Dynamic neural coding underlying feature-based learning across fronto-striatal circuits

By

Benjamin Voloh

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Psychology

January 31, 2020

Nashville, Tennessee

Approved:

Thilo Womelsdorf, PhD

Andrew Tomarken, PhD

Jeffrey Schall, PhD

Erin Calipari, PhD

Acknowledgments

If graduate school teaches anything, it's that you can't succeed alone. First and foremost, I would like to thank my supervisor, Dr. Thilo Womelsdorf, who is a seemingly bottomless well of energy and positivity. I have been given more opportunities in his lab than I could have hoped for. Thank you for your relentless support and sharing your invaluable knowledge.

I would also like to thank my committee member, Drs. Jeffrey Schall, Andrew Tomarken, and Erin Calipari, for not only pushing me to critically explore my own work and broaden my field of view, but also for warmly welcoming me into the Vanderbilt community.

The daily grind would not have been possible without my lab mates. Thank you all for the inspiring scientific discussions and managing to keep it fun. I would especially like to thank Ali Hassani, who didn't run away when the gears came off; Dr. Marcus Watson, who I spent many years in the trenches with; and Dr. Mariann Oemisch, who, always being one step ahead, was a constant source of inspiration and consultation. I would also like to thank everyone who collected the data I analyzed; this work would obviously not have happened without you.

To my partner Ioanna, factually, we should have never met, yet the stars aligned. Thank you for constantly surprising me, for being by my side through the thick of it, for motivating me to be a better person, and for your relentless honesty. I will never be able to thank you enough, but I'll keep trying.

To my oldest friends, Robbie, Zack, Chet, Ivan, and Ben2. Y'all are alright. Let's hang out some more.

And finally, to my family, who have provided me with a constant source of love and support. To my parents, every opportunity I've had is only possible because of your sacrifices, and to my sister, you inspire me with just how much a single person can juggle. Thank you. I would not have made it without you.

Table of Contents

List of Figures

List of Abbreviations

Abbreviation	*Comment*
LPFC	Lateral prefrontal cortex
ACC	Anterior cingulate cortex
STR	Striatum
SC	Superior colliculus
vmPFC	Ventro-medial prefrontal cortex
OFC	Orbito-frontal cortex
LFP	Local Field Potential

Chapter 1

Introduction

Gregory Bateson said that "a bit of information is a difference that makes a difference" (Bateson, 1972). In the brain, changes in neuronal firing patterns are sensed by neighbors in order to, ultimately, align the organism onto its goals. Yet neuronal dynamics are complex and multi-faceted; it is, as yet, unclear which specific variations in neuronal activity those self-same neighbors are sensitive to. To understand what may be conveyed by these fluctuating dynamics, it is thus instructive to consider the *purpose* of neurally-mediated information. One approach is to define a problem with a well-specified algorithmic solution, and then ask *what information* would help in the resolution of this problem (Marr, 1982). In so doing, we can determine specifically what differences in neural activation make a difference in goal-oriented behavior (Schall, 2004).

A particularly salient problem for any higher organism is to learn which of the cornucopia of features in the world are relevant to their current goals, and to flexibility adapt to changing reward contingencies. Via trial and error, they can learn what features are highly predictive of harm or benefit. To achieve this, they must maintain a memory for the values of those features, sense the outcomes arising from choices made on their bases, and recognize if those outcomes were not as expected. These three variables – value, outcome, deviations – are the comprising units of a particularly successful class of algorithms that can learn predictive value, namely, Reinforcement Learning (RL) algorithms.

In reinforcement learning, an agent repeatedly performs actions on the environment on the basis of its current best knowledge of what features can maximize reward. Observation of the ensuing outcome, and comparison with the original prediction, leads to the development of *prediction errors*. The prediction error is then used to update the feature values. After multiple iterations, the expected value will converge to its true value, the prediction will match the

outcome, the prediction error will be zero, and value-updating will have ceased. The core principles of RL algorithms can be captured by the original Rescorla-Wagner equations:

$$V(x, t+1) = V(x, t) + \alpha\delta$$
$$\delta = R - V(x, t)$$

Here, V(x,t) is the value function of option X at time t. The learning term $\alpha\delta$ is composed of a learning rate α, and a prediction error term δ. The learning rate determines how quickly an agent adapts its value function. The prediction error term δ is the difference between the expected reward V(x,t) and the observed reward R. The prediction error term can thus trigger learning; positive prediction errors result in increased future value assigned to option X, whereas a negative prediction error results in a lower value.

RL shows how learning may occur for different options, but it is agnostic to the identity of the choices, such as whole objects or their comprising features. Yet in both human and non-human primates, behavior is best explained by models that weight feature dimensions, not individual objects (Niv et al., 2015; Balcarras et al., 2016; Leong et al., 2017; Oemisch et al., 2019). Such weighting could be achieved by attentional mechanisms that highlight specific feature dimensions (Leong et al., 2017), or recently active representations that had led to specific choices (Roelfsema and van Ooyen, 2005). Feature-based learning may be particularly relevant when subjects have limited experience with stimuli, and may underlie the subsequent formation of holistic object representations as their expertise improves (Gauthier and Tarr, 1997; Bukach et al., 2006; Blair et al., 2009; Izquierdo et al., 2017).

Much of the evidence for how neurons encode the value of chosen stimuli, long-term memory for value, and prediction errors is derived from studies of neuronal firing rates. Yet the activity of individual neurons is highly variable, and is often correlated with band-limited, rhythmic network activity (Womelsdorf et al., 2014b; Womelsdorf and Everling, 2015; Voloh and Womelsdorf, 2016; Hahn et al., 2019). The goal of this thesis is two-fold; first, to delineate how neuronal spiking activity in the context of such oscillations may support the encoding of learning-

relevant information; and second, to lay the foundation for future work into neural coding in the context of intrinsic neuronal variability and complex naturalistic tasks.

Having defined the problem's scope and *what* information is necessary, I will first detail *where* such information is found, before detailing *how* it may be encoded.

1.1. Where: Brain-networks supporting feature-based learning

Signals related to value, reward history, prediction errors and outcome have been observed in a wide variety of areas, including frontal regions such as the orbitofrontal cortex (OFC), ventromedial prefrontal cortex (vmPFC), lateral prefrontal cortex (LPFC), anterior cingulate cortex (ACC), and parietal and primary sensory cortices, as well subcortical structures, most prominently the basal ganglia and its subdivision, the striatum (STR). In this work, we focus on the ACC, LPFC, and STR. These reason for this selection is justified below.

1.1.1 LPFC: Implementing goal-oriented attentional allocation

Functionally, the lateral prefrontal cortex comprises area 46 and part of 9, as well as area 8/8a (which also includes the frontal eye fields (FEF))(Passingham and Wise, 2012). Lesions to area 46 result in impairments in a delayed response task, where a target is primed, and after some delay – perhaps involving an occlusion of the stimuli – subjects must choose. This likely arises as a result of a failure to encode the (future) goal (a form of prospective memory) (Passingham and Wise, 2012). Lesions in this area also impair rule acquisition (Buckley et al., 2009). On the other hand, more posterior regions of the LPFC (including area 8, FEF) are more involved in orienting eye movements in light of current goals (presumably encoded by area 46). For example, in FEF, even when information about motor plans and visual information is multiplexed by single neurons, the majority of cells still encode the location of the target stimulus (Sato and Schall, 2003). Moreover, inactivation of area 8 results in saccades that are highly variable but still travel in the direction of the target (Passingham and Wise, 2012), indicating that goals are preserved and thus likely inherited from other areas.

1.1.2 ACC: Identification and adjustment when rules are maladaptive

ACC is heavily implicated in updating model of the environment, particularly when task contingencies change (Kolling et al., 2016; Shenhav et al., 2016). Macaques with lesions of the ACC are slower to switch new rules, driven mostly by a failure to persist in rewarded behavior (Kennerley et al., 2006). Such neurons can multiplex information about actions and values (Hayden and Platt, 2010). Moreover, neurons in ACC can encode reward history (Bernacchia et al., 2011), stimulus values (Kaping et al., 2011), prediction errors (Ito, 2003; Oemisch et al., 2019), and outcome surprise (or, unsigned prediction error) (Hayden et al., 2011; Oemisch et al., 2019), in line with evidence suggesting ACC integrates such information in order to signal when current strategies are maladaptive (Johnston et al., 2007), or to switch to counter-factual options (Fouragnan et al., 2019). In rodents, pharmacological inactivation of the ACC during a Morris Water Maze task results in impaired search strategies when trials are in conflict with previously learned patterns (Richards et al., 2014). These results thus indicate that the ACC has a prominent role in building models of the environment in order to evaluate outcomes and flexibility switch strategies when current ones have become maladaptive.

1.1.3 Midbrain: prediction error signaling via dopaminergic cells

Midbrain dopamine neurons have classically been thought to encode prediction error signals (Lee et al., 2012), following pathbreaking work by (Schultz et al., 1997). Dopaminergic cells in the ventral tegmental area (VTA) respond vigorously when an outcome is unexpected during Pavlovian conditioning. Importantly, after conditioning has occurred, the dopaminergic response shifts from the outcome period to the onset of the conditioned stimulus, suggesting dopaminergic cells may implement a form of Temporal-Difference learning (Schultz et al., 1997; Lee et al., 2012). These cells scale proportionally with the degree of prediction error (Fiorillo et al., 2003; Bayer and Glimcher, 2005).

Although this is a powerful account of the role of dopamine in learning, there are some complications (Redgrave and Gurney, 2006; Berridge, 2007). Dopaminergic cells can respond differentially to positive or negative reward (Matsumoto and Hikosaka, 2009; Thibeault et al.,

2019), perhaps because of differences in the source of information about negative/positive value (e.g. (Matsumoto and Hikosaka, 2007)). Indeed, a substantial number of cells show responses more consistent with uncertainty (unsigned prediction error), rather than prediction errors per se (Matsumoto and Hikosaka, 2009). This is in line with work showing that dopaminergic signaling may also represent the resolution of uncertainty (Bromberg-Martin and Hikosaka, 2009; White et al., 2019). Although dopaminergic responses show heterogeneity and their functional role is still under debate, it is nevertheless clear that they play a key role in learned behavior.

What is the source of visual information by which dopaminergic cells may evaluate expected visual stimuli? One source is the superior colliculus (SC; responsible for eye movement control)) (Comoli et al., 2003; Dommett et al., 2005; Takakuwa et al., 2017), which, considering its sensitivity to changes in luminance, is in an ideal position to signal the sudden appearance or disappearance of visual stimuli (Redgrave and Gurney, 2006). In the absence of a visual cortex, pharmacological inactivation of the SC abolishes visually-evoked response in dopaminergic cells (Comoli et al., 2003). Moreover, disinhibition of the SC results in visual responses in the dopaminergic cells (Dommett et al., 2005). Likewise, V1 may communicate more complex visual features (beyond just luminance changes), as evidenced by a study showing that similar visually-evoked responses in dopaminergic cells are apparent if SC is inactivated but V1 is left unperturbed (Takakuwa et al., 2018).

1.1.4 Striatum: action selection based on learned value

The STR is prominently involved in the encoding and learning of object values, as well as the generation of purposeful movements that depend on the value of these objects (Hikosaka et al., 2017). The STR is characterized by a few major subdivisions. The ventral striatum includes the shell and core (distinguished by the presence/ absence of calbindin receptors) (Haber 2016; 2006), as well as the dorsal sections of the caudate and putamen. The dorsal striatum includes the caudate and putamen. The caudate can be further subdivided into the head, body, and tail (Haber, 2016; Hikosaka et al., 2017). These subdivisions have different functional properties as it relates to learning. When stimulus-reward contingencies shift rapidly, neurons in the caudate

head, unlike the caudate tail, are highly responsive and able to differentiate objects. The inverse is true when object values are stable over time (Kim and Hikosaka, 2013). This suggests an anterior-posterior gradient in the caudate for short vs long term value encoding (Hikosaka et al., 2017). In addition to learning object values, the STR is also responsive to more abstract rule values (Ballard et al., 2018), and is likely involved in transfer learning (Dahlin et al., 2008).

The striatum has a distinct role in controlling saccadic eye movements (Hikosaka et al., 2017). The caudate contains diverse neuronal populations that are excited before both visually guided and memory guided saccade onset (Hikosaka et al., 1989). Importantly, these occur specifically in task-related contexts, and not during spontaneous saccades, suggesting the caudate is involved during goal-oriented visual processing. The caudate can further influence the SC via the substantia nigra pars compacta (Hikosaka et al., 2017). This is likely instantiated through a disinhibitory mechanism, thus allowing other sources to influence saccadic eye movements (Hikosaka et al., 1989)

1.1.5 Connectivity

Frontal areas, including area 24, 46, and 8, have strong and reciprocal connections with each other (Pandya et al., 1981; Vogt et al., 1987; Passingham and Wise, 2012; Yeterian et al., 2012; Barbas, 2015). Unlike other prefrontal areas, Area 24 also receive strong inputs from the hippocampus, and projects to the hypothalamus, suggesting a role in autonomic control as well (Passingham and Wise, 2012). ACC projections to the LPFC originate in deep layers and terminate in superficial layers, whereas the inverse is true of projections from the LPFC to the ACC (Barbas and Pandya, 1989; Barbas, 2015).

The cortex and striatum have well-defined and topographically preserved anatomical connections. At a rough glance, outputs of more rostral areas (including ACC) have terminals in the ventral striatum. More caudal areas, such as area 46, 8, and 9, have terminals that are more dorsal in the striatum, and motor areas have terminals that are more dorsal still (Haber, 2016; Hikosaka et al., 2017). Likewise, outputs from these regions of the striatum tend to have

reciprocal connections (via the thalamus) to the same cortical areas. Thus, at a rough approximation, the ventral/dorsal striatal axis maintains a gradient of limbic/cognitive/motor processing, and segregates these in cortico-striatal loops that can operate in parallel (Hikosaka et al., 2017).

A closer inspection, however, suggests that in addition to the high degree of segregation of cortical terminals in the striatum, there also exist hot spots where terminals can overlap (Haber 2006). Premotor cortex, SEF, 46, 9, 8a/FEF have large (focal) terminal fields that overlap within the dorsal caudate. On the other hand, terminals from 24c have a relatively large presence in the putamen (with smaller fields emerging in the caudate). Moreover, area 24b has terminal fields that extend from the ventral striatum and also up into the (ventral) caudate. It is important to note that area 24c, unlike 24b, tends to innervate premotor areas (premotor cortex and SEF), whereas 24b innervates more limbic areas (24c/vmPFC/OFC) (Vogt and Pandya, 1987; Calzavara et al., 2007; Haber, 2016). This is further evidence that the limbic/cognitive/motor divide evident at the level of the cortex is also preserved at the level of striatum.

The overlap in cortical terminals suggests that the striatum may subserve the integration of information across cortical areas. Such integration is further supported by striato-nigral loops (Haber et al., 2000; Haber, 2016). Striato-nigral loops have a conserved “motif”, whereby more ventral subdivisions of the substantia nigra receive only inputs, more dorsal subdivisions sending only outputs, and intermediate subdivisions having both. Importantly, the edges of this motif also have cells that project to or receive input from a topographically adjacent striatal area. To roughly visualize this, it means that information from the striatum arrives at ventral portions of the substantia nigra, travels dorsally, and is then output to the “next” striatal subdivisions. This suggests that striatal-nigral loops may be an anatomical substrate by which internal states provide contextual guidance to cognitive systems that ultimately influence action systems.

1.1.6 Summary

The ACC, LPFC, and striatum are heavily interconnected, yet remain functionally distinct. The anterior cingulate cortex helps in building models of the environment in light of current goals. The implementation of goals, on the other hand, is more the purview of the LPFC. The striatum encodes option values as they relate to specific actions. These computations are likely informed by goals inherited from the prefrontal cortex. In order to flexibly learn and act upon feature-reward associations, information contained in the individual modules of this highly recurrent system must be represented in such a way that it may be effectively and selectively transmitted. In the following, I will review evidence of how such information emerges in neural circuits.

1.2. How: Neural Coding for learning-relevant signals

1.2.1 Coding by individual neurons

Integration of learning-relevant information across the ACC, LPFC, and striatum critically depends on how that information is encoded, which specifies how a downstream reader might decode it. One of the main goals of neuroscience research is to specify the mapping between stimuli and neuronal responses. This is referred to as the encoding problem. By observing a particular neuronal response, it therefore becomes possible to reconstruct the stimulus. Alternatively, by observing a particular stimulus, it is possible to predict the neuronal response.

It is widely accepted that one of the main carriers of information in the brain is the spikes of individual neurons. In particular, information can be carried by the amount of spikes elicited by single neurons, referred to as a *rate code*. By observing how the rate changes as the stimulus changes (in a parametric way), it is possible to reconstruct neural "tuning curves". This has been a particularly effective strategy in delineating the responses to simple (Hubel and Wiesel, 1959; Wurtz, 2009), and higher order visual features such as motion (Britten et al., 1992, 1996), as well as value based tuning (Schultz et al., 1997; Fiorillo et al., 2003; Kennerley et al., 2009, 2011; Hayden and Platt, 2010; Kaping et al., 2011; Wallis and Kennerley, 2011; Oemisch et al., 2019; Sajad et al., 2019). Rate fluctuations can track changing stimulus properties to within 10 ms (Bair and Koch, 1996). One critical question, however, is the time period over which spikes are

integrated by a downstream reader. This can differ by cortical region (Heller et al., 1995; Runyan et al., 2017); for example, information is maximized in relatively short windows of 25 ms in area V1, and longer 50 ms windows in area IT (Heller et al., 1995).

An alternative hypothesis, following path-breaking work by (Segundo et al., 1963), is that information is transmitted by the precise timing or sequence of neuronal spikes, referred to as a *temporal code* (Thorpe, 1990; Richmond, 2009). Thus, the time-course of spike rate in two conditions may contain differences, even when the overall spike rate is equal between conditions (Gawne et al., 1996). For example, in the striatum, discounted value is encoded in the temporal domain (Falcone et al., 2019). Rate and temporal coding likely represents two extremes of a continuum of coding that depends on the temporal resolution of the downstream reader (Ratté et al., 2013). Shorter resolutions would favor a temporal code, whereas longer ones would favor a rate code (Panzeri et al., 2010; Ratté et al., 2013). Note that short resolutions do not directly imply temporal coding, because they may reflect a rate code operating at very fast time-scales (for example, when stimulus features vary quickly) (Borst and Theunissen, 1999).

Beyond temporal variability of spiking activity, intra-neuronal dynamics themselves are highly variable, and may permit multiplexing information, particularly if the complex morphology of neurons is taken into account. Compartment based models of neuronal populations have shown that combined activation of the cell soma and dendrites can result in different neuronal dynamics, such as Poisson firing or burst firing, that can concurrently carry separable streams of information, such as stimulus prediction or prediction errors (Körding and König, 2000; Larkum, 2013; Issa et al., 2016; Guerguiev et al., 2017; Naud and Sprekeler, 2018). The differential dynamics can depend on whether feedback and feedforward impinging inputs are coincident in time (Larkum, 2013).

1.2.2 From neurons to networks

The information contained by multiple neurons must be integrated across them. This offers a clue as to another potential code, namely, a *population code* (Pouget et al., 2000;

Averbeck et al., 2006). The simplest example of such a code is essentially the vector sum of the population of neurons, as is evident in encoding of reaching direction (Georgopoulos et al., 1986). It is important to note that the dimension that neurons are tuned to need not follow our intuitions of the task space, as is evident in cells whose activity signals the magnitude along individual dimensions of a 50-dimensional "face-space" (Chang and Tsao, 2017). However, neurons often show "mixed selectivity", in that they will respond to multiple task-relevant variables to varying degrees (Mante et al., 2013; Rigotti et al., 2013; Fusi et al., 2016) (but see (Hirokawa et al., 2019)). This allows for a robust, linear readout from the population as a whole (Fusi et al., 2016). Moreover, the specific sequence of activation of individual neurons can also convey information, such as sequence of place cell firing in the hippocampus conveying the direction and speed of locomotion (Huxter et al., 2008; Pastalkova et al., 2008). Population codes are attractive because variability between neurons can be "averaged out" (Pouget et al., 2000; Cohen and Maunsell, 2010). In the same vein, the loss of some finite subpopulation of neurons will not be detrimental because the population as a whole carries redundant information (Pouget et al., 2000; Pitkow and Angelaki, 2017). Indeed, up to 50% of the variability in a neuron's response can be predicted from its coupling with other neurons (Montijn et al., 2016), and the specific pattern of coupling can carry more information than simple averaging of the population as a whole (Stevenson and Kording, 2011; Runyan et al., 2017).

Population codes may also be more relevant when the task dimensionality is high (Gao and Ganguli, 2015). Indeed, a surprisingly low-number of neurons are necessary to explain a high degree of task variance, perhaps because simpler tasks do not tax neuronal resources enough to require higher dimensional representation (Gao and Ganguli, 2015; Pitkow and Angelaki, 2017). Indeed, coding patterns change when tasks or stimuli are naturalistic, especially when compared to more traditional stimuli (David et al., 2004; Felsen and Dan, 2005; Montijn et al., 2016). This problem is particularly salient for feature-based learning in naturalistic environments, where feature-dimensionality is high and thus must recruit additional mechanisms to constrain neuronal representation to task-relevant dimensions (Leong et al., 2017). Care must be taken in the construction and parameterization of high-dimensional stimulus probes, because truly

naturalistic stimuli are often too unconstrained for strong scientific inference (Rust and Movshon, 2005).

Propagation of information using many of the codes described above strongly depends on network architecture (Kumar et al., 2010; Hahn et al., 2019). If synaptic connectivity is high between neuronal groups, shared input gets amplified. However, this also means that random fluctuations will get amplified, particularly as they travel through successive processing stages, ultimately resulting in highly synchronous firing that can drown out the signal of interest (Mazurek and Shadlen, 2002; Reyes, 2003; Kumar et al., 2010; Hahn et al., 2019). In a population code, the precise pattern of correlations can either enhance or degrade information transmission (Abbott and Dayan, 1999).

In the above exposition of neural coding, three questions about the features defining a neural code stand-out. (1) How can downstream impact be maximized? (2) What is the timescale over which inputs are integrated? (3) How might information be selectively routed? As we shall see below, band-limited, population level oscillatory activity can provide a dynamic substrate for encoding and transmitting learning-relevant information.

1.2.3 Phase-of-firing supporting long-range encoding of information

Band-limited rhythmic activity emerges from specific circuit interactions that can often be localized to specific cortical layers (Womelsdorf et al., 2014b; Voloh and Womelsdorf, 2016). Typically, this activity falls within canonical bands, including delta (1-4 Hz), theta (4-10 Hz), alpha/beta (10-30), and gamma (30-100 Hz). Slower oscillations (particularly theta and alpha/beta) are prevalent in infragranular layers (5,6), whereas gamma is evident in granular layer 4 and supragranular layers 1 and 2/3 (Bollimunta et al., 2008; Wang, 2010; Buffalo et al., 2011; Godlove et al., 2014; Kerkoerle et al., 2014; Ninomiya et al., 2015; Bastos et al., 2018). This segregation of oscillatory frequency by layer may reflect the differential connectivity evident across cortical layers. For example, in early visual areas in macaques, gamma oscillations are evident in layer 4 (the main target of feedforward connections), and then spread to superficial

and deep layers, whereas slower alpha oscillations emerge in layer 1/2 and 5, the main targets of feedback connections (Kerkoerle et al., 2014). This concordance of oscillatory frequency and connectivity direction has been extended to higher order visual areas in parietal cortex and area 8a (Bastos et al., 2014), as well as eulaminate frontal areas (Bastos et al., 2018). It is important to recognize that agranular areas (such as the ACC) do not follow the same laminar organization as visual areas (Barbas, 2015). It is an open question if oscillatory frequency still follows laminar specific activation between areas. In one of the few studies probing oscillatory activation in frontal areas, Ninomiya and colleagues (Ninomiya et al., 2015) showed that slow and fast oscillations interacted at layer 3 of SEF, which the authors propose may be due to the propensity of SEF input and output connections to go through the same layer. Thus, this is consistent for a role for frequency-specific activity supporting integration, both within and across cortical areas.

Although ubiquitous, oscillatory activity emerges selectively during learning. A wide variety of studies have shown that ACC, LPFC, and STR coordinate via band-limited activity, particularly in the theta (4 -10Hz) and beta (10-30 Hz) band (Antzoulatos and Miller, 2014, 2016; Womelsdorf et al., 2014a; Voloh et al., 2015; Babapoor-Farrokhran et al., 2017). In ACC, theta activity prominently predicts behavioral adaptation following errors (Cavanagh et al., 2010), particularly when cognitive control demands are high (Womelsdorf et al., 2010a). Fast rhythmic activity in LPFC is structured via ACC theta during correct attentional allocation, but not during errors, suggesting ACC influences goal-oriented processing in LPFC via rhythmic activity (Voloh et al., 2015). In the beta band, synchrony among prefrontal ensembles emerges that is specific to learned rules (Buschman et al., 2012a). Following error commission, beta power scales with the degree of error, as well as the history of errors (Tan et al., 2014a). These studies thus suggest a role for theta and beta band activity encoding learning variables necessary for flexible control and adaptation.

Importantly, the alignment of oscillatory activity between distal sites can serve to selectively route information (Singer and Gray, 1995; Womelsdorf et al., 2007; Fries, 2015; Palmigiano et al., 2017; Hahn et al., 2019). If spikes from a sending region arrive at a low

excitability phase of a downstream region, they will be unable to overcome local inhibition, thus preventing the transmission of information. Alternatively, if they arrive at a high excitability phase, information can be successfully transmitted. Such a gating function has been demonstrated for macaque visual areas, where gamma-specific coordination of competing sites between area V1 and V4 during selective attention only occurs for stimuli inside, rather than outside, the receptive field (Bosman et al., 2012). The implication of this is that the frequency at which information is encoded thus naturally sets a time-window over which inputs can be integrated.

The fluctuating excitability of local circuits, evident as oscillatory activity, can be leveraged to encode information. Phase-of-firing coding can increase the informational content in spikes (Montemurro et al., 2008; Kayser et al., 2009), and can carry information related to stimulus orientation (Womelsdorf et al., 2012), objects in memory (Siegel et al., 2009), and spatial navigation parameters (Huxter et al., 2008). One way in which this could be achieved is through a rate-to-phase conversion (Mehta et al., 2002; Buzsáki, 2004; Fries et al., 2007). Neurons that are activated most strongly will be able to escape oscillatory inhibition earliest, thus firing on earlier phases.

Beyond the selection of single stimuli, oscillatory activity permits the encoding and decoding of multiplexed information (Akam and Kullmann, 2014). In *time-division multiplexing*, packets of information are ordered in time (Akam and Kullmann, 2014). Thus, a downstream reader can decode the packet information by determining the phase (relative to a reference oscillation) at which it arrived (Lisman and Jensen, 2013; Akam and Kullmann, 2014). In *frequency division multiplexing*, different information may be carried by different frequencies. A downstream reader can then decode information by evaluating changes in either the amplitude or phase, relative to a reference oscillation (Akam and Kullmann, 2010). Indeed, studies have shown that different oscillatory bands carry different information (Colgin et al., 2009; Buschman et al., 2012a; Womelsdorf et al., 2014b; Womelsdorf and Everling, 2015; Voloh and Womelsdorf,

2016), suggesting that frequency-specific oscillations may themselves reflect the segregation of task-relevant information (Ketz et al., 2015).

1.2.4 Summary

The activity of individual neurons shows a remarkable degree of variation that can be leveraged to encode information, including changes in overall intensity, changes in the timing of activation, or transitions between different dynamic states. The whole, however, is often greater than the sum of its parts; populations of neurons can more effectively convey information, either by introducing redundancy and reducing variation, or by being sensitive along orthogonal dimensions. The efficacy of neuronal codes, both in terms of the amount of information conveyed as well as the downstream impact, can be maximized when structured by band-limited rhythmic activity. While phase-of-firing coding has been observed in orientation tuning, object in memory, or stimulus identity, it is unknown how it might support more abstract features related to learning.

1.3. Thesis Overview

The focus of this thesis is to delineate how spiking activity of individual neurons structured by ongoing, population-wide rhythmic activity can facilitate the encoding of learning-relevant information. The thesis can be roughly divided into two parts. **Chapters 2 and 3** directly characterize and contrast neural coding schemes that are dependent on band-limited rhythmic activity. **Chapter 4 and 5** lay out the ground-work that would facilitate future research into neural coding.

Chapter 2 details how learning related signals are encoded across the fronto-striatal axis when firing rate modulations are structured by beta rhythms. **Chapter 3** is a follow-up study looking at whether outcome unexpectedness can be conveyed by the phase at which neurons fire. **Chapter 4** is motivated by the fact that differential dynamics within individual neurons can multiplex information. However, in order to identify how such differential dynamics may be

carried by network-level oscillations, we first seek to characterize how such dynamics are related to (local) oscillations in the beta (and theta) band. Finally, in **Chapter 5**, in recognition of the fact that the study of neural coding critically depends on task complexity, we develop a novel experimental suite that facilitates the development of complex, naturalistic tasks. This work includes three papers (published or submitted) in **Appendix C-E.**

Chapter 2

Phase of Firing Coding of Learning Variables across Prefrontal Cortex, Anterior Cingulate Cortex and Striatum during Feature Learning

2.1. Preamble

This chapter has been submitted to the journal *Neuron* (EMID:7d64234f88a6d571). It is also available as a preprint online (Voloh et al., 2019a), which can be accessed at https://doi.org/10.1101/2019.12.12.874859. I reproduce the manuscript here for clarity.

2.2. Abstract

The prefrontal cortex and striatum form a recurrent network whose spiking activity encodes multiple types of learning-relevant information. This spike-encoded information is evident in average firing rates, but finer temporal coding might allow multiplexing and enhanced readout across the connected the network. We tested this hypothesis in the fronto-striatal network of nonhuman primates during reversal learning of feature values. We found that neurons encoding current choice outcomes, outcome prediction errors, and outcome history in their firing rates also carried significant information in their phase-of-firing at a 10-25 Hz beta frequency at which they synchronized across lateral prefrontal cortex, anterior cingulate cortex and striatum. The phase-of-firing code exceeded information that could be obtained from firing rates alone, was strong for inter-areal connections, and multiplexed information at three different phases of the beta cycle that were offset from the preferred spiking phase of neurons. Taken together, these findings document the multiplexing of three different types of information in the phase-of-firing at an interareally shared beta oscillation frequency during goal-directed behavior.

2.3. Introduction

The lateral prefrontal cortex (LPFC) and anterior cingulate cortex (ACC) are key brain regions for adjusting to changing environmental task demands (Murray and Rudebeck, 2018; White et al., 2019). Both regions project to partly overlapping regions in the anterior striatum

(STR), which feeds back projections and thereby closes recurrent fronto-striatal loops (Haber and Knutson, 2010). Neurons in this recurrent network encode multiple learning variables during goal-directed behaviors, including the value of currently received outcomes, a memory of recently experienced outcomes, and a reward prediction error that indicates how unexpected currently received outcomes were given prior experiences (Hikosaka et al., 2019; Oemisch et al., 2019).

The multiplexing of outcomes, outcome history, and outcome unexpectedness (prediction errors) within the same neuronal population is likely realized by firing rate modulations in fronto-striatal brain areas (Fusi et al., 2016), but how this firing is temporally organized within the larger network is unresolved (Kumar et al., 2010; Akam and Kullmann, 2014; Panzeri et al., 2015). A large body of evidence has shown that ACC and LPFC synchronize their local activities at a characteristic beta oscillation frequency (Womelsdorf et al., 2014a, 2014b; Babapoor-Farrokhran et al., 2017; Voloh and Womelsdorf, 2017; Nácher et al., 2019), and that both areas engage in transient beta rhythmic oscillatory activity with the striatum during complex goal directed tasks (Howe et al., 2011; Leventhal et al., 2012; Antzoulatos and Miller, 2014; Feingold et al., 2015). However, whether this beta oscillatory activity is informative for learning and behavioral adjustment has remained unresolved (Womelsdorf and Everling, 2015; Spitzer and Haegens, 2017; Amemori et al., 2018). Prior studies have documented that beta activity emerges specifically during the processing of outcomes following correct trials during habit learning (Howe et al., 2011), and that following error trials overall beta activity is larger when the committed error is smaller (Tan et al., 2014a). However, these studies did not quantify whether neuronal spiking activity synchronizing to these beta oscillations contains learning relevant outcome information.

We therefore aimed to clarify how outcome related beta rhythmic activity relates to the behavioral learning of reward rules in ACC, LPFC, and STR. First, we quantify firing rate information about current outcomes, about prediction errors of these outcomes and the history of recent reward. These variables might be conveyed independently of network level beta

oscillatory activity. However, theoretical studies suggest that neuronal coding utilizing temporal organization can be efficient, high in capacity, and robust to noise (Kayser et al., 2009; Kumar et al., 2010; Kopell et al., 2011; Akam and Kullmann, 2014). Additionally, coding of information in the temporal activity pattern has been linked to mechanisms of efficient communication among neuronal groups, suggesting that coherently synchronized groups can exchange information by phase aligning their disinhibited activity periods (Singer and Gray, 1995; Womelsdorf and Fries, 2007; Hahn et al., 2014, 2019; Fries, 2015; Luczak et al., 2015; Palmigiano et al., 2017).

To this end, we recorded from LPFC, ACC, and STR while macaque monkeys engaged in trial-and-error reversal learning of feature reward rules. We found that during outcome processing, each area contained segregated ensembles of neurons encoding the current *Outcome*, the *Prediction Error* of those outcomes, and the recent *Outcome History* in their overall firing rates. A large proportion of rate coding neurons synchronized long-range to remote areas of the fronto-striatal network at a shared 10-25 Hz beta frequency range. We found that for those neurons that phase synchronized long-range, the three learning variables were encoded more precisely for spike output elicited at narrow oscillation phases in the beta band. This phase-of-firing gain of encoding exceeded the firing rate code and occurred at phases that were partly away from the neurons preferred spike phase. These findings document a multiplexing of learning variables through the phase of firing across the ACC, LPFC, and STR of nonhuman primates.

2.4. Methods

2.4.1 Animals

Data was collected from two adult, 9 and 7-year-old, male rhesus monkeys (*Macaca mulatta*) following procedures described in (Oemisch et al., 2019). All animal care and experimental protocols were approved by the York University Council on Animal Care and were in accordance with the Canadian Council on Animal Care guidelines.

2.4.2 Behavioral paradigm

Monkeys performed a feature-based reversal learning task that required covert attention to one of two stimuli based on the reward associated with the color of the stimuli. Which stimulus color was rewarded remained identical for ≥30 trials and reversed without explicit cue. The reward reversal required monkeys to utilize trial outcomes to adjust to the new color-reward rule. Details of the task have been described before (Oemisch et al., 2019). Each trial started when subjects foveated a central cue. After 0.5-0.9 sec, two black and white gratings appeared. After another 0.4 sec., the stimuli either began to move within their aperture in opposite directions (up-/downwards) or were colored with opposite colors (red/green or blue/yellow). After another 0.5-0.9 sec, they gained the color when the first feature was motion, or they gained motion when the first feature had been color. After 0.4-0.1 sec, the stimuli could transiently dim. The dimming occurred either in both stimuli simultaneously, or separated in time by 0.55 sec. Dimming represented the go-cue to make a saccade in the direction of the motion when it occurred in the stimulus with the reward associated color. The dimming acted as a no-go cue when it occurred in the stimulus with the non-rewarded color. A saccadic response was only rewarded when it was made in the direction of motion of the stimulus with the rewarded color. Motion direction and location of the individual colors were randomized within a block. Thus, the only feature predictive of reward within a block was color. Color-reward associations were constant for a minimum of 30 trials. Block changes occurred when 90% performance was reached over the last 12 trials, or 100 trials were completed without reaching criterion. The block change was uncued. Rewards were deterministic.

2.4.3 Electrophysiology

Extra-cellular recordings were made with 1–12 tungsten electrodes (impedance 1.2–2.2 MOhm, FHC, Bowdoinham, ME) in anterior cingulate cortex (ACC; area 24), prefrontal cortex (LPFC; area 46, 8, 8a), or anterior striatum (STR; caudate nucleus (CD), and ventral striatum (VS)) through a rectangular recording chambers (20 by 25 mm) implanted over the right hemisphere (Figure A-1). Electrodes were lowered daily through guide tubes using software-controlled precision micro-drives (NAN Instruments Ltd., Israel and Neuronitek, Ontario, Canada). Data

amplification, filtering, and acquisition were done with a multichannel acquisition system (Neuralynx). Spiking activity was obtained following a 300–8000 Hz passband filter and further amplification and digitization at 40 kHz sampling rate. Sorting and isolation of single unit activity was performed offline with Plexon Offline Sorter, based on analysis of the first two principal components of the spike waveforms. Experiments were performed in a custom-made sound attenuating isolation chamber. Monkeys sat in a custom-made primate chair viewing visual stimuli on a computer monitor running with a 60 Hz refresh rate. Eye positions were monitored using a video-based eye-tracking system (EyeLink, SRS Systems) calibrated prior to each experiment to a nine-point fixation pattern. Eye fixation was controlled within a 1.4°–2.0° radius window. During the experiments, stimulus presentation, monitored eye positions, and reward delivery were controlled via MonkeyLogic (www.brown.edu/Research/monkeylogic/). Liquid reward was delivered by a custom-made, air-compression controlled, and mechanical valve system. Recording locations were aligned and plotted onto representative atlas slices (Calabrese et al., 2015).

2.4.4 Data Analysis

Analysis was performed with custom Matlab code (Matlab 2019a), using functions from the fieldtrip toolbox (http://www.fieldtriptoolbox.org). For Elastic-net regression, the *glmnet* package in R was used (Friedman et al., 2010). Only correct and error responses were analyzed. Error responses included those where the responses were made to the incorrect target, or in the incorrect response window. Only trials from learned blocks were included, with a minimum of two blocks, unless otherwise indicated. Learned blocks were defined as ones where animals reached 90% correct responses within the last 10 trials within the block. Standard errors of the median were estimated via bootstrapping (200 repetitions, unless otherwise indicated).

Before continuing, we provide a coarse overview of the analysis steps. We first determined if experienced outcomes affected choices at the level of monkeys' behavior. After showing that experienced choices are in fact integrated, we next asked how this is reflected in the firing rate activity of individual neurons during a post-outcome period using a penalized GLM.

We use a data-driven clustering approach to assign functional labels to cells exhibiting similar sensitivities to experienced outcomes in their rate. On the basis of these functional labels, we extract a corresponding encoding metric for neurons in each functional cluster. We then analyze how the encoding metrics depend on time, or the phase of oscillatory activity in the local field potential. For the latter analysis, we use standard spectral decomposition techniques and spike-phase consistency measures to characterize how spikes and phases between distal electrodes are related. We quantify and compare differences in phase-dependent encoding in terms of (1) the degree of phase-dependent modulation of encoding, and (2) the phase at which encoding is maximal.

2.4.5 Behavioral analysis

To determine the timescale over which past outcomes are integrated, we used a binomial GLM:

$$\mathrm{Y} = \sum_{i}^{5} \beta_{t-i} X_{t-i}$$

Where Y was the current outcome, B_{t-i} is the influence of outcome X_{t-i} on trial t-i. Outcome for trial t-5 was defined as a nuisance variable that accounted for all responses occurring over very long time-scales (similar to (Walton et al., 2010; Rudebeck et al., 2017)).

2.4.6 Rate encoding of outcome history

To test how individual units integrated outcome history, we used a Poisson GLM:

$$\log(\lambda) = \sum_{i}^{5} \beta_{t-i} X_{t-i}$$

Where λ was the conditional intensity (spike count), B_{t-i} is the influence of outcome X_{t-i} on trial t-i. Firing counts on each trial were determined in a [0.1 0.7]s window after outcome onset (Oemisch et al., 2019). Neurons were included in the analysis if they were isolated for more than 25 (learned) trials across at least two blocks, and if they showed an overall firing rate of >1 Hz.

To mitigate issues of multi-collinearity, and extract only the most predictive regressors, we employed elastic-net regularization using the R package *glmnet* (Friedman et al., 2010). This procedure shrinks small coefficients to zero, and smoothly interpolates between ridge and lasso regularization by controlling a parameter alpha (with alpha=0 corresponding to ridge regression, and alpha=1 to lasso regression) (Friedman et al., 2010). We used an alpha of 0.95, which tends to select only one regressor in the presence of collinearity (as in pure lasso regression, (Tibshirani, 1996)), while at the same time avoiding issues with degeneracy if correlations among regressors are particularly strong (Friedman et al., 2010; Runyan et al., 2017). The optimal value of the shrinkage parameter (lambda) was the minimum as selected by 10-fold cross validation. To assess model stability and extract significant fits, we used a bootstrap approach, whereby trials were sampled with replacement 1000 times and the procedure was rerun. As the LASSO shrinks non-valuable predictors to zero, a model fit was said to be significant if at least one relevant regressor (outcome t-4 to t-0) was non-zero more than 95% of the time.

2.4.7 Functional clustering based on neural encoding

Our ultimate goal is to describe how encoding varies as a function of phase (and time). However, encoding cells showed variability in how they responded to experienced outcome (eg **Figure 2-2**). Thus, in order to properly evaluate changes in encoding in time and phase, we must first define populations of cells that encode similar types of information. To determine the putative function of significantly encoding units, we used a clustering approach via bootstrapped K-means. We clustered cells on the basis of their mean beta weights as determined by the penalized regression model (see above). As a preprocessing step, for units where the current outcome was negatively encoded (i.e. encodes errors), we flipped the sign of every coefficient in that model. This has the effect of erasing the directionality of any functional association, and thus collapses neurons with similar functions (for example, *Error* or *Correct* encoding units become *Outcome* encoding units). Cells were independently clustered for each area.

We clustered cells on the basis of their clustering stability. We opted for this method because clustering can be sensitive to initial conditions (Hastie et al., 2009). This involved three

steps: (1) choosing the optimal number of clusters *Nc*, (2) measuring clustering stability, and (3) performing the final clustering.

For steps 1-2, we used K-mean clustering with a cosine distance metric, which is insensitive to the magnitude of the vector and is instead concerned with the direction, unlike, for example, Euclidian distance. In other words, we clustered based on the *relative pattern* of beta weights of each cell, irrespective of differences in magnitude between cells.

To determine the optimal number of clusters, we extracted the Silhouette metric over many bootstrap iterations. Briefly, cells were sampled with replacement 1000 times and for each iteration, the optimal number of clusters was extracted where the silhouette was maximal. The overall optimal number of clusters *Nc* was the mode over all bootstrap iterations.

Next, we assessed the clustering stability of pairs of cells. To do so, we built a similarity matrix *S* via a bootstrap approach, where similarity was defined as the proportion of times that pairs of cells were clustered together. First, we resampled with replacement individual cells. Next, we ran K-means with cosine distance and *Nc* clusters. For units that were clustered together, their respective cell in the similarity matrix was incremented by one. Because bootstrapping could sample the same units twice, these pairs were ignored. Bootstrapping was run 100000 times.

To compute the final cluster assignment, we first formed a dissimilarity matrix D=1-S, before performing agglomerative clustering with Euclidian distance and *Nc* clusters.

2.4.8 Metrics for outcome, outcome history, and prediction error

The degree of encoding of *Outcome* ($E_{outcome}$), *Outcome History* ($E_{history}$), and *Prediction Error* (E_{pe}) was quantified on the basis of the GLM weights for trials -1 and 0:

$$E_{outcome} = abs(B_0)$$
$$E_{history} = abs(B_{-1} + B_0)$$
$$E_{PE} = abs(B_0 - B_{-1})$$

We refer to these generically as *Encoding Metrics*.

2.4.9 Latency analysis

To determine the latency of encoding for each functional cluster, we performed a time-resolved analysis **(Figure A-1C,D**). On the basis of our previous results showing that the outcome on trial 0 and -1 were most predictive **(Figure 2-1**), we used a simpler GLM of just the current and previous outcome. For the response variable, we calculated the spike density using a sliding Gaussian window, with a 200 ms window and 50 ms standard deviation. We performed this analysis [-0.4 0.7] around outcome onset. We thus obtained a time-resolved estimate of encoding.

To determine the latency of significant encoding, we looked at time points in the post-outcome period that were significantly different from the pre-outcome period. We thus determined, for each cell, when encoding exceed a threshold criteria in a time of interest. First, we z-score normalized each individual cell's encoding to the pre-outcome period ([-0.4 0] sec). Next, we asked, for each time point, whether the population response was significantly different from zero via a Wilcoxon signrank test. We then extracted the largest cluster mass of contiguous significant time points (at an alpha of 0.05, *see* (Maris and Oostenveld, 2007)) to find a *time-of-interest*. Finally, we extracted, for each individual cell, the time point where the area under the curve of the encoding metric in this time-of-interest reached 10% of the total. Thus, we obtain for each encoding cluster a distribution of latencies of when they started to show significant encoding of *Outcome*, *Outcome History* or *Prediction Error*.

2.4.10 Spectral decomposition and spike-LFP phase synchronization

To determine how encoding varied as a function of phase, we extracted the estimate of phase at the time of spikes, for a large frequency range. We first characterized the degree of spike-phase synchronization, described below.

We focused spike-phase analysis on pairs of distally recorded sites, thus obviating any concerns of spike energy bleeding into the LFP (Zanos et al., 2011). We considered frequencies from 6 Hz to 60 Hz. For frequencies from 6-30Hz, the resolution was 1 Hz, and above that it was 2Hz. For every frequency F, we determined the spike-LFP phase by extracting an LFP segment

centered on the spike of length 5/F (i.e. 5 cycles), as is standard to balance temporal and spectral resolution (Cohen, 2014). Spectral decomposition was done via an FFT after applying a Hanning taper. This procedure was applied separately to the pre-outcome period [-1 0]sec, and the post-outcome period [0.1 1] sec.

The strength of spike-LFP synchronization was quantified using the pairwise phase consistency (PPC), which is unbiased by the number of spikes (Vinck et al., 2011). The PPC is quantified on the basis of pairwise differences between spike-phases. If spikes tend to fire on specific phases, phase difference will be concentrated, and thus the PPC will take on a high value, whereas if spikes are distributed randomly relative to the LFP phase, phase differences will be random and the PPC will tend towards zero.

The effect size was determined as previously reported (Womelsdorf et al., 2014a; Voloh and Womelsdorf, 2017):

$$Effect\ size = \frac{1 + 2 * sqrt(PPC)}{1 - 2 * sqrt(PPC)}$$

This effect size can be interpreted as the relative increase in spike rate at the cell's preferred (mean) phase over its anti-preferred (opposite) firing phase. For example, a PPC value of 0.01 corresponds to a 1.5 times greater spike rate at the preferred phase.

We determined the frequency at which spike-LFP phase synchronization was significant by determining peaks in the PPC spectrum. A cell was said to synchronize to a particular frequency if the following criteria were met: (1) Peaks had to be above a threshold of 0.005, (2) show a minimum prominence of 0.005, and (3) show significant Rayleigh test (i.e. phase concentration).

To test for inter-areal differences in spike-beta synchronization, we extracted the maximal significant/prominent PPC peak in the [10 25] Hz band for cells that showed significant encoding. For those encoding cells that did not show significant PPC peaks, we extracted the

maximal PPC in this band instead. We tested for differences in synchronization strength using a one-way ANOVA, and report on pairwise comparisons after multiple comparison correction.

2.4.11 Phase-of-Firing dependent encoding of Outcome, Outcome History and Prediction Error

To determine if spikes falling on certain phases of the LFP were more informative, we re-ran the (reduced) GLM on phase-binned spikes, using only the previous and current outcomes (see latency analysis). We first aligned all spike-triggered-phases to the circular mean of their distribution. Phases were extracted from the frequency of the corresponding maximal peak in the [10 25] Hz band in the PPC. However, if spikes are phase locked to an LFP, the firing rate around the preferred phase will naturally be higher. Thus, we used non-equal bin sizes, adjusting the bin limits such that they contained approximately equal numbers of spikes, and phase-labelled individual spikes accordingly. We then re-ran the GLM analysis on spikes falling within a particular bin and computed the encoding metrics as described previously. To aid in comparison, we also fit the model using randomly permuted phases (thus preserving the over-all rate response structure). We ignored spike-LFP pairs where the GLM could not converge to a solution and threw a warning, or where the beta coefficients were above 20 (however, relaxing or tightening this constraint did not change the results).

To determine if the phase and degree of phase-dependent encoding, we fit a cosine function to the phase-binned encoding values (illustrated in **Figure 2-4A**) (Siegel et al., 2009; Womelsdorf et al., 2012). From this fit, we obtain three values: T (phase offset, or phase of cosine maxima), A (amplitude), and M (overall mean, or offset). The value T is thus the phase at which encoding is maximal, relative to the preferred firing phase. To compare the strength of encoding across functional clusters, we computed the *empirical phase-of-firing gain*:

$$PFG_E = 2 * \frac{A}{M}$$

This quantity represents the difference in encoding between the peak and trough relative to the overall encoding strength. A PFG value of 0 implies that phase-of-firing adds no information (corresponding to a pure rate code), whereas PFG=1 means that encoding between the peak and

trough is 100% stronger compared to the overall encoding strength. To determine if phase significantly added information above that of a rate code, we applied a randomization approach. For each cell, we first permuted the phase label of each spike, re-ran the GLM, re-fit a cosine and extracted the phase-of-firing gain. This procedure was repeated 50 times. The randomized phase-of-firing gain PFG_r was the mean of this distribution. We report on the "excess" phase-of-firing gain, defined as the difference between empirical and randomized phase-of-firing gain, which we refer to in the manuscript at the *Encoding Phase Gain (EPG):*

$$EPG = PFG_E - PFG_R$$

A positive value implies that encoding is modulated by phase above what would be expected by chance.

We tested the stability of encoding across phase bins for each neuron (with significant rate encoding) by determining the sign of the encoding metric (i.e. before taking the absolute value). We found that for the vast majority of cell-LFP pairs (~90%), the sign of the encoding metric was the same for all 6/6 phase bins as for the full model.

To test the frequency specificity of the *EPG*, we extended the above analysis to the larger 6-60 Hz frequency range (**Figure 2-4F**). We statistically tested the *EPG* across frequencies using the Wilcoxon signrank test. We also tested the degree to which our results may be influenced by cue-aligned activity. To this end, we first obtained the average evoked potential for each LFP channel, and subtracted this component from individual trials. We then performed all steps of the analysis again to compare the original *EPG* with the *EPG* free from potential cue-aligned biases.

To test whether the preferred firing phase or relative phase with maximal encoding was concentrated above what would be expected by chance, we used the circular Hodges-Ajne test **(Figure 2-5)**. To determine whether the phase showing maximal encoding differed from the preferred firing phase in each functional encoding cluster, we performed the Median test to test if the phase differed from zero (Zar, 2010) (**Figure 2-5B**).

We tested how the strength of phase synchronization related to the strength of phase-of-firing encoding by performing two analysis. First, we compared encoding in cells that showed significant spike-phase synchronization to those that did not. For non-synchronizing cells, we selected the center frequency with the maximal PPC in the [10 25] Hz range, and computed the *EPG* at this frequency. We compared *EPG* between locking and non-locking populations using the Kruskal Wallis test (**Figure 2-4G)**. Second, we asked whether spike-phase synchrony in different trial conditions contained similar information to that of the phase-of-firing. To this end, for each encoding cell, we compared trials that were predicted to have the maximal firing rate differences. For *Outcome* encoding, we compared correct *versus* error trials. For *Prediction Error* encoding, we compared correct trials following error *versus* following error trials following correct. For *Outcome History* cells, this was errors followed by errors *versus* correct outcomes followed by correct. We took the absolute difference of the PPC between the two conditions, and correlated it with the *EPG* of the respective cell using the Spearman rank correlation.

We also tested whether phase gain depended on the number of bins used to fit the cosine function. We performed the analysis for 4, 6, 8, and 10 bins. We used Spearman rank correlation to determine if *EPG* was related to the number of bins, and circular-linear correlation to associate the phase of maximal encoding with the number of bins (Zar, 2010).

2.4.12 Cell-type classification and analysis

To determine if phase-modulated encoding of information differed based on cell type, we focused the following analysis on highly isolated single units that showed encoding of learning-relevant variables and significant, prominent spike-beta locking. Detailed information is provided in (Oemisch et al., 2019). Briefly, to distinguish putative interneurons (narrow-spiking) and putative pyramidal cells (broad-spiking) in LPFC and ACC, we analyzed the peak-to-trough duration and the time for repolarization for each neuron. After applying Principal Component Analysis (PCA) using both measures, we used the first principal component to discriminate between narrow and broad-spiking cells. This allowed for better discrimination than using either measure alone. We confirmed that a two-Gaussian model fit the data better than a one-Gaussian

model using the Akaike and Bayesian Information Criterion (AIC, BIC). We then used the two-Gaussian model to define narrow and broad-spiking populations.

A similar analysis was applied to striatal units to distinguish putative interneurons from medium spiny projection neurons (MSN). Here, we use the peak-width and Initial Slope of Valley Decay (ISVD) (Berke, 2008; Lansink et al., 2010):

$$ISVD = \frac{V_t - V_{0.26}}{A_{PT}}$$

where V_T is the most negative value (trough) of the spike waveform, $V_{0.26}$ is the voltage at 0.26 ms after V_T, and A_{PT} is the peak-to-trough amplitude (Lansink et al., 2010). After PCA and two-Gaussian modelling (as described above), we defined two cut-off points. The first cutoff was the point at which the likelihood of narrow spiking cells was 3 times larger than the likelihood of broad-spiking cells, and vice-versa for the second cutoff.

We compared differences in *Encoding Phase Gain* between narrow and broad spiking neurons using the Kruskal-Wallis test, independently for each area. To clarify, we analyzed spike-LFP pairs here; thus, the same neuron may be included more than once.

2.5. Results

Animals performed a feature-based reversal learning task (Oemisch et al., 2019). Subjects were shown two stimuli with opposite colors and had to learn which of them led to reward (**Figure 2-1A**). The same color remained associated with reward for at least 30 trials before an uncued reversal switched the color-reward association (**Figure 2-1B**). During each trial, the subjects monitored the stimuli for transient dimming events in the colored stimuli. They received a fluid reward when making a saccade in response to the dimming of the stimulus with the reward associated color, while the dimming of the non-reward associated color had to be ignored. A correct, rewarded saccade to the dimming of the rewarded stimulus had to be made in the up- or downward direction of motion of that stimulus. This task required covert selective attention to one of two peripheral stimuli based on color, while the overt choice was based on the motion

direction of the covertly attended stimulus. In 110 and 51 sessions from monkey's HA and KE, respectively, we found that subjects attained plateau performance of on average 80.2% (HA: 78.8%, KE: 83.6%) within 10 trials after color reward reversal (**Figure 2-1C**). Using a binomial General Linear Model (GLM) to predict current choice outcomes (correct or error outcomes, excluding fixation breaks), we found that for both subjects, outcomes from up to three trials into the past significantly predicted the current choice's outcome (**Figure 2-1D**; Wilcoxon signrank test, $p<0.05$, multiple comparison corrected), closely matching previous findings (Walton et al., 2010).

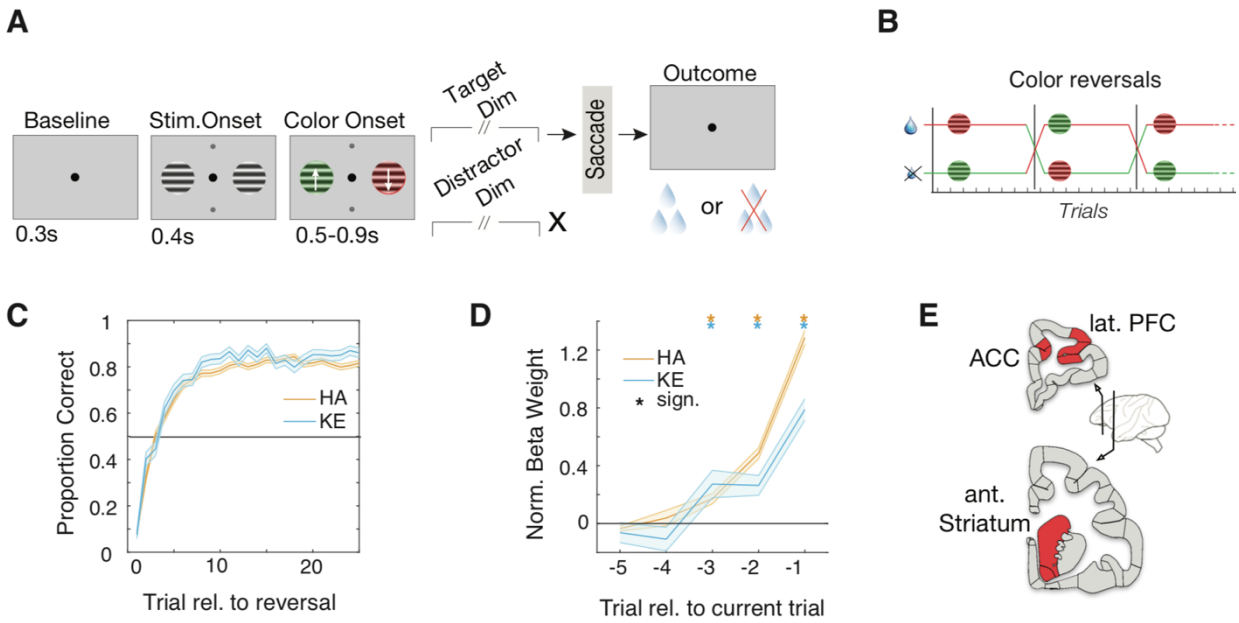

Figure 2-1. Subjects learned reversals

(A) Feature-based reversal learning task. Animals are presented with two black/white stimulus gratings to the left and right of a central fixation point. The stimulus gratings then become colored and started moving in opposite directions. Dimming of the stimuli served as a Go signal. At the time of the dimming of the target stimulus, the animals had to indicate the motion direction of the target stimulus by making a corresponding up or downward saccade in order to receive a liquid reward. Dimming of the target stimulus occurred either before, after or at the same time as the dimming of the distractor stimulus. **(B)** The task is a deterministic reversal learning task, whereby only one color is rewarded in a block. This reward contingency switches repeatedly and unannounced in a block-design fashion **(C)** Accuracy relative to block start for monkey HA (orange) and KE (blue). Shaded region represents the standard error. Subjects achieved plateau performance within 5-10 trials **(D)** Median beta coefficients from a binomial regression of current outcome as predicted by past outcomes. Shaded region represents the standard error. Outcomes up to three trials into the past predicted current outcome (Wilcoxon signrank. $p<0.05$, multiple comparison corrected). **(E)** Schematic depicting recorded brain areas.

