Title

Mechanisms of visual attention, and their relationships to expectation and navigation, in the human brain

Author

Slama, Sara Julia Katarina

Publication Date

2019

Peer reviewed|Thesis/dissertation

Mechanisms of visual attention, and their relationships to expectation and navigation, in the human brain

By

Sara Julia Katarina Slama

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Robert T. Knight, Chair
Professor Frederic E. Theunissen
Professor Michael A. Silver
Professor Aditya Guntuboyina

Fall 2019

Mechanisms of visual attention, and their relationships to expectation and navigation, in the human brain

Abstract

Mechanisms of visual attention, and their relationships to expectation and navigation, in the human brain

by

Sara Julia Katarina Slama

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor Robert T. Knight, Chair

The amount of sensory input received by the human brain far surpasses its capacity for conscious processing. Attention serves to gate incoming inputs to enable the receiver to select aspects of the environment for further processing based on its relevance to higher-level goals. When attention is impaired, behavior goes awry, and this often has serious consequences. While it is clear that attentional processes are critical for humans' ability to interact effectively with the environment, the lack of a clear and conclusive definition of attention itself has, at times, impeded research progress on attention in the neural and behavioral sciences.

In this dissertation, I examine the construct of human attention from both empirical and theoretical standpoints. In Chapter 1, I provide an introduction to the concept of attention, including diverging definitions. I also articulate the urgency of obtaining a clearer understanding of the neurocognitive processes that support attention. Chapter 2 presents an empirical study of a specific type of attention, known as visual search. The study uses rare intracranial recordings in humans to present evidence consistent with the view that visual search, a classical attention behavior, may need to be reinterpreted as navigation in visual space. In Chapter 3, I present a theoretical discussion of the distinctions and similarities between attention and expectation, including interpretational pitfalls in experimental designs aiming to directly compare the two. Chapter 4 provides closing thoughts with emphasis on future directions for this field of research.

Table of Contents

Acknowledgments

I thank my advisor, Bob Knight, for mentoring me toward scientific and personal success. Bob is a true renaissance advisor, who can point to scientific literature from the 1930s in one breath and offer advice about startup runway in the next. Bob sits down with his mentees to personally inspect intracranial data for epileptic artifacts, and he takes his lab camping every year. Working with Bob has been a true privilege, and I am deeply grateful. I also thank Dr. Donatella Scabini for imprinting a culture of respect for research participants, and care for colleagues, on our lab.

I thank my dissertation committee. The faculty member aside from Bob, who has devoted the most attention to my growth as a scientist and as a person, is Michael Silver. Michael also provided a wealth of constructive feedback on an earlier version of this dissertation, which made the current manuscript much stronger. I thank Frederic Theunissen for data analytic guidance and for his friendly style of being. I thank Aditya Guntuboyina for clever and kind advice throughout the years.

I thank the patients for their generous participation in my research study (Chapter 2). I thank Lisa Johnson, Arjen Stolk, Julia Kam, Randolph Helfrich, Yvonne Fonken, Francine Foo, Chris Holdgraf, Anaïs Llorens, and Colin Hoy for data collection. I thank Sujayam Saha for his intelligent and positive way of explaining statistical concepts.

I thank the entire Knight Lab from 2013 to 2019. I thank Anaïs Llorens for enabling me to zoom out on a scientific level, and for being a perceptive and caring friend on a personal level. I thank my podmate Athina for friendship and scientific advice. I thank Colin Hoy for thoughtful conversations about neuroethics, and for being an unspoken friend and ally. I thank Chris Holdgraf for supporting my python learning endeavors and for sharing valuable, and occasionally entertaining, personal advice. I thank Yvonne Fonken for her friendship, and for encouraging me to join the Knight lab. I thank Boaz Sadeh for being a kind and attentive listener to a near stranger early in my PhD. I thank my research assistants, Aysja Johnson, Meron Yemane, Jerren Chiu, Juliet Ernst, and Angela Tseng: I appreciate the work they did with me, and I am proud of their subsequent accomplishments.

I thank Afresh Technologies Inc. and the Open Philanthropy Project for mentoring in external consulting projects. I especially thank Dr. Volodymyr Kuleshov at Afresh Technologies for feedback pertaining to my machine learning (ML) work: Lessons learned from this experience fed directly back into my dissertation work in Chapters 2 and 4. Through making cost estimates for the Open Philanthropy Project, I acquired skills that enabled me to estimate the societal costs of attention deficits in Chapter 1.

I thank the HWNI community. I am especially grateful to Kati Markowitz, Tony Leonard, Candace Groskreutz, and Natalie Terranova for their hard work and personal dedication. I am grateful to Joni Wallis, Bruno Olshausen, Ken Nakayama, Mike DeWeese, Jesse Livezey, Mike Schachter, Mark D'Esposito, and many others, for valuable scientific conversations during chance encounters on walks across campus. I am grateful to my friends: Falk Lieder, for teaching me how to drive a car without knowing how to drive himself; Vasha Dutell, for being a great scientist and a great human; Natalia Bilenko, for inspiring me to come to Berkeley and for being a

no-nonsense friend; and Melissa Newton, for applying her intellect and maturity to solve many a personal conundrum.

I thank the communities that I was never officially a part of: I thank the Redwood Center, spearheaded by Bruno Olshausen, for intellectual enrichment, friendship, and community. I thank Charles Frye for the surreptitious signal processing tutorials. I thank Pentti Kanerva - for whom no amount of knowledge or accomplishments can interfere with his unshakable Finnish humility - for his friendship. I thank Robert Nishihara and his crowd of intellectually exciting and “to the point” friends for their company. I thank Yifan Wu, for hosting intellectually adventurous identi-teas, and for her honest friendship.

I thank my mother for her intelligence and spirit of defiance - and for the many years of hard work that she put into being a mother. I thank my father for his warm and humorous ability to “see” human beings as they are. I thank my brother Andreas for tutoring me in mathematical reasoning through his “mammoth phone calls”. I thank my sister Anna for applying her sharp intellect to entirely different questions: One can never get away with being insincere in her company. I thank my brother Kalle for his authenticity and deep insight. I thank my sister Susanna for her ruthless kindness. I thank my aunt Tina for being a patient listener to a growing bonus daughter. I thank my cousin Alina for being the (not so) little sister I always wanted.

I thank my housemates, Tilia, JP, Rowan, and Jim, for their support and patience in sharing their space with a PhD student in the midst of dissertation writing. I thank my previous housemates, Debbie, Pit, and Audrey, who made a gargantuan happy footprint on several years of my life. I thank Devon and Tobias for their quiet demeanor, and for bringing colorful thought-provoking adventures home. I thank my bonus housemate, Jarrod, for all the coffee, steak, and great conversations.

I thank Lucy (Luce) Stephenson for offering to read “1,700 pages” of this dissertation, and for actually reading the entire work. Wherever the grammar and punctuation are correct, I owe Luce my thanks. Any errors are my own.

I thank my partner, Michael (Meick) Eickenberg, who understands everything I share with him about science, but who insists on meeting me as a human being. When I want to dedicate my every thought to “serious pursuits”, Meick reminds me of my joy in knitting and singing, and makes me watch uncountable youtube videos.

I thank my voice teacher, Marika Kuzma, for helping me find my voice; my therapist, Jean, for enabling me to grow even through the middle years of my PhD; Yoga to the People, where I have practiced more days than not over my years in Berkeley; my yoga teacher, Nichol Chase, whose lessons make me feel capable of tackling new and hard problems; and the Berkeley Buddhist Monastery, for providing a serene, quiet space for meditation.

Chapter 1

Introduction: The brain, vision, attention, what it's all good for, and how my graduate research fits into the story

1.1. The brain and vision

"*Jo jo, jag vet, jag har ju varit på undersökning. Men, fast dom tog så många bilder så, int' hitta dom nånting. Så, jag tänkte att dom får nu fortsätta att snurra på därinne, bäst dom vill.*"

Karin Klåvus, konstnär, Närpes, Finland,
i svar på min fars utsago att hans dotter
studerar hjärnor i Amerika.

"*Yes yes, I know all about it. You see, I have been examined [neurologically]. But, even though they took so many pictures, they didn't find anything at all. So I figured that, I'll let them [the neurons] continue twirling about in there, as they wish."*

Karin Klåvus, Artist, Närpes, Finland, upon learning that Michael Slama's daughter studies brains in America.

The adult human brain contains approximately 86 billion neurons (Azevedo et al., 2009). It measures about 14 cm from ear to ear, and 17 cm from forehead to inion. As early as the 5th century B.C., the Greek physician Hippocrates realized that the brain is responsible for implementing our sensory percepts, emotions, and thoughts (Kandel et al., 2013).

Ever since the seminal work by two pioneers of neuroscience in the mid-20th century (Hubel & Wiesel, 1959; 1962; 1968), our understanding of the sense of *vision* has disproportionately informed our understanding of the brain. Because (1) the neural mechanisms of early visual cortex are relatively well understood, and (2) many neural computations generalize across cortical regions, the human visual system constitutes an excellent model for understanding brain function and the relevance of these mechanisms to disease (Yoon et al., 2013).

The strides made in our understanding of the early visual pathways make it tempting to believe that vision is now "solved". This belief is additionally fueled by advances in computer vision and artificial intelligence using architectures inspired by the human visual system (Krizhevsky et al., 2012; Szegedy et al., 2015; Simonyan & Zisserman, 2015; He et al., 2016). We may believe that we understand the mechanics of human vision since we can build machines whose visual behavior, to some extent, mimics that of humans. Yet it is important to recall that the human brain remains able to accomplish feats which machines cannot (Marcus, 2018). One example is object recognition generalized to different viewpoints (Pinto et al., 2011; DiCarlo et al., 2012; Marcus, 2018). It is also not true that we understand the mechanics of human vision, since even the function of cortical area V1 leaves much variance to be explained (Olshausen & Field, 2004). Finally, while most computational as well as neural models

of visual perception have a strictly feedforward architecture, a number of challenges to a simple serial model of vision have been raised over the past decades. One challenge lies in the fact that anatomical feedback projections to V1 greatly outnumber the number of feedforward projections (Felleman & Van Essen, 1991; Angelucci & Bullier, 2003).

Attention stands out as one aspect of human visual perception where much work remains to be done. What is attention, anyway?

1.2. Attention

"*Everyone knows what attention is.*"

William James (1890)

"*There is no such thing as attention.*"

Britt Anderson (2011)

Psychologists and neuroscientists alike have struggled to agree on a consistent definition of attention. Many argue that William James' definition from 1890 remains the clearest and most valid description of the phenomenon. He writes: "It is the taking possession of the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalisation, concentration of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others."

More recently, Chun and colleagues (2011) proposed an updated definition: "Given limited capacity to process competing options, attentional mechanisms select, modulate, and sustain focus on information most relevant for behavior." Importantly, Chun et al. (2011) introduce a *taxonomy* of attention by proposing a division of the concept into several subcomponents. At the broadest level, this taxonomy divides attention into (1) attention directed towards the external world versus (2) attention directed internally towards thoughts. Chun et al. (2011) further subdivide attention by sensory modality (e.g., visual, auditory, tactile, olfactory, and cross-modal), and by the properties of the attentional process itself (e.g., sustained, feature-based, and object-based attention). The topic of **Chapter 2** of this dissertation is a specific case of feature-based attention known as *visual search*.

Another recent push toward greater precision in the definition of attention has been proposed by Summerfield and Egner (2009). They highlight the distinct conceptual meaning of *attention* versus *expectation*. Attention, Summerfield & Egner argue, is a prioritization process that serves to select what is *important* (has "motivational relevance") in the external environment. Expectation, they posit, is a *prediction* process pertaining to what the subject believes is going to happen next ("prior likelihood"). In their definition, expectation is neutral to the importance of the expected object. Summerfield and Egner (2009) observe that with the burgeoning field of predictive coding (Huang & Rao, 2011), there is a set of research experiments which - based on experimental design alone - could be interpreted as engendering *either* an attention or an expectation process. According to Summerfield and Egner, a number of classical

experiments in the attention literature should be reinterpreted as *expectation* rather than *attention* experiments, or potentially as encompassing a component of both processes. These include any classical experiments where a cue instructs a subject to attend to a given spatial location in anticipation of an upcoming target (for example the Posner cueing paradigm; Posner, 1980)

This theoretical reinterpretation from *attention* to *expectation* opens the door for a new field of research, which aims to characterize *expectation* processes in equivalent amount of detail as *attention* processes have been characterized over the past decades. It appears that research results on expectation are beginning to parallel the very taxonomy of attention proposed by Chun et al. (2011), including expectation by sensory modality, spatial expectation, feature-based expectation (Summerfield & Egner, 2016); object-based expectation (Jiang et al., 2016), etc. Unfortunately, this labor will in many cases be redundant to the work that has already been conducted in the more mature field of attention research. In light of Summerfield and Egner's insights regarding the multiple processes elicited by common attention paradigms, a more fruitful way forward might be *reinterpretation* rather than *repetition* of existing results in the attention literature. In **Chapter 3** of this dissertation, I discuss a specific research study on object-based expectation (Jiang et al., 2016), and its interpretational pitfalls in comparison to the closely related field of object-based attention.

While we may not yet fully understand how to characterize and subdivide the heterogeneous construct of attention, it is clear that the set of neural processes collectively referred to as attention are of immense practical importance to the everyday lives of the 7.35 billion people who live in the world. Impairments in attention constitute some of the most debilitating symptoms of several brain disease, including: Alzheimer's disease (Li et al., 2012), Parkinson disease (Yarnall et al., 2011), stroke (Corbetta & Shulman, 2011), schizophrenia (Morris et al., 2013), major depression (Marazziti et al., 2010), and ADD/ADHD. Healthy individuals are also affected by fluctuations in attentional function: Subclinical lapses in attention cause lost productivity in school and at the workplace, and driving distracted carries a further cost of lives lost as a direct result of deadly attention lapses.

The annual cost of ADHD in the United States alone has been estimated to range between $143 and $266 billion (Doshi et al., 2012). This amounts to approximately 1% of the U.S. GDP of $19 trillion. In Europe, the cost of *childhood* ADHD alone, for the year 2010, was estimated at €2.5 billion euros ($2.8 billion). The total cost of all brain disorders was estimated at €798 billion ($890 billion, Olesen et al., 2012), corresponding to nearly 5% of GDP in the European Union. These estimates are associated with large error bars. However, it is clear that they are a gross *underestimate* of the net cost of attention impairments in that they exclude all clinical impairments beyond ADHD, all subclinical attentional impairments, and mismanaged attention in healthy people. Estimates for the financial burden of ADHD and other attentional impairments are more sparse for developing countries, and the prevalence of formal diagnoses and treatment availability varies. It is clear, however, that the prevalence of the disability itself (ADHD), does not respect national borders (de Graaf et al., 2008). These estimates demonstrate that a lower bound on the financial burden of impaired attention is on the order of billions of dollars per year.

This dissertation represents two steps toward greater clarity in our understanding of human attentional processes. **Chapter 2** is an empirical study, which suggests that some attentional processes may need to be reinterpreted as navigation behaviors. **Chapter 3** is a theoretical discussion paper, highlighting pitfalls in interpreting experimental results aimed at disentangling expectation and attention.

Chapter 2

Electrophysiological mechanisms of visual search in humans

2.0. Foreword

In this chapter, I describe the major experimental and data analysis work that I conducted during my PhD, providing novel results on the implementation of attention processes in the human brain. Richard Jimenez created the visualizations of electrode sites showing significant effects in Figures 2.11. and 2.15. I developed the statistical testing approach for identifying task-active and condition-selective electrode sites in collaboration with Sujayam Saha, based on previous work by Maris & Oostenveld (2007). Professor Robert Knight, Richard Jimenez, and I collaboratively conducted electrode localization of hundreds of electrodes by visual inspection: I am especially indebted to Professor Knight for this work, since it seems rare that a professor of his stature would be willing to sit down and do that kind of hard work with his mentees. I also owe thanks to Avgusta Shestyuk for advice regarding exploratory data analysis and data integrity checking, and to James Lubell for assistance with experiment programming. The data was collected at three sites around the world, and several researchers and clinicians contributed to that endeavor:

1. Oslo University Hospital, Norway: Anne-Kristin Solbakk, Tor Endestad, & Pål G. Larsson.
2. California Pacific Medical Center, San Francisco, USA: David King-Stephens, Kenneth D. Laxer, & Peter B. Weber.
3. UC Irvine, USA: Jack J. Lin.

Finally, I am infinitely grateful to the patients, who donated their time and data to make this research possible.

2.1. Abstract

Visual search is a fundamental human behavior, which has been proposed to include two component processes: inefficient search (Search) and efficient search (Pop-out). According to extant research, these two processes map onto two separable neural systems located in the frontal and parietal association cortices. In the present study, we use intracranial recordings from 23 participants to delineate the neural correlates of Search and Pop-out with an unprecedented combination of spatiotemporal resolution and coverage across cortical and subcortical structures. First, we demonstrate a role for the medial temporal lobe in visual search, on par with engagement in frontal and parietal association cortex. Second, we show a gradient of increasing engagement over anatomical space from dorsal to ventral lateral frontal cortex. Third, we replicate a previous study demonstrating nearly complete overlap in neural engagement across Search and Pop-out. We further demonstrate Pop-out selectivity (greater activity increases in Pop-out as compared to Search) in a distributed set of sites including

frontal cortex, contrary to the view that visual pop-out is implemented in low-level visual cortex or parietal cortex alone. Finally, we affirm a central role for the right lateral frontal cortex in Search.

2.2. Introduction

Visual search is is ubiquitous in everyday life, and is deployed in everything from driving to reading to airport security screening. Impairments in visual search ability have been documented in numerous brain disease including Alzheimer's disease, Parkinson disease, stroke, schizophrenia, and ADHD. Visual search is widely considered a classical attention behavior (Treisman & Gelade,1980; Wolfe, 2014; 2018). Accordingly, some search processes are thought to require deliberate allocation of attention to individual putative search targets in a serial fashion (inefficient search, hereafter referred to as Search), with response times (RTs) increasing with the number of distractors in the search display. In other search processes, attention is thought to be automatically captured by a salient item in a display (efficient search, hereafter referred to as Pop-out), rendering search times independent of the number of distractors. Pop-out is thought to rely on preattentive visual perceptual mechanisms which scan the entire visual field in parallel (Treisman & Gelade, 1980**;** Julesz, 1981; Wolfe, 2018; but see Nakayama & Martini, 2011). Visual search experiments overall are viewed as well-parameterized *attention* tasks. Depending on the design of the specific experiment, Search and Pop-out are thought to elicit attentive and pre-attentive processes, respectively.

<u>Putative medial temporal lobe engagement in Search and Pop-out</u>
The dominant model for the neural substrates of visual attention derives from human fMRI and PET studies as well as single-neuron research in non-human primates (Corbetta & Shulman, 2002). According to this framework, a bilateral dorsal network (the dorsal attention network, DAN) supports top-down attention, as deployed in Search. A right-lateralized ventral attention network (the ventral attention network, VAN) is implicated in bottom-up attention, as deployed in Pop-out (Corbetta & Shulman, 2002; Corbetta, Patel & Shulman, 2008). While these two networks are anatomically distinct, the main nodes of both networks are located in the frontal and parietal association cortices.

A variant of this frontoparietal attention framework derives from non-invasive EEG recordings in humans (Li et al., 2010) and from single-neuron recordings in non-human primates (Buschman & Miller, 2007). This work focuses on the respective roles of the frontal and parietal association cortices in Search versus Pop-out. The central prediction that follows from these studies is that Search preferentially engages frontal cortex whereas Pop-out preferentially engages parietal cortex.

A challenge to a pure frontoparietal model of visual search has recently been introduced in a series of publications in the non-human primate neuroscience field. These studies demonstrate that visual search of natural scenes rely extensively on medial temporal lobe (MTL) structures such as the entorhinal cortex (Killian et al., 2012; 2015). The explanation, according to Killian and colleagues, is that visual search

belongs to the family of *navigation* behaviors, not merely visual attention behaviors. In other words, they claim, visual search in primates can be conceptualized as navigation in visual space, analogous to navigation in physical space in rodents (Meister & Buffalo, 2016). Consistent with this, a recent study examining single-neuron responses in the human hippocampus and amygdala reported that individual neurons in these regions reflect target detection processes during visual search in natural images (Wang et al., 2018).

A second reason to predict MTL engagement in visual search, and especially in the visual pop-out phenomenon, stems from its association with novelty detection. Hedwig von Restorff discovered the relationship between novelty and memory nearly 90 years ago (von Restorff, 1933). Subsequent findings established that MTL memory structures are necessary for the detection and recollection of novel items (Knight, 1996; Parker et al., 1998). These findings raise the question of whether classical, well-parameterized experiments targeting Search and Pop-out also engage MTL structures.

Direct evidence for MTL engagement in Search and Pop-out in humans is lacking. This may be due to a general neglect of the contributions of these structures to attention in the cognitive neuroscience literature. In addition, the reduced signal-to-noise ratio in the MTL in fMRI recordings in humans impairs researchers' ability to draw conclusions about putative MTL involvement in various cognitive tasks. In this study, we address the question of MTL involvement in Search and Pop-out using direct intracranial recordings in humans.

<u>Do Search and Pop-out map neatly onto the frontal and parietal cortices, respectively?</u>
Returning to the question of Search and Pop-out as two separable perceptual processes, the idea that humans possess two visual systems goes back at least four decades (Treisman & Gelade, 1980; Julesz, 1981). These two putative perceptual systems have been proposed to map onto two distinct neural systems. Several models of the anatomical substrates for these two putative systems exist: I discuss three of these models below.

First, a set of findings from non-invasive EEG in humans, and electrophysiology in primates, suggest that Pop-out is implemented in parietal cortex (Li et al., 2010**;** Buschman & Miller, 2007, but see Nobre et al., 2002). This is consistent with reported salience maps in area LIP in monkeys (Gottlieb et al., 1998) and a broader role for parietal cortex in detecting external visual stimuli, holding such information in memory, and generating motor action plans (Mangano et al., 2015; Xu, 2018; Regev et al., 2018; Martin et al., 2019). By contrast, prefrontal cortex (PFC), including FEF and lateral frontal cortex are implicated in Search (Leonards et al., 2000; Buschman & Miller, 2007; Rossi et al., 2007; Li et al., 2010). Consistent with this view, these areas have been shown to contain topographic visual maps (Kastner et al., 2007; Silver & Kastner, 2009; Mackey et al., 2017).

Second, Corbetta's and Shulman's (2002; Corbetta et al., 2008) model of visual attention, as mentioned above, focuses on *networks* for goal-directed versus stimulus-driven attention, incorporating subregions of several cortical lobes. They propose that the DAN, which includes the intraparietal sulcus (IPs) and the frontal eye

fields (FEF), supports goal-directed attention as deployed in Search. According to Corbetta's and Shulman's framework, the VAN - including the temporoparietal junction (TPJ), parts of the middle frontal gyrus (MFG), inferior frontal gyrus (IFG), frontal operculum, and anterior insula - are involved in stimulus-driven attention as deployed in Pop-out.

Third, at least two theoretical models exist for the implementation of Pop-out. One model claims that the central mechanism enabling the visual pop-out phenomenon is computation of salience in early visual cortex, most notably V1 (Zhaoping, 2002; Zhaoping & Dayan, 2006; Zhaoping, 2019). An opposing model claims that Pop-out is, instead, a top-down phenomenon. This gives rise to the prediction that there should be neural correlates of Pop-out in higher-order cortical areas such as frontal cortex (Hochstein & Ahissar, 2002), consistent with the observation that PFC lesions causes decreased EEG novelty responses in humans (Knight, 1984).

As a fourth perspective, recent empirical evidence emphasizes the similarities rather than the differences between the tasks. Both fMRI work (Leonards et al., 2000) and a single intracranial study that has been conducted on the topic (Ossandon et al., 2012) suggest that the networks supporting Search and Pop-out are remarkably similar, and centered on fronto-parietal regions including the DAN. Only minute, local differences such as greater activation in lateral frontal cortex in Search were observed (Leonards et al., 2000), perhaps reflecting increased working memory demands. The fact remains, however, that Search versus Pop-out engender robust behavioral differences, and the origin of these differences must reside somewhere in the brain.

In the present study, we employ direct intracranial recordings of neural activity to map the anatomical sites at which Search and Pop-out converge and diverge in the human brain. Intracranial recording provides a method with improved spatiotemporal resolution compared to non-invasive neural recording methods in humans (Parvizi & Kastner, 2018). This superior temporal resolution enables us to observe and reject artifacts driven by RT differences across attention conditions, which may have confounded the results in previous studies using lower-resolution recording methods. We analyzed intracranially recorded voltage signals recorded from 1,321 electrocorticography (ECoG) grid and stereoelectroencephalography (SEEG) depth electrodes in 23 patients with medically refractory epilepsy, each of whom were undergoing diagnostic recording in preparation for potential resective surgery. Across patients, the sensors captured extensive areas of cortex as well as the MTL and the amygdala. This combination of high spatiotemporal resolution and broad coverage makes this study uniquely positioned to document where and how these two modes of searching our visual environment differ. To the extent that the tasks differ, we ask in which regions of the brain this occurs, and assess to what extent these differences overlap with existing models of the implementation of Search and Pop-out in the brain.

The patients completed a simple visual search task (Figure 2.10.a.), with two experimental conditions, Search and Pop-out. We focused on task-related activity increases, and condition-related modulations, in high-frequency activity (HFA; 80-150 Hz). HFA has been shown to correlate both with the fMRI BOLD signal (Logothetis et al., 2001; Mukamel et al., 2005; Nir et al., 2007) and with multi-neuron activity (Ray et al., 2008; Ray & Maunsell, 2011). The findings demonstrate a robust role for the MTL in

Search and Pop-out, and highlight sub-regional differences within the PFC, including different activation profiles between superior and inferior lateral frontal cortex.

2.3. Method

Participants

Thirty participants were enrolled in the study while undergoing treatment for medically refractory epilepsy at the Oslo University Hospital, Norway (N = 8), California Pacific Medical Center, USA (N = 7), and UC Irvine Medical Center, USA (N = 15). The Institutional Review Boards at each hospital as well as the Committee for Protection of Human Subjects at the University of California, Berkeley, approved the study procedures. Each patient provided written informed consent prior to participation. Electrode placement was determined solely by clinical needs.

Seven patients' datasets were excluded from the analyses. Two were excluded due to chance behavioral performance in the experiment. One was excluded for having fewer than 15 trials in one of the two experimental conditions, after applying trial exclusion criteria (see section *Removal of epochs* below). Two patients were excluded for showing no task-related activity beyond primary sensory areas in the recordings (see section *Task-active electrode selection: Identifying significant increases*), and two patients were excluded for excessive high-frequency noise.

Of the remaining 23 participants (35.04 ± 13.97 years; mean ± SD; 13 female), 21 were right-handed, and 2 were ambidextrous (see Table 2.1., Appendix A). For the patients who participated in the U.S. (n = 17), twelve were native English speakers; two were native Spanish speakers; one was bilingual in English and Spanish; one was an ASL speaker; and one was a native Vietnamese speaker. All patients whose data was collected at the Oslo University Hospital (n = 6) were native Norwegian speakers. Task instructions were provided in the participant's first language, either by the experimenters or with the help of a family member. Language was not a barrier to task comprehension in any of the subjects. All participants had normal or corrected to normal vision.

Stimuli

Participants completed a visual search task adapted from Li et al. (2010; 2013). The stimuli consisted of acute, isosceles triangles that were either red or green, and presented on a white background (Figure 2.10.a.) on a Windows laptop (15.6″ LCD screen). The stimuli were presented, and responses recorded, using E-Prime 2.0 software (Psychology Software Tools, Sharpsburg, PA).

The laptop was placed in front of participants at a viewing distance of 40-60 cm (approximately 16-24″). Each triangle had a base of 4.0 cm (3.8-5.7º visual angle) and a height of 4.5 cm (4.3-6.4º visual angle). On the stimulus display, the distance of each triangle from a central fixation cross was 6.0 cm (5.7-8.6º visual angle) horizontally and 4.0 cm vertically (3.8-5.7º visual angle), measured between the fixation cross and the center of each stimulus triangle.

To synchronize the neural and behavioral recordings for subsequent analyses, analog channels transmitted stimulus onset and offset signals from the stimulus

computer to the neural recording hardware. At the hospitals in the U.S., a photodiode was used to detect light changes on the stimulus monitor, and at Oslo University Hospital, a TRS ("Tip, Ring, and Sleeve", see: https://missionengineering.com/what-is-a-trs-cable/) connector transmitted an audio signal from the stimulus computer to the recording rig.

