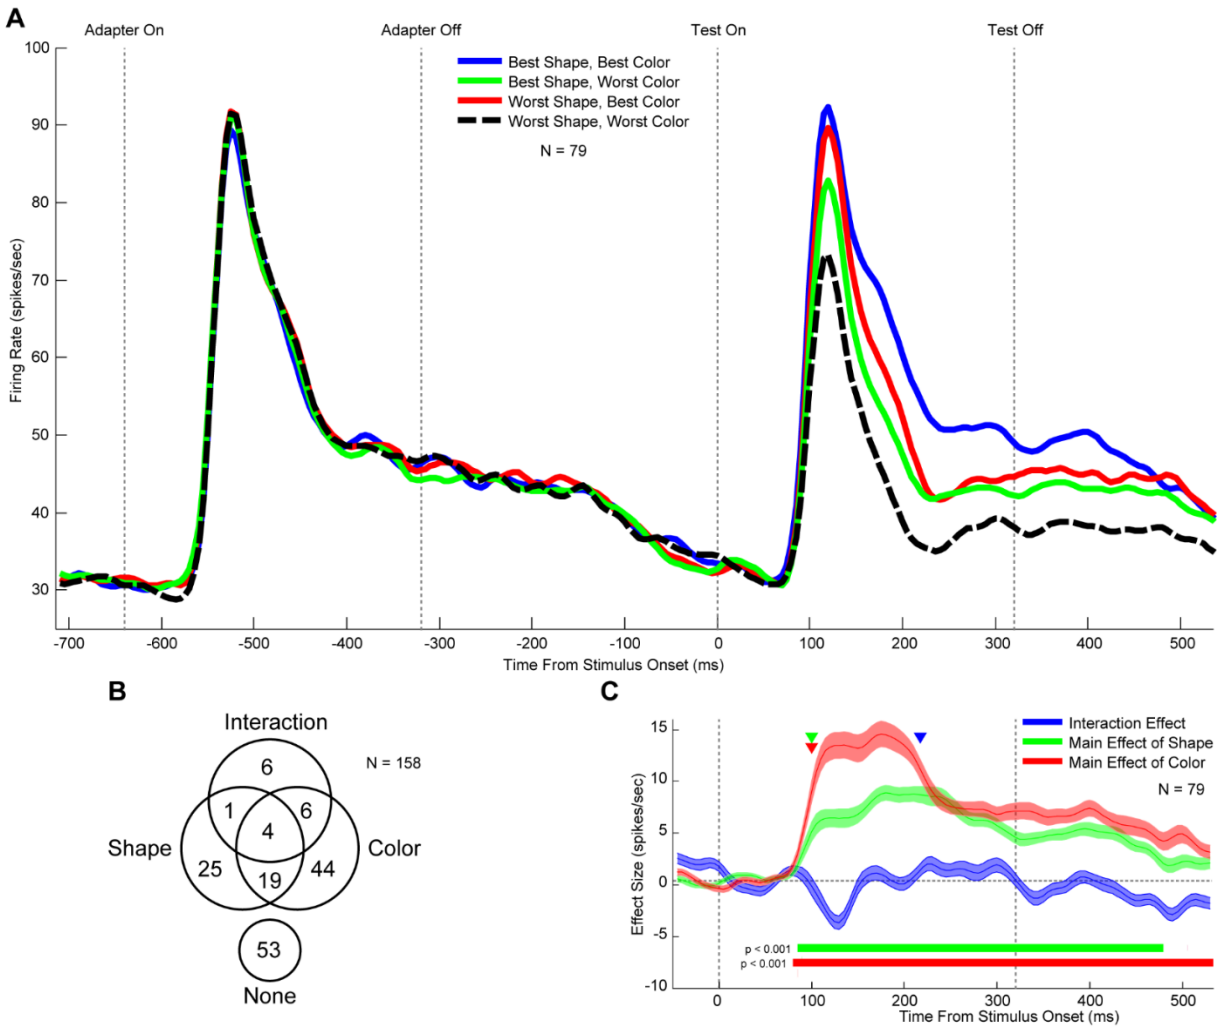

Figure S2.4. Responses to the test image selective for shape and color in Monkey 2. Significant main effects of shape were positive in 45/49 cases, indicating congruence with shape selectivity in responses to the adapter. Significant main effects of color were positive in 72/73 cases, indicating congruence with color selectivity in responses to the adapter. Interaction effects in 13/17 cases were positive and thus indicative of conjunction selectivity.

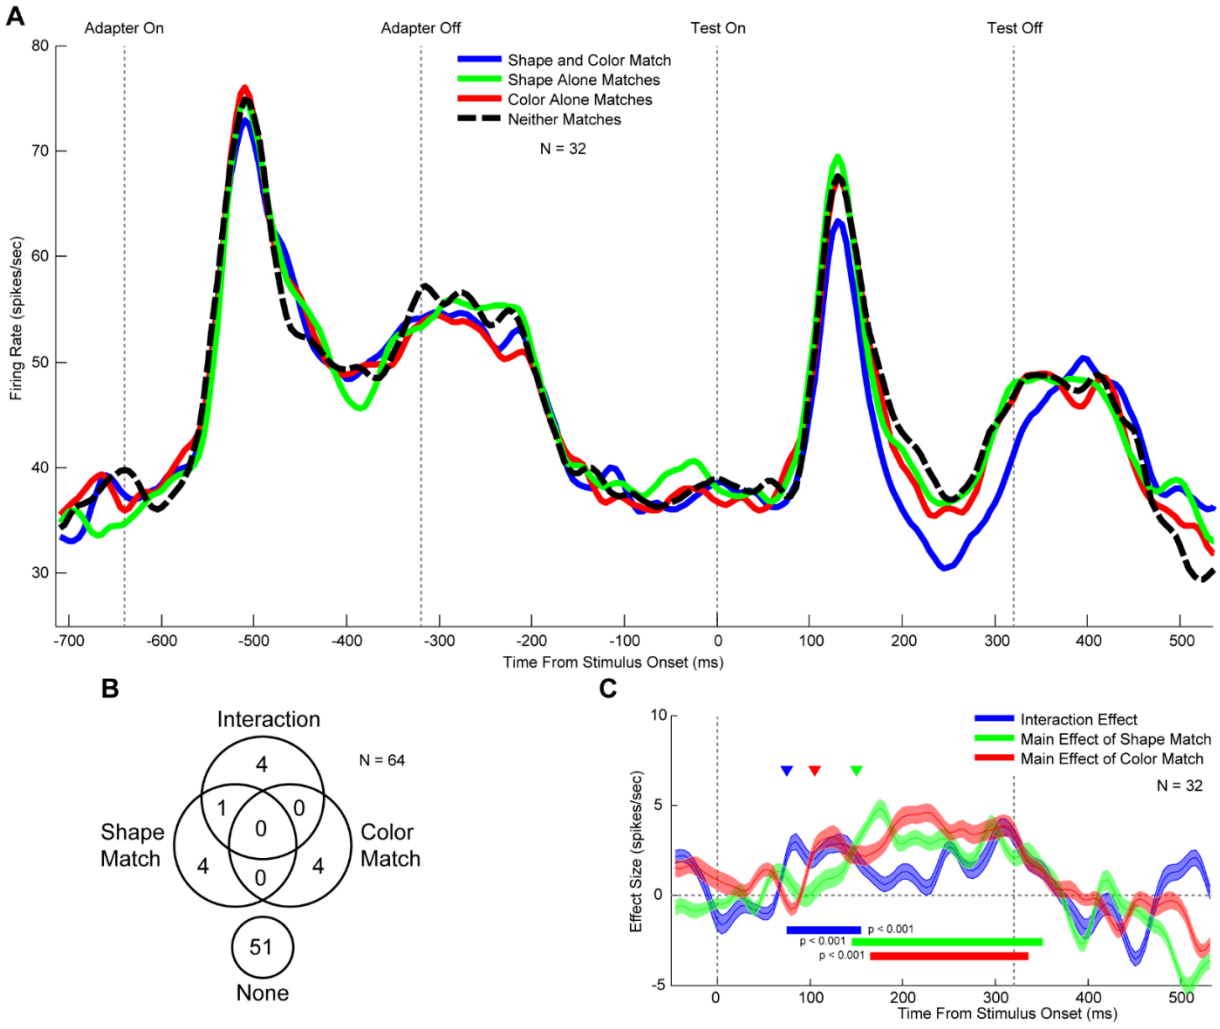

Figure S2.5. Suppression of the response to the test image in Monkey 1. Significant main effects of shape match were positive in 5/5 cases, indicating suppression. Significant main effects of color match were positive in 3/4 cases, indicating suppression. Interaction effects in 4/5 cases were positive and thus indicative of especially deep suppression when the test image was identical to the adapter.

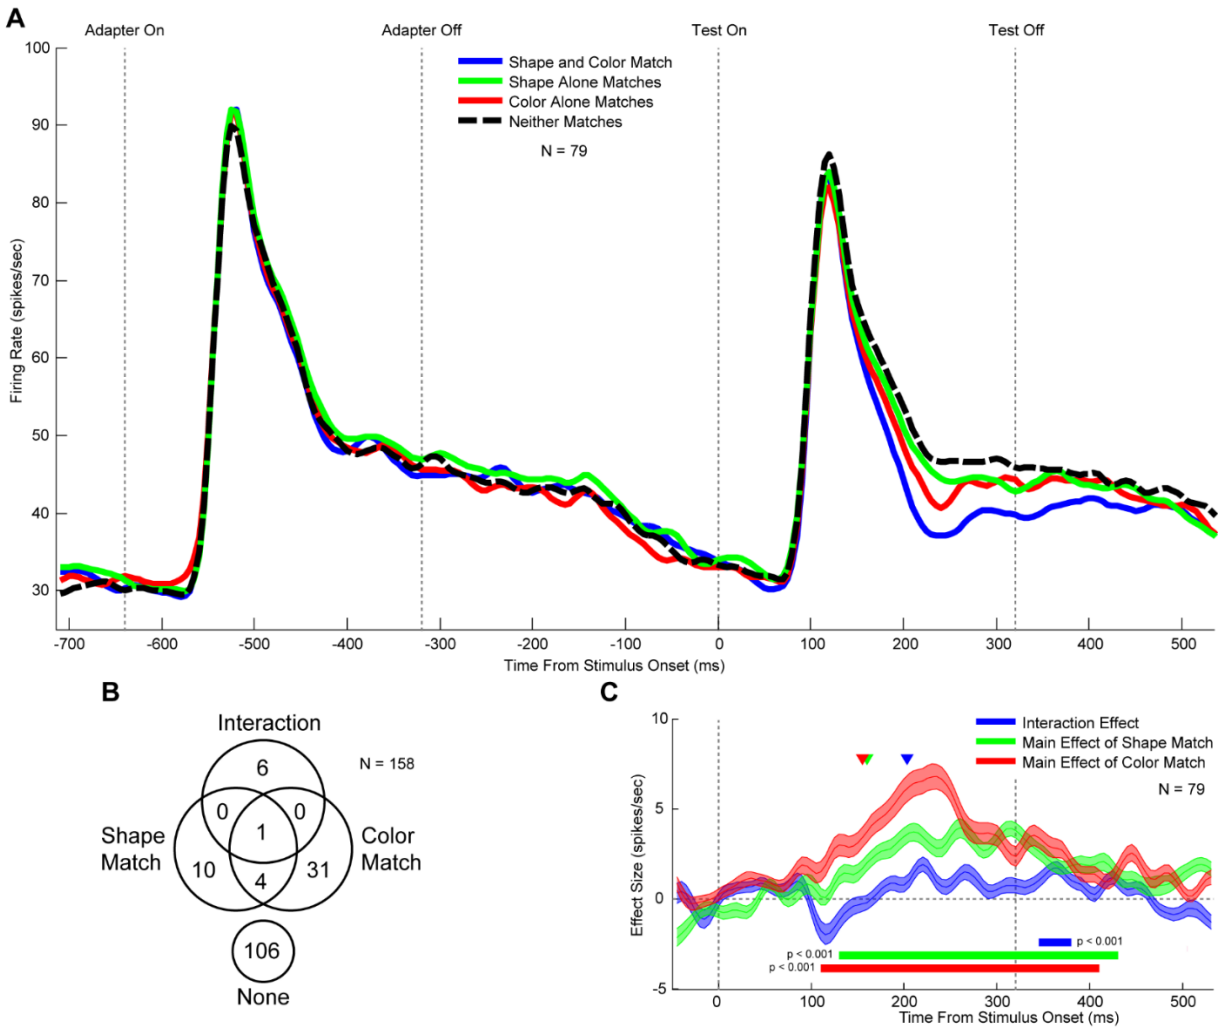

Figure S2.6. Suppression of the response to the test image in Monkey 2. Significant main effects of shape match were positive in 15/15 cases, indicating suppression. Significant main effects of color match were positive in 32/36 cases, indicating suppression. Interaction effects in 4/7 cases were positive and thus indicative of especially deep suppression when the test image was identical to the adapter.

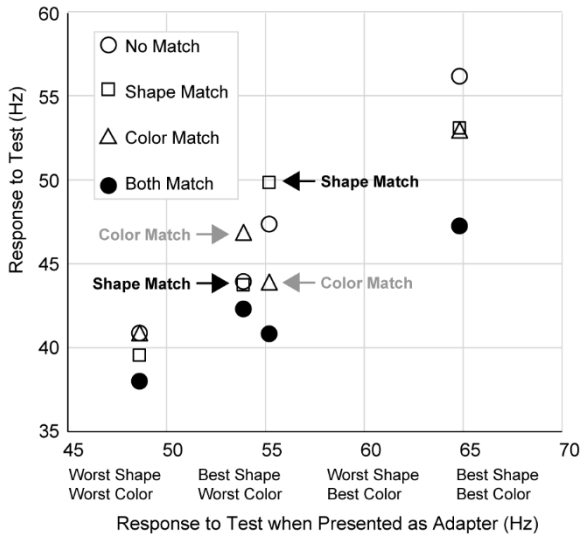

Figure S2.7. Dependence of repetition suppression on the neuron's preference for the repeated shape or color in Monkey 1. The response to the test when presented as adapter was nearly the same for "best shape worst color" (54.4 Hz) and for "worst shape best color" (54.7 Hz). The symbols have been offset symmetrically to left and right to allow visibility.

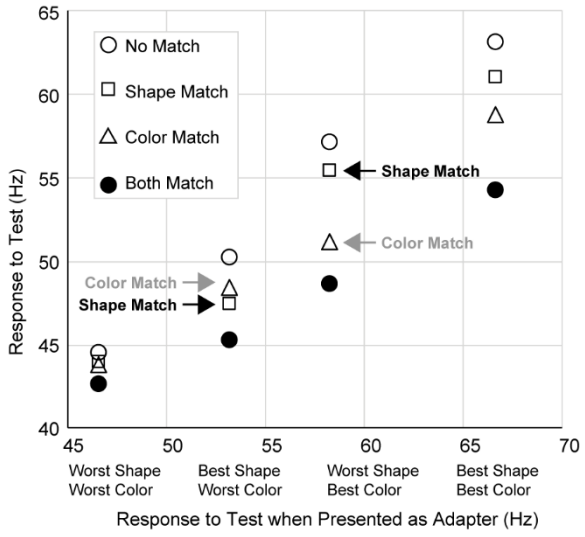

Figure S2.8. Dependence of repetition suppression on the neuron's preference for the repeated shape or color in Monkey 2.

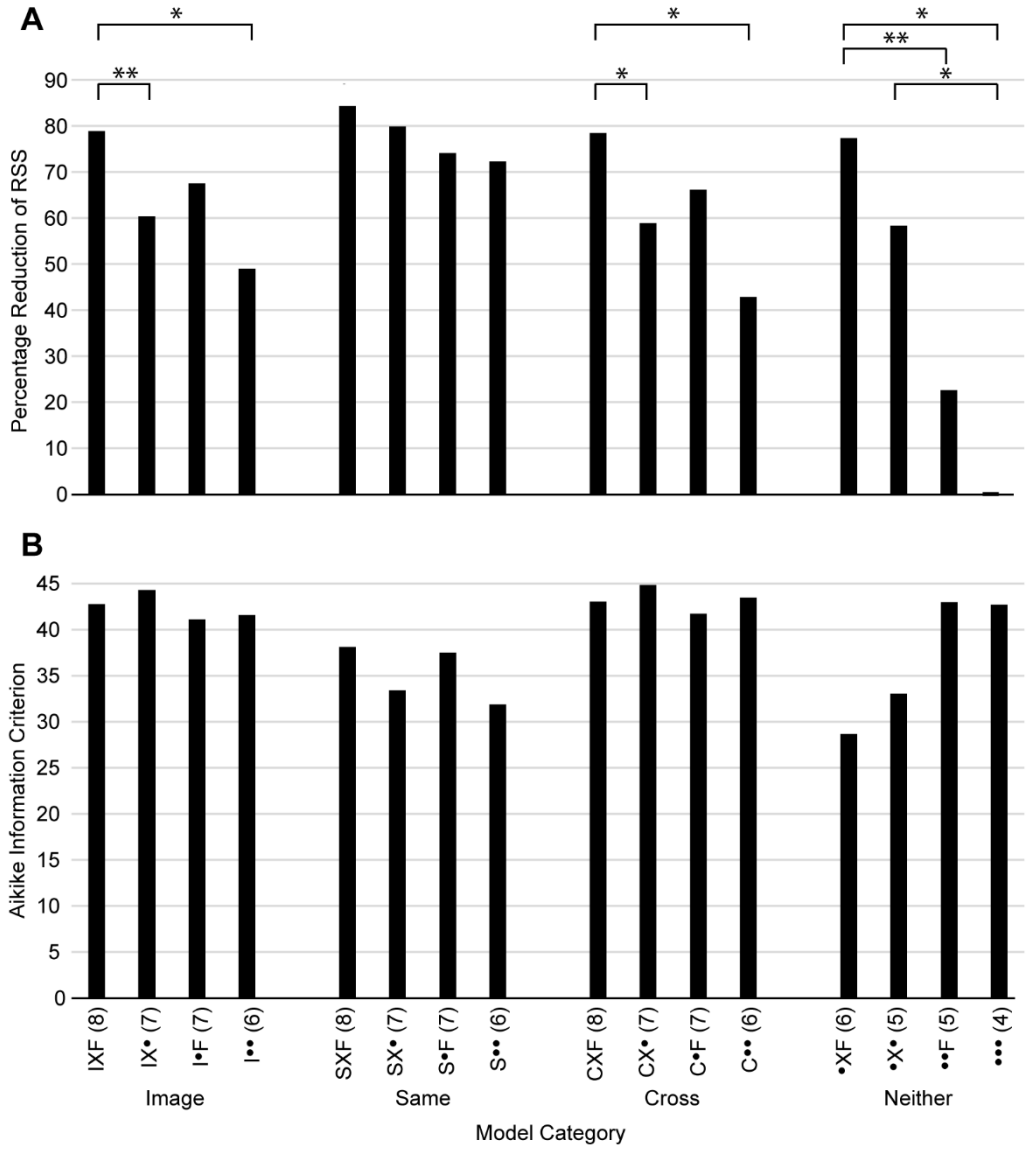

Figure S2.9. Performance measures for 16 models parametrically adjusted to provide the best fit to population data from Monkey 1 as shown in Figure S2.5A.

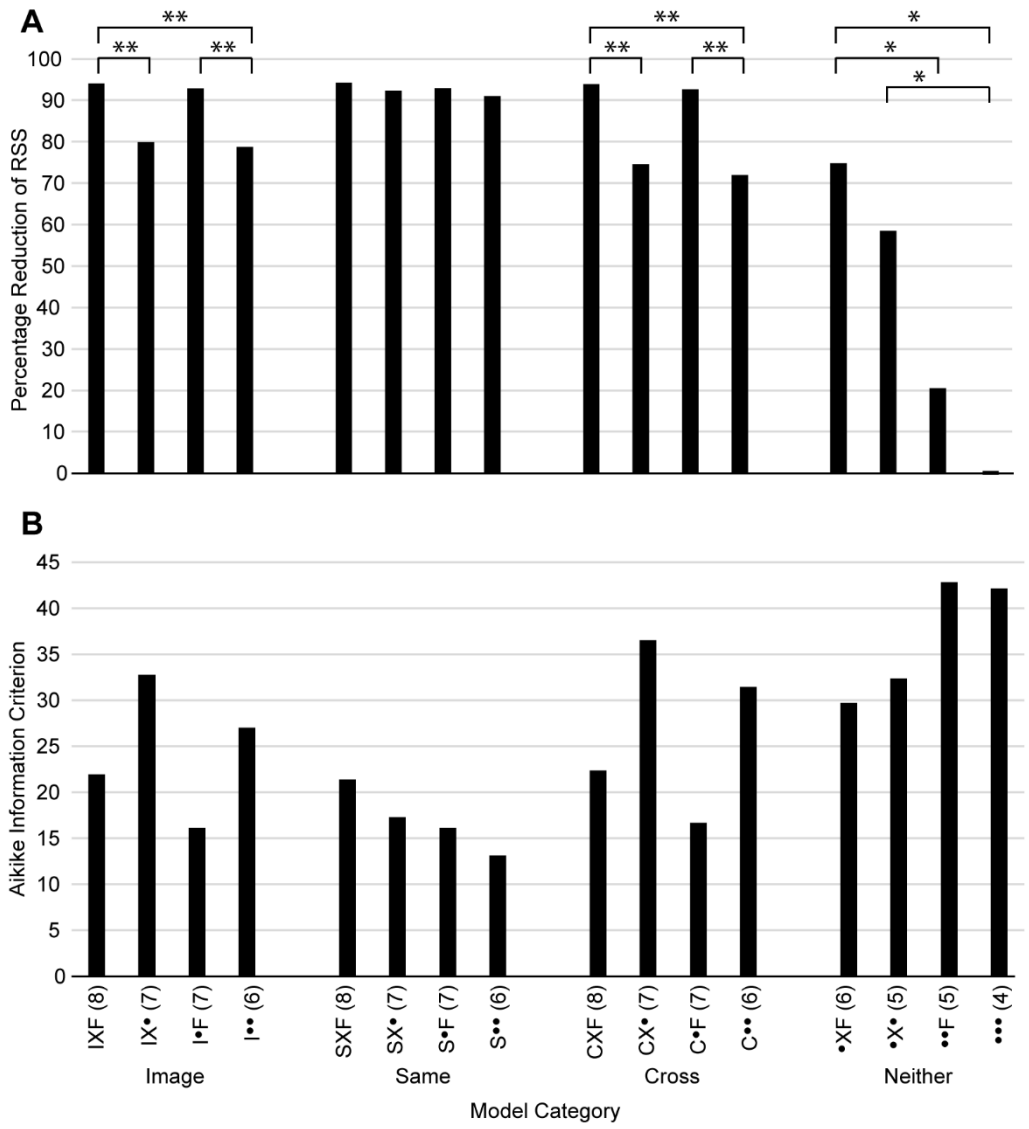

Figure S2.10. Performance measures for 16 models parametrically adjusted to provide the best fit to population data from Monkey 2 as shown in Figure S2.5A.

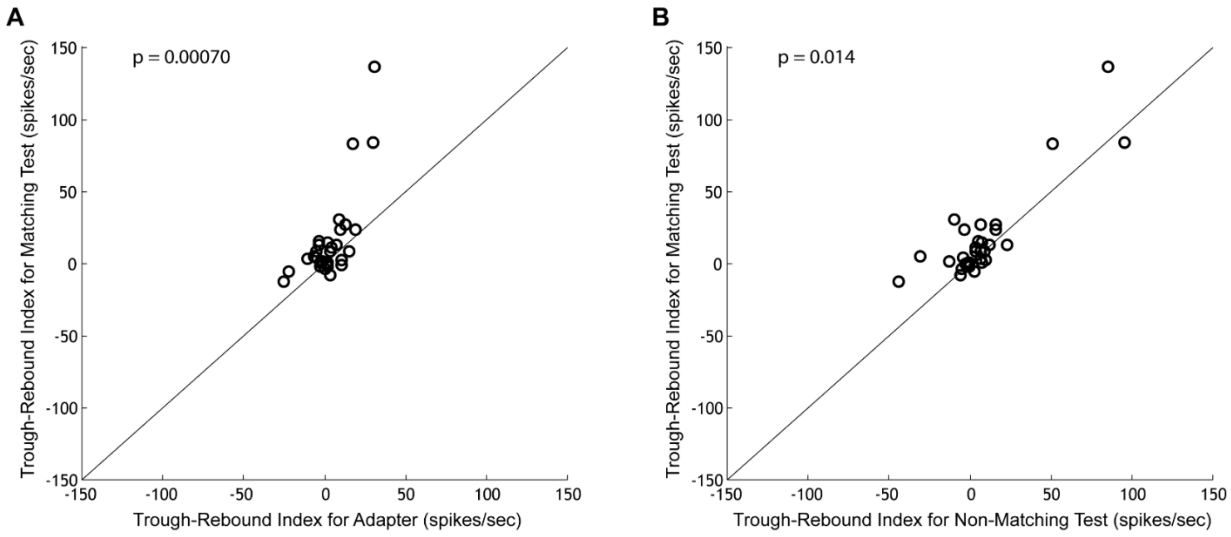

Figure S2.11. The trough-rebound response dynamic was most pronounced for a test image matching the adapter in both shape and color in Monkey 1. The indicated p-values are from a two-sided paired-sample signed rank test.