2.5.1 ACC, LPFC and STR Neurons Encode Outcomes, their History and their Prediction Error

To test how previous and current outcomes are encoded at the single neuron level, we analyzed a total of 1460 neurons, with 332/227 (monkey HA / KE) neurons in lateral prefrontal cortex (LPFC), 268/182 neurons in anterior cingulate cortex (ACC), and 221/230 neurons in anterior striatum (STR) (**Figure 2-1E, Figure A-1A**). These regions have previously been shown to encode outcome, outcome history, and prediction error information ((Ito, 2003; Bernacchia et al., 2011; Yamada et al., 2011; Lee et al., 2012; Shen et al., 2014; Hikosaka et al., 2017; Oemisch et al., 2019). We found multiple example neurons encoding different types of outcome variables. Some cells responded differently to correct versus erroneous trial outcomes irrespective of previous outcomes **(Figure 2-2A**), while others responded strongest when the current outcome deviated from the previous trials' outcome (signaling reward prediction error) **(Figure 2-2B**), or when the current outcome was similar to the previous trials' outcome, i.e. following a sequence of correct trials or a sequence of error trials **(Figure 2-2C**).

We quantified these types of outcome encoding using a LASSO Poisson GLM that predicted the spike counts during the outcome period (0.1-0.7 sec after reward onset) and extracted the characteristic patterns of beta weights across the past and current trial outcomes that distinguished different types of outcome encoding **(Figure 2-2D**). Neurons that encoded only the current trials' outcome showed large weights only for the current trial (*Outcome* encoding type). Neurons encoding a prediction error showed beta weights for previous trials that were opposite in sign to the current trial's outcome (*Prediction Error* encoding type). In neurons encoding the history of recent rewards, beta weights ramped up over recent trials toward the current trial outcome (*Outcome History* encoding type) (*see* also insets in **Figure 2-2A-C**). We found in a clustering analysis that these three types of outcome encoding were prevalent in each of the recorded brain areas **(Figure 2-2E-G, Figure A-1B**). Neurons encoding the current outcome were the most populous (~59%, 234/384 in monkey HA and 185/329 in monkey KE), followed by ~26% of neurons encoding *Prediction Errors* (64/231 in monkey HA and 39/170 in monkey KE) and ~32% of neurons encoding *Outcome History* (76/206 in monkey HA and 33/139 in monkey KE) **(Figure 2-2E-G**; χ^2 test, $p<0.05$). The relative frequency of these encoding types did not differ

between areas (χ^2 test, p=0.46). Across the population of neurons in ACC, LPFC and STR we found that *Outcome*, *Prediction Error* and *Outcome History* encoding emerged shortly after outcomes were received (*see* **Methods, Figure A-1C,D**, Wilcoxon signrank test, p<<0.05). Neurons encoding *Outcome, Prediction Error*, or *Outcome History* showed similar overall firing rates **(Figure A-1E;** Kruskal Wallis test, p_{LPFC}=0.27, p_{ACC}=0.58, p_{STR}=0.35).

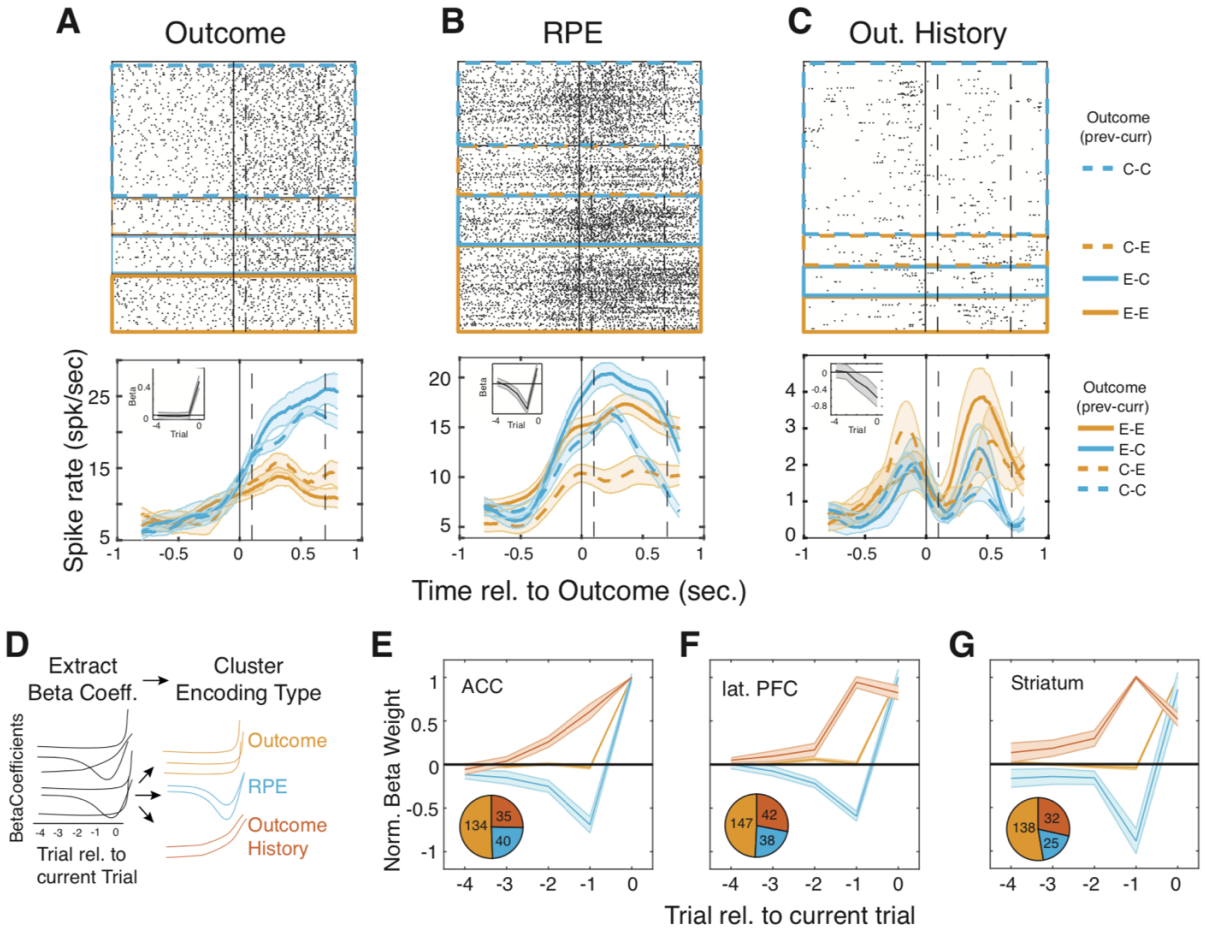

Figure 2-2. Outcomes, prediction errors, and outcome history is encoded across the fronto-striatal axis

(A-C) Examples of *Outcome (A)*, *RPE (B)*, and *Outcome History (C)* cells. (*top)* Raster plot for four separate trial conditions. Color denotes outcome on the current trial (blue=correct, yellow=incorrect), and line style to the previous trial (solid=error, dashed=correct). "C" is correct trials, "E" is error trials. The dotted vertical lines represent the time period used for the GLM analysis. (*bottom*) Time-resolved firing rate for the four different trial types. Insets in the top left depict the average GLM beta coefficients (*see* next panel, Methods) **(D)** Sketch of general approach. We regressed total spike counts in the [0.1 0.7] outcome period (dashed vertical lines, panels A-C) against outcomes up to 5 trials into the past using a penalized LASSO GLM. Beta coefficients were then clustered to group cells according to their most parsimonious functional designation. **(E-G)** Median normalized beta coefficients for three functional clusters in ACC (E), LPFC (F), and STR (G). Shaded regions represent the standard error. On

the basis of the pattern of beta weights (see text), each area exhibits "*Outcome*" cells (yellow), "*Outcome History*" cells (red), and "*RPE*" cells (blue). *(inset)* Relative frequency of each functional cluster. Outcome cells were the most populous in all three regions (Chi-squared test, p<<0.05).

2.5.2 Neuronal Synchronization at 10-25 Hz Beta Band across ACC, LPFC and STR

We found similar proportions and activation time courses of encoding neurons in ACC, LPFC and STR (Figure A-1**C,D**), which raised the question how these neuronal populations are connected. One possibility is that neuronal firing patterns are organized temporally, such that spikes in one area phase synchronize to neuronal population activity in remote areas. We assessed the degree of phase consistency of neuronal spikes with local field potential (LFP) fluctuations in distally recorded areas using the Pairwise-Phase Consistency metric (PPC) (**Figure 2-3A**; (Womelsdorf et al., 2010b; Vinck et al., 2011), *see* **Methods**). Across all (n=7938) spike-LFP pairs, we found a pronounced peak of phase synchronization in the beta band (10-25 Hz) with neurons firing on average ~1.15 times more spikes on their preferred (average) phase than at the opposite phase (**Figure 2-3B**). Neurons that encoded outcome variables in their firing rates showed similar phase synchrony as neurons not encoding reward outcome information (**Figure 2-3B**). Across all pairs, 55% (4320/7938) showed significant phase synchronization within the 10-25 Hz range (**Figure 2-3C;** Rayleigh test, p<0.05, *see* Methods for prominence criteria). Among those neurons that showed significant firing rate encoding of *Outcome*, *Prediction Error* or *Outcome History,* the most prevalent within-area beta synchrony was evident within the ACC, compared to LPFC (One-way ANOVA, p=0.014) or Striatum (p~0) (**Figure 2-3D**,**E**). Neurons in ACC were also more likely than expected by chance to show between-area spike-LFP synchrony, as compared to LPFC (p~0) and STR (p=0.042). There was also a trend for stronger between-area synchrony with spikes originating in STR, as compared to LPFC (p=0.058) (**Figure 2-3D**,**F**). Testing for the reciprocity of beta band phase synchrony showed that ACC spikes phase synchronized more strongly to beta in the LPFC than vice versa (p=0.047) **(Figure 2-3G**). Intra-areal LPFC and STR pairs showed statistically indistinguishable spike-phase synchrony strength (p=0.92), as did ACC and STR pairs (p=0.26). Similar findings were evident for the strength of inter-areal spike-LFP phase synchrony (**Figure A-2**). Both monkeys showed stronger synchronization for spikes originating in ACC compared to LPFC and STR spike output (**Figure A-3B**)**.** These results show that

neurons with firing rate information about *Outcome*, *Prediction Error*, or *Outcome History* synchronized their firing within and between areas at a characteristic 10-25 Hz frequency range.

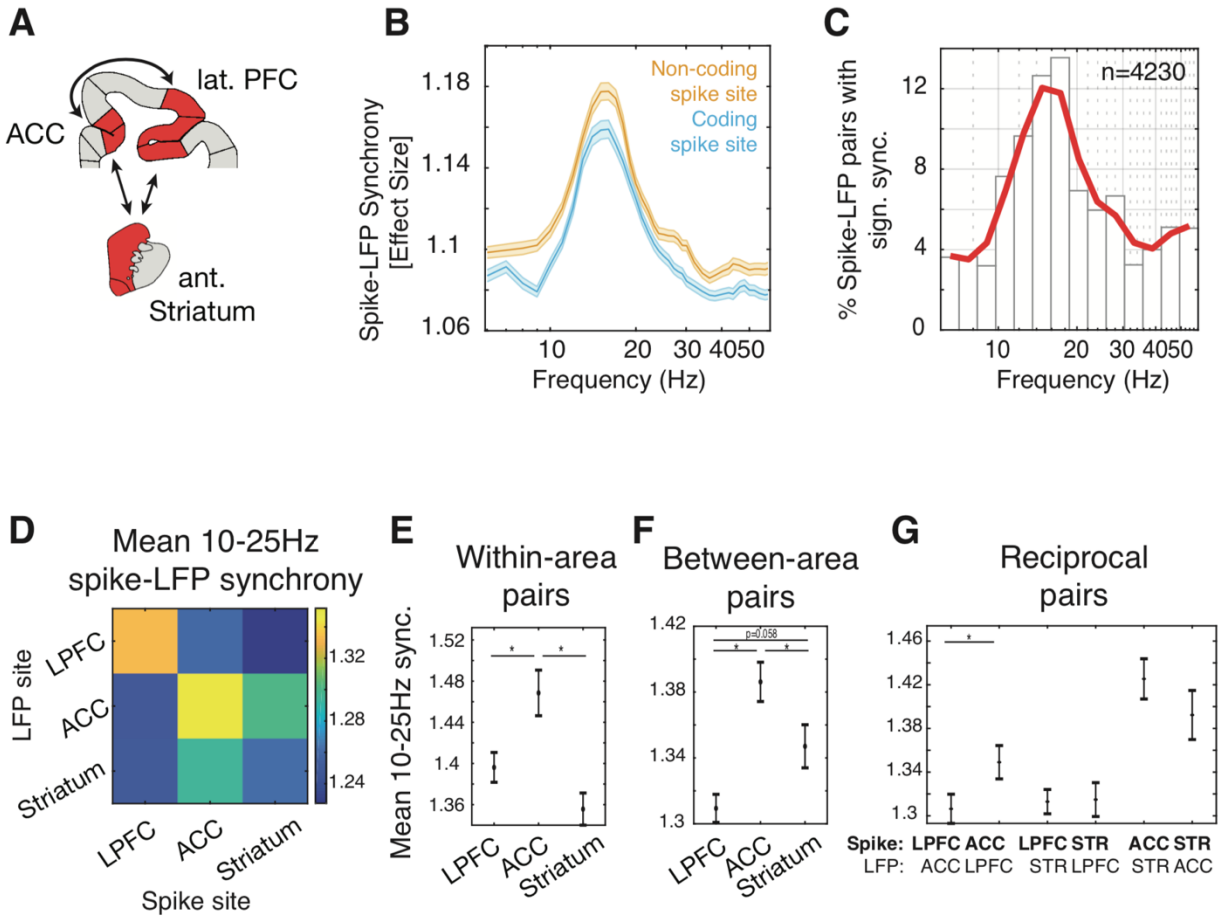

Figure 2-3. Spike-phase synchronization in beta band tends to occur in ACC

(A) We tested for functional connectivity between the ACC, LPFC, and STR**. (B)** Average PPC effect size between distal LFP and spikes at encoding (blue) and non-encoding (orange) spike sites. Shaded area represents the standard error. There is a prominent beta peak for both. X-axis is depicted on a log-scale for clarity. **(C**) Proportion of spike-LFP pairs that exhibited significant, prominent locking (*see* Methods). **(D)** Average inter-areal synchrony between all pairs of areas for signification, prominent beta in the [10-25] Hz range. Color denotes the average. **(E)** Mean and standard error contrasting spike-phase locking within different areas (diagonal in Fig 2-3D). ACC synchronizes more strongly that either LPFC or STR. **(F)** Same as (E) but for LFPs originating in other areas (i.e. summing across columns, less the diagonals, in Fig 2-3D). ACC synchronizes more strongly to distal beta, compared to LPFC or STR. **(G)** All pairwise comparison between regions. ACC spikes lock more strongly to LPFC beta than the inverse.

2.5.3 Phase-of-Firing at 10-25 Hz Encodes Outcome, Prediction Error and Outcome History

Neurons that synchronized to the LFP elicit more spikes at their mean spike-LFP phase, which we denote as the neurons' preferred spike phase. This preferred spike-phase might thus

be important to encode information shared among areas of the network (Womelsdorf et al., 2014b; Fries, 2015). We tested this hypothesis by quantifying how much *Outcome*, *Prediction Error*, and *Outcome History* information is available to neurons at different phase bins relative to their mean phase. If the phase-of-firing conveys information, then differences in spike counts between conditions should vary systematically across phases, as opposed to a pure firing rate code that predicts equal information for spike counts across phase bins (**Figure 2-4A**) (Womelsdorf et al., 2012; Hawellek et al., 2016). **Figure 2-4B** shows an example neuron exhibiting phase-of-firing coding (with spikes from ACC and beta phases from STR). This neuron exhibits large firing increases on error trials compared to correct trials, but only when considering spikes near its preferred spike phase; on the other hand, firing at the opposite phase showed a weaker difference. To quantify the phase-of-firing code for all three information types, we selected for each neuron the frequency within 10-25 Hz that showed maximal spike-LFP synchrony, subtracted the mean (preferred) spike phase to allow for comparison between neurons, and binned spikes on the basis of the LFP beta phases. To prevent an influence of overall firing rate changes between phase bins, we adjusted the width of each of the six phase bins to have (on average) an equal spike count across bins. We then fitted a GLM to the firing rates of each phase bin separately to quantify the *Outcome*, *Prediction Error*, and *Outcome History* encoding for each phase bin, and compared this phase specific encoding to a null distribution after randomly shuffling the spike phases (*see* **Methods**). **Figure 2-4C** illustrates example neurons for which the encoding systematically varied as a function of phase (*see also* **Figure A-4**). The example spike-LFP pair from **Figure 2-4B** encoded the trial *Outcome* significantly stronger than a rate code in spikes near [$-\pi/2$, $\pi/2$] radians relative to its preferred spike phase and weaker than a firing rate code at opposite phases (**Figure 2-4C**, *left*); Similar phase-of-firing encoding was evident for *Prediction Error* and *Outcome History* as independent variable (**Figure 2-4C***, middle and right panels)*.

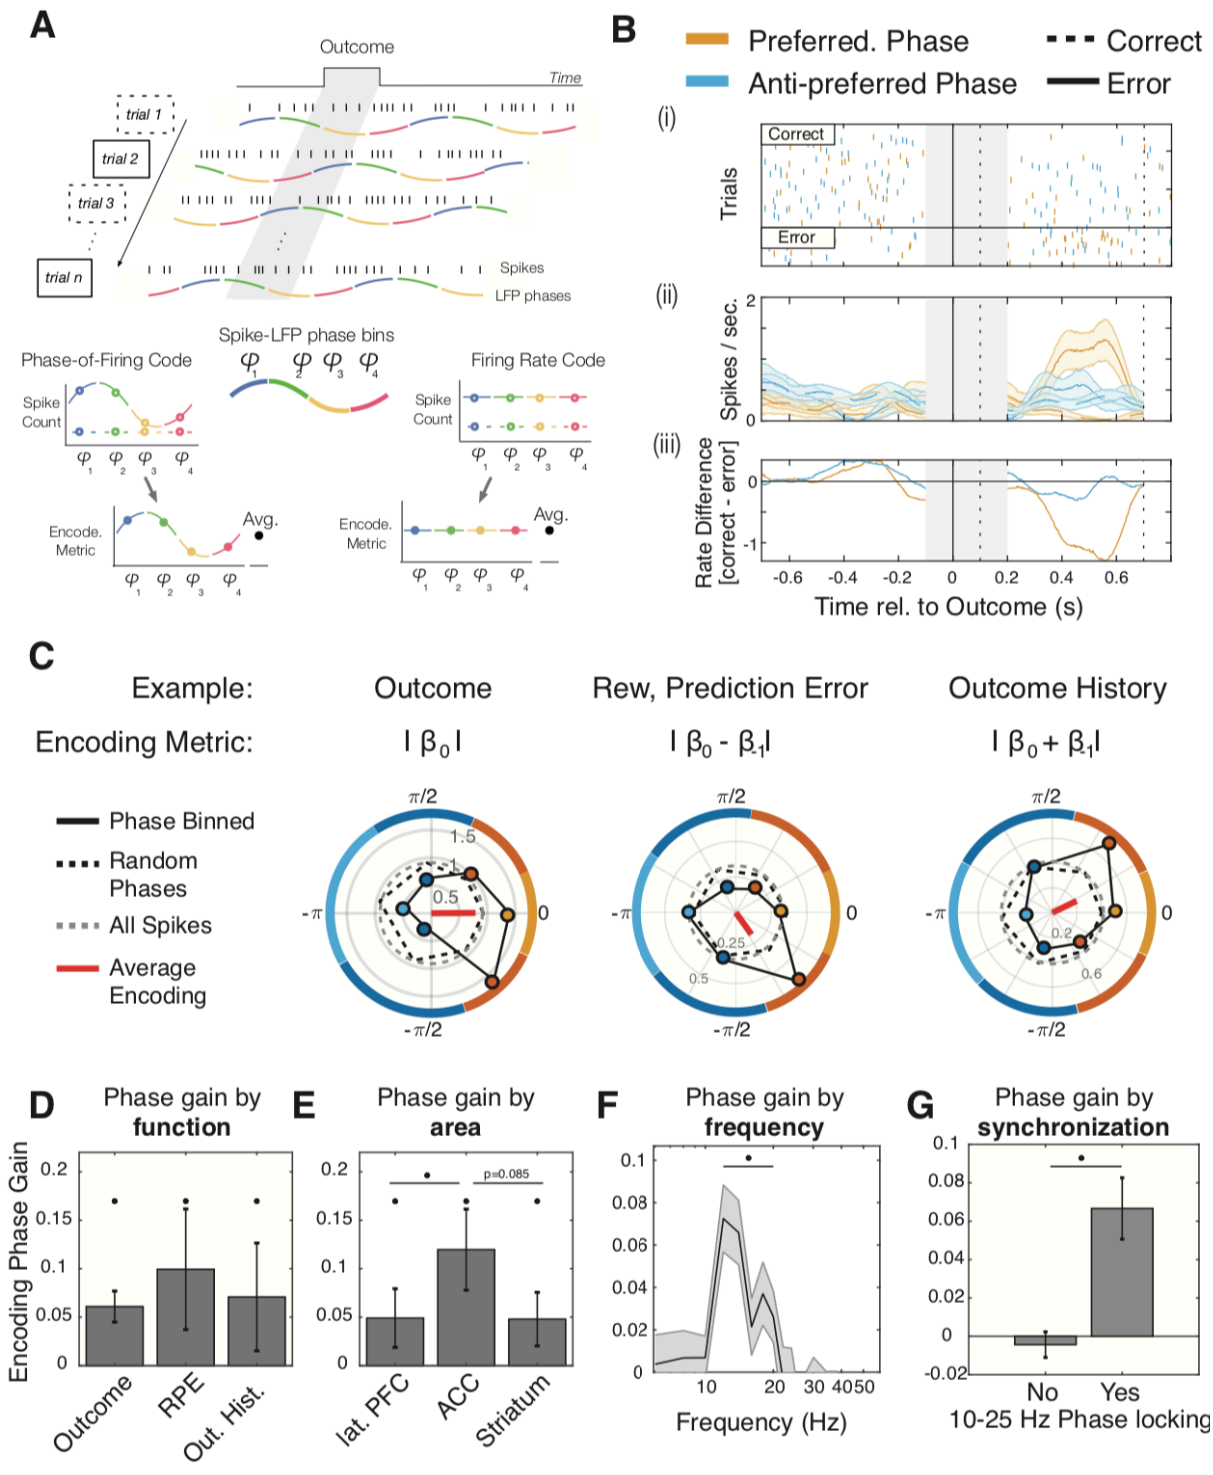

Figure 2-4. Encoded learning signals are modulated by beta phase

(A) Sketch of phase-dependent encoding analysis. Spikes were segregated according to the phase of the LFP, with bin limits adjusted such that an approximately equal number of spikes fell within each bin. Afterwards, a reduced GLM was run to extract the appropriate functional encoding metric for each phase bin. **(B)** Example cell showing phase-dependent outcome encoding. (i) Raster plot depicting the

firing of the cell on correct (top) or error (bottom) trials. Color represents firing on the maximal (orange) or minimal (yellow) encoding phase. The grey box between [-0.1 0.2] represents spikes that were not analyzed due to this window not being used for phase estimation. Vertical dashed lines depict the period used for GLM analysis. Horizontal black line splits correct from error trials. (ii) Spike density function for correct/error trials, using spikes from the maximal/minimal encoding phase. (iii) The difference in firing rate between outcomes for the maximal/minimal encoding phase is visualized as a difference. The difference is greater on preferred rather than anti-preferred phases for an extended period of time. **(C)** Example of phase-of-firing encoding for Outcome, RPE, and Outcome History cells. Colored dots reflect the encoding metric (derived from model coefficients, *see also* Methods) for a corresponding phase bin. Zero radians corresponds to the preferred (mean) firing phase. Numbers on concentric circles are the value of the encoding metric. The grey dotted line represents the encoding metric estimated using all spikes, whereas the black dotted line represents the average across many permutations of spike phases. The red line is the average direction. Colored border lines represent size of the bin. These cells show stronger encoding near the 0 rad. phase, and weaker encoding on opposite phases. **(D)** Median and standard error of the *EPG (Encoding Phase Gain)* in each functional cluster. All clusters showed some evidence of significant phase gain (Wilcoxon signrank test, two-sided, $p<0.05$). **(E)** Median phase gain for each area. ACC had a higher phase gain than LPFC or STR. **(F)** Median phase gain per frequency for encoding and locking cells (n=877). Phase gain was above chance at [10 20] Hz. **(G**) Median phase gain for spike-LFP pairs (with significant encoding) that showed significant synchrony in the 10-25 Hz beta band (n=877), vs. those that did not (n=2500). Non-locking cell showed significantly less phase gain compared to locking cells (Wilcoxon ranksum test, $p<0.05$).

We estimated the strength of phase modulation of rate encoding for each spike-LFP pair as the amplitude of a cosine fit to the phase-binned encoding metric and normalized it by the cosine amplitude obtained from fitting the phase-binned metric after randomly shuffling spike phases. We refer to this difference of the observed to the randomly shuffled phase modulation of encoding as the *Encoding Phase Gain* (*EPG*). Of the 877 spike-LFP pairs that significantly synchronized in the 10-25 Hz and encoded information in their firing rate, we found that 139 (16%) spike-LFP pairs showed significant phase-of-firing encoding, i.e. these pairs encoded significantly more information in the phase of firing than in their firing rate alone (randomization test, $p<0.05$). A significant *EPG* was evident for neurons whose average firing encoded *Outcome* (Wilcoxon signrank test, p~0), *Prediction Error*, ($p=0.025$), and *Outcome History* ($p=0.016$). The *EPG* did not differ between these three functional clusters (Kruskal Wallis test, $p=0.86$) (**Figure 2-4D**). Similarly, *EPG* was evident for spike-LFP pairs with the spiking neuron in ACC (p~0), in STR ($p=0.00028$), and in LPFC ($p=0.015$) (**Figure 2-4E**). EPG differed between areas (Kruskal-Wallis test, $p=0.016$), with ACC showing significantly larger *EPG* than LPFC ($p=0.004$), and a similar trend when compared to STR ($p=0.086$). *EPG* differences were more pronounced when selecting for

each encoding metric the 25% of spike-LFP pairs with the largest *EPG*. This selection revealed stronger EPG encoding of *Prediction Error* compared to *Outcome* (p~0) and *Outcome History* (p=~0). Large *EPG* were found significantly more likely in ACC than in LPFC (p=0.00066) or STR (p=0.032). These results were similar in each monkey (**Figure A-3C**). Encoding designation was largely stable across phase bins, with ~90% of spike-LFP pairs exhibiting similar beta coefficient signs across all phase bins. *EPG* did also not depend on the number of bins used (Spearman rank correlation, R=0.023, p=0.18). Likewise, *EPG* was similar for neurons that showed narrow (N) and broad (B) action potential waveforms that corresponds to putative distinct cell types with their encoding phase gain statistically indistinguishable in the ACC (N_N=70, mean=-0.026 ± 0.029; N_B=48, mean = -0.0047±0.06; Kruskal Wallis test for equal median, p=0.49), LPFC (N_N=85, mean=0.11 ± 0.04; N_B=54, mean = 0.0057±0.080; p=0.28), or STR (N_N=37, mean=0.014 ± 0.08; N_B=41, mean = 0.017 ± 0.11; p=0.40) (see **Methods**).

2.5.4 Phase-of-Firing Encoding Depends on Significant Synchronization in the Beta Band

We next tested how *EPG* related to the strength of synchronization. First, we found that *EPG* was significant at the population level and strongest in the same beta frequency band that showed the strongest spike-LFP synchronization (**Figure 2-4F**; Wilcoxon signrank test, p<0.05). Overall, *EPG* was most prevalent and significantly larger in spike-LFP pairs showing significant phase synchronization (**Figure 2-4G**; Kruskal-Wallis test, p~0). These results indicate that *EPG* was evident when neurons encoded *Outcome*, *Prediction Error* and *Outcome History* in their firing rate *and* when they synchronized at beta band frequencies.

We next asked whether the synchronizing phase that carried information was endogenously generated or whether it was externally triggered by the reward onset. We calculated the *EPG* with and without subtracting the reward-onset aligned evoked LFP response (*see* **Methods**). We found that the *EPG* was not different with (median=0.0704 ± 0.012 SE) versus without (median=0.067 ± 0.016 SE) subtraction of the time-locked, evoked potential, suggesting that the beta oscillation events providing informative phases were endogenously generated (Kruskal-Wallis test, p=0.90). We also tested whether LFP power variations influenced the phase-

of-firing encoding but found that the *EPG* did not correlate with beta band power variations (Spearman rank correlation, R=0.051, p=0.13). We did find, however, that *EPG* was positively correlated with the overall firing rate (Spearman rank correlation, R=0.14, p~0), though this did not vary by isolation quality (Kruskal Wallis test, p=0.15).

2.5.5 Preferred Spike Phase and Encoding Phase Differ for Prediction Error

The phase-of-firing encoding so far might be due to stronger spike-LFP synchronization in one task condition than in another condition (e.g. in error trials versus correct trials). Such site-specific selectivity of neuronal synchronization has been reported in previous studies (e.g. (Womelsdorf et al., 2010b; Salazar et al., 2012; Antzoulatos and Miller, 2016). To test this possibility, we compared spike-LFP synchronization (indexed with the PPC), in those two trial conditions that were predicted to have the maximal firing rate difference. For *Outcome* encoding we compared correct vs error trials; for *Prediction Error* encoding we compared correct trials following error *versus* error trials following correct trials; and for *Outcome History* encoding we compared correct trials following correct, or errors following errors. We then correlated the absolute difference in PPC in the beta band between two conditions with the phase-of-firing encoding. We found statistically indistinguishable levels of PPC between conditions for neurons encoding *Outcome* (Spearman rank correlation, R=0.038, p=0.35), *Prediction Error* (R=0.050, p=0.58), or *Outcome History* (R=-0.027, p=0.76).

This result suggests that the strength of phase-of-firing encoding does not follow simply from differentially strong spike-LFP phase synchronization. Rather, the spike-phase at which neurons maximally synchronize might not always coincide with the spike-phase at which encoding of task variables is maximal. Indeed, we often observed that the phase with maximal encoding was not at the zero-phase bin, i.e. it deviated from the preferred spike-phase (**Figure 2-4C; Figure A-4).** We tested this scenario by first calculating the preferred spike-phase for each neuron, and then quantifying the phase with maximal encoding relative to that phase. We found that all encoding neurons synchronized on average at similar phases, above what would be expected by chance (Hodges-Ajne test; *Outcome*, average phase: -0.28 ± 0.0034 SE radians, p~0;

RPE. average phase: 0.35 ± 0.0034 SE radians, p=0.00084; *Outcome History*. average phase: -0.68 ± 0.0045 SE radians, p=0.0013) (**Figure 2-5A**). The preferred spike-phase differed between the three encoding classes (Watson Williams test, p~0; each pairwise comparison showed: Watson-Williams test, p<0.02; **Figure 2-5B**).

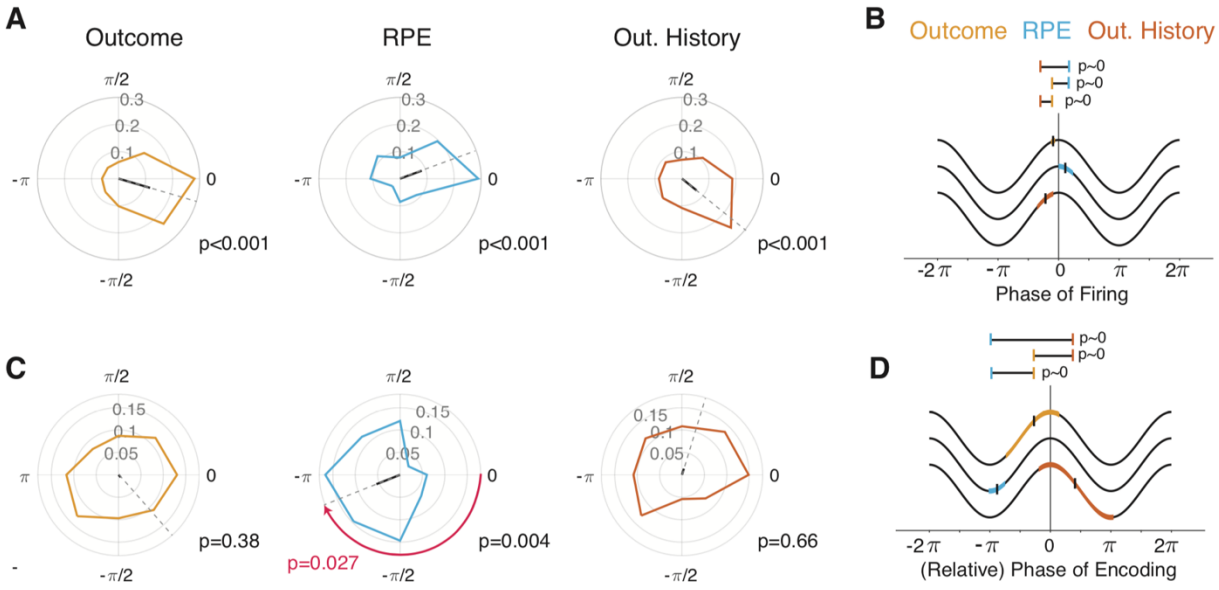

Figure 2-5.Preferred firing and encoding phase are dissociated in RPE cells

(A) Proportion of cells with a preferred phase of firing for each encoding cluster. Black lines depict the average phase. *Outcome (yellow), RPE (blue)* and *Outcome History* (*red*) cells showed strong evidence of phase concentration (Hodjes-Ajne test, p<0.05). **(B)** Mean preferred phase of firing for *Outcome*, *RPE*, and *Outcome History* clusters. Colored patches represent the 95% confidence interval about the mean, depicted as a black horizontal line. The distribution of preferred firing phases differed between all encoding clusters (Watson-Williams test, p<0.05). **(C)** Relative encoding phase (where 0 corresponds to the preferred phase). *Outcome* and *Outcome History* cells did not show an encoding phase preference, whereas RPE cells did (Hodges-Ajne test, p=0.0004). Importantly, for these RPE cells, this phase differed significantly from 0 (the preferred phase of firing; Median test, p=0.027) **(D)** Similar conventions as in (B) but for the encoding phase. The distribution of encoding phases differed between all three clusters.

Next we quantified for each cluster whether the phases showing maximal encoding were consistent across spike-LFP pairs (**Figure 2-5C**). To this end, we extracted the phase offset from our cosine fit, which represents the phase at which encoding was maximal relative to the preferred spike-phase. *Outcome* encoding neurons showed a preferred encoding phase that randomly varied across neurons (Hodges-Ajne test, p=0.38), as did *Outcome History* neurons

(p=0.66) (**Figure 2-5C**). In contrast, *Prediction Error* encoding neurons significantly encoded at similar phase-offsets relative to the neuron's synchronizing phases (Hodges-Ajne test, p=0.0004 at -2.76 ± 0.047 SE radians, corresponding to 27 ms at a 15 Hz oscillation cycle), which was significantly different than the mean spike phase (Median test, p=0.027). This effect was particularly pronounced for *Prediction Error* cells in ACC **(Figure A-5)**, and was consistent across both monkeys (**Figure A-3D)**. The phases showing maximal encoding differed between all three functional clusters (Watson-Williams test; p<0.001; **Figure 2-5D**). These result show that the preferred spike phase and the encoding phases are dissociated from one another for all types of information. Additionally, for *Prediction Error* information there was a systematic offset of maximal encoding at ~27 ms away from the preferred (mean) spike phase. There was no relation between the phase of maximal encoding and the number of bins used (Circular-Linear correlation, R=0.013, p=0.76).

We next validated that the dissociation of spike- and encoding- phases was not influenced by phase shifts due to differences in the peak oscillation frequencies within the beta band. The three sets of neuronal encoding clusters synchronized on average at the same center frequency of ~15 Hz (Kruskal Wallis, p=0.34), and showed maximal phase-of-firing encoding at similar frequencies (also ~15Hz) (Kruskal Wallis test, p=0.23). Moreover, the frequency showing strongest spike-LFP synchronization and the frequency showing maximal encoding phase gain matched closely (median frequency ratio: 1 ± 0.019 SE). This similarity of synchronization and encoding frequency did not differ on the basis of the functional designation (Kruskal-Wallis test, p=0.17), nor the area from which the spikes were sampled (Kruskal Wallis test, p=0.55).

2.6. Discussion

2.6.1 Summary

Here, we found a substantial proportion of neurons whose phase-of-firing in the beta frequency band conveyed significantly more information about three learning variables than their firing rate alone. This encoding phase gain was evident for spikes generated within the ACC, LPFC and STR of nonhuman primates in a [0.1 0.7] second period of outcome processing during

reversal learning performance. Phase-of-firing encoding was most prominent at the 10-25 Hz beta frequency at which spikes synchronized to the local fields across areas. However, the strength of spike-LFP phase synchronization could not explain the strength of the phase-of-firing encoding. Rather, maximal encoding occurred for many neurons at phases away from the preferred spiking phase. The dissociation of spiking- and encoding phases was particularly prominent for information about the *Prediction Error*.

Taken together, these results provide a first report of information multiplexing of learning variables at segregated phases of a beta oscillation that is synchronized across medial and lateral fronto-striatal loops. These findings suggest that oscillation phases are important carriers of information, above and beyond that of a phase-blind firing code. The gain of information through phase-specific firing provides an intriguing dynamic code that could link principles of efficient neuronal information transmission with the demands of representing multiple types of information in the same dynamical neural system.

2.6.2 Distributed Encoding of Learning Variables at a Shared Beta Rhythm Frequency

We found that three critical variables needed for adjusting behavior are represented in segregated neuronal populations not only in their firing rates, but in phase specific firing at a beta frequency that is shared among ACC, LPFC and STR. This finding suggests that the beta frequency could serve as a primary carrier for the fast distribution of learning related information within fronto-striatal networks (Womelsdorf and Everling, 2015; Spitzer and Haegens, 2017). Prior studies have shown that ACC, LPFC and STR causally contribute to fast learning of object values. With lesioned ACC, rhesus monkeys fail to use outcome history for updating values and show perseverative behaviors (Kennerley et al., 2006). Without LPFC, rhesus monkeys fail to recognize when a previously irrelevant object becomes relevant, as if they fail to calculate prediction errors needed for updating their attentional set (Rudebeck et al., 2017). When the anterior STR is lesioned, nonhuman primates tend to stick to previously learned behavior and show a lack of sensitivity to reward outcomes (Hikosaka et al., 2017, 2019). These behavioral lesion effects are consistent with the important role of each of these brain areas to track the history of recent

outcomes, registering newly encountered (current) outcomes, and calculating the unexpectedness of experienced outcomes (prediction error). Consequently, our finding of segregated neuronal ensembles encoding *Outcome*, *Prediction Error* and *Outcome History* complement a large literature that documents how these variables are represented in the firing of neurons across fronto-striatal areas.

What has been left unanswered, however, is how this firing rate information about multiple variables emerges at similar times and similar proportions across areas. Prior studies suggest that firing rate correlations between brain areas are relatively weak and poor candidates for veridical information transfer (Kumar et al., 2010; Hahn et al., 2019; Oemisch et al., 2019), while temporally aligning the spike output of many neurons to the phases of precisely timed, synchronized packets are a theoretically, particularly powerful means in affecting postsynaptic neuronal populations (Azouz and Gray, 2003; Fries, 2015; Luczak et al., 2015; Voloh and Womelsdorf, 2016; Hahn et al., 2019). Our findings support this notion of a temporal code using synchronized oscillations by showing that those neurons that carry critical information in their firing rates also tend to synchronize long-range between ACC, LPFC, and STR at a shared 10-25 Hz beta frequency. This beta frequency is thus a powerful candidate for distributed information transfer, because spike output of many neurons is concentrated at the same phase and thus activate postsynaptic membranes at similar times. This scenario of beta rhythmic information exchange within fronto-striatal networks is supported by previous nonhuman primate studies that demonstrated 10-25 Hz beta rhythmic synchronization during active task processing states between ACC and LFPC (Womelsdorf et al., 2014a; Voloh and Womelsdorf, 2017; Smith et al., 2019), between PFC and STR (Antzoulatos and Miller, 2014), between ACC and FEF (Babapoor-Farrokhran et al., 2017), between LPFC and FEF (Salazar et al., 2012; Antzoulatos and Miller, 2016), and between lateral PFC or FEF with posterior parietal cortex (Buschman and Miller, 2007; Buschman et al., 2012a; Dean et al., 2012; Salazar et al., 2012; Antzoulatos and Miller, 2016). Each of these studies have shown short-lived rhythmic long-range synchronization between distant brain areas during cognitive tasks at a similar ~15 Hz frequency. Our findings critically complement these studies by revealing that 10-25 Hz spike-LFP synchronization is prevalent not

only during cognitive processing, but also during the processing of outcomes after attention has been deployed and choices have been made. During this post-choice outcome processing, fronto-striatal circuits are likely to adjust their synaptic connection strength to minimize future prediction errors and improve performance (Leong et al., 2017; Hikosaka et al., 2019; Oemisch et al., 2019). Our results suggest that this updating utilizes beta rhythmic activity fluctuations during the post-choice outcome processing period.

2.6.3 Multiplexing of Information Through Phase-of-Firing Encoding

Our finding that spiking output carries separable types of information at different phases of the same oscillation frequency has potentially far-reaching implications. By finding that *Outcome, Prediction Error and Outcome History* were encoded at separate phases, the neuronal spiking activity effectively multiplexes independent information streams at different phases of beta synchronized firing. This stands in contrast to prior studies reporting that long-range beta rhythmic synchronization between LPFC, ACC or STR in the primate encoded relevant task variables via the strength of beta synchrony (Buschman and Miller, 2007; Buschman et al., 2012a; Dean et al., 2012; Salazar et al., 2012; Womelsdorf et al., 2014a; Antzoulatos and Miller, 2016; Babapoor-Farrokhran et al., 2017). For example, some prefrontal cortex neurons synchronize stronger at beta to posterior parietal areas when subjects choose one visual category over another (Antzoulatos and Miller, 2016), or when they maintain one object over another in working memory (Salazar et al., 2012). These findings are consistent with a communication-through-coherence schema where upstream senders are more coherent with downstream readers when they successfully compete for representation (Womelsdorf et al., 2007; Bosman et al., 2012; Fries, 2015). Yet it has remained unclear how such a scheme may operate when multiple items must be multiplexed and transmitted in the same recurrent network (Siegel et al., 2009; Kopell et al., 2011; Akam and Kullmann, 2014; Palminteri et al., 2015; Khamechian et al., 2019). Computationally, the multiplexing and the efficient transmission of information can operate in tandem when the temporal organization of activity is exploited at the sending and receiving site (Buzsáki, 2010; Kumar et al., 2010; Luczak et al., 2015; Hahn et al., 2019). Consequently, selective synchronization between distal sites could be leveraged to enhance

transmission *selectivity,* whereas temporally segregated information streams could enhance transmission *capacity* (McLelland and VanRullen, 2016). Our results resonate with this view by showing that neurons that synchronize long-range at one oscillation phase carries information of any of three learning variables at phases systematically offset from the synchronizing phase.

By finding evidence for such a temporal multiplexing in the beta frequency band we critically extend previous reports of phase encoding of information for object features, object identities, and object categories at theta, alpha and gamma frequencies (Siegel et al., 2009; Kayser et al., 2012; Turesson et al., 2012; Womelsdorf et al., 2012). In our study, the beta rhythmic phase-of-firing multiplexing applied to complex learning variables that were needed to succeed in the behavioral learning task. In particular, the presence of reward prediction error information provides a critical teaching signal that indicates how much synaptic connections should change to represent future value expectations more accurately (Leong et al., 2017; Oemisch et al., 2019). Our results suggest that this updating can utilize spike-timing dependent plasticity mechanisms that are tuned to firing phases ~27 ms away from the preferred synchronization phase in the beta frequency band. How such a temporal organization in the beta band is used in the larger fronto-striatal network will be an important question for future studies.

So far, evidence in humans and rodents suggested that processes linked to beta frequency activity during the evaluation of outcomes support the detection of errors and the updating of erroneous internal predictions (Howe et al., 2011; Tan et al., 2014b). In fact, there have been conflicting views on whether beta oscillations related to outcome signals are more likely to reflect a weighted integration of recent outcomes, or the unexpectedness of the current, observed outcome relative to recent outcomes (Howe et al., 2011; Tan et al., 2014a). Our findings reconcile these viewpoints by documenting that encoding of *Outcome History* weights and of *Prediction Errors* coexist in the same circuit at the same oscillation frequency in phase dependent firing of single neurons.

2.6.4 Neuronal Mechanisms underlying Phase-of-Firing Multiplexing

We found that the beta phase allowing maximal encoding of *Prediction Errors* was offset ~27 ms on average from the phase at which most spikes synchronized to the local fields. Such a dissociation of spike-phase and encoding-phase has been reported previously for the beta frequency band in parietal cortex, where maximal information of joint saccadic and joystick choice directions were best predicted by spike counts at ~50 degree away from the preferred beta spike phase (Hawellek et al., 2016). Such phase offsets underlying maximal encoding in parietal cortex as well as in ACC, LFC and STR in our study provide constraints on the possible circuit mechanisms that permit temporal segregation of inputs streams through phase specific oscillatory dynamics (Akam and Kullmann, 2014). One possible circuit mechanism that implements and utilizes multiplexed information streams through phase specific firing has been described and computationally modelled specifically for the low 10-20 Hz frequency range (Kopell et al., 2011; Gelastopoulos et al., 2019). This work suggests that distinct sets of pyramidal neurons can encode distinct input streams in their firing phases at 10-20 Hz beta phases when these inputs streams arrive with a phase offset to each other, e.g. when they arrive sequentially in time. According to this schema, a first input stream activates pyramidal neurons in deep cortical layers that feed information to superficial layers whose interlaminar inhibitory connections closes an interlaminar reverberant loop of activity. This interlaminar ensemble follows a beta activity rhythm due to cell specific dynamics that maintains the beta-phasic firing of active neurons (Kopell et al., 2011; Womelsdorf et al., 2014b). When a second input stream activates another set of pyramidal cells within the same beta rhythmic neural population, the input timing of that second stream was maintained at a different phase than the phase of the first activated ensemble (Gelastopoulos et al., 2019). The parallel coding of information at a common beta rhythm in these models provide a qualitative proof of concept about phase specific encoding of multiple types of inputs in larger beta rhythmic ensembles, and suggests their mechanistic realization (Gelastopoulos et al., 2019). Based on our results we predict that the set of neurons encoding current *Outcomes*, *Prediction Errors* and *Outcome History* will emerge earliest in deep ACC and LPFC layers that are associated with prominent beta rhythms (Bastos et al., 2018) and then transmit to superficial layer neurons to form an interlaminar beta rhythmic

ensemble multiplexing different types of information at the same shared beta frequency oscillation. Future work needs to specify how such multiplexing ensembles can transform inputs to generate novel representations.

In summary, we have documented that learning variables are encoded at separable phases of firing of neurons that synchronize long–range across primate fronto-striatal circuits. These phase encoding neurons also carried information in overall firing rate modulations which clarifies that an *asynchronous* rate code and a *synchronous* temporal code coexist in the same circuit (Kumar et al., 2010). By exploiting the temporal structure endowed in long–range neuronal synchronization our findings suggest how neuronal assemblies in one brain area could be read out from neural assemblies in distally connected brain areas (Singer and Gray, 1995). This phase-of-firing schema entails key features required from a versatile neural code including the efficient neural transmission and the effective representation of variables needed for adaptive goal directed behavior (Perkel and Bullock, 1968).

2.7. Acknowledgments

This research was supported by a grant from the Canadian Institutes of Health Research (T.W.) CIHR Grant MOP_102482 and by the National Institute Of Biomedical Imaging And Bioengineering of the National Institutes of Health under Award Number R01EB028161 (T.W.). The content is solely the responsibility of the authors and does not necessarily represent the official views of the Canadian Institutes of Health Research or the National Institutes of Health. We thank Dr. Andrew Tomarken for very useful methodological discussions.

2.8. Statement of Contributions

BV: Benjamin Voloh, MO: Mariann Oemisch; TW, Thilo Womelsdorf

Conceptualization, B.V. and T.W.; Methodology, B.V. and T.W.; Software, B.V. and T.W; Investigation, M.O. and T.W.; Writing – Original Draft, B.V. and T.W..; Writing – Review & Editing, all authors; Visualization, B.V. and T.W..; Supervision, T.W.

Chapter 3

Surprise Signaled by Spiking on Unexpected Phases of Beta Oscillations in Anterior Cingulate and Prefrontal Cortex

3.1. Abstract

Information about outcome unexpectedness, or surprise, is distributed across frontal circuits, and guides feature-value updating for the purpose of flexible learning and behavioral adjustment. A rich body of literature delineates how changes in average firing rates track the degree of surprise. During outcome processing, spikes synchronize to population level beta-rhythmic activity, suggesting that surprise may also be signaled by deviations from expected synchronizing phases. We tested this hypothesis in a network comprised of the anterior cingulate and prefrontal cortices of nonhuman primates during reversal learning of feature values. We observed that the phase of a beta frequency oscillation at which neurons were active was related to the degree of surprise. Importantly, learning speed across behavioral sessions was highest when neurons signaled surprise on anti-preferred phases, and when surprise was modulated more strongly as a function of phase. This effect was particularly pronounced for putative excitatory cells, rather than inhibitory cells, specifically during faster learning sessions. Taken together, our results provide a first account suggesting that spiking at unexpected phases signals unexpected outcomes in order to enhance learning performance.

3.2. Introduction

The anterior cingulate cortex (ACC) and lateral prefrontal cortex (LPFC) are key structures that are involved in learning feature values and flexibly guiding behavior to optimize reward. To achieve such flexibility, feature values can be updated during learning by comparing expected and observed outcomes (Balcarras et al., 2016; Leong et al., 2017). The discrepancy between these is a reward prediction error. Previous work has shown that changes in the firing rate of neurons can be used to signal such deviations (Schultz et al., 1997; Oemisch et al., 2019).

The implicit assumption of this work is that a downstream reader is sensitive to changes in the distributional properties (i.e. mean) of neuronal firing rates, constituting a firing rate code. Alternatively, information can be conveyed not only by firing rates, but also by the phase of an underlying population rhythm at which spikes occur (a phase-of-firing code) (Kayser et al., 2009; Womelsdorf et al., 2012; Onorato et al., 2019). In addition, we have recently observed that neurons that carry prediction error information tend to encode such information on the opposite phase at which they tend to fire (Voloh et al., 2019a). This suggests a novel form of prediction error encoding, namely, one based on the phase-of-firing. In this proposed framework, a downstream reader is sensitive to the *circular* distribution of firing phases; thus, deviations from expectations can be signaled by firing on the unexpected (anti-preferred) phase.

If anti-phase firing does indeed signal reward prediction errors, then they should have a concomitant role in learning. Indeed, band-limited activity, particularly in the beta (10-30 Hz) and gamma (30-60 Hz) range, are key rhythms linked to the feedforward and feedback flow of information (Bastos et al., 2014; Kerkoerle et al., 2014; Womelsdorf and Everling, 2015). Such rhythms are thus prime candidates for carrying prediction error information (Bastos et al., 2012). Indeed, a recent study has shown that stimulation at pro- or anti-phases of endogenous beta can lead to long term potentiation or depression, respectively, at synapses (Zanos et al., 2018), in line with earlier studies showing similar effects (Stanton and Sejnowski, 1989; Hyman et al., 2003; McCartney et al., 2004; Wespatat et al., 2004). Thus, aligning spike firing to specific phase of endogenous oscillations may subserve the reorganization of synaptic weights to promote learning, weakening those that lead to poor outcomes and strengthening those that lead to good outcomes.

To this end, we analyzed recordings in the LPFC and ACC while macaques performed a feature-based reversal learning task (Oemisch et al., 2019). We found that during outcome processing, the degree of unexpectedness – estimated using a best-fitting reinforcement learning model that weights feature values (Oemisch et al., 2019) – was associated with spiking on specific phases of a beta rhythm. The specificity and sensitivity of spike-phases signaling outcome

unexpectedness was associated with the speed of learning, and was particularly prominent for putative excitatory cells. These preliminary findings suggest that outcome unexpectedness can be signaled by spikes occurring on unexpected firing phases.

3.3. Methods

Analysis was performed in Matlab 2019a using custom software.

3.3.1 Task and Recordings

Task and behavior are as described in **Chapter 2.** This was a feature-based reversal learning task (Oemisch et al., 2019). Here, we focus analysis on the ACC (area 24) and LPFC (area 46, 8, 8a). Correct outcomes were those where the monkey reported the correct motion of the target stimulus. Error responses were made to the incorrect stimulus, or in the incorrect response window.

3.3.2 Expectation Maximization Algorithm

To identify which trial in a block monkeys showed statistically reliable learning, we modelled the trial-by-trial outcomes using the Expectation Maximization Algorithm (Smith et al., 2004, 2005, 2007). This framework describes the learning process as a hidden Markov process and is updated on each trial. The algorithm assumes the perspective of an ideal observer that takes into account all trial outcomes in a block in order to estimate the probability of a correct outcome on individual trials, and the attendant confidence intervals. We defined a learning trial within a block as the earliest trial where the lower confidence bound exceeded chance performance (p=0.5). We defined the learning speed *LS* as the average learning trial across blocks within a session.

3.3.3 Behavioral Modelling

We modelled subjects choice behavior using a feature-selective reinforcement learning algorithm with decay, which best described subjects behavior when compared to other

reinforcement learning algorithms or Bayesian models (Oemisch et al., 2019). The full details are described in (Oemisch et al., 2019), but we briefly describe here the most salient points.

This model assigns values to features that define each stimulus, which include the location (left, right), motion direction (up, down), and color (1, 2). Each of these features was assigned an index *i* and has an associated value V_i. After receiving an outcome R (0 for unrewarded trials, 1 for reward), values were updated according to:

$$V_{i,t+1} = V_{i,t} + \eta(R_t - V_{i,t}) \qquad Eq.1$$

For all three features *i* that corresponded to the *chosen* stimulus. The term $(R_t - V_{i,t})$ is the reward prediction error, scaled by the learning rate η. Feature values associated with the *unchosen* stimulus were updated according to the equation:

$$V_{i,t+1} = (1 - \omega)V_{i,t} \qquad Eq.2$$

Where ω is a decay parameter. Thus, if learning occurs, the value of the feature V_i will converge to the true value after many iterations.

3.3.4 Spectral decomposition and spike-LFP phase synchronization

We focused our analysis of spike-phase synchrony between pairs of distally recorded sites, thus obviating any concerns of spike energy bleeding into the LFP (Zanos et al., 2011). We considered frequencies from 6 Hz to 64 Hz. For frequencies from 6-30Hz, the resolution was 1 Hz, and above that it was 2Hz. For every frequency F, we determined the spike-LFP phase by extracting an LFP segment centered on the spike of length 5/F (i.e. 5 cycles). Spectral decomposition was done via an FFT after applying a Hanning taper. This procedure was applied to the post-outcome period [0.1 1] sec.

The strength of spike-LFP synchronization was quantified using the pairwise phase consistency (PPC), which is unbiased by the number of spikes (Vinck et al., 2011). The PPC is

quantified on the basis of pairwise differences between spike-phases. If spikes tend to fire on specific phases, phase difference will be concentrated, and thus the PPC will take on a high value, whereas if spikes are distributed randomly relative to the LFP phase, phase differences will be random and the PPC will tend towards zero.