Participants were instructed to locate a target triangle among four candidate triangles on the screen. The target triangle was defined by a color (red or green) and an orientation (1 of 8 orientations: 0º, 45º, 90º, 135º, 180º, 225º, 270º, or 315º). Patients indicated whether they found the target triangle on the left or right half of the screen by pressing the left versus right arrow key on the laptop keyboard, or the left versus right button of an external mouse, depending on the physical constraints of the recording room.

The trial sequence started with a green fixation cross (1000 ms, *Inter-Trial Interval*; notice that a 500-ms *Baseline Interval* was a subset of this interval), followed by the fixation cross turning black (500 ms, *Fixation Interval*). Next, the target triangle, which participants were required to hold in memory, was presented in the center of the screen (1000 ms, *Sample Interval*), followed by a fixation cross (500 ms, *Working Memory Interval;* Eimer, 2014), and ultimately the stimulus display consisting of four triangles including the target triangle (*Stimulus Display Interval*). The trial was terminated after the response or, if the participant did not respond, after 2,000 ms.

Each experiment included a total of 128 trials, divided into four blocks of 32 trials each. The task included two experimental conditions, Search and Pop-out. In Search, all four triangles had the same color (all red or all green), while in Pop-out, the target triangle had a different color than the three distractor triangles (red target with three green distractors, or green target with three red distractors). Half the trials (n=64) were Search trials. They were randomly interspersed with Pop-out trials, so that participants could not anticipate which condition an upcoming trial would belong to. Trial order was randomized separately for each participant.

Behavior analysis

The behavioral outcome measures were response accuracy (proportion correct) and response time (RT, *ms*). To compare behavioral performance between the two conditions, we computed a paired *t*-statistic between percent accuracy in the Pop-Out versus Search conditions, excluding non-response (time-out, >2000 ms from stimulus onset) trials. We assessed significance using a permutation test, where we shuffled condition labels 100,000 times, and then we compared the *t*-statistics obtained in the permutation tests to the observed *t*-statistics. The *p*-value was computed as: (1 + number of permuted statistics more extreme than the observed statistic) / (1 + number of permutations). We calculated the same test to compare median RTs between the two conditions across subjects, including only correct-response trials.

Neural data acquisition

Intracranial electrodes (ECoG or SEEG) were implanted for approximately one week in each patient to determine the epileptogenic focus. In eight patients, ECoG arrays, organized either in two-dimensional grids or one-dimensional strips, were placed on the cortical surface. In fifteen patients, SEEG depth electrodes, targeting subcortical structures, were implanted. A total of 1,321 electrodes were analyzed across the 23 participants.

At the Oslo University Hospital, ECoG and SEEG data were recorded using two 64-channel NicoletOne (Natus Neuro Inc., Middleton, WI, USA) amplifiers (in four patients) and a 256-channel ATLAS (Neuralynx, Bozeman, MT, USA) digital acquisition system (in one patient). For the SEEG cases (n = 5), the exposed electrode diameter was 0.8 mm, and the center-to-center distance between electrodes was 3.5 mm. In four of these patients, the digitization rate was 512 Hz (NicoletOne system). In one patient, it was 16,000 Hz (ATLAS system). ECoG data were sampled at 1,024 Hz (NicoletOne system). The electrodes were manufactured by DIXI Medical (Besançon, France).

At CPMC, ECoG data were recorded using a Nihon Kohden (Tokyo, Japan) Neurofax EEG-1200 digital acquisition system with 128/256-channel amplifier capacity. The digitization rate was 1,000 Hz. All five patients at this site were implanted with surface contacts (grids and strips). The electrodes were manufactured by the Ad-Tech Medical Instrument Corporation (Oak Creek, WI, USA).

At UC Irvine Medical Center, data were recorded using a Nihon Kohden digital acquisition system with a 128/256-channel amplifier at a digitization rate of 1,000/5,000 Hz, for both ECoG and SEEG cases. For the ten patients implanted with SEEG depth electrodes, the diameter of the electrode contacts was 0.9 mm, and the inter-electrode distance was 5.0 mm. The electrodes were manufactured by Integra Life Sciences (Plainsboro, NJ, USA).

For all patients recorded with surface electrodes, the exposed electrode diameter was 2.3 mm. The inter-electrode spacing was 10.0 mm for all ECoG grids in all patients, except a single 8x8 high-density grid in one patient, where the center-to-center spacing was 4 mm.

Anatomical reconstruction and visualization

For electrode localization, we first segmented each patient’s preoperative T1-weighted MRI scan using Freesurfer 5.3.0 (Dale et al., 1999). Next, we fused the MRI image with a post-implantation CT scan using the Fieldtrip toolbox (Oostenveld et al., 2011; Stolk et al., 2018). To correct for the displacement of electrodes and brain tissue due to pressure changes related to the patient's craniotomy, electrodes were realigned to the preoperative cortical surface (Hermes et al. 2010, Dykstra et al., 2012). We inspected the individual fusions for maximal interlocking between the CT and MR visually for quality assurance. In the case of two patients, whose native-space MRIs were missing, we co-registered the post-implantation CT scans to the stereotaxic Colin-27 brain template (Collins et al., 1998).

For group-level visualization, surface-based normalizations were conducted using Freesurfer via cortical gyrification patterns, and volume-based normalizations through fusion to the Colin-27 brain via overall geometry patterns. By doing so, we linked patient brains to their template homologs, allowing accurate comparison of regions of interest (ROIs). In SEEG datasets, for which each physiological time-series was calculated using a bipolar montage between pairs of electrodes, we calculated the position of each electrode as the mid-point between the two original electrode locations (Burke et al., 2013; Jafarpour et al., 2019). To visualize electrodes showing task-related effects, we used FieldTrip and custom Matlab scripts.

Assignment of electrode locations to anatomically defined regions of interest

A neurologist (RTK) identified the locations of the individual electrodes based on inspection of the fused MR and CT images, displayed in the patient's native space, using the BioImage Suite toolbox.

We identified electrodes as belonging to one of the following eight broad ROIs (see Table 2.2., Appendix A): Frontal cortex (not including sensorimotor cortex), parietal cortex (not including sensorimotor cortex), sensorimotor cortex, temporal cortex, occipital cortex, cingulate cortex, medial temporal lobe (MTL), and amygdala. Sensorimotor cortex, which was engaged due to the behavioral response (a button press), was not of primary interest in this study, and was not further examined. We additionally examined activations in subregions of those ROIs that had greater than 50 electrodes (Table 2.2., Appendix A).

In frontal cortex (Table 2.3., Appendix A), we examined the following five subregions, defined by gyral and sulcal landmarks: (1) the superior frontal gyrus and sulcus, (2) the middle frontal gyrus, (3) the inferior frontal gyrus and sulcus, (4) the orbitofrontal cortex, and (5) the medial prefrontal cortex.

In parietal cortex (Table 2.4., Appendix A), we separately examined activations in: (1) the superior parietal lobule including the intraparietal sulcus, (2) the inferior parietal lobule comprised of the supramarginal gyrus, angular gyrus, and temporoparietal junction, and (3) the precuneus.

In temporal cortex (Table 2.5., Appendix A), we examined the following seven subregions: (1) the insula, (2) the superior temporal gyrus, (3) the superior temporal sulcus, (4) the middle temporal gyrus, (5) the inferior temporal gyrus including the middle temporal sulcus, (6) the ventral stream including the lingual and fusiform gyri, and (7) the temporal pole.

In cingulate cortex (Table 2.6., Appendix A), we examined (1) the anterior cingulate cortex, (2) the midcingulate cortex, and (3) the posterior cingulate cortex.

Finally, in the medial temporal lobe (Table 2.7., Appendix A), we examined the (1) the hippocampal formation (HF), including the hippocampus and subiculum, (2) parahippocampal cortex, and (3) entorhinal cortex, including perirhinal cortex.

Preprocessing of neural data

We recorded the local field potential from a total of 503 ECoG and 818 SEEG electrodes in the 23 participants included in the final study sample. In subjects with an original sampling rate of 5,000 or 10,000 Hz, we resampled the signal to 1,000 Hz. The resulting datasets had sampling rates of 512, 1,000, or 1,024 Hz. To remove slow drift and high-frequency noise respectively, we high-pass filtered the signal at 1 Hz, and low-pass filtered it at 180 Hz. To remove line noise, we notch-filtered the signal at 60 Hz and harmonics for datasets recorded in the U.S., and at 50 Hz and harmonics for datasets recorded in Norway.

Each recording was visually inspected by a neurologist (RTK) for epileptic activity or poor signal quality (such as detached electrodes or high-frequency noise). Electrodes that reflected signal from epileptic tissue or white matter, or were consistently noisy during the recording, were removed from the dataset: 16.3% of surface electrodes (which did not include white matter electrodes) and 55.4% of depth electrodes (which included white matter electrodes) were excluded. Temporal epochs that showed epileptic activity spread were removed across all electrodes. To remove remaining shared noise sources from the data in the surface (ECoG) datasets, we applied the common average reference (subtracting the point-by-point average signal of the preprocessed dataset from each time point of retained electrodes) to all grids within patient. In depth (SEEG) datasets, we applied a bipolar reference (pairwise subtraction of adjacent electrode time-series).

Spectral decomposition

We extracted the analytic amplitude of the HFA signal through three steps. First, we bandpass-filtered the time-series from the complete recording in each electrode between 80 and 150 Hz using a zero-phase FIR filter (mne.io.Raw.filter function from the MNE toolbox, https://martinos.org/mne/stable/generated/mne.io.Raw.html). We then computed the Hilbert transform of the filtered signal (mne.io.Raw.apply_hilbert function from the MNE toolbox) yielding a complex time-series, of which we took the absolute value to obtain the instantaneous analytic amplitude. Finally, we low-pass filtered the HFA analytic amplitude at 10 Hz to facilitate detection of temporal variation on a time-scale of approximately 100 ms, following Haller et al. (2018).

Baseline normalization

To examine the relationship between the behavioral and neural data, we segmented the full recording time-series into trial epochs, time-locked to the onset of the sample triangle (start of the *Sample* interval). A full trial epoch ranged from 1,050 ms before the onset of the sample triangle to 4,000 ms after. We defined the *Baseline Interval* as the first 500 ms of the trial, i.e., -1,050 to - 550 ms relative to the onset of the sample triangle. During the baseline, the participant watched a black fixation cross at the center of the monitor. We ended the trial epoch at 4,000 ms, because this captures the longest

trials plus a 500-ms period of response and post-response activity. To facilitate detection of task-related changes in neural activity, we normalized the HFA analytic amplitude to the neutral *Baseline Interval*, by computing an adapted z-score as follows: From each sample point of the full trial epoch, we subtracted the mean, and divided by the standard deviation, of the pooled baseline interval (across all included baseline epochs within condition and within each electrode).

We then separated trials belonging to the Pop-Out versus Search condition, and conducted all subsequent analyses separately within each condition.

Removal of epochs

We excluded epochs from the dataset according to the following criteria: incorrect behavioral response, no response prior to trial timeout, epileptic activity or consistent noise in the raw time-series, RT outliers, and HFA activity outliers. RT outliers were defined as trials where the RT was more than 3 interquartile ranges (IQR, i.e., the 75th percentile - 25th percentile) lower than the 25th percentile or more than 3 IQR higher than the 75th percentile of the subject's own RT.

In order to minimize the number of trials with HFA signal artifacts in electrodes of interest, while maximizing the total number of trials retained across a patient's dataset, we used a two-step approach for identifying HFA activity outliers. First, we computed the set of active electrodes (see section, *Task-active electrode selection: Identifying significant increases* and Figure 2.1. below) in each patient's dataset prior to HFA outlier exclusion. Next, we excluded trials that showed a max-min range greater than 6 IQR above the 75th percentile (or greater than 6 IQR below the 25th percentile) in any one active electrode. Trials were excluded across each patient's full dataset. The maximum number of trials that were excluded in any one patient based on this criterion was 8 across the dataset, and 4 in any one condition. The smallest number of trials retained in any one subject and condition after all trial exclusion criteria was 21. Finally, we computed the set of active electrodes a second time, on the data where HFA outliers had been removed. The max-min range was selected as a metric of deviation in order to capture the time course shape of the most common HFA signal artifacts (rapid, high-amplitude modulations). The threshold of 6 IQR was chosen based on examination of two patients' datasets, and then generalized across all patients.

Task-active electrode selection: Identifying significant increases

To select the subset of electrodes that showed task-relevant changes in HFA activity for further analyses, we used a permutation test approach adapted from Maris & Oostenveld (2007; Figure 2.1.). We considered an electrode that showed a significant task-related increase relative to baseline in either condition to be task-active. We also identified electrodes that showed task-related decreases using an equivalent procedure, but these effects are not the main focus of the current paper. To make the comparison, we selected two intervals ("task interval" below) of equal length to the baseline: (1) stimulus onset to 500 ms following stimulus onset, (2) 500 ms before response up to response. The significance calculation for each electrode was computed as follows:

(1) Baseline-normalize both baseline and task intervals as outlined in the section, *Baseline normalization*, above.
(2) For both the baseline and task intervals, compute a one-sample *t*-test across trials at each time point.
(3) Define a cluster as the set of *t*-statistics associated with any set of two or more consecutive significant time points ($p < 0.05$; see Figure 2.1., grey shaded area in the center panels depicting *t*-statistics). Select all clusters.
(4) Compute the sum of *t*-values in each cluster.
(5) Select the cluster with the largest sum of *t*-values (in both the baseline and the task intervals).
(6) Subtract the largest sum of *t*-values in the baseline interval from the largest sum of *t*-values in the task interval, to obtain an observed statistic for the electrode.
(7) Randomly permute the task versus baseline labels 1000 times, and repeat the calculations in steps 1-6 above.
(8) Compute a *p*-value for each electrode as: (1 + number of permuted statistics larger than the observed statistic) / (1 + number of permutations).
(9) An electrode with a *p*-value < 0.025 one-tailed (FDR-corrected across all electrodes in all subjects) in either of the two intervals in either condition was considered significant. We selected a cutoff of $p = 0.025$ to Bonferroni-correct our threshold. We chose this cutoff because we computed increases and decreases as two separate, but parallel, one-tailed tests (see above). For physiological plausibility, we added a constraint that the duration of a cluster had to be, at a minimum, 50 ms.

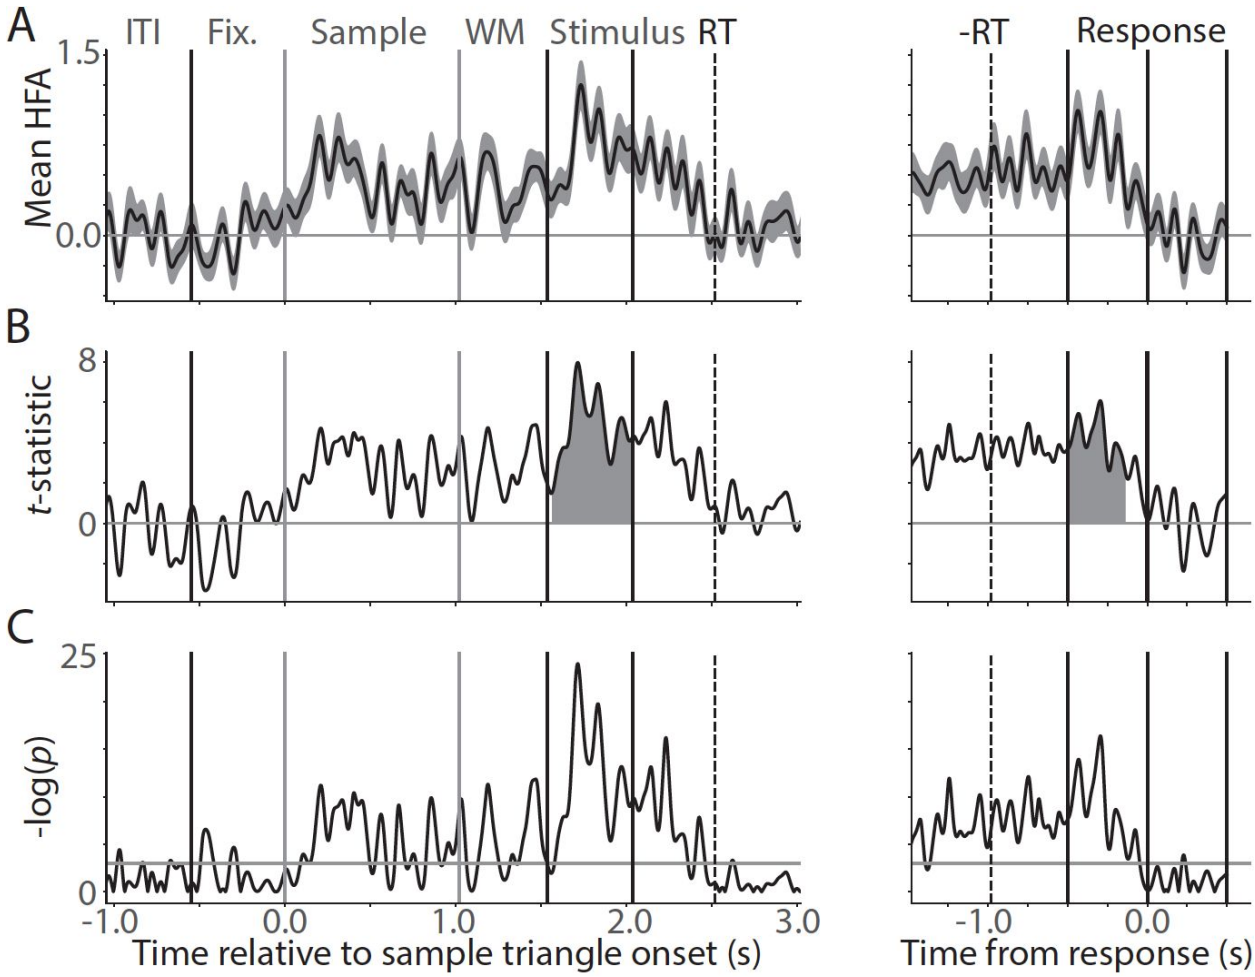

Figure 2.1. Example electrode and condition (Search) that was task-selective both in the first considered interval (stimulus onset to 500 ms, Stimulus) and the second considered interval (-500 ms to response). A) Mean, z-scored HFA signal, time-locked to sample triangle onset (left panel), and to the response (right panel). On the left panel, the first vertical line shows the offset of the baseline time-period. The second line shows the onset of the sample triangle. The third line shows the offset of the sample triangle. The fourth line shows stimulus onset. The fifth line shows the offset (500 ms) of the first task interval. The dotted line shows the median RT. On the right panel, the dotted line shows the median stimulus onset time. The second line shows the start of the second task interval (-500 ms relative to response). The third line shows the RT, and the fourth line marks 500 ms after the response. Grey labels indicate intervals between vertical lines. Black labels refer to time points (the vertical lines themselves). ITI, Inter-Trial Interval; Fix., Fixation Interval; Sample, Sample Interval; WM, Working Memory Interval; Stimulus, Stimulus Display Interval; RT, Response Time. B) One-sample *t*-statistics corresponding to the mean HFA traces on top. Shaded areas indicate significant clusters. The area of these clusters were used as the statistics considered in the permutation test. C) -log(*p*) values corresponding to the *t*-statistics in (B). The grey horizontal line shows $p = 0.05$. Any points above this line were significantly greater than the pooled baseline used to z-score the mean signal, as shown in (A).

Significance-testing the anatomical distribution of task-selective effects

We tested the distribution of task-active electrode sites for significant regional or hemispheric differences, using chi-square tests, according to the following approach: First, we tested if there were any regional differences in the distribution of task-active effects among the four broadly defined cortical lobes, cingulate cortex, MTL, and amygdala (Table 2.2., Appendix A). We similarly tested if there were differences between the subregions of the frontal cortex (Table 2.3., Appendix A), parietal cortex (Table 2.4., Appendix A), temporal cortex (Table 2.5., Appendix A), cingulate cortex (Table 2.6., Appendix A) and the MTL (Table 2.7., Appendix A), again using chi-square tests corrected for multiple comparisons using the FDR-correction. We did not examine subregion effects for the occipital lobe or the amygdala, due to comparatively low electrode counts in these regions. For any omnibus chi-square test (Tables 2-7) that was significant, we conducted follow-up pairwise chi-square tests to determine which ROIs or subregions were driving the overall effect. Any regions, for which the expected number of observations (electrodes) for either active or inactive status was less than 5, were excluded from the analysis. Specifically, the precuneus was excluded from the analysis of parietal subregions, and the temporal pole was excluded from the analysis of temporal subregions.

Identification of electrodes showing condition-based effects

We separately identified electrodes that showed greater activity increases in Pop-out than Search ("Pop-out electrodes") and electrodes that showed greater activity increases in Search than in Pop-out ("Search electrodes"). To accomplish this, we used a very similar procedure (Figures 2.2. and 2.3.) to the one employed for identifying task-active electrodes, except for the following modifications:

(1) We considered only electrodes that showed a significant task-related increase (see *Task-active electrode selection: Identifying significant increases* above).
(2) We excluded any electrodes located in sensorimotor cortex (localized according to the procedure described in the section, *Assignment of electrode locations to anatomically defined regions of interes*t) due to potential confounding effects of the behavioral response being given by button press. For example, it is conceivable that patients pressed the button with greater force in the easier and more confident Pop-out than Search condition, giving rise to spurious condition-related effects.
(3) Because we were interested in which electrodes showed a significantly larger *increase* relative to baseline in one condition versus the other, we zero-clipped the signal prior to making the comparison (using the function numpy.clip()), so that all compared time points were greater than or equal to zero. Had we omitted this modification, we may have inadvertently identified electrodes as showing a pop-out effect when that effect was in fact driven by a temporary *decrease* in the Search condition rather than an *increase* in the Pop-out condition.

(4) Instead of computing a one-sample *t*-test to identify significant deviations from 0 (step 2 in the task-active pipeline), we computed a two-sample *t*-test between the two conditions.
(5) Instead of using two intervals as in the task-active pipeline, we considered a single interval: from stimulus onset to median RT across the two conditions. The reasons for this were that (i) we were no longer restricted to using a time interval of equal length to the baseline; (ii) we observed that most condition-related changes were of stimulus-related type; (iii) using the median RT as a cutoff seemed to be a reasonable choice in order to avoid undue bias in either direction by the selection of interval length.

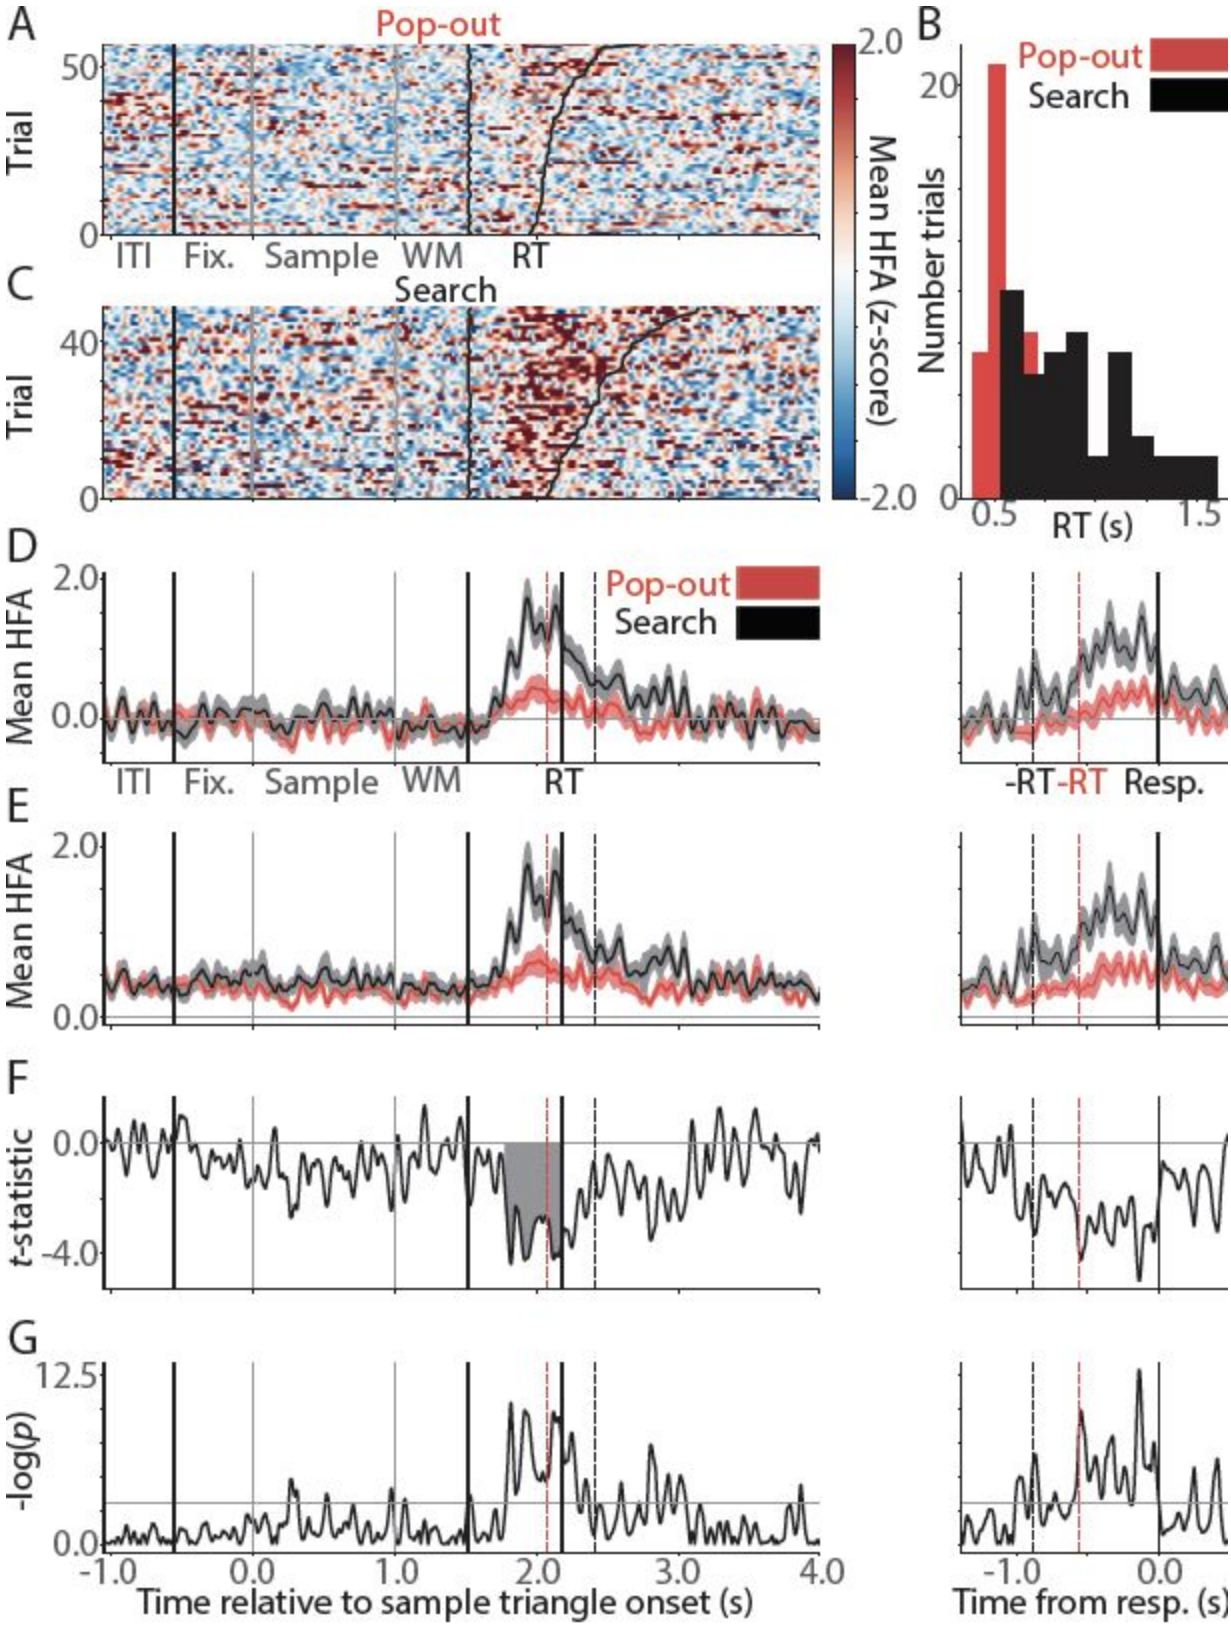

Figure 2.2. Example electrode that showed a search-selective effect. A) Stacked single-trial HFA plot for the Pop-out condition, where the trials are ordered by RT so that the fastest trial is

on the bottom, and the slowest trial is on the top. Grey labels indicate intervals between vertical lines, similarly to Figure 2.1. ITI, Inter-Trial Interval; Fix., Fixation Interval; Sample, Sample Interval; WM, Working Memory Interval. B) RT distributions for this patient: The RTs can affect the condition-based comparison. C) Stacked single-trial plot for the Search condition. The figure format corresponds to (A). D) Mean HFA activity, z-scored to a pooled baseline, and locked to sample (left) and response (right). The red dashed line on the left panel shows the median RT in Pop-out. The black dashed line shows the median RT in Search. The solid black line between the dashed lines show the median RT across the two conditions. In the right panel, the dashed lines show the median RTs in the two conditions relative to the response (with Pop-out in red and Search in black). E) The same HFA activity is shown as in (D), but here it is zero-clipped, to facilitate comparison between increases only in each condition. F) The two-sample *t*-statistic between the two conditions over time. The shaded area shows a contiguous cluster of significant differences between the two conditions. G) *p*-values corresponding to the *t*-statistics in (F). The horizontal line shows $p = 0.05$.