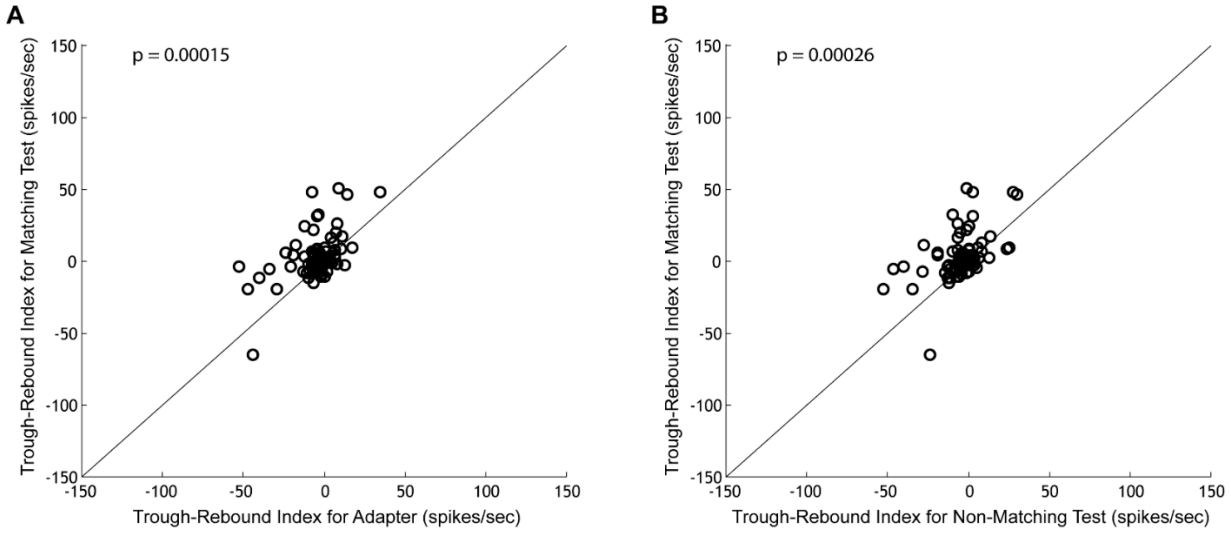

Figure S2.12. The trough-rebound response dynamic was most pronounced for a test image matching the adapter in both shape and color in Monkey 2. The indicated p-values are from a two-sided paired-sample signed rank test.

CHAPTER III

INDEPENDENT REPETITION SUPPRESSION IN MACAQUE AREA V2 AND INFEROTEMPORAL CORTEX

3.1 ABSTRACT

Neurons in area TE of macaque inferotemporal cortex exhibit repetition suppression. When an image is presented twice at the same location, first as adapter and then as test, the neuronal response to the test is reduced relative to the response to the adapter. This effect is image-specific as indicated by the fact that suppression is greatest when the test is identical to the adapter. It also generalizes across locations as indicated by the fact that repetition suppression still occurs when the adapter and the test are presented in different visual field quadrants, albeit at a slightly reduced strength. It is not known to what extent low level visual areas such as V2 exhibit the same properties of repetition suppression or indeed whether they exhibit repetition suppression at all under the conditions used to test it in TE. The aim of the study presented here was to determine to what extent V2 exhibits repetition suppression for large natural images and to determine to what extent effects observed in TE and V2 may be influenced by bottom-up and top-down processes. In both TE and V2 under the same conditions we tested repetition suppression for complex objects where we controlled both image identity and presentation location independently. In TE we replicated previous findings that repetition suppression is image-specific and survives changes in location. We found that V2 neurons show robust repetition suppression under the same stimulus conditions used to test TE neurons, although with some important differences. In V2 the spatial generalization was limited to images which appeared in close proximity within the same visual field quadrant. There was also a location-specific component whereby a different image presented at the same location induced a degree of

suppression. Like in TE, there was also a trough-rebound dynamic induced by image-specific repetition suppression in V2. Importantly, we found that repetition suppression effects in V2 precede those in TE by a number of measures, suggesting that repetition suppression for large complex images in V2 is not mediated by top-down influences from areas of higher order such as TE.

3.2 INTRODUCTION

Neurons in macaque inferotemporal cortex respond less strongly to a complex visual image when it is repeated than on its first presentation (*1*). This phenomenon is termed repetition suppression. Repetition suppression was first discovered (*2, 3*) in anterior area TE (*4*), a division of inferior temporal cortex on which most subsequent studies of the phenomenon have focused. In light of the fact that neurons in TE respond selectively to complex images, it is reasonable to speculate that repetition suppression arises within TE and is confined to TE and areas downstream from it. This idea is supported by the observation that repetition suppression survives spatial displacement between the first and second images in a sequence, which implies that the phenomenon depends at least in part on neurons with large receptive fields such as are typical of TE (*5-8*). It does not, however, rule out the possibility that repetition suppression occurs even in low-order areas.

Neurons in area V1 of cat and monkey and in a single study of area V2 in the cat, exhibit adaptation in response to prolonged stimulation with uniform displays rich in a particular feature, such as a drifting grating (*9, 10*). Accordingly, it is plausible that they might exhibit repetition suppression following brief presentation of a structured image. This possibility has been explored with mixed results in human fMRI-based studies (*11, 12*). Numerous studies have demonstrated direction-specific and orientation-specific adaptation in human V1 and V2 using

simple stimuli including drifting gratings (*13-20*). However, complex natural images, unlike gratings, are not optimized to drive subsets of direction- and orientation-specific neurons. Moreover, among neurons with small receptive fields, the typically large natural images are likely to elicit surround suppression as well as center excitation, with uncertain consequences for adaptation (*21*). Most reports based on monitoring blood-oxygen-level-dependent (BOLD) activation in humans during brief presentation of complex natural images describe suppression as being confined to visual areas of relatively high order, including lateral occipital cortex (LOC), the presumed homologue of macaque TE (*22-27*). At least one study has documented repetition suppression even in V1 and V2 but with the qualification that the depth of suppression is weaker than in areas of high order (*28*). However, repetition suppression in V1 and V2, as characterized at the level of the BOLD signal, must remain ambiguous at best inasmuch as it could occur at short latency (indicating an endogenous origin) or long latency (as the product of top-down transmission) and could reflect spiking activity or subthreshold synaptic events. fMRI is not well suited to making these distinctions.

To determine whether repetition suppression occurs in V2 and, if so, to determine its relation to repetition suppression in TE, we measured neuronal responses to sequential pairs of complex images in both areas under four conditions. On interleaved trials, the first and second items in a sequence were identical with regard to identity and location, identity alone, location alone or neither attribute. We found that repetition suppression does indeed occur in V2. On comparing V2 to TE with regard to the pattern of repetition suppression, we found clear differences with regard to timing (suppression latency was shorter in V2 than in TE) and location-dependence (suppression survived a larger spatial offset between adapter and test in TE than in V2). We

conclude that repetition suppression arises independently at multiple (and perhaps all) stages of the ventral stream. It is neither confined to nor dependent on TE.

3.3 MATERIALS AND METHODS

3.3.1 Subjects

The experimental subjects were three adult rhesus macaques (*macaca mulatta*) here designated M1 (male; 12 kg; laboratory designation Ol, used for IT and V2 experiments) and M2 (male; 6 kg; laboratory designation Sc, used for IT experiments) and M3 (male; 13 kg; laboratory designation Rs, used for V2 experiments). All procedures were in accordance with guidelines set forth by the United States Public Health Service Guide for the Care and Use of Laboratory Animals and were approved by the Carnegie Mellon University IACUC.

3.3.2 V2 Receptive Field Mapping

At the outset of each session involving V2, we located the center of the receptive field of the neuron under study by means of the following steps. We first crudely mapped it out with manually controlled stimuli. We then carried out automatic assessment of horizontal and vertical spatial selectivity within a 3° by 3° grid centered on the field's estimated location, using as stimuli bars with a width of 0.3° and a length of 8°. Ten horizontal bars centered horizontally on the grid spanned it in 0.3° vertical steps. Ten vertical bars centered vertically on the grid spanned it in 0.3° horizontal steps. Each bar contained high contrast content as if it had been cut from a square-wave grating with spatial frequency of 3.3 cycles/degree and an orientation of +45° or -45° relative to vertical. While the monkey maintained central fixation, the bars were presented for 300 ms each (150 ms at each of the two orientations with no interstimulus interval), with an interstimulus interval of 200 ms between locations, in random sequence until each location had been tested three times. Tuning for location with respect to each axis was clear from inspection

of the resulting post-stimulus-time histograms. In subsequent tests, we centered stimuli at the X coordinate of the vertical bar eliciting the strongest response and the Y coordinate of the horizontal bar eliciting the strongest response.

3.3.3 Task

Each trial began with presentation of a fixation spot centered at the midpoint of an LCD monitor with a 60 Hz refresh rate at a viewing distance of 32 cm. Once the monkey had attained fixation, the following displays appeared in sequence: fixation spot (320 ms), adapter stimulus (320 ms), fixation spot (320 ms), test stimulus (320 ms), fixation spot (320 ms). The monkey was required to maintain gaze on the fixation spot throughout the trial, within a window subtending 1.4-2.1°. Gaze was monitored with an ISCAN video-based eye tracking system. Upon successful completion of a trial, juice reward was delivered. Any fixation break terminated the trial and triggered onset of a checkerboard display which remained visible for two seconds or until the monkey had fixated it for a cumulative duration of 300 ms, whichever came first. Behavior was monitored and stimulus presentation and reward delivery were controlled by a PC running NIMH Cortex.

3.3.4 Stimuli

The fixation spot was a white disk 0.3° in diameter. The adapter and test stimuli were session-novel background free images of naturalistic objects. The horizontal or vertical extent of each image, whichever was greater, was 5.3°. The orthogonal dimension was close to this value, with the consequence that the aspect ratio was close to one. The library of available stimuli included 1,000 such images. Adapter and test stimuli could appear at one of two peripheral locations in the visual hemifield contralateral to the site of the recording electrode. Each stimulus appeared either above or below the horizontal meridian. For IT recordings, stimuli appeared 7.3°

from the vertical meridian and 7.1° either above or below the horizontal meridian. For V2 recordings, horizontal and vertical offsets were such that the lower-quadrant stimulus was centered on the receptive field of the recorded neuron and the upper-quadrant stimulus was at a location symmetric with respect to the horizontal meridian. The average horizontal offset from fixation was 1.8° with a range of 0.4° to 4.4° and the average vertical offset was 5.0° with a range of 3.6° to 5.9°. A 1.6° diameter invisible mask surrounded the fixation spot, preventing the stimulus from impinging on fixation. In sessions employing the "near" geometry, each image was cropped to a central disk 3° in diameter. The lower stimulus was centered on the receptive field of the recorded neuron and the upper stimulus was centered 3° higher. Before beginning any recording session, we assessed neuronal image selectivity by monitoring responses elicited by a randomly chosen subset of 16 novel images. We selected for use in a given session four images that elicited the strongest responses as judged by inspection of online raster and histogram displays. Arbitrarily numbering the images from 1 to 4, we constructed from them two dyads of stimuli. Dyad A consisted of images 1 and 2 and Dyad B consisted of images 3 and 4.

3.3.5 Session structure

From the two images in each dyad, 8 adapt-test sequences could be created. These were: Two in which the image repeated at the same location, two in which the image repeated at a different location, two in which a different image appeared at the same location and two in which a different image appeared at a different location. Each block of 16 successful trials contained one instance of each adapt-test sequence for each dyad. Sequencing was pseudo-random subject to the constraint that the adapter and test were drawn from different dyads on alternate trials. We adopted this design to minimize the potential for cross-trial carryover of image-specific repetition suppression. Trials terminated due to fixation-break were repeated to ensure that one

trial under each of the 16 conditions was completed successfully. A full run consisted of four successive blocks and thus of 64 successful trials encompassing four repetitions of each possible adapt-test sequence for each dyad.

3.3.6 Neurophysiological data collection

Cylindrical Cilux recording chambers with an inner diameter of approximately 2 cm (Crist Instrument Co. Inc., Hagerstown, MD) were implanted over the left hemispheres of each monkey over IT (M1 and M2) and over the left hemisphere (M1) and right hemisphere (M3) of V2. IT chambers were center approximately 16 mm anterior and 22 mm lateral to Horsley-Clarke zero. V2 chambers were centered approximately 18 mm posterior and 6 mm lateral to Horsley-Clarke zero.

Chamber placement was guided by pre-surgical T1-weighted structural MRI scans in a 4.7-Tesla scanner and was confirmed by post-surgical scans. The IT chambers gave access to cortex both in the ventral bank of the superior temporal sulcus and on the adjacent inferior temporal gyrus. The recording sites, as judged by reference to standard maps of cytoarchitecturally defined areas, lay within anterior area TE (*29*) and did not encroach on perirhinal cortex (*30*). The V2 chambers gave access to cortex on the posterior bank of the lunate sulcus. V2 chamber grid holes were functionally mapped to retinotopic coordinates and the line of inversion between V2 and V1 was located to determine the extent of access to V2 cortex. On each recording day, a cylindrical plug containing guide holes arranged at 1 mm spacing in a square grid was inserted in the chamber (Crist Instrument Co. Inc., Hagerstown, MD). A dura-penetrating guide tube was inserted through one guide hole and a single varnish coated tungsten microelectrode (FHC Inc., Bowdoin, ME) with an impedance of 0.1-5 MΩ at 1kHz was advanced through the guide tube by use of a hydraulic micromanipulator (M0-10; Narishige International Inc., East Meadow, NY).

Upon encounter with neurons giving phasic responses to visual stimuli, recording commenced. Threshold-crossing events, sampled at 40 KHz, were digitally recorded and stored for offline sorting by a Plexon MAP system (Plexon Inc., Dallas, TX).

3.3.7 Database

We sorted waveforms from each session using a PCA-based approach implemented by Plexon Offline Sorter (Plexon Inc., Dallas, TX). All further steps of analysis were carried out in Matlab on time-stamped action potential markers. We first analyzed data from each neuron to ensure that it met an arbitrary criterion for visual responsiveness. Only if a one-tailed t-test comparing mean firing rate 100-200 ms after adapt-stimulus onset to mean firing rate 100-200 before adapt-stimulus onset yielded an outcome of $p < 0.05$ did we include the neuron in the database for subsequent analyses.

3.3.8 Cluster-Based Permutation Test

To determine whether the post-image-onset time-varying population firing rates measured under two conditions were significantly different, we employed a nonparametric approach requiring neither arbitrary designation of a measurement window nor comparisons within multiple windows *(31)*. The starting point for this analysis was a table of mean firing rates of each neuron under each condition in each 5 ms bin spanning a window from 0 ms to 560 ms following image onset. In each bin, we carried out a paired T-test on the two distributions of firing rates. If the test yielded a p-value < 0.05, we tagged the bin as positive or negative according to which condition was associated with the higher firing rate. If the test yielded a p-value ≥ 0.05, we tagged the bin as zero. The T-test was used only as a means for imposing an arbitrary threshold and not to assess statistical significance. For each cluster of bins of uniform sign, either positive or negative, we computed the sum across those bins of the associated T-

statistics. The cluster with the greatest sum was classified as "best" for further steps of analysis. To test its statistical significance, we generated a permutation distribution, applying the above-described procedure 1,000 times to data in which the condition labels for each neuron had been randomly shuffled. We computed p as the fraction of iterations in which the best-cluster sum of T-statistics was greater than the best-cluster sum of T-statistics in the original data. We classified the original cluster as significant if $p < 0.05$. If this cluster was significant, we proceeded to assess the statistical significance of the observed cluster with the next highest sum of T-statistics. We computed p for this cluster as the fraction of cases in the previously generated permutation distribution for which the sum of T-statistics was greater than the observed value. We classified the cluster as significant if $p < 0.05$. We repeated this procedure until a non-significant result was obtained or no observed cluster larger than four bins remained.

3.3.9 Analysis of Suppression Timing

At the population level, we measured suppression as the difference in mean firing rate between trials in which the test image differed from the adapter and trials in which it was the same. Having smoothed the data in 1 ms bins with a 10 ms SD Gaussian kernel, we proceeded to compute the time of onset as the zero-intercept of a line passing through the peak of the curve and the most recent antecedent point at which the first derivative was zero, the second derivative was positive and the signal had not yet begun its consistent rise. At the level of the individual neuron, we computed the time to half-peak of the suppression versus time curve. The analysis was confined to trials in which adapter and test image appeared at the same location. It was based on data in 1 ms bins smoothed with a 10 ms SD Gaussian kernel. We defined the peak as the maximum value 50-300 ms post stimulus, and the time to half-peak as the earliest time at which the effect reached half this value in the window 50-300 ms following test image onset.

Neurons which reached half-peak before 50 ms were excluded from this analysis. Having computed time to half-peak for each neuron, we constructed a cumulative frequency plot of half-peak times for each area. We compared the cumulative distributions for V2 and TE using a Kolmogorov-Smirnov test (kstest2, Matlab, Mathworks).

3.4 RESULTS

3.4.1 Repetition suppression in TE

We monitored the activity of 103 visually responsive TE neurons while monkeys performed a task requiring them to maintain central fixation during successive peripheral presentation of an adapter and a test image (**Figure 3.1A**). Each image could appear in either the upper or the lower quadrant of the visual field contralateral to the recording hemisphere. Images were centered at 7.3° horizontal and 7.1° vertical eccentricity relative to fixation (gray squares in left panel of **Figure 3.1A**). For testing each neuron, we employed two pairs of images. To minimize cross-trial adaptation, images from pair 1 and pair 2 were presented on alternate trials. The adapter might be presented in either the upper or the lower quadrant. The test image was always presented in the lower quadrant. Trials involving each pair of images conformed to eight conditions obtained by crossing adapter identity (image A or B) with test image identity (same or different) and test image location (same or different) (**Figure 3.1D**). In describing the results, we will focus on data combined across the two monkeys used for each experiment. Key phenomena present in the combined data were evident in each monkey considered individually (**Figures S3.1-S3.6**).

Independently for conditions in which adapter and test appeared at the same location and at different locations, we asked whether repetition suppression occurred, as indicated by lower mean population firing rate when the test matched the adapter than when it differed from it.

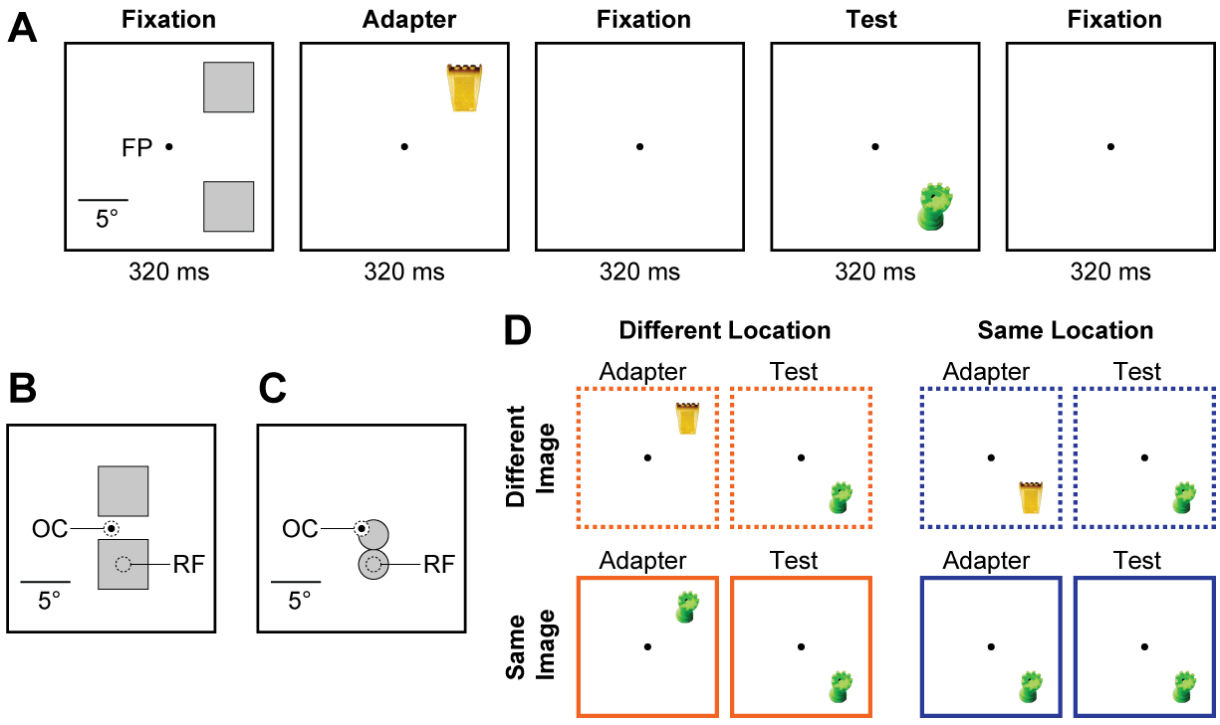

Figure 3.1. Task design. A. On each trial, while the monkey maintained gaze on a central fixation point (FP), two images were presented in succession in the visual field contralateral to the recording hemisphere. The gray squares indicate the two locations at which images could appear during recording in TE. In the trial illustrated here, the first (adapter) and second (test) images differed in both identity and location. **B.** During recording in V2, one of the possible locations was centered on the neuronal receptive field (RF) and the other possible location was symmetrically placed with respect to the horizontal meridian. An invisible occluder (OC) prevented impingement of images on the fixation point under conditions in which the RF was close to fixation. The location of the RF in this panel is the average across all neurons tested in this way. **C.** In a subset of late V2 sessions, images cropped to disks 3° in diameter were centered on the receptive field and 3° above it. **D.** On any given trial, the adapter and test could match or differ with respect to identity and could match or differ with respect to location. The test image always appeared in the lower quadrant. The borders of the panels (dashed or solid, blue or orange) indicate how curves representing population firing rate under the four conditions will be coded in **Figures 3.2A, 3.4A and 3.6A**.

Repetition suppression was present (**Figure 3.2A**) and was statistically significant ($p < 0.001$, cluster-based permutation test, **Figure 3.2C**) under both spatial conditions. However, it was deeper under the spatial-match condition than under the spatial-nonmatch condition (blue versus orange in **Figure 3.2C**). To determine whether these effects achieved significance (alpha = 0.05) at the level of individual neurons, we carried out two ANOVAs on data from each neuron, one

confined to same-location trials and the other to different-location trials. In each ANOVA, the repetition status of the test image (same or different) and its identity were factors and the dependent variable was firing rate 50-350 ms following test onset. Counts of cases in which the main effect of repetition status achieved statistical significance are presented in **Figure 3.2B**. Cases of significant suppression were significantly more numerous than expected by chance (chi-squared test, 1 df, with Yates correction) under both the same-location condition ($p = 5.6 E-16$), and the different-location condition ($p = 2.4 E-5$). The count of cases under the same-location condition was, however, significantly greater than under the different-location condition ($p = 0.0032$). The single-neuron results thus mirror the population-level results in indicating that repetition suppression occurs but is reduced in strength when the adapter and the test appear in different visual-field quadrants.