We determined the frequency at which spike-LFP phase synchronization was significant by determining peaks in the PPC spectrum. A cell was said to synchronize to a particular frequency if the following criteria were met: Peaks had to be (1) above a threshold of 0.005, (2) show a minimum prominence of 0.005, and (3) show significant Rayleigh test (i.e. phase concentration).

3.3.5 Relating learning speed with the phase of firing

To reiterate, the fundamental hypothesis underlying this study is that surprise may be signaled on anti-preferred firing phases (*see* **Figure 3-2**). If this is the case, then we may also expect that deviations from this would be detrimental to learning. To this end, we first modelled the relationship of surprise to spike phases, and then related these model parameters to learning speed.

As schematized in **Figure 3-2A**, spikes locking to distal LFP phases are defined by two characteristics. First, surprise may be signaled if the *degree* to which spikes lock to specific phases changes. Second, surprise may be signaled by spikes that occur away from the preferred firing phase. The preferred firing phase can be conceptualized as the phase at which a distal reader expects spikes to arrive. To this end, for each trial that was part of a learned block as defined by the EM algorithm (*see* above), we extracted (1) the surprise *S* (absolute RPE as derived from the RL model), and (2) the mean phase of spiking (thus obviating any concerns of firing rate differences affecting the results). We subtracted the mean firing phase of the cell (for each frequency) to extract the relative firing phase. We then used a linear-circular model to predict *S* as a function of the modal (relative) firing phase (Fisher, 1993; Shumway and Stoffer, 2017):

$$S = A * cos(x - \theta) \qquad Eq. 3$$

Where *A* represents the amplitude, *x* is the modal phase, and θ represents the phase offset (i.e. the phase at which *S* is maximal). Using a trigonometric identity, this can be written as:

$$S = A * \cos(\theta)\, cos(x) - A * sin(\theta)\, sin(x) \qquad Eq.4$$
$$S = \beta_1 cos(x) + \beta_2 sin(x)$$

Where:

$$\beta_1 = A * cos(\theta) \qquad Eq.5$$
$$\beta_2 = -A * sin(\theta) \qquad Eq.6$$

In its linear form, we can now apply standard regression techniques to estimate β_1 and β_2. Using these equations, we can solve for θ and *A*:

$$\theta = atan\left(\frac{\beta_2}{\beta_1}\right) \qquad Eq.7$$
$$A = sqrt(\beta_1^2 + \beta_2^2) \qquad Eq.8$$

The term A represents the spike-surprise *sensitivity*, whereas θ represents the spike-surprise *specificity*. We analyzed spike-LFP pairs where we could estimate the modal phase for at least 20 (learned) trials.

With these two parameters in hand, we then determined how spike-surprise was related to learning speed. To this end, we used a linear mixed effects model, with individual estimated parameters *A* and θ as fixed effects and subject as random effect:

$$LS = \beta_A A + \beta_{cos} cos(\theta) + \beta_{sin} sin(\theta) + (1|subject) \qquad Eq.9$$

This model thus describes how learning speed varies as a function of the spike-surprise sensitivity *A*, and specificity θ. To specify and clarify the influence of spike-surprise specificity on learning speed, we used equation 5-8 to extract (1) the degree of influence, $A_{learning}$, and (2) the phase at which this effect was maximal, $\theta_{learning}$. To assess significance, we extracted the minimum p-value associated with either the β_{cos} or β_{sin} terms.

Finally, to determine the frequency at which spike-surprise sensitivity influenced learning speed, we performed this analysis for all frequencies in the 6-64 Hz range.

3.3.6 Testing for uniqueness of spike-surprise

We further tested spike-surprise sensitivity and specificity in two important ways. First, we fit modal spike phases on each trial with the value of the chosen stimulus (instead of the surprise), and extracted parameters *A* and θ, representing spike-*value* sensitivity and specificity, respectively (Eqn. 5-8), before relating it to learning speed (see Eqn. 9). We then compared the efficacy of spike-value or spike-surprise sensitivity to predicting the session-wide learning speed with a (reduced) linear mixed effects model, using sensitivity as the fixed effect and monkey as the random effect. We then extracted the Akaike Information Criterion (AIC) for each model, and compared the difference to determine if surprise or value better predicted learning speed. In a separate analysis, we repeated this procedure but comparing spike-surprise and spike-value specificity.

In a second analysis, we tested if spike-surprise sensitivity/specificity were unique in time. To this end, we repeated the analysis in a [-1 0] sec. pre-outcome period, and extracted A_{pre} and θ_{pre}. We then asked if specificity or sensitivity changed between the pre- or post-outcome period. We describe the procedure for specificity first. Because phases may shift with frequency, we could not simply average across frequencies. To this end, we first took the pairwise difference in each frequency in a [10 25] Hz beta band. Then, for each frequency, we conducted a Hodges-Ajne omnibus test, and extracted a test statistic Z which was the average across frequencies. To assess significance, we performed a permutation test, by first permuting the label “pre” or “post”, and then repeating the analysis 1000 times. The significance of the test statistic Z was assessed in reference to this null distribution. For the sake of coherence, we performed the same analysis on spike-surprise sensitivity, but using a Wilcoxon signrank test instead.

3.3.7 Cell type classification

To determine if phase-modulated encoding of information differed based on cell type, we focused the following analysis on highly isolated single units that showed encoding of learning-relevant variables and significant, prominent spike-beta locking. Detailed information is provided in (Oemisch et al., 2019). Briefly, to distinguish putative interneurons (narrow-spiking) and putative pyramidal cells (broad-spiking) in LPFC and ACC, we analyzed the peak-to-trough duration and the time to repolarization for each neuron. After applying Principal Component Analysis (PCA) using both measures, we used the first principal component to discriminate between narrow and broad-spiking cells. This allowed for better discrimination than using either measure alone. We confirmed that a two-Gaussian model fit the data better than a one-Gaussian model using the Akaike and Bayesian Information Criterion (AIC, BIC). We then used the two-Gaussian model to define narrow and broad-spiking populations.

A similar analysis was applied to striatal units to distinguish putative interneurons from medium spiny projection neurons (MSN). Here, we use the peak-width and Initial Slope of Valley Decay (ISVD) (Berke, 2008; Lansink et al., 2010):

$$ISVD = \frac{V_t - V_{0.26}}{A_{PT}}$$

where V_T is the most negative value (trough) of the spike waveform, $V_{0.26}$ is the voltage at 0.26 ms after V_T, and A_{PT} is the peak-to-trough amplitude (Lansink et al., 2010). After PCA and two-Gaussian modelling (as described above), we defined two cut-off points. The first cutoff was the point at which the likelihood of narrow spiking cells was 3 times larger than the likelihood of broad-spiking cells, and vice-versa for the second cutoff.

We compared differences in *spike-surprise sensitivity* between narrow and broad spiking neurons using the Kruskal-Wallis test, independently for each area, while differences in spike-surprise specificity were assessed with a Watson-Williams test (Zar, 2010). To clarify, we analyzed spike-LFP pairs here; thus, the same neuron may be included more than once.

3.4. Results

3.4.1 Behavioral

Animals performed a feature-based reversal learning task **(Figure 3-1A;** (Oemisch et al., 2019)). On each trial, they were shown two stimuli, each of which were composed of three features, namely the color, location, and motion direction (**Figure 3-1B).** Animals were rewarded if they responded to the motion associated with a rewarded color. Color reward associations were reversed without cue (**Figure 3-1C).** Average subject performance aligned to block reversals indicated that they were able to re-learn color-reward contingencies (Figure 3-1E). Using ideal observer statistics, we determined that animals were able to learn the new reward association within 15.2 ± 0.97 / 18.2 ± 1.17 trials on average in each session (**Figure 3-1F;** monkey H/K).

To determine the effect of surprising outcomes on choice behavior, we modelled animal choice behavior using a reinforcement learning algorithm that updated feature values of chosen stimuli and decayed values associated with unchosen stimuli **(Figure 3-1G).** This model was chosen as it best described the animals' choices (Oemisch et al., 2019). As expected, we found that surprise (i.e. the absolute reward prediction error) was highest immediately following reversals, and decayed as the animals learned which color was rewarded **(Figure 3-1H).** This corroborates previous reports that information related to feature dimension as well as individual feature values contribute to performance (Niv et al., 2015; Balcarras et al., 2016; Leong et al., 2017; Oemisch et al., 2019).

Next, we were interested in discerning how outcome unexpectedness is related to the variability in session-wide learning performance (**Figure 3-1F).** To this end, we split the data into "fast" and "slow" sessions, defined as the top or bottom 50th percentile of learning performance, to determine the evolution of surprise during learning **(Figure 3-1I).** We observed that on slow sessions, surprise decayed more slowly as compared to fast sessions. Intriguingly, we also observed that surprise was higher on the first trial following the reversal in fast sessions (0.78 ± 0.014 SE), as compared to slow sessions (0.74 ± 0.16 SE; T-test, p=0.042; **Figure 3-1I inset**). This suggests that on slower sessions, animals were less attuned to surprising outcomes, particularly following reversals, and were thus slower to re-learn color-reward associations.

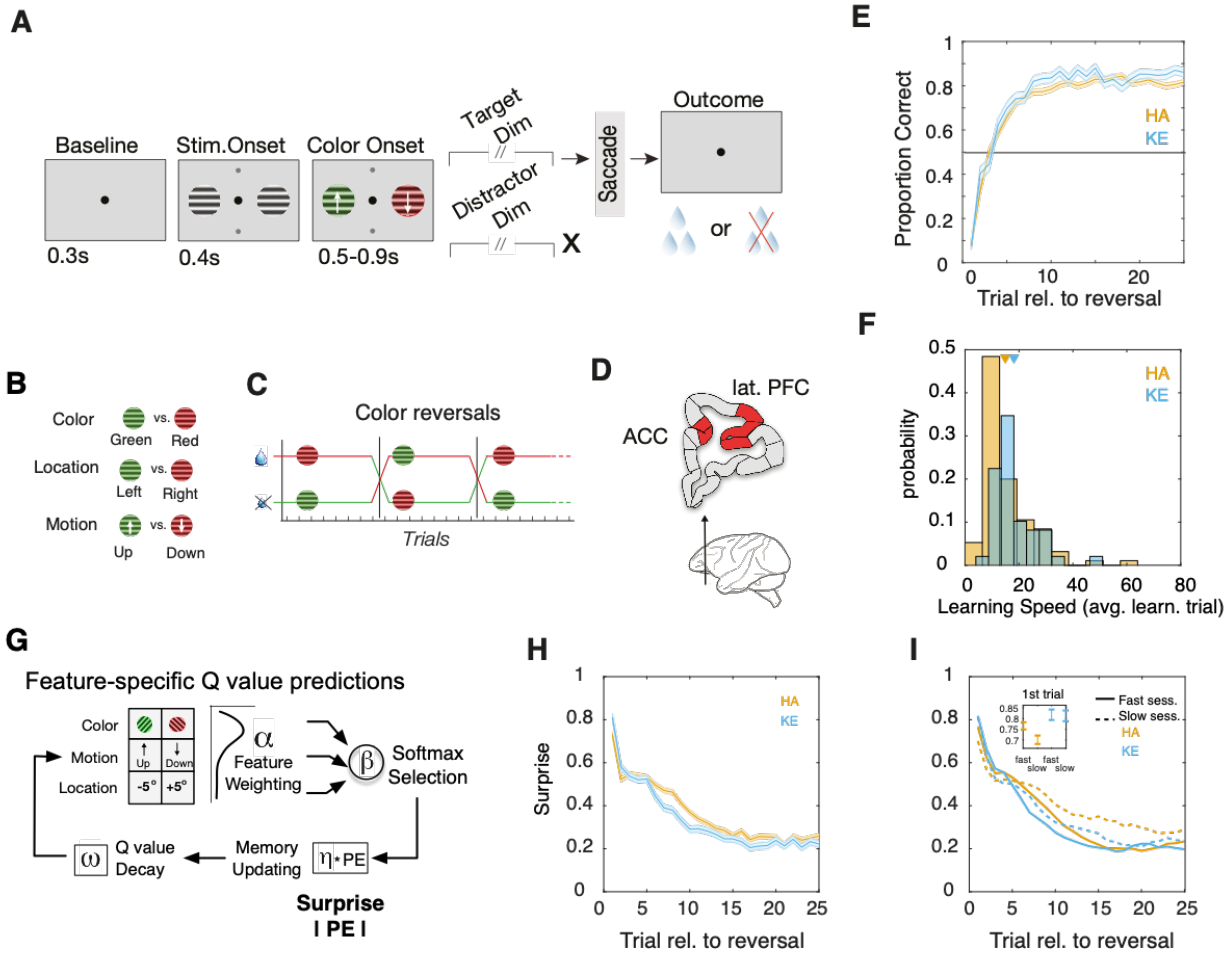

Figure 3-1. Subjects were less sensitive to surprising outcomes in slow learning sessions

(A) Feature-based reversal learning task. Animals are presented with two black/white stimulus gratings to the left and right of a central fixation point. The stimulus gratings then become colored and started moving in opposite directions. Dimming of the stimuli served as a Go signal. At the time of the dimming of the target stimulus the animals had to indicate the motion direction of the target stimulus by making a corresponding up or downward saccade in order to receive a liquid reward. Dimming of the target stimulus occurred either before, after or at the same time as the dimming of the distractor stimulus. **(B)** Each stimulus is defined by three features; color, location, and direction of motion. (**C)** The task is a deterministic reversal learning task, whereby only one color is rewarded in a block. This reward contingency switches repeatedly and unannounced in a block-design fashion. **(D)** Schematic depicting brain areas used in this present study. **(E)** Accuracy relative to block start for monkey HA (orange) and KE (blue). Shaded region represents the standard error. Subjects achieved plateau performance within 5-10 trials **(F)** Distribution of learning speed (mean learning trial) across sessions for each individual monkey**. (G)** Reinforcement learning algorithm weighting feature dimensions. Parameters alpha, beta, omega, and eta (α, β, ω, and η) represent the feature weighting, selection noise, decay rate, and learning rate, respectively. **(H)** Mean and standard errors of surprise (absolute PE) on each trial aligned to reversals. Both monkeys exhibit a gradual decrease in surprise as the block progresses **(I)** Same as (H), but split into fast (top 50%, solid line) and slow (bottom 50%, dotted line) learning sessions. Surprise decreased even more gradually on slow learning sessions. Surprise was also lower on the first trial following a reversal (inset).

3.4.2 Spikes synchronizing to distal beta signal surprise

The behavioral results suggest that the representation of surprise may be degraded on sessions with slower overall learning. In a previous study, we have shown that deviations from expectations are best signaled by spikes falling specific phases of a beta rhythm (**Chapter 2**; (Voloh et al., 2019a)). Importantly, we also observed that information gain was highest for phases that deviated away from the preferred firing phase of cells. Thus, we reasoned that that spike-phase synchrony may be a good candidate for signaling surprise, and that when the relationship between spikes and surprise was degraded, learning would be impeded.

Our hypothesis is schematized in **Figure 3-2A**. As animals' performance increases, outcome surprise decreases. There are three possibilities for how the spikes synchronizing to oscillatory activity may relate to the signaling of surprise. The null hypothesis is that there is no relationship; thus, surprise (as derived from the model) would be equivalent across spike-phases. Alternatively, if spikes do indeed signal surprise, then we may expect that spikes concentrated on specific phases tend to occur during particularly surprising (or unsurprising) trials. We term the degree to which spike phase predicts trial unexpectedness as *spike-surprise sensitivity*, corresponding to the "A" term in equation 3. The phase at which surprise is maximal relative to the preferred phase is termed the *spike-surprise specificity,* and corresponds to the θ in equation 3. Spike-surprise sensitivity and specificity are somewhat dependent; if there is no sensitivity, then there can be no specificity.

As a first step, for each cell-LFP pair (n=3600), we modelled the relationship of surprise (i.e. absolute RPE derived from the model) and the modal spike phase on each trial (*see* Methods). **Figure 3-2B** shows two example cell-LFP pairs with a characteristic relationship between the modal spike phase and surprise. Both examples show a high degree of *spike-surprise sensitivity*, with higher RPEs occurring on specific modal phases. However, they differ in their *spike-surprise specificity*; the example in the top panel signals surprise approximately on the anti-preferred firing phase, whereas the example on the bottom signals surprise specifically on the preferred phase.

We hypothesized that variation in the *spike-surprise sensitivity and specificity* may underlie variation in learning speed across sessions (**Figure 3-2C**). Either sensitivity or specificity may signal surprise. Decreased spike-surprise sensitivity would result in a poorer representation of outcome unexpectedness, thus resulting in slower learning rate. On the other hand, spike-surprise specificity may best signal surprise when spikes occur on phases that maximally deviate from the expected firing phase. Thus, spikes falling near the expected firing phase would result in slower learning, whereas spikes closer to unexpected phases would improve learning speed.

To test the hypothesis that spike-surprise sensitivity and/or specificity may underlie session-wide learning speed, we used a linear mixed effects model (**Figure 3-2D, equation 9**). Fixed effects included the parameters A and θ, corresponding to spike-surprise sensitivity and specificity, respectively, with subject as a random effect. We performed this analysis for LFP phases for each frequency in a 6-64 Hz range (**Figure 3-2E**). We found spike-surprise predicted learning speed in a ~10-25 Hz range, with particularly pronounced significant effects at frequencies centered around 14 and 25 Hz (**Figure 3-2E**). This suggests that spikes falling on specific phases may signal surprising outcomes, and thus improve learning performance, in line with our previous work (**Chapter 2;** (Voloh et al., 2019a)).

In a similar beta band, maximal spike-surprise sensitivity *A* predicted learning sessions that were ~5 trials faster on average (**Figure 3-2F).** The average sensitivity in the 10-25 Hz beta band was 0.077 ± 0.00052 SE. In a similar vein, spike-surprise specificity θ also predicted faster learning speed. Because there are two terms in our model related to this factor, we first visualized the effect in relation to a "test phase" corresponding to the anti-preferred phase of θ (**Figure 3-2G**). This revealed that learning speed was faster by about ~0.5 trials when spikes on surprising trials were concentrated on anti-preferred phases. **Figures 3-2H** show the full characterization in terms of the strength of the effect (**Figure 3-2Hi**), the phase at which it was maximal (**Figure 3-2Hii**), as well as the actual beta weights associated with the terms related to spike-surprise signaling (**Figure 3-2Hiii-iv**). In agreement with our analysis showing the anti-preferred phase resulted in *faster* learning, we can see that learning is predicted to be slower (**Figure 3-2Hi**) at phases near the cell *preferred* firing phase (**Figure 3-2Hii**). In summary, spikes locking to beta

oscillations showed spike-surprise sensitivity (**Figure 3-2F**) and specificity (**Figure 3-2G**) that predicted session-wide learning speed, suggesting a potentially novel mechanism for signaling surprise during feature-based reversal learning.

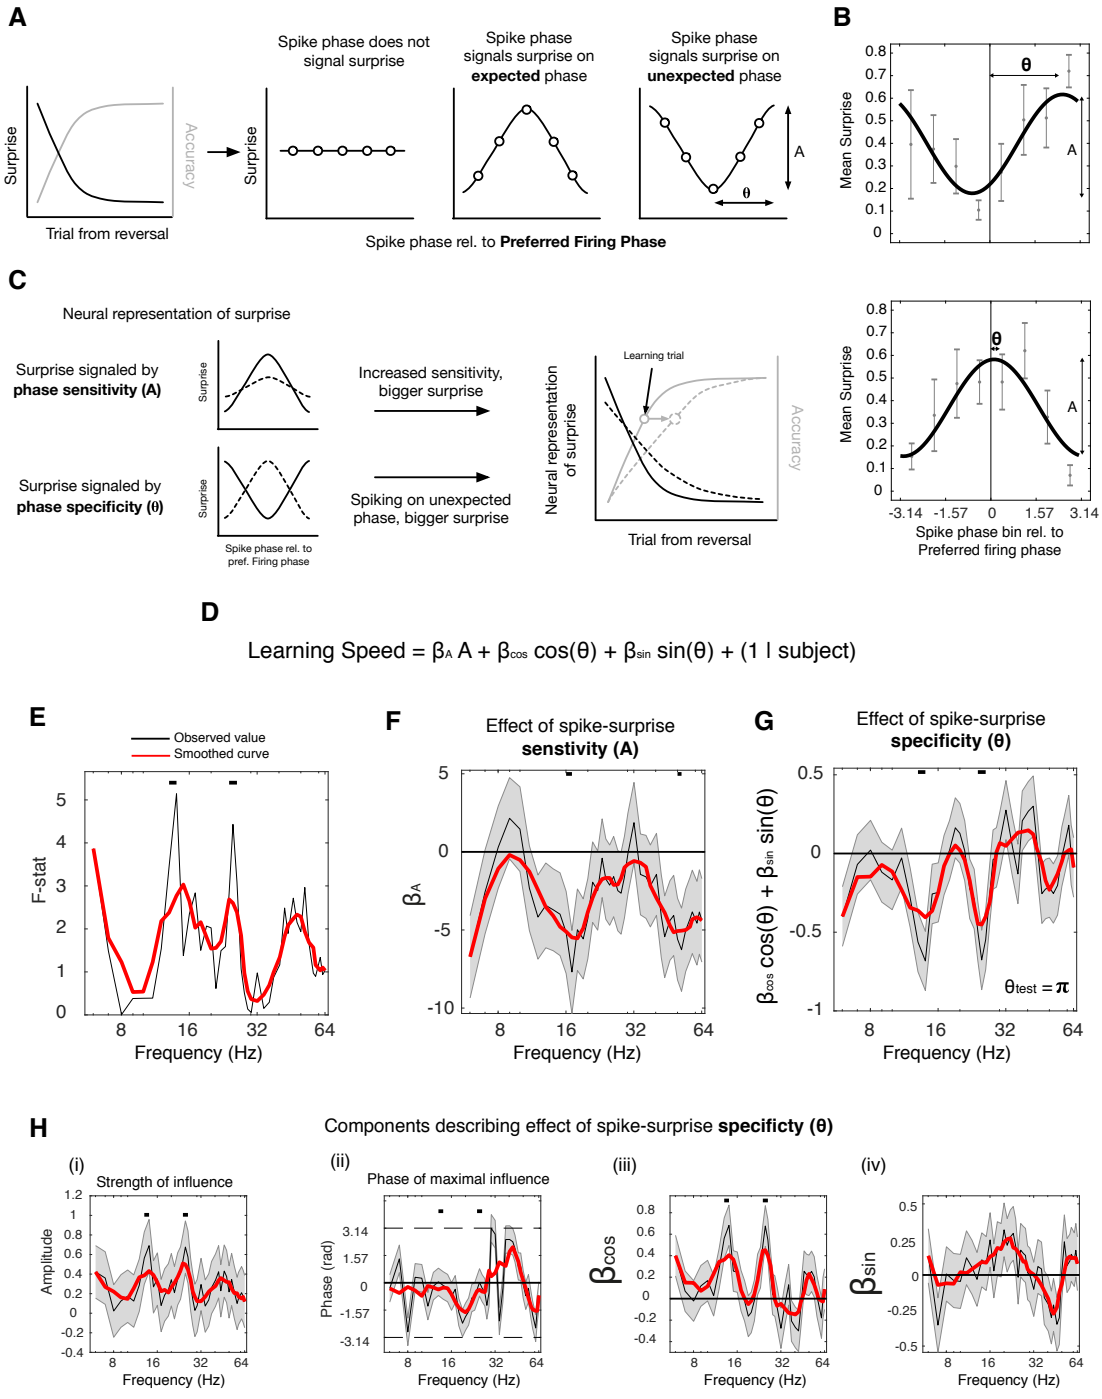

Figure 3-2. Spikes synchronizing to beta phases signaled surprise that predicted learning speed

(A) Possible relationship of spike-phases and surprise. As subjects learn rewarded features, accuracy increases and surprise decreases. The null hypothesis is that spikes are distributed randomly on surprising or unsurprising trials. Alternatively, spikes may concentrate on particular phases as a function of the degree of outcome unexpectedness. The degree to which spike phases signal surprise is the spike-surprise sensitivity *A.* The relative phase at which surprise is highest is the spike-surprise specificity θ. **(B)** Example neurons showing a high degree of spike-surprise sensitivity. Surprise values were binned according to the modal phase during the outcome processing period on that trial. We used 8 non-overlapping phase bins to visualize the effect. Grey bars represent the mean and standard error. The black overlaid curve represents the fitted cosine. Spike surprise *specificity* differed for these two cells, with high surprise signaled on anti-preferred phases (top) or preferred firing phases (bottom). **(C)** Spike-surprise sensitivity or specificity may signal surprise. Thus, degraded representation of surprise would result in slower learning, shifting the learning point to the right. **(D)** Linear-mixed effects model describing how the learning speed varies as a function of surprise signaled by spike phases**. (E)** F-statistic for the model in (D) with spike-phases estimated at each frequency. Black line is the actual observed value, and the red is smoothed curve (5 samples) visualizing the gross trends. Black horizontal bars above the line denote significant model fits ($p<0.05$). **(F)** Value and error associated with the beta coefficient for the spike-surprise sensitivity parameter A. **(G)** Output of terms corresponding to spike-surprise specificity θ, for a test input of $\theta=\pi$ (corresponding to anti-preferred phase firing. **(H)** Estimates of the strength of the spike-surprise specificity (i) and the phase at which these effects are maximal (ii). These are derived from beta coefficients, represented in (iii-iv).

3.4.3 Phase specificity signals surprise, not value

Up to this point, we have analyzed spike in relation to surprise, leaving open the question if spike phases are more attuned to other learning variable. One possibility is that spike-phases may be more related to the value of the chosen stimulus, which is highly correlated with surprise (Pearson correlation, R=-0.50, p=0). To this end, we first determined the spike-*value* sensitivity and specificity for each cell-LFP pairs, and then used these variables to predict learning speed (as above) (**Figure 3-3**). We found that this model could predict learning speed in the same beta band, as well as a theta (centered at 7Hz) and gamma (50 Hz) band **(Figure 3-3A)**. This effect was driven primarily by spike-value sensitivity (**Figure 3-3B**), but not specificity (**Figure 3-3C**), suggesting that sensitivity and specificity may signal two separate streams of information. To test this hypothesis, we fit a reduced model to spike-value or spike-surprise specificity/sensitivity (**Figure 3-3D**), and extracted the AIC to aid in model comparison. We found that in the beta band, spike-surprise, rather than spike-value, specificity better predicted learning speed (**Figure 3-3D, left**) in a beta band. The inverse was true for sensitivity (**Figure 3-3D, right**). This suggests that sensitivity and specificity may carry information related to value and surprise, respectively.

In line with these results, we found that spike-surprise sensitivity was indistinguishable between pre and post outcome periods (randomization test, p=0.78), whereas specificity differed between pre and post-outcome periods (p=0).

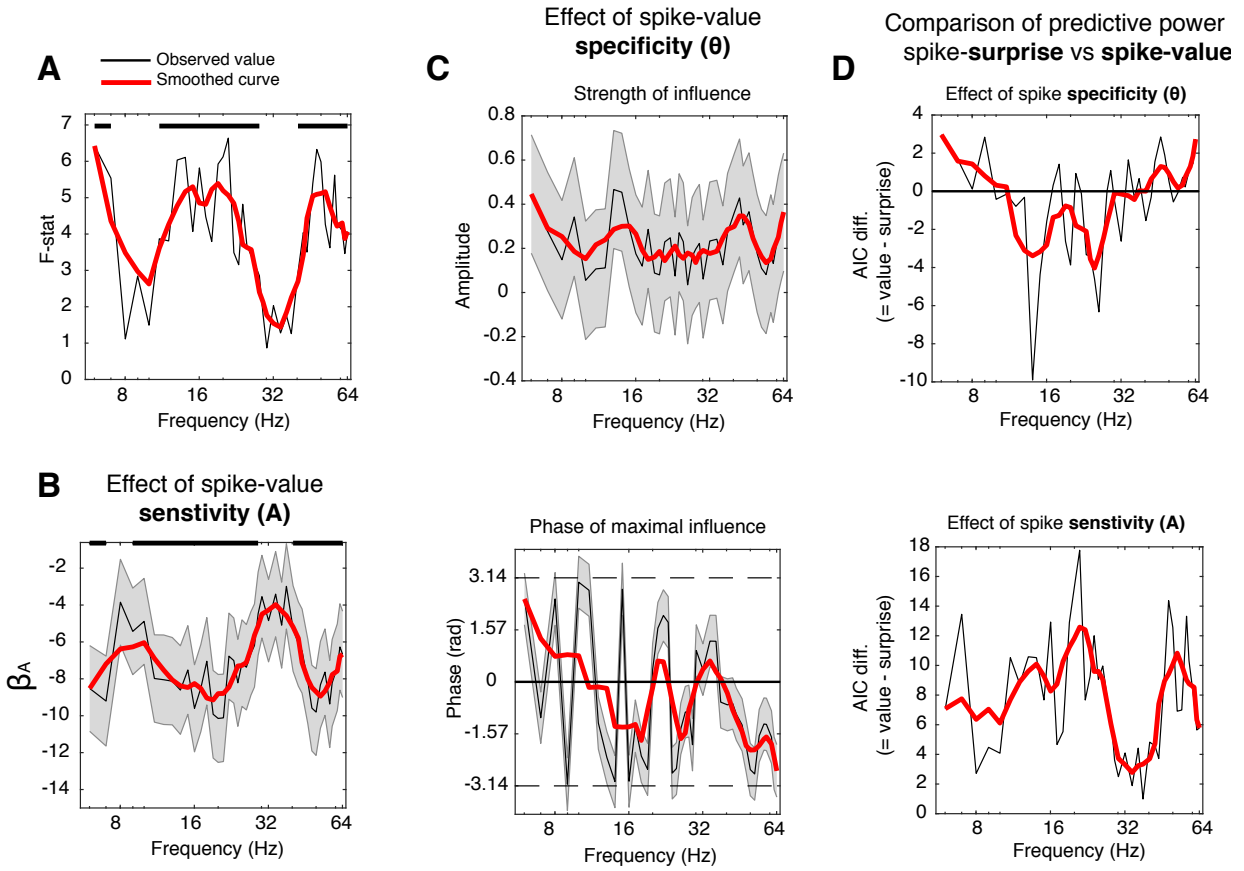

Figure 3-3. Surprise is better signaled by spike specificity, and value by sensitivity

(A) Same as in Figure 2D-E, but estimating sensitivity and specificity using the (model-derived) value of the chosen object. Significant effects are evident in theta, beta, and gamma bands (p<0.05). **(B)** Influence of spike-value sensitivity on learning speed. The effects in (A) can be traced to the effect of sensitivity. **(C)** The influence of spike-value specificity (see Eqn. 5-8), visualized in terms of the strength of the effect (top) and the phase at which the effect is maximal (bottom). There are no significant effects. **(D)** (top) Model comparison of spike-value vs surprise specificity in predicting learning speed. Spike-surprise specificity better predicted learning speed. (bottom) On the other hand, spike-value sensitivity better predicted learning speed.

3.4.4 Spike-phase synchrony and rate changes predict spike-surprise sensitivity

Previous studies on phase shifting have indicated that spikes are more likely to shift phase in relation a stimulus property when the overall level of spike-phase synchronization is low (Vinck

et al., 2010a). Intuitively, this makes sense; spike-phase locking would be maximal if spikes only lock to a single phase, thus preventing shits in relation to surprise. Inspired by this, we next tested the relationship of spike-phase synchronization with the spike-surprise sensitivity in a [10 25] Hz band corresponding to beta. We found that the two were negatively correlated (Spearman R=-0.11, p~0), indicating that when spike-phase synchronization is low, spikes span a wider range of phases and are more tightly linked to the degree of outcome unexpectedness. In a related analysis, we tested if changes in firing rate could predict spike-surprise sensitivity. To this end, we first correlated firing rates on individual trials in the outcome period with surprise, and extracted the Spearman Rho. We found that Rho and spike-surprise sensitivity were modestly but significantly correlated (Spearman R=0.039, p=0.019). Taken together, these results tentatively suggest that firing rate intensity may be converted into a phase-of-firing code (Buzsáki and Draguhn, 2004; Fries et al., 2007).

3.4.5 Putative excitatory cells exhibit stronger phase gain than putative inhibitory cells

Our results up to this point indicate that spikes synchronized to distal beta rhythms carry information related to outcome unexpectedness, and that the fidelity of this representation may effect learning speed. To gain insight into potential circuit mechanisms that may support this process, we classified cells into two distinct narrow and broad waveform groups (**Figure 3-4A** *see* Methods; (Ardid et al., 2015; Oemisch et al., 2019)). On the basis of our results above, we compared spike-surprise sensitivity between narrow (n=306) and broad (n=631) cell-LFP pairs, for fast and slow learning sessions (**Figure 3-4B**). We focused this analysis for spike-surprise signaling in the beta band for frequencies that showed prominent beta phase spike-phase synchrony (*see* **Methods**). We found that spike-surprise sensitivity was greater for broad cells than narrow cells for fast sessions (Unpaired T-test, p=0.019). On the other hand, sensitivity was statistically indistinguishable for slow sessions (p=0.26).

We next tested for differences in spike-surprise specificity **(Figure 3-4C).** In fast learning trials, the phase associated with maximal surprise differed between narrow and broad-waveform cells (Watson-Williams test, p~0), with broad cells showing specificity near anti-preferred phases (2.27

± 0.94 rad.), and narrow cells near preferred phases (0.28 ± 0.92 SE rad.). This difference was not present in slow learning sessions (p=0.72; Narrow cells, -0.35 ± 0.88 SE rad; Broad cells, -0.43 ± 0.92 SE rad). Taken together, these results suggest that during fast learning sessions, surprise is signaled by phase-specific firing of broad-waveform cells (putative excitatory cells).

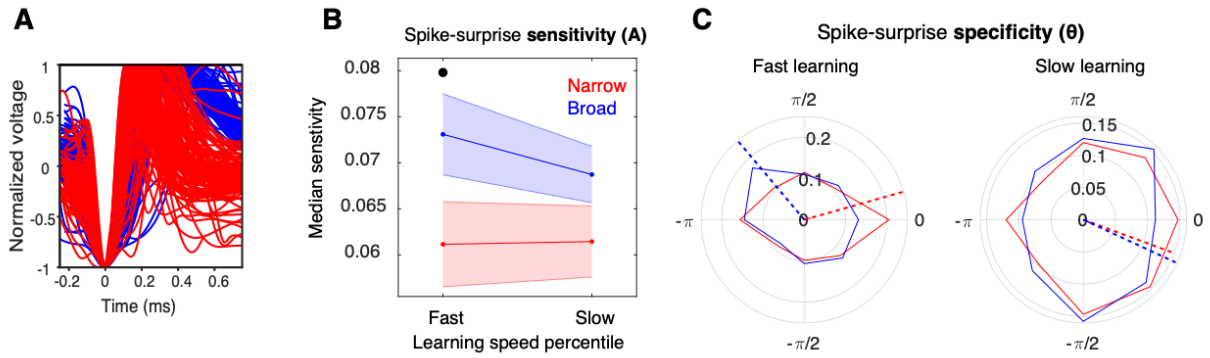

Figure 3-4. Broad-spiking neurons showed higher phase gain specifically for faster learning trials

(A) We classified neurons into broad (blue) and narrow (red) cell classes on the basis of the action potential waveforms. **(B)** Mean and standard error of spike-surprise sensitivity for narrow (red) and broad (red) cell classes, for fast (bottom 50% percentile) or slow (top 50% percentile) learning sessions. Note that we average across spike-LFP pairs here (so the same neuron may be selected more than once). Broad waveform cells showing higher sensitivity than narrow spiking cells during fast sessions (p=0.019), but not slow sessions (p=0.26). **(C)** Distribution of spike-surprise specificity for fast (left) and slow (right) learning sessions. Broad and narrow waveform cells signaled maximal surprise on different phases during fast learning session (p~0), but not slow learning sessions (p=0.72).

3.5. Discussion

In the current study, we found preliminary evidence that a characteristic relationship between the modal spike phase and outcome unexpectedness could predict session-wide learning speed in a feature-based reversal learning task. This effect was particularly pronounced in the beta band, for cell-LFP pairs across the ACC and LPFC. The relationship between spike phases and surprise was particularly pronounced for well-isolated, broad-waveform neurons, which correspond to putative excitatory cells. Taken together, these results provide preliminary evidence that the phase at which spikes occur conveys information that may be leveraged to quickly learn feature-reward associations.

We found here that faster learning speeds were related to increased spike-surprise sensitivity and specificity across ACC and LPFC circuits. This suggests that the phase of spiking may serve as a neural substrate signaling surprise for the purpose of updating value estimates, in line with evidence that ACC activity co-varies with relevant but not irrelevant surprise (O'Reilly et al., 2013). Surprise is directly related to an important computational principle, namely, unexpected uncertainty (Yu and Dayan, 2005). Unlike expected uncertainty, which is known variability in the environment (stochasticity), unexpected uncertainty is that which can be reduced through repeated observation (Gottlieb, 2012). Importantly, both expected and unexpected uncertainty are predicted to affect learning rates (Bach and Dolan, 2012). In the former case, decreased learning rates reduce the influence of one-off observations, while in the latter case, transient increases in unexpected uncertainty serves as a signal for exploration (Payzan-LeNestour and Bossaerts, 2012) and thus increase learning rates.

Because our task is deterministic, stochasticity plays no role in affecting learning rates. Moreover, because we study session-wide changes in learning rates, transient changes in unexpected uncertainty also cannot explain our result. This points to another source of uncertainty for cortical circuits, namely, estimation uncertainty, which represents uncertainty associated with internally generated beliefs (Yu and Dayan, 2005), which are also represented in the anterior cingulate and lateral prefrontal cortices (Payzan-LeNestour et al., 2013). Estimation uncertainty can also affect learning rates, though indirectly, by influencing the assessment of unexpected uncertainty (Payzan-Lenestour and Bossaerts, 2011). Our results suggest that the phase of spiking may thus be a source of estimation uncertainty.

What may alter estimation uncertainty on a session-wide level? One powerful driver is the overall level of motivation of the animal. Motivation can alter/transform the representation of task-relevant stimuli, a feature that has been termed "incentive salience" ((Robinson and Berridge, 2001; McClure et al., 2003). Incentive salience has been linked to the activity of midbrain dopaminergic neurons (Berridge, 2007), which send diffuse projections throughout the cortex, and specifically to the ACC and LPFC (Van Eden et al., 1987; Goldman-Rakic et al., 1990).

Importantly, a previous study has shown that dopamine administration at cortical synapses changes patterns of spike-phase synchrony in a manner that resembles their change during learning (Benchenane et al., 2010). In this study, the authors observed that coherence between prefrontal cortex spikes and hippocampal LFPs during choice points emerged after learning of a new rule. During these high coherence periods (after learning), as compared to low coherence periods (before learning), pyramidal but not inhibitory cells showed a greater degree of phase-concentration, as well as an almost 180^0 shift in their phase preference. Strikingly, they observed the same effects in anesthetized animals after (local) dopamine administration. Our results complement and extend these to a feature-based domain, and suggest that that alterations in spike-phase synchrony specifically signal surprise for the purpose of learning. We predict that such changes are promoted by dopaminergic tone originating in midbrain dopamine neurons.

In summary, we have shown that the sensitivity with which neurons lock to distal beta, and the specific phase at which they do so, predict session-wide changes in learning speed. These effects are particularly pronounced for putative excitatory cells rather than inhibitory cells during fast learning sessions. By exploiting deviations from expected firing phases, unexpected (i.e. surprising) spikes may serve to promote plasticity (Zanos et al., 2018). The specific timing of such spikes in relation to the reference beta oscillation may recruit spike-timing dependent plasticity mechanisms (Wespatat et al., 2004; Fell and Axmacher, 2011) that strengthen synapses leading to optimal outcomes.

3.6. Statement of contribution

BV, Benjamin Voloh; MV, Martin Vinck; MO, Mariann Oemisch; TW, Thilo Womelsdorf

Conceptualization, BV, MV, TW. Methodology, BV, MV, and TW; Investigation, MO, TW; Original draft – Writing: BV, MV, TW; Original draft – Editing, all authors; Visualization, BV, MV, TW; Supervision, TW

Chapter 4

Cell-Type Specific Burst Firing Interacts with Theta and Beta Activity in Prefrontal Cortex During Attention States

4.1. Preamble

This Chapter is a published work (Voloh and Womelsdorf, 2017), and can be accessed here: https://doi.org/10.1093/cercor/bhx287. I reproduce it here for clarity.

As reviewed in the introduction, bursting and nonbursting regimes of single cells can carry different information streams (Körding and König, 2000; Guergiuev et al., 2016; Naud and Sprekeler, 2018). Moreover, in a previous study in our lab, we had found that bursting preferentially synchronizes to distal beta and gamma rhythms (Womelsdorf et al., 2014a). This suggests that If burst activity carries unique learning-relevant information, it may be preferentially communicated via synchronized spiking in the beta and gamma band.

Unlike **Chapter 2 and 3,** which directly tested phase-of-firing coding during a reversal learning task, this chapter is based on the analysis of a cued spatial attention task. Moreover, we do not explicitly measure information contained in the analyzed neuronal activity. That said, the motivation for this study was to build upon these prior results and ask how burst firing interacts with *local* rhythms, which would allow stronger inference of putative mechanisms underlying burst-mediated rhythmic activity. Given the prevalence of beta-rhythmic coding that we observe (**Chapter 2 and 3**), we suspect that similar interactions may occur during learning. An interesting future direction would be compare burst-beta synchrony when features must be *associated* with rewards **(Chapter 2 and 3)**, rather than during the attentional *selection* of relevant features.

4.2. Abstract

Population-level theta and beta band activity in anterior cingulate and prefrontal cortices (ACC/PFC) are prominent signatures of self-controlled, adaptive behaviors. But how these

rhythmic activities are linked to cell-type specific activity has remained unclear. Here, we suggest such a cell-to-systems level linkage. We found that the rate of burst spiking events is enhanced particularly during attention states and that attention-specific burst spikes have a unique temporal relationship to local theta and beta band population level activities. For the 5-10Hz theta frequency range, bursts coincided with transient increases of local theta power relative to non-bursts, particularly for bursts of putative interneurons. For the 16-30Hz beta frequency, bursts of putative interneurons phase synchronized stronger than nonbursts, and were associated with larger beta power modulation. In contrast, burst of putative pyramidal cells showed similar beta power modulation as nonbursts, but were accompanied by stronger beta power only when they occurred early in the beta cycle. These findings suggest that in the ACC/PFC during attention states, mechanisms underlying burst firing are intimately linked to narrow band population level activities, providing a cell-type specific window into rhythmic inhibitory gating and the emergence of rhythmically coherent network states during goal directed behavior.

4.3. Introduction

Narrow band population-level theta and beta band activity emerge during goal directed behavior in anterior cingulate and prefrontal cortices (ACC/PFC). The occurrence and strength of theta and beta activity thereby closely relates to behavioral functions including successful attention (Sheth et al., 2012; Voloh et al., 2015), correct retrieval of task rules (Buschman et al., 2012b; Phillips et al., 2014), behavioral adjustment following errors (Womelsdorf et al., 2010b), or short-term maintenance of stimulus-response mapping rules (Salazar et al., 2012; Babapoor-Farrokhran et al., 2017) . These functional correlates of theta and beta activities emerge from the activation of cells and local circuits, but it has remained a fundamental open question which cell and circuit mechanisms are directly linked to these population level, narrow band activities (Kopell et al., 2014; Womelsdorf et al., 2014b).

Growing evidence suggests that population level theta and beta activities are not supported equally by all neurons in a circuit, but rather that distinguishable cell-types show

specific preferences to synchronize to the local electrical field at only a subset of narrow band frequencies (Hasenstaub et al., 2005; Womelsdorf et al., 2014b; Roux and Buzsáki, 2015). For example, in nonhuman primate prefrontal cortex, subsets of putative interneurons and putative pyramidal cells, defined by their narrow and broad action potential waveform shape, respectively, show unique synchronization preferences to only beta or theta activity during attentive states in the primate (Ardid et al., 2015). Direct optogenetic control of spiking activity has likewise shown cell specific preferences to synchronize to the local oscillatory activity, with different subtypes of interneurons and pyramidal cells linked to theta, beta or higher frequency activity (Cardin et al., 2009; Sohal et al., 2009; Stark et al., 2013; Kim et al., 2015, 2016; Bitzenhofer et al., 2017).

In addition to the cell-type, it has been documented that even for the same cell, not all spikes contribute similarly to population level rhythmic activity (Denker et al., 2011; Womelsdorf et al., 2014a) . This is particularly apparent for bursts, consisting of two or more spikes within a short (e.g. 5ms) time window. Previous studies suggests that bursts of pyramidal cells may directly index coordinated activity across larger recurrent networks (Larkum, 2013). In particular, pyramidal cell bursts can be a direct consequence of coincident arrival of dendritic and somatic synaptic inputs from diverse distant sources (e.g. (Mainen and Sejnowski, 1996; Larkum et al., 1999, 2007; Waters, 2004; Manita et al., 2015; Sherman et al., 2016)). Bursts also have an outsized role in shaping neural activity; interneuron bursts can induce large compound inhibitory postsynaptic potentials strong enough to silence connected pyramidal cells (Hilscher et al., 2017), bursts of projection cells have enhanced postsynaptic efficacy in driving targets compared to singleton spike events (Swadlow and Gusev, 2001), and induce more powerful long-term weight changes at their postsynaptic sites than singleton spikes (Birtoli and Ulrich, 2004; Bittner et al., 2015; Wilmes et al., 2016). Moreover, it is believed that burst spikes are generated by mechanisms that are distinct from those of singleton spikes, suggesting that bursts could form a unique information channel during neuronal information processes (Krahe and Gabbiani, 2004; Larkum, 2013; van Ooyen and van Elburg, 2014). Taken together, these characteristics assign

bursts a particular role in the local neural circuit to shape how input is transformed into effective output of the circuit (Sahasranamam et al., 2016).

Despite the possible importance of bursts to shape network processes, direct support for the role of burst firing in coordinated network activity during actual cognitive processes is sparse. In a recent report, we have documented that bursts firing of neurons in the ACC/PFC, but not isolated spikes of the same neurons, synchronized reliably to field activity in distant areas at narrow band theta, beta and gamma- band frequencies (Womelsdorf et al., 2014a). This study pointed to burst firing events as a unique signature of long-range network activity, but left unanswered how bursts interact in the local circuits in which the burst event occurs.

Here, we first build on these earlier results and show that burst rate and the proportion of burst firing of neurons in ACC/PFC show sustained increases during a selective attentional state. These burst rate increases emerged in neural circuits whose population activity is characterized by power spectral peaks in the theta and beta band. To connect burst firing with population level theta and beta band activity, we characterized spike-triggered LFP activity around burst spikes and non-burst spikes. We found that spikes constituting the beginning of a burst firing event coincide with transient increases in theta LFP power when compared to non-burst spikes. This burst specific theta power modulation was particularly apparent for bursts of putative interneurons that were identified by their narrow action potential waveform. Independent of the theta power burst relationship, we found for the beta frequency band that burst spikes synchronized stronger to phases of the beta cycle than non-burst spikes, but without concomitant modulation of beta power. These findings reveal cell-type specific relationships of burst firing with meso-scale network activity indexed by narrow-band LFP components.

4.4. Materials and Methods

4.4.1 Experimental Procedures

Experiments were conducted in two awake and behaving macaque monkeys as described in detail in (Kaping et al., 2011), following the guidelines of the Canadian Council of Animal Care

policy on the use of laboratory animals and of the University of Western Ontario Council on Animal Care. Extra-cellular field potential and action potential signals were recorded in each recording session from 1-6 tungsten electrodes (impedance 1.2-2.2 MΩ, FHC, Bowdoinham, ME) through standard recording chambers (19mm inner diameter) implanted over the left hemisphere in both monkeys. Electrodes were lowered through guide tubes with software controlled precision microdrives (NAN Instruments Ltd., Israel) on a daily basis, through a recording grid with 1 mm inter-hole spacing. Before the first recording session, anatomical 7T MRIs were obtained to visualize and reconstruct electrode as described in detail in (Kaping et al., 2011). Data amplification, filtering, and acquisition were done with a multi-channel processor (Map System, Plexon, Inc.), using headstages with unit gain.

The recording experiments were performed in a sound attenuating isolation chamber (Crist Instrument Co., Inc.) with monkeys sitting in a custom-made primate chair viewing visual stimuli on a computer monitor (85 Hz refresh rate, distance of 58 cm). The monitor covered 36º x 27º of visual angle at a resolution of 28.5 pixel/deg. Eye positions were monitored using a video-based eye-tracking system (ISCAN, Woburn, US, sampling rate: 120 Hz). Eye fixation was controlled within a 1.4-2.0 degree radius window. Stimulus presentation, monitored eye positions and reward delivery were controlled via the open-source software MonkeyLogic. Liquid reward was delivered by a custom made, air-compression controlled, mechanical valve system with a noise level during valve openings of 17 dB within the isolation booth.

4.4.2 Behavioral task.

Monkeys performed a covert selective attention, 2-forced choice discrimination task (**Figure 4-1A**). Following a 2 second intertrial interval, a small gray fixation point was presented centrally on the monitor. Monkeys had to direct their gaze and keep fixation onto that fixation point until a change-event of the target stimulus late in the trial. After 300 ms fixation, two black/white grating stimuli were presented to the left and right of the center and contained oblique movements of the grating within their circular aperture. After 0.4 s, each stimulus changed color to either black/red or black/green. After a variable time (0.05 to 0.75 s) the color

of the central fixation point changed to either red or green, which cued the monkeys to covertly shift attention towards the stimuli that had the same color as the attention cue. Monkeys maintained central fixation and sustained covert peripheral attention on the cued stimulus until it underwent a transient clockwise or counter-clockwise rotation, ignoring possible rotations of the non-attended (uncued) stimulus, which occurred in 50% of the trials. In order to obtain liquid reward, the monkeys had to discriminate the rotation by making up- or downward saccades for clockwise /counter-clockwise rotations (the mapping was reversed between monkeys). Following this overt choice and a 0.4 s waiting period the animals received fluid reward (for a detailed description, see (Kaping et al., 2011)). A key component of the task is that the location of covert spatial attention on one of the two colored stimuli (left or right) is distinct from the possible locations to which the animal made a saccade (up or down) to indicate the transient rotation of the attended stimulus.

4.4.3 Neuron isolation.

During recording, the spike threshold was adjusted such that there was a low proportion of multiunit activity visible against which we could separate single neuron action potentials in a 0.85 to 1.1 ms time window. Sorting and isolation of single unit activity was performed offline with Plexon Offline Sorter (Plexon Inc., Dallas, TX), using the separation of the first two to three principal components of the spike waveforms, and strictly limiting unit isolation to periods with temporal stability. For analysis, we selected the subset of 422 maximally isolated single units whose waveform principle components were clearly separated with a density profile separated from the density profiles from multiunit background activity and other simultaneously recorded waveforms. The first two principle components explained on average 73.37% (± 1.3 SE) of variance across all waveforms that crossed thresholds. To quantify the separation of the waveforms' first two principal component scores we calculated the Mahalanobis (ML) distance (using the Matlab function **mahal**). The ML distance metric uses the matrix of distances between data points to the mean, and the variance / covariance matrix to calculate the multivariate distances between points. We calculated the ML distance for the first two principal component scores of the spike waveforms of the recorded single unit relative to the scores of the waveform

of the multi activity and noise of the same recorded channel and found an average ML distance of 24.12 ±1.8 (for examples see **Figure B-1**).

4.4.4 Classifying cell types using spike waveform analysis.

For all well isolated neurons we normalized and averaged all action potentials (APs) and extracted the peak-to-trough duration and the time of repolarization as described in detail in (Ardid et al., 2015). The time for repolarization was defined as the time at which the waveform amplitude decayed 25% from its peak value. Across the average waveforms of the cells we calculated the Principal Component Analysis and used the first component (explaining 84.5 % of the total variance), as it allowed for better discrimination between narrow and broad spiking neurons, compared to any of the two measures alone. We used the calibrated version of the Hartigan Dip Test (Hartigan and Hartigan, 1985) to discarded unimodality for the first PCA component ($p < 0.01$) and for the peak to trough duration ($p < 0.05$) but not for the duration of 25% repolarization ($p > 0.05$). Additionally, we tested whether the distribution of the PCA score is better fit with two rather than one Gaussian. We applied Akaike's and Bayesian information criteria to test whether using extra parameters in the two-Gaussian model is justified. In both cases, the information criteria decreased (from -669.6 to -808.9 and from -661.7 to -788.9, respectively), confirming that the two-Gaussian model is better. We then used the two-Gaussian model and defined two cutoffs that divided cells into three groups. The first cutoff was defined as the point at which the likelihood to be a narrow spiking cell was 10 times larger than a broad spiking cell. Similarly, the second cutoff was defined as the point at which the likelihood to be a broad spiking cell was 10 times larger than a narrow spiking cell. This ensured across the whole population that 95% of cells (n = 401) were reliably classified: neurons at the left side of the first cutoff were reliably classified as narrow spiking neurons (18.7%, n = 79), neurons at the right side of the second cutoff were reliably classified as broad spiking neurons (76.5%, n = 323). The remaining neurons were left 'unclassified' as they fell in between the two cutoffs (4.7%, n = 20).

4.4.5 Data Analysis

All analyses were performed using Matlab (The Mathworks). Throughout, we used conservative, non-parametric tests to draw our conclusions, and report throughout on the median. For the figures, the standard error of the median was estimated with a bootstrap procedure. Briefly, the distribution of medians was generated with a bootstrap procedure, and the estimate of the standard error was the standard deviation of this distribution (Efron and Tibshirani, 1986).

All analyses were performed on correct trials. The baseline period was defined as activity that occurred 0.5 s. before cue onset. The attention period was defined as activity that occurred after cue onset, but before the first rotation of a stimulus, i.e. before target or distractor rotation (see **Figure 4-1A**). To prevent experimental artifacts from affecting analyses, we ignored trials were the LFP deflection was greater than 10 SD away from the mean for that trial (median percentage of discarded trials = 0.78 +/- 0.002%).

Burst events were defined as spikes that occurred with an interspike interval of ≤5 ms. All burst analyses were performed on the first spike of a burst event. Bursts are rare events occurring less frequently than individual spikes. Moreover, low spike numbers can result in highly variable estimates of phase consistency (Vinck et al., 2010b). To ensure sufficient number of spike/burst events for spectral analyses, we selected the subset of neurons that had (1) at least 30 burst spikes within the post-cue period used for analysis, and (2) with at least $\pm$ 0.5 s. of LFP data around the time of the spike. The $\pm$ 0.5 s. time window around spikes never overlapped with either the onset time of the centrally presented cue or the time of stimulus rotation. All LFP analyses were performed on the same channel as the spike. Based on these criteria, we analyzed 41 cells of the total of 422 recorded cells.

To prevent spike artifacts in the LFP, we first lowpass filtered the LFP at 100 Hz using a two-pass 4th order Butterworth filter. Next, to prevent spike-locked artifacts, we used an interpolation approach when analyzing spike-triggered effects (Ardid et al., 2015). For each spike-

centered LFP segment, a ±5ms section was excised around the spike (thus including the second spike of a burst event), and we cubically interpolated over this segment. As an additional control to prevent artificial biases of spike LFP interaction, we consider spectral contents in the low frequency ranges (<30 Hz). This is to prevent, in particular, improper phase estimation that arises at higher (>30 Hz) frequency ranges (Ardid et al., 2015).

To summarize, we selected those 41 cells that had (1) a minimum of 30 burst events and (2) had at least one second of LFP data around each event. This selection controls for any differences that might arise from changes in firing rate during attentional allocation (**Figure 4-1**), and allows for an estimate of LFP power and phase synchronization using windows without possible transient onset responses to the cue or rotation events of the stimuli. That said, we recognize that this results in low cell numbers. To address this, we use non-parametric significance tests throughout so that the results are not biased by any outliers.

4.4.6 Time-dependent change in burst proportion.

To compute the burst proportion, the time-resolved firing rate of each cell was computed via the peri-stimulus time histogram (PSTH). The PSTH was computed separately for the baseline period (-0.5-0 s.) and attention period from 0-2 s. The PSTH was calculated with variable time windows (50, 100, 200, 300, 400, 700 ms). The choice of window length did not affect the main results. Time points with insufficient trials (n<=20) for an accurate estimate were discarded from analysis.

The burst proportion was defined as the portion of burst spikes relative to all spikes:

$$Burst\ proportion = \frac{burst\ rate}{burst\ rate + nonburst\ rate} \qquad \text{(Eq. 1)}$$

This procedure is intrinsically normalized for changes in the burst/non-burst rate that may occur (1) in time, or (2) across cells. The change in burst proportion during the post-cue attention state relative to the pre-cue baseline was calculated via a burst proportion Attention Index:

$$AI_{burstProp} = \frac{BurstProp_{post} - BurstProp_{pre}}{BurstProp_{post} + BurstProp_{pre}} \qquad \text{(Eq. 2)}$$

This was computed independently for each cell. The non-parametric Wilcoxon sign-rank test was used to determine if the burst proportion increases above baseline. The relationship between $AI_{burstProp}$, non-burst rate, and burst rate with time was determined with the Spearman rank correlation, individually for each cell. These results were pooled to show if there was, on average, a monotonic increase (R>0) or decrease (R<0) in time. The overall population trend (increase or decrease) was determined with a χ^2 test on the proportion of cells that showed either an increase or decrease.