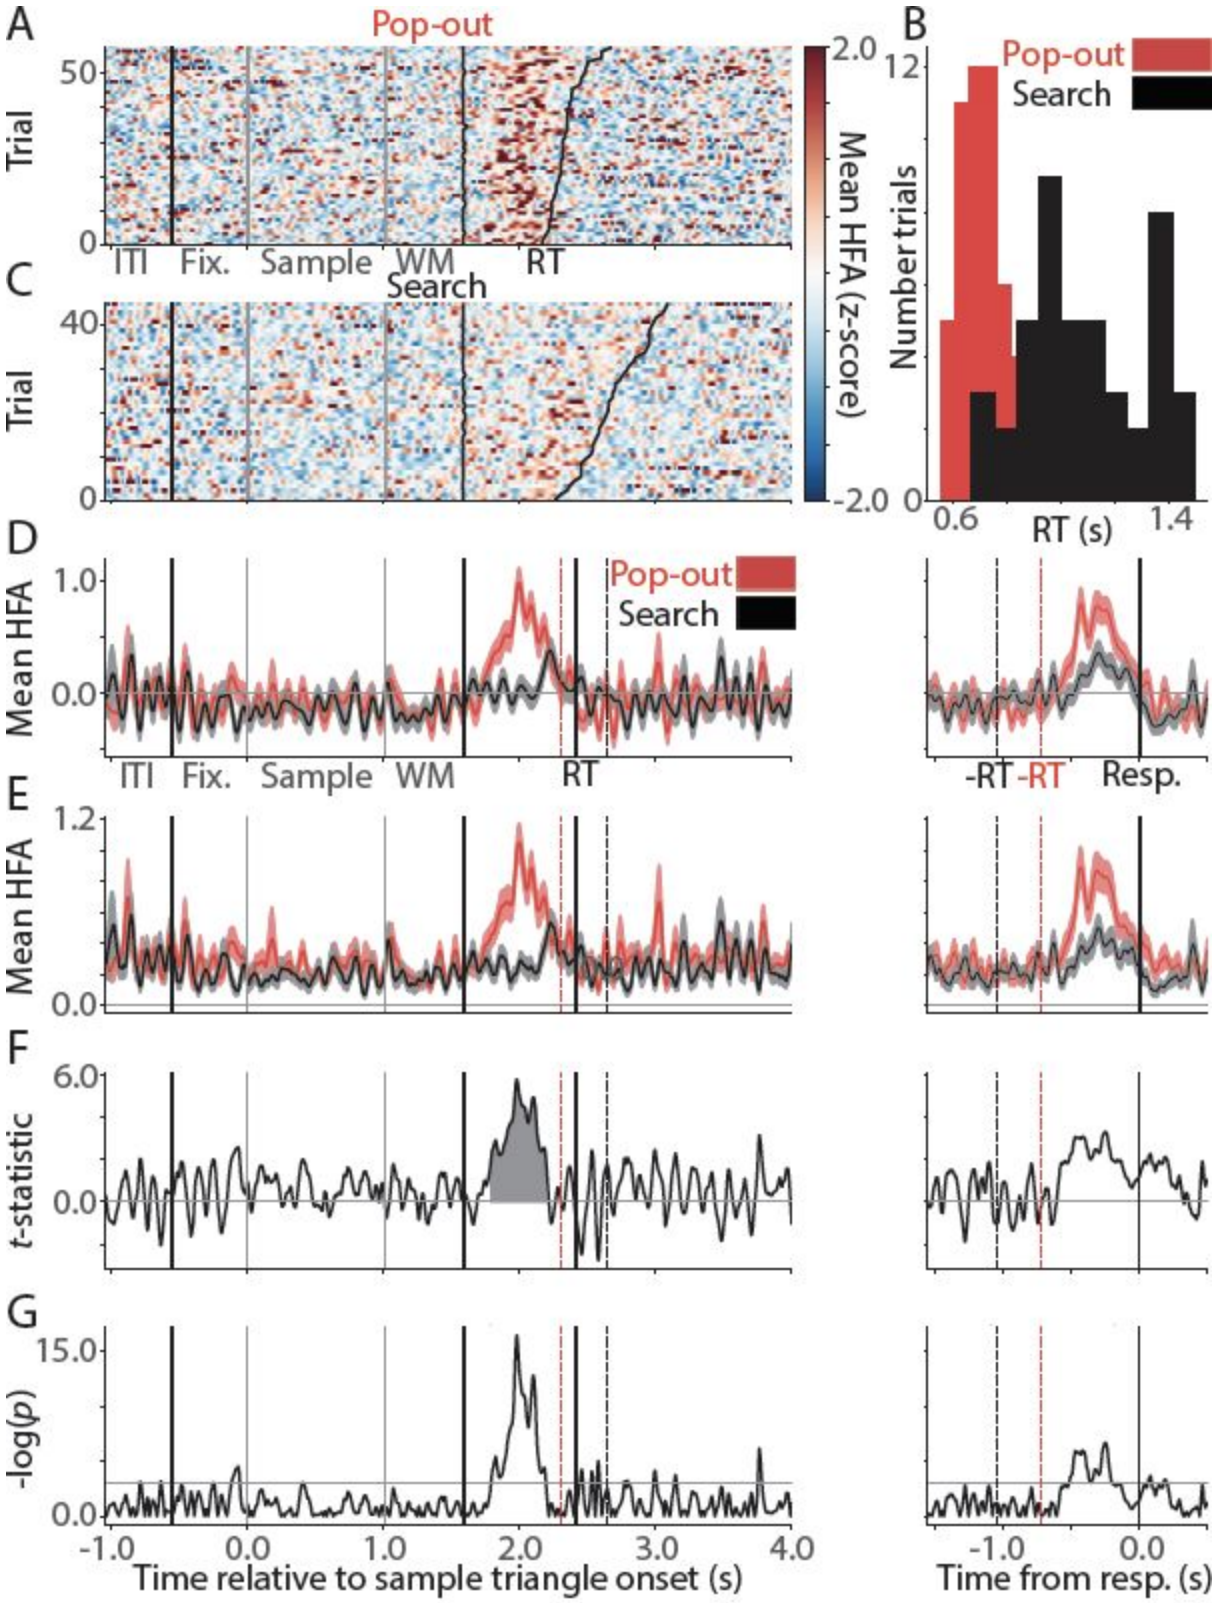

Figure 2.3. Example electrode that showed a Pop-out selective effect. The figure layout corresponds to that of Figure 2.2. Consistent with the task-active selection pipeline, we computed two one-tailed *t*-tests, one in each direction, at a *p*-value threshold of 0.025 for each

direction. We FDR-corrected the *p*-values across all included electrodes (those showing a task-selective increase); and included a minimum duration cutoff of 50 ms, identically to in the task-active pipeline above. Electrodes that showed a significant effect in both directions (n=3) were not considered to show a condition-related effect, and were not further analyzed.

While different RT distributions between the two conditions are an integral feature of the task, specific artifactual condition differences in the neural data can occur as a result of these behavioral differences between the two conditions. This is consistent with observations made in a previous intracranial study on visual search (Ossandon et al., 2012). We removed these artifactual effects through the following steps:

(1) We simulated three response types, which we predicted to occur in the HFA signal (previously reported in Ossandon et al., 2012; Haller et al., 2018): (i) stimulus-related responses; (ii) sustained responses; and (iii) response-related responses. We then tested how these response-types would behave given a representative patient's RT distributions in the two conditions in this task. See Figures 2.4-2.6. No artifacts driven by the known temporal characteristics of the ECoG/SEEG response (Ossandon et al., 2012; Haller et al., 2018, and simulated in Figures 2.5.-2.6.) are expected in the case of stimulus-related responses. Sustained-type responses can give rise to artifactual Search effects. Response-related responses can give rise to artifactual Pop-out or Search effects.
(2) We visually compared electrode traces that had been automatically labeled as showing a Pop-out or Search effect (using the pipeline described above) to the simulated artifactual responses. In cases where the signal time course of the selected traces matched a prototypical artifact time course, we labeled this effect as false. See Figures 2.7.-2.9. for example electrodes, for which the condition-effect labels were removed using this method.

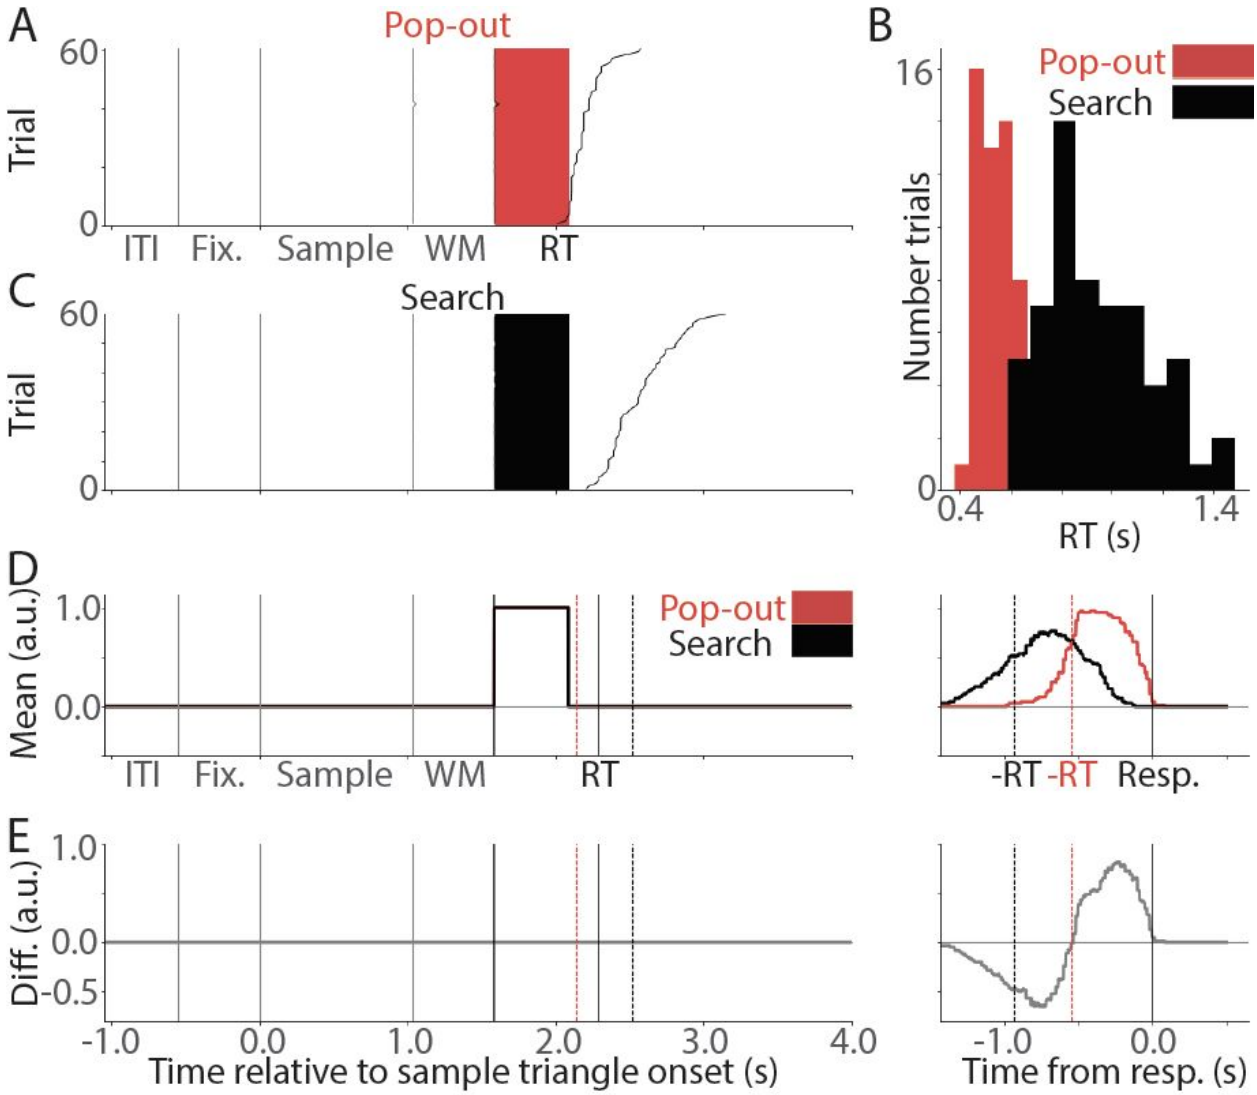

Figure 2.4. Simulated electrode showing stimulus-related HFA signal. A) Stacked single-trial plot for the simulated Pop-out condition, comparable to Figure 2.2.A. B) RT distributions for a representative patient, comparable to Figure 2.2.B. These RTs were used to generate the remaining figures in this plot (A, C-E). C) Stacked single-trial plot for the simulated Search condition, comparable to Figure 2.2.C. D) Mean of the single-trial traces, comparable to Figure 2.2.D. The left panel shows stimulus-locked traces, and the right panel shows response-locked traces. a.u., arbitrary units. E) Difference (Diff.) between the mean traces in the two simulated conditions, comparable to the outcome of a two-sample *t*-test as in Figure 2.2.F. The two panels illustrate that, in a stimulus-related electrode in which the two conditions do not differ in their response, the mean signal in the Pop-out and Search conditions will only differ if locked to the response (right panel). This does not affect the outcome of the selection of condition-related effects, which is based on stimulus-locked mean traces.

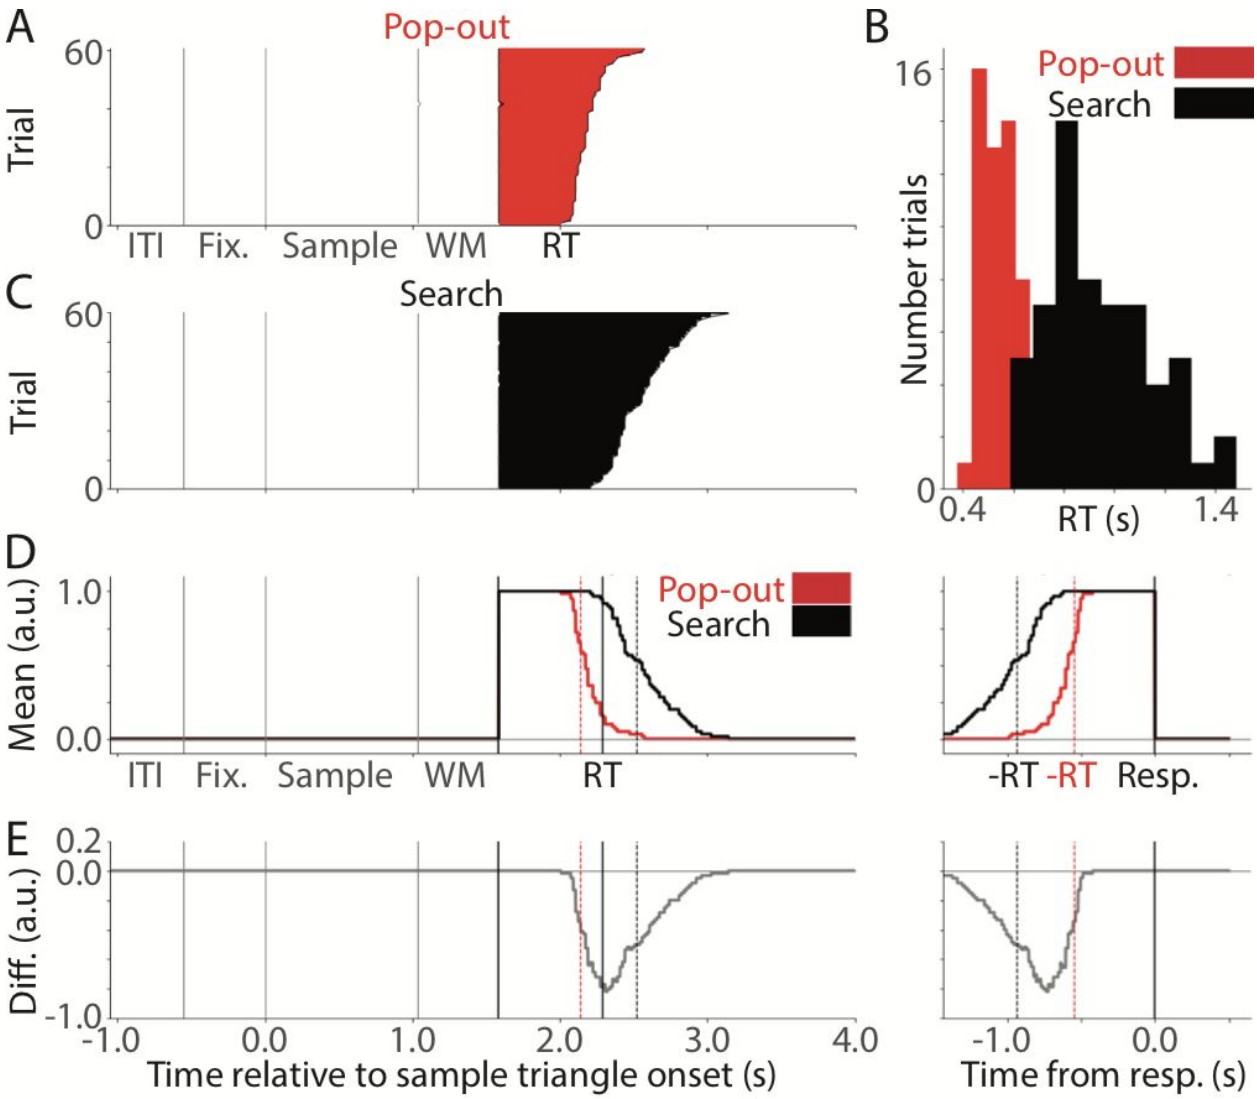

Figure 2.5. Simulated electrode showing sustained HFA signal. The figure layout corresponds to that of Figure 2.4. This illustrates that in a sustained-response electrode, in which the simulated response is governed only by the presence or absence of the stimulus screen, the mean signal in the Pop-out and Search conditions will differ such that a spurious Search effect can occur late in the trial interval. Figure 2.7. shows an example from the data, which features this type of artifact.

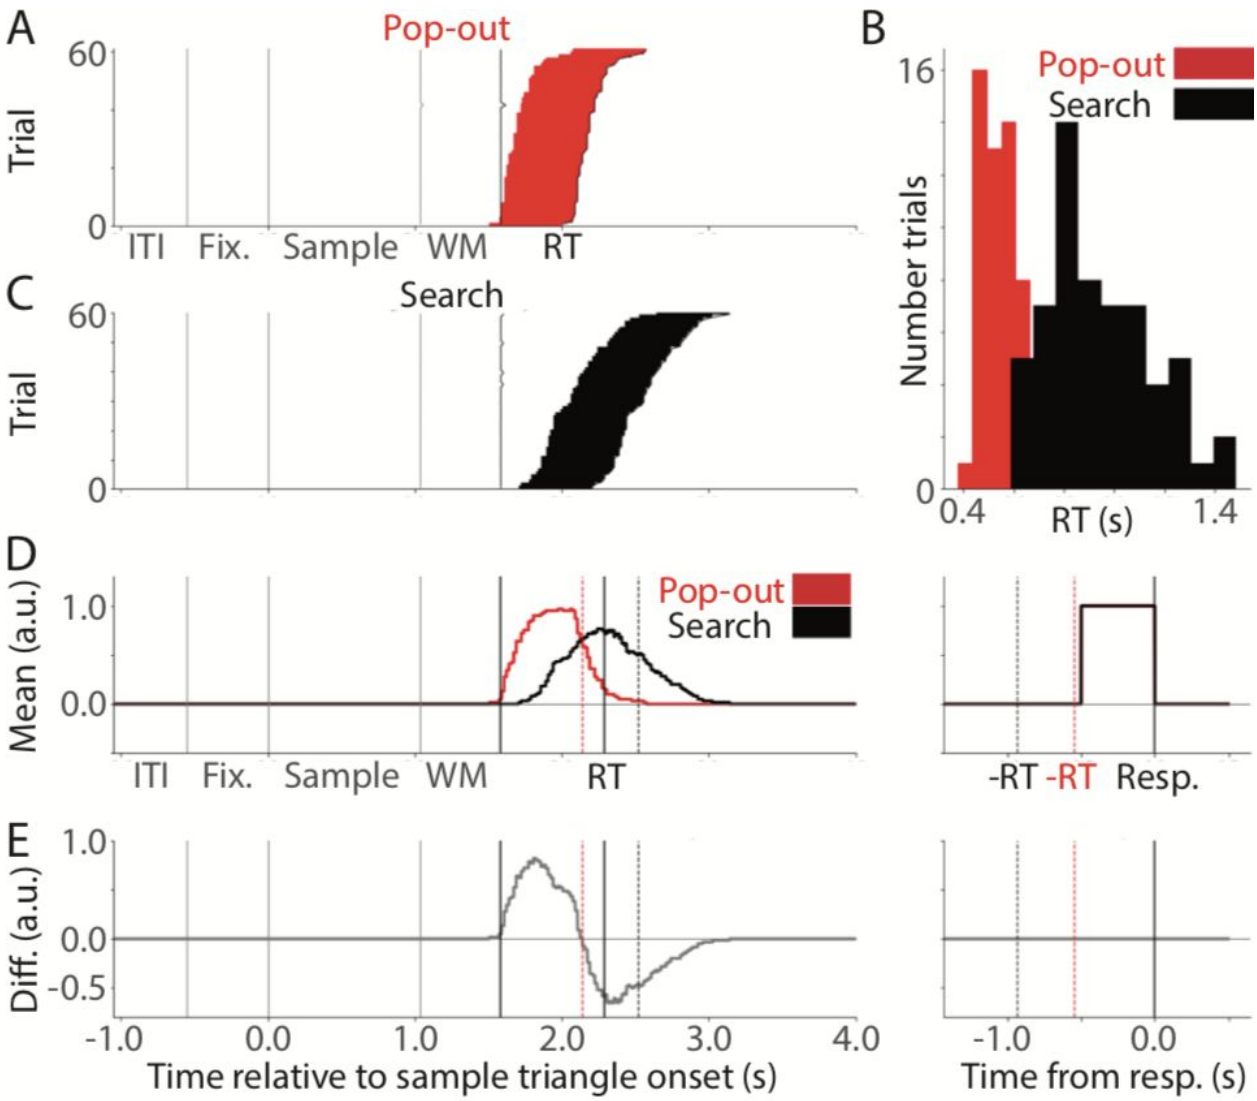

Figure 2.6. Simulated electrode showing response-related HFA signal. The figure layout corresponds to that of Figure 2.4. This figure illustrates that, in a response-related electrode, in which the simulated response begins 500 ms before the button press and continues until the time of the button press, the mean signal in the Pop-out and Search conditions will differ such that a spurious Pop-out effect can occur early in the trial interval, while a spurious Search effect can occur late in the trial interval. Figures 2.8. and 2.9. show examples from the data, in which these types of artifacts occur.

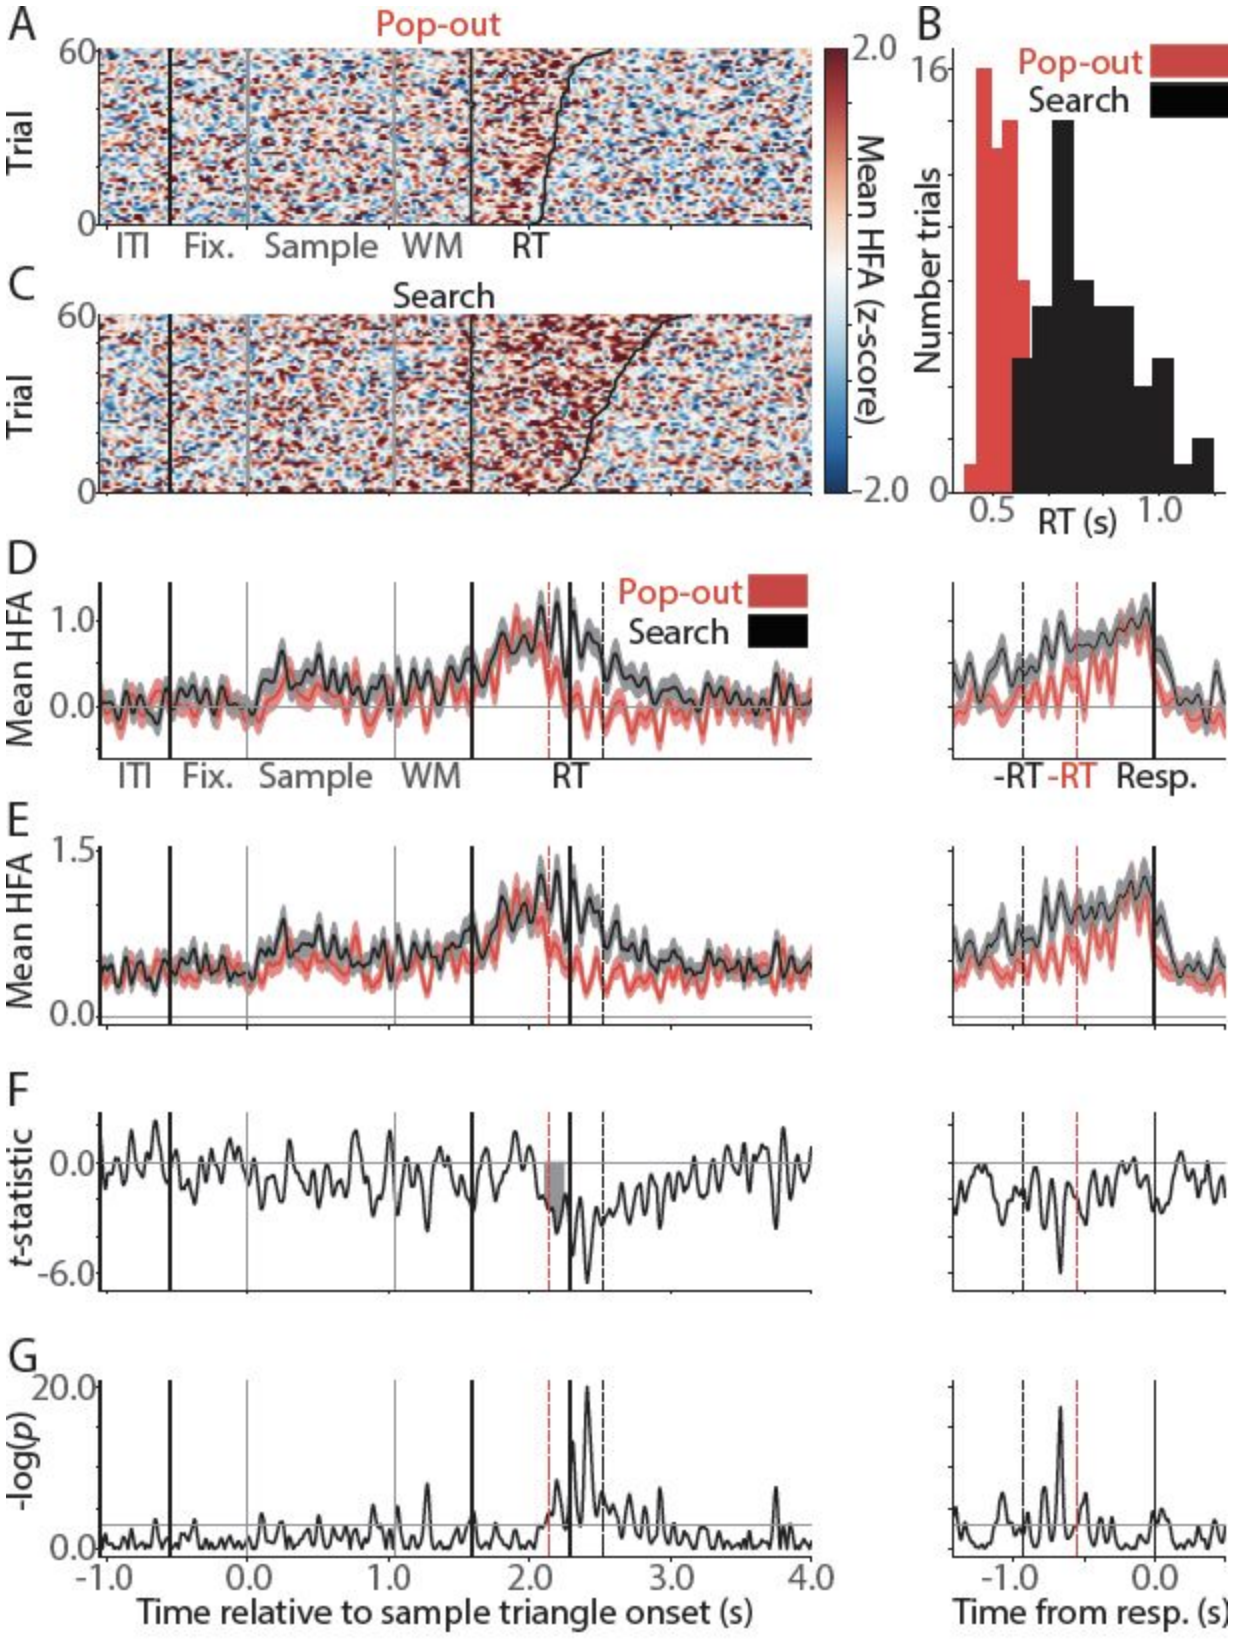

Figure 2.7. Electrode with sustained response pattern showing artifactual Search effect. The figure layout corresponds to that of Figure 2.2. This figure shows an example of an electrode,

which was automatically identified as showing a Search effect, but where this effect was artifactual due to the sustained response pattern of the electrode and the shape of the RT distribution in this patient. This electrode was labeled as not showing a condition-based effect, based on visual inspection. The observed artifact corresponds to the simulated artifact in Figure 2.5.