3.4.2 Timing of repetition suppression in TE

Suppression not only was weaker but also appeared to emerge at longer latency under the different-location than under the same-location condition. To quantify this effect, we measured the time relative to test-image onset of two points on each population curve in **Figure 3.2C**. We marked suppression-onset as the time at which a line fitted to the rising phase of the curve crossed zero. For this purpose, we defined the rising phase as extending from the time when the curve had a slope of zero and was beginning to rise to the time when it was at its maximum. Suppression onset began at 60 ms under the same-location condition and 100 ms under the different-location condition. We also computed the time to half-maximum height. This was 118

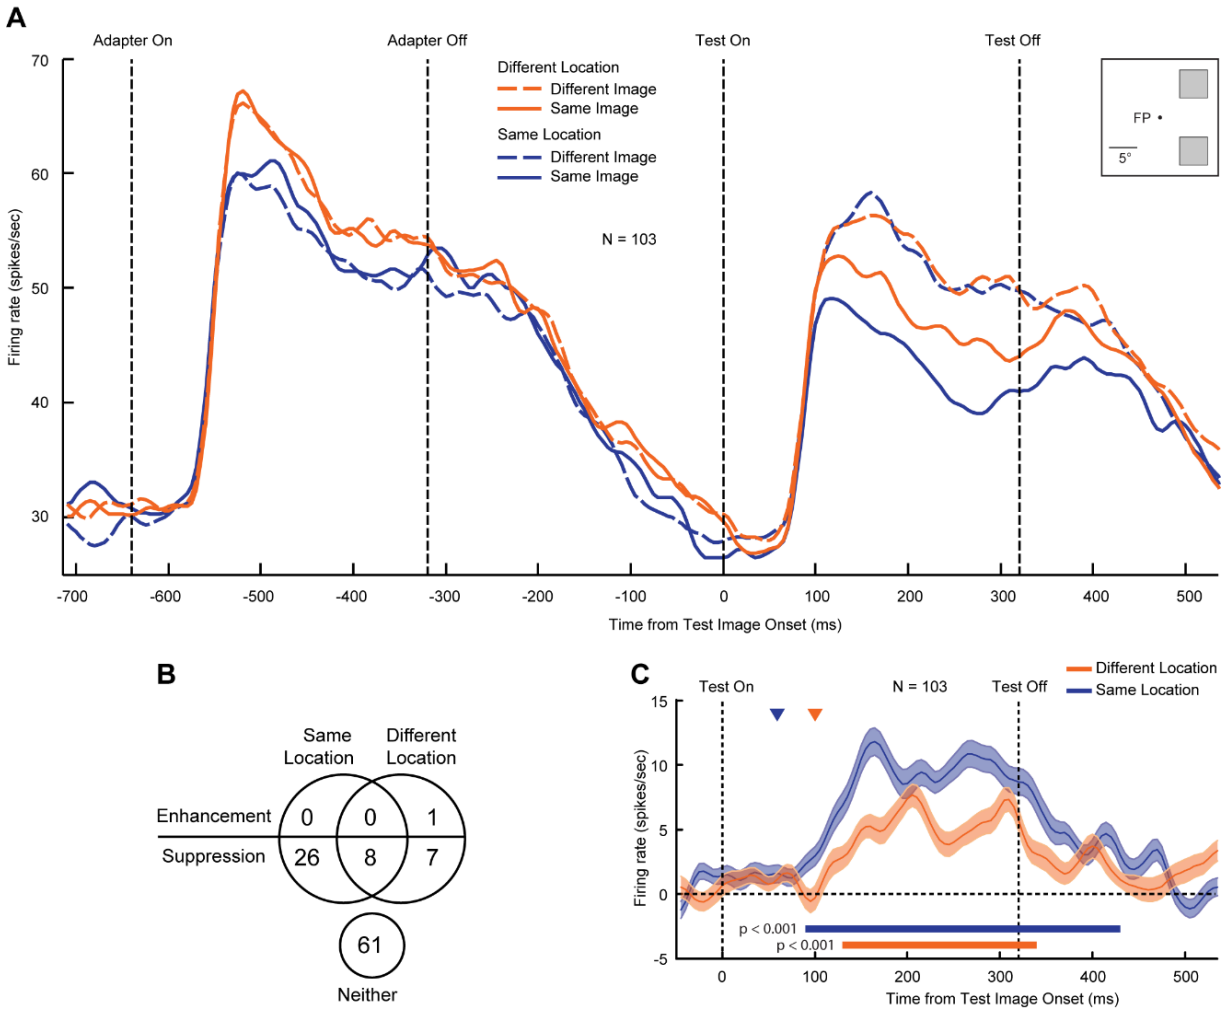

Figure 3.2. Repetition suppression in TE. A. Population mean firing rate as a function of time during the trial under four conditions determined by whether the adapter was in the upper quadrant (different location) or lower quadrant (same location) and whether the test image was identical to the adapter (same image) or not (different image). **B.** Counts of neurons in which firing rate under the same-image condition differed significantly from firing rate under the different-image condition. Counts are subdivided according to whether firing rate was lower under the same-image condition (suppression) or the different-image condition (enhancement). **C.** Repetition suppression as a function of time when adapter and test were at different locations (orange) and at the same location (blue). The measure of suppression is the population mean of firing rate under the different-image condition minus firing rate under the same image condition. Ribbon is ± standard error of the mean. The underlying bars indicate the periods during which the measures were significantly different from zero as determined by a cluster-based permutation test. The p-value for each test is juxtaposed to the corresponding bar. Each triangle indicates, for the plot of corresponding color, the estimated time of suppression onset (see Methods for details). Images in this experiment were placed at locations depicted by gray squares in **Figure 3.1A**. Curves were smoothed by convolution with a 10 ms Gaussian kernel.

ms under the same-location condition and 138 ms under the different-location condition. Thus, by both measures, different-location suppression was delayed relative to same-location suppression.

3.4.3 Impact of repetition on the dynamics of the visual response in TE

Repetition suppression was accompanied by an alteration of response dynamic. When the test matched the adapter with regard to both identity and location, the response to the test took on a trough-rebound configuration (**Figure 3.2A**, solid blue curve). To assess the significance of this effect, we computed, for each neuron under each condition of interest an index designed to assume a positive value in the event of rebound. This was the firing rate during a 50 msec window centered 400 msec following stimulus onset minus the firing rate in a 50 msec window centered 250 msec following stimulus onset. The trough-rebound index for the response to a test image precisely matching the adapter (mean = 2.7 spikes/sec) was significantly greater than corresponding index for the adapter (mean = -0.50 spikes/sec, $p = 0.0070$, two-sided paired-sample rank-sum test) (**Figure 3.3A**). It was also significantly greater than the index for a test image presented at the same location as the adapter but with a different identity (mean = -2.5 spikes/sec, p = 6.9 E-4) (**Figure 3.3B**). We conclude that image repetition genuinely induced a trough-rebound response dynamic.

3.4.4 Repetition suppression in V2

We recorded from 113 visually responsive V2 neurons using an equivalent procedure with the sole exception that the location of the upper-quadrant image was obtained by reflecting the receptive field center about the horizontal meridian (**Figure 3.1B**). We found that repetition suppression occurred in V2 (**Figure 3.4A**) at a statistically significant level ($p < 0.001$, cluster-based permutation test, **Figure 3.4C**) only under the same-location condition. To determine

whether this pattern was detectable at the level of individual neurons, we analyzed data from each neuron, using an approach identical to that applied to data from TE (**Figure 3.4B**). The only effect category in which the count of significant cases exceeded the number expected by chance was same-location image-specific suppression (p = 3.3 E-16, chi-squared test, 1 df, with Yates correction). The onset of same-location suppression, measured at the level of mean population firing rate, occurred at 42 ms following test-image onset. The time to half-maximum suppression was 112 ms.

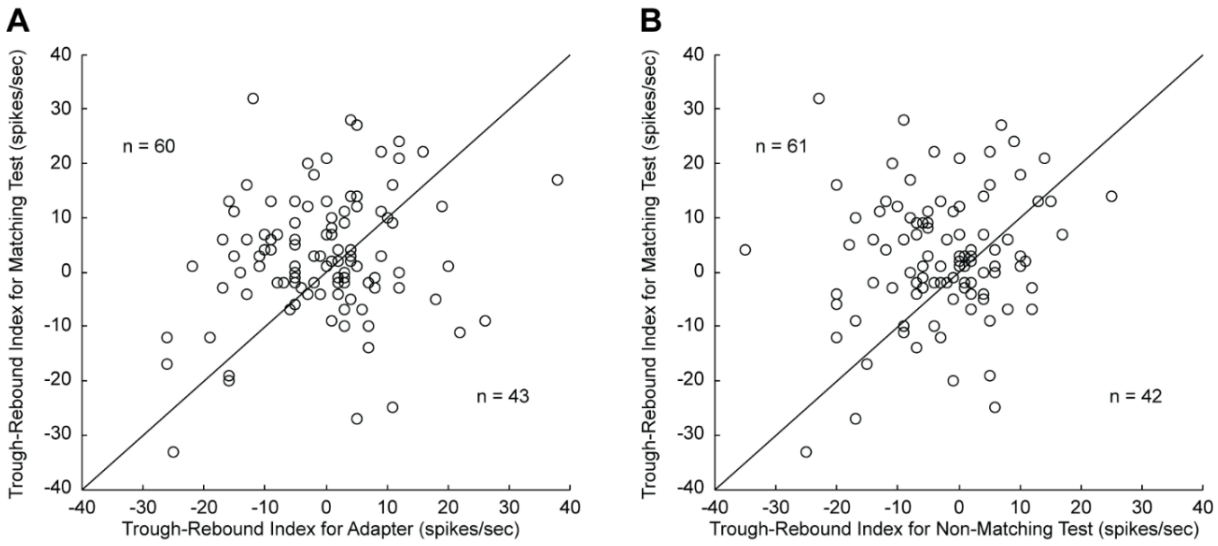

Figure 3.3. The trough-rebound dynamic in TE was enhanced when the test image matched the adapter. A. The trough-rebound index for the response to the test image matching the adapter is plotted against the trough-rebound index for the response to the adapter itself. **B.** The trough-rebound index for the response to the test image matching the adapter is plotted against the trough-rebound index for the response to the test image not matching the adapter. In each plot, each point represents data from one neuron. The numbers are the counts of points above and below the identity diagonal.

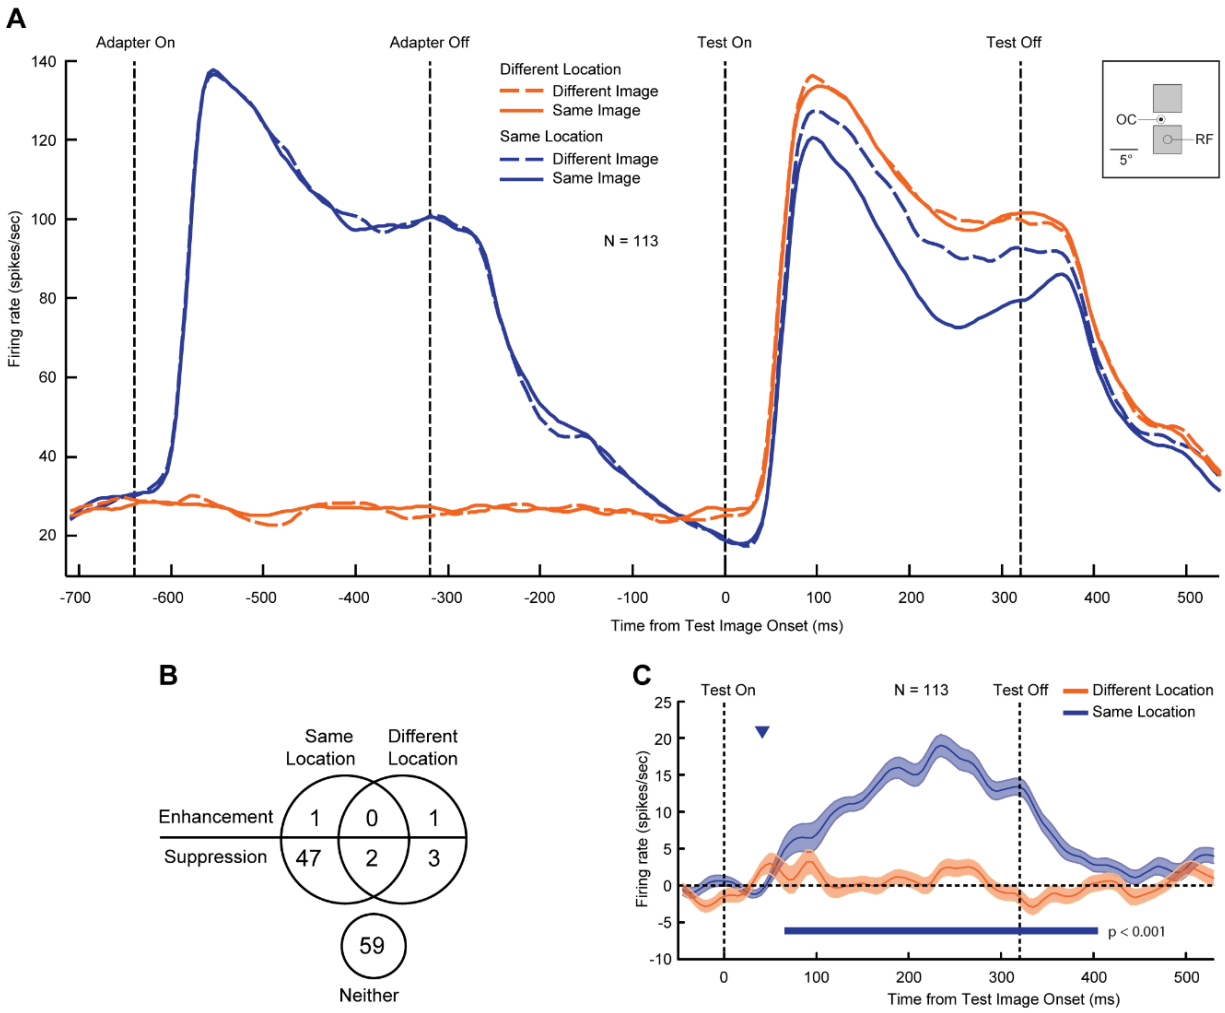

Figure 3.4. Repetition suppression in V2. A. Population mean firing rate under four conditions. **B.** Counts of neurons in which firing rate depended significantly on whether the test matched the adapter or did not. **C.** Mean suppression as a function of time following test image onset. Images in this experiment were placed at locations depicted by gray squares in **Figure 3.1B**. All conventions as in Figure 2.

3.4.5 Impact of repetition on the dynamics of the visual response in V2

The V2 population response to a test image matching the adapter in both identity and location had the form of a trough followed by a rebound (**Figure 3.4A**). The trough-rebound index for a test image matching the adapter in both identity and location (mean = 11.9 spikes/sec) was significantly greater than the corresponding index for the adapter (mean = 0.19

spikes/sec, p = 9.3 E-5, two-sided paired-sample rank-sum test) (**Figure 3.5A**). It was also significantly greater than the corresponding index for a test image presented at the same location as the adapter but with a different identity (mean = -0.11 spikes/sec, p = 8.8 E-5) (**Figure 3.5B**).

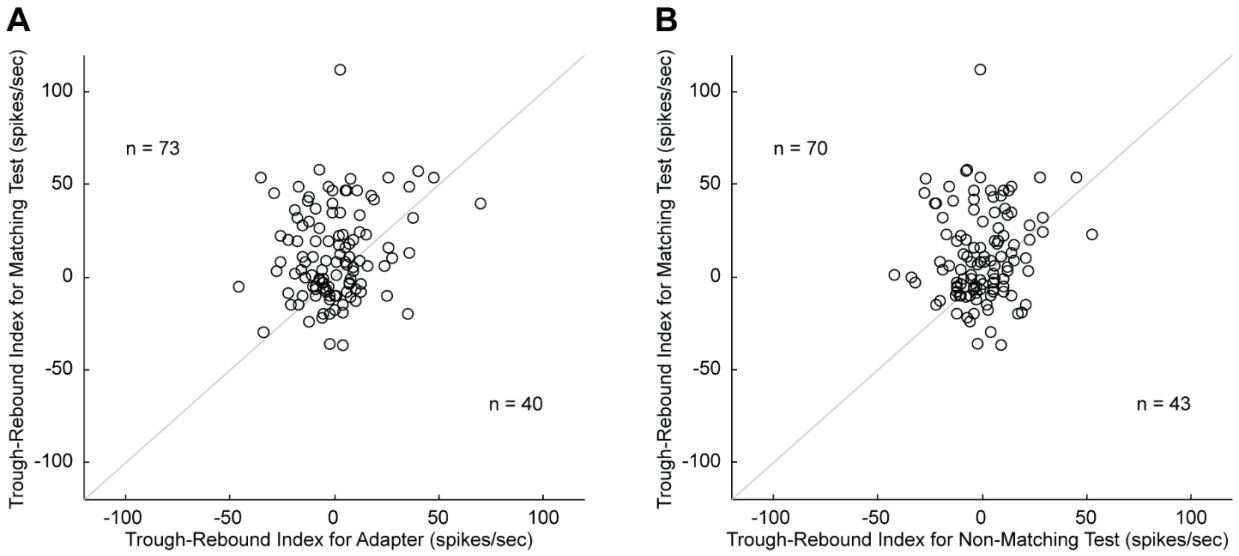

Figure 3.5. The trough-rebound dynamic in V2 was enhanced when the test image matched the adapter. A. The trough-rebound index for the response to the test image matching the adapter is plotted against the trough-rebound index for the response to the adapter itself. **B.** The trough-rebound index for the response to the test image matching the adapter is plotted against the trough-rebound index for the response to the test image not matching the adapter. All conventions as in **Figure 3.3**.

3.4.6 Timing of repetition suppression in V2 and TE

Analyses based on mean population firing rate, as described above, indicate that suppression developed at shorter latency (by 18 ms) and attained half-maximum magnitude earlier (by 6 ms) in V2 than in TE. The temporal offset between suppression in the two areas is evident upon direct comparison of plots representing suppression as a function of time (**Figure 3.6A**). To assess the statistical significance of the between-area timing difference, we measured time to half-maximum suppression in each neuron. Cumulative frequency plots representing the results of this analysis reveal a clear tendency for suppression to attain half-maximum magnitude earlier

in V2 than in TE (**Figure 3.6B**). The difference between the two distributions is statistically significant ($p = 0.0023$, Kolmogorov-Smirnov test). We conclude that suppression developed earlier in V2 than in TE.

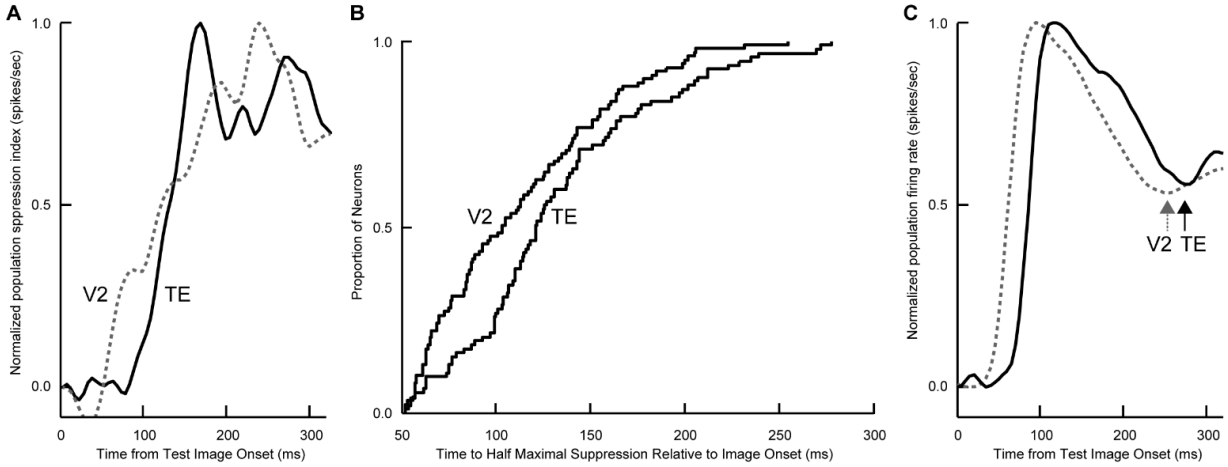

Figure 3.6. Repetition-related effects occur earlier in V2 than in TE. A. Population suppression as a function of time following test onset with adapter and test at the same location. The TE curve is replotted from the blue curve in **Figure 3.2C**. The V2 curve is replotted from the blue curve in **Figure 3.4C**. **B.** Cumulative frequency with respect to time following test onset at which neurons exhibited half-maximal suppression with adapter and test at the same location. The plots represent data from 93 TE neurons and 99 V2 neurons. Ten TE neurons and 14 V2 neurons were excluded due to failure to meet analysis criteria (see Methods). **C.** Trough-rebound dynamic when the test matched the adapter in identity and location. The TE curve is replotted from the solid blue curve in **Figure 3.2A**. The V2 curve is replotted from the solid blue curve in **Figure 3.4A**. Arrows indicate times of trough minima. The curves in **B-C** were normalized to the value at time zero and the maximum within a window spanning 0-350 ms.

To compare the timing of the trough-rebound patterns in V2 and TE, we directly compared plots representing normalized mean population firing rate as a function of time (**Figure 3.6C**). The results make clear that the trough in TE lagged the trough in V2 by a difference approximately equal to the difference in visual response latencies. The time at which the population firing rate attained a minimum at the base of the trough was 253 ms in V2 and 274 ms in TE, giving a displacement of 21 ms. To avoid basing quantitative comparison solely

on the time of the trough minimum, we carried out a further analysis based on the trough as a whole. We characterized the trough in V2 as a vector of firing rates extending in 1 ms bins from 100 to 350 ms post-stimulus-onset. We analyzed the Pearson's correlation between this vector and each vector in a family based on firing in TE from 100+t to 350+t ms where t varied in 1 ms steps from -50 to +50 ms. The correlation attained a maximum (r = 0.98) at t = 18 ms. We conclude that trough and rebound occurred earlier in V2 than in TE.

3.4.7 Location-specific suppression in V2

The preceding analyses do not address the possibility of pure location-specific suppression: a reduction of response strength occurring even when the adapter and the test are different images contingent on their appearing at the same location. In TE, pure location-specific suppression apparently does not occur, as indicated by the fact that the blue and orange dashed curves in **Figure 3.2A** are superimposed. In V2, it does appear to occur (the dashed blue curve is markedly lower than dashed orange curve in **Figure 3.4A**). However interpretation is clouded by the fact that an offset in firing rate was present even before test-image onset in consequence of the fact that firing fell below baseline following the response to an adapter presented in the neuronal receptive field. The lower firing rate following onset of the test image under the same-location condition might have arisen from a reduction in response strength but it might equally well have been the result of superimposing an unaltered response on a reduced baseline. To distinguish between these possibilities, we regressed an index of post-response-onset suppression (25-325 ms following test onset) on an index of pre-response-onset suppression (25 ms before to 25 ms after test onset) across the 113 recorded V2 neurons. The measure of pre-response-onset suppression was $Sup_{base} = (P_d-P_s)/(R_d+R_s)$ where P_d and P_s were the mean pre-response-onset (25 ms before to 25 ms after test onset) firing rates under different-location and same-location

conditions respectively and R_d and R_s were the mean post-response-onset (25-325 ms following test onset) firing rates. The measure of post-stimulus suppression was $Sup_{post} = (R_d-R_s)/(R_d+R_s)$. The best-fit linear function was $Sup_{post} = 0.13*Sup_{pre} + 0.079$ (**Figure 3.7**). Critically, the intercept was significantly greater than zero ($p < 1E-12$) whereas the slope was not ($p = 0.18$). We conclude that neurons in V2 exhibited genuine location-specific suppression, responding to the test at a reduced rate when it appeared at the same location as the adapter despite its possessing a different identity.

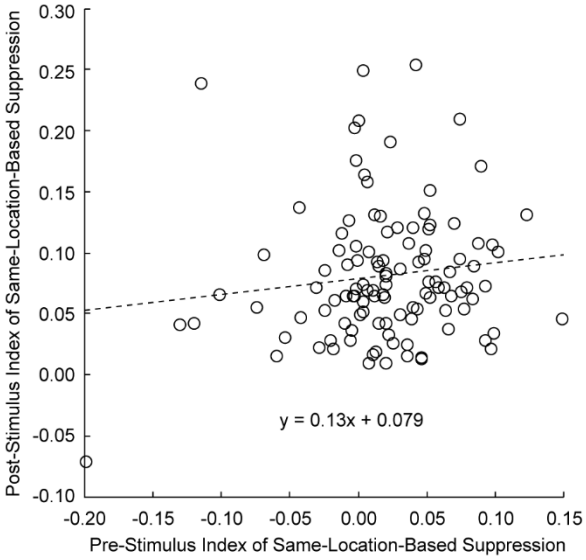

Figure 3.7. V2 neurons respond with reduced strength to a test image matching the adapter in location even when it does not share the adapter's identity. Post-stimulus suppression is plotted against pre-stimulus suppression for 113 neurons. The reduction in response strength is apparent in the tendency for observations to lie above zero. The reduction is not simply a result of the response riding on a suppressed pre-stimulus baseline as indicated by the fact that the intercept of the best-fit line lies above zero.