4.4.7 LFP power analysis.

We determined the dominant oscillatory components present in the LFP via spectral decomposition. For each trial, we set data after the time of the stimulus change to zero, thus analyzing activity only within the attention cue period void of possible on-responses to the rotation of the stimulus. We analyzed LFP power over a period [0.2, 2] sec after attention cue onset, thus preventing influences from cue-onset transients. Power was determined by Hanning tapering segments and performing a fast Fourier transform (FFT) over a 2-40 Hz range. We scaled power by the frequency in order to account for 1/f structure of LFP data, and normalized the range of each spectrum by [0 1] for comparison across cells.

To determine peaks in the spectral density plot, we used a peak detection algorithm based on Matlab's findpeaks algorithm. First, we smoothed this plot to prevent the influence of noise. Next, we found peaks that were a minimum of 4 points away from each other, and that were above a threshold defined as 50% of the difference between the maximum and minimum of the individual spectra. This procedure extracted oscillatory components with a peak in the spectral density plot. We corroborated the output of this algorithm via a visual inspection of the spectral plots.

We determined if theta and beta power changed with attention onset (0 – 0.5s.) in those LFPs (n=36) recorded on the same channel as the analyzed cells (n=41). We computed power as

(Eq. 3)

outlined above, and determined the normalized, pairwise change in power for each cell:

$$AI_{power} = \frac{Power_{post} - Power_{pre}}{Power_{post} + Power_{pre}}$$

Significance was assessed with the Wilcoxon sign-rank test.

4.4.8 Spike-triggered spectral analyses.

We analyzed the relationship of burst and nonburst spikes to the local field potential in four different ways, outlined next. An overview of the specific analysis steps for each of these analysis types is schematically presented in **Figure B-2**. For further analysis of burst related power modulation and synchronization, we defined the theta band as 5-10 Hz without including the 4 Hz bin, in order to comply with our criteria that there should be at least one second (i.e. 5 cycles at 5Hz, see below) of LFP data around each spike. Without this, more cells would not meet our criteria and thus reduce the pool of neurons for analysis.

4.4.9 Spike-triggered LFP power.

First, we estimated the spectral content centered around each (burst or nonburst) spike using functions from the fieldtrip toolbox (**Figure B-2B**). For each LFP segment around the spike, we calculated power for frequencies ranging from 5-30 Hz, with 0.5 Hz steps, with an adaptive 5-cycle window per frequency. Signals were transformed with a Hanning taper before FFT. We then averaged the power of individual LFP segments locked to bursts or non-bursts individually. Finally, we normalized the power across frequency and burst vs non-bursts to a range [0 1] (thus preserving relative difference in power across frequencies and spike types). We report on the median power across cells around bursts and non-bursts spikes. As well, we estimated the standard error of the median with a bootstrap procedure (see above). We determined if there was a difference in power depending on the spike type with a non-parametric pairwise Wilcoxon sign rank test.

4.4.10 Analysis of time-resolved spike-triggered LFP power.

We next assessed how bursts relate to theta/beta power aligned to spike onset with a time resolved approach (**Figure B-2C**). We computed power as indicated above, but over a shorter 3-cycle adaptive window. This window was slid from -0.2 to 0.2 sec relative to spike onset with a 1 ms time step. Based on our previous results, we analyzed theta and beta effects separately, by averaging the power in a 5-10 Hz, or a 16-30 Hz band, respectively. To account for differences between cells, we Z-score normalized the time-resolved power (preserving differences in burst vs non-burst aligned power).

We assessed if there was a significant difference in time-resolved power relative to bursts or non-bursts with a Wilcoxon sign rank test. To correct for multiple statistical comparisons, we used a cluster based permutation approach (Maris and Oostenveld, 2007). First, we identified the largest significant cluster mass based on temporally adjacent stretches where $p<0.05$. Next, we shuffled the condition label and cell identity before recalculating the largest significant cluster. We performed this procedure 200 times, and compared the observed cluster against the permutation distribution. We then adjusted observed p-values of the points within a cluster according to the p-value of the cluster permutation test.

4.4.11 Phase synchronization analysis.

In a third analysis, we assessed the degree of phase synchronization (**Figure B-2D**). We first extracted the angle of the individual LFP segments' spectra for each frequency (from 5-30 Hz). Next, we identified cells that showed significant phase synchronized to a preferred phase by computing the Rayleigh statistic across all LFP segments (i.e. for both burst and non-burst spikes). The Rayleigh statistic tests if phases significantly cluster in a mean direction. Cells were considered to synchronize significantly at theta or beta if they had a significant frequency bin in the respective frequency range. We determined if the proportion of BS or NS cells that significantly synchronized was different with a Z-test for proportion, separately for the theta and beta bands.

To assess if bursts locked more strongly than non-bursts to particular phases, we used the Pairwise Phase Consistency (Vinck et al., 2010b). This metric is not spuriously biased by differences in spike numbers. We computed the PPC for cells that synchronized significantly to either the beta or theta phases. To compare burst vs non-burst phase synchronization of the NS and BS cell populations, we averaged the phase consistency in the theta and beta bands, and assessed differences with the Wilcoxon signrank test.

The PPC can take on negative values for low sample numbers, which is uninterpretable (Vinck et al., 2010b). To this end, we converted raw PPC values to an effect size with the equation:

$$Effect\ Size = \frac{1 + 2 * sqrt(PPC)}{1 - 2 * sqrt(PPC)} \qquad \text{(Eq. 4)}$$

This effect size can be interpreted as the relative increase in spike rate at the cell's preferred firing phase. For example, a PPC value of 0.01 corresponds to a 1.5 times greater spike rate at the preferred phase.

We determined the average theta and beta phases at which NS and BS spikes occurred by taking, for each cell, the average phase in the theta and beta bands. We report in the main text the mean and 95% circular confidence interval (CI) for NS/BS cells at theta/ beta, computed using the CircStats toolbox (Berens, 2009). As well, to determine if NS and BS cells tended to fire at similar phases, we used the non-parametric Watson U2 two-sample differences in mean direction (Zar, 2010), using the *watsons_u2_perm_test.m* function found online (http://www.mathworks.com/matlabcentral/fileexchange/43543-watson-s-u2-statistic-based-permutation-test-for-circular-data). Finally, to determine if bursts and nonbursts lock to the same phases, we used the circular median test on the pairwise difference in phases between bursts and nonbursts (Zar, 2010). The null hypothesis was a median phase difference of zero, indicating the same preferred phase of firing for bursts and nonbursts across the population.

4.4.12 Phase-dependent power analysis.

With the fourth analysis approach, we identified the link between spike identity, phase of firing, and LFP power by calculating the phase-of-firing dependent power modulation (**Figure**

B-2E). We began this analysis using the LFP spectra previously computed over a 5 cycle adaptive window (as described above). For each LFP segment, we extracted the phase and power at each frequency. We next determined the preferred phase of firing of each cell (computed over all burst- and non-burst- aligned phases), subtracted this mean phase from the observed phases, and wrapped the transformed phases to the range [-pi pi], to obtain the phase relative to preferred phase of firing. This procedure allows comparison between cells, independent of individual cells' preferred phase of firing.

Next, we binned the LFP power of each segment according to the phase at which they occurred (using 6 equally spaced phase bins). We then averaged the phase-binned power in the frequencies of interest, namely the theta and beta frequency band. This procedure was performed separately for segments around burst and non-burst spikes. Finally, we z-score normalized the phase-binned power across phase bins and spike-identity.

As a first step, we determined if theta/beta was related to the phase at which the spike occurred. To this end, we took the median power in each phase bin across cell types, and fit a cosine of the form:

$$y = A * \cos(x + T) \qquad \text{(Eq. 5)}$$

where *A* is the amplitude of the cosine, and *T* is the phase shift. To determine whether power modulation was significantly different between bursts and non-bursts, we took the difference in power modulation between bursts and non-burst *($Ad = A_{burst} - A_{non\text{-}burst}$).* We also assessed the difference in phase shift, first by converting phases into the time domain, and then taking the difference *($Td = T_{burst} - T_{non\text{-}burst}$).* To assess significance, we randomly shuffled the spike identity and cell identity labels, recomputed the median power per phase bin, and refit the cosine. This procedure was repeated 1000 times to obtain a p- value for both the amplitude modulation, as well as the phase shift. This was done separately for BS and NS cells.

To determine if phase-dependent power depended on our choice of bin size, we recomputed the relevant statistics using bin size of 4, 5, 8, and 9. However, if the bin size were

too small, then the full dynamic range did not emerge. On the other hand, bin size that were too large resulted in unfilled phase bins in many of the cells. We report on results based on a bin size of 6 to try to balance the problems of a limited dynamic range and under-sampling. We note, however, that all results comparing burst vs. nonburst effects were qualitatively the same, regardless of the choice of bin size.

4.4.13 Testing for a relationship between spike-triggered LFP power and proportion of burst firing.

To ascertain if the burst proportion was related to LFP power modulations, we correlated the burst proportion and theta/ beta power. The burst proportion was calculated for a 0-2 s. period after cue onset. The average LFP power was calculated with a 5 cycle adaptive window (see above for more details) around the time of the (burst and non-burst) spikes. Power and burst proportion was Z-transformed across cells. We used a Spearman rank correlation to determine the relationship between them. We ignored outliers, defined as power greater than 5 STD.

We also tested whether the firing rate depended on local changes in power. For each LFP segment, we split the power in each frequency into 4 non-overlapping bins. Then, we computed the local firing rate in a [-0.1 0.1] s. time window locked to the spike. We averaged the local firing rate per power bin, calculating separately the grand average for the theta and beta frequency bands. Finally, to determine if the local firing rate differed as a function of the power bin, we performed a non-parametric Kruskal-Wallis test. This was done separately for theta and beta bands, and NS and BS cells.

4.4.14 Relationship of spike-train statistics and spike-triggered LFP power.

To test whether intrinsic spiking properties of the cell are related to local theta and beta activity we quantified the burstiness of the spike-train patterns of cells using the *local variability* measure (Shinomoto et al., 2009). Local variability describes the type of firing pattern that is most common to a neuron; LV values <1 indicate regular firing neuron, LV values ~1 indicate a Poisson process; LV values >1 indicate an irregular/ burst firing. We correlated LV with average theta/

beta power with the Spearman rank correlation to determine if any power related effects that we see might be related to the cell's propensity to burst.

4.4.15 Assessing changes in spike-triggered LFP activity with attention.

To test if LFP power and phase synchronization increased or decreased with attention, we compared activity in the attention period to the baseline period. To this end, we recomputed all relevant metrics in the baseline, but with an important caveat. Namely, we relaxed the constrain that spikes had to have 1 second of raw LFP data around them, instead replacing samples that fell outside the baseline period [-0.5 - 0]s with the average of the LFP segment. We note, however, that this means we are not comparing 'like-for-like' LFP activity and that by filtering data, one cannot exclude influences on the filtered LFP from outside the selected time period.

We assessed changes in spike-triggered LFP power in the theta and beta band. Thus, we computed spike-triggered power as outlined above in the baseline period. We then took the difference in power at each frequency, and computed AI_{power} (equation 3). We then averaged the result in the theta or beta band, and assessed significance with the Wilcoxon signrank test.

We then asked if the proportion of neurons that phase-locked to either theta or beta differed between the baseline and attention period. We computed the proportion of cells that locked to theta or beta phases in the baseline period, determined by the Rayleigh statistic, as above. We then asked if, for a particular frequency band, the proportion of neurons that showed significant phase-locking changed with attention. This was done separately for the theta and beta band, and for NS and BS cells. Significance was assessed with the Z-test for proportions, as described above.

4.5. Results

We recorded from different subfields of the rostral anterior cingulate cortex (areas 32 and 24) and lateral prefrontal cortex (areas 46, 9, posterior area 8), which we abbreviate as

ACC/PFC, of two rhesus macaques performing a color-cued spatial attention task (**Figure 4-1A**) (Kaping et al., 2011). Monkeys were cued to covertly shift attention to a target stimulus and sustain attention until the target stimulus transiently rotated clockwise or counterclockwise. Both monkeys performed the task well above chance (accuracy for monkey R: 75% STD: 8%; monkey M: 71%; STD: 11%) indicating that attention successfully shifted to the correct peripheral target stimulus following cue onset (Shen et al., 2014). We focus our analysis on correct trials, and restrict analysis to the time immediately following attention cue onset and for the time in the trial before either of the peripheral stimuli transiently changed its motion direction. This period contained attention modulated neurons in all recorded brain areas as described before (Westendorff et al., 2016).

During task performance, we recorded 422 single neurons that were well isolated from background activities (**Figure B-1**). To understand the relationship of burst firing (2 spikes within 5ms) to the local field potential during the attention state we extracted those 41 neurons (9.7%, 41/422) that showed sufficient bursts in the post-cue analysis window (see Methods), and compared the local field potential around burst and nonburst events (for an overview of the applied methods, *see* **Figure B-2***)*. Thus, the following analyses are agnostic to whether cells were characteristically "bursty". These cells were selected if they had at least thirty burst events across trials during a restricted time window in the selective attention state (from 0.5 s. after cue onset and until 0.5 s. before the first motion change of a peripheral stimulus). Thus, cell activity was not influenced by onset responses to the cue or by transient changes of the peripheral stimuli. 39% of these came from monkey M, and the remainder came from monkey R.

4.5.1 Burst firing probability increases following attention cue onset.

On average, bursts constituted 8.8% of all spike events. Across the 41 neurons, burst firing increased following attention cue onset relative to the baseline period (pre-cue, 0.63 $\pm$ 0.08 SE burst/s.; post-cue, 0.70 $\pm$ 0.14 SE burst/s), while the rate of non-burst firing decreased (pre-cue, 7.95 $\pm$ 1.49 SE spike/s.; post-cue, 6.87 $\pm$ 1.47 SE spike/s), rendering enhanced burst firing in ACC/PFC a signature of selective attention states. This was also evident in the temporal evolution

of the burst and non-burst rate after attention cue onset **(Figure 4-1B; *see* Figure 4-2 for examples)**. Relative to the pre-cue baseline time period, the proportion of burst- over non-burst firing significantly increased from ~0.2 s. after attention cue onset, reaching median burst attention index values of 0.1 to 0.2, which corresponds to ~20-50% more bursts compared to nonbursts in the attention period than in the pre-cue period (**Figure 4-1B**). 37 of 41 cells increased the proportion of burst firing during the selective attention state as indexed by Spearman rank correlations, which is a higher proportion than expected by chance (90%; χ^2-test, p<<<0.05). The change in burst proportion was composed of both increases in the rate of burst firing (59%; **Figure 4-1C**; χ^2-test, p=0.08), and decreases in non-burst firing (76%; **Figure 4-1C**; χ^2-test, p=0.001). Neurons showing increased burst proportion following attention cue onset were similarly likely to show decreases, increases or no change in non-burst firing (χ^2-test, p=0.48), suggesting that burst rate is modulated independently of non-burst firing rate **(***see* **Figure 4-2, Figure B-4** for example neurons**).**

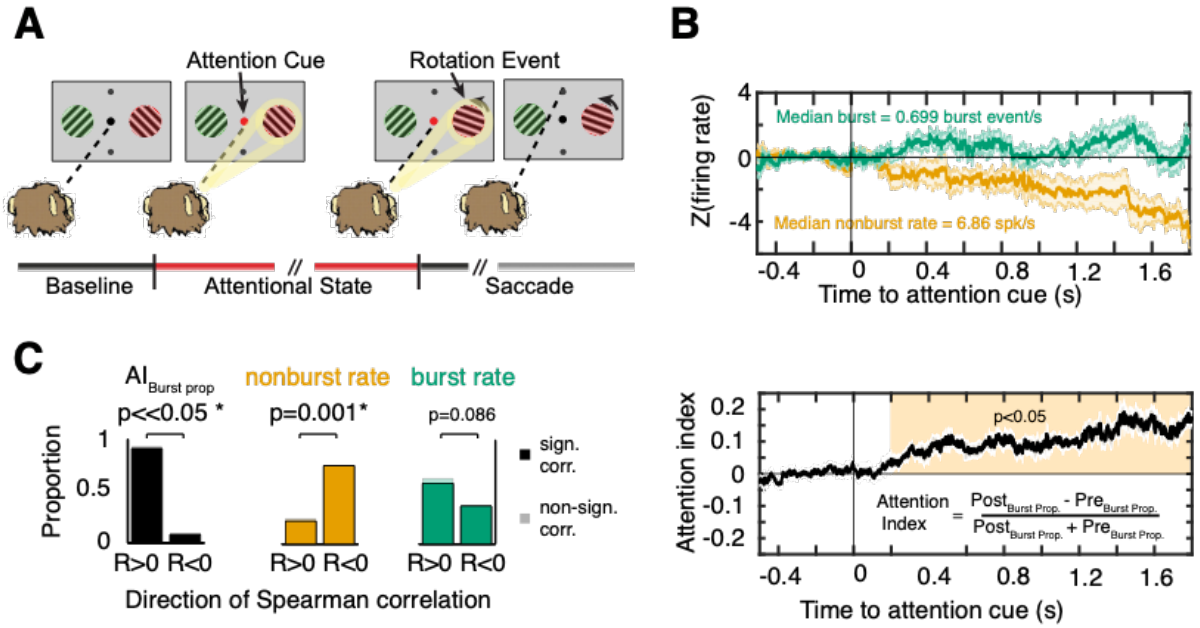

Figure 4-1. Attention task and increased proportion of burst firing following attention cue onset

(A) The task required continued central fixation starting in the Baseline Epoch (black). A color change of the fixation dot instructed to covertly shift attention to the color matching stimulus, marking the onset of an Attentional State (red bar). A rotation event in the cued stimuli had to be discriminated to receive reward by making an up-/downward saccade to clockwise/counterclockwise rotations (gray bar). Rotation events in the uncued stimulus (not shown) had to be ignored. (B) Top panel: Evolution of

baseline-normalized burst rate (green) and non- burst rate (yellow) around the time of attention cue onset (n = 41). Shading denotes standard error. Bottom panel: Burst proportion normalized relative to baseline, calculated as attention index. Gray shading denotes SE from bootstrap procedure and yellow shading shows time period with significant enhanced burst proportions (Wilcoxon sign-rank test, $P < 0.05$). (C) Summary of monotonic trends in time. The proportion of cells that showed a monotonic increase ($R > 0$) or decrease ($R < 0$) when correlating the relevant variable with time. Transparent bars signify cells that did not reach significance individually. Left: 37/41 cells exhibit a monotonic increase in burst proportion after attention cue onset (χ^2 test, $P \ll 0.001$). Middle: For 31/41 cells, the nonburst rate decreased in time (χ^2 test, $P = 0.001$). Right: Burst rate tended to increase in more cells after attention cue onset (n = 24/41, χ^2 test, $P = 0.08$).

Our extracellularly recorded neurons had action potential waveform shapes that reliably distinguished narrow spiking from broad spiking neurons based on their trough-to-peak ratio and their time to repolarization as reported previously (**Figure 4-3A,B**) (Ardid et al., 2015; Oemisch et al., 2015). Among the 41 cells selected for the burst analysis 12 (29.3%), 26 (63.4%), and 3 (7.32%) fell into the categories of narrow, broad, and unclassifiable neurons, respectively. This allowed analyzing burst rate changes separately for different neuron classes. Both BS cells (**Figure 4-3C**) and NS cells (**Figure 4-3D**) significantly increased their burst proportion following attention cue onset. Increased burst proportions were composed of increases in burst firing, as well as decreases in non-burst firing rate (**Figure 4-3C,D**).

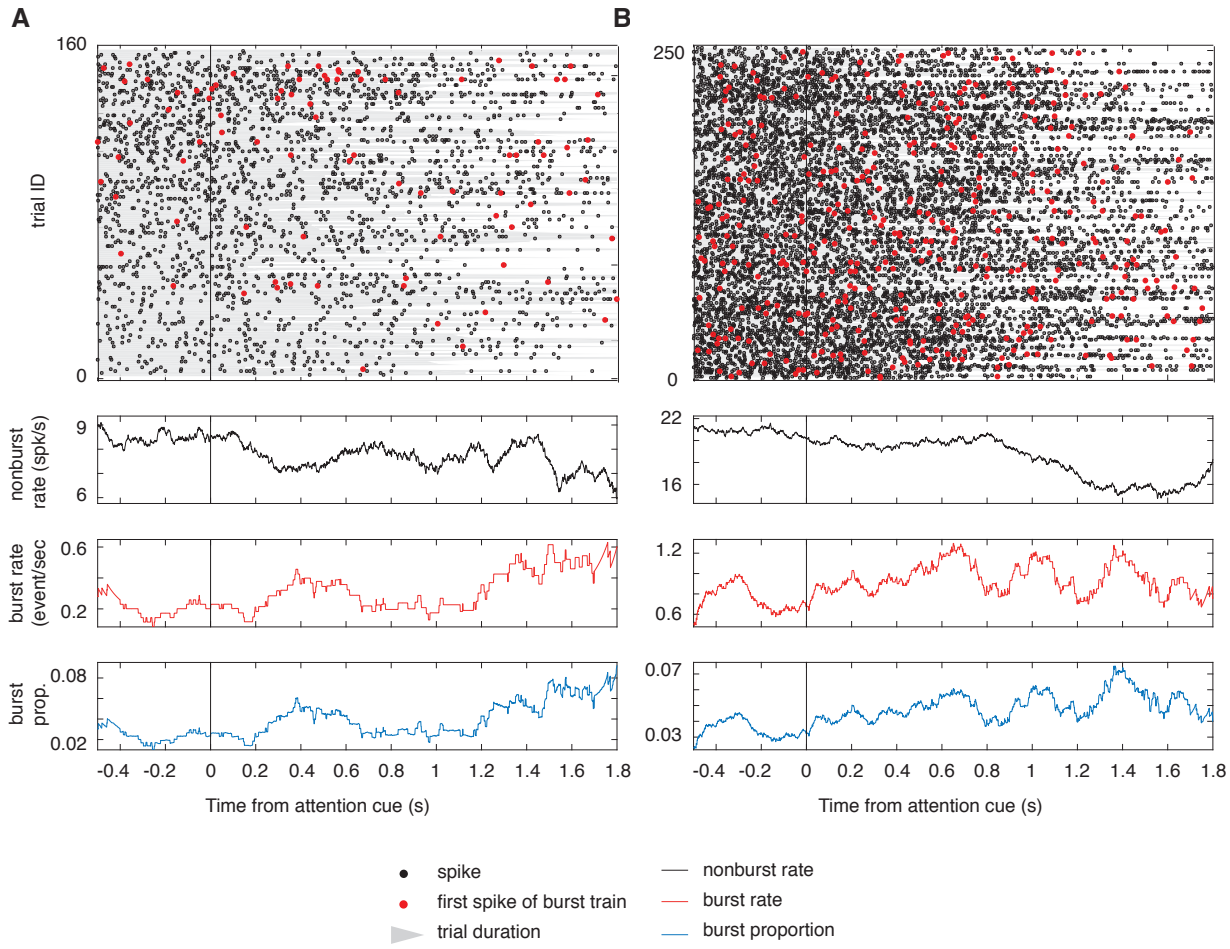

Figure 4-2. Example burst and nonburst spiking locked to attention cue onset

(A and B) (Top) Raster of spiking activity locked to attention cue onset. Black dots denote non- burst spikes, red dots denote the first spike of a burst train. The gray background patch visualizes the duration of each trial, highlighting that trials were of variable length. Bursts did not occur on every trial; however, even in long trials, they were prevalent late in the trial. “Middle top”: Nonburst rate aligned to attention cue onset, calculated with a 0.2 s. sliding window. The nonburst rate decreases with attention cue onset. “Middle bottom”: Burst rate locked to attention cue onset. The burst rate increases with attention cue onset. “Bottom”: The result is that the burst proportion increases with attention cue onset. (B) Same as (A) but for a different neuron. In these examples, the increased burst proportion appears to be driven by a concurrent increase in the burst rate, and a decrease in the nonburst rate (but, see **Figure B-4**)

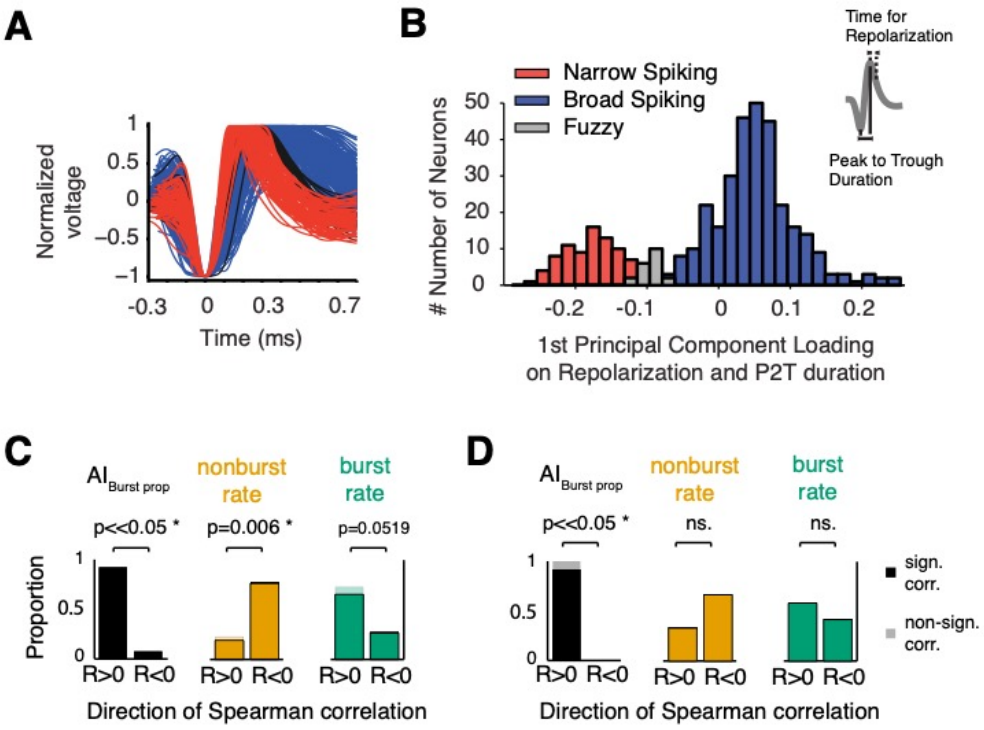

Figure 4-3. Cell-type specific modulation of burst and nonburst rate.

(A) Average normalized action potential waveforms across all recorded narrow-spiking (NS, red), broad-spiking (BS, blue), and unclassified (gray) cells. (B) Bimodal distribution of NS and BS cells (and unreliably classified cells in between) as indexed by PCR score that combines the peak-to-trough duration and time to 25% repolarization of action potential waveforms (see inset). (C and D) Proportion of (C) BS and (D) NS cells with increased ($R > 0$) and decreased ($R < 0$) burst proportion (left), nonburst rate (middle), and burst rate (right) after attention cue onset relative to baseline.

4.5.2 Relation of burst and non-burst spiking events to 5-10 Hz theta and 15-30 Hz beta band activity.

We first asked how the attention-specific burst events relate to local field activity of the neural circuit. To answer this question, we first identified the frequency ranges in the LFP showing the most prominent oscillatory activity during the attention period of the task (0.2 - 2 s.). We found that across the ACC/PFC, 74% (223/301) of the LFP recording sites had at least one clearly discernable power spectral peak in the theta or beta band indicative of periodically coordinated network activity (**Figure 4-4A-D; *see also*, Figure B-3** for examples of raw and normalized LFPs with spectral peaks). Of all recording sites, 25% (75/301) had power spectral densities with peaks within both theta and beta frequency bands, 19% (58/301) of LFPs showed power peaks at the

5-10 Hz theta frequency range without concomitant beta modulation, and 30% (90/301) of LFPs showed only a power peak within the 15-30 Hz beta frequency band (**Figure 4-4D**).

We asked if theta and beta power changed with attention compared to the baseline period. For this analysis, we selected the 0 – 0.5 s. time window after attention cue and compared power to a -0.5 – 0 s. wide time window before the cue. We found that across the LFP channels (n=36) recorded on the same channel as the selected cells, theta activity did not differ significantly before versus after the cue (median AI_{power}= -0.023 $\pm$ 0.014 SE, p = 0.36, consistent with (Voloh et al., 2015)), but beta power was significantly reduced (p = 1.56e-05, median AI_{power}= -0.047 $\pm$ 0.012 SE). Overall, we observed a reduction of theta and beta power over time, but power spectral peaks remained discernable throughout the attention period of the task indicating continued oscillatory activity despite reduced overall power (*see* **Figure B-10**).

The two main frequency ranges (theta and beta) with LFP power modulation were also apparent in spike-triggered LFPs across the 41 neurons recorded with sufficient number of bursts in the attention epoch (**Figure 4-5A**; for an example of burst and non-burst spike triggered LFP see **Figure B-5**). For the theta frequency band, the spike-triggered LFP triggered on the first spike of a burst showed significantly stronger theta power modulation than the spike-triggered LFP average for non-burst spikes (**Figure 4-5B**; Wilcoxon signrank test, p<0.02). There were no power differences for bursts versus non-bursts in the beta frequency range, though the dynamic range showed sufficient variability **(Figure B-6).**

Spike-triggered LFP power modulation so far was calculated in symmetric time windows centered on the time of the burst/non-burst event, leaving unanswered whether finer grained temporal analysis could reveal population level modulation preceding or following the burst/non-burst spiking events. To answer this question, we calculated LFP power in a sliding window analysis using adaptive 3 cycle windows every 1 ms around the time of the burst/non-burst event (see (Paz et al., 2008)). We found that across neurons, theta power was greater around burst spikes than non-burst spikes, starting as early as ~80ms before the burst firing event (**Figure**

4-5C; p<0.05, Wilcoxon signrank test, multiple comparison corrected). There was no difference in beta power modulation for burst/non-burst spikes (**Figure 4-5D**, **Figure B-7A**).

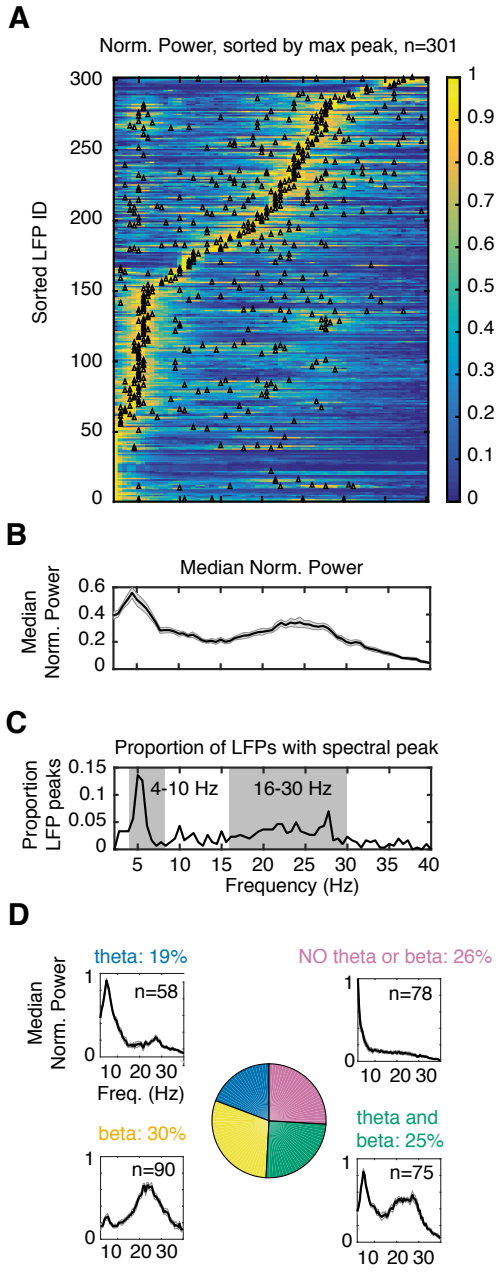

Figure 4-4. Theta and beta frequency components are the most prominent oscillatory signatures in LFP data.

(A) Individual LFP power spectra (n = 301, y-axis), sorted by the frequency of maximum power. Each power spectrum was normalized to account for 1/f noise, and scaled to the range [0 1] to compare across different LFPs. Triangles represent power spectral peaks with more than half-height amplitude (see "Materials and Methods" section). (B) Median LFP power spectrum, revealing peaks in the theta and beta frequency range. (C) The proportion of LFPs that had a spectral peak at each frequency of interest. Gray shading highlights the theta and beta frequency bands. (D) Theta and beta peaks are both evident in 25% of recording sites, theta peaks without beta are evident in 19% of sites, and 30% of sites show only a beta peak without theta component.

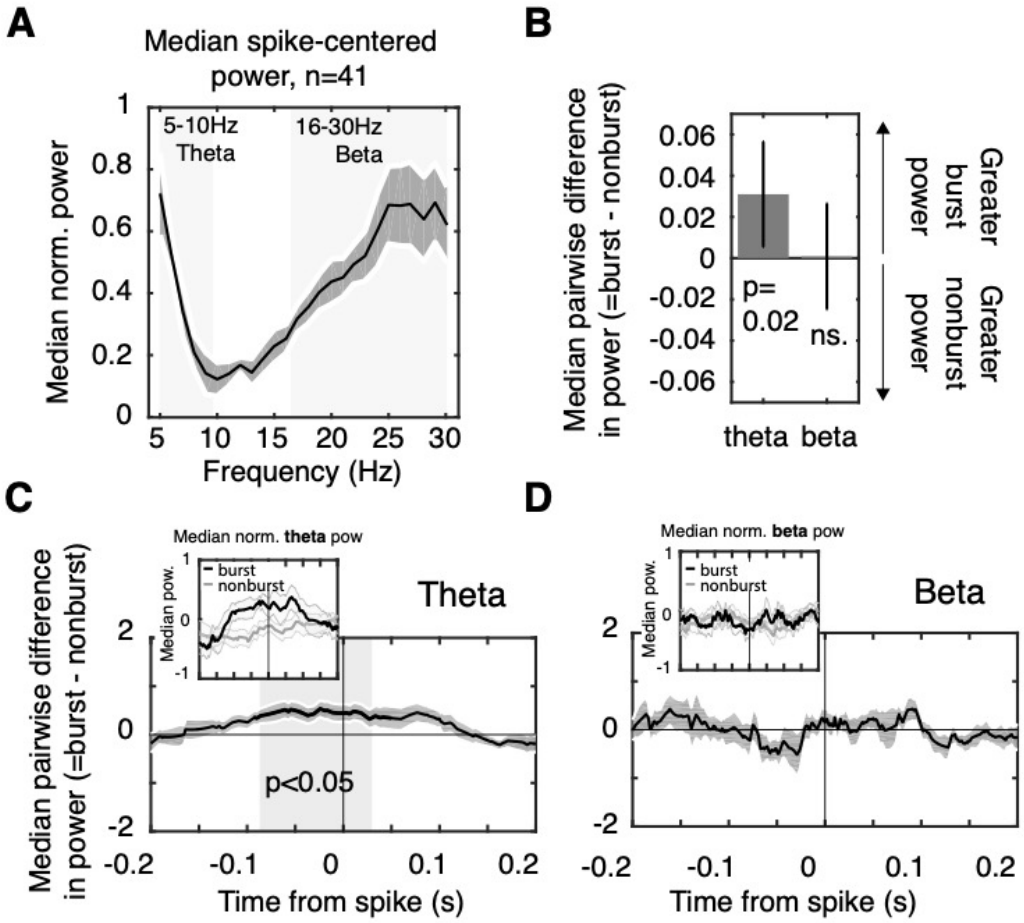

Figure 4-5. Theta and beta LFP oscillations are prevalent around spikes.

(A) Median average LFP spectra around the time of spike occurrences. LFP power was controlled for $1/f^{\alpha}$ noise and normalized to the range [0 1] before averaging. Dark gray patches are the standard error determined with a bootstrap procedure, and the light gray background highlights the theta and beta ranges. (B) Median difference in spike-triggered average LFP power across all (n = 41) cells in the theta and beta frequency band. Bars represent the standard error calculated with a bootstrap procedure. Theta LFP power is greater around bursts as compared to nonbursts, whereas beta power is equivalent. (C) Average pairwise difference in z-score normalized LFP theta power centered on bursts versus nonburst spikes. Significance at $P < 0.05$ (light gray shading) was assessed with a Wilcoxon sign-rank test, and multiple comparison corrected. Insets in the top left shows the time course of LFP power around bursts (black) and nonbursts (gray). Theta power is greater around bursts ~26 ms before spike onset. (D) Same format as (C) for the beta frequency band, showing no average power difference between burst and nonburst spikes.

Enhanced theta power around burst events was visible in the subset of neurons with broad spiking waveforms (**Figure 4-6A**, **Figure B-7B**), but did not reach significance at a α=0.05 level. In contrast, burst events for the subset of NS cells were associated with significantly

enhanced theta power when compared to non-burst spikes in a ~110msec time window starting 26ms prior to burst onset (**Figure 4-6B,** p<0.05, Wilcoxon signrank test, multiple comparison corrected, see also **Figure B-7C**). There was no consistent burst specific power modulation for BS or NS cells at the beta frequency band (**Figure 4-6C**,**D**).

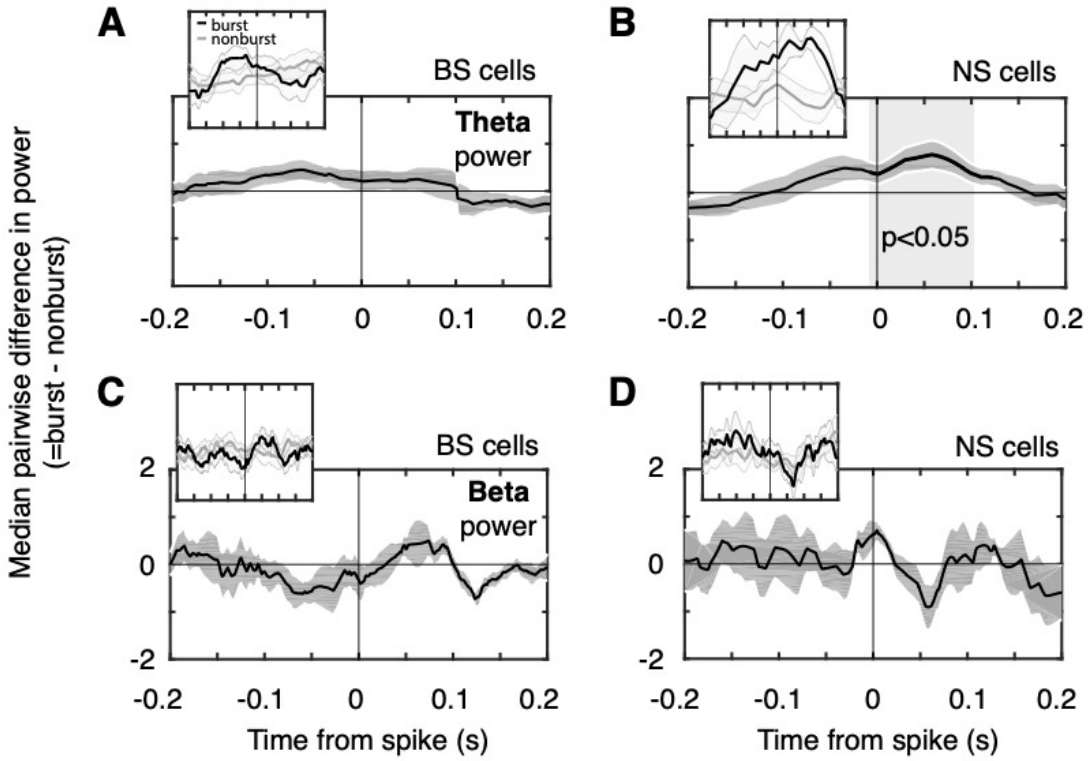

Figure 4-6. Time-resolved LFP power around bursts of NS and BS cells.

(A and B) Time-resolved LFP power differences in the theta frequency and for burst versus non- burst spikes of (A) BS and (B) NS cells. Insets in the top left shows the time course of LFP power around bursts (black) and nonbursts (gray). BS cells (n = 26) showed a nonsignificant increase in theta power before spike onset. NS cells (n = 12) theta power significantly increased ~57 ms after spike onset (shaded light gray area). (C and D) Same as A and B for beta frequency power. There is no change in beta power around bursts relative to nonbursts. Dark gray shadings denote SE computed with a bootstrap procedure.

4.5.3 Burst specific spike-LFP synchronization in the beta frequency band.

Spike-timing specific LFP power modulation indexes the strength of coherent network activity locked to the time of spikes, but does not indicate whether the spike events themselves synchronized to consistent phases of the narrow band LFP activity (Vinck et al., 2011). Across

cells, we observed multiple examples with significant phase synchronization at theta, beta, or both frequencies (**Figure 4-7A**). Across the whole population, both BS and NS neuron classes showed a large proportion of neurons with significantly phase synchronized burst- and non-burst spiking at the theta band (10/12 (83.3%) NS cells, and 14/26 (53.3%) BS cells), and showed a statistically equal likelihood to lock to theta phases (Z-test, p=.08). Similarly, NS and BS cell locked to the beta frequency band (10/12 (83.3%) NS cells, and 16/26 (61.5%) of BS cells) at a similar rate (Z-test, p=0.18). NS and BS cells that showed spike-phase synchronization tended to synchronize to similar phases in the theta band (43.7 ± 90^0 CI, 79.4 ± 90^0 CI, respectively; Watson's U2 test, p=0.9; CI = 95% confidence interval) and the beta band (-138 ± 39.8^0 CI, -144 ± 48.7^0 CI, respectively; Watson's U2 test, p=0.279).

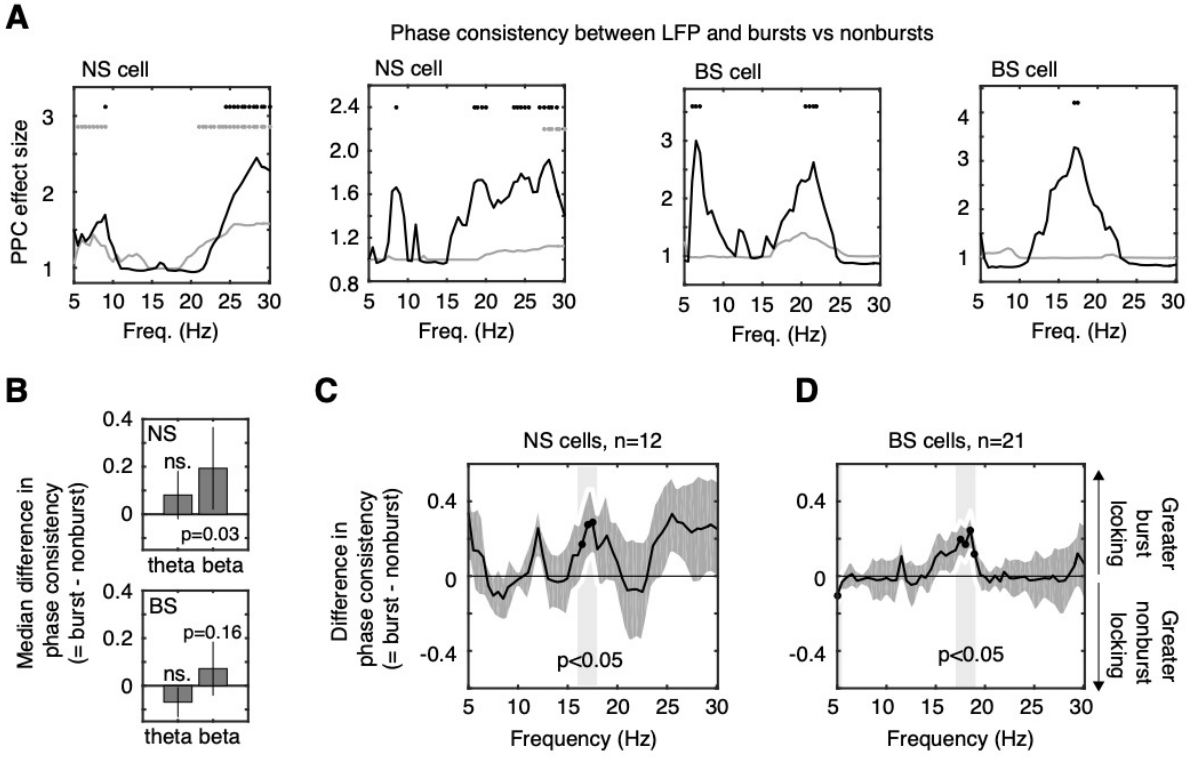

Figure 4-7. Increased phase locking to bursts is specific to the beta frequency band.

(A) Examples of bursts (black) and nonbursts (gray) phase locking to the local activity, determined with the pairwise phase consistency (PPC). Significant locking (Rayleigh test at alpha < 0.05) are marked with the appropriately colored dots. Note that even when nonbursts lock to theta or beta activity, locking to bursts is nonetheless stronger (compare panel 1 from 2). Beta-locked bursts tended to occur in a ~15–20 Hz range, though was also apparent in a higher >20 Hz range in NS cells (panels 1 and 2). (B) Median difference in phase locking in the theta and beta band, for NS cells (top panel) and BS cells (bottom

panel). Burst events of NS cells synchronize stronger to the beta frequency than nonburst spike events. (C and D) Median differences of burst versus nonburst to LFP phase locking for (C) NS cells and (D) BS cells. Note that both NS and BS cells show a narrow beta frequency band with significantly enhanced burst LFP locking over nonburst LFP locking.

For cells that did show significant phase synchronization at beta frequencies, burst spikes showed stronger synchronization than non-burst spikes (**Figure 4-7B**). Across the broader 16-30 Hz beta band bursts of NS cells showed significantly enhanced burst-LFP synchronization compared to non-burst-LFP synchronization (n=10, p = 0.02, Wilcoxon signed rank test). A similar statistical trend was visible for BS cells (n=16, p = 0.08, Wilcoxon signed rank test; **Figure 4-7B**). Finer grained spectral analysis revealed significantly enhanced burst over non-burst phase synchronization in a narrow ~16-18 Hz frequency range for both NS and BS cells (**Figure 4-7C,D**). Moreover, analysis of the pairwise difference in phases of bursts and nonbursts showed both tended to synchronize at similar beta phases, both in NS cells ($-135^0 \pm 39.6$ CI, $-172^0 \pm 47.7$ CI for nonbursts and bursts, respectively; Median test, p=1) and (to a lesser degree) in BS cells ($-130^0 \pm 52.3$ CI, $-175^0 \pm 50.4$ CI for nonbursts and bursts, respectively; Median test, p=0.08) (*see also*, **Figure 4-8G**).

4.5.4 Beta-band burst synchronization shows significant phase-dependent power modulation.

The previous analysis showed that bursts synchronized stronger than non-bursts to beta-rhythmic but not to theta-rhythmic LFP fluctuations, while overall spike-triggered LFP power was stronger for bursts in the theta band, but not in the beta band. This dissociation of phase synchrony and power modulation could entail that LFP power around the time of the burst varied independently of the phase at which the burst and non-burst spike occurred in the theta and beta cycles (Canolty et al., 2012). To test this possibility, we calculated the *phase-dependent power modulation* by grouping burst spikes and non-burst spikes into six non-overlapping LFP phase bins. For each cell with significant phase synchronization, we defined its average spike-LFP phase as the preferred phase of firing, and averaged LFP power for each of the six phase bins across cells (Womelsdorf et al., 2012). This analysis showed that in NS cells, burst spikes showed significantly stronger power modulation by the phase in the beta but not theta frequency band,

when compared to isolated non-burst spikes (**Figure 4-8A-C**; randomization test, p=0.019). In contrast to NS cells, bursts of BS cells show similar strength of phase dependent power modulation than non-burst spikes (randomization test, p=0.43) (**Figure 4-8B,F**). However, maximal beta power around bursts on average significantly led that of nonbursts (randomization test, p=0.024; **Figure 4-8E).** This was evident as a shift of the phase (relative to burst onset) that contained maximal LFP power by 5-10 ms (57.5^0) prior to the cells preferred phase of firing (**Figure 4-8F**). In contrast to these burst specific effects in the beta band, phase dependent power modulation in the theta band was similar for bursts and non-burst and NS/BS cell classes (**Figure 4-8A,D; Figure B-8**).

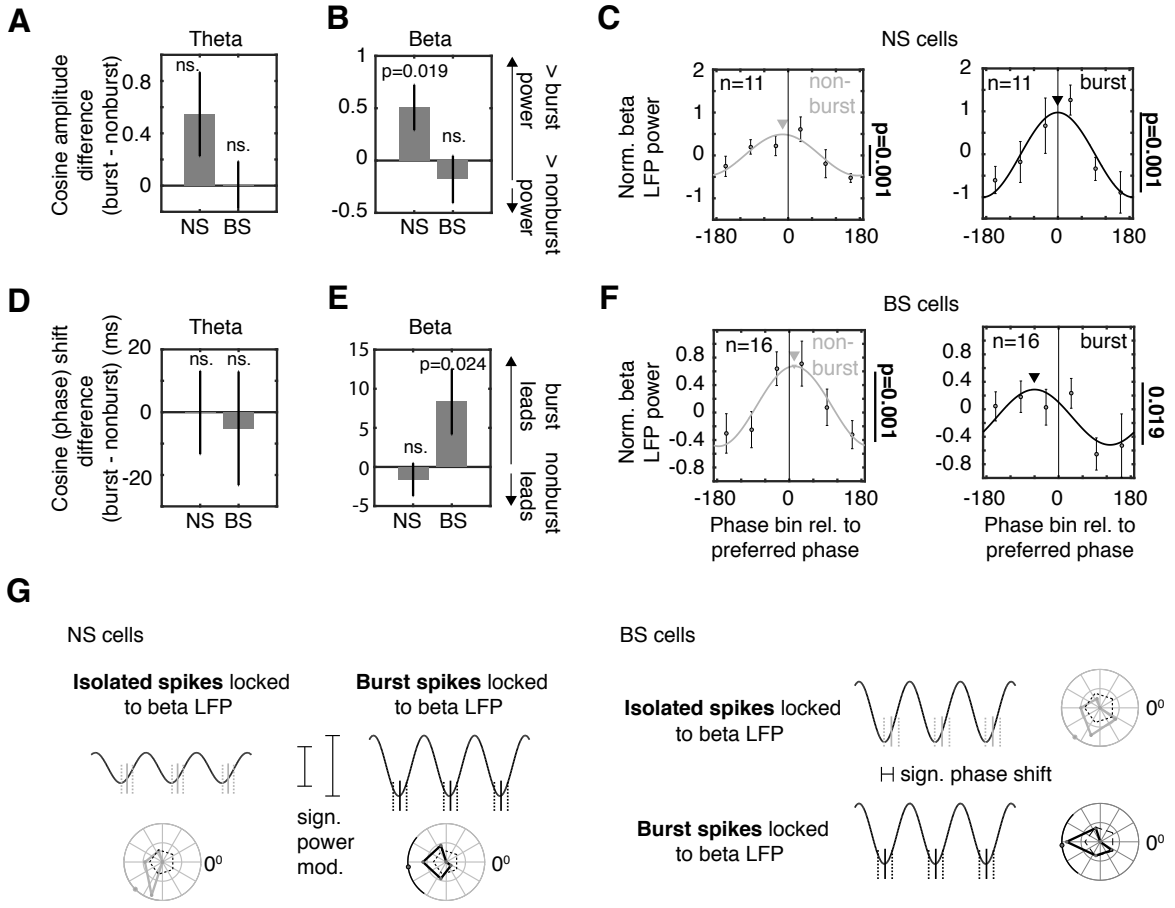

Figure 4-8. Phase-dependent power modulation of burst and nonburst events.

LFP power as a function of the phase of firing of burst and nonbursts, for only those NS and BS cells that showed significant phase locking in the respective frequency band. (A) Theta LFP power is similarly modulated by the phase for burst and nonburst spikes for both NS and BS cells. (B) Beta power is

stronger modulated by the bursts than nonburst LFP phases for NS cells, but not for BS cells. (C) Power (y-axis) significantly varied with the phase (x-axis) of nonburst spikes (left panel) and burst spikes (right panel) of NS neurons (randomization test, $P < 0.05$), relative to the preferred phase of firing. Peak power for spikes synchronizing near their preferred LFP phase. (D) Difference (burst vs. nonburst spike-LFP phases) of the phase (in milliseconds) at which power is maximal within a theta cycle, relative to the preferred phase of firing. The phase of maximal power modulation coincides with the preferred theta phase for both, NS and BS cells. (E) For the beta frequency band, bursts of BS cells that occurred prior to the cells preferred phase are associated with the maximum LFP beta power. The phase difference corresponds to ~5–10 ms. (F) Same as (C), but for BS neurons. Beta power is significantly modulated by the phase (randomization test, $P < 0.05$). Additionally, the burst-LFP phase with maximal power is significantly shifted relative to the preferred phase, preceding the preferred phase by ~57 degrees (randomization test, $P < 0.05$). (G) Summary sketch of beta modulation by burst specific phase of firing. Rose plots indicate the distribution of preferred firing phases—split by cell type and bursts versus singleton spikes—as well as the mean phase and 95% circular confidence intervals. (Left panel) In NS cells, both bursts and nonbursts occur near the same phase, but only the former leads to an increase in local beta power. (Right panel) On the other hand, in BS cells, for the same amount of beta power, bursts occur earlier in the cycle compared to nonbursts.

We corroborated that these results did not depend on the choice of bin size. We recomputed these statistics with a choice of 4, 5, 8, and 9 bins. In general, we found that for both the theta and beta band, in NS and BS cells, all results comparing burst vs nonburst effects were qualitatively the same. That is to say, beta power modulation was greater around bursts in NS cells, while peak beta power occurred on earlier phases during bursts in BS cells. Similarly, there was no difference in power modulation or phase shift in the theta band around bursts vs non-bursts.

In summary, these results show that in the beta frequency band, burst firing is associated with prominent local LFP power similar to the effects in the theta band, but the burst effect at beta depended on the phase at which the burst and non-bursts occur within the beta cycle. We graphically summarized these sets of results in **Figure 4-8G.**

4.5.5 No correlation of power modulation and overall firing rate or burst firing.

Results so far characterized those bursts occurring during the selective attentional state following cue onset (0.5 s. after cue onset and 0.5 s. prior to a change of either target or distractor stimulus). It might thus be possible that burst mediated LFP power modulation, or the burst

phase locking, is particularly prominent in neurons that show relatively larger modulation of burst firing. However, we found that local firing rate did not vary as a function of lower/higher theta or beta power, in either NS or BS cells (Kruskal-Wallis test, all p>0.2). Moreover, correlations of LFP power at theta or beta band did not correlate with burst proportion (**Figure B-9**).

4.5.6 LFP power is not apparently linked to intrinsic neuron properties.

The burst specific association with theta power and beta phases could be the result of neuron specific properties, or they may be better understood as network phenomena that emerge during active states irrespective of the intrinsic propensity of neurons to fire bursts. To address this issue, we computed the relative burstiness of the neurons' spiketrains using the *local variability (LV)* metric that yields higher and lower values the more bursty (*LV>1*) or more regular (*LV<1*) the firing pattern is (Shinomoto et al., 2009). We found that for the group of neurons selected for analysis, NS cells (n=12) exhibited an average LV of 0.751 ± 0.07 SE, and BS cells (n=26) showed 1.09 ± 0.146 SE, which reflect average values for the overall recorded cell population (Ardid et al., 2015), suggesting that the selected BS and NS neurons are not intrinsically bursty neurons. Moreover, intrinsic firing rate variability was not a predictor of the overall degree of spike-triggered oscillatory power in the theta band (Spearman correlation, R=-0.08, p=0.67) or beta band (R=0.10, p=0.529).

4.5.7 Overall spike-triggered LFP activity did not vary with attention demands.

To discern whether the burst specific effects of spike-LFP interactions were indexing an attention specific effect, we compared the post-cue effects to a pre-cue baseline epoch. For these analyses, we relaxed the constraint that each spike must have ±0.5 second of LFP data around it (*see* Methods), in the baseline period only. Instead, we replaced data points that occurred outside of the baseline [-0.5 0] period with the average of the remaining samples. We found that spike triggered theta or beta power on average did not vary from pre-attention cue to post-attention cue (theta band AI_{power} , $-1.5\times10^{-5} \pm 7.9\times10^{-3}$, p=0.66; beta band AI_{power} , $-3.7\times10^{-3} \pm 4.4\times10^{-3}$, p=0.22). Similarly, the proportion of neurons that significantly phase-locked to the LFP during the baseline vs attention period did not differ. For NS cells, 7/12 cells locked to theta in the baseline period, which was not different from the attention state (Z-test for proportion,

p=0.18), whereas 11/12 cells locked to the beta band, the same as in the attention state (p=0.58). The same was true for BS cells locking to theta (12/26, p=0.579), or beta (15/26, p=0.537) phases in the baseline period.