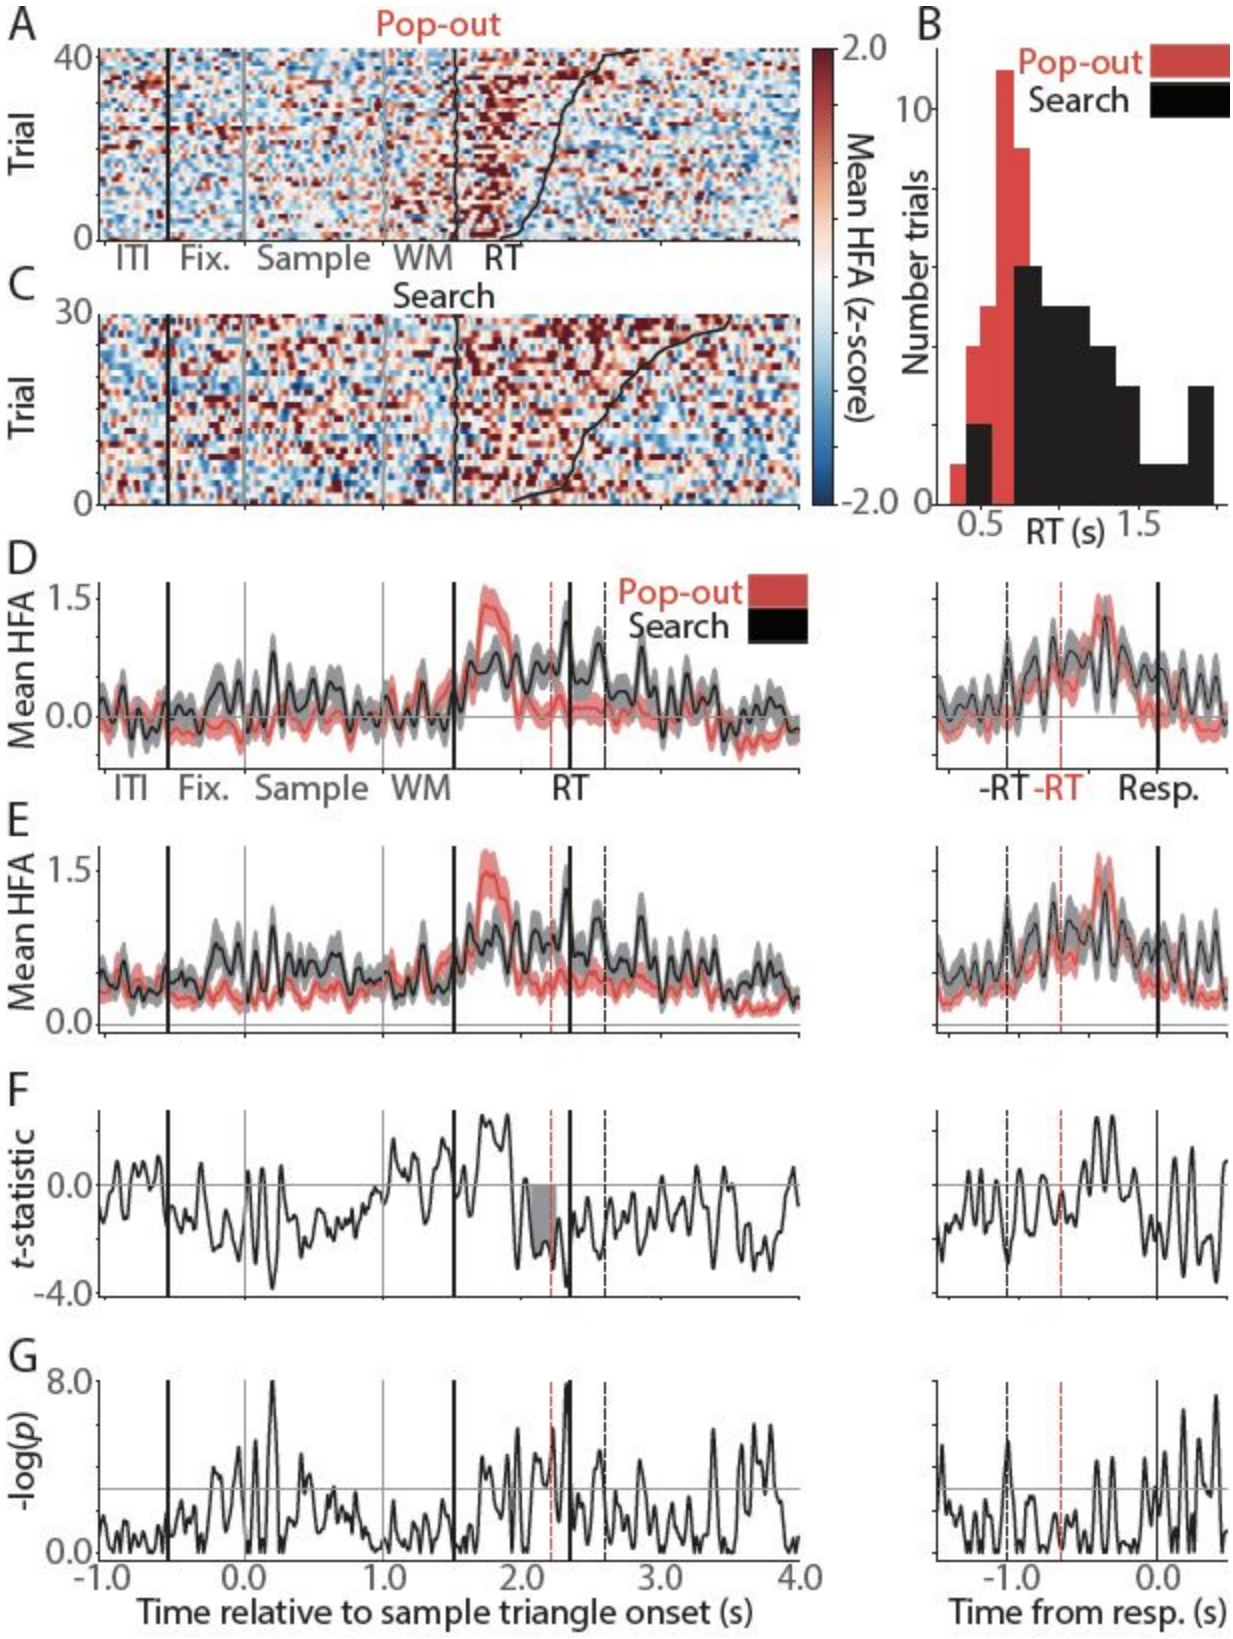

Figure 2.8. Electrode with response-related signal pattern showing artifactual Search effect. The figure layout corresponds to that of Figure 2.2. This figure shows an example of an

electrode, which was automatically identified as showing a Search effect, but where this effect was artifactual due to the response-related signal pattern of the electrode and the shape of the RT distribution in this patient: The response-locked panels on the right (D-E) show mean traces, which do not differ between conditions. This electrode was labeled as not showing a condition-based effect, based on visual inspection. The observed artifact corresponds to the simulated artifact in Figure 2.6. The shape of the *t*-statistic trace in the present figure (F) can be compared to the condition difference trace in Figure 2.6.

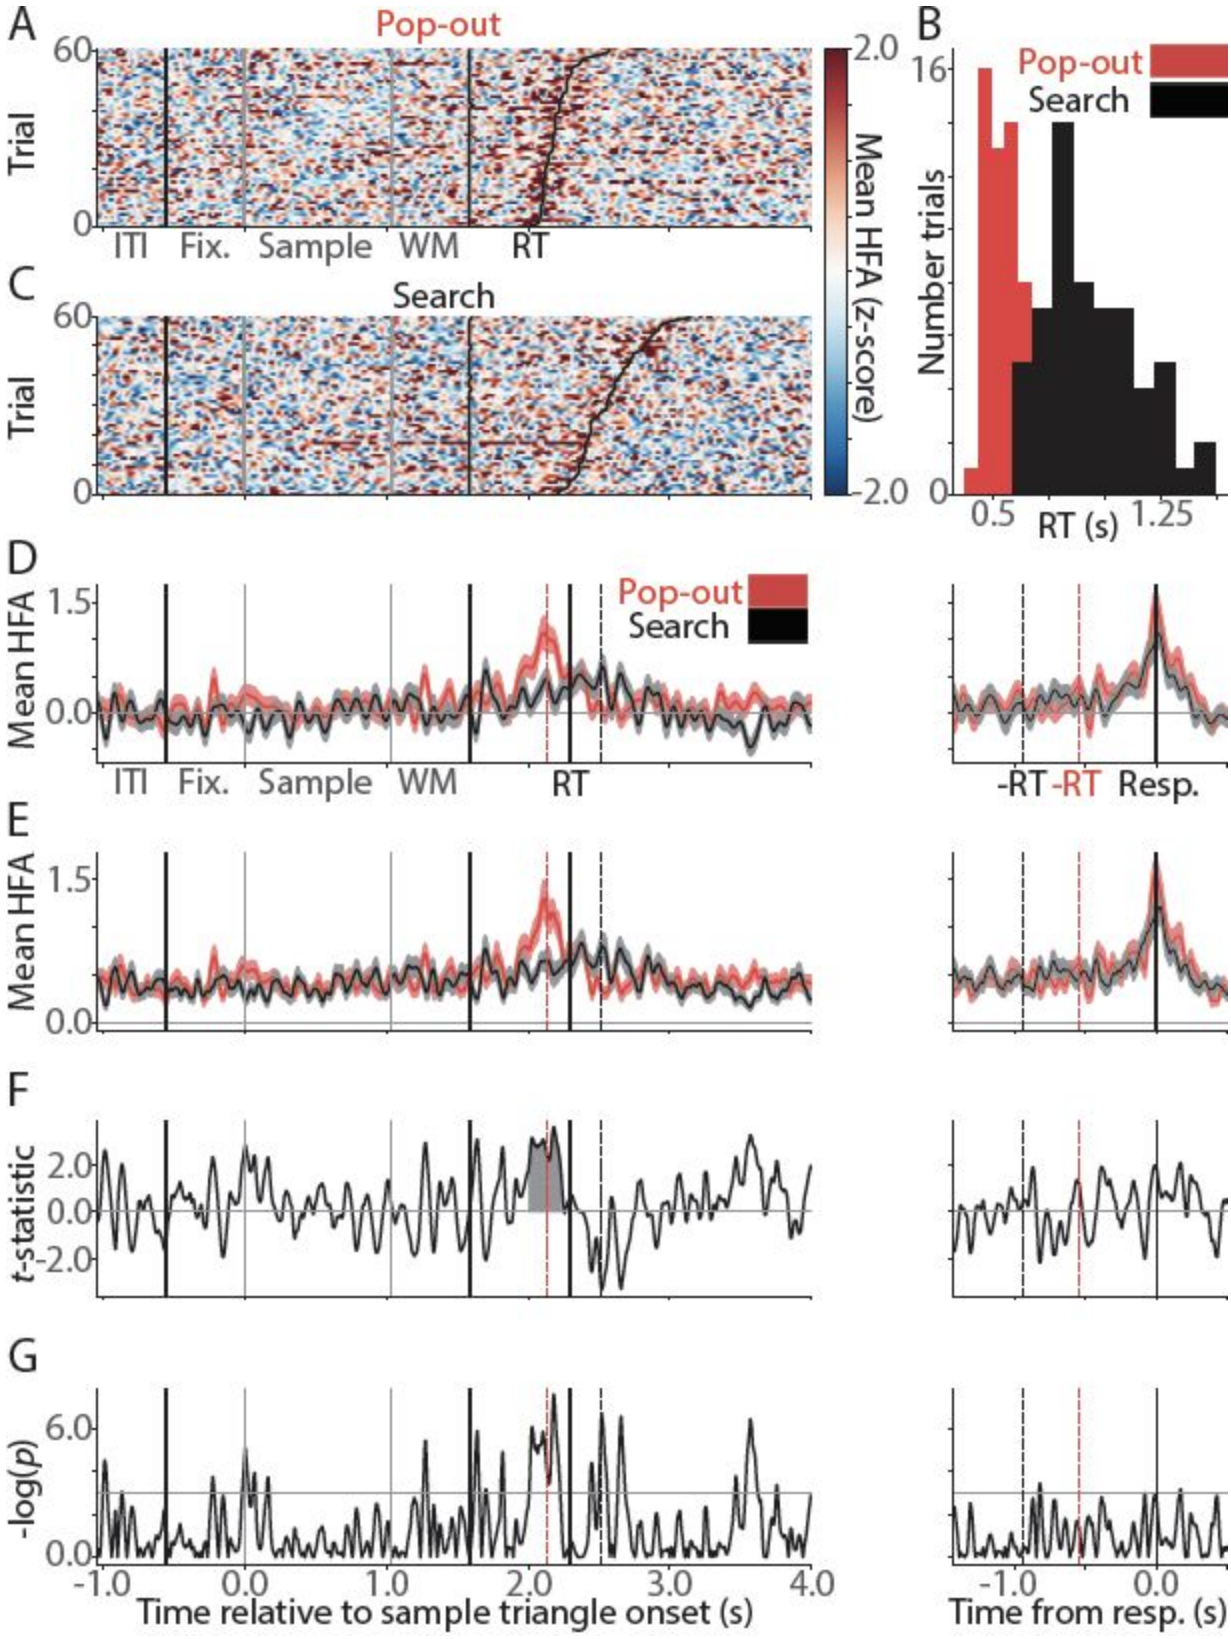

Figure 2.9. Electrode with response-related signal pattern showing artifactual Pop-out effect. The figure layout corresponds to that of Figure 2.2. This figure shows an example of an

electrode, which was automatically identified as showing a Pop-out effect, but where this effect was artifactual due to the response-related signal pattern of the electrode and the shape of the RT distribution in this patient: The response-locked panels on the right (D-E) show mean traces, which do not differ between conditions. Based on visual inspection, this electrode was labeled as not showing a condition-based effect. The observed artifact corresponds to the simulated artifact in Figure 2.6.

2.4. Results

Behavior

Consistent with past studies using similar tasks (Buschman & Miller, 2007; Li et al., 2010; Ossandon et al., 2012), subjects were more accurate ($t(22) = 4.73$, $p < 0.001$) and faster (showed faster RTs, $t(22) = -14.06$, $p < 0.001$) in Pop-out than in Search (Figure 2.10). The mean accuracy in the Pop-out condition was 96.0 % with a standard deviation of 3.9 %. The mean accuracy in the Search condition was 86.8 % with a standard deviation of 9.6 %. The mean RT (of the median RTs from each patient) in Pop-out was 680.2 ms with a standard deviation of 146.5 ms. The mean RT in Search was 964.8 ms with a standard deviation of 152.2 ms.

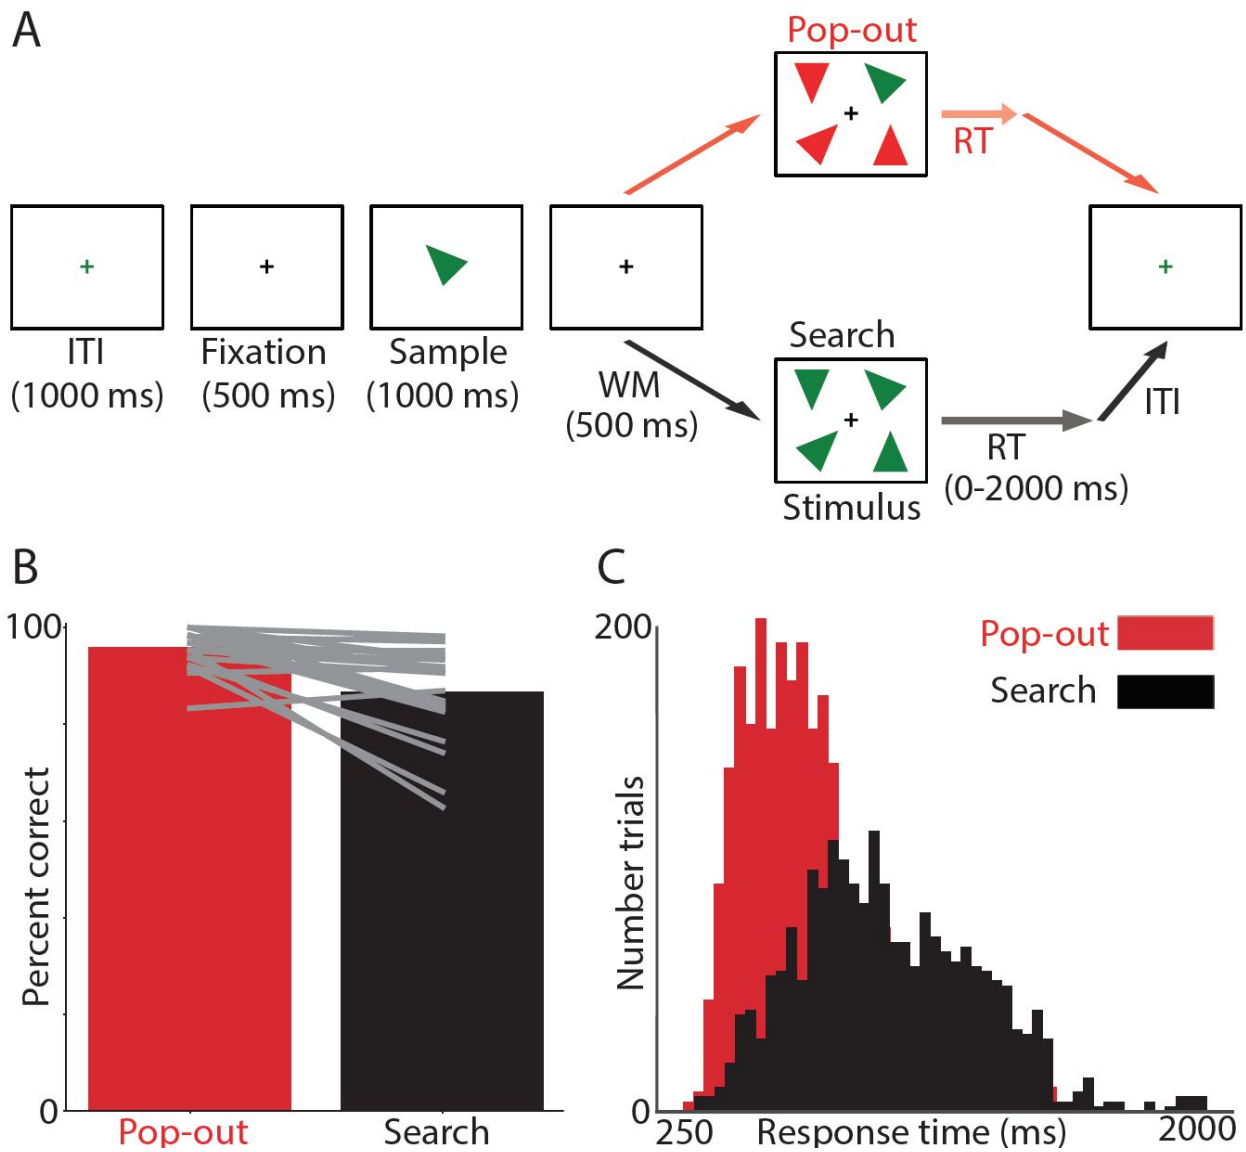

Figure 2.10. A) Visual search task adapted from Li et al. (2010). Patients searched for a target triangle ('Sample') of a given color and orientation. Upon locating the target triangle, they indicated its location on the left or right half of the screen by button press. ITI, Inter-Trial Interval; Fixation, Fixation Interval; Sample, Sample Interval; WM, Working Memory Interval; Stimulus, Stimulus Display Interval; RT, Response Time. B) Proportion correct responses in the Pop-out (red) and Search (black) condition across the 23 patients included in the study. Grey lines indicate individual subjects' performance. C) RT histograms from all correct trials for all participants (N=23), showing the Pop-out (red) and Search (black) conditions.

Task-selective responses by anatomically defined regions of interest

The total number of active sites, excluding electrodes implanted in sensorimotor cortex, was 363 out of 1,160, corresponding to 31% across the dataset (Table 2.2., Appendix A). Task-selective responses were detected in all considered brain regions, including all four lobes of the neocortex; the cingulate cortex; the medial temporal lobe; and the amygdala (Figure 2.11).

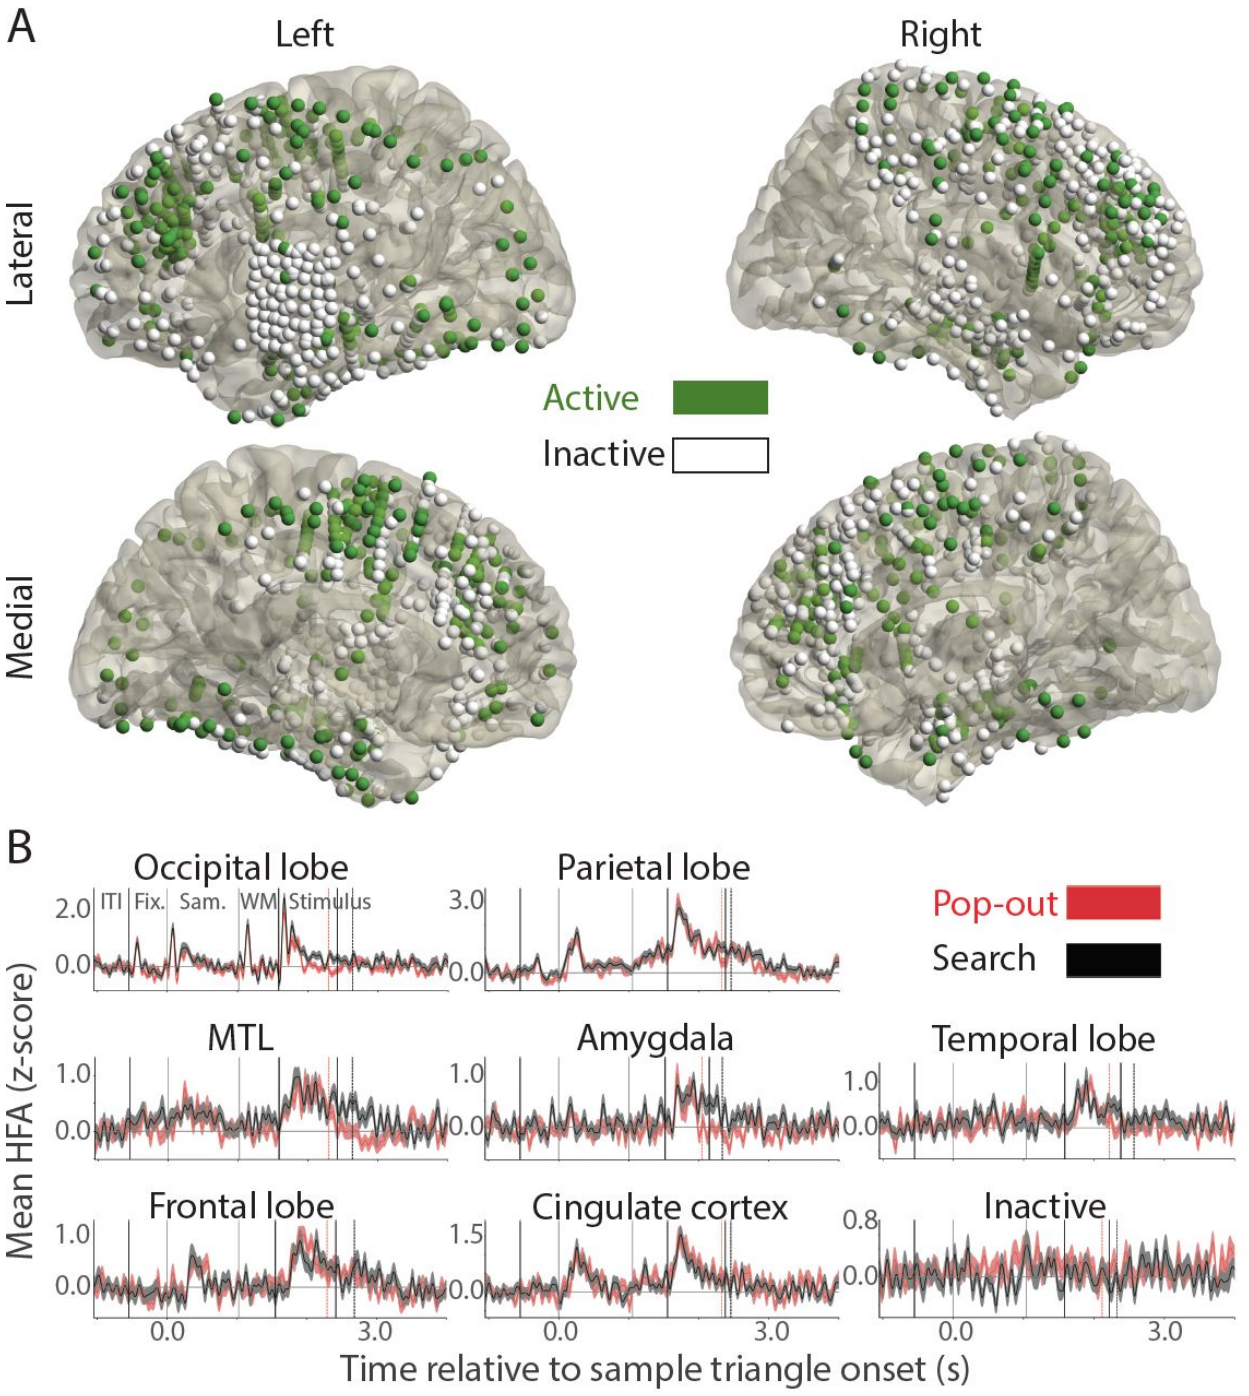

Figure 2.11. Task-selective electrode sites across the dataset (N=363 of 1,160). A) Task-active (green) and inactive (white) sites visualized on a common brain reconstruction (Colin-27). B) Example time courses from seven major regions of interest, and an example time course from an inactive electrode site. Pop-out time courses are plotted in red and Search time courses in black. The vertical red dashed line indicates the median RT in the Pop-out condition. The vertical black dashed line indicates the median RT in the Search condition. The various task epochs are indicated on the top left panel, and correspond to those shown in Figure 2.10.a: ITI, Inter-Trial Interval; Fix., Fixation Interval; Sam., Sample Interval; WM, Working Memory Interval; Stimulus, Stimulus Display Interval.