3.4.8 Repetition suppression in V2 with closely spaced images

In the preceding experiment, the adapter was outside not only the classic receptive field

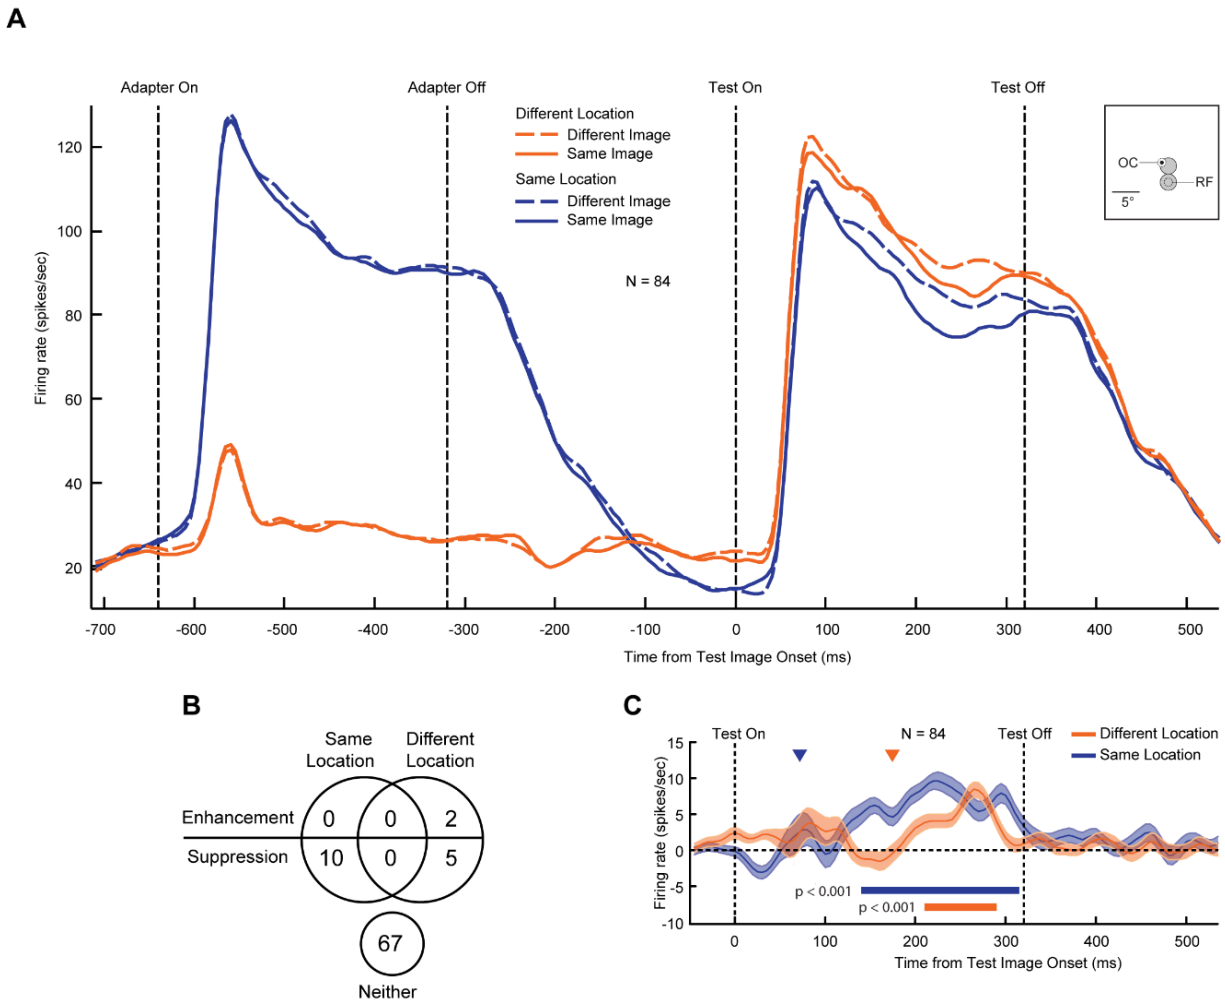

Figure 3.8. Repetition suppression in V2 with small closely spaced images. A. Population mean firing rate under four conditions. **B.** Counts of neurons in which firing rate depended significantly on whether the test matched the adapter or did not. **C.** Mean suppression as a function of time following test image onset. Images in this experiment were placed at locations depicted by gray disks in Figure 1B. All conventions as in **Figure 3.2**.

of the V2 neuron but also the near surround. To test whether spatial transfer of suppression might occur with the adapter in the near surround and the test image in the classic receptive field, we carried out an additional experiment in a late subset of recording sessions. In 84 visually

responsive V2 neurons, we carried out testing with 3° disks cropped from natural images presented either at the receptive field center or 3° above it (**Figure 3.1C**). We now observed image-specific suppression not only when adapter and test image were at the same location (solid blue curve below dashed blue curve in **Figure 3.8A**) but also when they were at different locations (solid orange curve below dashed orange curve in **Figure 3.8A**). Each effect was statistically significant ($p < 0.001$, cluster-based permutation test, **Figure 3.8C**). The counts of effects that achieved statistical significance in individual neurons are presented in **Figure 3.8B**. The frequency of significant effects is lower than in the main experiment. This might be due to the properties of the neurons sampled or to the reduction of stimulus size. The frequency of cases in which significant same-location image-specific suppression occurred nevertheless significantly exceeded the number expected by chance ($p = 0.0080$, chi-squared test, 1 df, with Yates correction). Suppression onset began at 72 ms under the same-location condition and 174 ms under the different-location condition, while the time to half-maximum height was 134 ms under the same-location condition and 236 ms under the different-location condition. Thus, by both measures, different-location suppression was delayed relative to same-location suppression.

3.5 DISCUSSION

Our results in TE are concordant with previous observations. A classic report described repetition suppression specific to the identity of the image as being of roughly equal strength under conditions in which adapter and test image were at the same location or separated by 5° (*8*). Two later reports indicated that suppression is present but reduced in strength when adapter and test image are at different locations; however these utilized a measure of suppression that confounded identity-specific with location-specific effects (*5, 7*). A recent report utilizing an identity-specific measure has established that suppression, although it does generalize across

locations, is indeed reduced in strength and also is prolonged in latency as the distance between adapter and test image increases (*6*). The present results confirm these observations with regard to both strength and timing.

Neuronal visual responses in the context of the repetition suppression paradigm take quite different forms in V2 and TE. In V2 alone, the neuronal response to the test image is reduced when it appears at the same location as the adapter even when its identity is different. In TE alone, image-specific suppression occurs even when the adapter and the test image are presented in a separate visual field quadrants. These discrepancies are surprising in the context of the widely held view that low-order (V1 and V2) and high-order (V4 and TE) ventral stream areas interact through bottom-up and top-down transmission. Bottom-up transmission is of obvious importance as indicated by the fact that surgical ablation of V1 renders neurons in TE visually unresponsive (*32*). Top-down transmission is widely thought to underlie the sensitivity of neurons in low-order areas to sophisticated attributes of visual displays (illusory contours, figure-ground segregation and oddball popout) presumably detected first in high-order areas (*33-41*). Top-down influences are also believed to mediate far-surround suppression (*42-44*) and the modulation of visual response strength by voluntary attention (*35, 38, 39, 41*). The argument that these effects depend on top-down transmission is indirect insofar as it is based on timing: sophisticated signals develop in low-order areas at comparatively long latencies and the latencies in question are longer than in areas of high order. However, direct evidence has come from a few studies demonstrating that activation of V4 influences neuronal activity in low-order cortex (*34*) and that response fluctuations in V4 are causally yoked to response fluctuations in low-order cortex (*33*). Against this backdrop, it is a challenge to explain how the reduction of firing rate in V2 on same-location trials could produce no trace in TE and how different-location image-

specific suppression in TE could produce no trace in V2. Obvious possibilities include the existence of a threshold nonlinearity in multi-synaptic connections between the two areas and precise balancing of excitatory and inhibitory influences.

The neuronal response to a repeated image displayed a trough-rebound dynamic in both TE and V2. This effect has been observed previously in TE but not in V2. Its occurrence in TE, although rarely commented on, is evident in post-stimulus-time histograms from numerous previous studies united by the practice of presenting adapter and test in rapid succession (*5, 7, 45-52*). We and others have observed an apparently identical trough-rebound dynamic in the responses of TE neurons to images rendered familiar by hundreds or thousands of exposures on previous days (*53, 54*). Although typical of responses to repeated and familiar images, the phenomenon is not unique to them: it can be elicited by presenting an image against the backdrop of an already present competing visual display (*55*). The trough-rebound pattern could arise from an increase of either excitatory or inhibitory feedback within a recurrent circuit with resonance at a frequency of around 5 Hz (*53, 55, 56*). The existence of such circuitry may help to explain recent reports to the effect that attention, in the presence of two separate visual targets, oscillates between them at a frequency of 4-6 Hz (*57-60*). The observation on which these reports are based - that the efficiency with which the images are processed varies in counterphase at 4-6 Hz, can be explained by competition within a resonant circuit confined to visual cortex without invoking attention as a top-down process. Propagation of the effect, if it occurs is, in fact, most likely bottom-up. This is suggested by the current observation that a repeated image elicits a trough-rebound response in V2 that leads the corresponding response in TE by around 20 msec. It is also concordant with the observation that 4 Hz oscillations recorded at the cortical surface in areas extending from V1 to TEO show apparent causality primarily in the posterior-to-

anterior direction (*61*). Oscillatory activity may be generated in low-order visual cortex and transmitted to areas of higher order or, alternatively, may arise independently in each area with phase in each area locked to the progressively later time of arrival of visual input.

Repetition suppression in V2 exhibits limited spatial generalization. An adapter in the surround suppresses the late phase of the response to a matching test image presented subsequently in the classical receptive field. This effect requires cross-talk not between two populations of neurons activated by the adapter but rather between two populations selective for the same image only one of which was activated by the adapter. It might be argued that unstable gaze allowed the adapter to stimulate the classical receptive field although the intention was to place it in the surround. This argument is nullified by the observation that the response to the adapter was very weak and of long latency as is typical for a surround stimulus. Moreover, even if the base of the adapting image and the apex of the test image overlapped to a minor degree, the features at base and apex were not consistently similar, as they would have to be for local feature-based cross-adaptation to occur.

The aim of this study was to compare repetition suppression in V2 and TE under conditions comparable to those employed in previous studies of TE involving briefly presented complex natural images. It thus differs from previous investigations of adaptation in low-order visual cortex, which typically have relied on prolonged exposure to simple stimuli such as drifting gratings. It is nevertheless worth considering the degree to which our results parallel the findings of classical adaptation studies (*9, 10, 62*). Adaptation has not previously been characterized in monkey V2 and has been demonstrated in only two studies of its presumed cat homologue, area 18, which demonstrated adaptation using brief presentation of bars (*63*) and prolonged exposure to gratings (*64*). Given the lack of V2 studies and paucity of area 18 studies, we will focus in our

comments on literature concerned with primary as well as secondary visual areas. We consider three key findings of the present study: the ability of natural images to elicit adaptation, the occurrence of adaptation when adapter and test are spatially offset and the impact of adaptation on the visual response dynamic.

Adaptation with natural images. We have found that V2 neurons show repetition suppression for natural images. No previous study of adaptation in low-order visual cortex has made use of such images. Instead, the standard approach has been to use stimuli designed to elicit maximal activity from a subpopulation of neurons, often but not always the neurons under study. The only significant departure from this practice has involved the use plaid displays, which arguably are not optimal for any subpopulation and thus represent a small step toward natural imagery. Exposure to a plaid for tens of seconds may induce adaptation partially specific to the plaid pattern as distinct from the grating components (*65-67*). This encourages the notion that adaptation as observed in the present study is specific to the particular combination of features in the complex image.

Spatial generalization. We have found that image-specific repetition suppression occurs with moderate strength and at long latency even when the adapter and the test are presented at a 3° center-to-center separation so that the adapter lies in the receptive field surround and the test in the receptive-field center. Previous studies assessing the impact of spatial displacement between the adapter and the test have not assessed image-specific adaptation. Rather, they have employed as adapter and test stimuli that are identical to each except with regard to location. One early anecdotal report indicated that an adapter in the receptive field surround reduced the response to a test subsequently presented in the classical receptive field (*68*); however, a more recent effort failed to reveal any such effect (*69*). It now appears that the main impact of presenting an adapter

in the receptive field surround is to reduce the ability of the same stimulus, on subsequent presentation in the surround, to induce surround suppression (*21, 70-73*). Even when adapter and test are placed in separate subfields within the classical receptive field of an area 17 simple cell, there appears to be little or no cross-adaptation (*74*). The failure to detect spatial generalization in previous experiments may reflect a difference between V1 and V2 but could also be ascribed to the fact that a weak effect confined in time may wash out in time-averaged response measures such as employed in the studies under consideration.

Response dynamic. Our observation that a repeated stimulus elicits a response with an enhanced trough-rebound dynamic has no clear parallel in the literature on adaptation in low-order visual cortex. The effect would not have emerged in most studies due to their reliance on time-averaged response measures. However, even in reports presenting population firing rate as a function of time, there is no clear indication of such an effect (*71, 72, 75*). This may indicate the absence in V1 of a trough-rebound pattern or, alternatively, the dependence of the effect on methods specific to our study, including presentation of the adapter and test in rapid succession and the use of images that exceed the background in mean luminance.

3.6 REFERENCES

1. R. Vogels, Sources of adaptation of inferior temporal cortical responses. *Cortex; a journal devoted to the study of the nervous system and behavior*, (2015).
2. G. C. Baylis, E. T. Rolls, Responses of neurons in the inferior temporal cortex in short term and serial recognition memory tasks. *Exp Brain Res* **65**, 614-622 (1987).
3. M. W. Brown, F. A. Wilson, I. P. Riches, Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res* **409**, 158-162 (1987).
4. K. S. Saleem, N. K. Logothetis, *Atlas of the Rhesus Monkey Brain in Stereotaxic Coordinates*. (Academic Press, Amsterdam, 2007).
5. W. De Baene, R. Vogels, Effects of Adaptation on the Stimulus Selectivity of Macaque Inferior Temporal Spiking Activity and Local Field Potentials. *Cerebral Cortex* **20**, 2145-2165 (2010).
6. F. Fabbrini, R. Vogels, Within- and between-hemifield generalization of repetition suppression in inferior temporal cortex. *J Neurophysiol* **125**, 120-139 (2021).

7. P. Kuravi, R. Vogels, Effect of adapter duration on repetition suppression in inferior temporal cortex. *Sci Rep* **7**, 3162 (2017).
8. A. Lueschow, E. K. Miller, R. Desimone, Inferior temporal mechanisms for invariant object recognition. *Cerebral cortex (New York, N.Y. : 1991)* **4**, 523-531 (1994).
9. A. Kohn, Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. *Journal of Neurophysiology* **97**, 3155-3164 (2007).
10. S. G. Solomon, A. Kohn, Moving sensory adaptation beyond suppressive effects in single neurons. *Current biology : CB* **24**, R1012-1022 (2014).
11. K. Grill-Spector, R. Henson, A. Martin, Repetition and the brain: neural models of stimulus-specific effects. *Trends in cognitive sciences* **10**, 14-23 (2006).
12. J. Larsson, S. G. Solomon, A. Kohn, fMRI adaptation revisited. *Cortex; a journal devoted to the study of the nervous system and behavior* **80**, 154-160 (2016).
13. G. M. Boynton, E. M. Finney, Orientation-specific adaptation in human visual cortex. *J Neurosci* **23**, 8781-8787 (2003).
14. F. Fang, S. O. Murray, D. Kersten, S. He, Orientation-tuned FMRI adaptation in human visual cortex. *J Neurophysiol* **94**, 4188-4195 (2005).
15. J. Larsson, S. J. Harrison, Spatial specificity and inheritance of adaptation in human visual cortex. *J Neurophysiol* **114**, 1211-1226 (2015).
16. J. Larsson, D. J. Heeger, M. S. Landy, Orientation selectivity of motion-boundary responses in human visual cortex. *J Neurophysiol* **104**, 2940-2950 (2010).
17. J. Larsson, M. S. Landy, D. J. Heeger, Orientation-selective adaptation to first- and second-order patterns in human visual cortex. *J Neurophysiol* **95**, 862-881 (2006).
18. H. A. Lee, S. H. Lee, Hierarchy of direction-tuned motion adaptation in human visual cortex. *J Neurophysiol* **107**, 2163-2184 (2012).
19. L. Montaser-Kouhsari, M. S. Landy, D. J. Heeger, J. Larsson, Orientation-selective adaptation to illusory contours in human visual cortex. *J Neurosci* **27**, 2186-2195 (2007).
20. S. O. Murray, C. A. Olman, D. Kersten, Spatially specific FMRI repetition effects in human visual cortex. *J Neurophysiol* **95**, 2439-2445 (2006).
21. S. C. Wissig, A. Kohn, The influence of surround suppression on adaptation effects in primary visual cortex. *J Neurophysiol* **107**, 3370-3384 (2012).
22. K. Grill-Spector *et al.*, Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* **24**, 187-203 (1999).
23. K. Grill-Spector, R. Malach, fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychol (Amst)* **107**, 293-321 (2001).
24. Z. Kourtzi, N. Kanwisher, Representation of perceived object shape by the human lateral occipital complex. *Science* **293**, 1506-1509 (2001).
25. Z. Kourtzi, N. Kanwisher, Cortical regions involved in perceiving object shape. *J Neurosci* **20**, 3310-3318 (2000).
26. S. O. Murray, E. Wojciulik, Attention increases neural selectivity in the human lateral occipital complex. *Nat Neurosci* **7**, 70-74 (2004).
27. H. Sawamura, S. Georgieva, R. Vogels, W. Vanduffel, G. A. Orban, Using functional magnetic resonance imaging to assess adaptation and size invariance of shape processing by humans and monkeys. *J Neurosci* **25**, 4294-4306 (2005).
28. M. Fritsche, S. J. D. Lawrence, F. P. de Lange, Temporal tuning of repetition suppression across the visual cortex. *J Neurophysiol* **123**, 224-233 (2020).
29. v. G. Bonin, P. Bailey, *The Neocortex of Macaca Mulatta*. (University of Illinois Press, Urbana, 1947).
30. W. A. Suzuki, Y. Naya, The Perirhinal Cortex. *Annual Review of Neuroscience* **37**, 39-53 (2014).

31. E. Maris, R. Oostenveld, Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* **164**, 177-190 (2007).
32. C. E. Rocha-Miranda, D. B. Bender, C. G. Gross, M. Mishkin, Visual activation of neurons in inferotemporal cortex depends on striate cortex and forebrain commissures. *J Neurophysiol* **38**, 475-491 (1975).
33. M. Chen *et al.*, Incremental integration of global contours through interplay between visual cortical areas. *Neuron* **82**, 682-694 (2014).
34. P. C. Klink, B. Dagnino, M. A. Gariel-Mathis, P. R. Roelfsema, Distinct Feedforward and Feedback Effects of Microstimulation in Visual Cortex Reveal Neural Mechanisms of Texture Segregation. *Neuron* **95**, 209-220.e203 (2017).
35. V. A. Lamme, P. R. Roelfsema, The distinct modes of vision offered by feedforward and recurrent processing. *Trends in neurosciences* **23**, 571-579 (2000).
36. T. S. Lee, D. Mumford, R. Romero, V. A. Lamme, The role of the primary visual cortex in higher level vision. *Vision research* **38**, 2429-2454 (1998).
37. T. S. Lee, M. Nguyen, Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences of the United States of America* **98**, 1907-1911 (2001).
38. T. S. Lee, C. F. Yang, R. D. Romero, D. Mumford, Neural activity in early visual cortex reflects behavioral experience and higher-order perceptual saliency. *Nat Neurosci* **5**, 589-597 (2002).
39. J. Poort *et al.*, The role of attention in figure-ground segregation in areas V1 and V4 of the visual cortex. *Neuron* **75**, 143-156 (2012).
40. J. Poort, M. W. Self, B. van Vugt, H. Malkki, P. R. Roelfsema, Texture Segregation Causes Early Figure Enhancement and Later Ground Suppression in Areas V1 and V4 of Visual Cortex. *Cerebral cortex (New York, N.Y. : 1991)* **26**, 3964-3976 (2016).
41. P. R. Roelfsema, M. Tolboom, P. S. Khayat, Different processing phases for features, figures, and selective attention in the primary visual cortex. *Neuron* **56**, 785-792 (2007).
42. A. Angelucci *et al.*, Circuits and Mechanisms for Surround Modulation in Visual Cortex. *Annu Rev Neurosci* **40**, 425-451 (2017).
43. S. Shushruth, J. M. Ichida, J. B. Levitt, A. Angelucci, Comparison of spatial summation properties of neurons in macaque V1 and V2. *J Neurophysiol* **102**, 2069-2083 (2009).
44. S. Shushruth *et al.*, Different orientation tuning of near- and far-surround suppression in macaque primary visual cortex mirrors their tuning in human perception. *J Neurosci* **33**, 106-119 (2013).
45. F. Fabbrini *et al.*, Probing the Mechanisms of Repetition Suppression in Inferior Temporal Cortex with Optogenetics. *Current biology : CB* **29**, 1988-1998.e1984 (2019).
46. D. A. Kaliukhovich, R. Vogels, Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral cortex (New York, N.Y. : 1991)* **21**, 1547-1558 (2011).
47. D. A. Kaliukhovich, R. Vogels, Stimulus repetition affects both strength and synchrony of macaque inferior temporal cortical activity. *J Neurophysiol* **107**, 3509-3527 (2012).
48. D. A. Kaliukhovich, R. Vogels, Divisive Normalization Predicts Adaptation-Induced Response Changes in Macaque Inferior Temporal Cortex. *J Neurosci* **36**, 6116-6128 (2016).
49. Y. Liu, S. O. Murray, B. Jagadeesh, Time Course and Stimulus Dependence of Repetition-Induced Response Suppression in Inferotemporal Cortex. *Journal of Neurophysiology* **101**, 418-436 (2009).
50. H. Sawamura, G. A. Orban, R. Vogels, Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. *Neuron* **49**, 307-318 (2006).