4.6. Discussion

Here we reported that neurons in ACC/PFC increased burst firing proportionally to non-burst firing when nonhuman primates engage in a selective attentional state. Burst firing events were a signature of covert attention for both narrow and broad spiking neurons. During the same attention state, we found that three quarters of recording sites showed prominent local field potential oscillatory peaks at a 4-10 Hz theta and/or a 16-30 Hz beta. Burst firing during the attention state had unique relationships to both theta- and beta-band population level activities. Within the theta frequency band, burst firing coincided with stronger theta power than non-burst firing. This theta effect was evident for bursts of broad spiking neurons, but it was strongest for bursts of narrow spiking neurons. Within the beta frequency band, bursts of narrow spiking neurons were more strongly synchronized to the phases of the beta cycle than isolated spikes of the same neurons. For broad spiking neurons, burst spikes were associated with strong beta power at phases preceding the preferred phase. This result contrasted to non-burst spikes that showed a cosine shaped drop off in power away from the preferred phase (**Figure 4-8F**). In summary, these results identify bursts as major signature of attentional states in nonhuman primate ACC/PFC and reveal burst specific modulation of local circuit field activity at those two oscillatory frequency bands that are closely associated with goal-directed controlled behavior (Phillips et al., 2014; Voloh et al., 2015; Babapoor-Farrokhran et al., 2017).

4.6.1 Increased burst firing in ACC/PFC characterizes attention states and long-range activated networks.

We found that burst firing in ACC/PFC increased shortly after a cue instructed subjects to deploy covert selective attention. The rate of burst firing rate increased at the same time as the firing of non-burst spikes decreased. This pattern of results renders ACC/PFC burst firing a unique characteristic of selective attentional processing states. This finding is consistent with the belief

that dendritic mediated burst firing is a reflection of neurons within circuits participating in recurrent network activity in larger brain networks (Larkum, 2013). According to this hypothesis, burst firing follows from coincident feedforward and feedback type synaptic inputs impinging on peri-somatic and distal dendritic regions of the burst firing neurons (Siegel et al., 2000; Larkum et al., 2004; Larkum, 2013). Recent in-vivo experiments have begun to support this hypothesis of coincident distal and dendritic activation to underlie a unique processing state reflected in burst firing patterns (Manita et al., 2015; Palmer et al., 2016).

During attentive states characterized by large-scale network coordination, enhanced dendritic activation in ACC/PFC could be the consequence of distant cortico-cortical axonal inputs, while perisomatic input could reflect prominent synaptic input from the thalamus and other subcortical sources (Barbas and Zikopoulos, 2007; Miller and Buschman, 2013; Barbas, 2015; Womelsdorf and Everling, 2015). For example, with regard to thalamo-prefrontal connectivity, anatomical studies have illustrated that layer 5 and 6 prefrontal neurons project to mediodorsal thalamic neurons that have widespread frontally directed feedback projections to superficial and middle cortical layers (Zikopoulos and Barbas, 2007; Xiao et al., 2009). This fronto-thalamic projection system thus closes a feedback loop that is associated with transient dendritic activation in superficial and middle frontal cortex layers that likely are accompanied by calcium transients. Some of these dendritic fields likely will be from layer 5 and 6 neurons that were projecting to the thalamus, closing a fast thalamocortical feedback loop (Thomson and Bannister, 2003) that could be contributing to the rhythmic modulation specific to the burst firing events we reported during attentional states.

We believe that our findings add critical support for the hypothesis that enhanced burst firing indexes such an active recurrent network state that underlies actual attentive processing in nonhuman primate long-range attention networks (Womelsdorf and Everling, 2015).

One important caveat of the burst-network hypothesis that needs to be addressed in future studies is that dendritic burst mechanisms have been described at the cellular level

exclusively for pyramidal cells and not interneurons (Larkum et al., 1999). Indeed, for the class of pyramidal cells, fast burst spike events are easier generated in neurons with larger dendritic trees (Mason and Larkman, 1990; Yang et al., 1996; van Ooyen and van Elburg, 2014). In contrast, we report that not only pyramidal cells, but also narrow spiking neurons that are putative interneurons, fire bursts that relate stronger to network states than their isolated spikes. One intriguing possibility to resolve this conundrum is to assume that some classes of interneurons are endowed with a burst firing mechanisms that co-localizes with the main spike mechanism in the soma and is independent from large dendritic trees (Krahe and Gabbiani, 2004). Such co-localized mechanisms exist and are believed to be more likely activated with enhanced barrages of synaptic inputs, characteristic of enhanced network activation (Krahe and Gabbiani, 2004). It might thus be possible that bursts characterize enhanced network activity states across neuron classes. It will be important in future studies to characterize the burst firing neuron types more precisely to infer which sub-classes of interneurons participate in attention specific burst firing.

The functional role of burst firing to network states critically depends on the definition of burst firing. Our previous work has shown that when bursts are defined using longer 20 ms time windows (instead of 5 ms windows), attentional modulation is absent (Womelsdorf et al., 2014a). Different burst time-scales point to different generating mechanisms, and thus perhaps different neuronal types (Chen and Fetz, 2005). Future analyses allowing for different burst definitions on the basis of classified neuronal types.

4.6.2 Putative interneuron bursts, neuronal synchronization and network oscillations.

We found that burst firing of putative interneurons is associated with stronger theta power than isolated interneuron spikes. The burst specific power modulation started shortly before the time of the burst event and lasted for ~100ms after the first burst spike. This finding is significant as it highlights that burst spikes have a unique relation to the strength of theta rhythmic activity of the local network. The acausal nature of this finding - with theta power increasing already shortly before the burst event - indicates not only that bursts are systematically locked to the ongoing low frequency rhythm (Ray and Maunsell, 2011), but that they could play a more active role to sustain, facilitate, or even initiate synchronized oscillations

in the local neural circuit. Such a role is consistent with recent optogenetic in-vitro studies that implicate interneuron bursts to exert a powerful inhibitory synchronization pulse to the surrounded pyramidal cell network (Berger et al., 2010; Hilscher et al., 2017). In particular, the burst firing of single Martinotti cells have been shown to impose compound inhibitory postsynaptic potentials strong enough to silence connected multiple pyramidal cells in the local circuit (Hilscher et al., 2017, **Figure B-6**). In these experiments, the burst-induced synchronized inhibition resets pyramidal cell activity, whose action potentials synchronize during the recovery from inhibition (Hilscher et al., 2017). Intriguingly, optogenetically induced rhythmic, low frequency (<20Hz), inhibitory pulsing of these burst firing interneurons not only initiated de-novo synchronized firing, but sustained rhythmically synchronized activation in the nearby pyramidal cell network (Hilscher et al., 2017). These widespread consequences of interneuron bursts have been documented specifically for Martinotti cells located close to layer 5. Modeling studies of the role of inhibitory bursts strongly support the potential of this class of interneurons to initiate and reset ongoing oscillatory activity by systematically silencing asynchronous pyramidal cell firing due to the temporally extended inhibitory potential of the burst spikes (Sahasranamam et al., 2016). It will be an important task for future studies to characterize more precisely which narrow spiking neurons in extracellular recordings may correspond to the Martinotti, low threshold firing cell type. For example, a previous study has suggested that at least three narrow spiking neurons are distinguishable in extracellular recordings with differences in the cells propensity to synchronize at theta versus beta activity (Ardid et al., 2015). In a similar vein, we believe that with sufficiently dense recording of layer 5 cell activity it will be possible to test whether bursts of interneuron subclasses are directly initiating or maintaining periodic activity at slow (<20Hz) frequencies in prefrontal brain networks during attention states.

4.6.3 Beta-synchronized burst firing may facilitate rapid changes in activation states.

We found that beta power is more strongly modulated during phase synchronized bursts, rather than nonbursts, for putative interneurons (**Figure 4-8B,C**), and that bursts of putative pyramidal cells coincide with strong phase dependent beta power over a shifted phase range of the beta cycle (**Figure 4-8E,F**). These findings reveal a novel link of burst specific firing events of

ACC/PFC neurons to phase synchronization in the beta frequency band and could provide important constraints for models of beta generation. For example, beta synchronized oscillations in the ACC/PFC are often not sustained, but briefly waxing and waning events of about three consecutive beta cycles (Murthy and Fetz, 1996; Feingold et al., 2015; Sherman et al., 2016). These brief beta events have been traced mechanistically to the coincident activation of dendritic and proximal inputs of a network of inhibitory and pyramidal cells (Sherman et al., 2016). One observation of this model is that brief beta events could emerge even if the distal dendritic input is not-rhythmic (Sherman et al., 2016), while the output of the beta generating circuits does carry beta rhythmicity that functionally couples frontal cortex with long-range targets (Cagnan et al., 2015; Feingold et al., 2015). Another observation is that decreasing the variability in the timing of inputs results in greater beta modulation (Sherman et al., 2016). These observations are consistent with two of our findings. First, beta power modulation in the local population surrounding NS cells could be the result of more strongly synchronized, less variable, inputs during bursts, as opposed to nonbursts. Likewise, putative pyramidal cells' bursts and nonburst events result in similarly high beta power modulation but burst phases span a wider phase range in the beta cycle than nonbursts. This result was evident in a leftward shift of the peak of the phase dependent power modulation and could indicate that burst spikes are elicited already at nonoptimal beta phases at a time when overall beta power has reached a sufficient level, as would be expected if it emerges from similarly synchronized inputs arriving at earlier phases (Sherman et al., 2016). Secondly, the burst spike output itself was linked strongly to beta activity in the local circuit as demonstrated here as burst-specific beta phase synchronization - both in putative interneurons and, to a lesser degree, putative pyramidal cells (**Figure 4-7B-D**). This is reminiscent of a previous study, demonstrating strong beta synchronization of bursts in one area to the LFP in other distant brain areas in ACC/PFC (Womelsdorf et al., 2014a). Taken together, the burst specific phase synchronization effects may reflect facilitated interactions of the local circuit with long-range connected areas that show rhythmic activity at a similar frequency, possibly contributing to a more beta synchronized spike output of the local circuit. This scenario predicts that long-range coherent network activity is supported by mechanisms generating burst

firing of single neurons during attention states (Larkum, 2013). This prediction awaits to be tested in future studies.

4.6.4 Burst spikes may actively contribute to the local field.

Beyond a role of burst firing for network activity inferred from power modulation and phase synchrony, burst firing mechanisms may also directly contribute to the local electrical field measured from sharp extracellular electrodes (Sanchez-Vives and McCormick, 2000; Buzsáki et al., 2012; Einevoll et al., 2013). We believe that such a potential direct influence on the LFP should not be understood as a confound, but as an important, possible window into the cellular mechanisms of local circuit formation. In particular, burst firing has been documented in-vivo to affect the extracellular local field by causing post-burst after-hyperpolarization of membranes of the burst firing neurons (discussed in (Buzsáki et al., 2012)). In pyramidal cells such burst induced hyperpolarization of cell membranes is mediated by Ca++ currents that can increase repolarizing K+ conductances in the somatic region (Hotson and Prince, 1980), or by activation of NMDA or Na+ spikes within dendritic compartments, leaving >15ms long traces of hyperpolarization (Nevian et al., 2007; Sjöström et al., 2008). It has been established that NMDA receptors in dendritic spines sense glutamate excitation and have a decay time constant in an estimated range of 10-100ms (Major et al., 2013). This NMDA decay time constant may relate to the finding in (rodent) mPFC slices that localized dendritic glutamate release not only triggers somatic burst firing, but also ≥100ms plateau potentials (Milojkovic et al., 2004).

Importantly, these burst related local dendritic effects of longer lasting after-hyperpolarization may affect the local extracellular field when bursts are coordinated in time (Buzsaki et al., 1988; Sanchez-Vives and McCormick, 2000; Buzsáki et al., 2012). This suggestion is consistent with our result showing that bursts spikes are significantly coordinated with population-level synchronized theta and beta band activities. We thus speculate that one component of the burst triggered LFP average may be attributable to a longer lasting (~100ms) hyperpolarization across burst firing neurons in the neural network. It will be an important question for future work to separate such direct field effects from more indirect interactions with

the field activity, similar to what has been started in visual cortex (Mitzdorf, 1985; Nauhaus et al., 2009; Ray and Maunsell, 2011; Teleńczuk et al., 2017).

4.6.5 Functional implications of burst specific network activity.

Beyond the neurophysiological implications of the reported burst-LFP interactions, our results suggest new avenues of study of a role for bursts in the formation or maintenance of task-relevant networks, already alluded to above. One possibility is that bursts are involved in the selection of relevant but distributed subnetworks. For example, in a sparse coding regime, only a minority of cells actively contribute to the formation of relevant subnetworks (Mizuseki and Buzsáki, 2013). Moreover, oscillatory activity can increase the sparseness of firing by restricting spike times to specific phases of the cycle (Jadi and Sejnowski, 2014). Our results suggest that bursts may be sufficient to modulate theta and beta activity in the local circuit, and may thus play a role in the selection of frequency-dependent subnetworks (Ketz et al., 2015; Womelsdorf and Everling, 2015). In support of this suggestion, we found that putative interneuron burst show stronger synchronization to the LFP at the beta frequency than nonbursts. Such a cell specific increase in burst synchronization might result in stronger pulsing of local circuit neurons. Following previous suggestions (Spitzer and Haegens, 2017), we believe that such an inhibitory based pulsing of local circuit activity could protect it from irrelevant inputs and effectively gate pyramidal cell output to carry relevant information at a beta-specific frequency band (Pesaran et al., 2008; Buschman et al., 2012b; Salazar et al., 2012; Bastos et al., 2014). A related, complementary possibility is that burst firing events are involved in the maintenance of selected subnetworks. This suggestion resonates with a recent study showing that brief bouts of beta activity were prominent during a maintenance phase of a working memory task, at a time and site where single cells did not encode the stimulus (Lundqvist et al., 2016). It might thus be possible that burst firing events relate to population level beta bouts in maintaining relevant information (or protecting it from novel sensory information) by increasing beta prominence (Womelsdorf et al., 2014b). Future studies should directly quantify the information content at burst-modulated sites to ascertain whether they have a primary role in selection or maintenance of band-specific subnetworks (Spitzer and Haegens, 2017).

Recent models also suggest that burst firing neurons connect to each other during synchronous burst firing to enable efficient synaptic strengthening among those neurons receiving similar dendritic inputs at similar times (Sjöström et al., 2008; Kaifosh and Losonczy, 2016; Wilmes et al., 2016). Such coincident activation of neurons switches on long-term potentiation mechanisms and could thus induce plasticity among synapses of those neurons participating in synchronous burst firing. According to this scenario, burst events might have a special role in the formation of networks of neurons participating in the same functional process. Our findings add to this suggestion by showing that burst events may be more strongly modulated by population level 'network level' rhythmic activities than nonburst singleton events. This enhanced burst-LFP relationship was not only evident specifically for bursts during the attentional state in prefrontal and anterior cingulate cortex, but became evident at those frequency bands that have been most prominently related to endogenously controlled, goal directed behaviors (Larkum, 2013; Fries, 2015; Womelsdorf and Everling, 2015). Taken together, we believe that mechanisms underlying burst spike generation will provide a direct window into the origin of cellular control of higher order attentional behaviors.

4.7. Funding

This research was supported by grants from the Canadian Institutes of Health Research (grant # MOP102482); the Natural Sciences and Engineering Research Council of Canada (grant #418430-2012); and the Ontario Ministry of Economic Development and Innovation (grant# ER11-08-140). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

4.8. Acknowledgments:

This research was supported by grants from the Canadian Institutes of Health Research (CIHR), the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Ontario

Ministry of Economic Development and Innovation (MEDI). We thank Stefan Everling for help with the electrophysiological recording and the surgical procedures, Dr. Daniel Kaping for help with the electrophysiological recording and Johanna Stucke for help with animal care. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

4.9. Statement of contribution

BV, Benjamin Voloh; TW, Thilo Womelsdorf

Conceptualization, BV, TW. Methodology, BV, TW; Investigation, TW; Original draft – Writing: BV, TW; Original draft – Editing, BV, TW; Visualization, BV, TW; Supervision, TW

Chapter 5

Measuring feature-based learning in naturalistic environments

5.1. Summary

The work comprising this chapter includes three manuscripts. I will provide a brief summary of the papers here.

To explore rigorously learning of feature-reward contingency in a naturalistic setting, we developed a novel stimulus set that define objects named *Quaddles* (**Appendix C;** (Watson et al., 2018)**;** https://doi.org/10.3758/s13428-018-1097-5). Quaddles are comprised of four feature dimensions that parametrically vary along four feature dimensions (body shape, arm shape, color, pattern). These stimuli are used in multiple experiments in our lab probing learning of changing feature-reward contingencies under varying levels of uncertainty and in response to changing contextual cues. The feature dimensions can be equated to have similar discriminability.

In order to display Quaddles in a naturalistic task context that allowed for 3D rendering of Quaddles, we developed the *Unified Suite for Experiments* (*USE;* **Appendix D;** (Watson et al., 2019)*;* https://doi.org/10.1016/j.jneumeth.2019.108374). USE is implemented in C# in the Unity3D game engine. It employs a hierarchical, state based architecture, thus ensuring that experimental control flows predictably.

Finally, prior studies have shown that learned stimulus values influence gaze behavior (e.g. (Blair et al., 2009; Ghazizadeh et al., 2016), the specific patterns of which evolve to resolve stimulus uncertainty (Yang et al., 2016). However, eye-tracking is prone to noise in suboptimal conditions, such as with heads-free eye tracking or with patient populations. To aid in the robust extraction and classification of gaze behavior in suboptimal conditions, we modified an existing algorithm to make use of robust statistics (**Appendix E)**. This work is submitted to the *Journal of*

Eye Movement Research, and is available as a preprint on *psyarxiv* ((Voloh et al., 2019b); https://doi.org/10.31234/osf.io/rzd6v)

5.2. Statement of Contribution

BV: Benjamin Voloh; MW: Marcus Watson; MN, Milad Naghizadeh; CT, Christopher Thomas; AH, Asif Hasan; SK, Seth Koenig; TW, Thilo Womelsdorf

Quaddles (Watson et al., 2018): Conceptualization (including feature enumeration, and object property desiderata) was made by BV, MW, TW, and MN. Generation scripts were made by MN under supervision of BV, MW, and TW. Discrimination experiment conceptualization was made by BV, MW, and TW. Coding of the experiment was done by BV and MW. Data collection was by BV, MN, and MW. Analysis was performed by MW, MN, and BV. First draft and subsequent revisions were performed by all authors.

USE (Watson et al., 2019): State-based logic, control flow, and conceptualization, MW and BV. Software development of the suite, MW, BV, CT, AH. Development of hardware, CT, with help from BV, MW, and TW. Manuscript, all authors.

MAD saccade (Voloh et al., 2019b): Conceptualization of saccade detection as outlier detection: BV. Analysis and code: BV and SK. Manuscript first draft – BV. Subsequent revisions, all authors.

Chapter 6

General Discussion

To understand how we are able to perform feature-based learning so well, it is critical to understand how pertinent information is represented, transformed, and acted upon. My thesis investigates how phase-of-firing coding may support encoding of learning relevant signals (**Chapter 1, 2**), and sets the groundwork for future work investigating at a micro (intracellular) (**Chapter 3**), and macro (environmental/behavioral) scale (**Chapter 4**).

In **Chapter 1**, I have shown that spikes modulated by the phase of beta oscillation can boost information related to outcome, reward history, and reward prediction error. Importantly, the firing phase and encoding phase were dissociated, particularly for neurons encoding reward prediction errors. In **Chapter 2**, we follow-up on this result to show that the degree of outcome unexpectedness was associated with specific phases of a beta oscillation. The degree and specific phase of this "spike-surprise" signal predicted session-wide learning speed. The observed beta rhythm originates in specific circuit motifs and is supported by differential neuronal dynamics (Wang, 2010; Womelsdorf et al., 2014b; Sherman et al., 2016). In **Chapter 3**, I contrast neuronal burst and nonburst firing, and find that bursts preferentially modulate theta and beta frequency band activity, dependent on the neuronal type. Finally, in **Chapter 4**, we develop a new experimental suite that allows future extension of these insights to naturalistic, dynamic tasks with high-dimensional features. These insights have important ramifications for theories of neural coding, particularly as it pertains to feature-based learning.

6.1. Implications for neural networks supporting feature-based learning

Our findings complement a large literature showing that outcome, outcome history, and prediction errors are encoded across the ACC, LPFC, and STR in neural firing rates (Ito, 2003; Kolling et al., 2016; Shenhav et al., 2016; Hikosaka et al., 2017; Rudebeck et al., 2017; Oemisch et al., 2019; Sajad et al., 2019). What has received less attention is how such information is propagated. Prior studies suggest that firing rates in the neocortex may be relatively weak

candidates for information transfer, particularly in feed-forward networks (Kumar et al., 2010), due to the amplification of random fluctuations in (even weakly) correlated neuronal groups that can drown out the signal of interest (Mazurek and Shadlen, 2002; Kumar et al., 2010). On the other hand, synchronizing spiking activity over small time windows are theoretically a powerful mechanism of affecting post-synaptic change (Azouz and Gray, 2003; Fries, 2015; Voloh and Womelsdorf, 2016; Hahn et al., 2019). Our findings support this suggestion by showing that neurons that carry learning-relevant information also tend to synchronize at a shared 10-25 Hz beta rhythm (**Chapter 1 and 2**).

An important question we could not address is the laminar specificity of these signals. A recent study of SEF – which is a similar agranular cortex as ACC – during a counter-manding task showed that information related to outcome showed a distinct laminar organization (Sajad et al., 2019). Neurons that positively encoded *rewarded* outcome ("Gain" neurons) were mostly found in deep layers, whereas those that positively encoded *unrewarded* outcomes ("Loss" neurons) were more likely found in superficial layers. Similarly, neurons sensitive to negative and positive prediction error were found in superficial and deep layers, respectively (Sajad and Schall, 2019), suggesting that superficial layers in agranular cortex may more generally convey "worse than expected" outcomes, while those in deep layers convey "better than expected" outcomes. An interesting future direction would be to disentangle our results according to the valence of encoding, in order to specify their laminar distribution in ACC and LPFC. Based on the aforementioned studies, we may expect that in ACC, phase-of-firing coding of positive outcomes and prediction errors is prevalent in deep layers, whereas negative outcomes and prediction errors would be located in superficial layers. These may then be communicated to LPFC via deep layer projection neurons terminating in superficial layers of LPFC (Barbas and Pandya, 1989; Barbas, 2015). Support for this suggestion comes from evidence that learning-relevant signals arise first in ACC before being detected in LPFC ((Oemisch et al., 2019; Voloh et al., 2019a); **Chapter 1**), and that functional connectivity, indexed as spike synchronization in the beta band, was more prevalent between spikes in ACC and beta in LPFC rather than the inverse direction

(Voloh et al., 2019a); **Chapter 1**). Further study is necessary to delineate the specific laminar circuitry underlying beta rhythmic information flow between ACC and LPFC (Nácher et al., 2019).

6.2. Implications of phase-of-firing multiplexed information

As described in the introduction, phase-of-firing coding has been observed for object features, object identities, and object categories (Kayser et al., 2009; Siegel et al., 2009; Turesson et al., 2012; Womelsdorf et al., 2012). **Chapter 1 and 2** extend these insights to complex learning variables.

Our study stands in contrast to prior reports that the beta-rhythmic synchronization between distal sites in the LPFC, ACC, or STR encode task relevant variables via the strength of synchrony (Buschman and Miller, 2007; Buschman et al., 2012a; Dean et al., 2012; Salazar et al., 2012; Womelsdorf et al., 2014a; Antzoulatos and Miller, 2016; Babapoor-Farrokhran et al., 2017). For example, beta synchrony is stronger when choosing one visual category over another (Antzoulatos and Miller, 2016), or when maintaining one object representation over another in working memory (Salazar et al., 2012). These studies are in line with the influential "communication through coherence" hypothesis (Fries, 2015), which posits that upstream senders compete for downstream representation by aligning oscillatory activity to the disinhibited phase of downstream regions. However, it remains unclear how such a scheme may operate when items must be multiplexed and transmitted (Akam and Kullmann, 2014). A competing proposition suggests that a larger phase range of the oscillatory cycle may be used, such that different information streams may be temporally segregated according to the phase of the underlying population rhythm (Lisman and Jensen, 2013; Akam and Kullmann, 2014). One resolution to this discrepancy is that both schemes may be employed, as shown in modelling work by (McLelland and VanRullen, 2016). They modelled circuits of excitatory-inhibitory neurons spread across higher and lower level areas, and varied the overall level of inhibition in either the lower or higher area in order to study the segregation of spiking in the higher level area. They found that increasing inhibition in the higher level area led to the selective representation of just one object, whereas as inhibition in the lower area led to temporally

segregated and multiplexed representation of multiple objects. Thus, selective coherence between distal sites could optimize transmission *selectivity*, whereas temporally segregated information could optimize transmission *capacity*. This is in line with our results showing that neurons that synchronize to distal beta also carried information at phases that were dissociated from the synchronizing phase.

That being said, individual neurons are rarely "carriers" that simply seek to faithfully transmit a received message; instead, they react to and transform such messages to form new informational packets (Perkel and Bullock, 1968; Kumar et al., 2010; Hahn et al., 2019). While our study addresses how three separate information streams may be multiplexed at the population level such that a downstream reader may parse them, it leaves open the question of how this information may then be integrated.

Future studies could clarify this in two ways. First, instead of assessing putative function on the basis of firing rates alone (as we have done here), information carried at specific firing phases of individual cells should be assessed. This would help delineate how beta-rhythmic activity facilitates multiplexing and/or integration of different (received) information streams. Second, it would be fruitful to chart out the pattern of connectivity of these separable populations in more detail. This would help clarify how phase-of-firing information is distributed across the network, and help specify how they are disentangled or read out by receiving populations.

6.3. Implications for neuronal mechanisms

We found information related to the degree of outcome unexpectedness was conveyed by the phases at which spikes occurred (**Chapter 1 and 2**). Prediction errors are critical signals that modulate synaptic connections that are effective at maximizing future rewards (Lee et al., 2012; Oemisch et al., 2019). Such synaptic reweighting may be achieved by spike-time dependent plasticity mechanisms (Caporale and Dan, 2008). Our results are consistent with a scenario in which such synaptic plasticity may be achieved by STDP mechanisms that are tuned to firing

phases that are ~27 ms (at 15 Hz) away from the preferred firing phases. Such longer time windows for synaptic plasticity may be supported by dopamine, which can increase the window for STDP efficacy by up to 45 ms (Zhang et al., 2009). This begs the question, however, of how events that are very distant temporally (such as choice and outcome) may be associated via a mechanism that operates in the millisecond range. One possibility is that synapses maintain an "eligibility trace" over many seconds, than can become strengthened or weakened with the action of specific neuromodulators (such as serotonin, dopamine, or norepinephrine), as has been demonstrated in rodent visual and medial prefrontal cortex slices (He et al., 2015). This does not preclude the existence of non-Hebbian plasticity rules that depend on the circuit architecture of specific areas (Suvrathan, 2019).

Beyond synaptic- and site-specific mechanisms that support structural plasticity, emerging evidence suggests that synchronizing network level oscillations can support the induction of transient plasticity. A recent study by Zanos and colleagues showed that stimulation at specific phases of transient beta oscillations induced short lived potentiation or depression in the beta-hosting circuits that outlived the oscillation itself (Zanos et al., 2018). Importantly, they also found similar (albeit weaker) effects at sites that were distal to beta-hosting sites, particularly when the two were more strongly synchronized. In light of these results, **Chapter 1 and 2** suggest that spikes signaling unexpected outcomes may induce transient plasticity in distal circuits specifically when they are synchronized in the beta band. The effect of this would be to stabilize (or degrade) those neuronal assemblies that would lead to reward on future trials (Womelsdorf and Hoffman, 2018).

Critically, this proposal would depend on the existence and continued life of beta-rhythmic activity in the downstream circuit. **Chapter 3** suggests that (at least in ACC and PFC), burst firing in inhibitory cells may play such a role. In support of this, we found that not only were bursts more tightly synchronized to (local) beta activity, but bursts that occurred near the preferred phase of putative inhibitory cells showed greater power. This is predicted by a model of neocortical beta (Sherman et al., 2016), where less variable inputs result in increased beta

amplitudes. Indeed, bursts in even one cell, if properly targeted, could have a powerful compound effect. Bursting of just one Martinotti cells induces a powerful inhibition on local pyramidal cells, which synchronize their output as they rebound from inhibition (Hilscher et al., 2017). Thus, beta-structured output that conveys information (**Chapter 1 and 2**) may recruit cell-specific bursting mechanisms that can sustain beta periodic activity in distal circuits (**Chapter 3**), for the purpose of synaptic re-weighting. In theory, this sustained activity would also allow for a more robust readout of incoming, beta-periodic information. This proposal remains to be tested.

6.4. Limitations of functional interpretations

Our interpretations of the phase-of-firing code **(Chapter 1)** are based in some part on the functional designation assigned on the basis of average firing rate differences in a limited [0.1 0.7] second time window after outcomes were observed. Thus, it may be possible that the functional designation may be further refined in a future analysis considering the full temporal response of these neurons (Richmond, 2009; Lowe and Schall, 2018). For example, we found that some functional clusters already encoded learning-related information before the onset of the actual outcome, suggesting that these neurons may encode some *expectation* of the reward. One future direction, then, is to determine the temporal specificity of phase-of-firing coding during feature-based learning, and contrast it with a rate code.

Beyond the specific temporal profile of activation, the functional interpretation of the identified encoding clusters depends on the tasks that are employed. For example, midbrain dopaminergic responses are consistent with a reward prediction signal in tasks with appetitive outcomes, but respond differently with negative (aversive) rewards (Matsumoto and Hikosaka, 2009), and may instead signal the degree of uncertainty associated with stimuli ((Bromberg-Martin and Hikosaka, 2009; White et al., 2019). Yet others debate whether dopamine neurons serve to convey prediction error information, rather than the motivational value (“incentive salience”) of learned stimuli (Robinson and Berridge, 2001; McClure et al., 2003; Berridge, 2007), or to increase response vigor in order initiate and sustain rewarding actions (Redgrave et al., 1999). Resolution of these questions requires taxing the same neural resources under different

task constraints (Gao and Ganguli, 2015; Pitkow and Angelaki, 2017). To this end, more naturalistic learning experiments would be helpful to clarify these questions, which can be facilitated by the experimental suite we have developed in **Chapter 4**. The suite could be the basis for development of a more diverse task battery, which manipulates both reward valence, as well as levels of uncertainty. This would help further clarify the role of phase-of-firing encoding of prediction errors, outcome, reward history (**Chapter 1**) or surprise and uncertainty (**Chapter 2**).

6.5. Limitations and implications for functional specialization

Another concern is the relatively large anatomical subdivisions used, where I collapse areas 46, 8, and 8a as "lateral prefrontal cortex", and ventral striatum and caudate as "striatum". Yet, as reviewed in the introduction, ACC and LPFC project to different subdivisions of the striatum (ventral striatum vs caudate, roughly). Ultimately, my choice comes down to one of priority; namely, the question of whether phase-of-firing supports encoding of learning-relevant variables. We split our analysis into these three broad subdivision in recognition of their gross functional, physiological, and anatomical differences (reviewed in the **Introduction**), whereas more fine-grained subdivisions would reduce out statistical power due to limited sampling issues. Based on the reviewed patterns of connectivity, it may the case that phase-of-firing coding would be more prevalent between areas with established anatomical connections, such as the ACC and ventral striatum, or the LPFC and caudate.

That said, it is still astonishing that we find similar distributions of learning-relevant information across three major subdivisions (**Chapter 1**). It is particularly surprising given that each area does seem to show some functional specialization (*see* **Introduction**). This likely reflects the highly recurrent nature of cortico-cortico and cortico-striatal connectivity (Womelsdorf and Everling, 2015; Hikosaka et al., 2017; Hunt and Hayden, 2017). This account implies that the necessary computations are repeated and distributed across many regions which become combined to make the ultimate decision (Hunt and Hayden, 2017). The observed functional specialization of these value-based decision making areas may arise as a function of their unique constellation of connectivity and inputs, rather than from unique informational

content (Womelsdorf and Everling, 2015). Connectivity-based specialization is difficult to measure because of the relative under-sampling of neuronal activity relative to the size of a given brain area (Gao and Ganguli, 2015) and the always present danger that observed effects are in fact mediated by unseen third variables (Bastos and Schoffelen, 2016). Though neuronal sampling continues to increase dramatically (Stevenson and Kording, 2011; Gao and Ganguli, 2015), and ever more powerful theoretical tools are being developed to extract complex task correlates from low-dimensional observations (Churchland et al., 2012; Agarwal et al., 2014; Haxby et al., 2014; Gao and Ganguli, 2015), careful consideration of spike latency and phase-of-firing can help track the flow of information across these regions.

6.6. Conclusion

In conclusion, this thesis provides new insights into neural encoding principles that underlie successful learning of rewarded feature values. My work has shown that learning-relevant variables are broadly distributed across key regions implicated in feature-based learning, and importantly, that informational content is enhanced on specific phases of beta-rhythmic population activity. Importantly, this implies that disruption of phase-of-firing coding could lead to impoverished learning performance. Aberrant oscillatory activity is prevalent in many psychiatric disorders, such as addiction, schizophrenia, and depression (Thut et al., 2011; Voytek and Knight, 2015), and pathological over- and/or under-coupling has been suggested to underlie these disorders (Voytek and Knight, 2015). Our work allows for the further possibility that even in the absence of coupling differences, psychiatric symptoms might arise due to the reduced efficacy of a phase-of-firing code. This work provides a valuable path forward in assessing interventional techniques that manipulate or rescue the coding capacity in pathological brain states.

References

Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. Neural Comput 11:91–101.

Agarwal G, Stevenson IH, Berényi A, Mizuseki K, Buzsáki G, Sommer FT (2014) Spatially distributed local fields in the hippocampus encode rat position. Science 344:626–630.

Akam T, Kullmann DM (2010) Oscillations and filtering networks support flexible routing of information. Neuron 67:308–320.

Akam T, Kullmann DM (2014) Oscillatory multiplexing of population codes for selective communication in the mammalian brain. Nat Rev Neurosci 15:111–122.

Amemori K ichi, Amemori S, Gibson DJ, Graybiel AM (2018) Striatal Microstimulation Induces Persistent and Repetitive Negative Decision-Making Predicted by Striatal Beta-Band Oscillation. Neuron 99:829-841.e6.

Antzoulatos EG, Miller EK (2014) Increases in functional connectivity between prefrontal cortex and striatum during category learning. Neuron 83:216–225.

Antzoulatos EG, Miller EK (2016) Synchronous beta rhythms of frontoparietal networks support only behaviorally relevant representations. Elife 5:1–22.

Ardid S, Vinck M, Kaping D, Marquez S, Everling S, Womelsdorf T (2015) Mapping of functionally characterized cell classes onto canonical circuit operations in primate prefrontal cortex. J Neurosci 35:2975–2991.

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci 7:358–366.

Azouz R, Gray CM (2003) Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. Neuron 37:513–523.

Babapoor-Farrokhran S, Vinck M, Womelsdorf T, Everling S (2017) Theta and beta synchrony coordinate frontal eye fields and anterior cingulate cortex during sensorimotor mapping. Nat Commun:13967.

Bach DR, Dolan RJ (2012) Knowing how much you don't know: A neural organization of uncertainty estimates. Nat Rev Neurosci 13:572–586.

Bair W, Koch C (1996) Temporal Precision of Spike Trains in Extrastriate Cortex of the Behaving

Macaque Monkey. Neural Comput 8:1185–1202.

Balcarras M, Ardid S, Kaping D, Everling S, Womelsdorf T (2016) Attentional Selection Can Be Predicted by Reinforcement Learning of Task-relevant Stimulus Features Weighted by Value-independent Stickiness. J Cogn Neurosci 28:333–349.

Ballard I, Miller EM, Piantadosi ST, Goodman ND, McClure SM (2018) Beyond Reward Prediction Errors: Human Striatum Updates Rule Values during Learning. Cereb Cortex 28:3965–3975.

Barbas H (2015) General Cortical and Special Prefrontal Connections: Principles from Structure to Function. Annu Rev Neurosci 38:269–289.

Barbas H, Pandya DN (1989) Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. J Comp Neurol 286:353–375.

Barbas H, Zikopoulos B (2007) The prefrontal cortex and flexible behavior. Neuroscientist 13:532–545.

Bastos AM, Loonis R, Kornblith S, Lundqvist M, Miller EK (2018) Laminar recordings in frontal cortex suggest distinct layers for maintenance and control of working memory. Proc Natl Acad Sci 115:201710323.

Bastos AM, Schoffelen J-M (2016) A Tutorial Review of Functional Connectivity Analysis Methods and Their Interpretational Pitfalls. Front Syst Neurosci 9:1–23.

Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ (2012) Canonical Microcircuits for Predictive Coding. Neuron 76:695–711.

Bastos AM, Vezoli J, Bosman CA, Schoffelen J-MJ-M, Oostenveld R, Dowdall JR, De Weerd P, Kennedy H, Fries P (2014) Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. Neuron 85:390–401.

Bateson G (1972) Form, Substance and Difference. In: Steps to an Ecology of Mind, pp 542. Chicago: University of Chicago Press.

Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. Neuron 47:129–141.

Benchenane K, Peyrache A, Khamassi M, Tierney PL, Gioanni Y, Battaglia FP, Wiener SI (2010) Coherent Theta Oscillations and Reorganization of Spike Timing in the Hippocampal-Prefrontal Network upon Learning. Neuron 66:921–936.

Berens P (2009) CircStat: A MATLAB toolbox for circular statistics. J Stat Softw 31:1–21.

Berger TK, Silberberg G, Perin R, Markram H (2010) Brief bursts self-inhibit and correlate the pyramidal network. PLoS Biol 8.

Berke JD (2008) Uncoordinated firing rate changes of striatal fast-spiking interneurons during behavioral task performance. J Neurosci 28:10075–10080.

Bernacchia A, Seo H, Lee D, Wang X-J (2011) A reservoir of time constants for memory traces in cortical neurons. Nat Neurosci 14:366–372.

Berridge KC (2007) The debate over dopamine's role in reward: The case for incentive salience. Psychopharmacology (Berl) 191:391–431.

Birtoli B, Ulrich D (2004) Firing Mode-Dependent Synaptic Plasticity in Rat Neocortical Pyramidal Neurons. J Neurosci 24:4935–4940.

Bittner KC, Grienberger C, Vaidya SP, Milstein AD, Macklin JJ, Suh J, Tonegawa S, Magee JC (2015) Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. Nat Neurosci 18:1–13.

Bitzenhofer SH, Ahlbeck J, Wolff A, Wiegert JS, Gee CE, Oertner TG, Hanganu-Opatz IL (2017) Layer-specific optogenetic activation of pyramidal neurons causes beta-gamma entrainment of neonatal networks. Nat Commun 8:14563.

Blair MR, Watson MR, Walshe RC, Maj F (2009) Extremely Selective Attention: Eye-Tracking Studies of the Dynamic Allocation of Attention to Stimulus Features in Categorization. J Exp Psychol Learn Mem Cogn 35:1196–1206.

Bollimunta A, Chen Y, Schroeder CE, Ding M (2008) Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques. J Neurosci 28:9976–9988.

Borst A, Theunissen FE (1999) Information theory and neural coding. Nat Neurosci 2:947–957.

Bosman CA, Schoffelen J-M, Brunet N, Oostenveld R, Bastos AM, Womelsdorf T, Rubehn B, Stieglitz T, De Weerd P, Fries P (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. Neuron 75:875–888.

Britten K, Shadlen M, Newsome W, Movshon J (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci 12:4745–4765.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between

behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci 13:87–100.

Bromberg-Martin ES, Hikosaka O (2009) Midbrain dopamine neurons signal preference for advance information about upcoming rewards. Neuron 63:119–126.

Buffalo EA, Fries P, Landman R, Buschman TJ, Desimone R (2011) Laminar differences in gamma and alpha coherence in the ventral stream. Proc Natl Acad Sci U S A 108:11262–11267.

Bukach CM, Gauthier I, Tarr MJ (2006) Beyond faces and modularity: the power of an expertise framework. Trends Cogn Sci 10:159–166.

Buschman TJ, Denovellis EL, Diogo C, Bullock D, Miller EK (2012a) Synchronous Oscillatory Neural Ensembles for Rules in the Prefrontal Cortex. Neuron 76:838–846.

Buschman TJ, Denovellis EL, Diogo C, Bullock D, Miller EK (2012b) Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. Neuron 76:838–846.

Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science (80-) 315:1860–1864.

Buzsáki G (2004) Large-scale recording of neuronal ensembles. Nat Neurosci 7:446–451.

Buzsáki G (2010) Neural syntax: cell assemblies, synapsembles, and readers. Neuron 68:362–385.

Buzsáki G, Anastassiou C a, Koch C (2012) The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. Nat Rev Neurosci 13:407–420.

Buzsaki G, Bickford RG, Ponomareff G, Thal LJ, Mandel R, Gage FH (1988) Nucleus basalis and thalamic control of neocortical activity in the freely moving rat. J Neurosci 8:4007–4026.

Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. Science 304:1926–1929.

Cagnan H, Duff EP, Brown P (2015) The relative phases of basal ganglia activities dynamically shape effective connectivity in Parkinson's disease. Brain 138:1667–1678.

Calabrese E, Badea A, Coe CL, Lubach GR, Shi Y, Styner MA, Johnson GA (2015) A diffusion tensor MRI atlas of the postmortem rhesus macaque brain. Neuroimage 117:408–416.

Calzavara R, Mailly P, Haber SN (2007) Relationship between the corticostriatal terminals from areas 9 and 46, and those from area 8A, dorsal and rostral premotor cortex and area 24c: An anatomical substrate for cognition to action. Eur J Neurosci 26:2005–2024.

Canolty RT, Ganguly K, Carmena JM (2012) Task-dependent changes in cross-level coupling

between single neurons and oscillatory activity in multiscale networks. PLoS Comput Biol 8:e1002809.

Caporale N, Dan Y (2008) Spike timing-dependent plasticity: a Hebbian learning rule. Annu Rev Neurosci 31:25–46.

Cardin JA, Carlén M, Meletis K, Knoblich U, Zhang F, Deisseroth K, Tsai L-H, Moore CI (2009) Driving fast-spiking cells induces gamma rhythm and controls sensory responses. Nature 459:663–667.

Cavanagh JF, Frank MJ, Klein TJ, Allen JJB (2010) Frontal theta links prediction errors to behavioral adaptation in reinforcement learning. Neuroimage 49:3198–3209.

Chang L, Tsao DY (2017) The Code for Facial Identity in the Primate Brain. Cell 169:1013-1028.e14.

Chen D, Fetz EE (2005) Characteristic membrane potential trajectories in primate sensorimotor cortex neurons recorded in vivo. J Neurophysiol 94:2713–2725.

Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy K V. (2012) Neural population dynamics during reaching. Nature 487:51–56.

Cohen MR, Maunsell JHR (2010) A Neuronal Population Measure of Attention Predicts Behavioral Performance on Individual Trials. J Neurosci 30:15241–15253.

Cohen MX (2014) Analyzing Neural Time Series Data: Theory and Practise, 1st ed. MIT Press.

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser M-B, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. Nature 462:353–357.

Comoli E, Coizet V, Boyes J, Bolam JP, Canteras NS, Quirk RH, Overton PG, Redgrave P (2003) A direct projection from superior colliculus to substantia nigra for detecting salient visual events. Nat Neurosci 6:974–980.

Dahlin E, Neely AS, Larsson A, Backman L, Nyberg L (2008) Transfer of Learning After Updating Training Mediated by the Striatum. Science (80-) 320:1510–1512.

David S V., Vinje WE, Gallant JL (2004) Natural stimulus statistics alter the receptive field structure of V1 neurons. J Neurosci 24:6991–7006.

Dean HL, Hagan MA, Pesaran B (2012) Only Coherent Spiking in Posterior Parietal Cortex Coordinates Looking and Reaching. Neuron 73:829–841.

Denker M, Roux S, Lindén H, Diesmann M, Riehle A, Grün S (2011) The local field potential reflects surplus spike synchrony. Cereb Cortex 21:2681–2695.

Dommett E, Coizet V, Blaha CD, Martindale J, Lefebvre V, Walton N, Mayhew JEW, Overton PG, Redgrave P (2005) How visual stimuli activate dopaminergic neurons at short latency. Science (80-) 307:1476–1479.

Efron B, Tibshirani R (1986) Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy. Stat Sci 1:54–75.

Einevoll GT, Kayser C, Logothetis NK, Panzeri S (2013) Modelling and analysis of local field potentials for studying the function of cortical circuits. Nat Rev Neurosci 14:770–785.

Feingold J, Gibson DJ, DePasquale B, Graybiel AM (2015) Bursts of beta oscillation differentiate postperformance activity in the striatum and motor cortex of monkeys performing movement tasks. Proc Natl Acad Sci U S A 112:13687–13692.

Fell J, Axmacher N (2011) The role of phase synchronization in memory processes. Nat Rev Neurosci 12:105–118.

Felsen G, Dan Y (2005) A natural approach to studying vision. Nat Neurosci 8:1643–1646.

Fiorillo CD, Tobler PN, Schultz W (2003) Discrete Coding of Reward Dopamine Neurons. Science (80-) 299:1898–1902.

Fisher NI (1993) Statistical Analysis of Circular Data. Cambridge University Press.

Fouragnan EF, Chau BKH, Folloni D, Kolling N, Verhagen L, Klein-Flügge M, Tankelevitch L, Papageorgiou GK, Aubry JF, Sallet J, Rushworth MFS (2019) The macaque anterior cingulate cortex translates counterfactual choice value into actual behavioral change. Nat Neurosci 22:797–808.

Friedman J, Hastie T, Tibshirani R (2010) Regularization Paths for Generalized Linear Models via Coordinate Descent. J Stat Softw 33:1–22.

Fries P (2015) Rhythms for Cognition: Communication through Coherence. Neuron 88:220–235.

Fries P, Nikolić D, Singer W (2007) The gamma cycle. Trends Neurosci 30:309–316.

Fusi S, Miller EK, Rigotti M (2016) Why neurons mix: high dimensionality for higher cognition. Curr Opin Neurobiol 37:66–74.

Gao P, Ganguli S (2015) On simplicity and complexity in the brave new world of large-scale

neuroscience. Curr Opin Neurobiol 32:148–155.

Gauthier I, Tarr MJ (1997) Becoming a “Greeble” expert: Exploring mechanisms for face recognition. Vision Res 37:1673–1682.

Gawne TJ, Kjaer TW, Richmond BJ (1996) Latency: Another potential code for feature binding in striate cortex. J Neurophysiol 76:1356–1360.

Gelastopoulos A, Whittington MA, Kopell NJ (2019) Parietal low beta rhythm provides a dynamical substrate for a working memory buffer. Proc Natl Acad Sci U S A 116:16613–16620.

Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. Science 233:1416–1419.

Ghazizadeh A, Griggs W, Hikosaka O (2016) Object-finding skill created by repeated reward experience. J Vis 16:17.

Godlove DC, Maier A, Woodman GF, Schall JD (2014) Microcircuitry of agranular frontal cortex: testing the generality of the canonical cortical microcircuit. J Neurosci 34:5355–5369.

Goldman-Rakic PS, Lidow MS, Gallager DW (1990) Overlap of dopaminergic, adrenergic, and serotoninergic receptors and complementarity of their subtypes in primate prefrontal cortex. J Neurosci 10:2125–2138.

Gottlieb J (2012) Attention, Learning, and the Value of Information. Neuron 76:281–295.

Guergiuev J, Lillicrap TP, Richards BA (2016) Towards deep learning with segregated dendrites. arXiv Prepr 1610:1–29.

Guerguiev J, Lillicrap TP, Richards BA (2017) Towards deep learning with segregated dendrites. Elife 6:1–29.

Haber SN (2016) Corticostriatal circuitry. Neurosci 21st Century From Basic to Clin Second Ed:1721–1741.

Haber SN, Fudge JL, McFarland NR (2000) Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. J Neurosci 20:2369–2382.

Haber SN, Knutson B (2010) The reward circuit: Linking primate anatomy and human imaging. Neuropsychopharmacology 35:4–26.

Hahn G, Bujan AF, Frégnac Y, Aertsen A, Kumar A (2014) Communication through Resonance in

Spiking Neuronal Networks. PLoS Comput Biol 10.

Hahn G, Ponce-Alvarez A, Deco G, Aertsen A, Kumar A (2019) Portraits of communication in neuronal networks. Nat Rev Neurosci 20:117–127.

Hartigan JA, Hartigan PM (1985) Change-Points The Dip Test of Unimodality Nonparametric Regression Analysis. Ann Stat 13:70–84.

Hasenstaub A, Shu Y, Haider B, Kraushaar U, Duque A, McCormick DA (2005) Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. Neuron 47:423–435.

Hastie T, Tibishirani R, Friedman J (2009) The Elements of Statistical Learning Data Mining , Inference , and Prediction, 2nd ed.

Hawellek DJ, Wong YT, Pesaran B (2016) Temporal coding of reward-guided choice in the posterior parietal cortex. Proc Natl Acad Sci U S A 113:13492–13497.

Haxby J V, Connolly AC, Guntupalli JS (2014) Decoding Neural Representational Spaces Using Multivariate Pattern Analysis.

Hayden BY, Heilbronner SR, Pearson JM, Platt ML (2011) Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. J Neurosci 31:4178–4187.

Hayden BY, Platt ML (2010) Neurons in Anterior Cingulate Cortex Multiplex Information about Reward and Action. J Neurosci 30:3339–3346.

He K, Huertas M, Hong SZ, Tie XX, Hell JW, Shouval H, Kirkwood A (2015) Distinct Eligibility Traces for LTP and LTD in Cortical Synapses. Neuron 88:528–538.

Heller J, Hertz JA, Kjær TW, Richmond BJ (1995) Information flow and temporal coding in primate pattern vision. J Comput Neurosci 2:175–193.

Hikosaka O, Ghazizadeh A, Griggs W, Amita H (2017) Parallel basal ganglia circuits for decision making. J Neural Transm 125:515–529.

Hikosaka O, Kim HF, Amita H, Yasuda M, Isoda M, Tachibana Y, Yoshida A (2019) Direct and indirect pathways for choosing objects and actions. Eur J Neurosci 49:637–645.

Hikosaka O, Sakamoto M, Usui S (1989) Functional properties of monkey caudate neurons. I. Activities related to saccadic eye movements. J Neurophysiol 61:780–798.

Hilscher MM, Leão RN, Edwards SJ, Leão KE, Kullander K (2017) Chrna2-Martinotti Cells Synchronize Layer 5 Type A Pyramidal Cells via Rebound Excitation Bacci A, ed. PLOS Biol 15:e2001392.

Hirokawa J, Vaughan A, Masset P, Ott T, Kepecs A (2019) Frontal cortex neuron types categorically encode single decision variables. Nature:doi.org/10.1038/s41586-019-1816-9.

Hotson JR, Prince D a (1980) A calcium-activated hyperpolarization follows repetitive firing in hippocampal neurons. J Neurophysiol 43:409–419.

Howe MW, Atallah HE, McCool A, Gibson DJ, Graybiel AM (2011) Habit learning is associated with major shifts in frequencies of oscillatory activity and synchronized spike firing in striatum. Proc Natl Acad Sci U S A 108:16801–16806.

Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. J Physiol 148:574–591.

Hunt LT, Hayden BY (2017) A distributed, hierarchical and recurrent framework for reward-based choice. Nat Rev Neurosci 18:172–182.

Huxter JR, Senior TJ, Allen K, Csicsvari J (2008) Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus. Nat Neurosci 11:587–594.

Hyman JM, Wyble BP, Goyal V, Rossi C a, Hasselmo ME (2003) Stimulation in hippocampal region CA1 in behaving rats yields long-term potentiation when delivered to the peak of theta and long-term depression when delivered to the trough. J Neurosci 23:11725–11731.

Issa EB, Cadieu CF, Dicarlo JJ (2016) Evidence that the ventral stream codes the errors used in hierarchical inference and learning.

Ito S (2003) Performance Monitoring by the Anterior Cingulate Cortex During Saccade Countermanding. Science (80-) 302:120–122.

Izquierdo A, Brigman JL, Radke AK, Rudebeck PH, Holmes A (2017) The neural basis of reversal learning: An updated perspective. Neuroscience 345:12–26.

Jadi MP, Sejnowski TJ (2014) Cortical oscillations arise from contextual interactions that regulate sparse coding. Proc Natl Acad Sci 111:6780–6785.

Johnston K et al. (2007) Top-Down Control-Signal Dynamics in Anterior Cingulate and Prefrontal Cortex Neurons following Task Switching. Neuron 53:453–462.

Kaifosh P, Losonczy A (2016) Mnemonic Functions for Nonlinear Dendritic Integration in Hippocampal Pyramidal Circuits. Neuron 90:622–634.

Kaping D, Vinck M, Hutchison RM, Everling S, Womelsdorf T (2011) Specific Contributions of Ventromedial, Anterior Cingulate, and Lateral Prefrontal Cortex for Attentional Selection and Stimulus Valuation Behrens T, ed. PLoS Biol 9:e1001224.

Kayser C, Ince RAA, Panzeri S (2012) Analysis of slow (theta) oscillations as a potential temporal reference frame for information coding in sensory cortices. PLoS Comput Biol 8:1–13.

Kayser C, Montemurro MA, Logothetis NK, Panzeri S (2009) Spike-phase coding boosts and stabilizes the information carried by spatial and temporal spike patterns 1 . Supplemental Material and Methods. Neuron 61:597–608.

Kennerley S, Behrens T, Wallis J (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. Nat Neurosci 14:1581–1589.

Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the Frontal Lobe Encode the Value of Multiple Decision Variables. J Cogn Neurosci 21:1162–1178.

Kennerley SW, Walton ME, Behrens TEJ, Buckley MJ, Rushworth MFS (2006) Optimal decision making and the anterior cingulate cortex. Nat Neurosci 9:940–947.

Kerkoerle T Van, Self MW, Dagnino B (2014) Alpha and gamma oscillations characterize feedback and feed- forward processing in monkey visual cortex. Proc Natl Acad Sci 111:14332–14341.

Ketz N a., Jensen O, O'Reilly RC (2015) Thalamic pathways underlying prefrontal cortex–medial temporal lobe oscillatory interactions. Trends Neurosci 38:3–12.

Khamechian MB, Kozyrev V, Treue S, Esghaei M, Daliri MR (2019) Routing information flow by separate neural synchrony frequencies allows for "functionally labeled lines" in higher primate cortex. Proc Natl Acad Sci U S A 116:12506–12515.

Kim H, Ährlund-Richter S, Wang X, Deisseroth K, Carlén M (2016) Prefrontal Parvalbumin Neurons in Control of Attention. Cell 164:208–218.

Kim HF, Hikosaka O (2013) Distinct Basal Ganglia Circuits Controlling Behaviors Guided by Flexible and Stable Values. Neuron 79:1001–1010.

Kim T, Thankachan S, McKenna JT, McNally JM, Yang C, Choi JH, Chen L, Kocsis B, Deisseroth K, Strecker RE, Basheer R, Brown RE, McCarley RW (2015) Cortically projecting basal forebrain

parvalbumin neurons regulate cortical gamma band oscillations. Proc Natl Acad Sci 112:201413625.

Kolling N, Wittmann MK, Behrens TEJ, Boorman ED, Mars RB, Rushworth MFS (2016) Value, search, persistence and model updating in anterior cingulate cortex. Nat Neurosci 19:1280–1285.

Kopell N, Whittington MA, Kramer MA (2011) Neuronal assembly dynamics in the beta1 frequency range permits short-term memory. Proc Natl Acad Sci U S A 108:3779–3784.

Kopell NJ, Gritton HJ, Whittington M a, Kramer M a (2014) Perspective Beyond the Connectome : The Dynome. Neuron 83:1319–1328.

Körding KP, König P (2000) A learning rule for dynamic recruitment and decorrelation. Neural Networks 13:1–9.

Krahe R, Gabbiani F (2004) Burst firing in sensory systems. Nat Rev Neurosci 5:13–23.

Kumar A, Rotter S, Aertsen A (2010) Spiking activity propagation in neuronal networks: Reconciling different perspectives on neural coding. Nat Rev Neurosci 11:615–627.

Lansink CS, Goltstein PM, Lankelma J V., Pennartz CMA (2010) Fast-spiking interneurons of the rat ventral striatum: Temporal coordination of activity with principal cells and responsiveness to reward. Eur J Neurosci 32:494–508.

Larkum M (2013) A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. Trends Neurosci 36:141–151.

Larkum ME, Senn W, Lüscher HR (2004) Top-down dendritic input increases the gain of layer 5 pyramidal neurons. Cereb Cortex 14:1059–1070.

Larkum ME, Waters J, Sakmann B, Helmchen F (2007) Dendritic spikes in apical dendrites of neocortical layer 2/3 pyramidal neurons. J Neurosci 27:8999–9008.

Larkum ME, Zhu JJ, Sakmann B (1999) A new cellular mechanism for coupling inputs arriving at different cortical layers. Nature 398:338–341.

Lee D, Seo H, Jung MW (2012) Neural Basis of Reinforcement Learning and Decision Making. Annu Rev Neurosci 35:287–308.

Leong YC, Radulescu A, Daniel R, DeWoskin V, Niv Y, Partners T (2017) Dynamic Interaction between Reinforcement Learning and Attention in Multidimensional Environments. Neuron

93:451–463.

Leventhal DK, Gage GJ, Schmidt R, Pettibone JR, Case AC, Berke JD (2012) Basal ganglia beta oscillations accompany cue utilization. Neuron 73:523–536.

Lisman J, Jensen O (2013) The theta-gamma neural code. Neuron 77:1002–1016.