The proportions of active electrodes differed significantly between these areas, $\chi 2$ (6, N = 1,160) = 41.84, $p < .001$ (Figures 2.11. and 2.12.). The occipital lobe, with the most limited coverage (17 electrodes), showed the strongest activation profile (65%), as expected given the visual nature of the task. The proportion active electrodes in the occipital lobe was significantly larger than in the temporal lobe, $\chi 2$ (1, N = 386) = 14.70, $p < .001$ (FDR-corrected); and the cingulate cortex, $\chi 2$ (1, N = 163) = 9.54, $p = .008$ (FDR-corrected).

A)

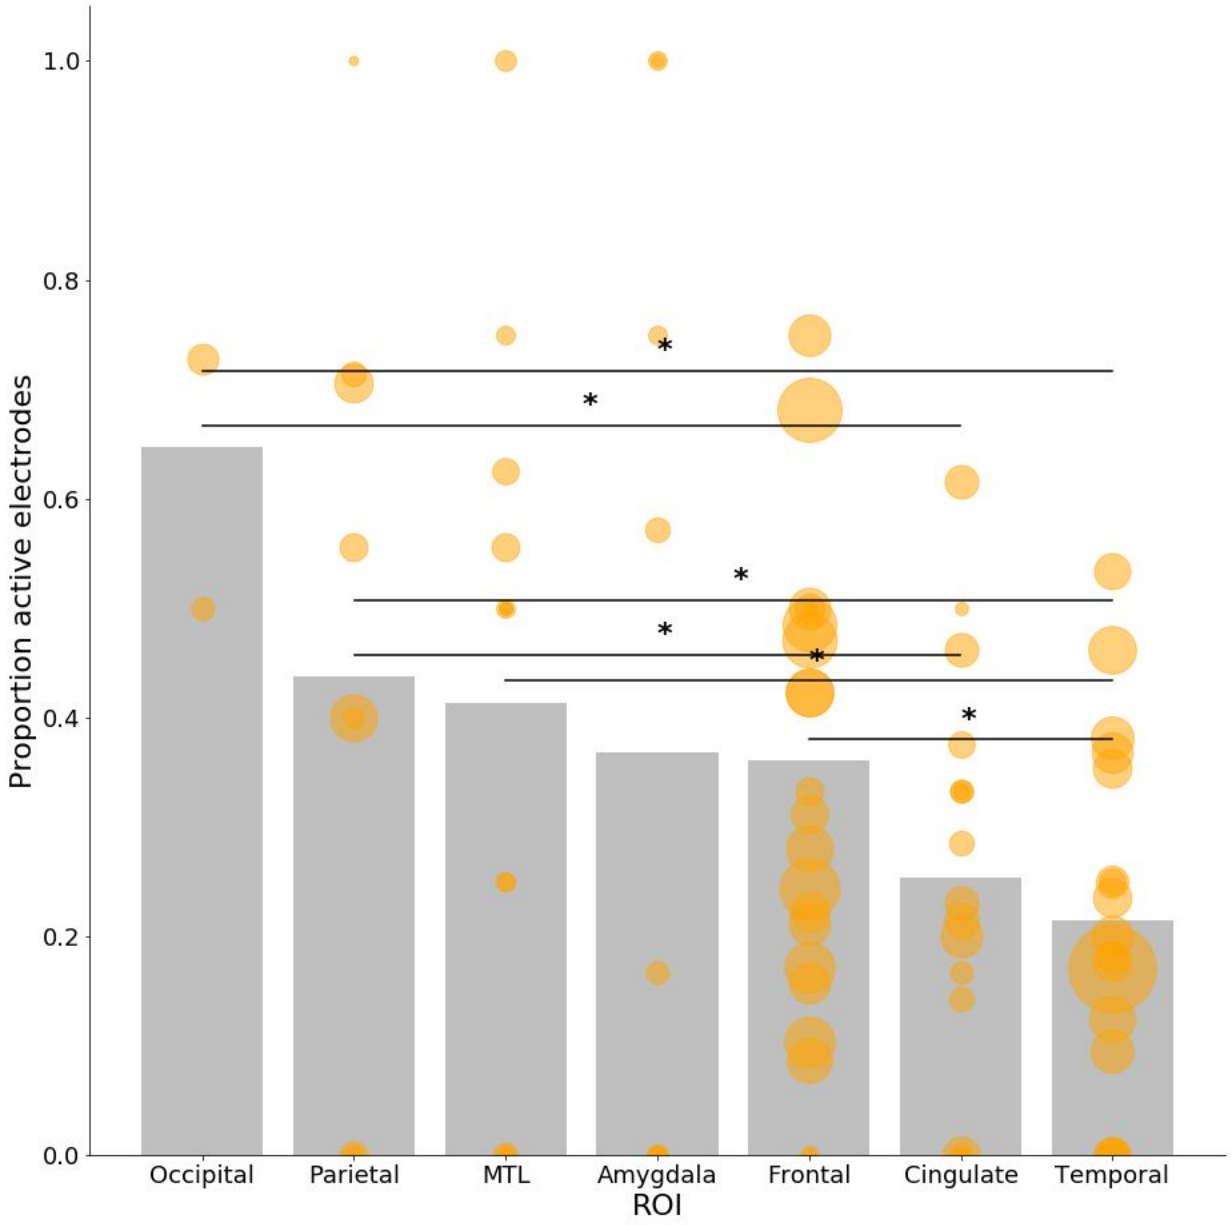
1.0
0.8
0.6
0.4
0.2
0.0
Proportion active electrodes
*
*
*
*
*
*
Occipital
Parietal
MTL
Amygdala
Frontal
Cingulate
Temporal
ROI

B)

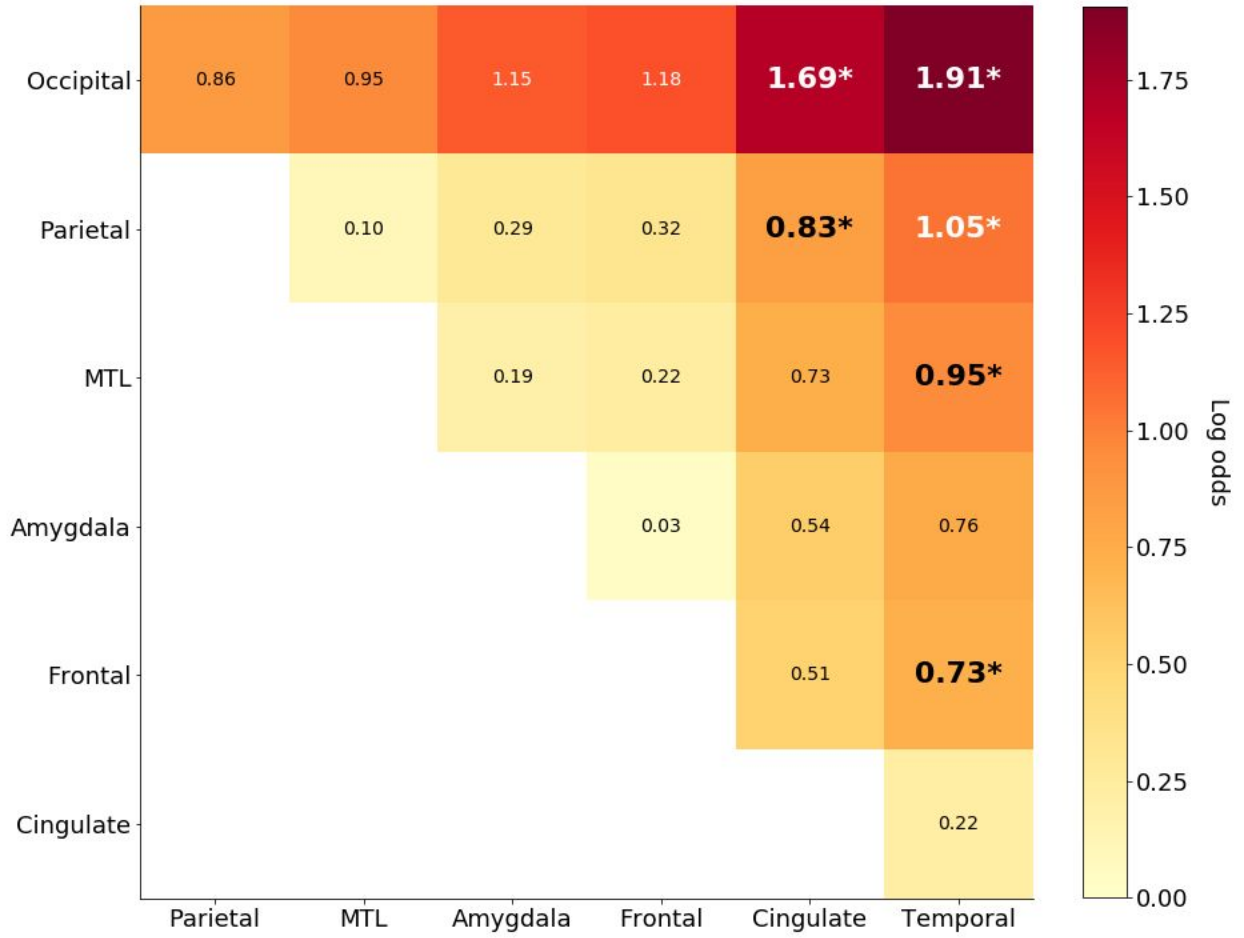

Figure 2.12. Distribution of task-active sites across the major ROIs. A) Proportion active electrodes by ROI. Bar height shows proportion active electrodes for a given ROI across the dataset. Circle height shows proportion active electrodes for individual subjects, who had coverage in that ROI. Circle area is proportional to the number of considered electrodes in that subject. Black asterisks indicate a significant pairwise comparison ($p < 0.05$). B) Effect sizes of pairwise comparisons between the major ROIs, corresponding to the pairwise chi-square test statistics reported. The effect sizes are illustrated as the log odds ratio between the ROI on the y-axis (indicating rows, on the left) versus the ROI on the x-axis (indicating columns, at the bottom): The proportion active electrodes in the ROI on the y-axis is always larger than the corresponding ROI on the x-axis. Statistics with an asterisk indicate a significant pairwise chi-square test.

Parietal cortex, extensively implicated in visual attention tasks, showed 44% active electrodes (Tables 2.2. and 2.4.; Figures 2.11. and 2.12.). The proportion active electrodes in parietal cortex was significantly larger than in the temporal lobe, $\chi 2$ (1, N = 449) = 16.16, $p < .001$ (FDR-corrected) and the cingulate cortex, $\chi 2$ (1, N = 226) = 7.24, $p = .025$ (FDR-corrected). We further examined three subregions of parietal cortex: the precuneus, the inferior parietal lobule (IPL), and the superior parietal lobule (SPL). In the precuneus, six of eight electrodes (75%) were active; due to the low electrode

count, we did not further analyze this proportion. We compared the proportion active electrodes between the IPL (33% or 18/55 electrodes) and SPL (65% or 11/17 electrodes), and found a significantly greater proportion of active electrodes in SPL, $\chi 2$ (1, N = 72) = 4.27, p = .039. The log odds ratio of this effect was 1.33. We note, however, that all SPL electrodes were derived from a single subject and the right hemisphere alone.

In prefrontal cortex (PFC; Tables 2.2. and 2.3.; Figures 2.11.-2.12.), 36% of electrodes were active. This proportion was significantly larger than in the temporal lobe, $\chi 2$ (1, N = 821) = 20.28, p < .001 (FDR-corrected), but did not differ significantly from parietal sites. The proportion active electrodes were not uniformly distributed across frontal cortex, $\chi 2$ (4, N = 452) = 23.85, p < .001 (FDR-corrected; Figure 2.13.). In lateral PFC, we observed a gradient of increasing activations from superior to inferior frontal gyri. In the superior frontal gyrus (SFG), we observed a proportional activation rate of 31%. In the middle frontal gyrus (MFG), 43% of electrodes were active. In the inferior frontal gyrus (IFG), the activation rate was 53%, with a larger proportion located in the left hemisphere (68% activation in the left hemisphere vs. 48% in the right hemisphere). The proportion active electrodes in IFG was significantly larger than that in SFG, $\chi 2$ (1, N = 202) = 8.92, p = .009 (FDR-corrected). We observed proportionally fewer task-related increases in OFC (19%) in spite of extensive coverage (75 electrodes). This proportion was significantly smaller than that in IFG, $\chi 2$ (1, N = 148) = 17.96, p < .001 (FDR-corrected), and MFG, $\chi 2$ (1, N = 211) = 11.32, p = .004 (FDR-corrected). In medial PFC (mPFC), 31% (12/39) of electrodes were active. This proportion did not differ significantly from other frontal subregions.

A)

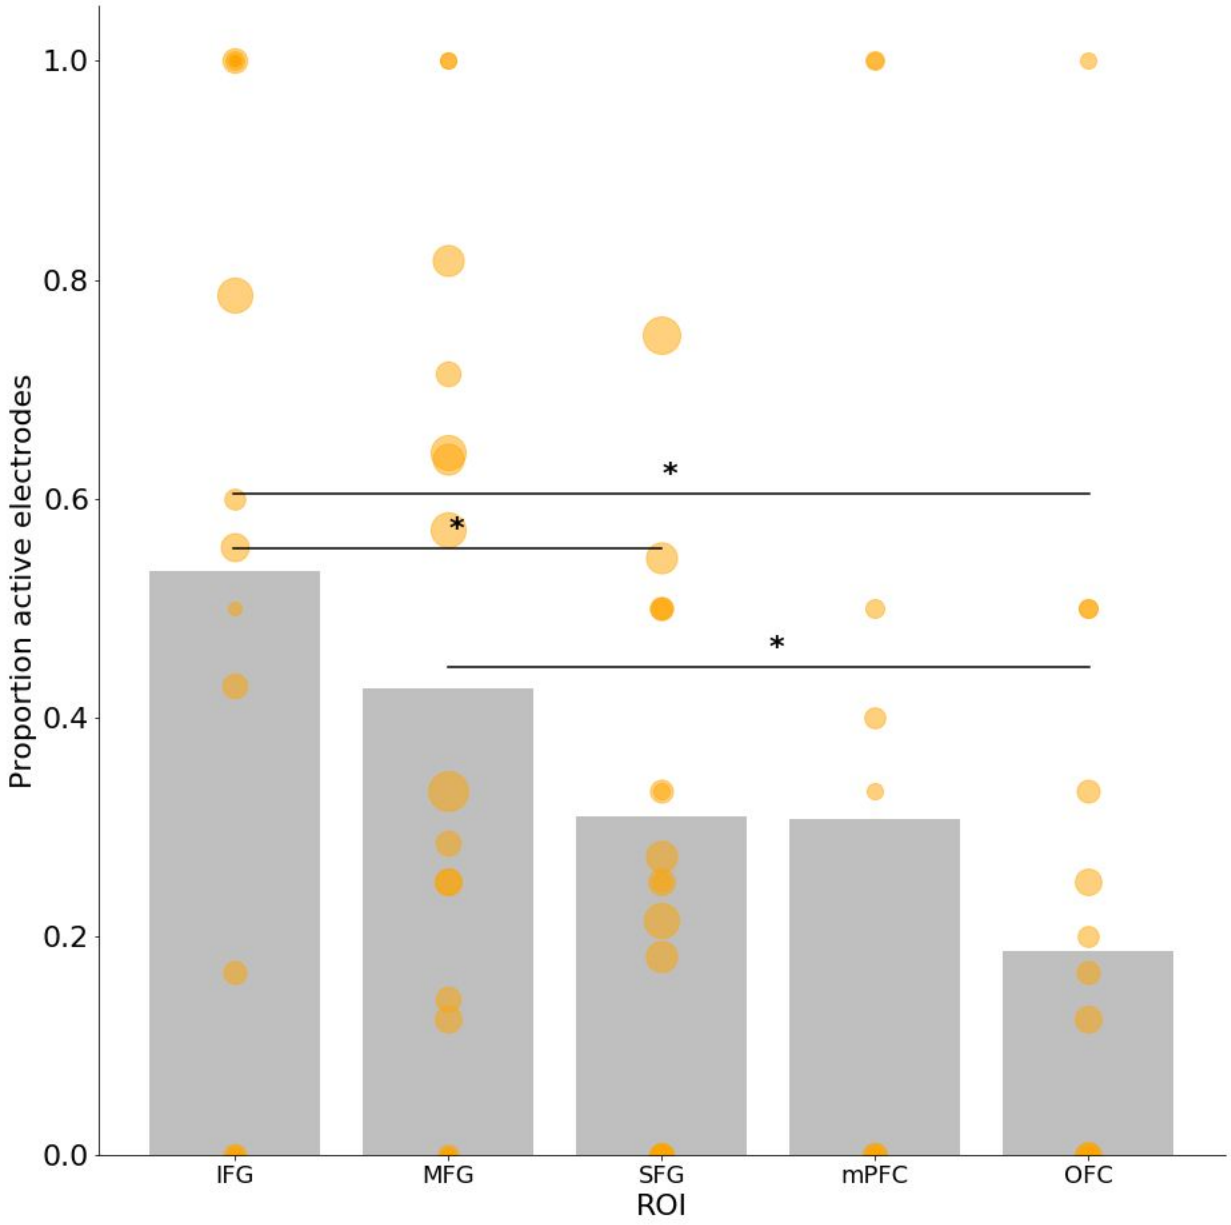
1.0
0.8
0.6
0.4
0.2
0.0
Proportion active electrodes
*
*
*
IFG
MFG
SFG
mPFC
OFC
ROI

B)

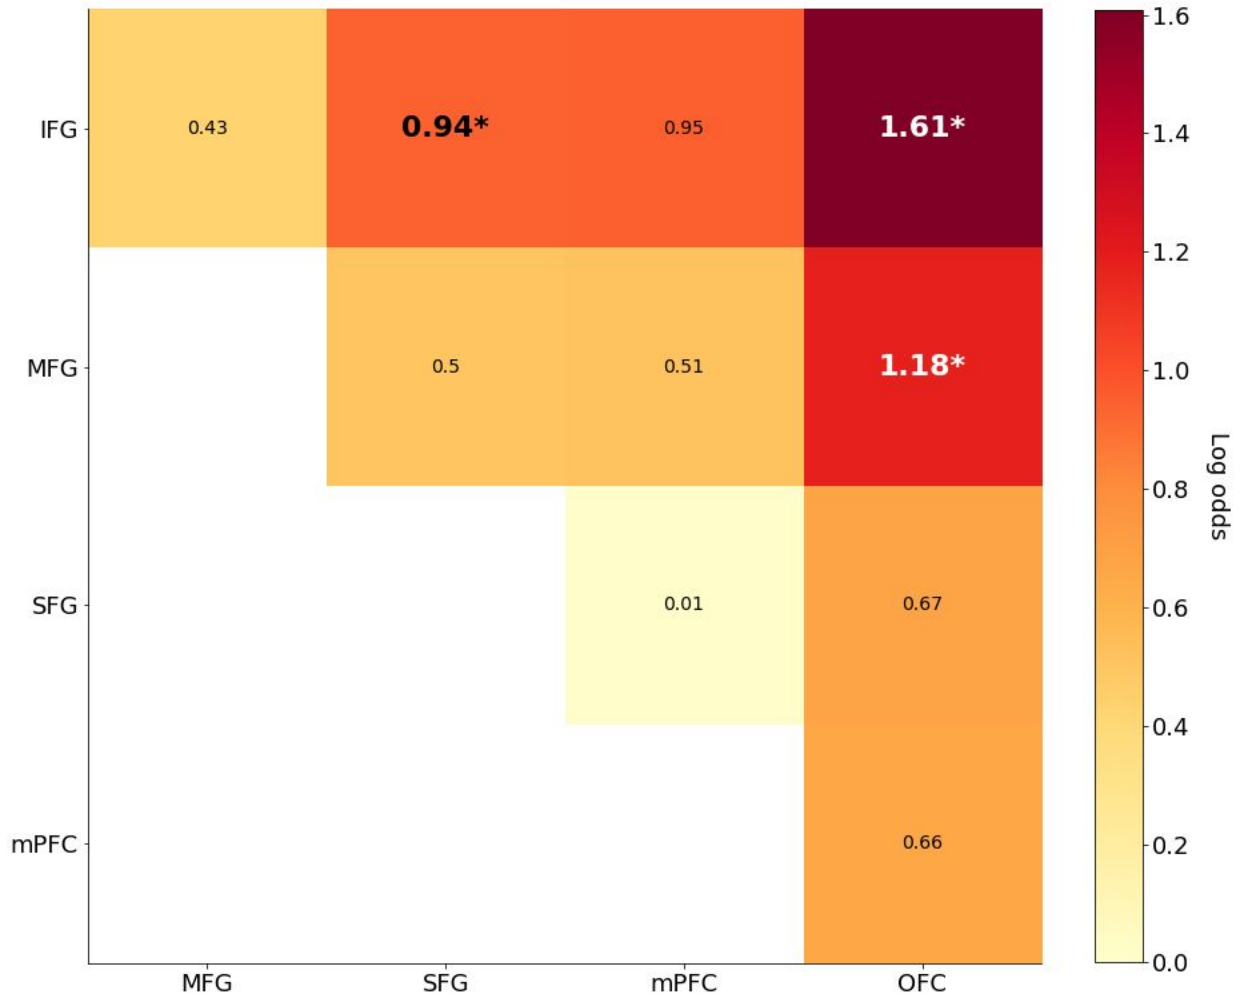

Figure 2.13. Distribution of task-active sites across subregions of frontal cortex. The layout of the figure corresponds to that of Figure 2.12. A) Proportion active electrodes by frontal subregion. B) Effect sizes of pairwise comparisons between frontal subregions.

The MTL (Tables 2.2. and 2.7., Appendix A; Figures 2.11. and 2.12.) showed a proportional activation profile at 41% (24/58 electrodes), which is indistinguishable from frontal cortex. Similarly to frontal cortex, this proportion was significantly larger than that in temporal cortex, $\chi 2$ (1, N = 427) = 9.86, p = .008 (FDR-corrected). We detected no significant sub-regional differences within the MTL. In the hippocampus, we observed 41% (9/22) active electrodes. In parahippocampal cortex, 32% (7/22) of electrodes were active. Eight of fourteen (57%) entorhinal cortex electrodes were active.

The amygdala (Table 2.2., Appendix A; Figures 2.11.-2.12.) showed 37% active electrodes (14/38). This proportion did not differ significantly from the other ROIs.

The cingulate cortex (Tables 2.2. and 2.6., Appendix A; Figures 2.11.-2.12.) was engaged only at 25% in spite of extensive coverage (37/146 electrodes). This proportion was significantly smaller than that in the occipital and parietal cortices, as reported above. An omnibus chi-square test for significant regional differences between cingulate

subregions was significant, $\chi 2$ (1, N = 146) = 7.12, p = .047 (FDR-corrected); however, the follow-up pairwise chi-square tests were not significant. We observed 33% (9/27) and 35% (18/52) active electrodes in the posterior (PCC) and midcingulate cortex (MCC), respectively. In the anterior cingulate cortex (ACC), we found only 15% active electrodes in spite of extensive coverage (10/67).

Temporal lobe engagement (Tables 2.2. and 2.5., Appendix A; Figures 2.11., 2.12., and 2.14.) was comparatively low, at 21% (79/369 electrodes). The proportion active electrodes in temporal cortex was significantly smaller than that in occipital, frontal, and parietal cortices, as well as the MTL, as reported above. Within the temporal lobe, we observed a gradient of increasing activations from dorsal to ventral areas (Figure 2.14.). This observation is consistent with the view that the ventral stream is central for object detection while lateral temporal cortex would not be expected to be engaged in visual search. The ventral stream areas (comprised of the lingual and fusiform gyri) showed 41% active electrodes (18/44). This proportion was significantly larger than that in STG, $\chi 2$ (1, N = 88) = 10.24, p = .010 (FDR-corrected), and STS $\chi 2$ (1, N = 126) = 13.51, p = .004 (FDR-corrected). The STG and STS showed only 9% and 11% active electrodes, respectively, in spite of extensive coverage in both areas (4/44 electrodes in STG and 9/82 electrodes in STS). The MTG showed 20% (16/79) active electrodes; ITG showed 22% active electrodes (13/59). In the insula, 26% (12/46) of electrodes were active. Of 15 electrodes localized to the temporal pole, 7 were active (47%); we did not consider the temporal pole in significance testing, due to low electrode count.