51. K. Vinken, H. P. Op de Beeck, R. Vogels, Face Repetition Probability Does Not Affect Repetition Suppression in Macaque Inferotemporal Cortex. *J Neurosci* **38**, 7492-7504 (2018).
52. N. P. Williams, C. R. Olson, Contribution of Individual Features to Repetition Suppression in Macaque Inferotemporal Cortex. (In preparation).
53. S. Lim, Mechanisms underlying sharpening of visual response dynamics with familiarity. *Elife* **8**, (2019).
54. T. Meyer, C. Walker, R. Y. Cho, C. R. Olson, Image familiarization sharpens response dynamics of neurons in inferotemporal cortex. *Nat Neurosci* **17**, 1388-1394 (2014).
55. J. E. Rollenhagen, C. R. Olson, Low-Frequency Oscillations Arising From Competitive Interactions Between Visual Stimuli in Macaque Inferotemporal Cortex. *Journal of Neurophysiology* **94**, 3368-3387 (2005).
56. S. Moldakarimov, J. E. Rollenhagen, C. R. Olson, C. C. Chow, Competitive Dynamics in Cortical Responses to Visual Stimuli. *Journal of Neurophysiology* **94**, 3388-3396 (2005).
57. L. Dugué, P. Marque, R. VanRullen, Theta oscillations modulate attentional search performance periodically. *J Cogn Neurosci* **27**, 945-958 (2015).
58. I. C. Fiebelkorn, Y. B. Saalmann, S. Kastner, Rhythmic sampling within and between objects despite sustained attention at a cued location. *Current biology : CB* **23**, 2553-2558 (2013).
59. A. N. Landau, P. Fries, Attention samples stimuli rhythmically. *Current biology : CB* **22**, 1000-1004 (2012).
60. A. N. Landau, H. M. Schreyer, S. van Pelt, P. Fries, Distributed Attention Is Implemented through Theta-Rhythmic Gamma Modulation. *Current biology : CB* **25**, 2332-2337 (2015).
61. G. Spyropoulos, C. A. Bosman, P. Fries, A theta rhythm in macaque visual cortex and its attentional modulation. *Proceedings of the National Academy of Sciences of the United States of America* **115**, E5614-e5623 (2018).
62. M. Snow, R. Coen-Cagli, O. Schwartz, Adaptation in the visual cortex: a case for probing neuronal populations with natural stimuli. *F1000Res* **6**, 1246 (2017).
63. S. B. Nelson, Temporal interactions in the cat visual system. I. Orientation-selective suppression in the visual cortex. *J Neurosci* **11**, 344-356 (1991).
64. N. A. Crowder *et al.*, Relationship between contrast adaptation and orientation tuning in V1 and V2 of cat visual cortex. *J Neurophysiol* **95**, 271-283 (2006).
65. A. Aschner, S. G. Solomon, M. S. Landy, D. J. Heeger, A. Kohn, Temporal Contingencies Determine Whether Adaptation Strengthens or Weakens Normalization. *J Neurosci* **38**, 10129-10142 (2018).
66. M. Carandini, H. B. Barlow, L. P. O'Keefe, A. B. Poirson, J. A. Movshon, Adaptation to contingencies in macaque primary visual cortex. *Philos Trans R Soc Lond B Biol Sci* **352**, 1149-1154 (1997).
67. M. Carandini, J. A. Movshon, D. Ferster, Pattern adaptation and cross-orientation interactions in the primary visual cortex. *Neuropharmacology* **37**, 501-511 (1998).
68. L. Maffei, A. Fiorentini, S. Bisti, Neural correlate of perceptual adaptation to gratings. *Science* **182**, 1036-1038 (1973).
69. N. J. Priebe, M. M. Churchland, S. G. Lisberger, Constraints on the source of short-term motion adaptation in macaque area MT. I. the role of input and intrinsic mechanisms. *J Neurophysiol* **88**, 354-369 (2002).
70. S. Durand, T. C. Freeman, M. Carandini, Temporal properties of surround suppression in cat primary visual cortex. *Vis Neurosci* **24**, 679-690 (2007).
71. C. A. Patterson, J. Duijnhouwer, S. C. Wissig, B. Krekelberg, A. Kohn, Similar adaptation effects in primary visual cortex and area MT of the macaque monkey under matched stimulus conditions. *J Neurophysiol* **111**, 1203-1213 (2014).

72. C. A. Patterson, S. C. Wissig, A. Kohn, Distinct effects of brief and prolonged adaptation on orientation tuning in primary visual cortex. *J Neurosci* **33**, 532-543 (2013).
73. B. S. Webb, N. T. Dhruv, S. G. Solomon, C. Tailby, P. Lennie, Early and late mechanisms of surround suppression in striate cortex of macaque. *J Neurosci* **25**, 11666-11675 (2005).
74. S. G. Marlin, R. M. Douglas, M. S. Cynader, Position-specific adaptation in simple cell receptive fields of the cat striate cortex. *J Neurophysiol* **66**, 1769-1784 (1991).
75. J. A. Westerberg, M. A. Cox, K. Dougherty, A. Maier, V1 microcircuit dynamics: altered signal propagation suggests intracortical origins for adaptation in response to visual repetition. *J Neurophysiol* **121**, 1938-1952 (2019).

3.7 Individual monkey figures

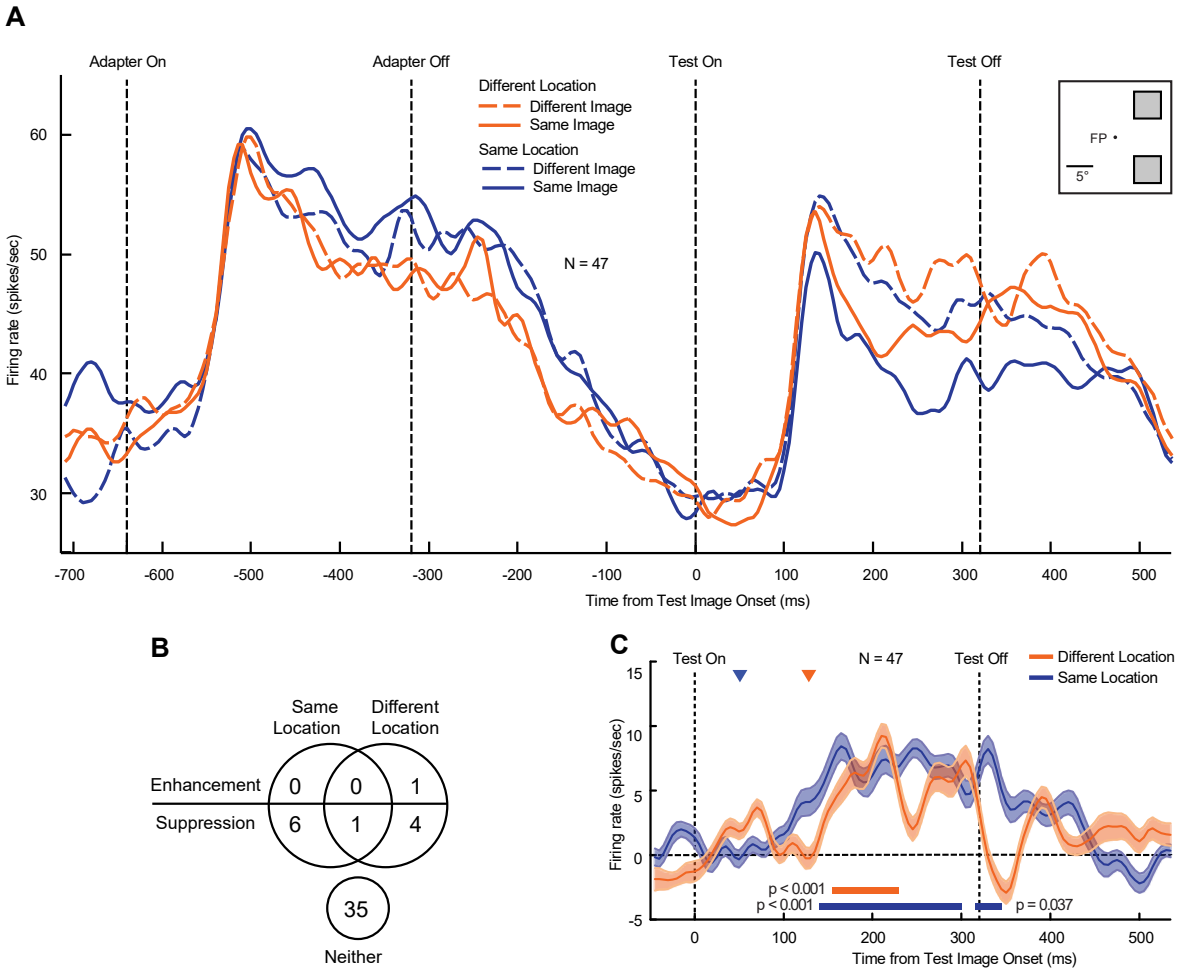

Figure S3.1. Repetition Suppression in TE in Monkey 1. All conventions as in **Figure 3.2**.

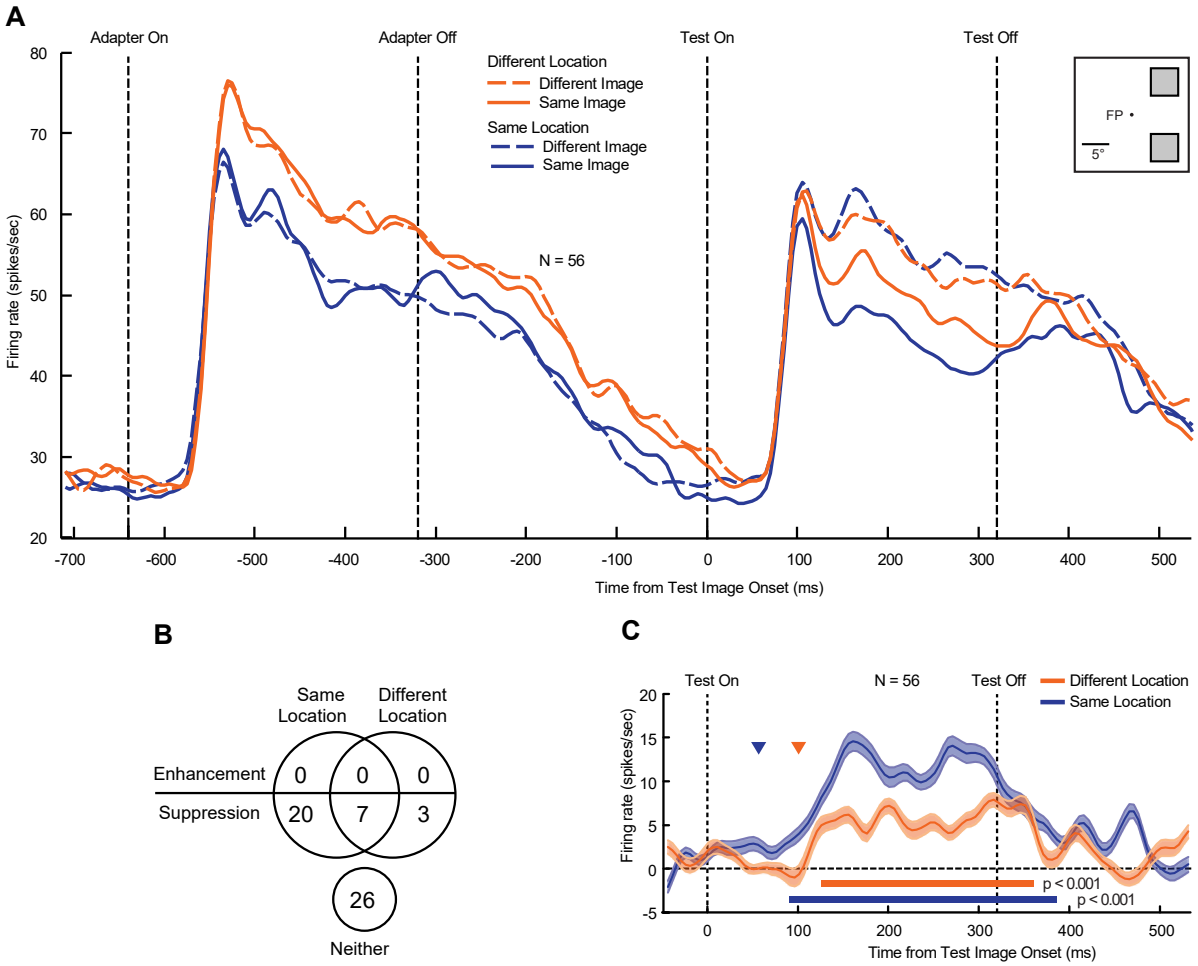

Figure S3.2. Repetition Suppression in TE in Monkey 2. All conventions as in **Figure 3.2**.

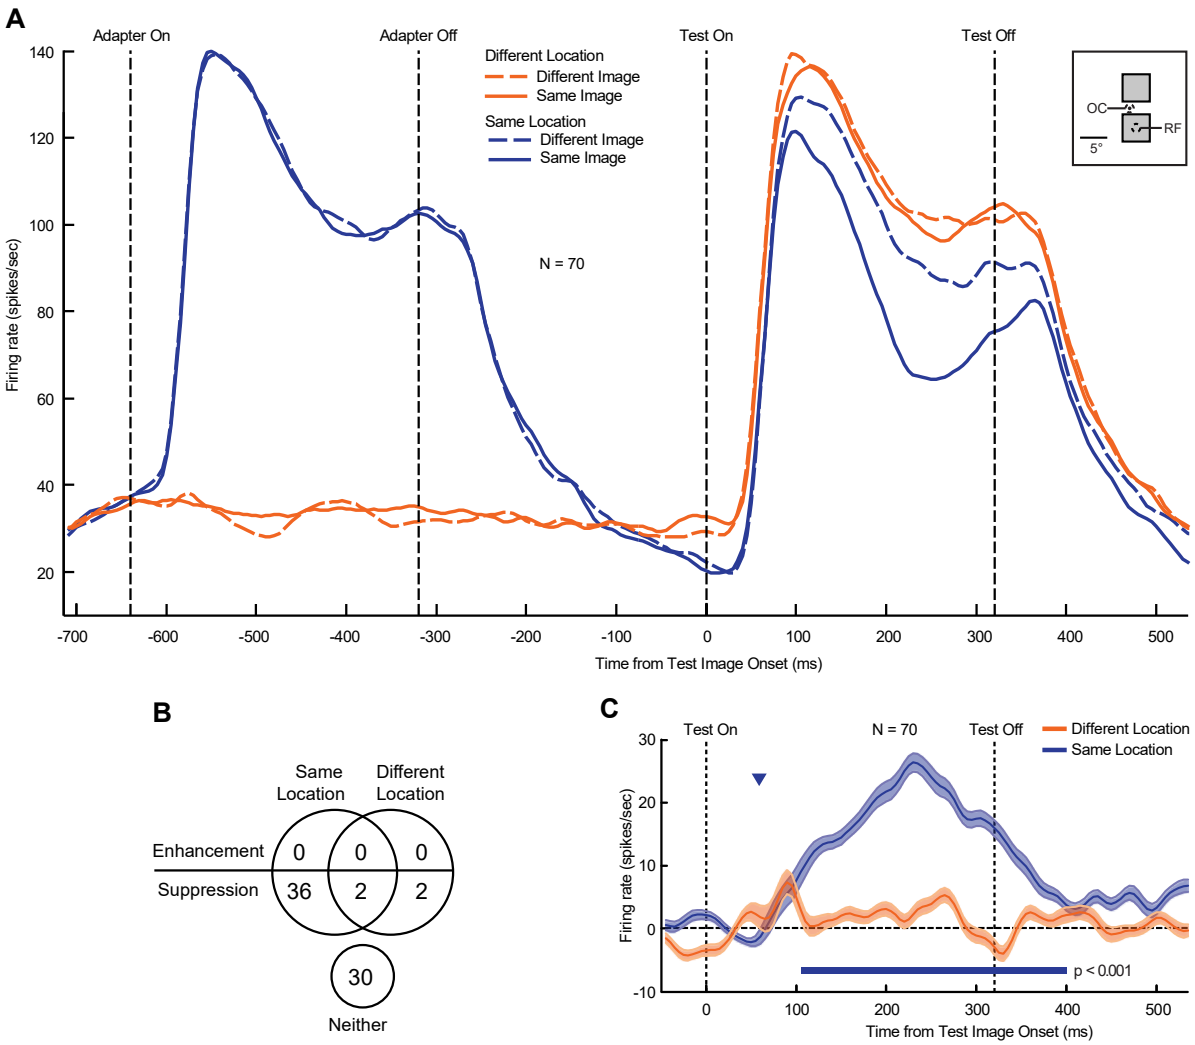

Figure S3.3. Repetition Suppression in V2 in Monkey 1. All conventions as in **Figure 3.4**.

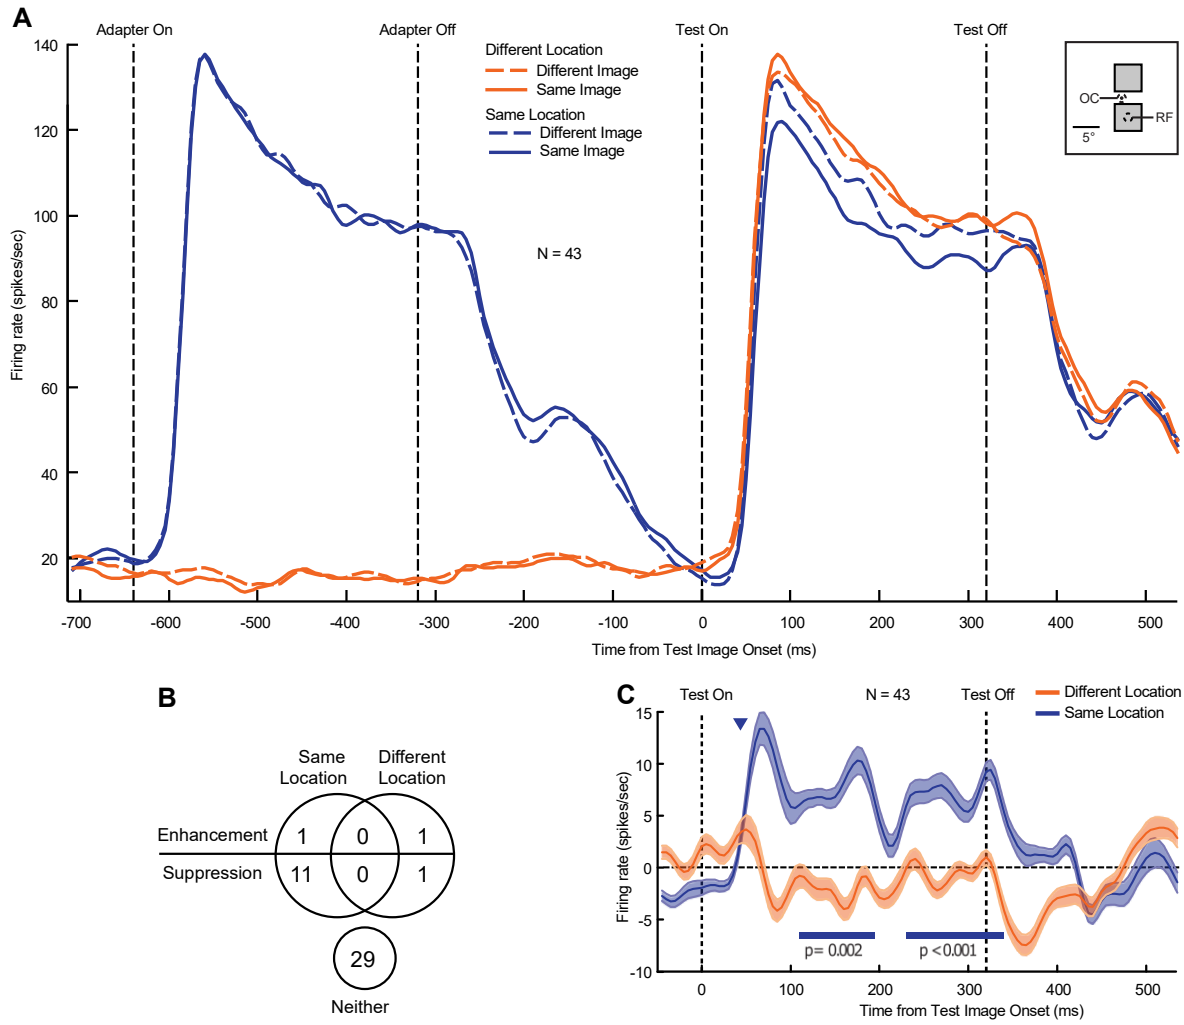

Figure S3.4. Repetition Suppression in V2 in Monkey 3. All conventions as in **Figure 3.4**.

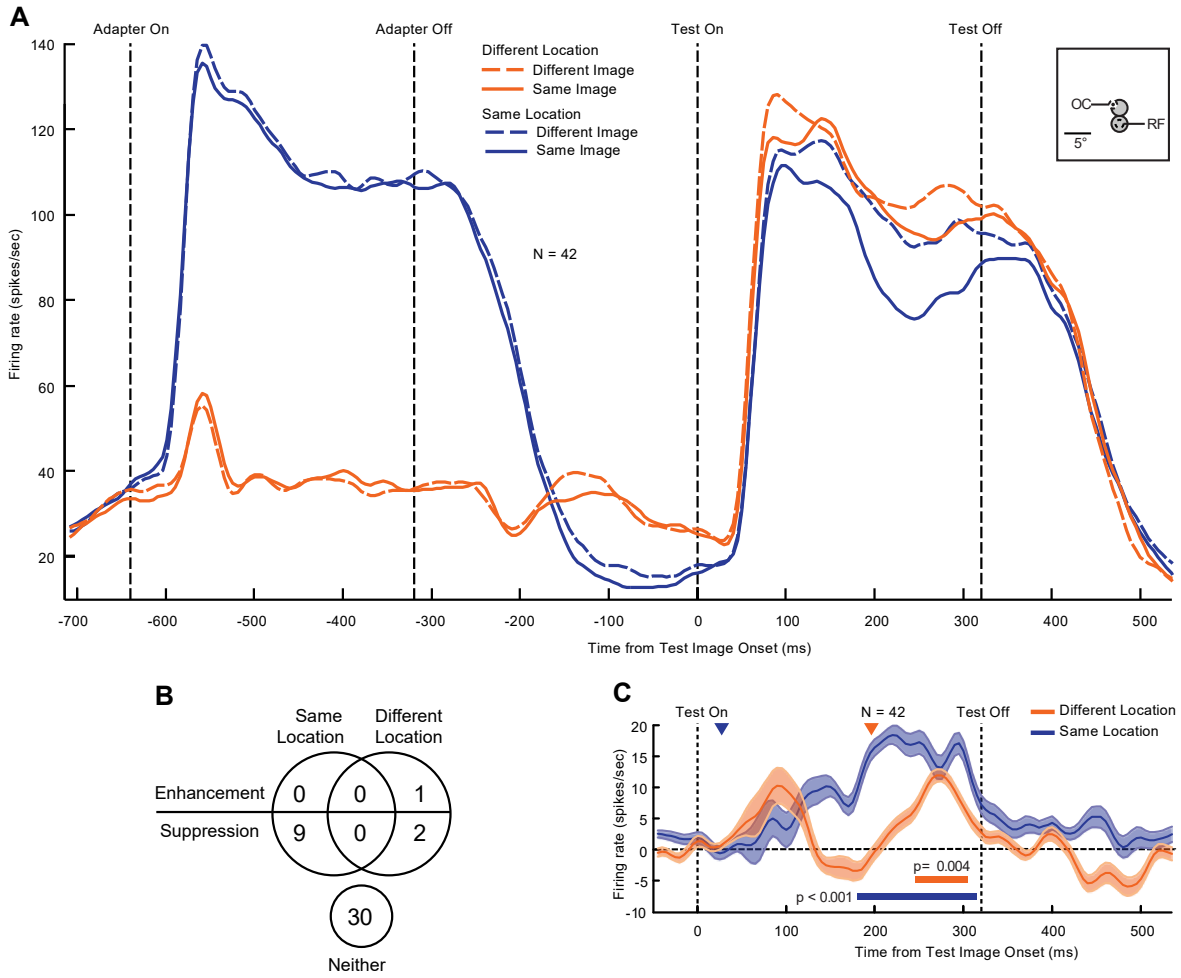

Figure S3.5. Repetition suppression in V2 with small closely spaced images in Monkey 1. All conventions as in **Figure 3.8**.

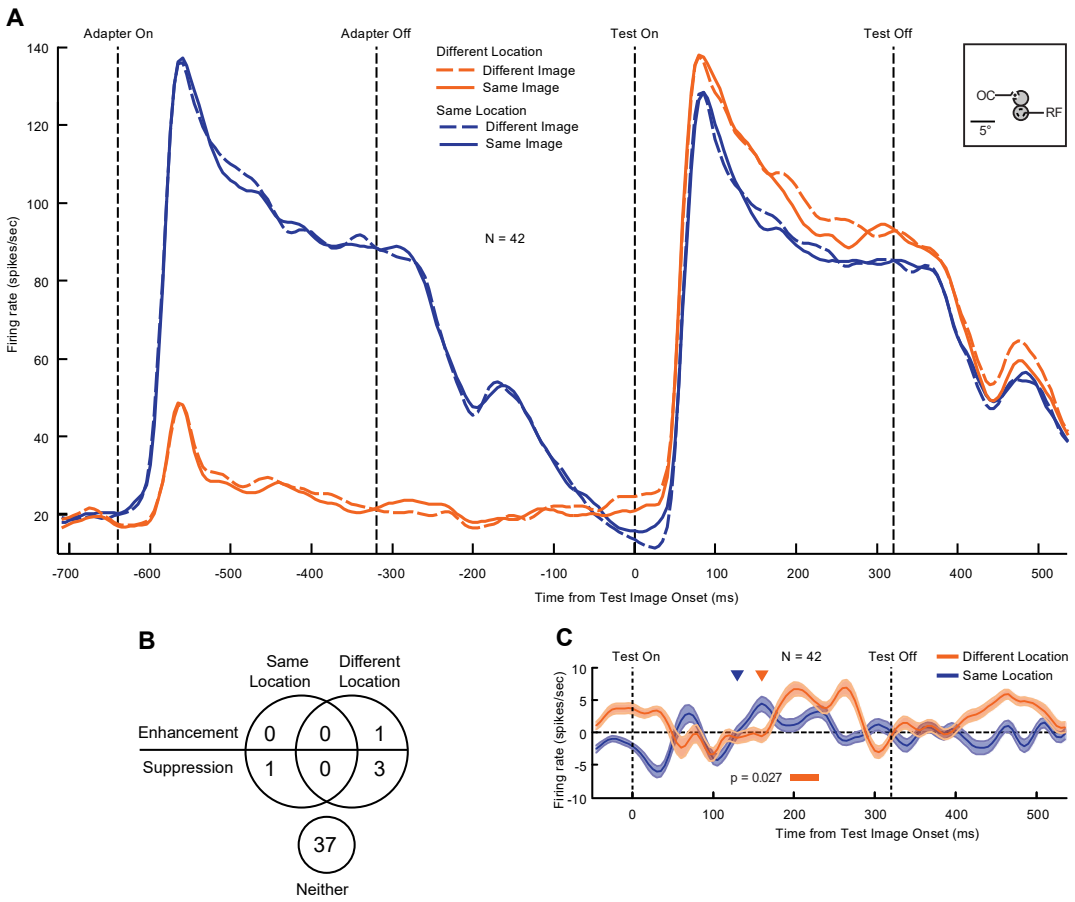

Figure S3.6. Repetition suppression in V2 with small closely spaced images in Monkey 3. All conventions as in **Figure 3.8**.