Lowe KA, Schall JD (2018) Functional categories of visuomotor neurons in macaque frontal eye field. eNeuro 5.

Luczak A, McNaughton BL, Harris KD (2015) Packet-based communication in the cortex. Nat Rev Neurosci 16:745–755.

Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, Miller EK (2016) Gamma and Beta Bursts Underlie Working Memory. Neuron 90:152–164.

Mainen ZF, Sejnowski TJ (1996) Influence of dendritic structure on firing pattern in model neocortical neurons. Nature 382:363–366.

Major G, Larkum ME, Schiller J (2013) Active properties of neocortical pyramidal neuron dendrites. Annu Rev Neurosci 36:1–24.

Manita S, Suzuki T, Homma C, Matsumoto T, Odagawa M, Yamada K, Ota K, Matsubara C, Inutsuka A, Sato M, Ohkura M, Yamanaka A, Yanagawa Y, Nakai J, Hayashi Y, Larkum ME, Murayama M (2015) A Top-Down Cortical Circuit for Accurate Sensory Perception. Neuron 86:1304–1316.

Mante V, Sussillo D, Shenoy K V., Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503:78–84.

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

Marr D (1982) Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. San Francisco, USA: Freeman.

Mason A, Larkman A (1990) Correlations between morphology and electrophysiology of pyramidal neurons in slices of rat visual cortex. II. Electrophysiology. J Neurosci 10:1415–1428.

Matsumoto M, Hikosaka O (2007) Lateral habenula as a source of negative reward signals in dopamine neurons. Nature 447:1111–1115.

Matsumoto M, Hikosaka O (2009) Two types of dopamine neuron distinctly convey positive and negative motivational signals. Nature 459:837–841.

Mazurek ME, Shadlen MN (2002) Limits to the temporal fidelity of cortical spike rate signals. Nat Neurosci 5:463–471.

McCartney H, Johnson AD, Weil ZM, Givens B (2004) Theta reset produces optimal conditions for long-term potentiation. Hippocampus 14:684–687.

McClure SM, Daw ND, Read Montague P (2003) A computational substrate for incentive salience. Trends Neurosci 26:423–428.

McLelland D, VanRullen R (2016) Theta-Gamma Coding Meets Communication-through-Coherence: Neuronal Oscillatory Multiplexing Theories Reconciled. PLoS Comput Biol 12:4–10.

Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. Nature 417:741–746.

Miller EK, Buschman TJ (2013) Cortical circuits for the control of attention. Curr Opin Neurobiol 23:216–222.

Milojkovic BA, Radojicic MS, Goldman-Rakic PS, Antic SD (2004) Burst generation in rat pyramidal neurones by regenerative potentials elicited in a restricted part of the basilar dendritic tree. J Physiol 558:193–211.

Mitzdorf U (1985) Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. Physiol Rev 65:37–100.

Mizuseki K, Buzsáki G (2013) Preconfigured, Skewed Distribution of Firing Rates in the Hippocampus and Entorhinal Cortex. Cell Rep 4:1010–1021.

Montemurro MA, Rasch MJ, Murayama Y, Logothetis NK, Panzeri S (2008) Phase-of-Firing Coding of Natural Visual Stimuli in Primary Visual Cortex. Curr Biol 18:375–380.

Montijn JS, Meijer GT, Lansink CS, Pennartz CMA (2016) Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective. Cell Rep 16:2486–2498.

Murray EA, Rudebeck PH (2018) Specializations for reward-guided decision-making in the primate ventral prefrontal cortex. Nat Rev Neurosci 19:404–417.

Murthy VN, Fetz EE (1996) Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. J Neurophysiol 76:3949–3967.

Nácher V, Hassani SA, Womelsdorf T (2019) Asymmetric effective connectivity between primate anterior cingulate and lateral prefrontal cortex revealed by electrical microstimulation. Brain Struct Funct 224:779–793.

Naud R, Sprekeler H (2018) Sparse bursts optimize information transmission in a multiplexed neural code. Proc Natl Acad Sci 115.

Nauhaus I, Busse L, Carandini M, Ringach DL (2009) Stimulus contrast modulates functional connectivity in visual cortex. Nat Neurosci 12:70–76.

Nevian T, Larkum ME, Polsky A, Schiller J (2007) Properties of basal dendrites of layer 5 pyramidal neurons: a direct patch-clamp recording study. Nat Neurosci 10:206–214.

Ninomiya T, Dougherty K, Godlove DC, Schall JD, Maier A (2015) Microcircuitry of agranular frontal cortex: contrasting laminar connectivity between occipital and frontal areas. J Neurophysiol 113:3242–3255.

Niv Y, Daniel R, Geana A, Gershman SJ, Leong YC, Radulescu A, Wilson RC (2015) Reinforcement learning in multidimensional environments relies on attention mechanisms. J Neurosci 35:8145–8157.

O'Reilly JX, Schüffelgen U, Cuell SF, Behrens TEJ, Mars RB, Rushworth MFS (2013) Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. Proc Natl Acad Sci U S A 110:E3660--9.

Oemisch M, Westendorff S, Azimi M, Hassani SA, Ardid S, Tiesinga P, Womelsdorf T (2019) Feature-specific prediction errors and surprise across macaque fronto-striatal circuits. Nat Commun 10:1–15.

Oemisch M, Westendorff S, Everling S, Womelsdorf T (2015) Interareal Spike-Train Correlations of Anterior Cingulate and Dorsal Prefrontal Cortex during Attention Shifts. J Neurosci 35:13076–13089.

Onorato I, Neuenschwander S, Hoy J, Lima B, Rocha K-S, Broggini AC, Uran C, Spyropoulos G, Womelsdorf T, Fries P, Niell C, Singer W, Vinck M (2019) A distinct class of bursting neurons

with strong gamma synchronization and stimulus selectivity in monkey V1. bioRxiv:583955.

Palmer L, Murayama M, Larkum M (2016) Dendritic integration in vivo. In: Dendrites, pp 399–428. Oxford University Press.

Palmigiano A, Geisel T, Wolf F, Battaglia D (2017) Flexible information routing by transient synchrony. Nat Neurosci 20:1014–1022.

Palminteri S, Khamassi M, Joffily M, Coricelli G (2015) Contextual modulation of value signals in reward and punishment learning. Nat Commun 6:8096.

Pandya DN, Van Hoesen GW, Mesulam MM (1981) Efferent connections of the cingulate gyrus in the rhesus monkey. Exp Brain Res 42:319–330.

Panzeri S, Brunel N, Logothetis NK, Kayser C (2010) Sensory neural codes using multiplexed temporal scales. Trends Neurosci 33:111–120.

Panzeri S, Macke JH, Gross J, Kayser C (2015) Neural population coding: combining insights from microscopic and mass signals. Trends Cogn Sci 19:162–172.

Passingham RE, Wise SP (2012) The Neurobiology of the Prefrontal Cortex: Anatomy, Evolution, and the Origin of Insight. Oxford University Press.

Pastalkova E, Itskov V, Amarasingham A, Buzsáki G (2008) Internally generated cell assembly sequences in the rat hippocampus. Science 321:1322–1327.

Payzan-Lenestour E, Bossaerts P (2011) Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. PLoS Comput Biol 7.

Payzan-LeNestour É, Bossaerts P (2012) Do not bet on the unknown versus try to find out more: Estimation uncertainty and “unexpected uncertainty” both modulate exploration. Front Neurosci 6:1–6.

Payzan-LeNestour E, Dunne S, Bossaerts P, O’Doherty JP (2013) The Neural Representation of Unexpected Uncertainty during Value-Based Decision Making. Neuron 79:191–201.

Paz R, Bauer EP, Paré D (2008) Theta synchronizes the activity of medial prefrontal neurons during learning. Learn Mem 15:524–531.

Perkel DH, Bullock TH (1968) Neural coding: a report based on an NRP work session. Neurosci Res Program Bull 6:219–334.

Pesaran B, Nelson MJ, Andersen R a (2008) Free choice activates a decision circuit between

frontal and parietal cortex. Nature 453:406–409.

Phillips JM, Vinck M, Everling S, Womelsdorf T (2014) A Long-Range Fronto-Parietal 5- to 10-Hz Network Predicts “Top-Down” Controlled Guidance in a Task-Switch Paradigm. Cereb Cortex 24:1996–2008.

Pitkow X, Angelaki DE (2017) Inference in the Brain: Statistics Flowing in Redundant Population Codes. Neuron 94:943–953.

Pouget A, Dayan P, Zemel R, House A (2000) INFORMATION PROCESSING WITH POPULATION CODES. 1.

Ratté S, Hong S, DeSchutter E, Prescott S a. (2013) Impact of neuronal properties on network coding: Roles of spike initiation dynamics and robust synchrony transfer. Neuron 78:758–772.

Ray S, Maunsell JHR (2011) Network Rhythms Influence the Relationship between Spike-Triggered Local Field Potential and Functional Connectivity. J Neurosci 31:12674–12682.

Redgrave P, Gurney K (2006) The short-latency dopamine signal: A role in discovering novel actions? Nat Rev Neurosci 7:967–975.

Redgrave P, Prescott TJ, Gurney K (1999) Is the short-latency dopamine response too short to signal reward error? Trends Neurosci 22:146–151.

Reyes AD (2003) Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. Nat Neurosci 6:593–599.

Richards B a, Xia F, Santoro A, Husse J, Woodin M a, Josselyn S a, Frankland PW (2014) Patterns across multiple memories are identified over time. Nat Neurosci 17:981–986.

Richmond BJ (2009) Information Coding. Encycl Neurosci:137–144.

Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. Nature 497:585–590.

Robinson TE, Berridge KC (2001) Incentive-sensitization and addiction. Addiction 96:103–114.

Roelfsema PR, van Ooyen A (2005) Attention-gated reinforcement learning of internal representations for classification. Neural Comput 17:2176–2214.

Roux L, Buzsáki G (2015) Tasks for inhibitory interneurons in intact brain circuits. Neuropharmacology 88:10–23.

Rudebeck PH, Saunders RC, Lundgren DA, Murray EA (2017) Specialized Representations of Value in the Orbital and Ventrolateral Prefrontal Cortex: Desirability versus Availability of Outcomes. Neuron 95:1208-1220.e5.

Runyan CA, Piasini E, Panzeri S, Harvey CD (2017) Distinct timescales of population coding across cortex. Nature 548:92–96.

Rust NC, Movshon JA (2005) In praise of artifice. Nat Neurosci 8:1647–1650.

Sahasranamam A, Vlachos I, Aertsen A, Kumar A (2016) Dynamical state of the network determines the efficacy of single neuron properties in shaping the network activity. Sci Rep 6:26029.

Sajad A, Godlove DC, Schall JD (2019) Cortical microcircuitry of performance monitoring. Nat Neurosci 22:265–274.

Sajad A, Schall JD (2019) Microcircuitry of agranular cortex : Laminar organization of signals for the feedback related negativity. In: Society for Neuroscience. Chicago, USA.

Salazar RF, Dotson NM, Bressler SL, Gray CM (2012) Content-Specific Fronto-Parietal Synchronization During Visual Working Memory. Science (80-) 338:1097–1100.

Sanchez-Vives M V, McCormick DA (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. Nat Neurosci 3:1027–1034.

Sato TR, Schall JD (2003) Effects of Stimulus-Response Compatibility on Neural Selection in Frontal Eye Field get or distractors fall in the receptive field, but the late phase of the activity reliably differentiates the target from distractors. We have hypothesized that the time. Neuron 38:637–648.

Schall JD (2004) On Building a Bridge Between Brain and Behavior. Annu Rev Psychol 55:23–50.

Schultz W, Dayan P, Montague PR (1997) A Neural Substrate of Prediction and Reward. 275:1593–1600.

Segundo JP, Moore GP, Stensaas LJ, Bullock TH (1963) Sensitivity of neurons in Aplysia to temporal pattern of arriving impulses. J Exp Biol 40:643–667.

Shen C, Ardid S, Kaping D, Westendorff S, Everling S, Womelsdorf T (2014) Anterior Cingulate Cortex Cells Identify Process-Specific Errors of Attentional Control Prior to Transient Prefrontal-Cingulate Inhibition. Cereb Cortex.

Shenhav A, Cohen JD, Botvinick MM (2016) Dorsal anterior cingulate cortex and the value of control. 19:4–6.

Sherman MA, Lee S, Law R, Haegens S, Thorn CA, Hämäläinen MS, Moore CI, Jones SR (2016) Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice. Proc Natl Acad Sci:201604135.

Sheth S a, Mian MKM, Patel SRS, Asaad WFW, Williams ZM, Dougherty DD, Bush G, Eskandar EN (2012) Human dorsal anterior cingulate cortex neurons mediate ongoing behavioural adaptation. Nature 488:2–6.

Shinomoto S et al. (2009) Relating neuronal firing patterns to functional differentiation of cerebral cortex. PLoS Comput Biol 5.

Shumway RH, Stoffer DS (2017) Time Series Analysis and Its Applications. Cham: Springer International Publishing.

Siegel M, Körding KP, König P (2000) Integrating Top-Down and Bottom-Up Sensory Processing by Somato-Dendritic Interactions. J Comput Neurosci 8:161–173.

Siegel M, Warden MR, Miller EK (2009) Phase-dependent neuronal coding of objects in short-term memory. Proc Natl Acad Sci U S A 106:21341–21346.

Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. Annu Rev Neurosci 18:555–586.

Sjöström PJ, Rancz EA, Roth A, Häusser M (2008) Dendritic excitability and synaptic plasticity. Physiol Rev 88:769–840.

Smith AC, Frank LM, Wirth S, Yanike M, Hu D, Kubota Y, Graybiel AM, Suzuki WA, Brown EN (2004) Dynamic Analysis of Learning in Behavioral Experiments. J Neurosci 24:447–461.

Smith AC, Stefani MR, Moghaddam B, Brown EN (2005) Analysis and design of behavioral experiments to characterize population learning. J Neurophysiol 93:1776–1792.

Smith AC, Wirth S, Suzuki WA, Brown EN (2007) Bayesian analysis of interleaved learning and response bias in behavioral experiments. J Neurophysiol 97:2516–2524.

Smith EH, Horga G, Yates MJ, Mikell CB, Banks GP, Pathak YJ, Schevon CA, McKhann GM, Hayden BY, Botvinick MM, Sheth SA (2019) Widespread temporal coding of cognitive control in the human prefrontal cortex. Nat Neurosci 22:1883–1891.

Sohal VS, Zhang F, Yizhar O, Deisseroth K (2009) Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. Nature 459:698–702.

Spitzer B, Haegens S (2017) Beyond the Status Quo: A Role for Beta Oscillations in Endogenous Content (Re-) Activation. Eneuro 4:ENEURO.0170-17.2017.

Stanton PK, Sejnowski TJ (1989) Associative long-term depression in the hippocampus induced by hebbian covariance. Nature 339:215–218.

Stark E, Eichler R, Roux L, Fujisawa S, Rotstein HG, Buzsáki G (2013) Inhibition-Induced theta resonance in cortical circuits. Neuron 80:1263–1276.

Stevenson IH, Kording KP (2011) How advances in neural recording affect data analysis. Nat Neurosci 14:139–142.

Suvrathan A (2019) Beyond STDP — towards diverse and functionally relevant plasticity rules. Curr Opin Neurobiol 54:12–19.

Swadlow H a, Gusev a G (2001) The impact of "bursting" thalamic impulses at a neocortical synapse. Nat Neurosci 4:402–408.

Takakuwa N, Kato R, Redgrave P, Isa T (2017) Emergence of visually-evoked reward expectation signals in dopamine neurons via the superior colliculus in V1 lesioned monkeys. Elife 6:1–17.

Takakuwa N, Redgrave P, Isa T (2018) Cortical visual processing evokes short-latency reward-predicting cue responses in primate midbrain dopamine neurons. Sci Rep 8:1–9.

Tan H, Jenkinson N, Brown P (2014a) Dynamic neural correlates of motor error monitoring and adaptation during trial-to-trial learning. J Neurosci 34:5678–5688.

Tan H, Zavala B, Pogosyan A, Ashkan K, Zrinzo L, Foltynie T, Limousin P, Brown P (2014b) Human subthalamic nucleus in movement error detection and its evaluation during visuomotor adaptation. J Neurosci 34:16744–16754.

Teleńczuk B, Dehghani N, Le Van Quyen M, Cash SS, Halgren E, Hatsopoulos NG, Destexhe A (2017) Local field potentials primarily reflect inhibitory neuron activity in human and monkey cortex. Sci Rep 7:40211.

Thibeault KC, Kutlu MG, Sanders C, Calipari ES (2019) Cell-type and projection-specific dopaminergic encoding of aversive stimuli in addiction. Brain Res 1713:1–15.

Thomson AM, Bannister AP (2003) Interlaminar Connections in the Neocortex. Cereb Cortex 13:5–14.

Thorpe SJ (1990) Spike arrival times: A highly efficient coding scheme for neural networks. Parallel Process neural Syst Comput:91–94.

Thut G, Schyns PG, Gross J (2011) Entrainment of perceptually relevant brain oscillations by non-invasive rhythmic stimulation of the human brain. Front Psychol 2:1–10.

Tibshirani R (1996) Regression Shrinkage and Selection via the Lasso. J R Stat Soc Ser B 58:267–288.

Turesson HK, Logothetis NK, Hoffman KL (2012) Category-selective phase coding in the superior temporal sulcus. Proc Natl Acad Sci U S A 109:19438–19443.

Van Eden CG, Hoorneman EMD, Buijs RM, Matthijssen MAH, Geffard M, Uylings HBM (1987) Immunocytochemical localization of dopamine in the prefrontal cortex of the rat at the light and electron microscopical level. Neuroscience 22:849–862.

van Ooyen A, van Elburg RAJ (2014) Dendritic Size and Topology Influence Burst Firing in Pyramidal Cells. In: The Computing Dendrite, pp 381–395. Springer New York.

Vinck M, Lima B, Womelsdorf T, Oostenveld R, Singer W, Neuenschwander S, Fries P (2010a) Gamma-phase shifting in awake monkey visual cortex. J Neurosci 30:1250–1257.

Vinck M, Oostenveld R, Van Wingerden M, Battaglia F, Pennartz CM a (2011) An improved index of phase-synchronization for electrophysiological data in the presence of volume-conduction, noise and sample-size bias. Neuroimage 55:1548–1565.

Vinck M, van Wingerden M, Womelsdorf T, Fries P, Pennartz CM a (2010b) The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization. Neuroimage 51:112–122.

Vogt BA, Pandya DN (1987) Cingulate cortex of the rhesus monkey: II. Cortical afferents. J Comp Neurol 262:271–289.

Vogt BA, Pandya DN, Rosene DL (1987) Cingulate cortex of the rhesus monkey: I. Cytoarchitecture and thalamic afferents. J Comp Neurol 262:256–270.

Voloh B, Oemisch M, Womelsdorf T (2019a) Phase of Firing Coding of Learning Variables across Prefrontal Cortex, Anterior Cingulate Cortex and Striatum during Feature Learning.

Voloh B, Valiante TA, Everling S, Womelsdorf T (2015) Theta–gamma coordination between anterior cingulate and prefrontal cortex indexes correct attention shifts. Proc Natl Acad Sci U S A 112:8457–8462.

Voloh B, Watson M, Koenig S, Womelsdorf T (2019b) MAD saccade: statistically robust saccade threshold estimation. psyarxiv:2–7.

Voloh B, Womelsdorf T (2016) A Role of Phase-Resetting in Coordinating Large Scale Neural Networks During Attention and Goal-Directed Behavior. Front Syst Neurosci 10:18.

Voloh B, Womelsdorf T (2017) Cell-Type Specific Burst Firing Interacts with Theta and Beta Activity in Prefrontal Cortex During Attention States. Cereb cortex 28:4348–4364.

Voytek B, Knight RT (2015) Dynamic Network Communication as a Unifying Neural Basis for Cognition, Development, Aging, and Disease. Biol Psychiatry 77:1089–1097.

Wallis JD, Kennerley SW (2011) Contrasting reward signals in the orbitofrontal cortex and anterior cingulate cortex. Ann N Y Acad Sci 1239:33–42.

Walton ME, Behrens TEJ, Buckley MJ, Rudebeck PH, Rushworth MFS (2010) Separable Learning Systems in the Macaque Brain and the Role of Orbitofrontal Cortex in Contingent Learning. Neuron 65:927–939.

Wang X (2010) Neuophysiological and Computational Principles of Cortical Rhythms in Cognition. Physiol rev 90:1195–1268.

Waters J (2004) Boosting of Action Potential Backpropagation by Neocortical Network Activity In Vivo. J Neurosci 24:11127–11136.

Watson MR, Voloh B, Naghizadeh M, Womelsdorf T (2018) Quaddles: A multidimensional 3-D object set with parametrically controlled and customizable features. Behav Res Methods 51:2522–2532.

Watson MR, Voloh B, Thomas C, Hasan A, Womelsdorf T (2019) USE: An integrative suite for temporally-precise psychophysical experiments in virtual environments for human, nonhuman, and artificially intelligent agents. J Neurosci Methods 326:108374.

Wespatat V, Tennigkeit F, Singer W (2004) Phase sensitivity of synaptic modifications in oscillating cells of rat visual cortex. J Neurosci 24:9067–9075.

Westendorff S, Kaping D, Everling S, Womelsdorf T (2016) Prefrontal and anterior cingulate

cortex neurons encode attentional targets even when they do not apparently bias behavior. J Neurophysiol 116:796–811.

White JK, Bromberg-Martin ES, Heilbronner SR, Zhang K, Pai J, Haber SN, Monosov IE (2019) A neural network for information seeking. Nat Commun 10:5168.

Wilmes KA, Sprekeler H, Schreiber S (2016) Inhibition as a Binary Switch for Excitatory Plasticity in Pyramidal Neurons. PLoS Comput Biol 12:1–27.

Womelsdorf T, Ardid S, Everling S, Valiante TA (2014a) Burst Firing Synchronizes Prefrontal and Anterior Cingulate Cortex during Attentional Control. Curr Biol 24:2613–2621.

Womelsdorf T, Everling S (2015) Long-Range Attention Networks: Circuit Motifs Underlying Endogenously Controlled Stimulus Selection. Trends Neurosci.

Womelsdorf T, Fries P (2007) The role of neuronal synchronization in selective attention. Curr Opin Neurobiol 17:154–160.

Womelsdorf T, Hoffman K (2018) Latent Connectivity: Neuronal Oscillations Can Be Leveraged for Transient Plasticity. Curr Biol 28:R879–R882.

Womelsdorf T, Johnston K, Vinck M, Everling S (2010a) Theta-activity in anterior cingulate cortex predicts task rules and their adjustments following errors. Proc Natl Acad Sci 107:5248–5253.

Womelsdorf T, Lima B, Vinck M, Oostenveld R, Singer W, Neuenschwander S, Fries P (2012) Orientation selectivity and noise correlation in awake monkey area V1 are modulated by the gamma cycle. Proc Natl Acad Sci U S A 109:4302–4307.

Womelsdorf T, Schoffelen J-M, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P (2007) Modulation of neuronal interactions through neuronal synchronization. Science 316:1609–1612.

Womelsdorf T, Valiante T, Sahin NT, Miller KJ, Tiesinga P (2014b) Dynamic circuit motifs underlying rhythmic gain control, gating and integration. Nat Neurosci 17:1031–1039.

Womelsdorf T, Vinck M, Leung LS, Everling S (2010b) Selective theta-synchronization of choice-relevant information subserves goal-directed behavior. Front Hum Neurosci 4:210.

Wurtz RH (2009) Recounting the impact of Hubel and Wiesel. J Physiol 587:2817–2823.

Xiao D, Zikopoulos B, Barbas H (2009) Laminar and modular organization of prefrontal projections

to multiple thalamic nuclei. Neuroscience 161:1067–1081.

Yamada H, Inokawa H, Matsumoto N, Ueda Y, Kimura M (2011) Neuronal basis for evaluating selected action in the primate striatum. Eur J Neurosci 34:489–506.

Yang CR, Seamans JK, Gorelova N (1996) Electrophysiological and morphological properties of layers V-VI principal pyramidal cells in rat prefrontal cortex in vitro. J Neurosci 16:1904–1921.

Yang SCH, Lengyel M, Wolpert DM (2016) Active sensing in the categorization of visual patterns. Elife 5:1–22.

Yeterian EH, Pandya DN, Tomaiuolo F, Petrides M (2012) The cortical connectivity of the prefrontal cortex in the monkey brain. Cortex 48:58–81.

Yu AJ, Dayan P (2005) Uncertainty, neuromodulation, and attention. Neuron 46:681–692.

Zanos S, Rembado I, Chen D, Fetz EE (2018) Phase-Locked Stimulation during Cortical Beta Oscillations Produces Bidirectional Synaptic Plasticity in Awake Monkeys. Curr Biol 28:2515-2526.e4.

Zanos TPT, Mineault PPJ, Pack CC (2011) Removal of spurious correlations between spikes and local field potentials. J Neurophysiol 105:474–486.

Zar JH (2010) Biostatistical Analysis, 5th ed. Prentice Hall.

Zhang JC, Lau PM, Bi GQ (2009) Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses. Proc Natl Acad Sci U S A 106:13028–13033.

Zikopoulos B, Barbas H (2007) Parallel driving and modulatory pathways link the prefrontal cortex and thalamus. PLoS One 2.

Appendices

Appendix A. Supplemental information for Chapter 1

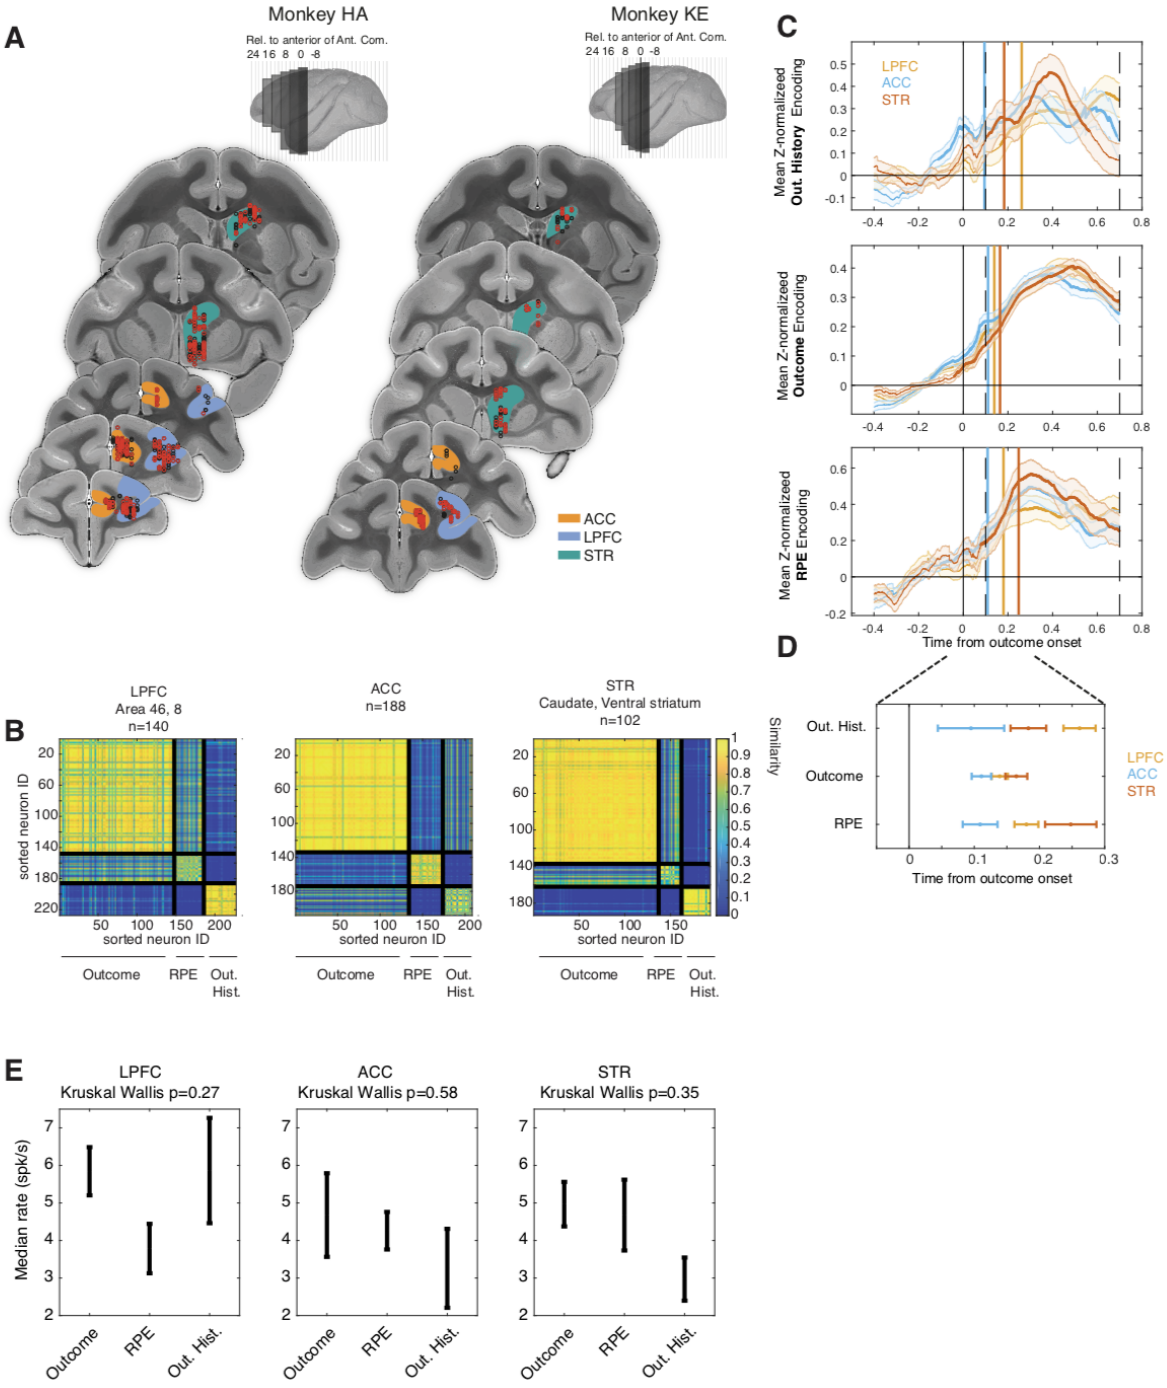

Figure A-1. Recording sites and functional clustering.

(A) All recording units used in the analysis. Units were collapsed across the anterior-posterior axis into equally spaced bins for each monkey. (*top*) Slice bin limits are visualized in the top, lateral view. (*bottom*) Units depicted on a representative atlas slice . Red dots represent encoding cells, and black represents non-encoding cells. Colored outlines correspond to the ACC (orange), STR (green), and LPFC (blue). Encoding cells were broadly distributed across the fronto-striatal axis. (**B**) Three encoding clusters emerge across the fronto-striatal axis. Pairwise similarity between pairs of neurons for each region. Neurons were sorted such that those in the same cluster were adjacent. (**C**) Time resolved, z-score normalized encoding metrics relative to the [-0.4 0] pre-outcome period, separated for *Outcome History* clusters (top), *Outcome* clusters (middle), and *RPE* clusters (bottom), and for LPFC (yellow), ACC (blue) and STR (red). At each time point, we assessed if encoding is above the baseline period (Wilcoxon signrank test). The bolded lines represent the largest contiguous mass where encoding was above baseline. This region represents the time-of-interest over which the latency was calculated for each individual cell. Latency was defined as the point at which 10% area-under-the-curve for the TOI was reached. Vertical lines depict the median latency for each cluster. **(D)** Median and standard deviation for each cluster of cells. All clusters showed significant encoding after the outcome onset. **(E)** Median and standard error of firing rate of encoding clusters for each of three regions. Fire rate differences were similar within all clusters (Kruskal Wallis, p>0.05).

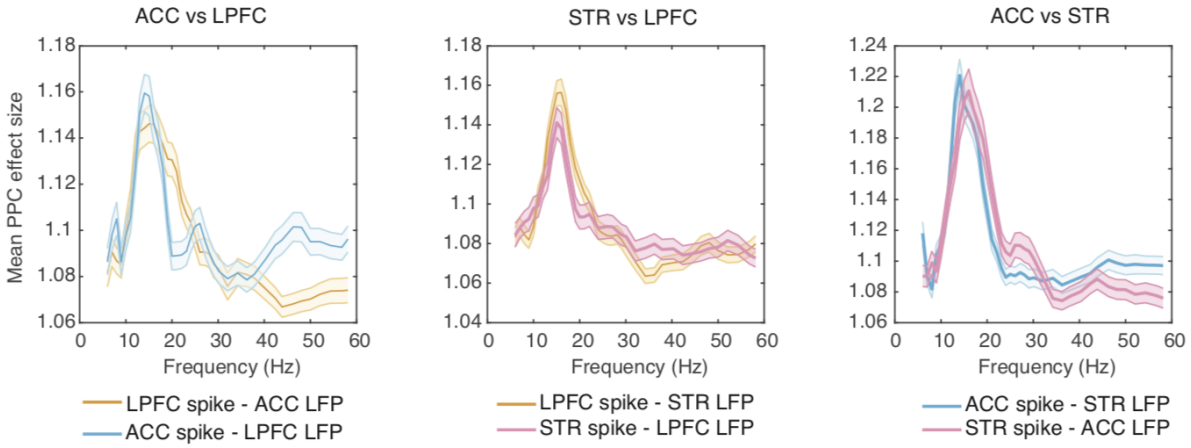

Figure A-2. Inter-areal synchronization

Mean and standard deviation of spike-LFP phase synchronization (PPC) between ACC and LPFC (left), STR and LPFC (middle) and between ACC and STR (right).

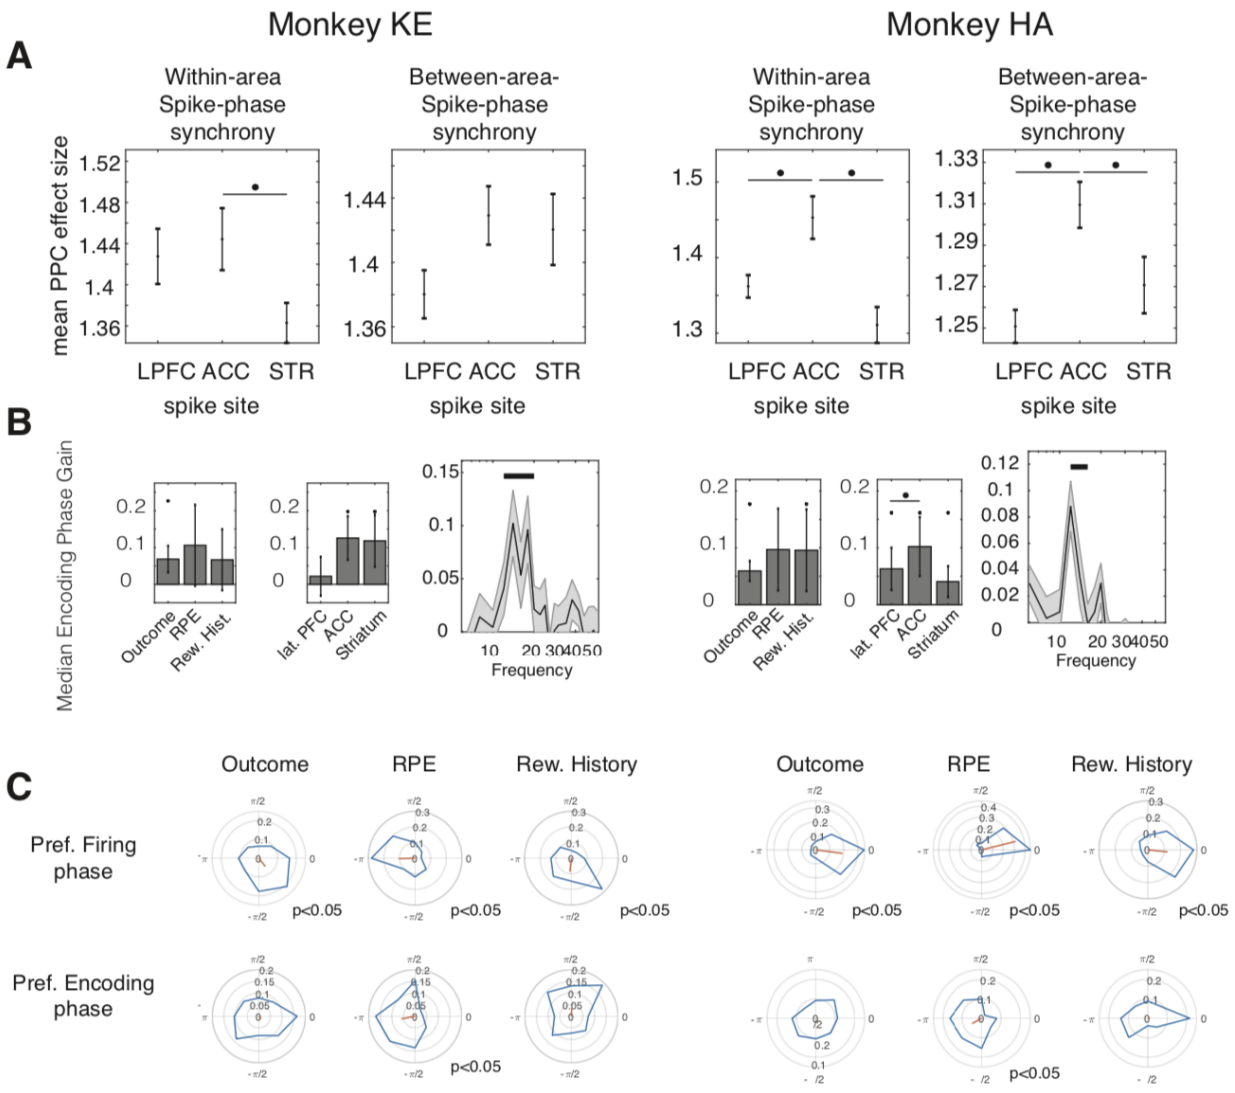

Figure A-3. Summary of main results for individual monkeys

Individual results for monkey KE (left) and HA (right). **(A)** Average spike-LFP phase synchronization for spike-LFP pairs within and between areas. (**B**) Average Encoding Phase Gain for each encoding metric (*left*), brain area *(middle*), and across frequencies (*right)* (**C**) Polar histograms of the preferred firing phase (upper panels) and maximal encoding phase (bottom panels) for each encoding metric.

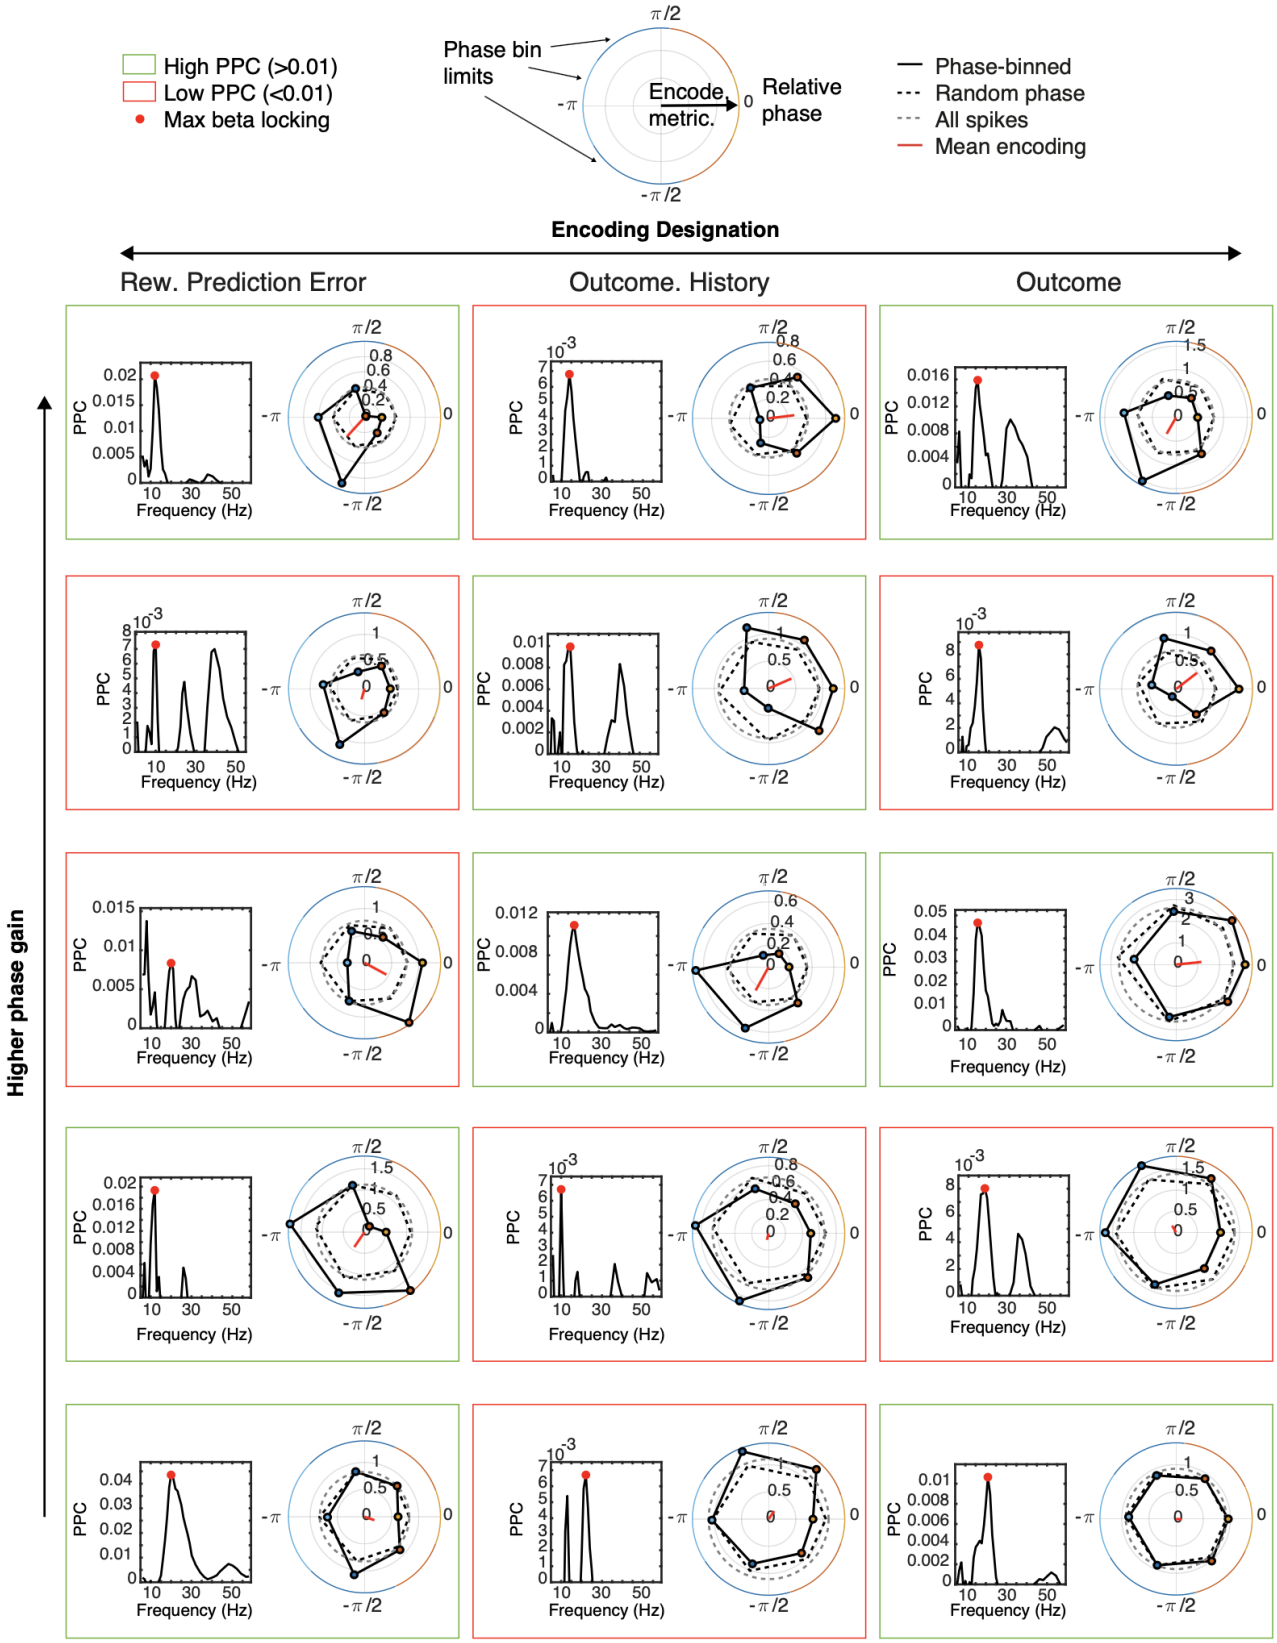
High PPC (>0.01)
Low PPC (<0.01)
Max beta locking
Phase bin limits
Encode metric.
Relative phase
Phase-binned
Random phase
All spikes
Mean encoding
Encoding Designation
Rew. Prediction Error
Outcome. History
Outcome
Higher phase gain
PPC
Frequency (Hz)

Figure A-4. Encoding Phase Gain examples

Columns depict the functional designation. Rows are ordered according to the relative phase gain, with lower phase gain at the bottom, and high phase gain at the top. The spike-phase consistency is depicted on the left, with the maximal significant locking in the [10 25] Hz beta band signified with the red dot. The corresponding phase-dependent encoding is depicted on the right. 0 corresponds to the preferred firing phase. Numbers on concentric circles are the value of the encoding metric. The grey dotted line represents the encoding metric estimated using all spikes, whereas the black dotted line represents the average across many spike-phase randomizations. The red line is the average direction. Colored border lines represent size of the bin. The colored box represents examples with a high (green) or low (red) degree of synchrony.

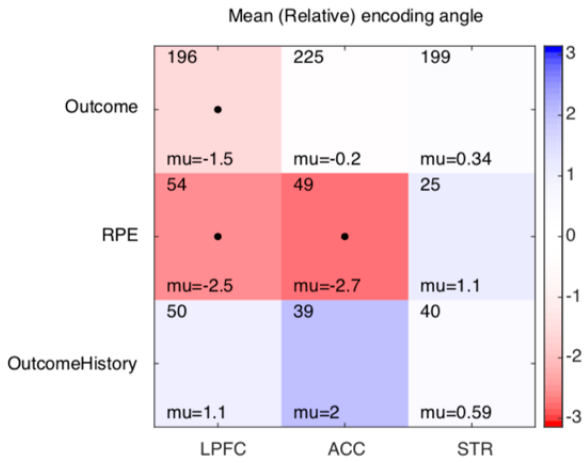

Figure A-5. Preferred encoding for each area and function

Mean preferred encoding phase for *Outcome*, *RPE*, and *Outcome History* (y-axis) cells in the ACC, LPFC, and STR (x-axis). Color represents the (relative) encoding. Black dots represent significant phase concentration (Hodge-Ajne test, $p<0.05$). The number of spike-LFP pairs that went into each cell is depicted in the top left, and the mean phase is on the bottom left.

Appendix B. Supplemental information for Chapter 3

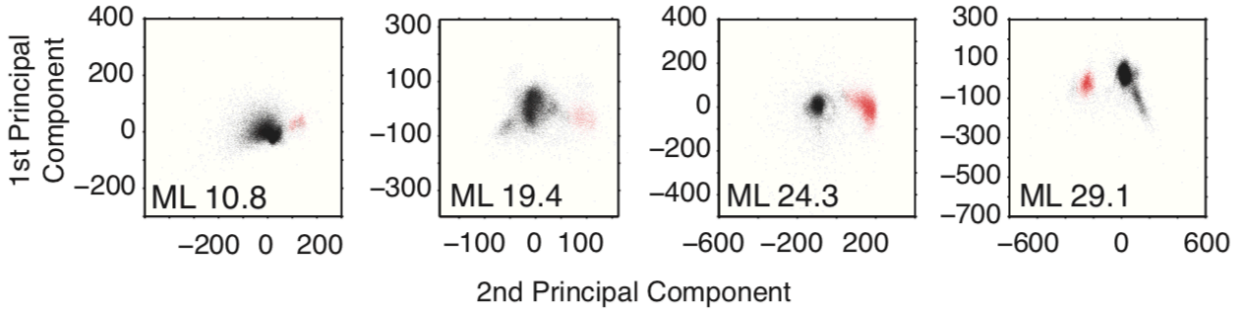

Figure B-1. Scatterplots of the scores of the 1st and 2nd principal component for waveforms belonging to a single neurons (red) and multiunit and noise waveforms (black) that crossed the threshold.

The inset shows the Mahalanobis (ML) distance quantifying the separation of the scores of the red and black data points. The average ML distance across the neuron population was 24.12 ±1.8. The spike isolation procedure also considered the 3rd principal components, the density of the distribution of waveform PCR scores, and the temporal stability of the isolated waveform.

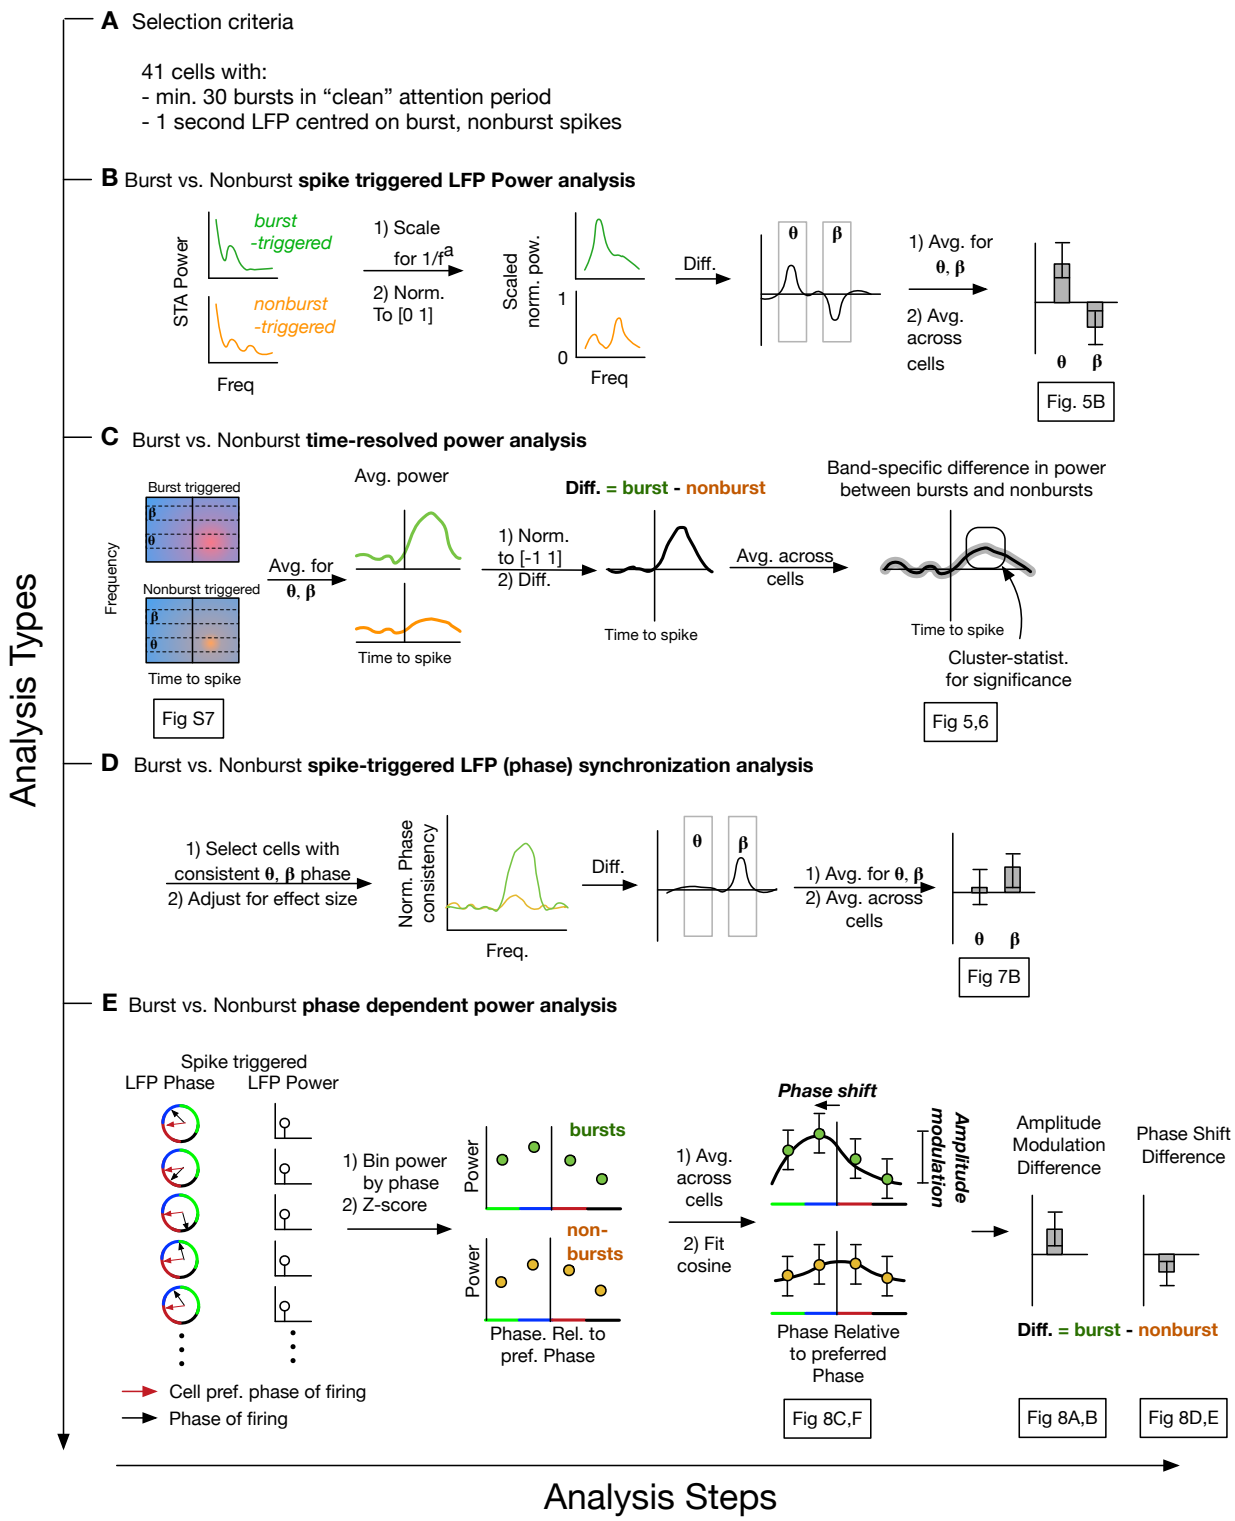

Figure B-2. Summary sketch of all spike-triggered methods and analyses

Summary of all main analyses performed comparing theta and beta LFP activity aligned to burst or nonburst spikes. From top to bottom, analyses are ordered as presented in the main text. From left to

right are the steps taken to generate the final results. The relevant figure that is presented in the main text is highlighted with a box. See methods for more details.

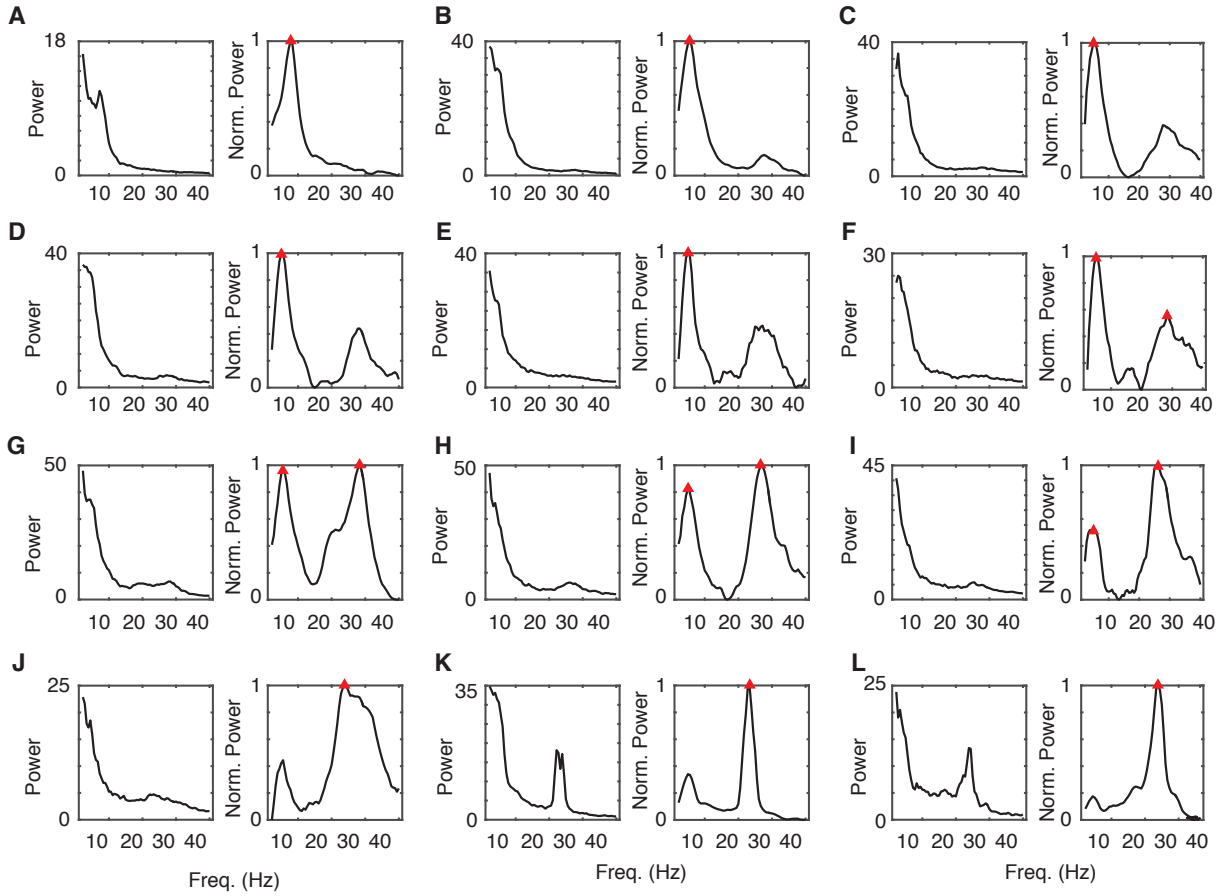

Figure B-3. Example LFP spectra showing theta and beta peaks

(A-L) *Left:* Example raw LFP power spectra for those sites selected for further cell-specific analysis. *Right:* spectra after scaling for 1/f noise and normalizing to the range [0 1]. Red triangle denote peaks as detected by the peak detection algorithm. Subplots are ordered according to whether theta or beta peaks are dominant, starting with subplot (A), where only theta is evident, and ending with subplot (L), where only beta is evident.