A)

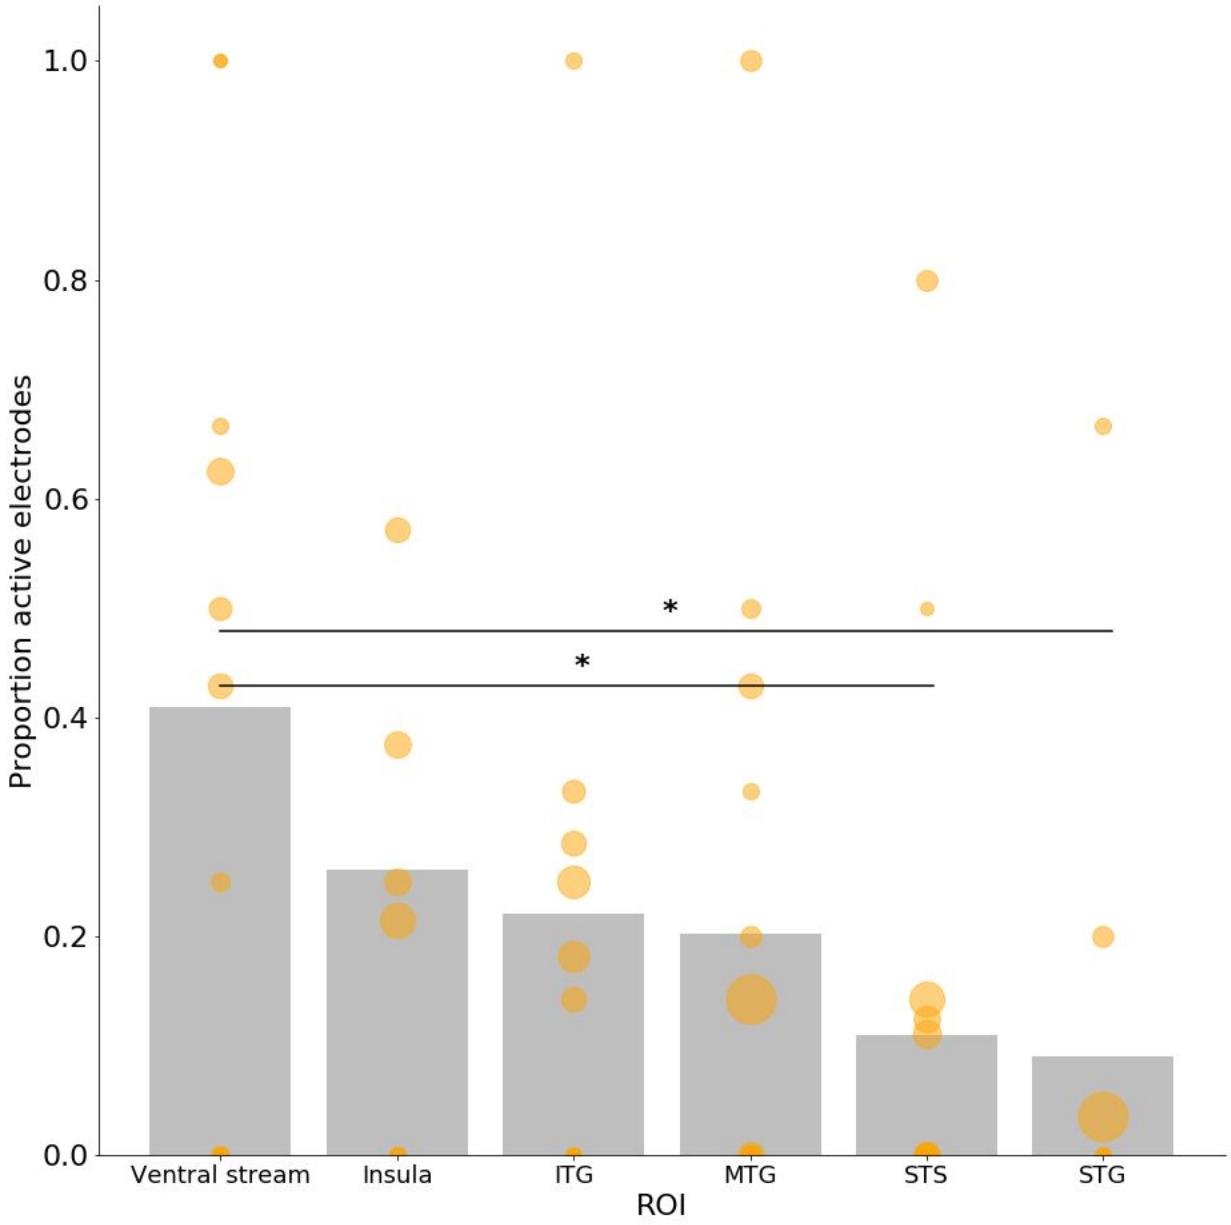
1.0
0.8
0.6
0.4
0.2
0.0
Proportion active electrodes
*
*
Ventral stream
Insula
ITG
MTG
STS
STG
ROI

B)

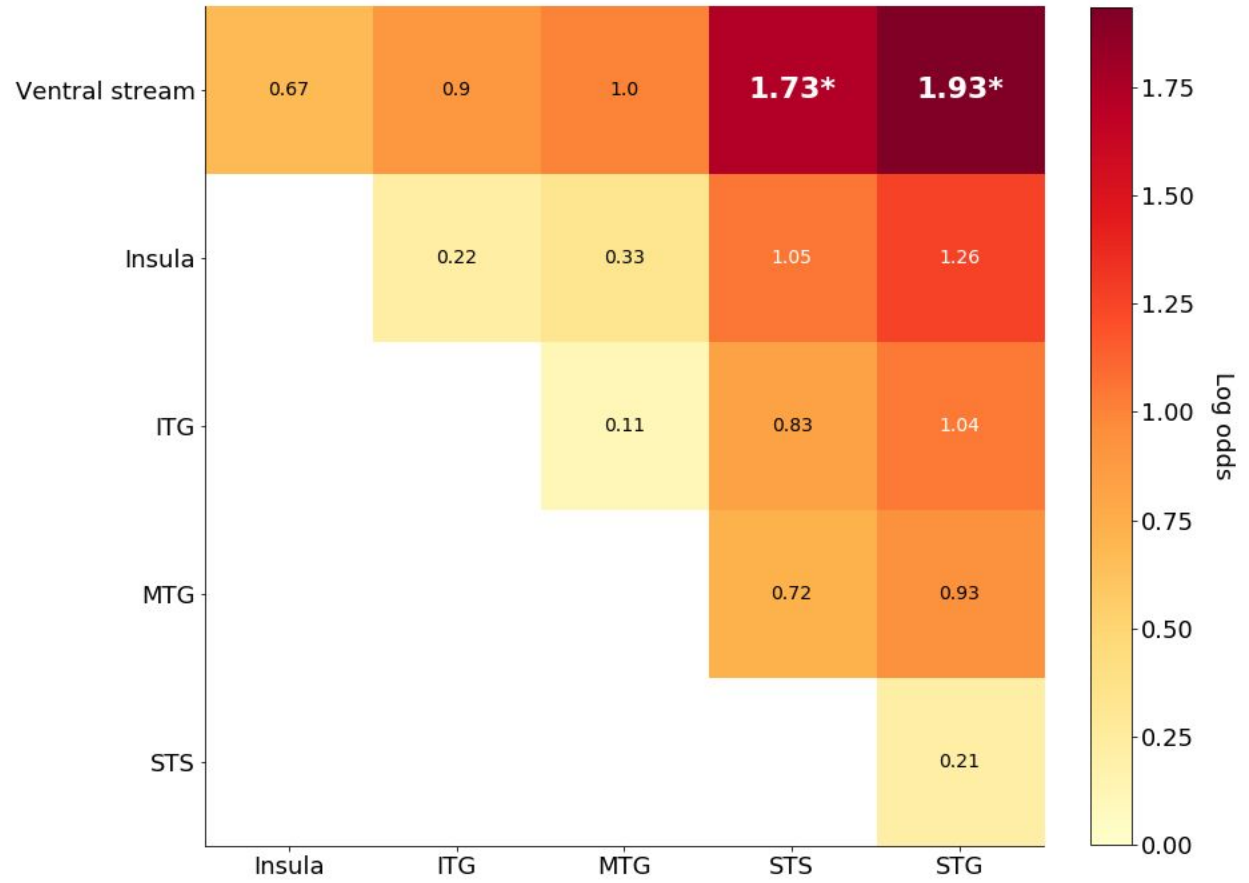

Figure 2.14. Distribution of task-active sites across subregions of temporal cortex. The layout of the figure corresponds to that of Figure 2.12. A) Proportion active electrodes by temporal subregion. B) Effect sizes of pairwise comparisons between temporal subregions.

In sum, occipital, frontal, parietal cortex showed an elevated proportion active electrodes compared to temporal cortex. The MTL showed a proportional activation rate on par with that in frontal and parietal cortex. As illustrated in Figures 2.12.-2.14., none of these effects were driven by a single, anomalous subject.

Condition-selective anatomical sites

Among the 363 electrodes which showed task-related increases outside of sensorimotor areas (Table 2.2., Appendix A, and Figure 2.15.), 30 sites (8.3%; 13 of 23 patients) showed condition-related modulations (*Pop-out* > *Search* or *Search* > *Pop-out*).

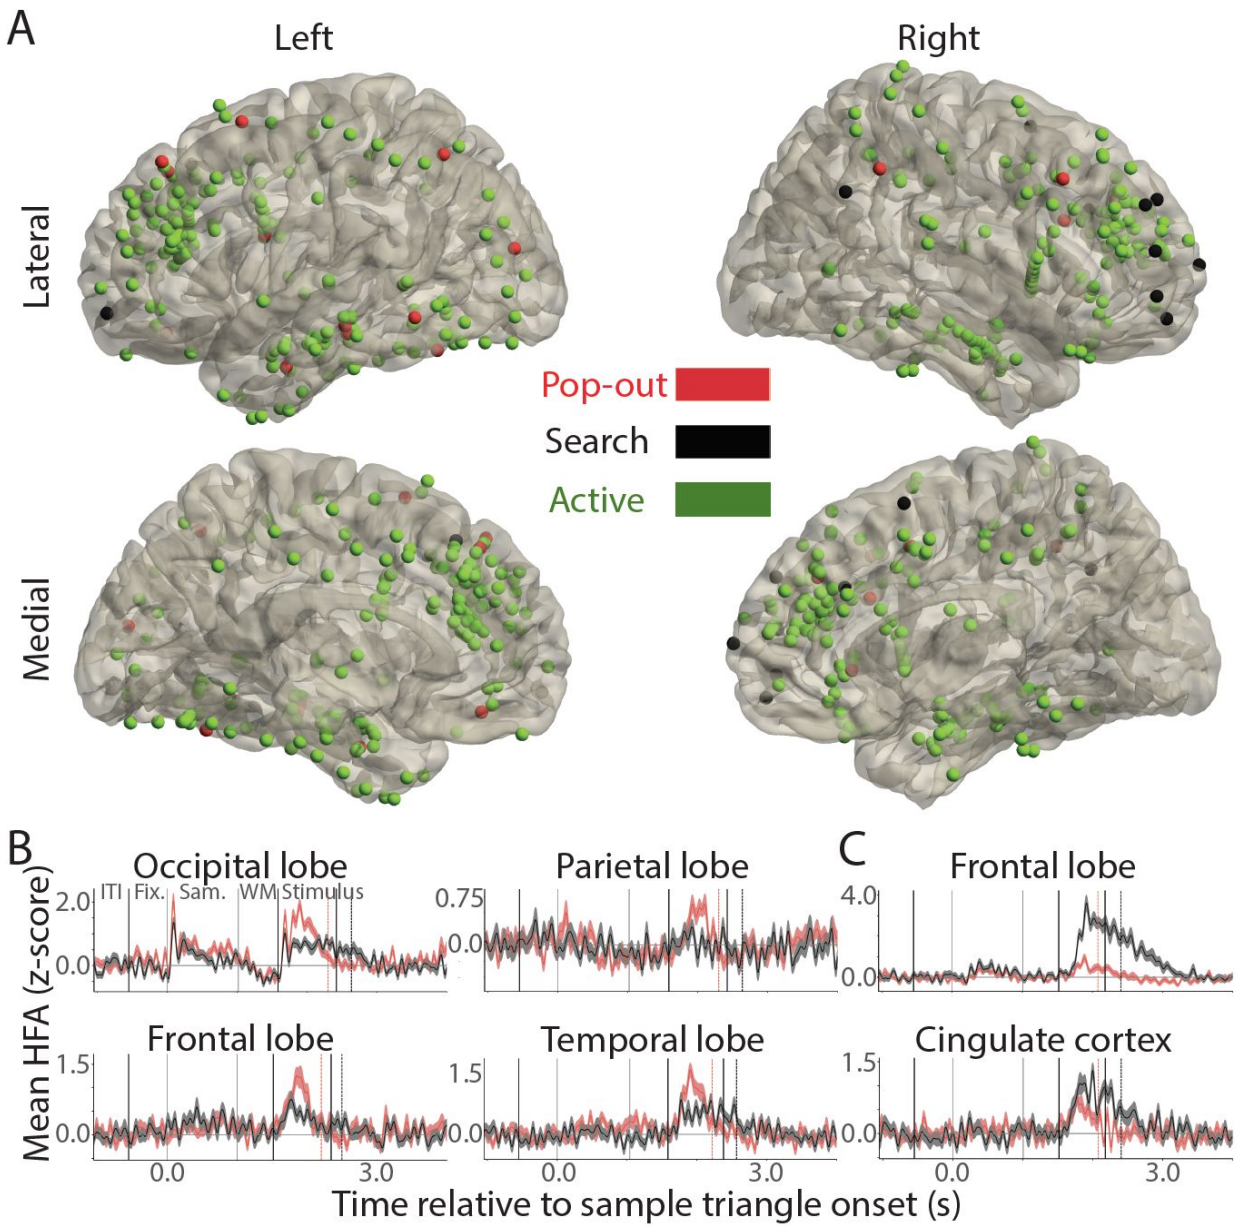

Figure 2.15. Condition-selective electrode sites. A) Pop-out (red) and Search-selective (black) sites are shown, along with task-active (green) sites for which the two conditions did not differ. The brain template (Colin-27) is identical to that shown in Figure 2.11. B) Example time courses of electrode sites that showed a significant Pop-out selective effect. The mean signal in Pop-out is plotted in red, and Search in black. C) Example time courses of electrode sites that showed a significant Search-selective effect. Colors of traces are the same as in (B).

This proportion of condition modulations is slightly smaller than previously reported (Ossandon et al., 2012), possibly owing to different methods and thresholds for determining statistical significance, as well as our explicit exclusion of condition modulations arising from specific RT-related artifacts (see *Method* section, and below, for more details). Condition-related effects were found in all four cortical lobes as well as cingulate cortex. Because of the small number of sites showing condition-related

modulations and their wide anatomical distribution, we do not statistically test these effects as a function of region, but rather report their observed anatomical locations directly.

Electrodes showing greater activity increases in the Pop-out than Search condition (n = 19 electrodes; 10 patients) were found in the parietal lobe (n = 4 electrodes; 4 patients; SPL, IPL, and SMG); frontal lobe (n=8 electrodes; 5 patients; SFG, SFS, MFG, IFG, OFC); temporal lobe (n=6 electrodes; 4 patients; MTG, ITG, insula, and fusiform gyrus), and occipital lobe (n=1 electrode; lateral occipital cortex).

Electrodes showing greater activity increases in the Search than the Pop-out condition (n = 11 electrodes; 5 patients) were predominantly found in frontal cortex (9 electrodes; 4 patients). Of these frontal Search sites, all but one (mPFC) were found in lateral frontal cortex. The remaining two Search electrodes (N=2 patients) were in cingulate cortex (ACC and mid-cingulate cortex).

We removed specific artifactual condition differences, which occurred as a result of the different RTs distributions in the two conditions (see Method for details). Consistent with our simulations (Figures 2.4.-2.6.), we observed three types of artifacts in the data: (1) spurious Search effects in electrodes with sustained activity (Figure 2.7.); (2) spurious Search effects in electrodes with response-related activity (Figure 2.8.); and (3) spurious Pop-out effects in electrodes with response-related activity (Figure 2.9.). We observed 14 electrodes showing the first type of artifact, located in all four cortical lobes, as well as the MTL and amygdala. Six electrodes showed the second type of artifact; these were located in frontal, occipital, and temporal cortex. Fifteen electrodes were characterized by the third type of artifact, and were found in frontal, parietal, and temporal cortex, as well as the cingulate cortex and MTL. It is worth noting that the total proportion of condition effects automatically detected in this study prior to artifact exclusion (65/363 electrodes or 17.9%) are consistent with those previously reported by Ossandon et al. (2012); Ossandon and colleagues also mention the presence of RT-related artifacts in their observed condition differences, but do not explicitly exclude them by type.

In sum, we found condition differences between Pop-out and Search at equivalent rates as those previously reported. Pop-out effects were found in all four cortical lobes, including frontal cortex. Search effects were found in lateral frontal cortex and cingulate cortex.

2.5. Discussion

We used intracranial neural recordings with high spatiotemporal resolution in humans to document two hitherto unknown properties of the visual search network in humans: First, there are marked regional differences in prefrontal engagement in this task, including an increasing gradient of activations from superior to inferior lateral frontal cortex. Second, the medial temporal lobe is engaged on par with the frontoparietal cortices in this classical attention task. We discuss these findings below.

Behavior results

Consistent with past research examining the neural correlates of serial search and pop-out, we observed faster RTs and higher accuracy in the Pop-out than the Search condition (Li et al., 2010; Ossandon et al., 2012; Yan et al., 2016). Moreover, the distributions of RTs in the two conditions were similar to those previously reported (Figure 2.b.; Buschman & Miller, 2007; Li et al., 2010; Ossandon et al., 2012)

Cortical distribution of task-active sites

Strong activation profile in visual areas. Across cortical areas, we observed the strongest activation profile in the occipital lobe (Figure 2.11) as would be expected in a visual experiment. Similarly, we observed strong activation in the ventral stream (Figure 2.14.), consistent with its role in low-level vision and object detection (DiCarlo & Cox, 2007; DiCarlo et al., 2012).

Fronto-parietal network engagement. Beyond visual cortex, the frontal and parietal cortices were robustly active, as predicted by fronto-parietal models of attention allocation in both Pop-out and Search (Corbetta & Shulman, 2002; Buschman & Miller, 2007; Li et al., 2010; and Ossandon et al., 2012). The robust parietal activations in the SPL and IPL (areas adjacent to the IPs), which we observed, were certainly also expected (Gottlieb et al., 1998; Silver et al., 2005; Silver & Kastner, 2009; Li et al., 2010).

Regional differences within prefrontal cortex. The lateral convexity of the human PFC is known to be involved in working memory and controlled attention (Knight et al., 1995; Curtis & D'Esposito, 2003). It is specifically thought to be required for holding representations of sensory stimuli in mind over short periods of time (seconds) while competing distractor items are considered and discarded. This is necessary in the Search condition of the present experiment, and possibly also in the Pop-out condition albeit to a lesser extent. It therefore comes as no surprise that we observed robust activation across the lateral frontal lobe (Figure 2.13.).

Importantly, we observe differences in the proportional activation across subregions of lateral frontal cortex, whereby the IFG was most strongly active, followed by MFG and SFG, in a gradient of decreasing activation from ventral to dorsal subregions. These observations may be interpreted from the perspective of preferential engagement of the ventral attention network in this experiment (Corbetta & Shulman, 2002; Corbetta et al., 2008), though it is not clear why the aggregate activations across both conditions should show this pattern. From the perspective of a working memory role for lateral frontal cortex, several hypotheses exist for the functional specialization of its subregions (reviewed in Curtis & D'Esposito, 2003). Some theories make a “material-specific” distinction whereby dorsal subregions are involved in the maintenance of *spatial* information, while ventral subregions are responsible for the maintenance of *object-related* information. Other models make a “process-specific” distinction whereby dorsal subregions support higher-order control functions like manipulation of items in working memory, while ventral subregions are engaged in “simpler” operations, such as encoding and retrieval of items into and from memory. Our

data cannot disentangle these propositions. It does, however, suggest that a dichotomy between dorsal and ventral subregions of lateral frontal cortex does not capture the complexity of prefrontal function, as we observe a continuum of increasing activity from the superior to middle to inferior frontal gyri in this classical visual search task.

We observed proportionally less activation in the OFC region (Figure 2.13.), consistent with a vast set of work documenting different roles for the lateral frontal cortex versus OFC in behavior (reviewed in Hartikainen & Knight, 2003), along with a prominent role for ventral frontal regions in the default mode network (Shulman et al., 1997; Raichle et al., 2001; Raichle, 2015). Hence, we might expect to find decreases rather than increases in the OFC in this task (Ossandon et al., 2011; Raccah et al., 2018). Lateral temporal cortex (STG and STS), also a component of the default mode network, similarly showed a very low proportional activation rate in our data (Figure 2.14.).

The medial temporal lobe is engaged on par with frontal and parietal cortex. We observed robust MTL activation during the search interval (Figure 2.11.) in this experiment: The proportion active electrodes in the MTL was indistinguishable from the frontal and parietal cortices. We detected active electrode sites in entorhinal cortex, hippocampus and, to a lesser extent, the parahippocampal cortex. This result is consistent with previous work demonstrating a role for the entorhinal grid cell system in free-viewing of natural scenes in stationary macaques (Killian et al., 2012; Killian et al., 2015) as well as evidence demonstrating visual working memory impairment in patients with MTL lesions when verbal rehearsal strategies are prevented (Olson et al., 2006). To the best of our knowledge, this is the first time that MTL engagement, including hippocampus and entorhinal cortex, has been demonstrated in stationary humans performing a controlled, classical visual search paradigm. One previous study has reported individual target-selective cells in the human hippocampus and amygdala while searching a display of naturalistic images (Wang et al., 2018), but no equivalent result has been reported that includes the entorhinal cortex in a classical Search versus Pop-out paradigm. The implication of our findings is that visual attention researchers now need to add MTL structures to the well-known cortical ROIs, such as the IPs, FEF, and dlPFC, for disentangling the mechanisms that enable top-down and bottom-up visual attention in visual search. Additionally, classical “attention” effects in serial search and pop-out may need to be re-examined through through the lens of viewing these search behaviors as navigation behaviors in visual space.

Adjacent to the MTL, we observed an activation rate of 37% in the amygdala. This effect may be driven by its role in detecting novelty and salience (Rutishauser et al., 2006); the observation that its activity is modulated by visual fixations in free-viewing natural images (Minxha et al., 2017); and its role in a broader limbic network including the MTL.

In sum, this work confirms a central role for the frontoparietal attention network in visual search. In addition, we observe prominent regional differences within the PFC, highlighting a clear division of labor between OFC versus lateral PFC, and a graded activation profile across lateral PFC. Moreover, we provide evidence for a central role for the MTL, including the entorhinal cortex and the amygdala, in this classical Search versus Pop-out paradigm. The implementation of visual search in humans appears to

engage widely distributed brain regions: It is not merely the purview of a few isolated regions in the parietal and frontal cortices as previously reported.

Condition-selective effects

Most task-active sites are equally engaged in Pop-out and Search. We find that the vast majority of task-active sites do not show significant condition modulations, consistent with past research documenting a largely overlapping network of regions in the frontal and parietal cortices, which are engaged in both tasks (Leonards et al., 2000; Ossandon et al., 2012; and Katsuki & Constantinidis, 2014). Beyond shared mechanisms in the frontoparietal attention network, it makes sense that the two tasks share large amounts of neural infrastructure for shared subprocesses of the visual search task, including low-level feature detection (edges and colors), object detection, decision-making (whether to select the left or right response), motor planning, etc. We nonetheless detected a sparse, yet robust, set of condition-selective sites, the anatomical distributions of which are noteworthy. We describe these below.

Search-selective sites are concentrated in lateral frontal cortex. Consistent with a framework that proposes that frontal cortex has a critical role in facilitating Search as opposed to Pop-out, (Li et al., 2010; Buschman & Miller, 2007), we find Search-selective sites exclusively in the frontal and cingulate cortices. Specifically, the majority of Search-selective electrodes were located in lateral frontal cortex. While the number of electrodes (n=9) showing significant search-selective effects were few, they occurred across several patients (N=4), and the effects were among the largest we observed in this dataset (see Figure 2.15.c).

This result is consistent with numerous past studies, reporting increased lateral frontal cortex engagement in Search but not Pop-out across a range of recording modalities. fMRI studies have shown increased BOLD signal in regions of lateral frontal cortex in Search tasks relative to Pop-out (SFS: Leonards et al., 2000; MFG and IFG: Anderson et al., 2007; 2010). Other fMRI work have shown a more general role for lateral frontal cortex (MFG) in guiding top-down visual attention (Gazzaley et al., 2007).

In a macaque lesion study, Rossi et al. (2007) showed that ablating the lateral surface of right PFC impaired search performance in a difficult search condition, in which the search cue was often switched across trials, but not in a color pop-out similar to the present study. Partially convergent evidence was reported by Iba and Sawaguchi (2003), who found that a muscimol injection to dlPFC of rhesus macaques caused an equal impairment to a conjunction (difficult) and pop-out search.

TMS studies have, similarly, demonstrated a causal role for dlPFC in enabling performance in Search, but not Pop-out, in humans (Kalla et al., 2009). One of these TMS studies was conducted on an identical Search and Pop-out paradigm as the one employed in the present study (Yan et al., 2016).

A plausible functional explanation for this preferential engagement of lateral frontal cortex in Search is its greater demands on working memory as compared to Pop-out. Indeed, Anderson and colleagues (2010), used a joint working memory (verbal and spatial) and visual search paradigm to demonstrate that the right MFG and IFG were engaged both in working memory and serial search. They also showed that, on a

behavioral level, simultaneous engagement in a working memory task, impaired serial search performance. Hence, our results, which show sparse but robust preferential engagement of the lateral frontal cortex in Search, fit into a broader body of literature. This body of research converges to demonstrate that Search places greater demands on working memory, and that these greater working memory demands are reflected in increased neural activity in lateral frontal cortex.

Pop-out selective sites are distributed across the cortical hierarchy, and include lateral frontal cortex. A number of brain regions have been claimed to be the most critical area for visual pop-out. Several empirical studies highlight a central role for parietal cortex, especially areas adjacent to the IPs (Gottlieb et al., 1998; Buschman & Miller, 2007; Li et al., 2010; Wardak et al., 2010; and Yan et al., 2016).

In contrast, a prominent computational theory (Zhaoping, 2002; Zhaoping & Dayan, 2006; Zhaoping, 2019) and related neural evidence from macaques (Yan et al., 2018), emphasizes the importance of early salience computations in primary visual cortex (V1). Zhaoping's theory does incorporate the possibility of feedback from higher-order visual areas, such as V2, V3, V4, and IT cortex (Zhaoping, 2019). It does not, however, include any putative involvement of higher-order cortical areas, such as the parietal or frontal cortices. It should be noted that salience maps have been observed in several cortical areas beyond V1, including: area V4 in the occipital lobe (Burrows & Moore, 2009), area LIP in parietal cortex (Gottlieb et al., 1998), and the FEF in the frontal lobe (Thompson & Bichot, 2005). Visual salience effects have also been reported in subcortical regions including the pulvinar nucleus of the thalamus (Robinson & Petersen, 1992), and the superior colliculus (White et al., 2017).

In sharp contrast to these bottom-up theories, which depict visual pop-out as reflecting anatomically early detection of visual salience, Hochstein and Ahissar (2002), proposed the Reverse Hierarchy Theory (RHT) for vision. According to this theory, the visual pop-out phenomenon should be implemented in high-order cortical areas such as frontal cortex, where large receptive fields can account for the broad spread of visual attention necessary for parallel detection of a salient object. This is consistent with viewing Pop-out as a high-level object detection task (Hochstein & Ahissar 2002; see also Nakayama & Martini, 2011), wherein the pop-out effect is invariant to object size on the retina, unlike the view of Pop-out as detection of low-level feature anomalies based on lateral inhibition in V1 (Zhaoping, 2002; Zhaoping & Dayan, 2006; Zhaoping, 2019). Consistent with RHT, two macaque studies using fMRI and single-neuron recordings respectively, have reported lateral frontal cortex responses to visual pop-out (Wardak et al., 2010 using fMRI; Katsuki & Constantinidis, 2012 using single-neuron recordings).

Finally, a theory for cortical engagement in visual pop-out has been proposed by Corbetta & Shulman (2002), and extended by Corbetta et al. (2008). This theory proposes that bottom-up capture of attention by a salient stimulus, as in visual pop-out, is accomplished by a distributed network of cortical regions, notably all in the right hemisphere, including the TPJ as well as parts of the MFG, IFG, frontal operculum and anterior insula (the right-lateralized ventral attention network, VAN).

Our data reveals a distributed set of sites selective for Pop-out, in partial agreement with each of the extant theories. We observed several Pop-out selective

sites in the parietal lobe as predicted by past research (though not all were adjacent to the IPs). A single electrode site in the occipital lobe was selective for Pop-out; we refrain from interpreting this result in light of the limited coverage in this area. Temporal lobe pop-out selectivity can be understood from the perspective of the importance of ventral temporal regions for object detection (DiCarlo & Cox, 2007; DiCarlo et al., 2012). Similarly, Pop-out selectivity in ventral temporal cortex has previously been observed in other color pop-out paradigms (Wardak et al., 2010; Ossandon et al., 2012).

Our most striking finding is the presence of Pop-out selective sites in frontal cortex. These sites were detected not merely in areas at or near the FEF, which has a known role in detecting salience and planning eye movements, but also in lateral frontal cortex and OFC. This result is consistent with the view that top-down processing has a prominent role in implementing visual pop-out (Hochstein & Ahissar, 2002; Nakayama & Martini, 2011). It is inconsistent with a view of visual pop-out as exclusively implemented in V1 and adjacent occipital lobe areas (Zhaoping, 2002; Zhaoping & Dayan, 2006; Zhaoping, 2019), and suggests that such theories need to be substantially expanded. Our result may explain Iba and Sawaguchi's (2003) observation that a muscimol injection to dlPFC of monkeys caused impairments in both pop-out and serial search. Future causal investigations of the role of lateral frontal cortex in Pop-out and Search may find differential impairment in the two conditions if the sites of the temporary lesion or stimulation is selected with greater anatomical precision within dlPFC than has been possible in the past. Furthermore, our result is consistent with two previous functional studies in macaques which also reported the presence of Pop-out selective sites in lateral frontal cortex (Wardak et al., 2010; Katsuki & Constantinidis, 2012). To the best of our knowledge, this is the first time that such sites have been detected in humans.

Finally, while our data is consistent with the view that Pop-out is implemented in a distributed set of sites across cortex, they do not fit neatly with the framework espousing a right-lateralized ventral attention network for salience-driven orienting of attention (Corbetta & Shulman, 2002; Corbetta et al., 2008), since we also observe robust Pop-out selective effects in the left hemisphere (Figure 2.15.). In sum, our results are at odds with any exclusivist claims about regional engagement in Pop-out, such as the view that Pop-out is solely a bottom-up phenomenon, which does not require frontal cortex. We have demonstrated that the known presence of salience maps in low-level visual regions, as well as parietal cortex and the frontal eye fields, is not mutually exclusive with the simultaneous rapid engagement of a high-level cognitive region, the lateral frontal cortex, in visual pop-out.

Conclusions

The present study documents two previously unreported, putative neural mechanisms subserving visual search in humans. First, lateral PFC engagement shows regional differences with greater activation in inferior than superior regions, and clearly diminished proportional activation in OFC. Second, we establish a central role for the medial temporal lobe - including the hippocampus, entorhinal cortex, and amygdala - in

visual search in humans. This study therefore represents a major step towards a more complete understanding of the neural mechanisms enabling this fundamental perceptual behavior.

Chapter 3

How Does Expectation Shape Object-Based Attentional Selection?

3.0. Foreword

In this chapter, I provide a theoretical contribution to our understanding of attention and expectation, and the extent to which these two constructs converge. The overarching question asked here is “Is expectation attention?” The chapter is based on the following publication:
Slama, S. J. K., & Helfrich, R. F. (2017). How Does Expectation Shape Object-Based Attentional Selection? *The Journal of Neuroscience*, *37*(17), 4427–4429. https://doi.org/10.1523/jneurosci.0414-17.2017

I am grateful to my co-author Randolph, who helped me navigate this business of publishing a paper. I am further indebted to my advisor, Bob, for supporting my publishing this paper despite the resulting time away from my main research.

Finally, I am grateful to the following highly insightful people. Each of these individuals, in large and small ways, inspired me to think better: Falk Lieder, Jesse Livezey, Robert Nishihara, Sebastian Musslick, Jacob Miller, and Rika Antonova.