CHAPTER IV

SEPARATE CONTRIBUTIONS OF IMAGE CONTENT IN THE CENTER AND IN THE SURROUND TO REPETITION SUPPRESSION IN MACAQUE VISUAL AREA V2

4.1 ABSTRACT

Monkey inferotemporal cortex neurons in area TE respond with declining strength to repeated presentations of a large and complex natural image. This phenomenon – repetition suppression – has often been assumed to arise at the level of TE because TE neurons possess the large receptive fields and sophisticated stimulus selectivity necessary for recognizing the image as a repetition. However, we recently discovered that neurons in V2 exhibit repetition suppression under conditions identical to those employed in studies of TE. This raises the question: How do V2 neurons, with classical receptive fields encompassing only a small fraction of the image, recognize it as a repetition? One possibility is that they are sensitive to repetition of image content not only in the classical receptive field but also in the receptive field surround. To assess this possibility, we monitored neuronal responses to sequential displays in which we controlled independently the repetition of elements in the center and in the surround of a large natural image. Each stimulus consisted of a 3° disk centered on the classical receptive field and an adjoining annulus with an outer diameter of 8°. The disks and annuli were taken from different natural scenes. The display on each trial consisted of an adapter (320 ms), a delay (320 ms) and a test (320 ms). Across trials, we independently varied the relation of the adapter to the test with respect to the identity of the disk component (same or different) and the annulus component (same or different). We found that repetition of the central disk was sufficient to produce repetition suppression but that suppression was enhanced by simultaneous repetition of the annulus. Suppression arising from repetition of the annulus occurred in a relatively late phase

of the response, in accordance with the idea that it might have been mediated by horizontal connections. The contribution of the surround to repetition suppression in these experiments cannot be explained in terms of adaptation in classical grating-based adaptation experiments. Repetition of surround content in the current paradigm leads to reduced response strength whereas repetition of surround content in grating-based adaptation experiments leads to enhanced response strength.

4.2 INTRODUCTION

Neurons in area TE of macaque inferotemporal cortex respond with declining strength to repeated presentations of a large and complex natural image. This phenomenon, termed repetition suppression, has often been assumed to arise at the level of TE in inferotemporal cortex because inferotemporal neurons possess the large receptive fields and sophisticated stimulus selectivity necessary for recognizing the image as a repetition. We have recently shown that V2 neurons exhibit image-specific repetition suppression under the same conditions used in studies of TE. That is, to large and complex natural objects. TE neurons typically have large receptive fields (*1, 2*), often spanning multiple quadrants of visual space, and they are sharply tuned to the complex stimulus properties of natural objects with the consequence that they response selectively to only a few objects out of a reasonably large set of even very similar objects (*3, 4*). V2 neurons have smaller receptive fields typically restricted to within a single quadrant of the visual field (*5*) and they appear to respond with reasonable strength to the majority of natural images, unlike TE neurons. In addition to their smaller receptive fields, V2 neurons are thought to primarily represent visual features such as orientation, color and other low-level stimulus properties without particular selectivity to the unique object identities resulting from combinations of these features (*5, 6*). This raises the question: How do V2

neurons, with classical receptive fields encompassing only a small fraction of the image, recognize it as a repetition? One possibility is that they are sensitive to repetition of image content not only in the classical receptive field but also in the extraclassical receptive field surround.

Most studies of the contribution of image content outside the classical receptive field of neurons in low order visual cortex have focused on V1 with the basic finding that image content outside the classical excitatory receptive field has a suppressive effect on the response of the neuron (*7-13*). These findings have also been extended to V2, with studies using the same stimulus parameters used in V1 showing that the responses of V2 neurons to content in the surround follow the same basic suppressive center-surround interactions as what was found in V1 (*14*). These studies have also mainly focused on the initial response of the neuron and we know of only one study which has tested the effect of stimulus repetition using large stimuli which extend beyond the classical receptive field, recording from V1 neurons in anesthetized macaques (*15*). This study found, in general, strong suppression when the classical receptive field was adapted and weak suppression and in many cases enhancement when the classical receptive field as well as the surround was adapted, and that this adaptation depended on the orientation of the adapter relative to the preferred orientation of the neuron, such that suppression was strongest at the preferred orientation and enhancement was strongest at the orthogonal orientation. This is consistent with the idea that adaptation weakens the response of the adapted neuron and that if stimulating the suppressive surround is inhibiting the response of the neuron, then adapting the surround will lessen the inhibition, in the case of this study leading to an overall enhancement of the response. All of these studies have only used disks and annuli of grating patches or drifting gratings as stimuli, and so the effect of repetition of content in the

surround for natural images in V2 neurons of awake animals remains very much an open question.

Until recently, studies of low-order visual areas such as V1 and V2 have for the most part been confined to using simple stimuli designed to elicit maximal responses from individual neurons, and have manipulated stimulus properties such as contrast, color, orientation and spatial frequency, which are known to modulate the responses of neurons in these areas in a predictable way (*6, 16, 17*). More recently, it has been shown that many V2 neurons, unlike V1 neurons, can be driven more strongly by naturalistic texture patches than by their preferred grating patch (*18, 19*). These findings, as well as our own observations in V2 using natural images, suggests that V2 neurons may be more sensitive to complex image properties than previously appreciated. However, the fact that higher firing rates can be elicited by texture patches or natural images and thus that they may be more preferred stimuli for V2 neurons than grating patches does not make clear what the tuning properties of V2 neurons are for particular aspects of complex images nor how sharply tuned V2 neurons are for features of natural images. These findings suggest that the responses of neurons in low order areas such as V2 to natural images and in particular with respect to natural image content in the surround, which is ubiquitous under natural viewing conditions, are worthy of further study.

To determine the contribution of content in the surround of a large complex image to repetition suppression in V2, we monitored neuronal responses to sequential displays in which we controlled independently the repetition of natural scene content in the central disk and in an adjoining annulus comprising content in the surround of a compound image constructed by adjoining the two components. We found that repetition of the central disk was sufficient to produce repetition suppression but that suppression was enhanced by simultaneous repetition of

the annulus. Suppression arising from repetition of the annulus occurred in a relatively late phase of the response, beginning over 150 ms after stimulus onset, in accordance with the idea that it might have been mediated by horizontal connections. The contribution of content in the surround of the image to repetition suppression under these conditions cannot be explained in terms of adaptation in classical grating-based experiments. Repetition of surround content in the current paradigm led to reduced response strength whereas repetition of surround content in grating-based adaptation experiments led to an enhancement of the response. We also recorded responses on trials in which either only the disk or only the annulus was presented and could either match or not match the preceding disk or annulus. On these trials, we saw evidence for repetition suppression only when the disk was presented, but saw no evidence for repetition suppression when only the annulus was presented, making the finding of enhanced suppression when both the disk and the annulus are repeated even more surprising.

4.3 MATERIALS AND METHODS

4.3.1 Subjects

The experimental subjects were two adult rhesus macaques (*macaca mulatta*) here designated M1 (male; 12 kg; laboratory designation Ol) and M2 (male; 13 kg; laboratory designation Rs). All procedures were in accordance with guidelines set forth by the United States Public Health Service Guide for the Care and Use of Laboratory Animals and were approved by the Carnegie Mellon University IACUC.

4.3.2 Receptive Field Mapping

At the outset of each session, we located the center of the receptive field of the neuron under study by means of the following steps. We first crudely mapped it out with manually controlled stimuli. We then carried out automatic assessment of horizontal and vertical spatial

selectivity within a 3° by 3° grid centered on the field's estimated location, using as stimuli bars with a width of 0.3° and a length of 8°. Ten horizontal bars centered horizontally on the grid spanned it in 0.3° vertical steps. Ten vertical bars centered vertically on the grid spanned it in 0.3° horizontal steps. Each bar contained high contrast content as if it had been cut from a square-wave grating with spatial frequency of 3.3 cycles/degree and an orientation of +45° or -45° relative to vertical, with orientation determined by the preference of the neuron. While the monkey maintained central fixation, the bars were presented for 300 ms each, with an interstimulus interval of 200 ms, in random sequence until each had appeared three times. Tuning for location with respect to each axis was clear from inspection of the resulting post-stimulus-time histograms. In subsequent tests, we centered stimuli at the X coordinate of the vertical bar eliciting the strongest response and the Y coordinate of the horizontal bar eliciting the strongest response.

Once the X and Y coordinates of the receptive field of the recorded neuron were determined, two additional tests of stimulus selectivity were conducted. First, we tested orientation selectivity using a set of 5 orientations at progressive 30° offsets, starting from 0° vertical. These stimuli were full contrast square wave grating patches of 3.3 cycles/degree and 6° in diameter. They were presented in random order until each orientation had been presented 3 times. Tuning for orientation was determined by inspection of online rasters and histograms, with the orientation giving the highest average response being selected as the preferred orientation. Following this test of orientation selectivity, we presented grating disks and annuli of the same contrast and spatial frequency, at the neuron's preferred orientation, which varied in outer diameter (for the disks) or inner diameter (for the annuli) from 1° to 6° in 1° increments,

with the outer diameter of the annuli being set to 8 degrees, to further map the response field of the neuron.

4.3.3 Task

Each trial began with presentation of a fixation spot centered at the midpoint of an LCD monitor with a 60 Hz refresh rate at a viewing distance of 32 cm. Once the monkey had attained fixation, the following displays appeared in sequence: fixation spot (320 ms), adapter stimulus (320 ms), fixation spot (320 ms), test stimulus (320 ms), fixation spot (320 ms). The monkey was required to maintain gaze on the fixation spot throughout the trial, within a window subtending 1.4-2.1° throughout the trial. Gaze was monitored with an ISCAN video-based eye tracking system. Upon successful completion of a trial, juice reward was delivered. Any fixation break terminated the trial and triggered onset of a checkerboard display which remained visible for two seconds or until the monkey had fixated it for a cumulative duration of 300 ms, whichever came first. Behavior was monitored and stimulus presentation and reward delivery were controlled by a PC running NIMH Cortex.

4.3.4 Stimuli

The fixation spot was a white disk 0.3° in diameter. The stimuli were disks of 3° diameter and annuli of 3° inner diameter and 8° outer diameter, each cut from one of 8 unrelated natural scenes. On different trials, either only disks were used, only annuli were used, or disks and annuli appeared together to form a single scene patch 8° in diameter composed of a chimera of disk and annulus together. Disks and annuli that appeared together were never taken from the same scene. All stimuli were centered on the receptive field of the recorded neuron. The average horizontal offset from fixation was ~1.8° with a range of ~0.4° to 4.4° and the average vertical offset was ~5.0° with a range of ~3.6° to 5.9°. A 1.6° diameter invisible mask surrounded the

fixation spot, preventing the stimulus from impinging on fixation. Before beginning any recording session, we assessed neuronal image selectivity by monitoring responses elicited by the set of 8 natural scenes. We selected for use in a given session four scenes that elicited the strongest responses as judged by inspection of online raster and histogram displays. Arbitrarily numbering the selected scenes from 1 to 4, we cut disks 1 and 2 from scenes 1 and 2 and cut annuli 1 and 2 from scenes 3 and 4.

4.3.5 Session structure

From the two disks and two annuli, 16 adapt-test sequences could be created on trials where the disk and annulus appeared together. Among these were four cases in which the disk alone repeated, four in which the annulus alone repeated, four in which both components repeated and four in which neither component repeated. On trials when only the disk component was present, 4 adapt-test sequences could be created. Among these were 2 in which the disk repeated and 2 in which it did not. On trials when only the annulus component was present, 4 adapt-test sequences could be created, 2 in which the annulus repeated and 2 in which it did not. Each block of 24 successful trials contained one instance of each adapt-test sequence. Sequencing within a block was pseudo-randomized. Trials terminated due to fixation-break were repeated to ensure that one trial under each of the 24 conditions was completed successfully. A full run consisted of four successive blocks and thus of 96 successful trials encompassing four repetitions of each possible adapt-test sequence.

4.3.6 Neurophysiological data collection

Cylindrical Cilux recording chambers with an inner diameter of approximately 2 cm (Crist Instrument Co. Inc., Hagerstown, MD) were implanted over the left hemisphere (M1) and

right hemisphere (M2) of V2. Chambers were centered approximately 18 mm posterior and 6 mm lateral to Horsley-Clarke zero.

Chamber placement was guided by pre-surgical T1-weighted structural MRI scans in a 4.7-Tesla scanner. The chambers gave access to cortex on the posterior bank of the lunate sulcus. Chamber grid holes were functionally mapped to retinotopic coordinates and the line of inversion between V2 and V1 was located to determine the extent of access to V2 cortex. On each recording day, a cylindrical plug containing guide holes arranged at 1 mm spacing in a square grid was inserted in the chamber (Crist Instrument Co. Inc., Hagerstown, MD). A blunt guide tube was inserted through one grid hole until contact with the dura matter was made and a single varnish coated tungsten microelectrode (FHC Inc., Bowdoin, ME) with an impedance of 0.1-5 MΩ at 1kHz was advanced through the guide tube by use of a hydraulic micromanipulator (M0-10; Narishige International Inc., East Meadow, NY). Upon encounter with neurons giving phasic responses to visual stimuli, recording commenced. Threshold-crossing events, sampled at 40 KHz, were digitally recorded and stored for offline sorting by a Plexon MAP system (Plexon Inc., Dallas, TX).

4.3.7 Database

We sorted waveforms from each session using a PCA-based approach implemented by Plexon Offline Sorter (Plexon Inc., Dallas, TX). All further steps of analysis were carried out in Matlab on time-stamped action potential markers. We first analyzed data from each neuron to ensure that it met an arbitrary criterion for visual responsiveness. Only if a one-tailed t-test comparing mean firing rate 100-200 ms after adapt-stimulus onset to mean firing rate 100-200 before adapt-stimulus onset yielded an outcome of $p < 0.05$ did we include the neuron in the database for subsequent analyses.

4.3.8 Cluster-Based Permutation Test

To determine whether the post-image-onset time-varying population firing rates measured under two conditions were significantly different, we employed a nonparametric approach requiring neither arbitrary designation of a measurement window nor comparisons within multiple windows (*20*). The starting point for this analysis was a table of mean firing rates of each neuron under each condition in each 5 ms bin spanning a window from 0 ms to 560 ms following image onset. In each bin, we carried out a paired T-test on the two distributions of firing rates. If the test yielded a p-value < 0.05, we tagged the bin as positive or negative according to which condition was associated with the higher firing rate. If the test yielded a p-value ≥ 0.05, we tagged the bin as zero. The T-test was used only as a means for imposing an arbitrary threshold and not to assess statistical significance. For each cluster of bins of uniform sign, either positive or negative, we computed the sum across those bins of the associated T-statistics. The cluster with the greatest sum was classified as "best" for further steps of analysis. To test its statistical significance, we generated a permutation distribution, applying the above-described procedure 1,000 times to data in which the condition labels for each neuron had been randomly shuffled. We computed p as the fraction of iterations in which the best-cluster sum of T-statistics was greater than the best-cluster sum of T-statistics in the original data. We classified the original cluster as significant if $p < 0.05$. If this cluster was significant, we proceeded to assess the statistical significance of the observed cluster with the next highest sum of T-statistics. We computed p for this cluster as the fraction of cases in the previously generated permutation distribution for which the sum of T-statistics was greater than the observed value. We classified the cluster as significant if $p < 0.05$. We repeated this procedure until a non-significant result was obtained or no observed cluster larger than four bins remained.

4.3.9 Analysis of Suppression Timing

At the population level, we measured suppression as the difference in mean firing rate between the no match condition and the condition of interest. Having smoothed the data in 1 ms bins with a 10 ms SD Gaussian kernel, we proceeded to compute the time of onset as the time at which the response reached half-height. We defined the peak as the maximum value 50-300 ms post stimulus, and the time to half-height as the earliest 1 ms bin in which the effect reached half the peak value in the window 50-300 ms following test image onset.

4.4 RESULTS

We collected data from 78 visually responsive V2 neurons (35 in M1 and 43 in M2) while monkeys performed a task in which they were required to maintain central fixation while a sequence of adapter and test images were displayed peripherally in the visual hemifield contralateral to the recording chamber centered on the receptive field of the recorded neuron (sequence in **Figure 4.1A**). For testing each neuron, we employed two disks and two annuli, cut from different natural scenes, which could be combined to form 4 unique disk-annulus stimuli. Disks and annuli that appeared in the same session were never taken from the same scene, so there was never continuity of features across the disk-annulus boundary. This was done to ensure that the level of discontinuity of features between the disk and annulus was controlled across stimuli. There were three trial types with respect to which components appeared on the given

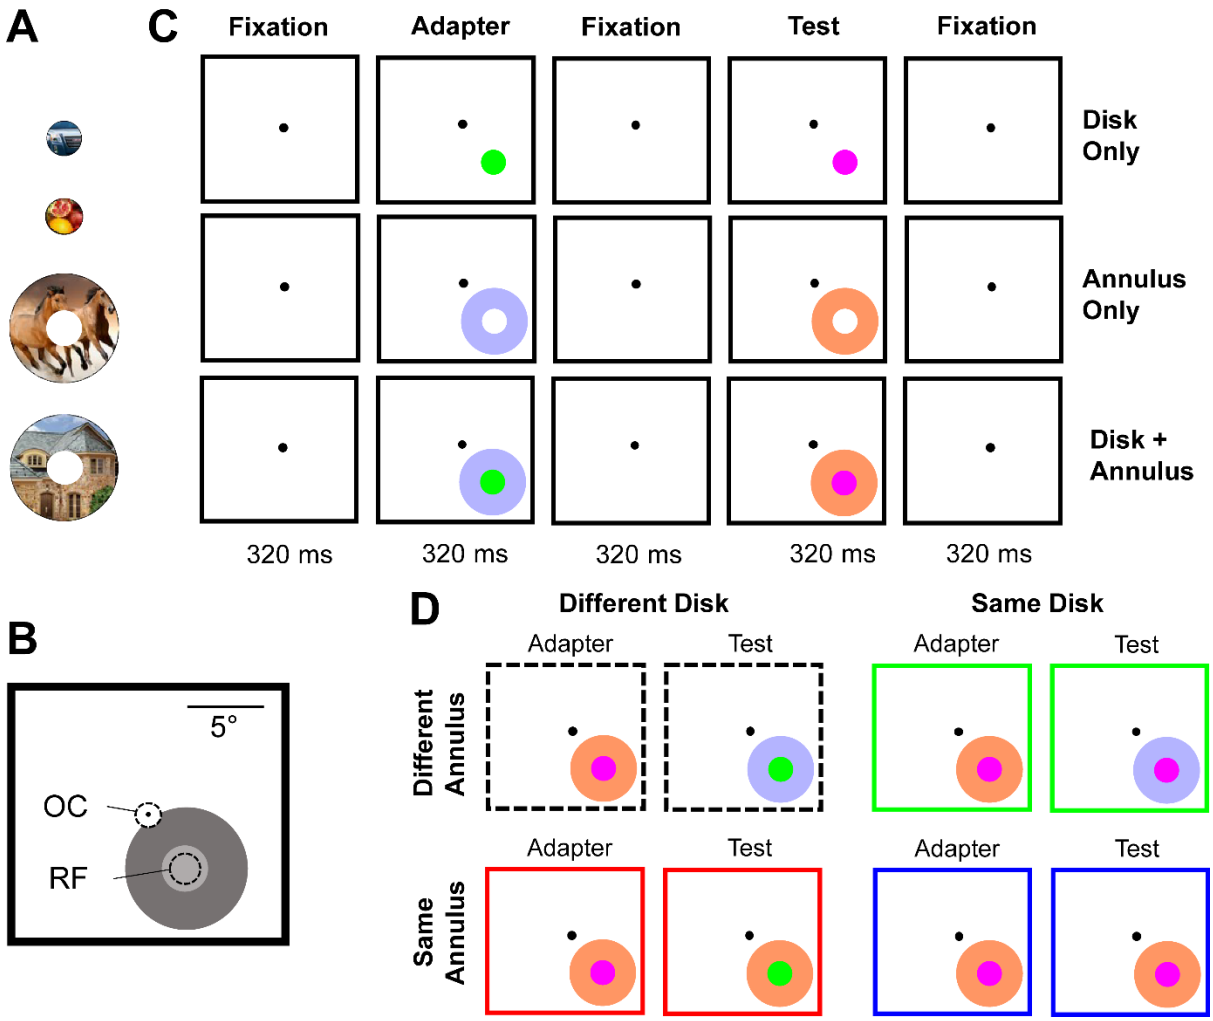

Figure 4.1. Experimental design. A. For each recording session, 2 disks and 2 annuli were cut from unrelated natural scenes. **B.** Both the disk (light grey; diameter = 3°) and annulus (dark grey; inner diameter = 3°; outer diameter = 8°) were centered on the receptive field of the neuron being recorded (RF: dashed circle). An invisible occluder (OC; diameter = 1.6°) prevented stimuli from impinging on the fixation point. **C.** On each trial, two displays (adapter followed by test) were presented in sequence in the visual field contralateral to the recording hemisphere. On a given trial, the displays might consist of a disk alone, an annulus alone or a combination of disk and annulus. **D.** On disk + annulus trials, there were four possible match conditions determined by disk match status (same as adapter or different) and annulus match status (same as adapter or different). The panels depicting each condition are bordered by a color identical to the color of the curve depicting mean firing rate under that condition in **Figure 4.3A**.

trial: There were trials in which only the disk component appeared (**Figure 4.1A,** top row), trials in which only the annulus component appeared (**Figure 4.1A,** middle row), and trials in which the disk plus the annulus component appeared (**Figure 4.1A,** bottom row). Trials involving only the disk or only the annulus conformed to 4 sequences obtained by crossing test identity (A or B) with adapter identity (A or B). For trials on which the disk and annulus appeared together there were 16 possible sequences. The test might match the adapter in only the disk component, only the annulus component, both components or neither component (**Figure 4.1C**). In describing the results, data is combined across the two monkeys.