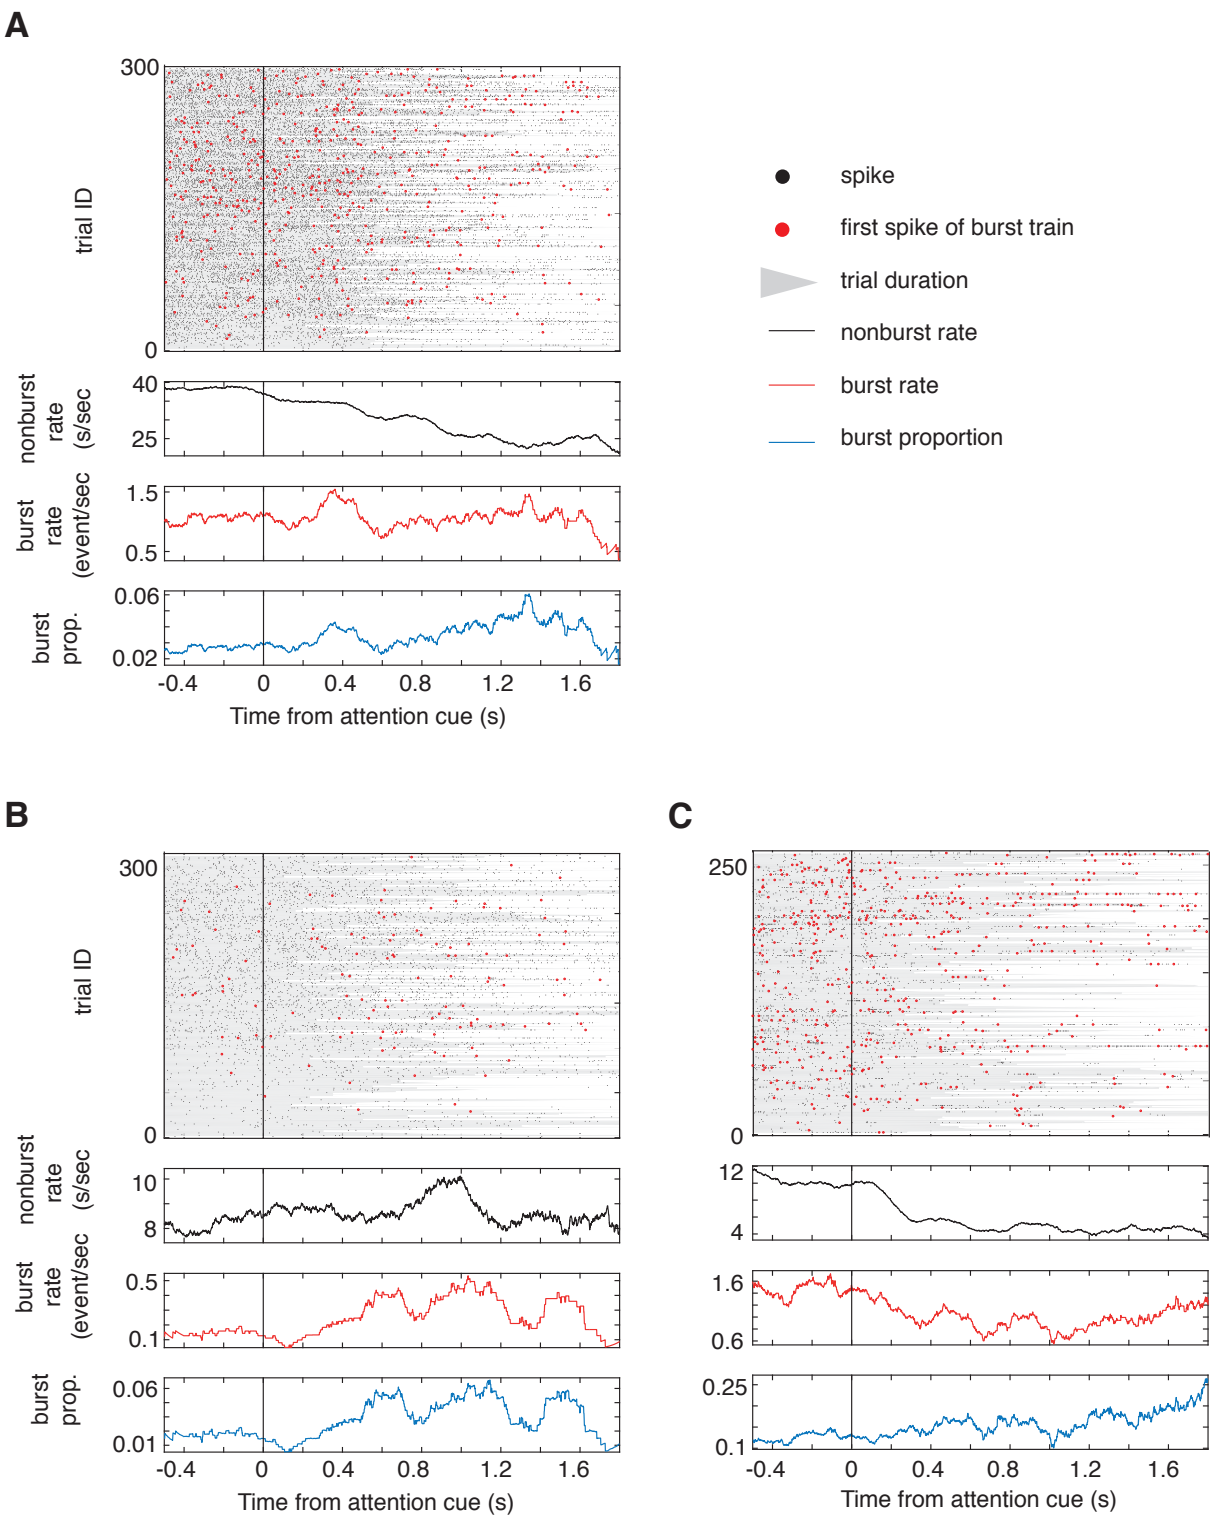

Figure B-4. Example neurons showing that nonburst rate is modulated independently of the burst rate

Complementary to Figure 2 in the main text. **(A-C)** *(Top)* Raster of spiking activity locked to attention cue onset. Black dots denote nonburst spikes, red dots denote the first spike of a burst train. The grey background patch visualizes the duration of each trial, highlighting that trials were of variable length.

Bursts did not occur on every trial; however, even in long trials, they were prevalent late in the trial. *(Middle top)* Nonburst rate aligned to attention cue onset, calculated with a 0.2 s. sliding window. *(Middle bottom)* Burst rate locked to attention cue onset. (*Bottom)* Burst proportion locked to attention cue onset. **(A)** Note that the nonburst rate decreases dramatically while the burst rate remains relatively unchanged, compared to baseline. The result is that the burst proportion increases with attention cue onset. **(B)** Both the burst and nonburst rate appear to increase above baseline. However, as evidenced by the increased burst proportion, the burst rate increases more so than the nonburst rate. **(C)** On the other hand, both the burst rate and nonburst rate can decrease, but the burst proportion can nevertheless increase above baseline.

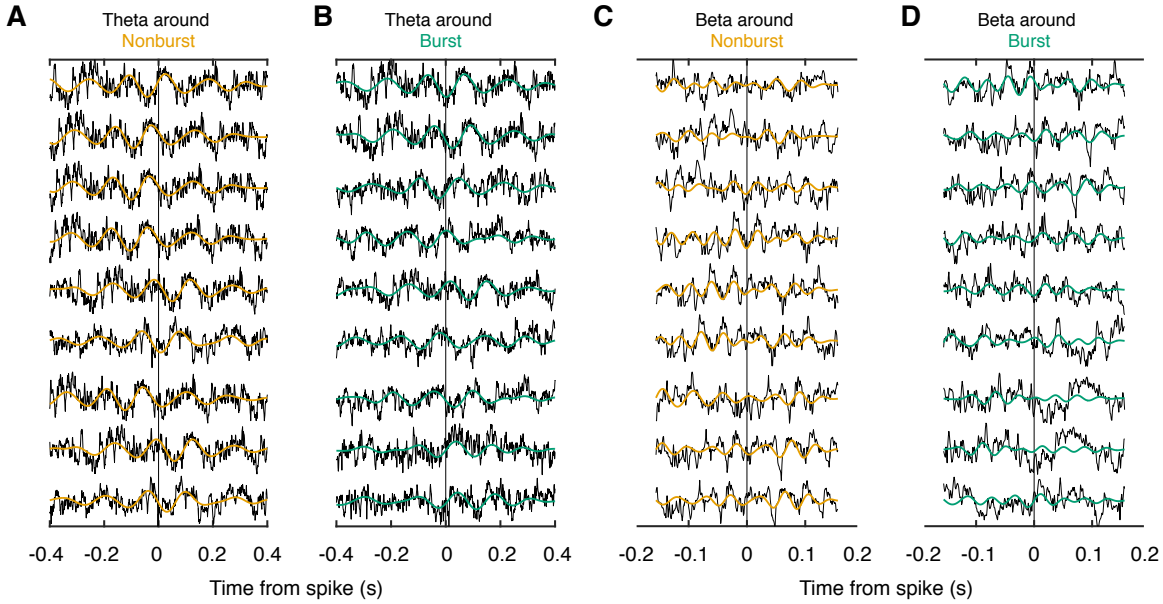

Figure B-5. LFP traces around example narrow spiking (NS) cell

(A-B) Example raw LFP traces around (A) nonbursts and (B) bursts, band-pass filtered for 5-10Hz theta for visualization. **(C-D)** Same as (A-B), but visualized for beta (16-30 Hz).

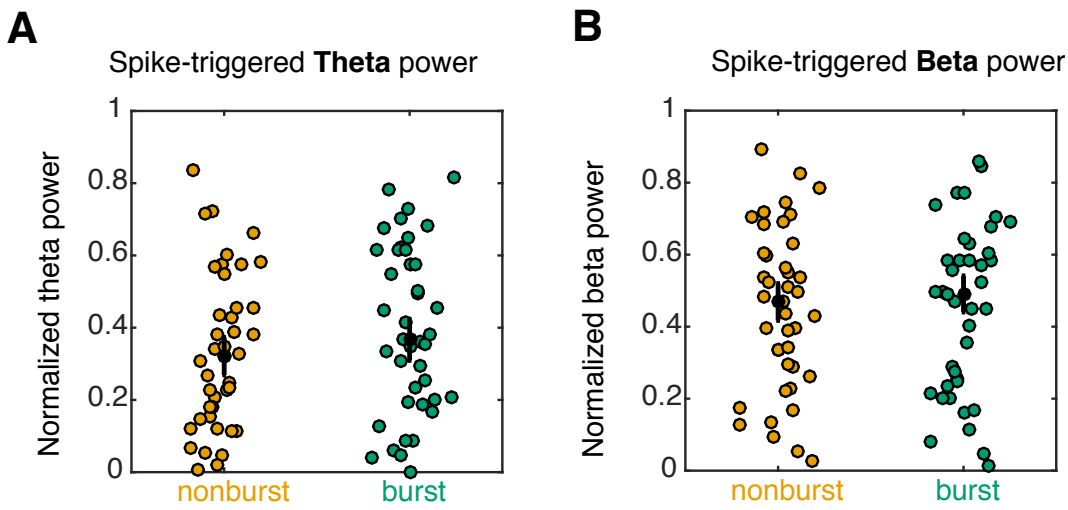

Figure B-6. Normalized theta and beta power locked to burst and non-burst events

Companion to Figure 4. **(A)** Normalized theta (*see* methods) for all cells (n=41), locked to bursts (green), or nonbursts (orange). **(B)** Same as (A) but for beta power

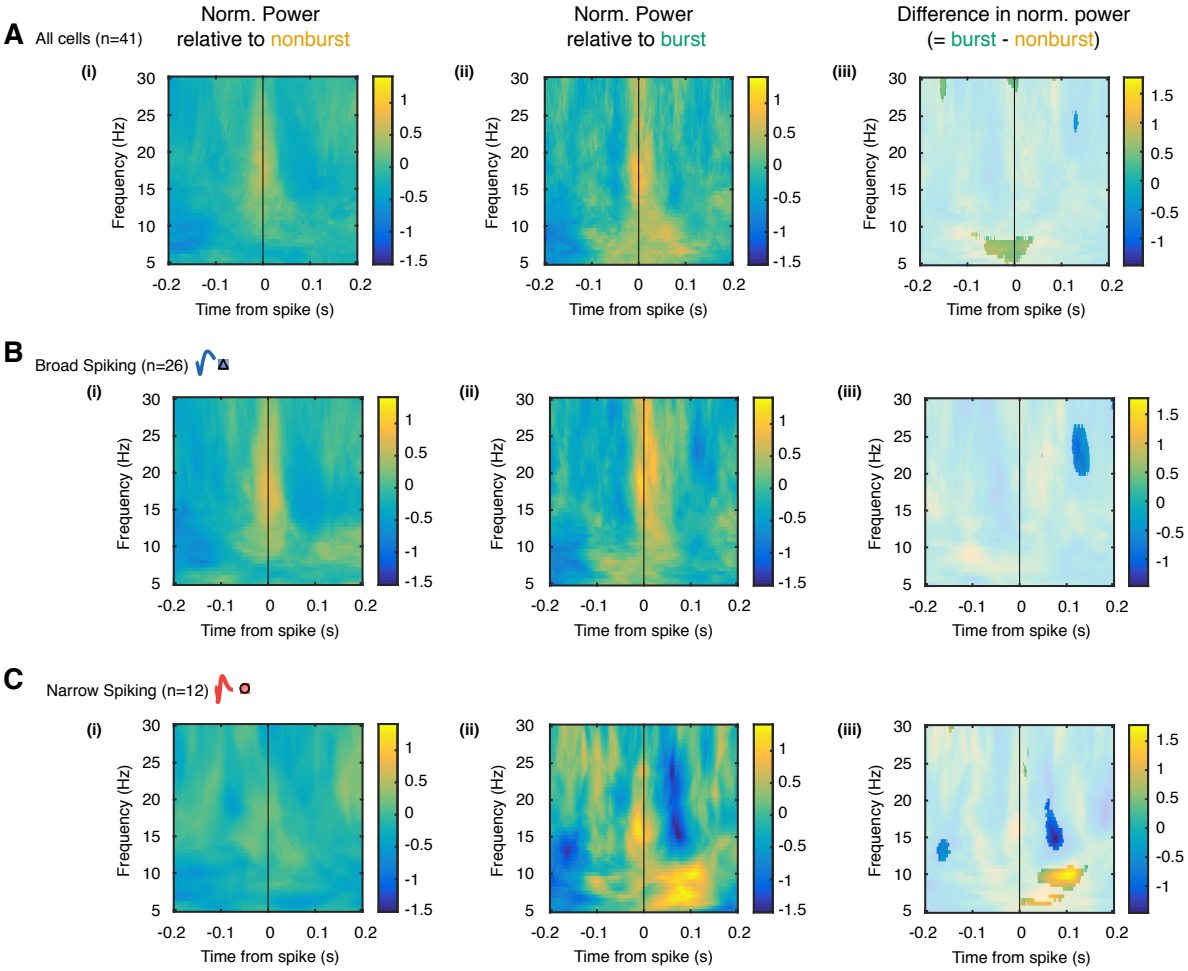

Figure B-7. Full time-frequency spectra aligned to the burst and non-burst spikes for all cells and for narrow spiking (NS) and broad spiking (BS) cells separately.

This figure complements **Figure 4-5,6** of the main text that shows the average spike triggered average power. Power was z-score normalized for each cell individually such that differences in LFP power locked to bursts vs non-bursts were preserved. (**A-C**) Power locked to (A) all cells and separately for (B) BS cells and (C) NS cells. Median power across cells locked to (i) non-bursts and (ii) bursts. (iii) The difference between burst and non-burst locked LFP power. Faded pixels denote non-significant time-frequency bins ($p > 0.05$). No multiple comparison was applied. (A) Across all cells, there is a consistent increase in theta power around the time of spike onset. (B) BS cells show no difference in theta power locked to bursts or non-bursts. (C) In NS cells, bursts show stronger theta power than non-bursts ~100 ms after spike onset. There is a concurrent decrease in (slow) beta power.

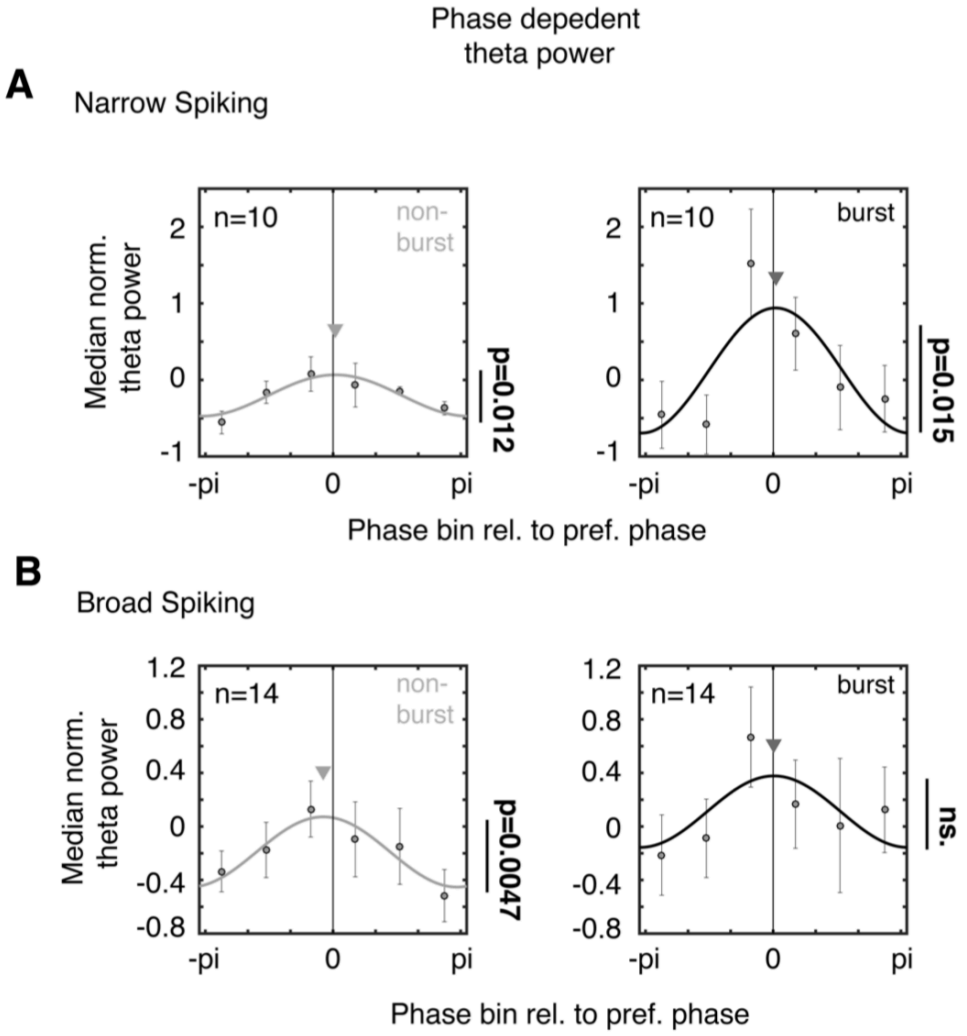

Figure B-8. Phase-dependent power modulation of bursts and non-bursts in the theta frequency range for narrow spiking (NS) and broad spiking (BS) cells.

Cells have been selected if they showed significant theta phase locking. (**A**) Normalized LFP power (y-axis) as a function of the spike phase (x-axis) for NS cells. Power is significantly phase modulated (randomization test, $p<0.05$) with peak power for spikes synchronizing near their preferred LFP phase (with 0 defined as the preferred, average phase to which the spikes synchronized to the LFP in the theta band). The *left panel* showed phase modulated power for non-burst spikes and *right panel* shows results for burst spikes. (**B**) Same as A for non-bursts (*left*) and burst (*right*) spikes of BS cells. Theta power is significantly amplitude modulated by the non-burst and burst phase (randomization test, $p<0.05$) with the exception of bursts from BS cells.

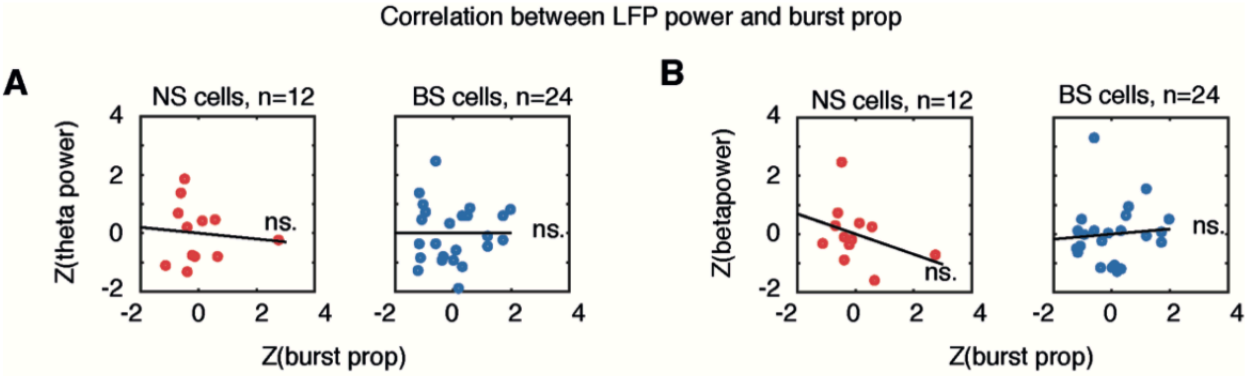

Figure B-9. Burst proportion is not related to the strength of theta or beta power modulation around bursts and non-burst spikes.

(**A**) Spearman correlation between burst proportion and spike-aligned theta power for NS cells (red, *left panel*) and BS cells (blue, *right panel*). Burst proportion and theta power have been z-transformed, and outliers (Z>5) ignored. Black line is the line of best fit. Burst proportion does not co-vary with theta power in either NS or BS cells. (**B**) Same as (A), but for beta power. Beta power and burst proportion are not related.

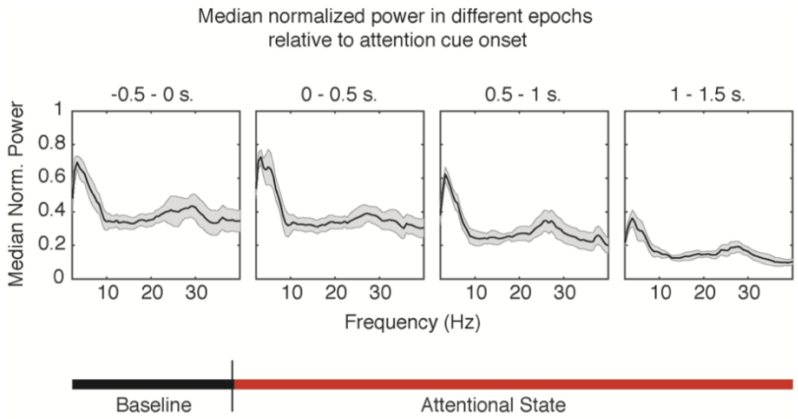

Figure B-10. Theta and beta peaks are evident throughout the attentional state.

Median LFP power locked to attention cue onset in non-overlapping, 0.5s. time bins. Individual spectra were normalized for 1/f noise and scaled to the range [0 1]. In the attentional state, although overall power was reduced, theta and beta oscillatory activity was still evident at the population level as peaks in the spectra.

Appendix C. Quaddles: A multidimensional 3-D object set with parametrically controlled and customizable features

Behavior Research Methods (2019) 51:2522–2532
https://doi.org/10.3758/s13428-018-1097-5

Quaddles: A multidimensional 3-D object set with parametrically controlled and customizable features

Marcus R. Watson [1] · Benjamin Voloh [2] · Milad Naghizadeh [1] · Thilo Womelsdorf [1,2]

Published online: 7 August 2018

Abstract

Many studies of vision and cognition require novel three-dimensional object sets defined by a parametric feature space. Creating such sets and verifying that they are suitable for a given task, however, can be time-consuming and effortful. Here we present a new set of multidimensional objects, *Quaddles*, designed for studies of feature-based learning and attention, but adaptable for many research purposes. Quaddles have features that are all equally visible from any angle around the vertical axis and can be designed to be equally discriminable along feature dimensions; these objects do not show strong or consistent response biases, with a small number of quantified exceptions. They are available as two-dimensional images, rotating videos, and FBX object files suitable for use with any modern video game engine. We also provide scripts that can be used to generate hundreds of thousands of further Quaddles, as well as examples and tutorials for modifying Quaddles or creating completely new object sets from scratch, with the aim to speed up the development time of future novel-object studies.

Keywords Stimulus set · Feature space · Multidimensional objects

In many experiments in the cognitive sciences, participants must view three-dimensional (3-D) stimuli, or 2-D projections of 3-D stimuli, that they have not encountered before. Such novel object sets have been used in studies of such phenomena as object recognition and discrimination (e.g., Biederman & Gerhardstein, 1993; Bülthoff & Edelman, 1992; Chuang, Vuong, & Bülthoff, 2012; Gauthier, James, Curby, & Tarr, 2003; Harman & Humphrey, 1999; Harman, Humphrey, & Goodale, 1999; Hayward & Tarr, 1997; Richler, Wilmer, & Gauthier, 2017; Tarr, Bülthoff, Zabinski, & Blanz, 1997; Wong & Hayward, 2005), perception and attention to different object properties (Arnott, Cant, Dutton, & Goodale, 2008; Cant & Goodale, 2007), memory for objects (Humphrey & Khan, 1992; Knutson, Hopkins, & Squire, 2012; Mercer & Duffy, 2015), facial perception and recognition (e.g., Gauthier & Tarr, 1997; Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999; Wong, Palmeri, & Gauthier, 2009), category representation (e.g., Wallraven, Bülthoff, Waterkamp, van Dam, & Gaissert, 2014; Williams, 1998), conditioned fear responses (e.g., Barry, Griffith, Vervliet, & Hermans, 2016; Bennett, Vervoort, Boddez, Hermans, & Baeyens, 2015; Scheveneels, Boddez, Bennett, & Hermans, 2017), linguistic demonstratives and gestures (Cooperrider, 2015), and emotional influences on perception (Estes, Jones, & Golonka, 2012). Some sets of novel objects have also been presented in articles specifically written to encourage their adoption by other researchers (Barry, Griffith, De Rossi, & Hermans, 2014; Buffat et al., 2014), or simply hosted online (Harris, 2015) for other researchers to use. Figure 1 shows representative exemplars of these sets.

As the use of novel objects in research has become more commonplace, there has been a parallel rise in studies in which participants have engaged and interacted with complex, continually changing virtual environments. Such dynamic tasks, presented on traditional monitors or stereoscopic displays, enable the presentation of much richer stimuli and the collection of much richer data streams than more traditional, static tasks. Dynamic tasks have been used to investigate the processes underlying phenomena such as the mechanisms of spatial navigation in humans and other animals (Bohil, Alicea, & Biocca,

✉ Marcus R. Watson
watsonmr@yorku.ca

Thilo Womelsdorf
thilo.womelsdorf@vanderbilt.edu

[1] Department of Biology, Centre for Vision Research, York University, Toronto, Ontario, Canada

[2] Department of Psychology, Vanderbilt University, Nashville, TN, USA

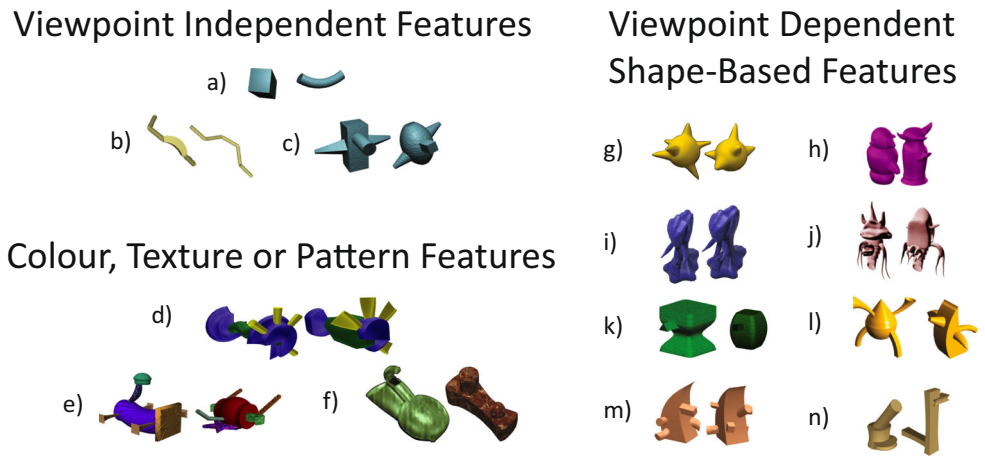

Fig. 1 Representative exemplars from a number of previously reported sets of novel objects. Most do not have features that are viewpoint-independent or that are defined by elements other than shape, both of which qualities we needed for studies with multidimensional view-invariant features. No set has both. (a) geons (Biederman & Gerhardstain, 1993; Hayward & Tarr, 1997), (b) strings (Biederman & Gerhardstein, 1993; Bülthoff & Edelman, 1992; Tarr et al., 1997), (c) multi-geons (Biederman & Gerhardstein, 1993); (*d*) Yadgits (Harris, 2015), (e) Fribbles (Barry et al., 2014; Williams, 1998); (*f*) nonsense objects (Cant & Goodale, 2007; Humphrey & Khan, 1992); (g) amoeboids (Bülthoff & Edelman, 1992; Wong & Hayward, 2005), (h) Greebles (Gauthier & Tarr, 1997), (i) YUFOs (Gauthier et al., 2003), (j) Sheinbugs (Richler et al., 2017), (*k*) two-part objects (Hayward & Tarr, 1997), (l) pair-wise similar objects (Biederman & Gerhardstein, 1993), (m) geons with occlusion (Wong & Hayward, 2005), and (*n*) Ziggerins (Wong et al., 2009). Exemplar pairs g, m, and *n* were taken from http://ww2.psy.cuhk.edu.hk/~mael/Stimuli.html; a, b, c, *d*, e, h, i, *k*, and l were taken from http://wiki.cnbc.cmu.edu/Novel_Objects; j was taken from http://gauthier.psy.vanderbilt.edu/resources/; and *f* was taken from Cant and Goodale (2007). Some stimuli were edited to remove background color in the figure

2011; Ekstrom et al., 2003; Weisberg, Schinazi, Newcombe, Shipley, & Epstein, 2014); multisensory integration in the determination of one's own location (Ehrsson, 2007; Lenggenhager, Tadi, Metzinger, & Blanke, 2007); memory retrieval (Watrous, Tandon, Conner, Pieters, & Ekstrom, 2013); priority in attention, gaze, and memory (Aivar, Hayhoe, Chizk, & Mruczek, 2005; Jovancevic, Sullivan, & Hayhoe, 2006); the temporal organization of gaze in realistic tasks (Johnson, Sullivan, Hayhoe, & Ballard, 2014); subliminal cueing (Aranyi et al., 2014; Barral et al., 2014); and the development of brain–computer interfaces (Leeb et al., 2007).

Both dynamic tasks and novel 3-D object sets, then, have become standard tools in the cognitive science repertoire. We are not aware of any published work that combines the two streams, but we anticipate that this will rapidly become commonplace, as more researchers become aware of the power and flexibility these tools enable, without a corresponding sacrifice in experimental control. Our laboratory has begun running such studies, in which we examine attentional and oculomotor changes as participants learn about a novel object set in a dynamic environment (Watson, Voloh, Naghizadeh, Chen, & Womelsdorf, 2017). With so many sets of novel objects freely available (see Fig. 1), it came as a surprise that we could not find a multidimensional set that met our requirements. Instead, we had to design our own and test their suitability for our task, a much more difficult and time-consuming project than we had originally anticipated, and one that we hope to make substantially easier for future researchers.

In the present article we review this novel object set, named *Quaddles* in reference to the four feature dimensions that define the object space. In addition, we describe (and provide links to) tools that allow the creation of thousands of parametrically varied objects using preexisting features, which enable researchers to design custom features quickly and relatively easily. Finally, we present the results of a feature detection task showing how sensitivity, response bias, and detection efficiency to different feature values of a particular set of Quaddles can be quantified, establishing that, for the most part, the feature values are equally detectable and do not produce strong response biases.

Introducing Quaddles

The experimental task for which Quaddles were designed has participants moving freely using a joystick around a realistic virtual 3-D environment and choosing between objects in that environment. Their object selections are either rewarded or not, on the basis of the particular feature values of the objects, and participants have to learn through trial and error which feature values are associated with reward. Our requirements for these objects were that they have:

- An aesthetically pleasing appearance
- Multiple feature dimensions, including nonshape dimensions
- Multiple feature values along each feature dimension
- Feature values that are roughly equally perceptible and do not show strong response biases

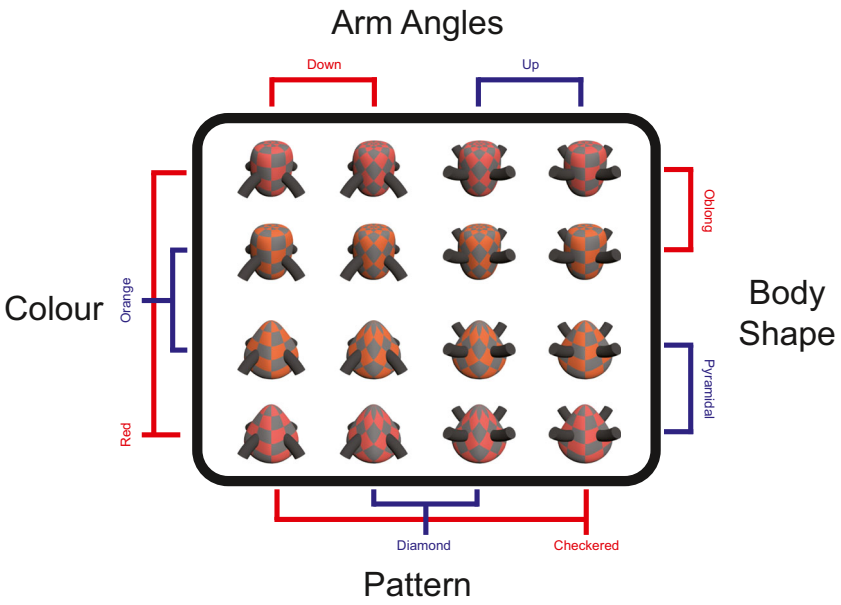

Fig. 2 The primary set of 16 Quaddles used in the feature identification study described in this article

- Vertical symmetry
- Features that can all be clearly and simultaneously viewed from any angle around the vertical axis
- The ability to be exported to any commonly used image format (PNG, GIF, JPG, etc.)
- The ability to be exported to any common video game engine (Unity 3D, Unreal Engine, etc.)

No previous set of novel objects we are aware of meets all these criteria. Very few of these sets have viewpoint-independent features that can all be viewed simultaneously from any angle around the object, and most have purely shape-based feature dimensions (see Fig. 1). Furthermore, we are only aware of one object set for which feature similarity has been quantified (Barry, De Rossi, & Hermans, 2014), and balancing perceptibility across different feature dimensions has only been attempted for much simpler stimuli (Kaldy, Blaser, & Leslie, 2006)—hence, the need for a new set of objects.

The primary set of Quaddles are defined by four feature dimensions: body shape, arm angle, surface color, and surface pattern, each of which has two possible values (Fig. 2). The set can easily be extended by adding new feature values (Figs. 3 and 4), by using intermediate (morphed) values between the existing values (Fig. 3), or by adding or removing dimensions (Fig. 4). All Quaddles shown in these figures are freely available from our website (http://accl.psy.vanderbilt.edu/resources/analysis-tools/) in FBX file format. High-definition pictures of each object from a variety of perspectives are also available in JPG formats with black, white, gray, or green backgrounds, as well as in PNG format with transparent backgrounds. Videos of the objects rotating against black, white, and gray backgrounds are also available in mp4 format. Finally, the same website contains a detailed manual and links to a github script repository containing code that can generate all the objects shown here, as well as tens of thousands of others, and that can easily be modified to create custom object sets.

In the remainder of the article, we summarize the methods for creating Quaddles (more complete details can be found in the manual, hosted at http://accl.psy.vanderbilt.edu/resources/analysis-tools/), and present the results of a feature detection study quantifying sensitivity, response bias, and efficiency to the different feature values. In this task, participants were cued with two feature values prior to being shown a single Quaddle that contained only one of these values, and they had to report which of the two values was present. We calculated independent measures of sensitivity (d') and response bias (criterion) for each of the feature values, using standard signal detection measures (Macmillan & Creelman,

Fig. 3 Examples of morphing, using two different feature values on each of the four feature dimensions

Fig. 4 Additional objects showing some possible variations on the basic Quaddle body plan that can be easily generated via scripting. The top row shows changes in the number of arms, the addition of a "hat" feature that allows two body shapes to be stacked on top of each other, and the addition of feature values such as the hairy surface texture, whereas the bottom two rows show objects with all the same feature dimensions, but alternative feature values

2005). Sensitivity on this task quantifies the ability to correctly report a feature as being present or absent, whereas response bias quantifies the likelihood of reporting a particular feature overall, either correctly or incorrectly. The desired results were indistinguishable sensitivities to the different feature values and bias scores that would cluster around 0, indicating no consistent biases across participants. We also calculated efficiency scores that combined accuracy and response time, controlling for speed–accuracy trade-offs (Smilek, Enns, Eastwood, & Merikle, 2006; Townsend & Ashby, 1983). Furthermore, we quantified the stability of the efficiency scores using consistency metrics, both across blocks within individual participants and between participants. To prefigure our results, this analysis demonstrated that feature value detection efficiency is fairly consistent within individual participants, but this consistency is substantially weaker across participants. Furthermore, these can be used as covariates in analyses of other tasks using Quaddles, thus controlling for any effects that differential sensitivities, strong response biases, or different efficiencies might produce on these tasks.

Method

Stimulus generation

Quaddles were generated using Autodesk 3DS Max software. The primary set, used to generate all results presented in this article, is defined by four feature dimensions (body shape, branch angularity, pattern, and color), each of which can take on two possible feature values (e.g., body shape can be pyramidal or oblong), giving a total of eight feature values and 16 possible objects (Fig. 2). The scripts we provide online include options for incorporating further feature dimensions and values into stimulus sets, allowing the creation of hundreds of thousands more objects, but we do not describe these in detail here (full instructions are given in the manual).

Textures, which define the surface colors and patterns, are imported from PNG files created in a Matlab script. The neutral gray color is the same for all objects, whereas the other colors are chosen within the CIE L*c* h* space such that the L* and c* values (luminance and saturation, respectively) are held constant, but h* values (hue) vary by 15°, meaning that there is a small difference in hue between the two colors, but not in the other components of the colors. Textures are applied to object surfaces using standard UV mapping options: a cylindrical wrap for pyramidal bodies and a spherical wrap for oblong ones (different wraps were chosen because they resulted in smaller artifacts at the top and bottom of objects).

Quaddle bodies are initially generated as spheres and then molded into the desired body shapes using freeform deformation (Sederberg & Parry, 1986), in which a lattice of control points is added to the object and manipulated to create the desired shape. Thus, all body shapes are morphs of each other, allowing for intermediary shapes as desired. Each Quaddle has four arms, initially generated as straight cylinders and then morphed into the desired shape, thus also supporting intermediary values. The same is true of both the hues used to define object color, and their surface patterns. This means it is easy to create objects chosen from anywhere within the feature space

defined by the four feature dimensions (Fig. 3). Given any two objects, one can also create videos of the morph between them, or even objects that morph in real time in a 3-D environment.

Object generation was automated using a 3DS Maxscript that creates and saves complete object sets. An optional function allows JPEG, PNG, or other image files to be generated of every object created from any distance, height, and rotation. Experimenters also have the option of saving videos of the objects rotating 360° from any perspective.

For illustration purposes, we generated two more feature values along each dimension, and generated partial morphs of the objects along each dimension (Fig. 3). We also generated a number of further objects with varied numbers of arms, amalgamated body shapes stacked on top of each other, and various new textures and shapes (Fig. 4). Making new Quaddles in this way is quite easy using simple modifications of our existing scripts, allowing the powerful and flexible generation of new object sets.

Experimental procedures

The York University Office of Research Ethics approved the present study as conforming to the standards of the Canadian Tri-Council Research Ethics guidelines (Certificate # 2016-214). Ten participants (mean age 28 years ± 3.8 *SE*; six male, four female) took part in the study. One was excluded from further analyses due to chance performance. They were seated approximately 60 cm from an LED monitor with a 60-Hz refresh rate, with heads unrestrained. The entire study, including an instructional tutorial, took approximately 1 h. The task was coded in the Unity game engine.

For the duration of the experiment, participants viewed a diamond-shaped arena from one of its vertices (Fig. 5). At the start of each trial, the floor of the arena changed to one of 20 different textures, chosen at random, which were also used in the learning task for which we designed the Quaddles. The floor textures were chosen from a large, free database of textures (https://share.allegorithmic.com), and included a wide variety of hues, contrasts, spatial frequencies, and semantic information. After 200 ms, two cues appeared, each showing an iconic representation of one of the eight stimulus feature values (2 feature values × 4 feature dimensions). After a further 250 ms (± 50-ms jitter), a single Quaddle was displayed at the center of the screen for 250 ms, subtending 3.5°–4.0° of visual angle at a 60-cm viewing distance. Participants had to quickly decide whether the single Quaddle contained the feature value of the left or right iconic image cue. A mask pattern was then flashed over the Quaddle location for 50 ms, after which both it and the Quaddle were removed. The cues remained on screen until participants had responded by pressing either the "Z" or the "/" key on a standard keyboard (indicating that the left or right cue, respectively, was accurate), or until 2 s had elapsed, whichever was quicker. If participants did not respond within 2 s, the game was paused and they were asked to respond more quickly on future trials. After response, feedback was presented for 500–600 ms in the form of a colored border around the chosen cue, with green indicating correct and red indicating incorrect. Each Quaddle had one of the two cued feature values, but not the other, and participants were tasked with reporting which of the two cues was present. After feedback the cues were removed, and the next trial started immediately (see Fig. 5).

Participants were given approximately 5 min of training on a slower version of the task prior to starting, and they were instructed to respond as quickly and accurately as possible. A single block contained 512 trials, consisting of 32 trials for each of the 16 Quaddles. Each of the four feature values

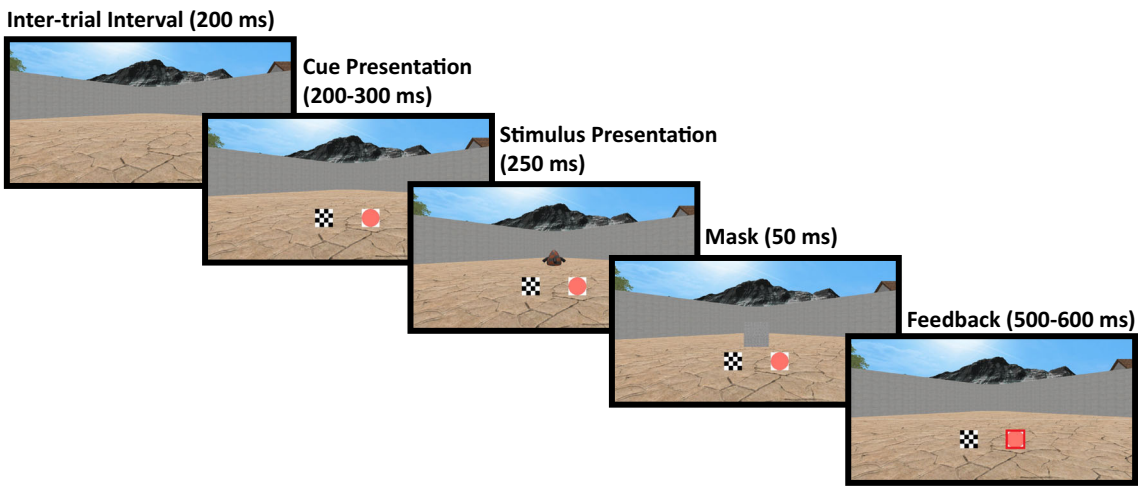

Fig. 5 Overview of a feature detection trial. On this trial, the participant incorrectly responded that the presented Quaddle had a reddish color, instead of a checkered pattern, so their incorrect choice was outlined in red during the feedback phase. Had they chosen correctly, their choice would have been outlined in green

present on a given Quaddle was presented as a valid cue eight times, twice with each of the four feature values not found on the same Quaddle as the invalid cue, on both the left and the right side. Every 32 trials, each of the 16 Quaddles was shown twice, once with the valid cue on the right, and once on the left, but in all other respects the cues, Quaddles, and side of the valid cue were randomized. After a block, participants were given an optional break. Most participants ran through three blocks in approximately 60 min, but three participants were only able to finish two due to time constraints.

Several pilot versions of the study were run. After each, we adjusted the object feature values to try to eliminate any gross response biases. We present the results only for participants run using the final set of feature values, which had the most unbiased performance across feature dimensions.

Results

One participant was excluded from analyses due to chance accuracy. For the remaining participants, all trials in which a given feature value was a valid cue were grouped together, as were all trials in which it was an invalid cue. This resulted in 16 groups of trials, within which the accuracy and response time on correct trials were averaged. There were substantial interindividual differences, with individual accuracies for particular feature values ranging from 54% to 99%, and correct response times ranging from 500 to 1,180 ms.

For each of the eight feature values, we calculated a hit rate (the accuracy on all trials in which the target was validly cued) and a false alarm rate (1 – accuracy on all trials in which the target was invalidly cued). This effectively treated our experiment as eight independent yes–no tasks, where reporting a given feature value as present corresponded to a "yes" response and reporting any other value as present corresponded to "no," allowing for the calculation of standard sensitivity (d') and bias (c, criterion) measures (Macmillan & Creelman, 2005), shown in Fig. 6. A c of 0 indicates no bias, negative numbers indicate a preference for a given alternative, and positive numbers indicate a preference against a given alternative (thus, Fig. 6b shows very minor biases for most features, but stronger biases for participants to report checkered patterns and not to report downward-bending arms). Individual sensitivities for particular feature values ranged from 1.01 to 4.42, and the absolute values of biases ranged from 0 to .44.

The mean sensitivity across all participants and all feature values was 2.59, which lay between Quartiles 1 and 2 for all feature values. The checkered and diamond patterns were clearly different from the other feature values, in that they had the two lowest median sensitivities, which were the only median sensitivities below the first quartiles of any of the other feature values, as well as third quartiles that were below two of the other values' median sensitivities (orange color and

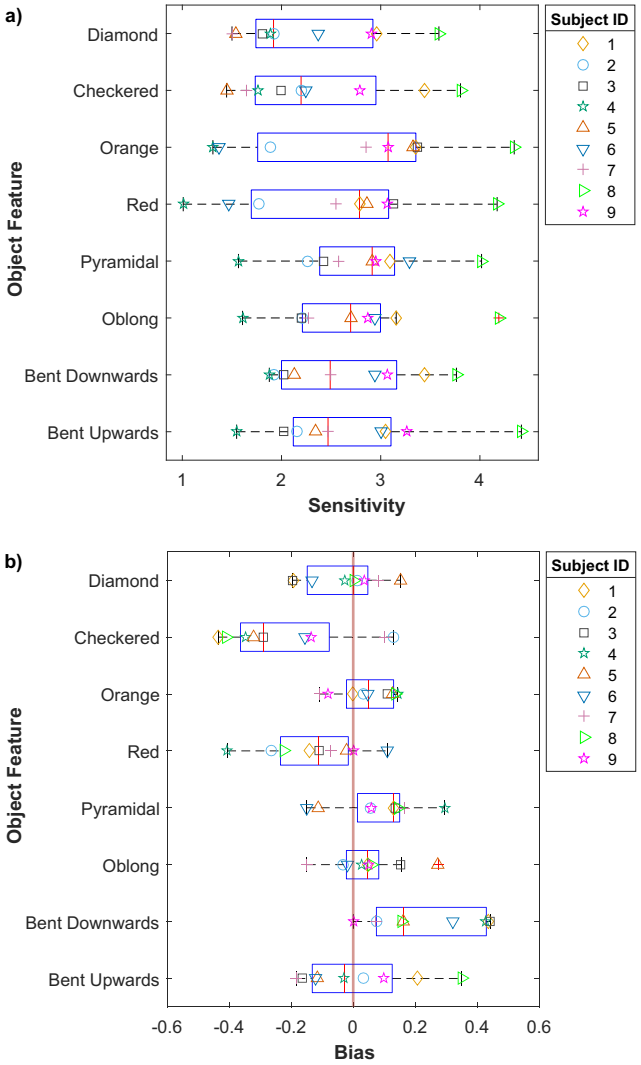

Fig. 6 Feature detection study sensitivities (a) and response biases (b) for each of the eight feature values. Higher sensitivity scores indicate more accurate detection of the feature in question; negative bias indicates a tendency to report the feature in question; positive bias shows a tendency not to report it; and 0 indicates no bias

pyramidal shape). All other median sensitivities were between the first and third quartiles for all other feature values, with a single exception (the oblong shape's third quartile was below the median sensitivity of the orange color). Thus, participants were consistently less sensitive to pattern than to the other feature dimensions. Consistent with this, trials with either valid or invalid pattern cues had the lowest median response times (not shown), indicating that the values of the pattern dimension were particularly difficult to discriminate and that the lower sensitivity was not simply due to a speed–accuracy trade-off.

The response biases were clustered around 0, with a median bias for one feature value (diamond pattern) at exactly 0, three below 0, and four above. The absolute value of all medians was less than .20, with the exception of the checkered pattern, which had a bias of – .29. All feature values produced both positive

and negative biases in different participants, with the exception of downward-bent arms, which had no bias or a positive bias for all participants. Thus, in general, feature values did not produce strong and consistent response biases, with the exceptions of checkered patterns and downward-bent arms.

Another method of quantifying performance on two-alternative forced choice tasks such as these, which has the advantage of producing a single number for each participant while controlling for a speed–accuracy trade-off, is to compute *efficiency scores* from the raw accuracies and response times (Smilek et al., 2006; Townsend & Ashby, 1983), which we used to quantify interindividual differences in performance and to examine the degree to which individuals' detection efficiency was self-consistent across blocks and consistent between participants (Figs. 7 and 8).

There were 56 cue pairs (8×8 feature values, but the same value cannot be paired with itself), for which accuracies and response times were combined by dividing the response time on correctly answered trials in which these cues were presented by the mean accuracy on these trials (Smilek et al., 2006; Townsend & Ashby, 1983). Doing this corrects for speed–accuracy trade-offs in an intuitive way: When accuracy is perfect, efficiency will be identical to the mean response time, and as accuracy decreases, the combined score increases, making efficiency scores similar to response time in that smaller values indicate better performance. This assumes that response time and accuracy are linearly correlated, which was supported in the present data ($r = -.53, p < .001$). These scores were then transformed into z scores using the mean and standard deviation of the efficiency scores across all trials for each participant. These 56 scores were combined into 28 for each participant by averaging over the valid–invalid and invalid–valid trials for each feature value pair, which was also justified, on the basis of a strong correlation between the two sets of scores ($r = .83, p < .001$). Finally, these 28 feature value pair scores were combined into eight feature value scores, by

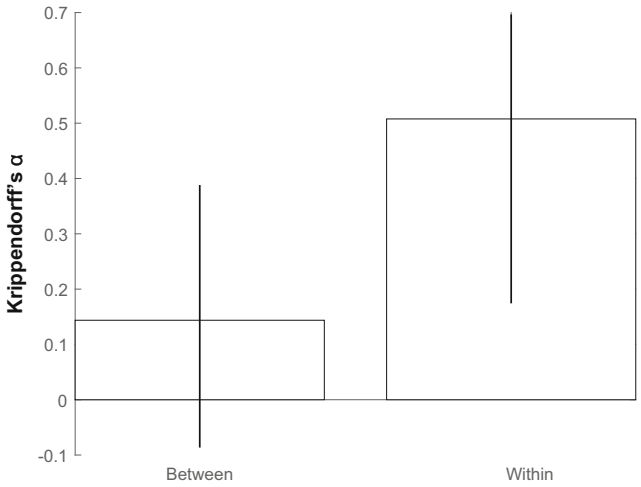

Fig. 8 Krippendorff's alpha values for normalized efficiency score ranks for feature values between and within participants. The between-participant scores are close to 0, indicating that any given individual's feature value preferences are close to random. Within-participant scores, on the other hand, are higher, indicating that individual preferences are stable across multiple blocks. Error bars indicate 95% confidence intervals, calculated using 10,000 bootstrap samples

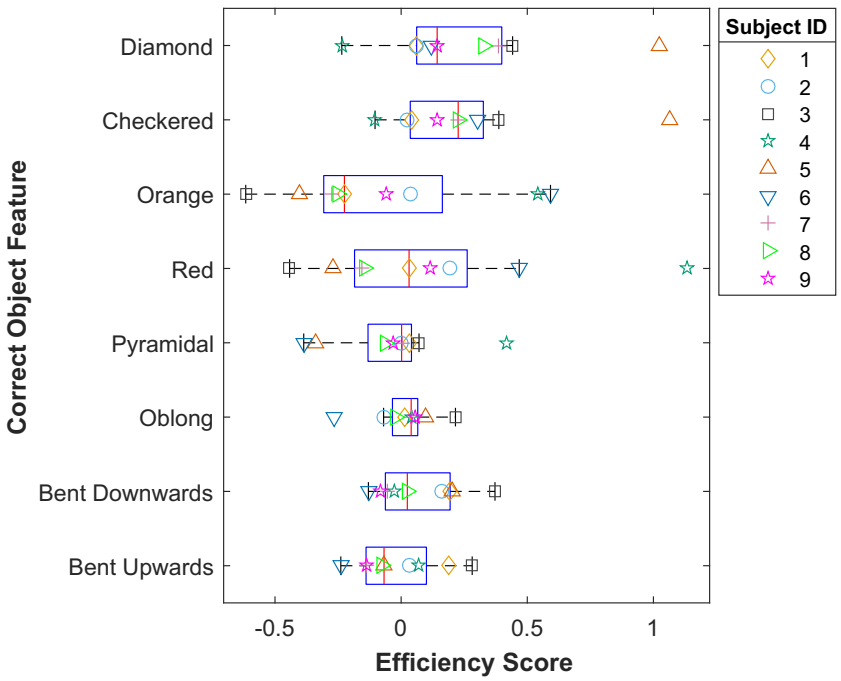

Fig. 7 Normalized efficiency scores on the feature identification task. Lower numbers indicate better performance. 0 indicates each participant's mean efficiency across all trials, and 1 indicates a standard deviation away from this mean

averaging all trials in which a given feature value was presented as a cue with any of the other seven feature values. A score of 0 indicated the mean efficiency across all feature values for each participant, and 1 (or – 1) indicated a standard deviation away from this mean, with lower scores indicating greater efficiency.

These normalized efficiency scores are shown in Fig. 7. Their median values across all participants are close to 0, with no scores being more than 0.25 *SD*s away from the mean, indicating that most of the variance in response times and accuracies on the present task was due to factors other than response biases to particular feature values. All the 95% confidence intervals for these scores included 0, indicating little, if any, difference between the efficiencies for any feature values. However, once again the two surface pattern feature values have the two highest scores (.23 and .14), consistent with the lowered sensitivity scores to these values (Fig. 6).