3.1. How Does Expectation Shape Object-Based Attentional Selection?

In 1984, Duncan (1984) demonstrated that visual attention prioritizes whole objects: subjects performed better when reporting two features on the same object than when reporting one feature from each of two adjacent objects. This marked the beginning of the field of object-based attention research (Egly et al., 1994). Later studies showed that the neural (fMRI) representation of multiple attributes belonging to a target object is enhanced relative to those of a control object (O’Craven et al., 1999), and that object-based attention has a temporal dimension in sampling the target and control objects at different frequencies (Fiebelkorn et al., 2013).

The view of attention as a homogeneous construct has recently been called into question. Attention is now viewed as an umbrella term for multiple subprocesses [feature-based attention (Treisman and Gelade 1980); spatial attention (Posner, 1980); top-down vs bottom-up attention (Buschman and Miller, 2007); and sustained attention (Rosenberg et al., 2016)]. Object-based attention (Egly et al., 1994), as described above, is one of these subprocesses. This heterogeneity in the attention literature has prompted a number of scientists (Chun et al., 2011) to advocate for taxonomizing the construct to facilitate greater precision in future research.

In a series of recent articles, Summerfield and Egner (2009, 2016) propose that not only is the attention literature heterogeneous in consisting of multiple subcomponents, but that some experiments designed to investigate attention may confound attention itself with the conceptually distinct construct of expectation. Summerfield and Egner (2009, 2016) postulate that expectation is driven by information

about the probability of an upcoming event while attention is driven by information about its relevance. In contrast to the straightforward conceptual difference between attention and expectation, the difference in experimental design required to induce each of these two cognitive processes is subtle. In both attention and expectation experiments, a cue provides information about an upcoming target with some probability, and the presence of this cue facilitates performance (Summerfield and Egner, 2016). Summerfield and Egner (2009, 2016) argue that if the cue provides information about the probability of occurrence of an upcoming event and the task of the subject is simply to indicate whether the event occurred, then the experiment elicits a neurocognitive expectation process. If the experiment is instead designed so that the cue indicates which of multiple dimensions is task relevant and the task of the subject is to respond to some feature within the relevant dimension, then the experiment elicits an attention mechanism (Summerfield and Egner, 2009; 2016). Summerfield and Egner (2009, 2016) further speculate that expectation and attention might rely on distinct neural mechanisms: an expected event in many cases results in decreased neural processing, while an attended event results in increased processing (Summerfield and Egner, 2009). This recasting of portions of the attention literature as expectation enables interpretation of experimental results from the perspective of predictive coding, a computational theory of how the brain facilitates perception from sensation by minimizing the prediction error between expected and received sensory input (Huang and Rao, 2011).

If the central claim by Summerfield and Egner (2009, 2016), that a subtle adaptation to existing attention experiments induces a distinct but parallel process to attention, is correct, then this raises the question of whether expectation, like attention, is also a heterogeneous construct. This theory creates the opportunity to branch out new expectation subfields in parallel to the more mature subdomains of attention research (e.g., feature based; spatial; top-down vs bottom-up; object-based; and sustained expectation). The field of feature-based expectation has already been introduced (Summerfield and Egner, 2016). Another possible new subfield could be object-based expectation, paralleling the existing field of object-based attention (Jiang et al., 2016). An open question within this proposed framework is whether the expectation statuses of individual features interact to form object-level expectation.

In a recent article in The Journal of Neuroscience, Jiang et al. (2016) used a computational model to simulate three competing models for how expectations about one object feature affects expectations about another (Jiang et al., 2016, their Fig. 3), as follows: (1) expectations do not spread from one object feature to the other; (2) a prediction error in one feature spreads to the other feature to render the whole object unexpected (the reconciliation hypothesis); and (3) the expectation status of one feature repels the expectation status of the other feature, thereby promoting the perceptual inference that the two features belong to separate objects (the segregation hypothesis).

The authors performed a behavioral experiment in which subjects were presented with a cloud of moving dots, which were either red or green (color dimension) and were moving either up or down (motion dimension; Jiang et al., 2016, their Fig. 1a). A trial-by-trial auditory cue generated an expectation of the upcoming feature values with 75% validity: the timbre signaled the upcoming color, and the pitch direction

signaled the upcoming motion direction. When subjects were instructed to allocate sustained attention to one feature (color), the trial-by-trial expectation cue was associated with a behavioral benefit in response times not only to the attended feature, but also to the unattended feature (motion; Jiang et al., 2016, their Fig. 1c). This effect was interpreted as evidence for cross-feature spread of expectation in support of the reconciliation hypothesis.

To examine the neural substrates of this effect, Jiang et al. (2016) performed the same experiment using fMRI. They analyzed their data using intersubject multivoxel pattern analysis (for review, see Haxby et al., 2014), comparing the voxel-by-voxel pattern of activation between conditions in different regions of the brain. In one analysis, a linear support vector machine classifier was trained to distinguish activity patterns in early visual cortex resulting from two conditions in which expectation was consistent (color and motion both expected vs color and motion both unexpected). In a second analysis, the same classifier was trained to distinguish between two conditions where expectation was inconsistent between features (color expected and motion unexpected vs color unexpected and motion expected). The classifier showed greater accuracy when separating the two consistent conditions than when separating the two inconsistent conditions (Jiang et al., 2016, their Fig. 4f,g). This result matched the predictions of the reconciliation hypothesis and neither of the alternative hypotheses (Jiang et al., 2016, their Fig. 4a– c). The authors concluded that objects are the unit of selection, not only for attention, but also for expectation.

In considering the behavior results, readers with a psychology background will note an analog in the classical Stroop (1935) and Simon effects (Simon and Wolf, 1963): performance on experimental tasks that require subjects to keep track of a conjunction of inconsistent features is worse than performance on tasks that require subjects to keep track of only a single feature. These behavioral effects indicating interference by inconsistent feature conditions are common in the psychology literature and need not be interpreted as object-level perceptual selection.

In considering the fMRI results, we note that, to use a relative comparison of the performance of two classifiers to draw conclusions about brain function, the two classifiers must have an equal chance at separating the data, aside from the hypothesis to be tested. If we assume that early visual cortex contains neural ensembles that separately encode color and motion, then, from a predictive coding perspective (Rao and Ballard, 1999; Alink et al., 2010), the color ensemble will be highly active when the color is unexpected and the motion ensemble will be highly active when the motion direction is unexpected. The classifier for the two consistent conditions must then separate the state of cortex in which both ensembles are highly active from the state in which neither ensemble is highly active. This could in principle reduce to a comparison as to whether visual cortex is active versus inactive. The classifier for the inconsistent conditions, however, must separate the two intermediate states where the color ensemble is highly active and the motion ensemble is not from the state where the motion ensemble is highly active and the color ensemble is not. This would entail learning to separate complex spatiotemporal patterns of activity in the motion and color ensembles, a task that is more difficult than separating high from low activity. This scenario could explain the result of the classifier comparison without making inferences

about cross-feature spread of expectation in early visual cortex. Therefore, the presented conclusion may not be the only interpretation of these data.

A simple control could have resolved this issue: the inclusion of a second object, such as an additional spatially separate cloud of dots. Experiments on object-level spread of attention (Duncan, 1984; Egly et al., 1994; O'Craven et al., 1999; Fiebelkorn et al., 2013) typically include a control object for the purpose of demonstrating that attention does not spread there. A control object would provide a stronger test of the three hypotheses: no spread to a control object would support the object-based interpretation, but, without a control object, the results are difficult to interpret. For instance, they might indicate spatial spread of expectation.

Future experiments on object-based expectation would further benefit from considering the role of time, a feature that has recently started to gain traction in the attention literature (Buschman & Kastner, 2015; Denison et al., 2017). From the perspective of a perceptual decision-making framework (Kayser et al., 2010), it is possible that probability- and relevance-driven selection interact over the time course of a perceptual decision. It is known that object-based attention alternates between a cued and uncued location on the same object at a faster frequency (8 Hz) than between a cued location on one object and an uncued location on a different object (4 Hz; Fiebelkorn et al., 2013). Rhythmic sampling on subsecond timescales has also been demonstrated for spatial attention (Landau & Fries, 2012), raising the possibility that rhythmic brain activity may support multiple forms of attention. One hypothesis is that expectation may align the phase of ongoing rhythmic attentional sampling to optimize cortical excitability at the time of an expected event. Neural oscillations may disentangle these rapid endogenous prioritization processes during feature integration (Helfrich & Knight, 2016). Delineating the spatiotemporal progression of these rapid mechanisms will necessitate the use of recording methods with greater temporal resolution than fMRI, such as electroencephalography, magnetoencephalography, and electrocorticography. In the meantime, the present study (Jiang et al., 2016) is a step in the right direction toward understanding the multiple processes that are collectively called attention.

Chapter 4

Closing remarks

4.1. Conclusions

This dissertation asked two overarching questions pertaining to the broad construct of attention: In **Chapter 2**, I asked: *'Is attention navigation?*' In **Chapter 3**, I asked '*Is expectation attention?*' In response to the first question, this dissertation work provides evidence consistent with the view that one attentional behavior, visual search, has a navigation component; this is based on the robust observed MTL activations. The answer to the second question is 'maybe'. To find out, we need to create experiments that truly disentangle attentional prioritization processes from purely predictive processes without confounds pertaining to the type of attention or expectation elicited.

4.2. Alternative analysis approaches

In **Chapter 2**, I applied a cluster-based permutation test approach to identify task-selective electrodes and condition-based effects. The starting point was a simple 100-ms threshold of significant activity increases. In addition to the arbitrary selection of a duration threshold, this approach was very sensitive to filtering choices. I used a machine learning approach (time-resolved logistic regression classifier; Figure 4.1.), to demonstrate that task-relevant information still remained at "inactive" electrode sites when using the100-ms cutoff method.

A)

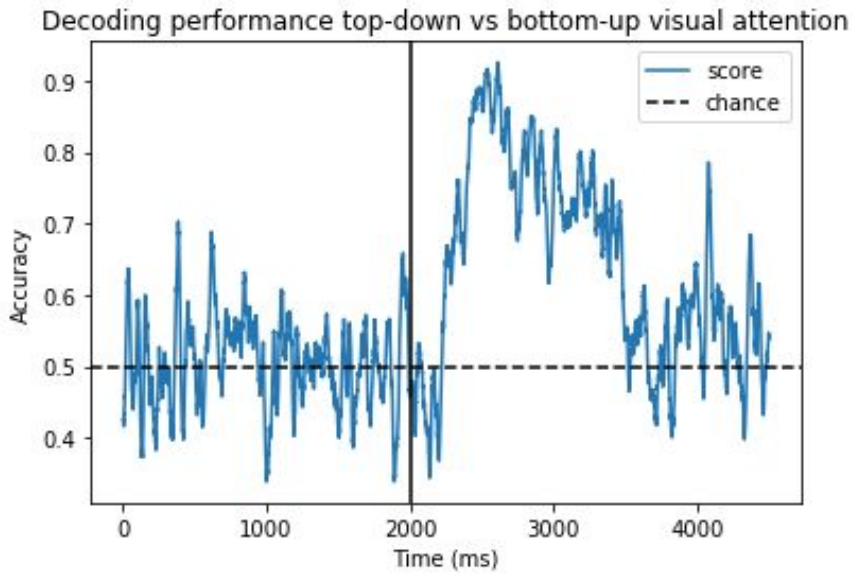

B)

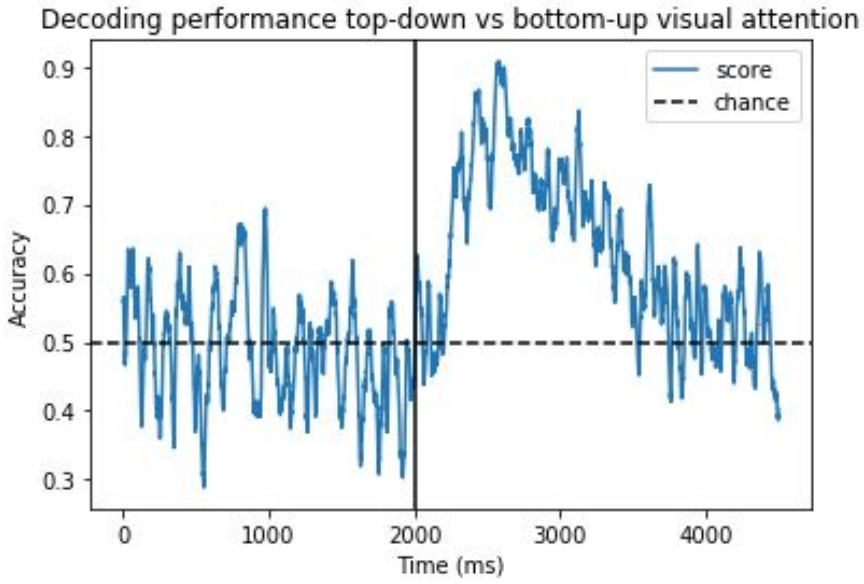

C)

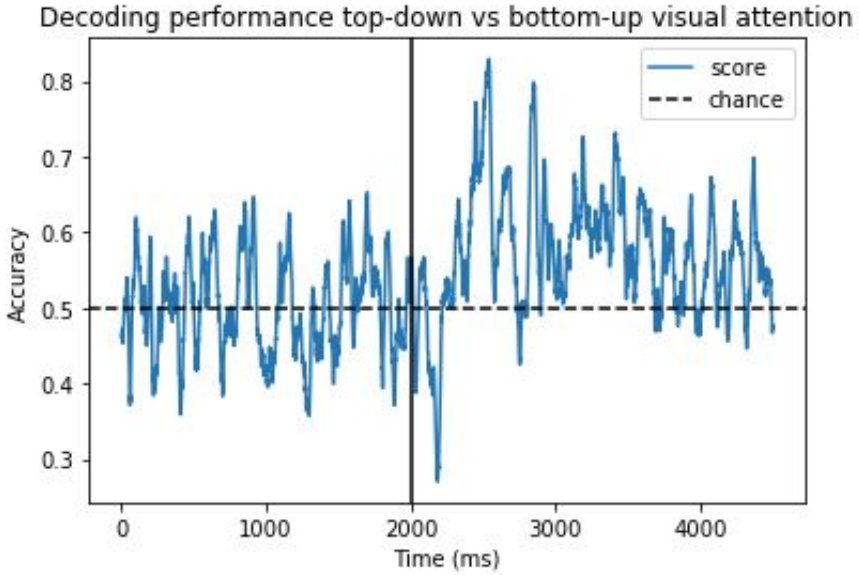

Figure 4.1. Decoding performance over time of Search vs. Pop-out for one patient. Decoding was conducted using a logistic regression analysis at each time-step with the HFA signal across electrodes as the predictors, and attention condition as the outcome. Performance (ordinate axis) was evaluated using the ROC-AUC metric (Receiver Operating Characteristic, Area Under the Curve). A) Classification performance using all electrodes. B) Classification performance including only “active” electrodes according to an early task-selective criterion (100 ms consecutive significance). C) Classification performance including only “passive electrodes” according to the same criterion. This analysis demonstrated that there was remaining information about the experimental conditions in electrodes, which had previously been excluded as not being task-active (“passive”), thereby motivating a less stringent pipeline.

This helped motivate the need for a more sensitive detection algorithm for task-active electrodes. The cluster-based permutation test, which I adapted from Maris & Oostenveld (2007) in collaboration with Sujayam Saha, was a substantial improvement over the 100 ms threshold, as it was able to detect brief but consistent activity increases. Nevertheless, if I were to start my dissertation work from the beginning today, I might take a different approach.

First, the experiment includes intervals beyond the "neutral" baseline and the task interval of interest, where patients actively search for a target. These other intervals could be exploited to obtain a richer characterization of the time courses of the HFA responses at various sites. For example, an electrode that shows an activity increase only during the search interval, should perhaps not be viewed as equivalent to an electrode that responds to the search display *and* the sample triangle, or an electrode that responds to the search display, sample triangle, and a color change in the fixation cross. Hence, in a future analysis, it may be of interest to featurize the task in greater detail, and use a regression approach, perhaps complemented by dimensionality reduction and clustering, to obtain a richer description of the response profiles in the dataset.

Second, the cluster-based permutation test alone is not well-suited to handling condition differences with different response times, as discussed in the *Method* section of **Chapter 2**. A simple solution, which I experimented with, is to compare average activity on a single-trial basis between the search display onset and the response. However, this can give rise to other artifactual condition differences in electrodes with equivalent response profiles between the two conditions: For example, consider an electrode site with sharp stimulus-locked responses in both conditions. Such a site would be labeled as showing a Pop-out effect simply as a consequence of the shorter trials over which we would be averaging with a mean-based approach. In Chapter 2, I explicitly simulate several types of artifacts that occur with the cluster-based pipeline, and exclude these artifacts based on visual inspection. A more objective approach to obtain equivalent results would be to use some variant of dynamic time warping to align the signals in the two conditions prior to computing condition differences. Dynamic time-warping itself is prone to overfitting, since it allows remapping of each individual time point. A more appropriate approach might be to apply various constraints to the types of alignment allowed by the algorithm: I have conducted some early explorations of linear time-warping (Williams et al., 2019) for this purpose and, in a future analysis, this may become my preferred method.

Third, the choice of re-referencing can highly affect the results of intracranial analyses including, at worst, injecting activity patterns from one area into another, thereby giving rise to misleading conclusions (Ludovic Bellier, personal communication). If I were to start over, I would compare the effects of different referencing schemes on my data, and thereby ascertain that the conclusions do not depend on a single referencing scheme.

4.3. Immediate follow-up to the current work

In **Chapter 2**, I describe insights gained from a first set of analyses of an extremely rich dataset of intracranial recordings. This same dataset can be used to ask many more questions. I describe my immediate analysis plans for the data below.

First, I will investigate potential lateralization effects in the existing results. There are reasons to believe that both the cortical (Corbetta et al. 2002; 2008; 2011) and medial temporal lobe (Jacobs et al., 2010; Miller et al., 2018) effects would be right-lateralized given the spatial component of the task.

Second, I will investigate onset times across the cortical hierarchy within the pop-out condition. This would provide a richer understanding of activity propagation in this fascinating mode of searching the visual environment: Does it begin from the top (i.e., frontal cortex)? From the bottom (i.e., visual and parietal cortex)? Or might it originate in a subcortical source, such as the medial temporal lobe or thalamic nuclei?

Third, I will investigate the anatomical distribution of HFA decreases across cortex. Early exploratory analyses show robust decreases in several areas. Based on previous work, I predict that these decreases will overlap with the default mode network (Ossandon et al., 2011). A competing theory, however, suggests that cortical activity decreases occur in a more distributed network of regions, and includes local decrease effects with complex spectrotemporal response profiles (Ramot et al., 2012). A third line of work suggests that frontal HFA decreases may play a precise role in implementing temporal prediction of upcoming stimuli (Durschmid et al., 2018). This interpretation may be relevant to the search task employed in Chapter 2, where the Search condition is believed to elicit a serial visual process that unfolds over time as participants inspect and reject candidate objects as distractors. This serial process may engender temporal prediction mechanisms in the brain.

4.4. Broader questions to guide future work

Following the current work, a number of future directions would be exciting for myself or others to explore. One direction is that of learning how to search visual space during development, and how potential interactions between search and reward might play into that process. Perhaps learning how to search visual space can be understood from the perspective of a reinforcement learning framework.

Another interesting future set of experiments could involve parametrically varying the search task on a continuum between search (locating an object in space based on its features) and object identification (being able to name an object based on its features) to assess whether the neural processes responsible for object localization versus object identification also form a continuum (Nakayama & Martini, 2011).

Search processes occur in many domains beyond vision. Numerous algorithms are known in computer science and AI for efficiently searching space depending on various constraints of the space and the task. It would be interesting to explore to what extent known search algorithms might be able to explain search behavior in the vision and navigation domains, and to what extent such search behavior maps onto more abstract domains such as internal search of memory (e.g., Kanerva, 2009).

Relatedly, the relationship between visual search and navigation, which is gaining increasing favor in the literature (Meister & Buffalo, 2016), may offer a path toward bridging the roles of the MTL in navigation versus in memory for abstract concepts. This framework may pave the way toward a concrete approach to investigating the cognitive map (Tolman, 1948; Schiller et al., 2015). This direction seems especially promising in light of a parallel literature that is beginning to document relationships between eye movements and non-visual cognition (Benedek et al., 2017).

These, and many more directions of investigation, hold a great deal of promise for improving our understanding of search and its broader role in human cognition.

"*[...] 'tis as 'tis, why it might have been worse, and I feel my thanks accordingly.*"
Far From the Madding Crowd by Thomas Hardy (1874)

References

Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L (2010) Stimulus predictability reduces responses in primary visual cortex. J Neurosci 30:2960 –2966.

Anderson, B. (2011). There is no such thing as attention. *Frontiers in Psychology*, *2*(246), 1–8. https://doi.org/10.3389/fpsyg.2011.00246

Anderson, E. J., Mannan, S. K., Husain, M., Rees, G., Sumner, P., Mort, D. J., McRobbie, D., Kennard, C. (2007). Involvement of prefrontal cortex in visual search. *Experimental Brain Research*, *180*(2), 289–302. https://doi.org/10.1007/s00221-007-0860-0

Anderson, E. J., Mannan, S. K., Rees, G., Sumner, P., & Kennard, C. (2010). Overlapping functional anatomy for working memory and visual search. *Experimental Brain Research*, *200*(1), 91–107. https://doi.org/10.1007/s00221-009-2000-5

Angelucci, A., & Bullier, J. (2003). Reaching beyond the classical receptive field of V1 neurons: Horizontal or feedback axons? *Journal of Physiology Paris*, *97*(2–3), 141–154. https://doi.org/10.1016/j.jphysparis.2003.09.001

Azevedo, F.A., Carvalho, L.R., Grinberg, L.T., Farfel, J.M., Ferretti, R.E., Leite, R.E., Jacob Filho, W., Lent, R., & Herculano-Houzel, S. (2009). Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Neurology*, *513*(5), 532–541.

Benedek, M., Stoiser, R., Walcher, S., & Körner, C. (2017). Eye behavior associated with internally versus externally directed cognition. *Frontiers in Psychology*, *8*(JUN), 1–9. https://doi.org/10.3389/fpsyg.2017.0109

Burrows, B. E., & Moore, T. (2009). Influence and limitations of popout in the selection of salient visual stimuli by area V4 neurons. *Journal of Neuroscience*, *29*(48), 15169–15177. https://doi.org/10.1523/JNEUROSCI.3710-09.2009

Buschman, T. J., & Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*, *315*, 1860–1862. https://doi.org/10.1126/science.1138071

Buschman T.J., & Kastner, S. (2015) From behavior to neural dynamics: an integrated theory of attention. Neuron 88:127–144.

Burke, J. F., Zaghloul, K. A., Jacobs, J., Williams, R. B., Sperling, M. R., Sharan, A. D., & Kahana, M. J. (2013). Synchronous and asynchronous theta and gamma activity during episodic memory formation. *Journal of Neuroscience*, *33*(1), 292–304. https://doi.org/10.1523/jneurosci.2057-12.2013

Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*, 73–101. https://doi.org/10.1146/annurev.psych.093008.100427

Collins, D. L., Zijdenbos, A. P., Kollokian, V., Sled, J. G., Kabani, N. J., Holmes, C. J., & Evans, A. C. (1998). Design and Construction of a Realistic Digital Brain Phantom. *IEEE Transactions on Medical Imaging*, *17*(3), 463–468.

Corbetta, M., Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: from environment to theory of mind. *Neuron*, *58*(3), 306–324. https://doi.org/10.1016/j.neuron.2008.04.017

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215. https://doi.org/10.1038/nrn755

Corbetta, M., & Shulman, G. L. (2011). Spatial Neglect and Attention Networks. *Annual Review of Neuroscience* (Vol. 34). http://doi.org/10.1146/annurev-neuro-061010-113731

Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*(9), 415–423. https://doi.org/10.1016/S1364-6613(03)00197-9

Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical Surface-Based Analysis: I. Segmentation and Surface Reconstruction. NeuroImage, 9(2), 179–194. https://doi.org/https://doi.org/10.1006/nimg.1998.0395

de Graaf, R., Kessler, R. C., Fayyad, J., ten Have, M., Alonso, J., Angermeyer, M., Borges, G., Demyttenaere, K., Gasquet, I., de Girolamo, G., Haro, J. M., Jin, R., Karam, E. G., Ormel, J., Posada-Villa, J. (2008). The prevalence and effects of adult attention-deficit/hyperactivity disorder (ADHD) on the performance of workers: results from the WHO World Mental Health Survey Initiative. *Occupational and Environmental Medicine*, *65*(12), 835–842. https://doi.org/10.1136/oem.2007.038448

Denison R. N., Heeger D. J., & Carrasco, M. (2017) Attention flexibly trades off across points in time. Psychon Bull Rev. Advance online publication. Retrieved April 3, 2017. doi:10.3758/s13423-016-1216-1.

DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, *11*(8), 333–341. https://doi.org/10.1016/j.tics.2007.06.010

DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How does the brain solve visual object recognition? *Neuron*, *73*(3), 415–434. https://doi.org/10.1016/j.neuron.2012.01.010

Doshi, J. A., Hodgkins, P., Kahle, J., Sikirica, V., Cangelosi, M. J., Setyawan, J., Erder, M. H., & Neumann, P. J. (2012). Economic Impact of Childhood and Adult Attention-Deficit/Hyperactivity Disorder in the United States. *Journal of the American Academy of Child & Adolescent Psychiatry*, *51*(10), 990-1002.e2. https://doi.org/10.1016/j.jaac.2012.07.008

Duncan J (1984) Selective attention and the organization of visual information. J Exp Psychol 113:501–517

Dürschmid, S., Reichert, C., Hinrichs, H., Heinze, H.-J., Kirsch, H. E., Knight, R. T., & Deouell, L. Y. (2018). Direct Evidence for Prediction Signals in Frontal Cortex Independent of Prediction Error. *Cerebral Cortex*. https://doi.org/10.1093/cercor/bhy331

Dykstra, A. R., Chan, A. M., Quinn, B. T., Zepeda, R., Keller, C. J., Cormier, J., Madsen, J. R., Eskandar, E. N., Cash, S. S. (2012). Individualized localization and cortical surface-based registration of intracranial electrodes. *NeuroImage*, *59*, 3563–3570. https://doi.org/10.1016/j.neuroimage.2011.11.046

Egly, R., Driver, J., & Rafal, R.D. (1994) Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. J Exp Psychol 123:161–177.

Eimer, M. (2014). The neural basis of attentional control in visual search. *Trends in Cognitive Sciences*, *18*(10), 526–535. https://doi.org/https://doi.org/10.1016/j.tics.2014.05.005

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

Fiebelkorn, I. C., Saalmann, Y. B., & Kastner, S. (2013) Rhythmic sampling within and between objects despite sustained attention at a cued location. Curr Biol 23:2553–2558.

Gottlieb, J. P., Kusunoki, M., & Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, *391*(6666), 481–484. https://doi.org/10.1038/35135

Gazzaley, A., Rissman, J., Cooney, J., Rutman, A., Seibert, T., Clapp, W., & D'Esposito, M. (2007). Functional interactions between prefrontal and visual association cortex contribute to top-down modulation of visual processing. *Cerebral Cortex*, *17*(SUPPL. 1), 125–135. https://doi.org/10.1093/cercor/bhm113

Haller, M., Case, J., Crone, N. E., Chang, E. F., King-Stephens, D., Laxer, K. D., Weber, P. B., Parvizi, J., Knight, R. T., & Shestyuk, A. Y. (2018). Persistent neuronal activity in human prefrontal cortex links perception and action. *Nature Human Behavior*, *2*(1), 80–91. https://doi.org/10.1038/s41562-017-0267-2.Persistent

Hardy, T. (1874). *Far From the Madding Crowd*. London: Smith, Elder & Co.

Hartikainen, K. M., & Knight, R. T. (2003). Lateral and Orbital Prefrontal Cortex Contributions to Attention. In J. Polich (Ed.), *Detection of Change* (pp. 99–116). New York: Springer Science and Business Media. https://doi.org/10.1007/978-1-4615-0294-4_6

Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014) Decoding neural representational spaces using multivariate pattern analysis. Annu Rev Neurosci 37:435– 456.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Vol. 2016-Decem, pp. 770–778). https://doi.org/10.1109/CVPR.2016.90

Helfrich, R. F., & Knight, R. T. (2016) Oscillatory dynamics of prefrontal cognitive control. Trends Cogn Sci 20:916 –930.