4.4.1 Repetition suppression for disk alone and annulus alone

We first asked whether repetition suppression was present, as indicated by lower mean population firing rate when the test matched the adapter than when it did not, on trials when either the disk or the annulus appeared in isolation. To quantify the level of suppression when the test matched the adapter on trials in which only the disk was present and separately on trials when only the annulus was present, we subtracted the mean population firing rate on match trials from the mean population firing rate on no match trials. The no match trials acted as a baseline in which we would expect no image-specific suppression to occur. This yielded two time-varying indices in which suppression was expressed as a positive value. For trials in which the disk appeared alone, repetition suppression was present (**Figure 4.2A**) and was statistically significant ($p < 0.001$, cluster-based permutation test, **Figure 4.2B**). Suppression on trials in which the annulus appeared alone was very weak, if present at all, and did not achieve statistical significance by the same cluster-based permutation test as that employed on data from the disk alone trials. To measure the timing of the onset of suppression, we calculated the time to half-

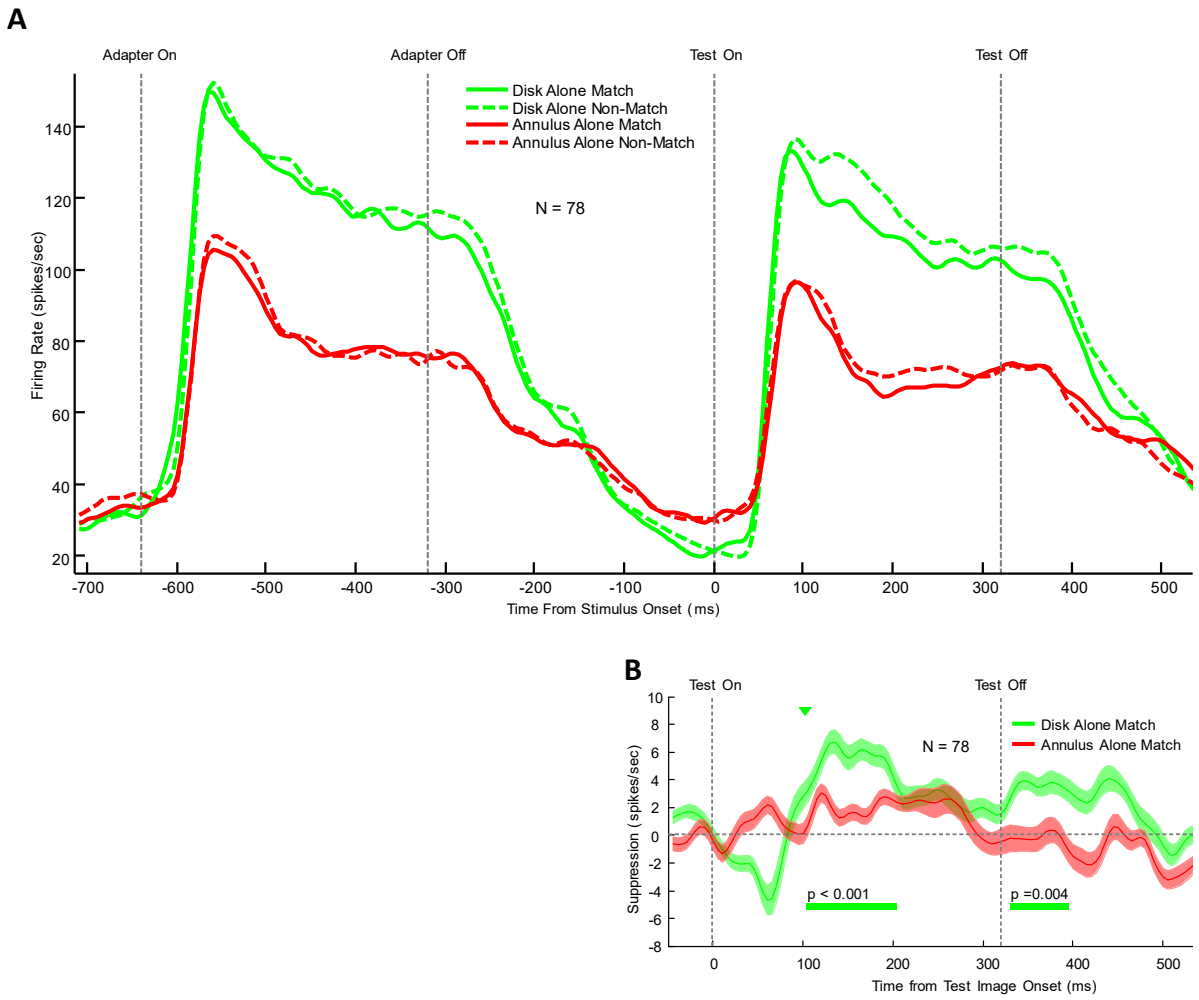

Figure 4.2. Repetition suppression for disk-alone and annulus-alone conditions. A. Population mean firing rate as a function of time during disk-alone trials (green) or annulus-alone trials (red). For each display type, the test could match the adapter (solid curve) or differ from it (dashed curve). **B.** Population mean firing rate on nonmatch trials minus population mean firing rate on match trials under the disk-alone condition (green) and the annulus-alone condition (red). Repetition suppression is reflected in positive values. Ribbons represent ± standard error of the mean. The underlying horizontal bars of corresponding color denote periods during which suppression was significantly greater than zero as indicated by a cluster-based permutation test. The p-value for each significant result is appended to the bar. The green triangle indicates the time to half-height (103 ms) of disk-alone repetition suppression. Curves were smoothed with a 10 ms Gaussian kernel.

height of the suppression signal. The onset of suppression for the test image on trials when the disk appeared alone and was a repetition of the adapter occurred at 103 ms post stimulus based

on the time to half-height measure.

4.4.2 Repetition Suppression for disk plus annulus

To characterize the separate contributions of the central (disk) and surround (annulus) components of a large complex image to repetition suppression in V2, we analyzed the population response as a function of whether the test matched the adapter in only the disk component, only the annulus component, both components or in neither component on trials in which the disk and annulus were presented together. There was clear evidence for repetition suppression when both components matched the adapter, as well as when only the disk matched, but not when only the annulus matched (**Figure 4.3A**). To quantify the level of suppression under the three different match conditions we subtracted the mean population firing rate under each from the no match condition which acted as a baseline in which we would expect no image-specific suppression to occur. This yielded three time-varying indices in which suppression was expressed as a positive value. The effect of disk match and both match suppression achieved statistical significance in the cluster-based permutation test (underlying bars in **Figure 4.3B**) but the effect of annulus match did not and remained indistinguishable from the no match condition. The repetition suppression effect was noticeably stronger when both the disk and the annulus matched the adapter than when only the disk was a match, despite the fact that there was no effect of repetition of the annulus when the disk was not a repetition. To quantify this effect we computed one additional time-varying index by subtracting the mean firing rate on trials in which both the disk and the annulus were the same as the adapter (both match) from the mean firing rate on trials in which the disk was the same but the annulus was not (disk match). This gave a measure of the suppression induced by adding annulus repetition to disk repetition. This effect achieved significance by the cluster-based permutation test (underlying bar, **Figure 4.3C**).

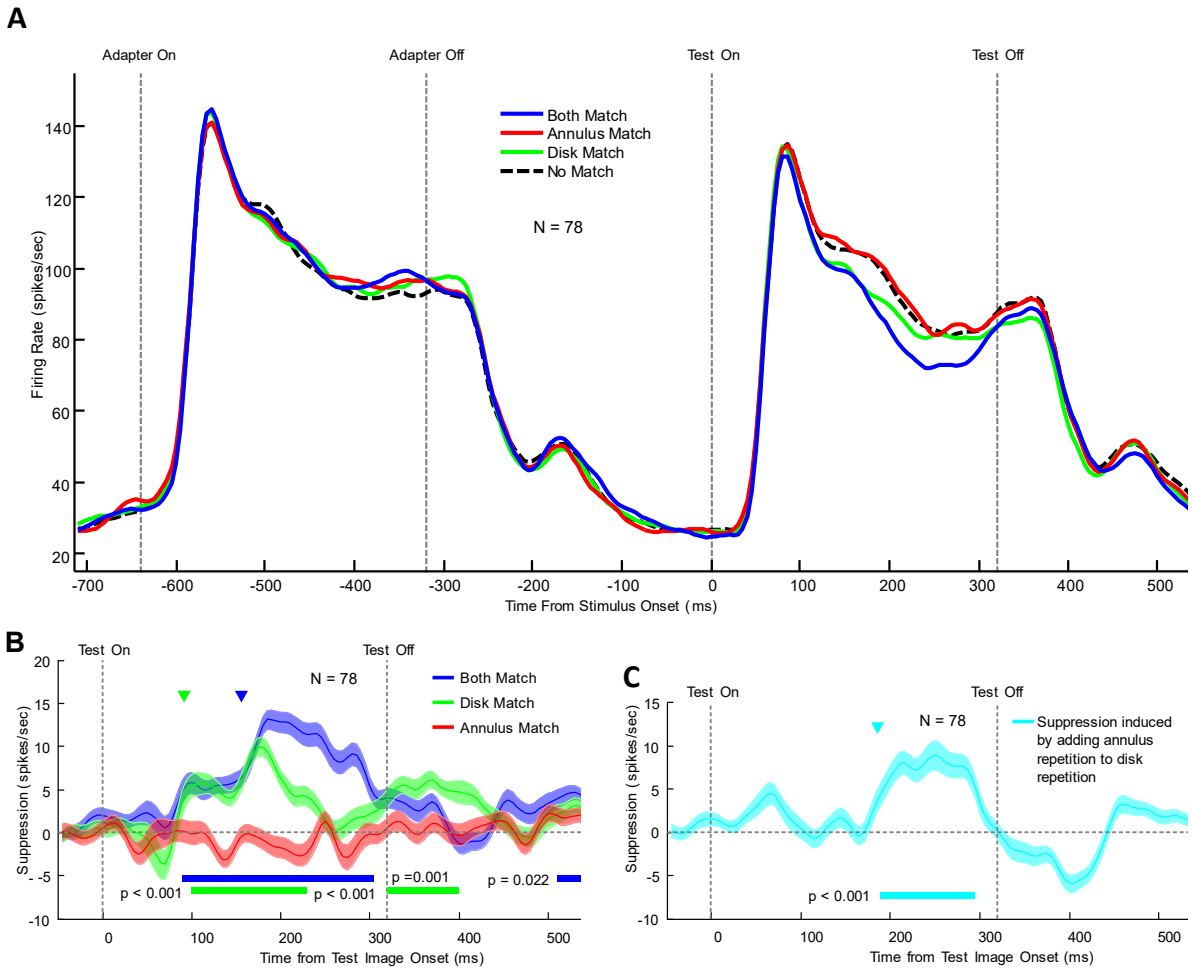

Figure 4.3. Repetition suppression for displays containing both disk and annulus. A. Population mean firing rate as a function of time during trials sorted according to the match status of disk and annulus: both match (blue), disk matches (green), annulus matches (red) or neither matches (black dash). **B.** Population mean firing rate on nonmatch trials (dashed black in A) minus population mean firing rate on trials conforming to three match conditions: both match (blue in A), disk match (green in A), annulus match (red in A). Repetition suppression is reflected in positive values. Green and blue triangles indicate the time to half-height of repetition suppression under the disk-match (92 ms) and both-match (155 ms) conditions respectively. **C.** Population mean firing rate on disk-match trials (green curve in A) minus population mean firing rate on both-match trials (blue in A). Late enhancement of repetition suppression under the both-match condition achieved significance in a cluster-based permutation test (underlying horizontal bar with appended p-value). Triangle indicates time to half-height of the enhancement of repetition suppression under the both-match condition (185 ms). All other conventions as in **Figure 4.2**.

We also computed the time to half-height for the suppression induced by the annulus matching when the disk was also a match. The time to half-height for this effect was 185 ms, markedly later than that when only the disk was repeated on trials in which the disk and the annulus appeared together (92 ms) and on repetition trials when the disk appeared in isolation (103 ms).

4.5 DISCUSSION

This study sought to untangle the contributions of the central component and the surround component of large complex images to repetition suppression effects which we have previously been observed for natural object stimuli in V2. The greater suppression we observed on trials in which both the disk and the annulus matched the adapter than on trials in which only the disk matched suggest that both components play a role in the suppression generated for repetitions of large complex objects in V2. Prior to this study, it was unclear what impact the content in the surround of the large complex images we used to study repetition suppression in V2 neurons might have been. It could have been the case that repetition of content outside the classical receptive field would lead to enhancement or at least reduced suppression as adaptation to large gratings did in one study of V1 neurons (*15*), although replicating this outcome in our current study would have seemed unlikely given the robust repetition suppression we saw previously for large natural objects. Based only on our previous results in V2 using natural object stimuli, it also could have been the case that the repetition suppression we observed was only due to the image content in the classical receptive field and that features of an object which fell outside the receptive field had no impact on suppression. We have shown here that neither of these assumptions holds true when tested and that content in the surround of the image does contribute to repetition suppression, leading to stronger suppression when the content in the surround as well as the center repeats, compared to when only content in the center repeats. This

demonstrates that the phenomenon observed in V1 with grating stimuli does not transfer to V2 with complex stimuli. The presentation duration of the adapter also differed between the V1 study and our own, with the V1 study using long adapter durations of 5 or 40 seconds compared to our brief 320 ms presentations, which could contribute to the differing results. Whether or not V2 neurons show the same pattern of repetition suppression effects for gratings and for longer adapter durations and conversely how V1 neurons would respond to brief presentations of gratings or natural images are all questions for further study.

We also recorded responses on trials in which both the adapter and the test contained either only the disk or only the annulus. One feature of note based on adapter responses on these trials is that the annulus alone generated a significant response, although it was not nearly as strong as the response to the disk alone. This might seem surprising given that stimulation of the receptive field surround in V1 and V2 is known to generate a suppressive signal which reduces the response of the neuron, however these suppressive effects are typically measured when the classical receptive field is also stimulated (*8, 14*). In the comparatively few studies which present data on the response to annuli alone without the presence of a stimulus in the excitatory receptive field, they find that annuli alone do generate excitatory responses, although they are much weaker than when the receptive field center is stimulated with an optimal or near-optimal image (*7, 10, 18*). When only the classical receptive field is stimulated it generates a robust excitatory response. When stimulation of the surrounding area of visual space is added to stimulation of the classical receptive field the response is reduced compared to when only the classical receptive field is stimulated. When only the surround is stimulated it generates a weak but excitatory response. This highlights the complexity of the interactions between stimulation within and outside of the classical receptive field. The response to the annulus alone was especially robust in

our data when compared to examples present in the literature. To attempt to quantify the relative strength of the annulus response for studies in which it was measured in a comparable way to our study, we took as a measure of relative strength of annulus alone response the ratio of disk alone response to annulus alone response (disk:annulus), such that a lower value of disk:annulus represented a comparatively stronger annulus response, and a disk:annulus value of 1 would mean that the response to the annulus was as strong as the response to the disk. In a study by Cavanaugh et al. in V1 neurons of anesthetized macaques using grating stimuli on a uniform luminance display their disk:annulus ratio was ~3.5 to 3.6 based on a few example single neurons (*7*). In a study by Gieselmann & Thiele in a population of V1 multi-units recorded in awake macaques using grating stimuli the disk:annulus ratio of the mean population response was ~2.4, showing a more robust annulus response than in the Cavanaugh study (*10*). In a third study conducted by Ziemba et al. in a population of V2 neurons in anesthetized macaques using naturalistic texture patches their disk:annulus ratio was ~2.9 based on data from an example neuron, somewhere between the Cavanaugh et al. and Gieselmann & Thiele results (*18*). In our own data our disk:annulus ratio was 1.9, a higher relative annulus response than these few examples in the literature by anywhere from 26 to 90%. Factors which could account for the difference include that: 1) We were recording from V2 and not V1 (annulus responses were higher in anesthetized V2 than in anesthetized V1). 2) That we were recording in awake animals (annulus responses were higher in awake V1 than in anesthetized V1). 3) Our stimuli were generally of higher luminance than the background display, and not controlled for luminance, as is common in studies using natural stimuli in higher order visual areas (annulus responses were higher in V1 when luminance was not controlled than when mean luminance was held constant). It could be that the increase in luminance of our stimuli compared to the background led to an

increase in the firing rates observed, but that this luminance dependent increase in firing rate would likely be stimulus nonspecific and therefore would not contribute to image-specific repetition suppression. 4) We were using patches of natural scenes and not gratings or synthetically generated texture patches, although there are no parallel comparisons for this factor. The examples that exist in the literature and which are summarized here used trial averaged firing rates for disk and annulus size tuning curves, often normalized, and did not present population peristimulus time histograms as we do. It is therefore difficult to get a full sense of the response to annuli alone from what currently exists in the literature. Our goal with the current experiment was to determine the contribution of surround image content to repetition suppression in V2 and not to map the surround responses in detail, although given the current gaps in the literature this is an area worthy of further study.

In the current study suppression did occur when a small image centered on the receptive field was repeated, but no significant suppression occurred when an annulus directly surrounding the receptive field was repeated. The assumption that only content in the classical receptive field should contribute to repetition suppression might at first appear to be approximately true based on these observations alone, however these results do not speak directly to the question of whether or not content in the surround of a large image can contribute to repetition suppression. The fact that no measurable repetition suppression was present when the annulus alone was repeated, as well as when the annulus repeated but the disk did not when they were presented together, makes the finding of enhanced suppression when both the disk and annulus are a match even more surprising. It is as though there is a gating mechanism at play whereby repetition of content in the surround leads to suppression, contingent on content in the center also being repeated, and otherwise not.

An alternative explanation for why we see a contribution of the annulus to repetition suppression could be that the annuli we used in these experiments simply encroached on the classical receptive fields of the recorded neurons. This explanation might seem plausible based on the unusually high responses to the annulus presented alone in our data, however it is unlikely for two reasons: 1) Receptive field sizes in parafoveal V2 are on average 1-2° in diameter, smaller than the inner diameter (3°) of the annuli used (*14*). 2) If this were the case and the annulus fell on the classical receptive field of the recorded neuron, repetition suppression arising from the annulus matching would be expected to also be present when the annulus alone matched the adapter and when the annulus matched in the presence of a nonmatching disk, however this was not the case.

We also analyzed the timing of repetition suppression in this study. We used as a measure of the relative onset of repetition suppression the time to half-height of the time varying suppression signal under each condition. The onset of repetition suppression when the disk was repeated both on trials when it appeared alone (103 ms) and with a non-matching annulus (92 ms) were in rough agreement with each other. The timing of repetition suppression for the condition in which both the disk and the annulus were a match to the adapter had a markedly later onset of 155 ms. When only the contribution of suppression induced by adding annulus repetition to disk repetition was considered, the time to half-height occurred 185 ms after stimulus onset, an even more pronounced difference. The suppression generated when only the disk matched also fell off quickly, whereas when both components matched suppression remained high longer into the late phase of the response. The fact that the contribution of the annulus matching to repetition suppression in V2 has a delayed time course is consistent with the idea that it could be mediated in part by neurons adjacent to the recorded neuron, with classical

receptive fields in the surround of the recorded neuron, and that the added delay is due to increased processing and/or conduction time.

4.6 REFERENCES

1. C. G. Gross, D. B. Bender, C. E. Rocha-Miranda, Visual receptive fields of neurons in inferotemporal cortex of the monkey. *Science* **166**, 1303-1306 (1969).
2. A. Lueschow, E. K. Miller, R. Desimone, Inferior Temporal Mechanisms for Invariant Object Recognition. *Cerebral Cortex* **4**, 523-531 (1994).
3. M. Ito, H. Tamura, I. Fujita, K. Tanaka, Size and position invariance of neuronal responses in monkey inferotemporal cortex. *J Neurophysiol* **73**, 218-226 (1995).
4. R. Desimone, T. D. Albright, C. G. Gross, C. Bruce, Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* **4**, 2051-2062 (1984).
5. R. Gattass, A. P. Sousa, M. Mishkin, L. G. Ungerleider, Cortical projections of area V2 in the macaque. *Cerebral cortex (New York, N.Y. : 1991)* **7**, 110-129 (1997).
6. K. R. Gegenfurtner, D. C. Kiper, S. B. Fenstemaker, Processing of color, form, and motion in macaque area V2. *Vis Neurosci* **13**, 161-172 (1996).
7. J. R. Cavanaugh, W. Bair, J. A. Movshon, Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *J Neurophysiol* **88**, 2530-2546 (2002).
8. A. Angelucci *et al.*, Circuits for local and global signal integration in primary visual cortex. *J Neurosci* **22**, 8633-8646 (2002).
9. S. Shushruth *et al.*, Different orientation tuning of near- and far-surround suppression in macaque primary visual cortex mirrors their tuning in human perception. *J Neurosci* **33**, 106-119 (2013).
10. M. A. Gieselmann, A. Thiele, Comparison of spatial integration and surround suppression characteristics in spiking activity and the local field potential in macaque V1. *Eur J Neurosci* **28**, 447-459 (2008).
11. S. Durand, T. C. Freeman, M. Carandini, Temporal properties of surround suppression in cat primary visual cortex. *Vis Neurosci* **24**, 679-690 (2007).
12. B. S. Webb, N. T. Dhruv, S. G. Solomon, C. Tailby, P. Lennie, Early and late mechanisms of surround suppression in striate cortex of macaque. *J Neurosci* **25**, 11666-11675 (2005).
13. A. Angelucci *et al.*, Circuits and Mechanisms for Surround Modulation in Visual Cortex. *Annu Rev Neurosci* **40**, 425-451 (2017).
14. S. Shushruth, J. M. Ichida, J. B. Levitt, A. Angelucci, Comparison of spatial summation properties of neurons in macaque V1 and V2. *J Neurophysiol* **102**, 2069-2083 (2009).
15. S. C. Wissig, A. Kohn, The influence of surround suppression on adaptation effects in primary visual cortex. *J Neurophysiol* **107**, 3370-3384 (2012).
16. J. B. Levitt, D. C. Kiper, J. A. Movshon, Receptive fields and functional architecture of macaque V2. *J Neurophysiol* **71**, 2517-2542 (1994).
17. N. A. Crowder *et al.*, Relationship between contrast adaptation and orientation tuning in V1 and V2 of cat visual cortex. *J Neurophysiol* **95**, 271-283 (2006).
18. C. M. Ziemba, J. Freeman, J. A. Movshon, E. P. Simoncelli, Selectivity and tolerance for visual texture in macaque V2. *Proceedings of the National Academy of Sciences of the United States of America* **113**, E3140-3149 (2016).
19. J. Freeman, C. M. Ziemba, D. J. Heeger, E. P. Simoncelli, J. A. Movshon, A functional and perceptual signature of the second visual area in primates. *Nat Neurosci* **16**, 974-981 (2013).

20. E. Maris, R. Oostenveld, Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* **164**, 177-190 (2007).

CHAPTER V

GENERAL DISCUSSION

5.1 Summary of Results

5.1.1 General Summary

Repetition suppression is characterized by a reduction in the response to a subsequent presentation of a stimulus compared to an initial presentation, and is most robust at short latencies, with suppression falling off at increasing delays. It is a robust neurophysiological phenomenon observed in both humans and non-human primates across many different experimental paradigms and recording modalities (*1-7*). It has been most closely studied in regions of the temporal lobe such as area TE in monkeys (*8*). Despite its robust and reproducible nature, both the mechanisms which give rise to repetition suppression and its functional consequences remain poorly understood. The overarching goals of the studies presented here in Chapters II - IV have been twofold; 1) to characterize repetition suppression within TE with the aim of answering several outstanding questions regarding the systems level mechanisms which give rise to repetition suppression in TE as well as the functional consequences of repetition suppression in TE for downstream areas and behavior, and 2) to push the study of repetition suppression beyond TE to other visual areas at different levels of the hierarchy, which have been understudied with regard to repetition suppression, specifically area V2 which resides at a lower order of the hierarchy compared to TE and which provides input to and receives feedback connections from TE (*9-11*).

The results and findings specific to each experiment and each brain area under study will be discussed in the following sections and cover how they relate to a general understanding of

repetition suppression and more specifically to its mechanisms and their implications for sensory processing and behavior.

5.1.2 Summary of Results for Experiments Performed in TE

In one experiment in TE focusing on the contribution of individual features to repetition suppression, we manipulated stimulus color and stimulus shape independently in the standard repetition suppression paradigm described previously. In this experiment we found that:

1. Repetition suppression in area TE of inferotemporal cortex operates on the level of features and not at the level of the image as a whole.
2. There is no measurable effect of firing rate fatigue in the recorded neurons.
3. The tuning of a neuron to particular features governs the level of suppression when those features are repeated.
4. Neurons in TE are uniformly distributed with respect to the strength of image-specific repetition suppression, with no evidence for separate suppressing and non-suppressing populations. This has consequences for how repetition suppression could be read out by downstream areas and used to guide behavior.
5. Image repetitions generate a stronger trough-rebound dynamic than non-repetitions.
6. Suppression is slightly stronger when both features match than the sum of suppression for each of the individual features.

5.1.3 Summary of Results for Experiments Comparing Repetition Suppression in TE and V2

In another set of experiments conducted in both TE and V2 using the same stimulus parameters in the standard paradigm commonly employed in the study of TE and described previously, we were able to make comparisons between repetition suppression effects observed

in TE and V2. In this set of experiments we manipulated image identity and image location independently with respect to which element was repeated at test. In TE we found that:

1. Image-specific repetition suppression occurred even when the adapter and test appeared at different locations.
2. Image-specific repetition suppression was stronger when both the adapter and test image appeared at the same location.
3. Image-specific repetition suppression in the different location condition had a delayed onset compared to the same location condition.
4. The trough-rebound dynamic was stronger when the same image appeared at the same location than when a different image appeared at the same location; in other words, it accompanied only image-specific repetition suppression.

Analyzing the pattern of repetition suppression effects in V2 under the same experimental conditions, we found a different pattern of results. In V2 we found that:

1. Unlike in TE, image-specific repetition suppression only occurred when the adapter and test appeared at the same location, and did not occur when they appeared at different locations.
2. Unlike what was found in TE, there was also a location-specific repetition suppression effect whereby the response to the test image was weaker when it appeared at the same location as the adapter even when it was a different image.
3. Similar to what was found in TE, the trough-rebound dynamic was stronger when the same image appeared at the same location than when a different image appeared at the same location; in other words, it accompanied only image-specific repetition suppression.

Analyzing the pattern of repetition suppression effects in a subset of recorded V2 neurons under slightly different experimental conditions in which the different location adapter appeared in the receptive field surround of the recorded neuron rather than in the upper visual field quadrant, we made the following additional observations:

1. Modest image-specific repetition suppression does occur when the adapter and test appeared at different locations when the adapter appeared at a location directly adjacent to the receptive field of the recorded neuron.
2. This image-specific suppression occurred at a longer latency than other repetition suppression effects observed in V2 or TE.

In analyses directly comparing the repetition suppression signal in TE and V2 we found:

1. At the population level, the image-specific repetition suppression signal in V2 precedes the signal in TE.
2. At the level of individual neurons, the onset of the image-specific repetition suppression signal in V2 neurons consistently preceded the signal in TE neurons.
3. At the population level, the deepest part of the trough of the response dynamic in V2 preceded the trough in TE by approximately the same amount of time as the difference in onset of the visual response between the two areas.

Taken together, these results provide little evidence that repetition suppression is due to either predominantly top-down or predominantly bottom-up influences in either V2 or TE, but rather that it could arise largely independently in each cortical area, perhaps due in part to lateral connections within an area.

5.1.4 Summary of Results for Experiments Performed in V2

We performed an additional experiment in V2, again using the standard repetition suppression paradigm typically used in the study of TE. In this experiment the goal was to determine the separate contributions of image content in the center of an image and content in the surround of an image for images which extended beyond the classical receptive field of a typical V2 neuron, such as the relatively large complex natural images typically used to study repetition suppression in TE and those used in our study comparing responses in TE and V2 under the same stimulus conditions. To do this, we manipulated independently the repetition status of image content in the center and in the surround of a compound image consisting of two separate components which were contiguous. Looking at responses in V2 under these conditions we found that:

1. Content in the surround of an image contributes to repetition suppression, but only when the content in the center is also repeated.
2. The contribution of content in the surround to repetition suppression occurs at an increased latency, possibly indicating that it is mediated by lateral connections which could be coming from neurons whose receptive fields are centered in the image surround.

5.2 Repetition suppression in V2

It is interesting in its own right that we observed such robust repetition suppression effects in V2 under these stimulus conditions, which have previously been used to induce repetition suppression in TE and therefore were tailored to the properties of TE neurons i.e. preference for complex naturalistic images and large receptive fields, and not specifically tailored to the response properties of V2 neurons. We had some reason to believe that we might

see repetition suppression under these conditions in V2, as suppressive effects due to stimulus adaptation have been observed in V2 (*12, 13*), albeit under very different experimental conditions. It was, however, unclear what form these effects might have taken.

The images used for experiments conducted in V2 were not quantified based on their physical properties, and so it is not possible to determine to what extent specific suppression effects were due to how well the stimuli matched the response properties of the recorded V2 neurons. The extent to which low level image properties match the tuning of a V2 neuron to these properties, and how this is related to the level of repetition suppression observed, are interesting questions and could be explored in more detail in future studies.

Taken together, these results highlight the need to study repetition suppression at the level of individual neurons and in multiple brain areas to form a more complete picture of how it is generated across the visual system. In the following sections, I will revisit what is known about repetition suppression in light of our current findings.

5.3 Generalization of repetition suppression in TE to changes in position and to new images resembling the adapter

The results presented in Chapter III (**Figure 3.2**) confirm previous findings (*2, 14-16*) that in TE, repetition suppression does generalize across spatial locations, although the strength of suppression is reduced compared to repetitions that appear at the same location. We did not test manipulations of image size. We have extended these findings from TE to V2, showing that under certain conditions V2 neurons also show repetition suppression that generalizes across spatial locations, although to a much more limited degree than that seen in TE neurons.

It has been shown that repetition of similar but non-identical, parametrically manipulated images also generates repetition suppression (*2, 17*), although the principles governing the level of suppression observed at different distances in parameter space are as yet undetermined. Our findings presented in Chapter II shed some light on these principles. We have shown that a major principle governing the level of suppression is the degree of feature overlap between the adapter and the test image (**Figures 2.4, 2.5A, 2.6**). Furthermore we have shown that at the level of the individual neuron, the degree to which the features being repeated match the preferences of the neuron also governs the level of suppression such that suppression will be stronger for the features that the neuron prefers and weaker when non-preferred features are repeated (**Figure 2.6**). Therefore the level of cross-suppression between two similar but non-identical images is related to the features which the two images share in common.

5.4 Relation of repetition suppression to identity non-specific suppression

In all studies presented here we have attempted to eliminate the confound between image-specific and image non-specific suppression by always comparing the matching test image to a non-matching test rather than to the adapter. It appears that in most contexts some identity non-specific suppression does occur, and while we have attempted to remove its influence as a nuisance variable, its close association with image-specific repetition suppression is worth further consideration. Are similar mechanisms involved, or are they two distinct processes? Is there an interaction between identity-specific and identity non-specific suppression? The ability to effectively experimentally disentangle these phenomena opens the door to addressing these questions in more detail which could be the work of future studies. Interestingly, we did observe a robust form of image non-specific repetition suppression in V2 that did not appear to occur in TE. In V2 we observed robust suppression when a completely

different image appeared at the same location compared to when different images were presented at different locations. This effect was not observed in TE. Is this location-specific repetition suppression due the same or different underlying mechanisms as image-specific suppression? Why is it present in V2 but absent in TE? Does it have to do with the different receptive field sizes in the two areas? These are questions for future studies.

5.5 Stimulus specificity of the enhancement of oscillatory amplitude that accompanies repetition suppression

As has been noted in prior studies of repetition suppression in TE (*2, 16-21*), we have found evidence that repetition suppression is accompanied by an enhancement in oscillator amplitude at approximately 5 Hz. In addition to confirming this basic finding, we have also demonstrated that the enhancement of oscillatory amplitude is image-specific in that it occurs more on same-image tests than for different-image tests. We have also determined that it occurs consistently across individual neurons (**Figures 2.7, 3.3**). In our experiments in which stimulus location was manipulated, the enhancement in oscillatory amplitude was confined to same-location trials in addition to being image-specific, as the same image presented at a different location did not induce a strong oscillatory dynamic. We further extended these findings in Chapter III by demonstrating that there is an image-specific enhancement in repetition induced oscillatory amplitude in V2 (**Figure 3.5**) that is similar to what we observed in TE. The fact that 5 Hz oscillations are only present in same-image trials suggests that the mechanism which gives rise to these oscillations, whether it is primarily inhibitory or excitatory in nature, is due to image-specific repetition suppression, although the same or a similar mechanism generating 5 Hz oscillations may be involved in other processes.

It is interesting that we observed the same phenomena in V2 that we found in TE and which has been previously noted by others in TE. Theta oscillations have a rich history of study in the hippocampus, in which the hippocampal theta rhythm has been implicated in many studies as being a substrate for memory encoding and recall (*22-24*). It is possible that theta oscillations in areas of cortex proximal to the hippocampus such as parahippocampal and perirhinal cortex could be transmitted from the hippocampus and be driven by hippocampal theta. Given the relative proximity of TE to the hippocampus, and previous theories of repetition suppression that implicated it in recognition memory, it could have been supposed that the increase in oscillatory amplitude accompanying repetition suppression was somehow related to hippocampal theta and recognition memory. The fact that we also observe an increase in oscillatory amplitude related to image-specific repetition suppression in V2 makes any interpretation of a 5 Hz repetition signal in TE being related to hippocampal theta or memory recall be put into serious doubt, given how much further V2 is from the hippocampus than TE, and the fact that the trough of the oscillatory response in V2 leads the trough in TE. It is much more likely given these observations that resonance in local reciprocally inhibitory circuits in each area gives rise to the increase in oscillatory amplitude independently, or that they propagate from early to late visual cortical areas. The broader function of these changes in response dynamic remains to be determined.

5.6 Relation of repetition suppression to familiarity suppression

The fact that repetition suppression was observed in V2 under the same stimulus conditions used in studies of TE, combined with the fact that V2 also exhibits familiarity suppression under similar conditions when the same images are repeated over the course of weeks of exposure (*25*), is congruent with the idea that repetition suppression could be mechanistically related to the development of familiarity suppression, and that this could play

out across visual cortex rather than being confined to a particular area or areas. In addition, the fact that the increase in oscillatory amplitude is also observed for familiar stimuli is suggestive that they could share some underlying mechanisms. The findings from our studies leave these interpretations in play.

5.7 Relation of repetition suppression to recognition memory

Our findings presented in Chapter II show that TE neurons appear homogeneous with regard to their level of suppression (**Figure 2.8**). This has implications for how downstream areas might read out the signal of repetition suppression present in TE. As previous studies have suggested, neurons in downstream areas could unambiguously detect a repeat compared to an ineffective stimulus on the basis of firing rate if there existed two separate populations of neurons in TE, one showing consistently strong repetition suppression and the other not (*26, 27*). Previous work has claimed support for this theory (*19*), but none has tested it as thoroughly as we have here. Given our evidence for a homogeneous population in TE with respect to repetition suppression, it is unclear how repetition suppression could be read out unambiguously in downstream areas, and therefore the question of how repetition suppression may impact behavior, and specifically the recognition of previously experienced stimuli, requires further investigation.

5.8 Repetition suppression and the existence of separate suppressing and non-suppressing subpopulations

As mentioned previously, prior theories of repetition suppression have suggested that separate suppressing and non-suppressing populations could unambiguously signal the repetition status of an image and that this could be used by downstream neurons to decode stimulus recency. Findings presented in Chapter II (**Figures 2.8,2.9**) have shown that neurons are

homogeneous with respect to their level of repetition suppression. This has implications for several theories of repetition suppression which will be discussed in detail in the following section, but in particular that it is unlikely that firing rates in subpopulations of TE neurons could be used as a recency code.

5.9 Theories on potential mechanisms of repetition suppression

A very prominent theory paper on repetition suppression by Grill-Spector et al. 2006 (*28*) outlined three main models of repetition suppression that have been previously proposed in the literature, which I will briefly review here in light of our current findings:

Fatigue model of repetition suppression. Our findings in Chapter II are at odds with fatigue-based models of repetition suppression. Both in terms of our own modeling of our empirical data, which did not require incorporating neuronal fatigue to obtain a good fit to the data, and an empirical investigation of the data on a trial-by-trial level, in which we found no evidence for neuronal fatigue. It is therefore unlikely that fatigue plays a major role in repetition suppression. It is possible that the response to a test that is completely different from the adapter nevertheless appears slightly reduced and that this could be due to a more fatigue-like mechanism, but this is a separate question. This form of identity non-specific suppression may be unrelated to repetition suppression and therefore could arise due to different mechanisms. By employing similar methods as we have here and removing the confounds between identity-specific and identity-nonspecific suppression by careful experimental control, one may be able to look at non-specific suppression more thoroughly in an investigation focused on that phenomenon in the future.

Sharpening model of repetition suppression. In sharpening models, the key feature is that the responses of the majority of neurons are reduced while some smaller subset of neurons, typically those that respond especially strongly, are preserved. This is thought to have the effect of

increasing the sparseness of the population response and biasing the population response towards neurons which represent the stimulus especially well. This has the consequence that there will be neurons which show strong suppression and those which show weak suppression or no suppression at all. As shown in Chapter II (**Figure 2.9**) we found no evidence for separate suppressing and non-suppressing populations, and therefore found no evidence in support of the sharpening model.

Facilitation model of repetition suppression. The main prediction made by facilitation models is that processing is faster for repeated stimuli. This could either be in terms of a more rapid onset of the response or take the form of a shorter duration of the response such that processing "finishes" sooner. This model has the feature that it naturally fits well with behavioral priming results which show faster responding to repeated stimuli, however as discussed in Chapter I the relationship between repetition suppression and repetition priming is a tenuous one. We did not find evidence for the facilitation model in terms of faster onset times for repeated stimuli. Our analysis of the timing of the onset of repetition suppression is not inconsistent with a facilitation model that predicts shorter duration of the response, as we found a relatively late onset for repetition suppression which was not robust in the initial processing phase but peaked at around 200 milliseconds post-stimulus, however following this strong suppression the response rebounded and did not continue to fall after the initial strong suppression at about 200 milliseconds, with the result that response offset times and return to baseline firing was not markedly faster for repeated stimuli. This pattern of response is not directly predicted by a basic facilitation model, however the model could be refined to incorporate a trough-rebound dynamic. It is unclear whether the trough-rebound dynamic we observe is due to response facilitation or some other mechanism and therefore this relationship deserves further investigation.

One additional class of models was not covered in detail by Grill-Spector et al. but is worth mentioning here:

Correlation model of repetition suppression. In general, the studies presented in Chapters III & IV give the impression that repetition suppression is generated, at least to some degree, independently in each visual area, at least as far as TE and V2 are concerned. This could be due to within-region, lateral connections (*29*). This possibility predicts the contrasting effects seen in V2 and TE. This type of within-area lateral interaction is compatible with the mechanistic framework proposed by correlation-based models of repetition suppression (*12, 13, 30, 31*) in which lateral connections are the most plausible biological substrate, for which there is some experimental evidence (*32*). Models featuring lateral connections suggest within-region neural mechanisms such as lateral inhibition, focusing on the role of local connections. These models support the possibility that lateral connections and not purely top-down or bottom-up processes play a role in generating different responses to stimulus repetition in different regions. The functional consequences of repetition suppression predicted by these models has not been thoroughly tested at the level of single neurons in awake animals. These studies pose challenges but are worth exploring in future work to determine whether different local connectivity in different brain areas gives rise to the differing response to stimulus repetition that we have observed.

5.10 Future directions

The studies presented here have expanded our understanding of repetition suppression. Importantly, they can also point to future approaches which can address the gaps that remain in our understanding of repetition suppression.

5.10.1 Behavioral consequences of repetition suppression

What are the behavioral consequences of repetition suppression? A potential candidate is recognition memory, that is the ability to determine whether or not an image is novel or has been repeated. This could be investigated by having simultaneous trial by trial readouts of recognition performance and neural repetition suppression, and to see if suppression is reduced on trials in which a repeat is not recognized. This approach could also be applied to priming by taking reaction time as the behavioral outcome measure. The study of repetition suppression in the context of behavior could be applied to many domains, such as visual search, since the behavioral correlates of repetition suppression are not yet known. The key to the success of this approach will be careful experimental design, controlling for known confounds such as whether the image in question is a target or a distractor and the amount of time that has passed since it was last viewed.

5.10.2 Systems-level mechanisms of repetition suppression

Determining the systems-level mechanisms governing repetition suppression, such as further arbitrating between the various models of repetition suppression outlined in the previous section, will likely be aided by recent improvements in recording technology. Specifically, the ability to record simultaneously from large populations of neurons with well-defined spatial relationships to each other, over longitudinal timescales by using high density chronically implanted microelectrode arrays. The knowledge of the spatial relationship between recorded neurons will give clues as to the roles of lateral connections in repetition suppression and could be aided by causal modeling. Having a large population response at single trial resolution will allow for looking at spike count correlations between neurons and evidence for sharpening of the response (increased sparseness) which will provide more evidence for or against models such as

the sharpening model or correlation-based models. The use of chronic implants will allow for the use of many more individual stimuli and exposures than are possible with acute penetrations and therefore much more of the possible stimulus space can be explored and many more lags and timescales of suppression can be tested.

5.10.3 Synaptic mechanisms of repetition suppression

What are the synaptic mechanisms involved in repetition suppression? Are they driven by synaptic depression, LTD in excitatory neurons, LTP in inhibitory neurons, or some combination of the above or other mechanisms? A full understanding of these mechanisms will likely require the combination of multiple experimental approaches. The combination of genetic tools which allow for efficient in vivo labeling of neuronal subtypes in combination with in vivo activity imaging such as 2-photon microscopy, the use of activity indicators such as c-fos or microscopy guided electrophysiology which could read out the activity in the known labeled neurons will give valuable information about the cell types involved and specifically what role inhibitory neurons play in repetition suppression. The use of highly targeted synthetic receptor agonists and antagonists such as 3-(2-carboxypiperazin-4-yl)-propyl-1-phosphonic acid (CPP) which can selectively block NMDA mediated LTP (*33*) will also likely be critical for understanding the synaptic mechanisms of repetition suppression. All of this must also be done in awake and intact animals to maintain the standard repetition suppression paradigm and be able to draw inferences from these approaches back to previous work. These methods are currently best suited for use in rodents, but method development in marmosets means all these tools may soon be available in a primate model. While the use of genetic and molecular methods has been undergoing rapid development in recent years, systems neuroscience has pointed the way towards elucidating synaptic mechanisms by providing inferences about when and where to look and what the most

likely candidates might be. It is the synergy of these approaches that will ultimately lead to a full understanding.

5.11 Concluding remarks

We live in a chaotic world, and it's possible our nervous systems have evolved in a state where novelty and unpredictability are the norm. The question of how the brain responds in the relatively special case when the same stimulus is encountered again then becomes an interesting one. How is this information then used to better guide behavior, presumably the goal of all sensory processing? It is clear that stimulus patterns that repeat are indeed processed differently from more novel patterns. Much progress has been made in terms of understanding the basic properties of how the brain responds to consistent input, and yet despite our current understanding of repetition suppression it is clear that more work remains to be done. Repetition suppression appears to be a ubiquitous process that occurs throughout the cortex if not the entire nervous system. It is a complex phenomenon and cannot be explained by simple mechanisms. Given this complexity it is likely due to multiple underlying mechanisms at different levels of processing within the nervous system, only some of which are beginning to be elucidated. Its functional consequences are not well understood although there are competing theories for how it could improve sensory representations and guide behavior. The work presented here has pushed our understanding of repetition suppression further, and hopefully will light the way forward for other fruitful explorations in the future of the questions which still remain. Reaching a more complete understanding of repetition suppression has been a challenge for neuroscience but given that it appears to be such a fundamental aspect of sensory processing, unlocking its secrets should provide new insights into the fundamental computational principals that govern the brain.

5.12 REFERNECES

1. Z. Kourtzi, N. Kanwisher, Representation of perceived object shape by the human lateral occipital complex. *Science* **293**, 1506-1509 (2001).
2. W. De Baene, R. Vogels, Effects of Adaptation on the Stimulus Selectivity of Macaque Inferior Temporal Spiking Activity and Local Field Potentials. *Cerebral Cortex* **20**, 2145-2165 (2010).
3. G. C. Baylis, E. T. Rolls, Responses of neurons in the inferior temporal cortex in short term and serial recognition memory tasks. *Exp Brain Res* **65**, 614-622 (1987).
4. M. E. Nagy, M. D. Rugg, Modulation of Event-Related potentials by word repetition: The effects of Inter-Item lag. *Psychophysiology* **26**, 431-436 (1989).
5. A. Todorovic, F. van Ede, E. Maris, F. P. de Lange, Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *Journal of Neuroscience* **31**, 9118-9123 (2011).
6. A. F. d. C. Hamilton, S. T. Grafton, Repetition suppression for performed hand gestures revealed by fMRI. *Hum Brain Mapp* **30**, 2898-2906 (2009).
7. M. W. Brown, F. A. Wilson, I. P. Riches, Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res* **409**, 158-162 (1987).
8. K. S. Saleem, N. K. Logothetis, *Atlas of the Rhesus Monkey Brain in Stereotaxic Coordinates*. (Academic Press, Amsterdam, 2007).
9. H. Nakamura, R. Gattass, R. Desimone, L. G. Ungerleider, The modular organization of projections from areas V1 and V2 to areas V4 and TEO in macaques. *J Neurosci* **13**, 3681-3691 (1993).
10. R. Gattass, A. P. Sousa, M. Mishkin, L. G. Ungerleider, Cortical projections of area V2 in the macaque. *Cerebral cortex (New York, N.Y. : 1991)* **7**, 110-129 (1997).
11. J. S. Baizer, L. G. Ungerleider, R. Desimone, Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J Neurosci* **11**, 168-190 (1991).
12. A. Kohn, Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. *Journal of Neurophysiology* **97**, 3155-3164 (2007).
13. S. G. Solomon, A. Kohn, Moving sensory adaptation beyond suppressive effects in single neurons. *Current biology : CB* **24**, R1012-1022 (2014).
14. F. Fabbrini, R. Vogels, Within- and between-hemifield generalization of repetition suppression in inferior temporal cortex. *J Neurophysiol* **125**, 120-139 (2021).
15. A. Lueschow, E. K. Miller, R. Desimone, Inferior Temporal Mechanisms for Invariant Object Recognition. *Cerebral Cortex* **4**, 523-531 (1994).
16. P. Kuravi, R. Vogels, Effect of adapter duration on repetition suppression in inferior temporal cortex. *Sci Rep* **7**, 3162 (2017).
17. Y. Liu, S. O. Murray, B. Jagadeesh, Time Course and Stimulus Dependence of Repetition-Induced Response Suppression in Inferotemporal Cortex. *Journal of Neurophysiology* **101**, 418-436 (2009).
18. K. Vinken, H. P. Op de Beeck, R. Vogels, Face Repetition Probability Does Not Affect Repetition Suppression in Macaque Inferotemporal Cortex. *J Neurosci* **38**, 7492-7504 (2018).
19. H. Sawamura, G. A. Orban, R. Vogels, Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. *Neuron* **49**, 307-318 (2006).
20. D. A. Kaliukhovich, R. Vogels, Divisive Normalization Predicts Adaptation-Induced Response Changes in Macaque Inferior Temporal Cortex. *J Neurosci* **36**, 6116-6128 (2016).

21. D. A. Kaliukhovich, R. Vogels, Stimulus repetition affects both strength and synchrony of macaque inferior temporal cortical activity. *J Neurophysiol* **107**, 3509-3527 (2012).
22. A. Jeewajee, C. Lever, S. Burton, J. O'Keefe, N. Burgess, Environmental novelty is signaled by reduction of the hippocampal theta frequency. *Hippocampus* **18**, 340-348 (2008).
23. B. C. Lega, J. Jacobs, M. Kahana, Human hippocampal theta oscillations and the formation of episodic memories. *Hippocampus* **22**, 748-761 (2012).
24. J. Winson, Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science* **201**, 160-163 (1978).
25. G. Huang, S. Ramachandran, T. S. Lee, C. R. Olson, Neural Correlate of Visual Familiarity in Macaque Area V2. *J Neurosci* **38**, 8967-8975 (2018).
26. L. Li, E. K. Miller, R. Desimone, The representation of stimulus familiarity in anterior inferior temporal cortex. *J Neurophysiol* **69**, 1918-1929 (1993).
27. E. K. Miller, L. Li, R. Desimone, Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* **13**, 1460-1478 (1993).
28. K. Grill-Spector, R. Henson, A. Martin, Repetition and the brain: neural models of stimulus-specific effects. *Trends in cognitive sciences* **10**, 14-23 (2006).
29. A. Kohn, J. A. Movshon, Adaptation changes the direction tuning of macaque MT neurons. *Nat Neurosci* **7**, 764-772 (2004).
30. H. Barlow, P. Foldiak, The computing neuron. *Adaptation and decorrelation in the cortex*, 54-72 (1989).
31. K. De Meyer, M. W. Spratling, A model of non-linear interactions between cortical top-down and horizontal connections explains the attentional gating of collinear facilitation. *Vision research* **49**, 553-568 (2009).
32. T. Hosoya, S. A. Baccus, M. Meister, Dynamic predictive coding by the retina. *Nature* **436**, 71-77 (2005).
33. E. W. Harris, A. H. Ganong, D. T. Monaghan, J. C. Watkins, C. W. Cotman, Action of 3-((+/-)-2-carboxypiperazin-4-yl)-propyl-1-phosphonic acid (CPP): a new and highly potent antagonist of N-methyl-D-aspartate receptors in the hippocampus. *Brain Res* **382**, 174-177 (1986).