Hermes, D., Miller, K. J., Noordmans, H. J., Vansteensel, M. J., & Ramsey, N. F. (2010). Automated electrocorticographic electrode localization on individually rendered brain surfaces. *Journal of Neuroscience Methods*, *185*, 293–298. https://doi.org/10.1016/j.jneumeth.2009.10.005

Hochstein, S., & Ahissar, M. (2002). View from the Top: Hierarchies and Reverse Hierarchies in the Visual System. *Neuron*, *36*, 791–804.

Huang, Y. & Rao, R. P. (2011) Predictive coding. Wiley Interdiscip Rev Cogn Sci 2:580–593.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology*, *148*(1), 574–591.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *Journal of Physiology*, *160*(1), 106–154.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, *195*(1), 215–243.

Iba, M., & Sawaguchi, T. (2003). Involvement of the dorsolateral prefrontal cortex of monkeys in visuospatial target selection. *Journal of Neurophysiology*, *89*(1), 587–599. https://doi.org/10.1152/jn.00148.2002

Jacobs, J., Korolev, I., Caplan, J. B., Ekstrom, A. D., Litt, B., Baltuch, G., Fried, I., Schulze-Bonhage, A., Madsen, J. R. & Kahana, M. J. (2010). Right-lateralized brain

oscillations in human spatial navigation. *Journal of Cognitive Neuroscience*, *22*(5), 824–836. https://doi.org/10.1162/jocn.2009.21240

Jafarpour, A., Griffin, S., Lin, J. J., & Knight, R. T. (2019). Medial Orbitofrontal Cortex, Dorsolateral Prefrontal Cortex, and Hippocampus Differentially Represent the Event Saliency. *Journal of Cognitive Neuroscience*, *31*(6), 874–884. https://doi.org/10.1162/jocn

James, W. (1890). *Principles of psychology*. H. Holt and Company.

Jiang, J., Summerfield, C., & Egner, T. (2016) Visual prediction error spreads across object features in human visual cortex. J Neurosci 36:12746–12763.

Julesz, B. (1981). Textons, the elements of texture perception, and their interactions. *Nature*, *290*(5802), 91–97. https://doi.org/10.1038/290091a0

Kalla, R., Muggleton, N. G., Cowey, A., & Walsh, V. (2009). Human dorsolateral prefrontal cortex is involved in visual search for conjunctions but not features: A theta TMS study. *Cortex*, *45*(9), 1085–1090. https://doi.org/10.1016/j.cortex.2009.01.005

Kandel, E. R., Schwartz, J. H., Jessell, T. M., Siegelbaum, S. A., & Hudspeth, A. J. (2013). *Principles of Neural Science* (5th ed.). New York: The McGraw-Hill Companies, Inc.

Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive Computation*, *1*(2), 139–159. https://doi.org/10.1007/s12559-009-9009-8

Kastner, S., DeSimone, K., Konen, C. S., Szczepanski, S. M., Weiner, K. S., & Schneider, K. A. (2007). Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *Journal of Neurophysiology*, *97*(5), 3494–3507. https://doi.org/10.1152/jn.00010.2007

Katsuki, F., & Constantinidis, C. (2012). Early involvement of prefrontal cortex in visual bottom up attention. *Nature Neuroscience*, *15*(8), 1160–1168. http://doi.org/10.1038/nn.3164

Katsuki, F., & Constantinidis, C. (2014). Bottom-up and top-down attention: Different processes and overlapping neural systems. *Neuroscientist*, *20*(5), 509–521. https://doi.org/10.1177/1073858413514136

Kayser, A.S., Buchsbaum B.R., Erickson D.T., & D'Esposito M. (2010) The functional anatomy of a perceptual decision in the human brain. J Neurophysiol 103:1179–1194.

Killian, N. J., Jutras, M. J., & Buffalo, E. A. (2012). A map of visual space in the primate entorhinal cortex. *Nature*, *491*(7426), 761–764. https://doi.org/10.1038/nature11587

Killian, N. J., Potter, S. M., & Buffalo, E. A. (2015). Saccade direction encoding in the primate entorhinal cortex during visual exploration. *Proceedings of the National Academy of Sciences*, 201417059. https://doi.org/10.1073/pnas.1417059112

Knight, R. T. (1984). Decreased response to novel stimuli after prefrontal lesions in man. *Electroencephalography and Clinical Neurophysiology*, *59*, 9–20.
Knight, R. T., Grabowecky, M. F., & Scabini, D. (1995). Role of human prefrontal cortex in attention control. In H. H. Jasper, S. Riggio, & P. S. Goldman-Rakic (Eds.), *Epilepsy and the functional anatomy of the frontal lobe* (pp. 21–31). New York: Raven Press.

Knight, R. T. (1996). Contribution of human hippocampal region to novelty detection. *Nature*, *383*, 256–259.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems*, 1097–1105.

Landau, A. N., & Fries, P. (2012) Attention samples stimuli rhythmically. Curr Biol 22:1000–1004.

Leonards, U., Sunaert, S., Van Hecke, P., & Orban, G. A. (2000). Attention mechanisms in visual search - An fMRI study. *Journal of Cognitive Neuroscience*, *12* (SUPPL. 2), 61–75. http://doi.org/10.1162/089892900564073

Li, L., Gratton, C., Fabiani, M., & Knight, R. T. (2013). Age-related frontoparietal changes during the control of bottom-up and top-down attention: An ERP study. *Neurobiology of Aging*, *34*, 477–488.
https://doi.org/10.1016/j.neurobiolaging.2012.02.025

Li, L., Gratton, C., Yao, D., & Knight, R. T. (2010). Role of frontal and parietal cortices in the control of bottom-up and top-down attention in humans. *Brain Research*, *1344*, 173–184. https://doi.org/10.1016/j.brainres.2010.05.016

Li, R., Wu, X., Fleisher, A. S., Reiman, E. M., Chen, K., & Yao, L. (2012). Attention-related networks in Alzheimer's disease: a resting functional MRI study. *Human Brain Mapping*, *33*(5), 1076–1088. https://doi.org/10.1002/hbm.21269

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. Nature, 412, 150–157. https://doi.org/10.1038/35084005

Mackey, W. E., Winawer, J., & Curtis, C. E. (2017). Visual field map clusters in human frontoparietal cortex. *ELife*, *6*, 1–23. https://doi.org/10.7554/elife.22974

Mangano, G. R., Oliveri, M., Turriziani, P., Smirni, D., Zhaoping, L., & Cipolotti, L. (2015). Repetitive transcranial magnetic stimulation over the left parietal cortex facilitates visual search for a letter among its mirror images. *Neuropsychologia*, *70*, 196–205. https://doi.org/10.1016/j.neuropsychologia.2015.03.002

Marazziti, D., Consoli, G., Picchetti, M., Carlini, M., & Faravelli, L. (2010). Cognitive impairment in major depression. *European Journal of Pharmacology*, *626*(1), 83–86. https://doi.org/10.1016/j.ejphar.2009.08.046

Marcus, G. (2018). Deep Learning: A Critical Appraisal, 1–27. arXiv:1801.00631.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–190. https://doi.org/10.1016/j.jneumeth.2007.03.024

Martin, A. B., Yang, X., Saalmann, Y. B., Wang, L., Shestyuk, A., Lin, J. J., Parvizi, J., Knight, R. T., & Kastner, S. (2019). Temporal Dynamics and Response Modulation across the Human Visual System in a Spatial Attention Task : An ECoG Study, *39*(2), 333–352.

Meister, M. L. R., & Buffalo, E. A. (2016). Getting directions from the hippocampus: The neural connection between looking and memory. *Neurobiology of Learning and Memory*, *134*, 135–144. https://doi.org/https://doi.org/10.1016/j.nlm.2015.12.004

Miller, J., Watrous, A. J., Tsitsiklis, M., Lee, S. A., Sheth, S. A., Schevon, C. A., Smith, E. H. Sperling, M. R., Sharan, A., Asadi-Pooya, A. A., Worrell, G. A., Meisenhelter, S., Inman, C. S., Davis, K. A., Lega, B., Wanda, P. A., Das, S. R., Stein, J. M., Gorniak, R., Jacobs, J. (2018). Lateralized hippocampal oscillations underlie distinct aspects of human spatial memory and navigation. *Nature Communications*, *9*(1). https://doi.org/10.1038/s41467-018-04847-9

Minxha, J., Mosher, C., Morrow, J. K., Mamelak, A. N., Adolphs, R., Gothard, K. M., & Rutishauser, U. (2017). Fixations Gate Species-Specific Responses to Free Viewing of Faces in the Human and Macaque Amygdala. *Cell Reports*, *18*(4), 878–891. https://doi.org/10.1016/j.celrep.2016.12.083

Morris, R., Griffiths, O., Le Pelley, M. E., & Weickert, T. W. (2013). Attention to irrelevant cues is related to positive symptoms in schizophrenia. *Schizophrenia Bulletin*, *39*(3), 575–582. https://doi.org/10.1093/schbul/sbr192

Mukamel, R., Gelbard, H., Arieli, A., Hasson, U., Fried, I., & Malach, R. (2005). Coupling between neuronal firing, field potentials, and fMRI in human auditory cortex. *Science*, *309*, 951–954. https://doi.org/10.1126/science.1110913

Nakayama, K., & Martini, P. (2011). Situating visual search. *Vision Research*, *51*(13), 1526–1537. https://doi.org/10.1016/j.visres.2010.09.003

Nir, Y., Fisch, L., Mukamel, R., Gelbard-Sagiv, H., Arieli, A., Fried, I., & Malach, R. (2007). Coupling between Neuronal Firing Rate, Gamma LFP, and BOLD fMRI Is Related to Interneuronal Correlations. *Current Biology*, *17*, 1275–1285. https://doi.org/10.1016/j.cub.2007.06.066

Nobre, A. C., Coull, J. T., Walsh, V., & Frith, C. D. (2002). Brain Activations during Visual Search: Contributions of Search Efficiency versus Feature Binding. *NeuroImage*, *18*(1), 91–103. https://doi.org/https://doi.org/10.1006/nimg.2002.1329

O'Craven, K.M., Downing, P.E., & Kanwisher, N. (1999) fMRI evidence for objects as the units of attentional selection. Nature 401:584 –587.

Olesen, J., Gustavsson, A., Svensson, M., Wittchen, H.-U., & Jönsson, B. (2012). The economic cost of brain disorders in Europe. *European Journal of Neurology : The Official Journal of the European Federation of Neurological Societies*, *19*(1), 155–162. https://doi.org/10.1111/j.1468-1331.2011.03590.x

Olshausen, B. A., & Field, D. J. (2004). What is the other 85% of V1 doing? In T. J. Sejnowski & L. van Hemmen (Eds.), *Problems in Systems Neuroscience* (pp. 182–211). Oxford, United Kingdom: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195148220.003.0010

Olson, I. R., Moore, K. S., Stark, M., & Chatterjee, A. (2006). Visual working memory is impaired when the medial temporal lobe is damaged. *Journal of Cognitive Neuroscience*, *18*(7), 1087–1097. https://doi.org/10.1162/jocn.2006.18.7.1087

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, *2011*. https://doi.org/10.1155/2011/156869

Ossandon, T., Jerbi, K., Vidal, J. R., Bayle, D. J., Henaff, M.-A., Jung, J., Minotti, L., Bertrand, O., Kahane, P., & Lachaux, J.-P. (2011). Transient Suppression of Broadband Gamma Power in the Default-Mode Network Is Correlated with Task Complexity and Subject Performance. *The Journal of Neuroscience*, *31*(41), 14521–14530. https://doi.org/10.1523/JNEUROSCI.2483-11.2011

Ossandon, T., Vidal, J. R., Ciumas, C., Jerbi, K., Hamame, C. M., Dalal, S. S., Bertrand, O., Minotti, L., Kahane, P., & Lachaux, J.-P. (2012). Efficient "Pop-Out" Visual Search Elicits Sustained Broadband Gamma Activity in the Dorsal Attention Network. *Journal of Neuroscience*, *32*(10), 3414–3421. https://doi.org/10.1523/JNEUROSCI.6048-11.2012

Parker, A., Wilding, E., & Akerman, C. (1998). The von Restorff effect in visual object recognition memory in humans and monkeys: The role of frontal/perirhinal interaction. *Journal of Cognitive Neuroscience*, *10*(6), 691–703. https://doi.org/10.1162/089892998563103

Parvizi, J., & Kastner, S. (2018). Promises and limitations of human intracranial electroencephalography. *Nature Neuroscience*, *21*(4), 474–483. https://doi.org/10.1038/s41593-018-0108-2

Pinto, N., Barhomi, Y., Cox, D. D., & DiCarlo, J. J. (2011). Comparing state-of-the-art visual features on invariant object recognition tasks. In *IEEE Workshop on Applications of Computer Vision (WACV)* (pp. 463–470). https://doi.org/10.1109/WACV.2011.5711540

Posner, M. I. (1980) Orienting of attention. Q J Exp Psychol 32:3–25.

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *PNAS*, *98*(2), 676–682.

Raichle, M. E. (2015). The Brain's Default Mode Network. *Annual Review of Neuroscience*, *38*(1), 433–447. https://doi.org/10.1146/annurev-neuro-071013-014030

Ramot, M., Fisch, L., Harel, M., Kipervasser, S., Andelman, F., Neufeld, M. Y., Kramer, U., Fried, I., Malach, R. (2012). A Widely Distributed Spectral Signature of Task-Negative Electrocorticography Responses Revealed during a Visuomotor Task in the Human Cortex. *Journal of Neuroscience*, *32*(31), 10458–10469. https://doi.org/10.1523/jneurosci.0877-12.2012

Rao, R. P., & Ballard, D. H. (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci 2:79–87.

Ray, S., Crone, N. E., Niebur, E., Franaszczuk, P. J., & Hsiao, S. S. (2008). Neural correlates of high-gamma oscillations (60-200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *The Journal of Neuroscience*, *28*(45), 11526–11536. https://doi.org/10.1523/JNEUROSCI.2848-08.2008

Ray, S., & Maunsell, J. H. R. (2011). Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biology*, *9*(4), 1–15. https://doi.org/10.1371/journal.pbio.1000610

Regev, T. I., Winawer, J., Gerber, E. M., Knight, R. T., & Deouell, L. Y. (2018). *Human posterior parietal cortex responds to visual stimuli as early as peristriate occipital cortex*. *European Journal of Neuroscience* (Vol. 48). https://doi.org/10.1111/ejn.14164

von Restorff, H. (1933) Psychol. Forsch. 18: 299. https://doi.org/10.1007/BF02409636

Robinson, D. L., & Petersen, S. E. (1992). The pulvinar and visual salience. *Trends in Neurosciences*, *15*(4), 127–132.

Rosenberg, M. D., Finn, E. S., Scheinost, D., Papademetris, X., Shen, X., Constable, R. T., & Chun M. M. (2016) A neuromarker of sustained attention from whole-brain functional connectivity. Nat Neurosci 19:165–171.

Rossi, A. F., Bichot, N. P., Desimone, R., & Ungerleider, L. G. (2007). Top Down Attentional Deficits in Macaques with Lesions of Lateral Prefrontal Cortex. *Journal of Neuroscience*, *27*(42), 11306–11314. https://doi.org/10.1523/JNEUROSCI.2939-07.2007

Rutishauser, U., Mamelak, A. N., & Schuman, E. M. (2006). Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron*, *49*(6), 805–813. https://doi.org/10.1016/j.neuron.2006.02.015

Schiller, D., Eichenbaum, H., Buffalo, E. A., Davachi, L., Foster, D. J., Leutgeb, S., & Ranganath, C. (2015). Memory and Space: Towards an Understanding of the Cognitive Map. *Journal of Neuroscience*, *35*(41), 13904–13911. https://doi.org/10.1523/jneurosci.2618-15.2015

Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., & Petersen, S. E. (1997). Common blood flow changes across visual tasks: II. Decreases in cerebral cortex. *Journal of Cognitive Neuroscience*, *9*(5), 648–663. https://doi.org/10.1162/jocn.1997.9.5.648

Silver, M. A., Ress, D., & Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *Journal of Neurophysiology*, *94*(2), 1358–1371. https://doi.org/10.1152/jn.01316.2004

Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences*, *13*(11), 488–495. https://doi.org/10.1016/j.tics.2009.08.005.

Simon, J. R., & Wolf, J. D. (1963) Choice reaction time as a function of angular stimulus-response correspondence and age. Ergonomics 6:99–105.

Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR* (pp. 1–14). Retrieved from http://arxiv.org/abs/1409.1556

Stolk, A., Griffin, S., Van Der Meij, R., Dewar, C., Saez, I., Lin, J. J., Piantoni, G., Schoffelen, J. M., Knight, R. T., Oostenveld, R. (2018). Integrated analysis of

anatomical and electrophysiological human intracranial data. *Nature Protocols*, *13*(7), 1699–1723. https://doi.org/10.1038/s41596-018-0009-6

Stroop, J. R. (1935) Studies of interference in serial verbal reactions. J Exp Psychol 18:643–662.

Summerfield C., & Egner, T. (2009) Expectation (and attention) in visual cognition. Trends Cogn Sci 13:403–409.

Summerfield C., & Egner, T. (2016) Feature-based attention and feature-based expectation. Trends Cogn Sci 20:401–404.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A. (2015). Going Deeper with Convolutions. In *The IEEE Conference on Computer Vision and Pattern Recognition* (pp. 1–9). https://doi.org/10.1002/jctb.4820

Thompson, K. G., & Bichot, N. P. (2005). A visual salience map in the primate frontal eye field. *Progress in Brain Research*. https://doi.org/10.1016/S0079-6123(04)47019-8

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97–136. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7351125

Wang, S., Mamelak, A. N., Adolphs, R., & Rutishauser, U. (2018). Encoding of Target Detection during Visual Search by Single Neurons in the Human Brain. *Current Biology*, *28*(13), 2058–2069. https://doi.org/https://doi.org/10.1016/j.cub.2018.04.092

Wardak, C., Vanduffel, W., & Orban, G. A. (2010). Searching for a salient target involves frontal regions. *Cerebral Cortex*, *20*(10), 2464–2477. https://doi.org/10.1093/cercor/bhp315

White, B. J., Kan, J. Y., Levy, R., Itti, L., & Munoz, D. P. (2017). Superior colliculus encodes visual saliency before the primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(35), 9451–9456. https://doi.org/10.1073/pnas.1701003114

Williams, A. H., Poole, B., Maheswaranathan, N., Dhawale, A. K., Fisher, T., Wilson, C., D., Brann, D. H., Trautmann, E., Ryu, S., Shusterman, R., Rinberg, D., Ölveczky, B. P., Shenoy, K. V. & Ganguli, S. (2019). Discovering precise temporal patterns in large-scale neural recordings through robust and interpretable time warping. *BioRxiv Preprint*. https://doi.org/http://dx.doi.org/10.1101/661165

Wolfe, J. (2014). Approaches to Visual Search: Feature Integration Theory and Guided Search. In S. Kastner & A. C. Nobre (Eds.), *The Oxford Handbook of Attention* (1st ed., pp. 11–55). Oxford, United Kingdom: Oxford University Press.

Wolfe, J. M. (2018). Visual Search. In J. T. Wixted (Ed.), *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* (4th ed., pp. 1–55). John Wiley & Sons, Inc. http://doi.org/10.1002/9781119170174.epcn213

Xu, Y. (2018). The Posterior Parietal Cortex in Adaptive Visual Processing. *Trends in Neurosciences*, *41*(11), 806–822. https://doi.org/10.1016/j.tins.2018.07.012

Yan, Y., Wei, R., Zhang, Q., Jin, Z., & Li, L. (2016). Differential roles of the dorsal prefrontal and posterior parietal cortices in visual search: A TMS study. *Scientific Reports*, *6*(January), 1–9. https://doi.org/10.1038/srep30300

Yan, Y., Zhaoping, L., & Lia, W. (2018). Bottom-up saliency and top-down learning in the primary visual cortex of monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(41), 10499–10504. https://doi.org/10.1073/pnas.1803854115

Yarnall, A., Rochester, L., & Burn, D. J. (2011). The interplay of cholinergic function, attention, and falls in Parkinson's disease. *Movement Disorders : Official Journal of the Movement Disorder Society*, *26*(14), 2496–2503. https://doi.org/10.1002/mds.23932

Yoon, J. H., Sheremata, S. L., Rokem, A., & Silver, M. A. (2013). Windows to the soul: vision science as a tool for studying biological mechanisms of information processing deficits in schizophrenia. *Frontiers in Psychology*, *4*(681), 1–15. https://doi.org/10.3389/fpsyg.2013.00681

Zhaoping, L. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences*, *6*(1), 9–16.

Zhaoping, L., & Dayan, P. (2006). Pre-attentive visual selection. *Neural Networks*, *19*(9), 1437–1439. https://doi.org/10.1016/j.neunet.2006.09.003

Zhaoping, L. (2019). A new framework for understanding vision from the perspective of the primary visual cortex. *Current Opinion in Neurobiology*, *58*(Box 1), 1–10. https://doi.org/10.1016/j.conb.2019.06.001

Appendices

Appendix A: Tables

Table 2.1. Participant information.

SID	Sex	Age	Ha.	SEEG/ ECoG	He.	ROIs	Number included (implanted) contacts	Number active contacts	Number included (compl.) trials Pop-Out	Number included. (compl.) trials Search
S1	F	28	R	ECoG	R+L	T, MTL	22 (22)	5	57 (64)	53 (64)
S2	F	63	R	SEEG	R+L	T, F, MTL, cing., amy.	51 (128)	19	48 (64)	43 (64)
S3	M	27	R	SEEG	R+L	F, P, T, MTL, amy.	40 (101)	7	51 (64)	50 (64)
S4	F	34	R	SEEG	R+L	F, T, MTL, cing., amy.	61 (126)	18	48 (67)	38 (61)
S5	M	35	R	SEEG	R	F, cing.	36 (79)	17	58 (64)	42 (64)
S6	M	52	R	SEEG	R	F, P, cing.	75 (128)	33	40 (50)	38 (46)
S7	F	23	R	ECoG	R+L	F, T	54 (72)	17	59 (64)	56 (64)
S8	F	36	R	ECoG	R+L	F, P, cing.	48 (48)	20	56 (64)	52 (64)
S9	M	32	R	ECoG	R+L	F, T, MTL, cing.	53 (64)	12	49 (64)	33 (64)
S10	F	69	R	ECoG	R+L	F, P, T, MTL	35 (39)	11	59 (64)	29 (64)
S11	M	25	A	ECoG	R+L	F, P, T, MTL, cing.	59 (80)	11	53 (64)	42 (64)
S12	F	21	R	ECoG	R+L	F, P, cing.	107 (120)	47	57 (64)	49 (64)

S13	F	58	R	SEEG	R+L	F, T, MTL, cing., amy.	53 (122)	35	57 (64)	50 (64)
S14	M	34	A	SEEG	R+L	F, T, MTL, cing., amy.	37 (113)	5	62 (64)	60 (64)
S15	F	22	R	SEEG	R+L	F, T, cing.	65 (148)	30	42 (64)	30 (64)
S16	M	41	R	SEEG	R+L	F, T, MTL, cing., amy.	49 (142)	11	51 (64)	38 (64)
S17	F	32	R	SEEG	R+L	F, T, MTL, ACC, amy.	35 (91)	8	54 (64)	31 (64)
S18	F	26	R	ECoG	L	P, T, O, MTL	125 (156)	33	58 (64)	45 (64)
S19	M	23	R	SEEG	R+L	F, T, cing.	78 (160)	37	61 (64)	60 (64)
S20	F	50	R	SEEG	R+L	F, T, MTL, cing.	31 (106)	9	34 (64)	21 (64)
S21	M	27	R	SEEG	R+L	F, T, cing., amy.	68 (120)	15	48 (64)	38 (64)
S22	M	23	R	SEEG	R	F, T, O, MTL, cing., amy.	42 (90)	13	44 (64)	36 (64)
S23	F	25	R	SEEG	R+L	F, T, MTL, cing., amy.	97 (180)	42	56 (64)	43 (64)

Table header: SID, subject ID; Ha., handedness; He., hemispheric coverage; ROIs, coverage by ROI; Number included (compl.) trials, number included (completed) trials.
Sex: F, female; M, male. Handedness: R, right; L, left; A, ambidextrous. Hemispheric coverage: R, right; L, left; R + L, right and left hemispheres. Coverage by ROI: F, frontal; P, parietal; T, temporal; O, occipital; MTL, medial temporal lobe; cing., cingulate; amy., amygdala.

Table 2.2. Distribution of task-selective effects, by large cortical ROIs.

Anatomical ROI	**Percent active electrodes (active/included)**		
	left	right	total
Frontal cortex	**44%** (75/171)	**31%** (88/281)	**36%** (163/452)
Parietal cortex	**44%** (12/27)	**43%** (23/53)	**44%** (35/80)
Temporal cortex	**20%** (46/227)	**23%** (33/142)	**21%** (79/369)
Occipital cortex	**73%** (8/11)	**50%** (3/6)	**65%** (11/17)
Cingulate cortex	**22%** (16/72)	**28%** (21/74)	**25%** (37/146)
Medial temporal lobe	**48%** (15/31)	**33%** (9/27)	**41%** (24/58)
Amygdala	**27%** (4/15)	**43%** (10/23)	**37%** (14/38)
Sensorimotor cortex	**66%** (57/87)	**45%** (33/74)	**56%** (90/161)
Total	**36%** (233/641)	**32%** (220/680)	**34%** (453/1321)
Total excluding sensorimotor cortex	**32%** (176/554)	**31%** (187/606)	**31%** (363/1160)

Table 2.3. Distribution of task-selective effects within frontal (non-motor) cortex.

Frontal ROI	**Percent active electrodes (active/included)**		
	left	right	total
Superior frontal gyrus and sulcus	**29%** (12/42)	**32%** (28/87)	**31%** (40/129)
Middle frontal gyrus	**55%** (32/58)	**33%** (26/78)	**43%** (58/136)
Inferior frontal gyrus and sulcus	**68%** (17/25)	**46%** (22/48)	**53%** (39/73)
Orbitofrontal cortex	**19%** (6/32)	**19%** (8/43)	**19%** (14/75)
Medial prefrontal cortex	**57%** (8/14)	**16%** (4/25)	**31%** (12/39)
Total	**44%** (75/171)	**31%** (88/281)	**36%** (163/452)

Table 2.4. Distribution of task-selective effects within parietal (non-motor) cortex.

Parietal ROI	**Percent active electrodes (active/included)**		
	left	right	total
Superior parietal lobule	**0%** (0/0)	**65%** (11/17)	**65%** (11/17)
Inferior parietal lobule	**38%** (9/24)	**29%** (9/31)	**33%** (18/55)
Precuneus	**100%** (3/3)	**60%** (3/5)	**75%** (6/8)
Total	**44%** (12/27)	**43%** (23/53)	**44%** (35/80)

Table 2.5. Distribution of task-selective effects within temporal cortex.

Temporal region of interest	**Percent active electrodes (active/included)**		
	left	right	total
Insula	**15%** (3/20)	**35%** (9/26)	**26%** (12/46)
Superior temporal gyrus	**9%** (3/35)	**11%** (1/9)	**9%** (4/44)
Superior temporal sulcus	**9%** (4/44)	**13%** (5/38)	**11%** (9/82)
Middle temporal gyrus	**20%** (10/50)	**21%** (6/29)	**20%** (16/79)
Inferior temporal gyrus	**23%** (8/35)	**21%** (5/24)	**22%** (13/59)
Ventral stream	**42%** (13/31)	**38%** (5/13)	**41%** (18/44)
Temporal pole	**42%** (5/12)	**67%** (2/3)	**47%** (7/15)
Total	**20%** (46/227)	**23%** (33/142)	**21%** (79/369)

Table 2.6. Distribution of task-selective effects within cingulate cortex.

Cingulate ROI	**Percent active electrodes (active/included)**		
	left	right	total
Anterior cingulate cortex	**18%** (6/33)	**12%** (4/34)	**15%** (10/67)
Midcingulate cortex	**26%** (7/27)	**44%** (11/25)	**35%** (18/52)
Posterior cingulate cortex	**25%** (3/12)	**40%** (6/15)	**33%** (9/27)
Total	**22%** (16/72)	**28%** (21/74)	**25%** (37/146)

Table 2.7. Distribution of task-selective effects within the medial temporal lobe.

MTL ROI	**Percent active electrodes (active/included)**		
	left	right	total
Hippocampus	**38%** (5/13)	**44%** (4/9)	**41%** (9/22)
Parahippocampal cortex	**50%** (6/12)	**10%** (1/10)	**32%** (7/22)
Entorhinal cortex	**67%** (4/6)	**50%** (4/8)	**57%** (8/14)
Total	**48%** (15/31)	**33%** (9/27)	**41%** (24/58)