As well as quantifying efficiency scores for particular feature values, we wanted to quantify how consistent these scores were, both between and within participants. This was accomplished using a tool from the content analysis field, *Krippendorff's alpha*, or *Kα* (Hayes & Krippendorff, 2007; Krippendorff, 2011). Kα, which indicates the reliability of multiple sets of scores for a number of items, ranges between – 1 and 1, where 1 indicates perfect consistency, 0 indicates a completely random distribution of scores across sets, and – 1 indicates perfectly systematic disagreement (Krippendorff, 2008). Generally speaking, Kα is used to measure the consistency of questionnaires or other rating instruments, in which case a high value (e.g., .80 or higher) is desirable. However, in the present case, values approaching 0 indicated a *lack of* consistent bias toward particular feature values, as we desired for our object set.

To calculate Kα, the normalized efficiency scores were transformed to rank orders. For the between-participants Kα, this was done across all trials to give a single set of scores for each participant. Kα was calculated using a freely available Matlab script (Eggink, 2012), and standard errors and confidence intervals were calculated using a bootstrap method, as recommended by Zapf, Castell, Morawietz, and Karch (2016), save that we used 10,000 samples due to the small number of participants, and used bias-corrected and accelerated confidence intervals, which provide more accurate estimates of the true interval (DiCiccio & Efron, 1996). The resulting between-participant consistency was low (Kα = .14, *SE* = .17), and its 95% confidence interval included 0. For the within-participants Kα, we calculated a separate set of efficiency scores for each block performed by each participant, found the rank-ordering of these scores, then calculated a single across-block Kα for each participant using these rank orders and averaged these values across participants, using a 10,000-sample bootstrap to calculate standard errors and confidence intervals. This showed a substantial degree of within-participant consistency (Kα = .51, *SE* = .12). Thus, individuals have reasonably consistent efficiencies to the different feature values over time, although across individuals these efficiencies are much closer to randomly distributed (see Fig. 8).

Discussion

This article presented the multidimensional set of Quaddle objects, suitable for use both in studies involving navigation through visually appealing virtual 3-D worlds, or for more static studies that require images or videos of multidimensional objects. The results of a simple feature discrimination task showed that participants are roughly equally sensitive to the feature values along each dimension (Fig. 6a), with the exception of the two pattern dimensions; that there are not strong and consistent response biases to most of these feature values (Fig. 6b); and that although individuals' response efficiencies for particular feature values are somewhat consistent across time, there is little consistency across participants (Figs. 7 and 8). These results suggest that Quaddles can be used "off the shelf" in a wide array of tasks that require balanced feature values, and they point to modifications that could further improve this balance (e.g., by increasing the size of individual checkers or diamonds, or increasing the contrast between their dark and light elements, to increase the discriminability of the patterns). They can even be 3-D-printed, should an experiment be truly real-world, and can generate stereoscopic images for use with virtual reality or augmented reality experiments.

Limitations of the discrimination study

Although the results of our feature detection task are encouraging, it is important to acknowledge their limitations. First, there is substantial individual variability. For a study in which it is critical that each feature value be equally discriminable by all participants (for some arbitrarily small value of "equally"), it might be necessary to produce a much larger set of objects with many intermediate feature values, and run participants on a complex adaptive staircase task (see, e.g., Anderson & Johnson, 2006; Klein & Macmillan, 2001; Kujala & Lukka, 2006; Treutwein, 1999), resulting in a personalized set of objects for each participant. Producing such intermediate objects is possible with relatively simple modifications of our existing scripts. Developing such a staircase task would require careful consideration of the specific requirements of the experiment in question.

Our study presents objects at a single distance. In a study in which object distances vary, such as any involving navigation through a 3-D world, different features will become more or less discriminable at different distances. Equating discriminability across multiple distances would make for a much longer and more complex study than was feasible in our time frame.

Finally, our study presented objects within a single arena that does not change, save for the floor, which changes drastically across trials. This is simply because the study for which we developed these objects involves a single arena with floors that change across trials. The surround of an object can have powerful effects on feature discrimination, but our study does not control for these effects, as we reasoned that given the wide variety of floors we present, their effects would be essentially random. Once again, any experimenter for whom this is a critical concern will have to run another set of feature detection studies, modifying the objects and environment as needed.

We believe it is apparent that *controlling for all factors that affect feature detection is not feasible*, because one can only measure detection in a given experimental context. We have mentioned three factors that we did not control for, which likely interact in highly non-linear ways: individual preferences, object distance, and visual background. Even if their interactions are completely linear, controlling for all of them simultaneously would require a very complex task and a very large number of participants. Furthermore, there are certainly other relevant contingencies that we have not outlined here. Indeed, preliminary results from our laboratory suggest that in a dynamic object selection task in which the rewarded feature values are learned through trial and error, Quaddle shapes may have an advantage over the other feature dimensions (data not shown). Sensitivity, response bias, and efficiency scores need to be measured in their respective task environments, consistent with the insight that multiple factors determine whether an object guides attention in visual search tasks (Wolfe & Horowitz, 2017).

Experimenters will have to determine to what degree their particular task requires controlling for different factors that might affect feature value discrimination and detection, and design their objects and tasks accordingly. Alternatively, instead of controlling for such factors, one could simply quantify their differential effects, and use sensitivity, bias, or efficiency scores as covariates in statistical analyses, to be partialed out from the main effects of the respective studies.

Possibilities for further customization

Figures 2, 3, and 4 show several ways in which the basic Quaddle feature space can be manipulated, but there are many other ways in which researchers might change Quaddles for their own purposes. For example, they might wish to remove the vertical symmetry of some, or all, feature dimensions, so that manipulating or navigating around objects would be an important part of identifying them, as is the case with most of the objects shown in Fig. 1, as well as with many, if not most, real-world objects. Similarly, it might also be of interest to systematically vary the discriminability or salience of different feature dimensions, and to quantify this variance using a feature detection task. This enables the role of feature bias to be directly studied, as opposed to minimized as with the present object set. Such changes would require trivial modifications to the existing scripts. Aside from these examples, many other possibilities of course exist for future studies to implement additional changes.

Concluding remarks

With this article, we introduced a new object set, characterized its discriminability, and provided tools to facilitate its use in a wide range of possible future studies. This novel set of 3-D objects has normed, parametric features, suitable for a wide range of tasks; open online access to the examples and tools allows researchers to rapidly create custom object sets suitable for other studies. This pragmatic aspect resonates well with the spirit of recent toolkits for video game engines that have streamlined the development and running of dynamic experiments (Doucet, Gulli, & Martinez-Trujillo, 2016; Jangraw, Johri, Gribetz, & Sajda, 2014). The properties of Quaddles make them a suitable set of novel objects for future studies using more realistic and complex tasks, and the scripts we provide can significantly reduce the time necessary to develop different objects for such tasks.

Author note This work was supported by Grant MOP 102482 from the Canadian Institutes of Health Research (to T.W.) and by the Natural Sciences and Engineering Research Council of Canada, Brain in Action CREATE-IRTG program (to M.R.W. and T.W.). The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of this article. The authors thank Hongying Wang for technical support, and Isabel Gauthier for comments on a draft version of the manuscript. The study described herein was approved by the York University Office of Research Ethics (Certificate # 2016-214).

References

Aivar, M. P., Hayhoe, M. M., Chizk, C. L., & Mruczek, R. E. B. (2005). Spatial memory and saccadic targeting in a natural task. *Journal of Vision*, *5*(3), 177–193. https://doi.org/10.1167/5.3.3

Anderson, A. J., & Johnson, C. A. (2006). Comparison of the ASA, MOBS, and ZEST threshold methods. *Vision Research*, *46*, 2403–2411. https://doi.org/10.1016/j.visres.2006.01.018

Aranyi, G., Kouider, S., Lindsay, A., Prins, H., Ahmed, I., Jacucci, G., . . . Cavazza, M. (2014). Subliminal cueing of selection behavior in a virtual environment. *Presence: Teleoperators and Virtual Environments*, *23*, 33–50. https://doi.org/10.1162/PRES_a_00167

Arnott, S. R., Cant, J. S., Dutton, G. N., & Goodale, M. (2008). Crinkling and crumpling: An auditory fMRI study of material properties. *NeuroImage*, *43*, 368–378.

Barral, O., Aranyi, G., Kouider, S., Lindsay, A., Prins, H., Ahmed, I., . . . Cavazza, M. (2014). Covert persuasive technologies: bringing subliminal cues to human–computer interaction. In A. Spagnolli, L. Chittaro, & L. Gamberini (Eds.), Persuasive technology (pp. 1–12). Berlin, Germany: Springer International Publishing.

Barry, T. J., Griffith, J. W., De Rossi, S., & Hermans, D. (2014). Meet the Fribbles: Novel stimuli for use within behavioural research.

Frontiers in Psychology, *5*, 103. https://doi.org/10.3389/fpsyg.2014.00103

Barry, T. J., Griffith, J. W., Vervliet, B., & Hermans, D. (2016). The role of stimulus specificity and attention in the generalization of extinction. *Journal of Experimental Psychopathology*, *7*, 143–152.

Bennett, M., Vervoort, E., Boddez, Y., Hermans, D., & Baeyens, F. (2015). Perceptual and conceptual similarities facilitate the generalization of instructed fear. *Journal of Behavior Therapy and Experimental Psychiatry*, *48*, 149–155. https://doi.org/10.1016/j.jbtep.2015.03.011

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182. https://doi.org/10.1037/0096-1523.19.6.1162

Bohil, C. J., Alicea, B., & Biocca, F. A. (2011). Virtual reality in neuroscience research and therapy. *Nature Reviews Neuroscience*, *12*, 752–762. https://doi.org/10.1038/nrn3122

Buffat, S., Chastres, V., Bichot, A., Rider, D., Benmussa, F., & Lorenceau, J. (2014). OB3D, a new set of 3D objects available for research: a web-based study. *Frontiers in Psychology*, *5*, 1062. https://doi.org/10.3389/fpsyg.2014.01062

Bülthoff, H. H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences*, *89*, 60–64.

Cant, J. S., & Goodale, M. (2007). Attention to form or surface properties modulates different regions of human occipitotemporal cortex. *Cerebral Cortex*, *17*, 713–731.

Chuang, L. L., Vuong, Q. C., & Bülthoff, H. H. (2012). Learned non-rigid object motion is a view-invariant cue to recognizing novel objects. *Frontiers in Computational Neuroscience*, *6*, 26. https://doi.org/10.3389/fncom.2012.00026

Cooperrider, K. (2015). The co-organization of demonstratives and pointing gestures. *Discourse Processes*, *53*, 632–656. https://doi.org/10.1080/0163853X.2015.1094280

DiCiccio, T. J., & Efron, B. (1996). Bootstrap confidence intervals. *Statistical Science*, *11*, 189–228.

Doucet, G., Gulli, R. A., & Martinez-Trujillo, J. C. (2016). Cross-species 3D virtual reality toolbox for visual and cognitive experiments. *Journal of Neuroscience Methods*, *266*, 84–93. https://doi.org/10.1016/j.jneumeth.2016.03.009

Eggink, J. (2012). kriAlpha [Matlab script]. Retrieved August 30, 2017 from https://www.mathworks.com/matlabcentral/fileexchange/36016-krippendorff-s-alpha

Ehrsson, H. H. (2007). The experimental induction of out-of-body experiences. *Science*, *317*, 1048. https://doi.org/10.1126/science.1142175

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*, 184–188. https://doi.org/10.1038/nature01964

Estes, Z., Jones, L. L., & Golonka, S. (2012). Emotion affects similarity via social projection. *Social Cognition*, *30*, 584–609.

Gauthier, I., James, T. W., Curby, K. M., & Tarr, M. J. (2003). The influence of conceptual knowledge on visual discrimination. *Cognitive Neuropsychology*, *20*, 507–523.

Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: Exploring mechanisms for face recognition. *Vision Research*, *37*, 1673–1682. https://doi.org/10.1016/S0042-6989(96)00286-6

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature Neuroscience*, *2*, 568–573.

Harman, K. L., & Humphrey, G. K. (1999). Encoding "regular" and "random" sequences of views of novel three-dimensional objects. *Perception*, *28*, 601–615.

Harman, K. L., Humphrey, G. K., & Goodale, M. (1999). Active manual control of object views facilitates visual recognition. *Current Biology*, *9*, 1315–1318.

Harris, J. (2015). Yadgits. Retrieved August 24, 2017, from http://wiki.cnbc.cmu.edu/Novel_Objects, last modified May 14, 2015

Hayes, A. F., & Krippendorff, K. (2007). Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures*, *1*, 77–89.

Hayward, W. G., & Tarr, M. J. (1997). Testing conditions for viewpoint invariance in object recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 1511–1521. https://doi.org/10.1037/0096-1523.23.5.1511

Humphrey, G. K., & Khan, S. C. (1992). Recognizing novel views of three-dimensional objects. *Canadian Journal of Psychology*, *46*, 170–190.

Jangraw, D. C., Johri, A., Gribetz, M., & Sajda, P. (2014). NEDE: An open-source scripting suite for developing experiments in 3D virtual environments. *Journal of Neuroscience Methods*, *235*, 245–251. https://doi.org/10.1016/j.jneumeth.2014.06.033

Johnson, L., Sullivan, B., Hayhoe, M., & Ballard, D. (2014). Predicting human visuomotor behaviour in a driving task. *Philosophical Transactions of the Royal Society B*, *369*, 20130044. https://doi.org/10.1098/rstb.2013.0044

Jovancevic, J., Sullivan, B., & Hayhoe, M. (2006). Control of attention and gaze in complex environments. *Journal of Vision*, *6*(12), 9: 1431–1450. https://doi.org/10.1167/6.12.9

Kaldy, Z., Blaser, E. A., & Leslie, A. M. (2006). A new method for calibrating perceptual salience across dimensions in infants: the case of color vs. luminance. *Developmental Science*, *9*, 482–489.

Klein, S. A., & Macmillan, N. A. (2001). Threshold estimation: The state of the art. *Perception & Psychophysics*, *63*, 1277–1278.

Knutson, A. R., Hopkins, R. O., & Squire, L. R. (2012). Visual discrimination performance, memory, and medial temporal lobe function. *Proceedings of the National Academy of Sciences*, *109*, 13106–13111. https://doi.org/10.1073/pnas.1208876109

Krippendorff, K. (2008). Systematic and random disagreement and the reliability of nominal data. *Communication Methods and Measures*, *2*, 323–338. https://doi.org/10.1080/19312450802467134

Krippendorff, K. (2011). Computing Krippendorff's alpha-reliability. Retrieved August 28, 2017, from http://repository.upenn.edu/asc_papers/43

Kujala, J. V., & Lukka, T. J. (2006). Bayesian adaptive estimation: The next dimension. *Journal of Mathematical Psychology*, *50*, 369–389.

Leeb, R., Friedman, D., Müller-Putz, G. R., Scherer, R., Slater, M., & Pfurtscheller, G. (2007). Self-paced (asynchronous) BCI control of a wheelchair in virtual environments: A case study with a tetraplegic. *Computational Intelligence and Neuroscience*, *2007*, 79642. https://doi.org/10.1155/2007/79642

Lenggenhager, B., Tadi, T., Metzinger, T., & Blanke, O. (2007). Video ergo sum: Manipulating bodily self-consciousness. *Science*, *317*, 1096–1099. https://doi.org/10.1126/science.1143439

Macmillan, N. A., & Creelman, D. C. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

Mercer, T., & Duffy, P. (2015). The loss of residual visual memories over the passage of time. *Quarterly Journal of Experimental Psychology*, *68*, 242–248. https://doi.org/10.1080/17470218.2014.975256

Richler, J. J., Wilmer, J. B., & Gauthier, I. (2017). General object recognition is specific: Evidence from novel and familiar objects. *Cognition*, *166*, 42–55.

Scheveneels, S., Boddez, Y., Bennett, M. P., & Hermans, D. (2017). One for all: The effect of extinction stimulus typicality on return of fear. *Journal of Behavior Therapy and Experimental Psychiatry*, *57*, 37–44.

Sederberg, T. W., & Parry, S. R. (1986). Free-form deformation of solid geometric models. In D. C. Evans & R. J. Athay (Eds.), SIGGRAPH '86: Proceedings of the 13th Annual Conference on Computer

Graphics and Interactive Techniques (pp. 151–160): New York, NY: ACM.

Smilek, D., Enns, J. T., Eastwood, J. D., & Merikle, P. M. (2006). Relax! Cognitive strategy influences visual search. *Visual Cognition*, *14*, 543–564.

Tarr, M. J., Bülthoff, H. H., Zabinski, M., & Blanz, V. (1997). To what extent do unique parts influence recognition across changes in viewpoint. *Psychological Science*, *8*, 282–289. https://doi.org/10.1111/j.1467-9280.1997.tb00439.x

Townsend, J. T., & Ashby, F. G. (1983). Stochastic modeling of elementary psychological processes. New York, NY: Cambridge University Press.

Treutwein, B. (1999). Adaptive psychophysical procedures. *Vision Research*, *35*, 2503–2522.

Wallraven, C., Bülthoff, H. H., Waterkamp, S., van Dam, L., & Gaissert, N. (2014). The eyes grasp, the hands see: Metric category knowledge transfers between vision and touch. *Psychonomic Bulletin & Review*, *21*, 976–985. https://doi.org/10.3758/s13423-013-0563-4

Watrous, A. J., Tandon, N., Conner, C. R., Pieters, T., & Ekstrom, A. D. (2013). Frequency-specific network connectivity increases underlie accurate spatiotemporal memory retrieval. *Nature Neuroscience*, *16*, 349–356. https://doi.org/10.1038/nn.3315

Watson, M. R., Voloh, B., Naghizadeh, M., Chen, S., & Womelsdorf, T. (2017). Information sampling and object selection strategies demonstrate the learning and exploitation of feature relevance. In *2017 Neuroscience Meeting Planner* (Program No. 801.30). Washington, DC: Society for Neuroscience.

Weisberg, S. M., Schinazi, V. R., Newcombe, N. S., Shipley, T. F., & Epstein, R. A. (2014). Variations in cognitive maps: Understanding individual differences in navigation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 669–682. https://doi.org/10.1037/a0035261

Williams, P. (1998). Representational organization of multiple exemplars of object categories (Working paper, University of Massachusetts at Boston). Retrieved August 23, 2017, from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.5.8336&rep=rep1&type=pdf

Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behavior*, *1*, 0058. https://doi.org/10.1038/s41562-017-0058

Wong, A. C.-N., & Hayward, W. G. (2005). Constraints on view combination: Effects of self-occlusion and differences among familiar and novel views. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 110–121. https://doi.org/10.1037/0096-1523.31.1.110

Wong, A. C.-N., Palmeri, T. J., & Gauthier, I. (2009). Conditions for face-like expertise with objects: Becoming a Ziggerin expert—but which type? *Psychological Science*, *20*, 1108–1117. https://doi.org/10.1111/j.1467-9280.2009.02430.x

Zapf, A., Castell, S., Morawietz, L., & Karch, A. (2016). Measuring inter-rater reliability for nominal data—Which coefficients and confidence intervals are appropriate. *BMC Medical Research Methodology*, *16*, 93. https://doi.org/10.1186/s12874-016-0200-9

Appendix D. USE: An integrative suite for temporally-precise psychophysical experiments in virtual environments for human, nonhuman, and artificially intelligent agents

Journal of Neuroscience Methods 326 (2019) 108374

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

journal homepage: www.elsevier.com/locate/jneumeth

USE: An integrative suite for temporally-precise psychophysical experiments in virtual environments for human, nonhuman, and artificially intelligent agents

Watson Marcus R[a,*], Voloh Benjamin[b], Thomas Christopher[b], Hasan Asif[c], Womelsdorf Thilo[a,b,c]

[a] Department of Biology, Centre for Vision Research, York University, Toronto, ON, M6J1P3, Canada
[b] Department of Psychology, Vanderbilt University, Nashville, TN, 37240 United States
[c] Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN, 37240 United States

ARTICLE INFO

Keywords:
Psychophysics
Software
Cognition
Translational neuroscience
Behavioral control
Unity 3D
Stimulus presentation
Vision
Virtual reality
Video games
3D
Virtual environment
Experimental design

ABSTRACT

Background: There is a growing interest in complex, active, and immersive behavioral neuroscience tasks. However, the development and control of such tasks present unique challenges.
New Method: The Unified Suite for Experiments (*USE*) is an integrated set of hardware and software tools for the design and control of behavioral neuroscience experiments. The software, developed using the Unity video game engine, supports both active tasks in immersive 3D environments and static 2D tasks used in more traditional visual experiments. The custom USE SyncBox hardware, based around an Arduino Mega2560 board, integrates and synchronizes multiple data streams from different pieces of experimental hardware. The suite addresses three key issues with developing cognitive neuroscience experiments in Unity: tight experimental control, accurate sub-ms timing, and accurate gaze target identification.
Results: USE is a flexible framework to realize experiments, enabling (i) nested control over complex tasks, (ii) flexible use of 3D or 2D scenes and objects, (iii) touchscreen-, button-, joystick- and gaze-based interaction, and (v) complete offline reconstruction of experiments for post-processing and temporal alignment of data streams.
Comparison with Existing Methods: Most existing experiment-creation tools are not designed to support the development of video-game-like tasks. Those that do use older or less popular video game engines as their base, and are not as feature-rich or enable as precise control over timing as USE.
Conclusions: USE provides an integrated, open source framework for a wide variety of active behavioral neuroscience experiments using human and nonhuman participants, and artificially-intelligent agents.

1. Introduction

1.1. Static and active tasks

Participants in most traditional psychology or neuroscience experiments are presented with impoverished stimuli and tightly-restricted response options. This approach, which for the sake of brevity we refer to as *static*, maximizes experimental control, constrains possible interpretations of results, allows for the comparatively easy creation of experimental tasks, and is a critical part of the reductionist approach that has led to many of the exceptional successes of the cognitive sciences. However, researchers are increasingly concerned with understanding behaviour and neural processing in tasks and contexts that are more complex and naturalistic, an approach we refer to as *active*. Here we present an integrated suite of software and hardware designed to aid in the development, control and analysis of active cognitive neuroscience experiments, the *Unified Suite for Experiments* (*USE*).

Given the success of static tasks and their dominance in the cognitive sciences, why should researchers bother with active tasks? We provide three answers to this question. First, there is the common concern about ecological validity and the generalizability of results, namely that results from simple static tasks do not necessarily generalize and may be misleading when taken to provide insight into real-world behaviour and neural activity (Kingstone et al., 2008; Schmuckler, 2001; Chaytor and Schmitter-Edgecombe, 2003). For example, there are substantial quantitative and qualitative differences in biases for gaze to be directed towards others' eyes that depend on whether participants are viewing static pictures, watching movies, or are involved in genuine interactions with other humans (Risko et al.,

* Corresponding author.
E-mail address: watsonmr@yorku.ca (M.R. Watson).

https://doi.org/10.1016/j.jneumeth.2019.108374
Received 9 January 2019; Received in revised form 24 June 2019; Accepted 24 July 2019
Available online 25 July 2019
0165-0270/

2012). By facilitating the generation of both static and active tasks using the same underlying task structure, USE enables the direct investigation of differences between them.

Second, more naturalistic stimuli and flexible possibilities of control can be more immersive and dramatically more motivating for human and nonhuman primate subjects (Bennett et al., 2016; Bouchard et al., 2011; Slater and Wilbur, 1997; Witmer and Singer, 1998; Youngblut, 2007). As one testament to this, over the decades where video games developed more realistic visuals and more flexible and responsive controls, they developed from an obscure hobby to a multi-billion dollar industry with more consumers and larger revenues than film-making (The NDP Group, 2009; Shanley, 2017). As another testament, more realistic computer game environments can lead to reliable increases of learning outcomes in variety of contexts (e.g. Mayer, 2018).

Finally, active tasks enable the collection and exploration of a wide array of precise and multi-modal data about complex behaviors, which are necessary to generate hypotheses for understanding real-world behaviors (Kingstone et al., 2008). Despite more than two centuries of formal psychological research and millennia of informal speculation, there is no well-established body of fine-grained data about human action in most tasks. By providing such fine-grained data on realistic tasks, active tasks can produce data that enable both hypothesis-testing and exploratory or observational work, sometimes even in the same experiment.

1.2. Developing active tasks with USE

Active tasks may be appealing, but they are more challenging to develop, control, and analyze than static tasks. USE was developed to make these challenges more manageable. The practicality and scope of USE make it a versatile alternative to experimental creation and control suites such as the Psychtoolbox (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997), PsychoPy (Peirce, 2007, 2008), or MonkeyLogic (Asaad and Eskandar, 2008a, b; Asaad et al., 2013; Hwang et al., 2019), with the specific focus of creating, controlling and analyzing tasks that have the complexity, dynamism, and visual fidelity typical of video games. USE also provides unique solutions to common challenges of experimental control (see 5.2) that set it apart from other active experiment creation suites (Brookes et al., 2018; Doucet et al., 2016; Jangraw et al., 2014).

USE provides an integrated solution for development, timekeeping, and analysis components required for typical experiments (Fig. 1). The remainder of this paper delineates each of these components. For the development component, a set of scripts for the video game engine Unity enables the development of nested hierarchies of experimental control, as well as tools for common experimental requirements (data collection and recording, communication with other programs and equipment, etc). For the timekeeping component, the suite incorporates the *USE SyncBox*, an Arduino-based timing and I/O hardware system, to relay signals and codes between experimental hardware, and to maintain a central time record. This can send either event codes or simple pulses to other experimental equipment, and track the current monitor status using light sensors placed on the monitor, which allows the detection of any skipped or stuck frame (currently only in post-processing). Finally, for the analysis component, the suite includes a set of Matlab scripts for offline data parsing and timestream synchronization.

USE allows flexible experimental protocols built around (i) active or static experiments, (ii) 2D and 3D displays, and (iii) touchscreen, joystick or keyboard/button press interfaces. In addition, it provides an interface for artificial agents that thereby can be tested with the identical settings used for experiments in humans or nonhuman primates. We anticipate that the key strength of USE is to facilitate the design and realization of active studies, though it can also be used to generate static tasks, as we have shown in prior studies (Oemisch et al., 2017; Watson et al., in press), enabling the comparison of active and static task variants.

Links to the key components of USE, as well as manuals, tutorials, and example experiments, can be found on our website (http://accl.psy.vanderbilt.edu/resources/analysis-tools/unifiedsuiteforexperiments).

1.3. Goals of USE

In developing USE, we aimed to satisfy multiple criteria. The goal was to develop a system that was:

- *Temporally accurate and precise* - all data should synchronize at millisecond precision.
- *Modular* - all software and hardware components are developed so that they can be used independently, or in different combinations.
- *Generic* - specific components can be easily adapted to multiple purposes.
- *Algorithmic* - each component of the system is explained in a principled way to facilitate implementation in other research contexts using other hardware (e.g. a different microcontroller board) or software (e.g. a different video game engine).
- *Translational* –functionally identical protocols can easily be generated for use with different groups of humans, non-humans, or artificially intelligent agents.
- *Offline reconstructable* - every monitor frame displayed during an experiment can be recreated at will in offline analysis and combined with synchronized information from any other data stream to reconstruct the experiment.
- *Cost efficient* - the software is all free, the total cost of the custom hardware is under $500, and experiments can be developed and run using any modern computer and consumer-grade monitor.
- *Multi-platform* - experiments can be developed on Mac OSX or Windows computers, and run on any modern computer (including Linux), smartphone, tablet, or gaming console, and can also be developed for the web.
- *Portable* - only a single computer, a small box for the arduino, two light sensors and a small number of cables are required for complete experimental control.
- *Practical* - the suite solves the key challenges in active experiment implementation and organization, and does so in ways that are intended to be as user-friendly as possible.
- *Open source* - all components of development and analysis software and I/O data streaming should be freely available under open source licenses.

One criteria we did *not* aim to satisfy was that developing experiments be a completely novice-friendly process. In our experience, suites which attempt to do so end up sacrificing generality, flexibility, and power. Thus, designing an experiment in USE, while it is much easier than doing so from scratch, does require a level of familiarity with both Unity and general programming principles. Actually running experiments created with USE, however, does not require any special knowledge and can be simple enough to be the responsibility of the undergraduates who are tasked with running more traditional studies in many laboratories (see 4.2.7). That said, we expect as more labs become involved in the development of active tasks, many programmatic components that define the actionable environment will be re-used as they are developed and shared through open-source collaboration, increasing the ease of experimental design.

1.4. The Unity game engine

The experiment development and control software components of USE are implemented in the Unity *game engine*. Game engines are development environments that implement various functionalities that games commonly require, such as the rendering of 2D and 3D graphics, physics simulation, realistic lighting, animation, sound management, etc. Unity is a free (but not open source) engine that runs on Windows,

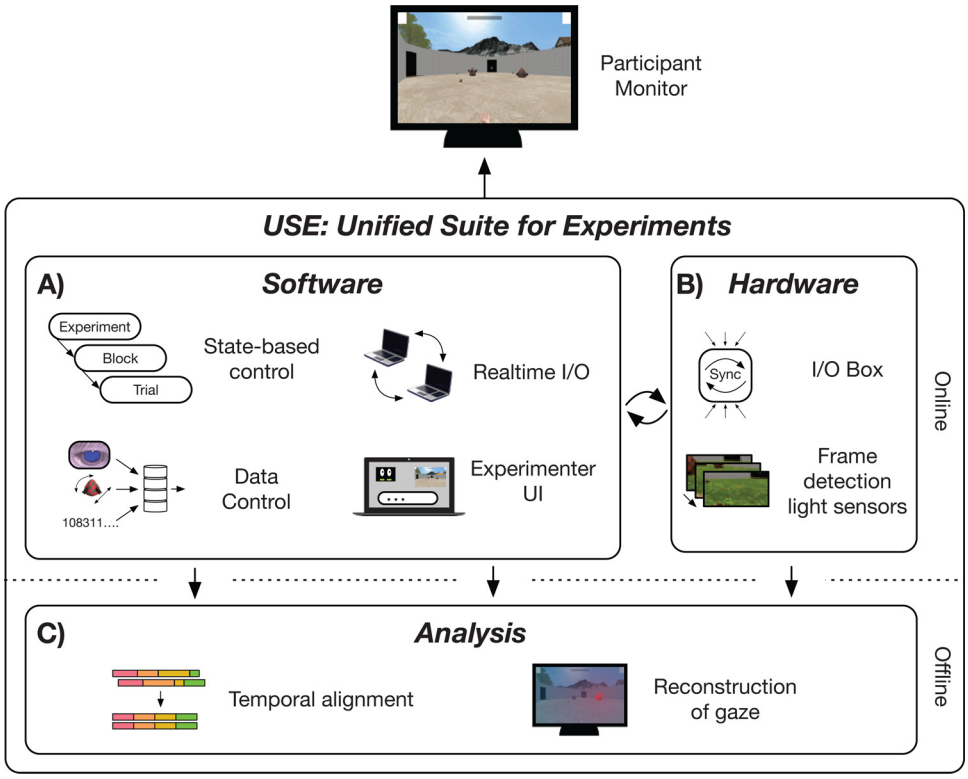

Fig. 1. Unified Suite for Experiments allows strong experimental control in immersive tasks. Immersive, dynamic tasks present unique challenges to experimental control. To address these, USE includes novel contributions for (A) software, (B) Hardware, and (C) analysis tools. **(A)** USE employs a novel state-based control system that underpins its strong experimental control. Realtime I/O is based on current industry standards. USE also allows experimentalists to flexibly define and collate data occurring at multiple experimental timescales, and online visualization of subject's performance. **(B)** The USE SyncBox is a newly developed, dedicated, and inexpensive machine that allows synchronization of different data streams (such as joystick data, eyetracker data, and game data), and can communicate with other hardware. As well, custom built photo-diode solution allows for the precise detection of physical frame onsets. **(C)** USE also contains a number of post-hoc analysis scripts, including a way to align all data sources, and reconstruct subject's gaze.

Macintosh, and Linux computers, supports building games for all major computers, phones, tablets, and game consoles, and has built-in support for stereoscopic presentation. Games made with Unity were downloaded over 5 billion times in Q3 2016 alone (Unity Technologies, 2017a). Some recent games of note made with Unity include Cites: Skylines (Colossal Order, 2015), Her Story (Barlow, 2015), Kerbal Space Program (Squad, 2015), Pokémon GO (Niantic, 2016), and Super Mario Run (Nintendo EDP, 2016). Aside from its price, Unity is attractive to small and medium-sized developers due to its relative ease of use, full feature set, comprehensive tutorials, and active online forums (Unity Technologies, 2017b). USE scripts are all written in C#, which as of 2019 is the only language supported by Unity.

1.5. Specific challenges of active tasks

Building an experimental suite on top of a game engine presents several unique challenges. First, game engines must ensure that all required commands are run in time for frame rendering to be carried out, but the precise execution order below the frame rate is often irrelevant, and commands can be spread across many scripts, making it difficult to follow their interactions. Command execution order in USE is guaranteed by a novel state-based system (cf. Wagner et al., 2006) that enables the generation of nested hierarchies of control, such as the experiment-block-trial structure common in cognitive studies (section 4.2). As an example of the general utility of state-based control systems for experiment development, a state-based system has been used to create electrical engineering experiments in Unity exists (Liao and Qu, 2013), but its use case is sufficiently distinct from cognitive neuroscience that it is untransferable.

Second, just as control below the level of a frame can be difficult to achieve, accurate timing below the level of a frame is, for most practical purposes, unavailable within Unity. Indeed, it is difficult even to ascertain frame onset times, complicating synchronization of displays with other experimental hardware such as neural acquisition devices. To solve this, USE incorporates a newly designed, dedicated SyncBox hardware, which acts as a central timekeeper, a generic communication/synchronization device, and also uses data from light sensors to track the current frame status (section 4.3).

Finally, for studies that incorporate eyetracking, identifying gaze targets in a 3D scene rendered on a 2D screen is difficult, particularly when objects have complex shapes and the subject freely navigates through the 3D environment. In most static tasks, an Area of Interest (*AOI*) is defined around each object, with a larger radius than the object itself to account for measurement imprecision and the spatial extent of the fovea. This is much more difficult and computationally-intensive in active tasks, as objects' two-dimensional silhouettes on the screen can be difficult to determine, and can change drastically from frame to frame. Our solution involves a novel method named *Shotgun Raycast*, which detects all objects whose two-dimensional silhouette falls within a specified number of degrees of visual angle from a gaze point on the screen (section 4.4).

These novel contributions extend Unity's robust physics and rendering capabilities, and make it a suitable platform for behavioral neuroscience research.

1.6. Enabling translational research: human, non-human, and artificially intelligent agents

Our laboratory uses healthy human undergraduates, neurosurgery candidates, macaque monkeys, and artificial (reinforcement) learning agents to run tasks with a variety of different input mechanisms and recording devices. It was critical for us that USE be able to generate comparative task structure and data across all these groups. USE enables this in several ways. First, experimental protocols can be adjusted to meet the needs of different participant groups, such that different protocols using the same functional logic can be run from the same base scripts with small changes to configuration files (see 4.2.3). Second, an *input broker* enables entirely different inputs (for example, touchscreen presses, button clicks, fixations, or selection by an artificial agent) to produce the same effect in a trial (see 4.2.4). Finally, an artificial intelligence (AI) wrapper enables two-way communication with different learning agents (see 4.2.5), similar to existing AI testing platforms (e.g.

Beattie et al., 2016; Leibo et al., 2018; OpenAI, 2018). If desired, full-resolution screenshots can be sent on every frame to visually-guided agents that can learn fully active tasks, as have proved influential in recent years (e.g. Mnih et al., 2015; Wang et al., 2018), as well as more traditional models that operate on simple digital vectors that represent static features of interest in the scene (e.g. Kruschke, 1992)

1.7. Online resources

Three Supplemental files are hosted with this paper:

- A review of the revised gaze classification algorithm used in our task (see 4.4.3 and 5.4).
- A detailed description of the timing tests performed on the USE SyncBox (see 5.1).
- A description of the normalization procedure used to produce Fig. 6.

As previously noted, all USE manuals, scripts, schematics, and tutorial are available via the USE website (http://accl.psy.vanderbilt.edu/resources/analysis-tools/unifiedsuiteforexperiments/).

1.8. Outline

In the remainder of the paper we outline an experiment that provides specific examples of the various components of USE in action (4.1), before turning to an overview of USE's experimental design and control software (4.2), timing and communication hardware (4.3), and offline analysis capabilities (4.4). We then present results demonstrating the robust timing and control capabilities of USE (5) and discuss the implications of this work (6).

2. Methods

2.1. Example experiment

Static tasks built with USE have already been published (Oemisch et al., 2017; Watson et al., in press), but to demonstrate more of the suite's capabilities, we briefly present an active learning task, in which participants must navigate through a 3D scene, interact with objects, and try to learn a reward structure. A full overview of the task and results will be presented in another manuscript (Watson et al., in press), here we focus on reporting sufficient detail to show USE's capabilities.

Figs. 2a and b display the task environment and trial protocol. In brief, participants must choose between two objects on each trial, one of which is rewarded and one of which is not, and they must learn the rules that govern reward through trial-and-error. Trials consist of participants navigating through an arena using a joystick, picking up one of the two objects by walking over it, and taking it to a door where they receive visual feedback on choice accuracy. The two objects are selected from a pool of 16 (Fig. 2c, cf. Watson et al., in press), which are parametrically-defined by four feature dimensions (shape, colour, pattern, and arm direction) of two possible values each. In each block there are two rules, both of which reward a single feature value, but each of which only applies in a single context (symbolized by the arena floor). For example, objects with pyramidal body shape might be rewarded on grass floors, while objects with downward-bending arms rewarded on marble floors (Fig. 2d). Once participants have achieved satisfactory performance on a block, new rules are generated, and a new block begins. Participants are explicitly informed of each block change, and told that they will now have to discover a new pair of rules.

50 participants ran in the study for approximately one hour each. Two participants were excluded due to chance performance. Gaze was tracked using a Tobii TX300 combined head-free eyetracker and OLED monitor (Tobii Technology, Inc.). Fig. 2e demonstrates that over the course of a block they learned to make more accurate choices, to make these choices more rapidly, and to preferentially fixate the rewarded object over the unrewarded one (see 4.4.3 for details of how fixation targets are determined in USE).

Fig. 2F surveys the display configuration, structure, and response mode of several of the other experiments that have already been coded in USE.

2.2. USE: Unity Suite for Experiments

2.2.1. States and Control Levels

Most software commands in USE are controlled by an architecture of *States* and *Control Levels*, which allow experiments to be defined as a series of *hierarchical finite state machines* (Wagner et al., 2006). This State-Level architecture is flexible enough to support any standard experimental hierarchies such as experiment-block-trial, as well as more complex structures. Here, we explain their abstract functional role and give examples of their use in the experiment described in 4.1

A State is an object (in the object-oriented programming sense) that defines the operation of the experiment during a period of time (Fig. 3a). These are organized in Control Levels that group States that operate at similar levels of abstraction in the experiment (Fig. 3b). In the example experiment, a Trial Level groups together the States that define the various trial epochs, while a Block Level groups together States that define the block sequence, and passes control to the lower Trial Level as needed, and an even higher Main Level groups the States that govern aspects of the experiment outside the Block (Fig. 3c). Thus, States can be used to define both the finest-grained and the most general parts of the experiment.

Each State specifies commands that are run every frame while it is active, and initialization and termination commands to be run at the start of its first and the end of its last frame, respectively. For example, the Fixation State of the example experiment's Trial Level determines what happens when the participant stands in front of a door before entering the arena. At the start of the first frame in which an experiment enters a new State, a StateInitialization method runs. For the Fixation State, this method turns a circle on the door blue, cueing the participant to look at it. Within each frame of a State, a number of method groups run, controlling Unity's update cycle (for full details of the update cycle, see the USE manual). For the Fixation State, these methods control the size of the blue circle, which shrinks as it is fixated, and the position of the door, which opens after the circle has completely shrunk. The update cycle is followed by StateTerminationCheck methods that verify whether any end conditions for the state have been reached. For the Fixation State, there is a single condition: is the door fully open? If one of the end conditions has been met, a StateTermination method group is run at the end of the last frame of the state, and the experiment transitions to a new State on the following frame. If no end conditions have been met, the current State's update cycle methods run again on the following frame. For the Fixation State, the StateTermination method starts opening the door, revealing the main arena behind it, and on the following frame the Trial Level transitions to the Explore State.

States are grouped together in Control Levels (Fig. 3b). Any State in a Control Level can transition to any other State in the same Level. These transitions are defined by TerminationSpecifications, each of which includes (a) a StateTerminationCheck, (b) a StateTermination, (c) a successor State, and (d) the StateInitialization of the successor State that will be run. By defining the States that make up a Level, and the desired transitions between them, an experimenter has defined the finite state machine that constitutes that Level.

Like States, Control Levels have Initialization and Termination methods. LevelInitialization and LevelTermination methods run once each, at the start of the first frame (prior to any StateInitializations) and end of the last frame (after any StateTerminations) in the Control Level, respectively.

2.2.2. Hierarchical State/Level organization

A critical aspect of USE is that *Control Levels can be children of States*,

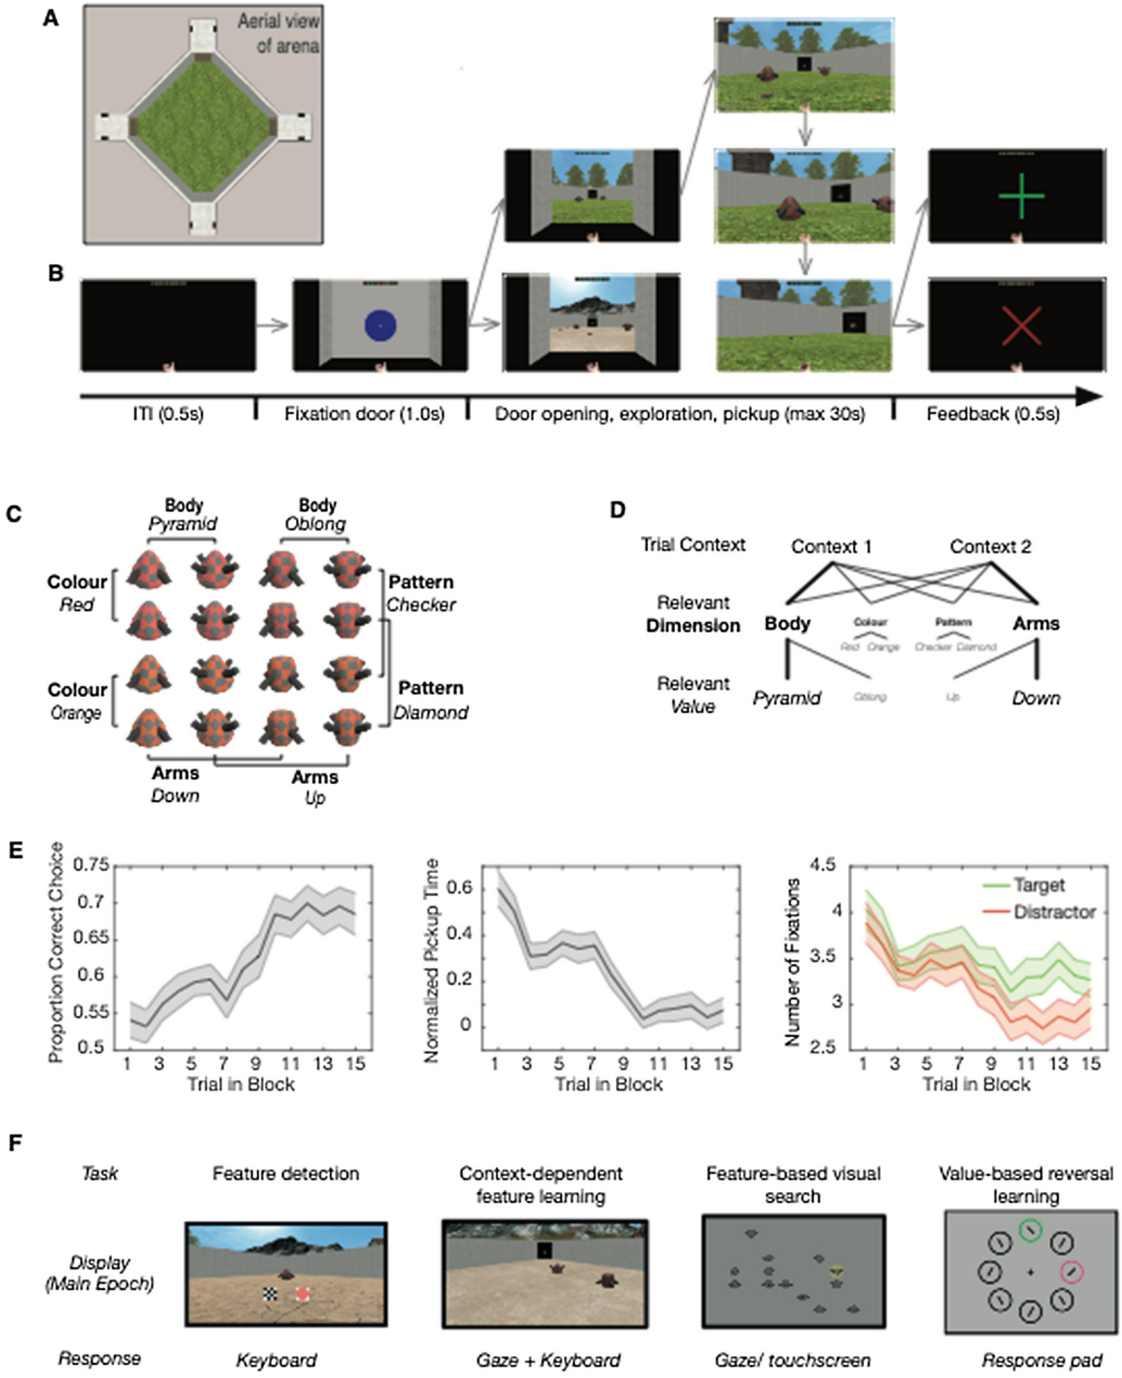

Fig. 2. Example USE experiment. Details of one experiment coded in USE, and brief sketches of several others **(A)** An overhead view of the task arena. **(B)** An example trial sequence. Trials are separated by an ITI where most of the screen is black, but the avatar's hand and their reward history is visible. Participants begin in one of the four corridors at the corners of the arena, facing a closed door. Fixating a white dot at the centre of a larger blue circle on the door opens it, revealing the arena. Participants then navigate towards one of the two objects in the arena and pick it up by walking over it. Finally, they take the object to any of the other corridor doors, where they are given visual feedback on their choice accuracy. **(C)** Set of objects used in the experiment (Watson et al., in press). Each object is defined by four feature dimensions (body shape, color, arm position, and pattern), each of which can take on two specific values (e.g. pyramidal or oblong, red or orange, etc). **(D)** An example reward structure that participants must learn in a given block. Objects are presented in one of two contexts, defined by the floor of the arena. Each context has an associated relevant feature dimension, and one rewarded feature value in that dimension. In this example, the objects presented in context 1 are rewarded if their body (feature dimension) is pyramidal (feature value). Thick lines and large font highlight the path towards the rewarded feature value, whereas small lines and font denote a path towards unrewarded values. **(E)** Accuracy, time from door opening to object pickup (normalized by dividing by the distance from door to object, and z-scoring within each participant), and number of fixations pre-pickup to rewarded target objects and unrewarded distractor objects. Shaded areas indicate standard error of the mean. **(F)** Other examples of experiments coded in USE. (Top) These tasks can have richly detailed or simple stimuli, in various degrees of interaction between player and environment. (Bottom) USE can support different inputs as required by the specific experiment.

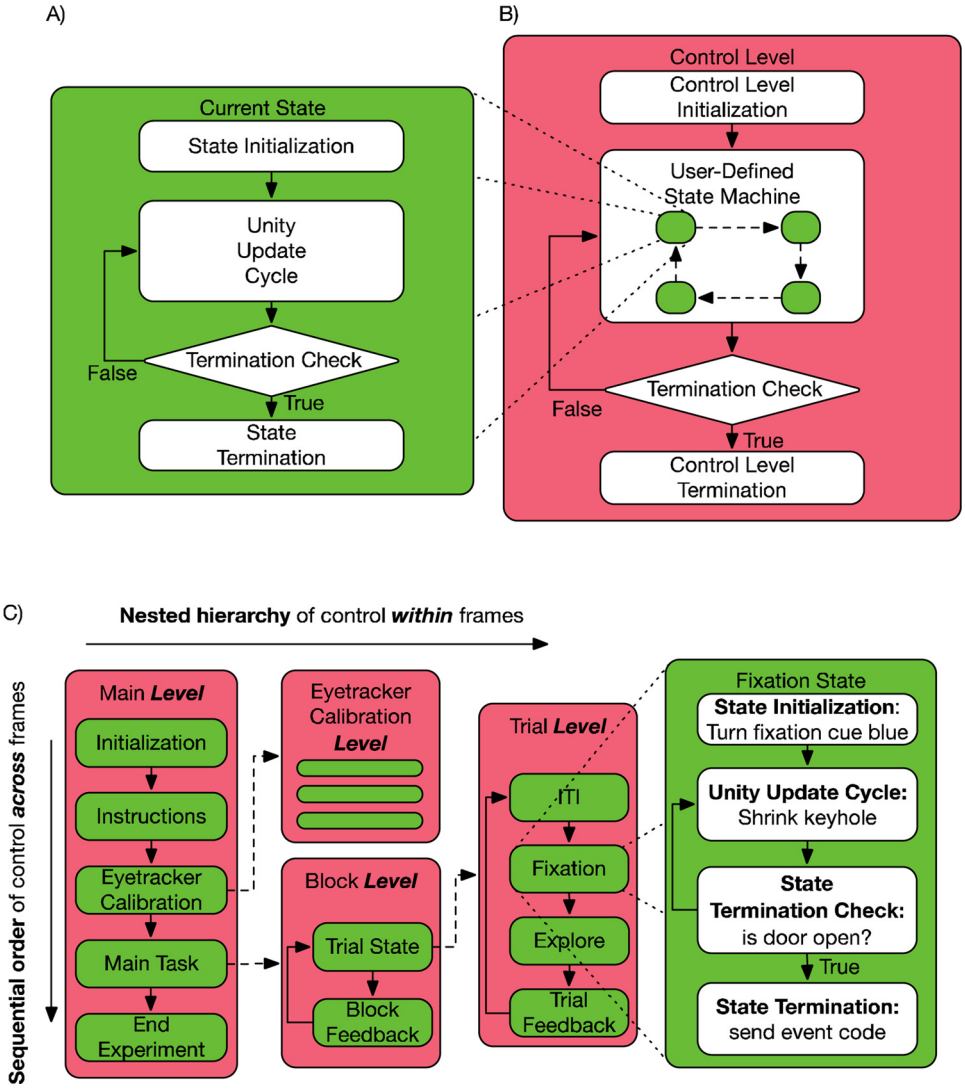

Fig. 3. State/Level organization allows for flexible and strong experimental oversight. Depiction of (A) State and (B) Control Level architecture and their interaction in (C) an example experiment. White rounded rectangles denote groups of user-defined methods, diamonds represent Boolean conditions. Dashed lines with arrowheads represent the relationship of parent States and their child Control Levels. Dotted lines from all four corners of a State illustrate the zoomed-in view of the internal components of the state itself (not of child Control Levels). **(A)** A State contains an Initialization method group that runs only once at the start of the first frame in which the State is active. The Unity update cycle runs each frame until a Termination Check returns True, at which point the Termination method is run. (Actual States are somewhat more complex than illustrated, as it is possible to have an arbitrary number of Initializations, Termination Checks, and Terminations.) **(B)** A Control Level defines the States that can transition to each other. An Initialization method runs once the first time a Control Level becomes active. Only one state is active at any one time, and its Update/Termination Check cycle runs during each frame, after which the Control Level's Termination Check is run, and if it returns true the Control Level's Termination method runs. **(C)** Control flow of an example experiment, illustrating that Control Levels can be children of States. (This is only an illustrative portion of the full task hierarchy.) This allows nested control within frames, but also guarantees control across frames. The Main Level is the highest level of control. After completing an Initiation and Instruction State, the Level enters a Calibration State, which contains a lower-level Sequence that itself has multiple states. After completing this, the Main Task state is run, which also has a nested Block Level, which itself has a Trial State and associated Trial Sequence. Here, we explicitly illustrate the Fixation State within the Trial Sequence, and the processes operating during this State.

allowing Control Levels to pass control to sub-Levels. This guarantees that commands run within the experimenter-defined order, ensuring that the subject's experience is as the experimenter had intended.

The State/Level architecture governing the example experiment (4.1) is visualized in Fig. 3c. The Main Level is the top level of the hierarchy, which defines high-level States such as setup, eyetracker (gaze) calibration, and the task itself. The eye tracker State has an associated Eyetracker calibration Level, which itself has multiple States associated with it (presentation of the calibration dot, analysis of the estimated gaze location, experimenter acceptance or rejection of the calibration results, etc). Once this has been completed, we progress to the Main Task state. This has embedded within it a Block Level, which itself has a Trial State that passes control to a Trial Level. In this example, there are four trial States – inter-trial interval, fixation, exploration, and feedback — each of which will include its own initialization, update and termination methods. We have described these in detail for the Fixation State above.

Once all trials in a block have been completed, the Block Level's TerminationCheck verifies if there are blocks remaining in the experiment. If so, the Block Level loops back to the beginning, and a new round of trials begin, but if not, the Block Level ends and the Main Level shifts to an End Experiment State, which performs various housekeeping functions before closing the application.

We anticipate that the large majority of experiments will use at least a two-level hierarchy of Block and Trial Level, and that more complex experiments will benefit from a top-level Main Level, as used in the example experiment. State and Level logic provides a flexible way of designing experimental hierarchies, while maintaining precise control over the order of commands during Unity's primary update loops.

2.2.3. *Selective reuse of States for different experimental protocols*

It is quite common, particularly in translational research, for modifications of the same basic experimental protocol to be used in different experiments. For example, non-human primates typically require a fluid or pellet reward in addition to visual feedback for correct performance, while artificial agents might require only an abstract numeric representation of stimuli, without any visual input at all. Such modifications are often implemented by copying and pasting large chunks of code between experiments, which inevitably results in unintended differences, as later changes do not propagate across the different versions. Another solution is to gate portions of code using *if* or *switch* statements, which quickly makes code extremely hard to follow as the

number of cases mounts.

USE simplifies this process by enabling experimenters to define more States than are necessary for a single experiment, and to select which of these States will be included in a given Control Level at runtime simply by listing their names when defining States. Thus, it is easy to generate highly similar experiments that differ only in a few trial or block States, guaranteeing that the logic they share will remain constant even after later edits to code.

2.2.4. Input broker

Changes to experimental protocols often require very different actions to have similar effects. For example, our laboratory has run studies where objects are selected by keyboard button presses, touchscreen touches, fixations, or with the output of an artificial learning agent (*see* Fig. 2f). A dedicated input broker is used to handle all of these cases and can be customized to support whatever input methods are needed for a given experiment.

2.2.5. AI wrapper

Artificial learning agents implemented in any language can interact with USE and play through the same experiments as human or animal participants, using the suite's AI wrapper. These agents can act based on representations of the environment that consist either of a numerical vector (e.g. different object feature values would be encoded as different numbers), or a screenshot of the current frame. To integrate an AI with a USE task, three core functions must be implemented: 1) Reset – starts/restarts the task, and configures the wrapper to use numerical representation or screenshots to represent the environment. 2) Step - moves the environment to the next step, and outputs its numerical or image representation. 3) Act - takes an action as an input parameter and returns the result of applying this action on the environment, including reward value and indicators of whether the current trial, block, or experiment has ended.

The USE suite includes a Unity-hosted TCP server to serve AI player requests, and a python library that implements a TCP client. Any python-based artificial learning agent that can operate using the Step and Act functions described above can interface with the client, and thus run any tasks developed in USE. The tutorial includes a simple python agent as a demonstration of these capabilities.

2.2.6. Data control

USE incorporates a generic DataController class that enables the flexible collection and writing of as many data streams as may be desired, and writes these to text files. These can include variables of any type, and each DataController object is independently controlled, meaning that data can be collected and written at different frequencies. For example, it makes sense to collect positions of the camera and moveable objects every frame, while trial accuracy might only be updated once every few seconds, and block-level information might be generated every few minutes.

Importantly for timing and later analysis purposes, in our studies we generate FrameData files, which are updated every frame and include the positions, rotations and sizes of all objects for which these values can change (see 4.4.2), as well as the current expected state for each of two flickering patches beneath the light sensors used for timing alignment (see 4.4.1).

2.2.7. Initialization screen and experimenter view

When running a USE experiment, our laboratory employs a two-monitor setup, one for the experimenter and one for the participant (see Fig. 4). The experimenter's screen allows various factors to be specified both at the start of an experiment and during its runtime. These are intended to allow individuals who may not have the programming experience to develop a task to nevertheless be able to control it at runtime, as needed in many laboratories.

At the start of an experiment, a customizable initialization screen is displayed. This can include file selection dialog panels enabling the selection of configuration files, as well as text boxes, Boolean check boxes, and numeric sliders, to specify information that might be desired for the experiment (e.g. subject ID, condition, or duration of different trial States). Each of these can be set to display the previously-chosen value as a default, or some function on this (e.g. subject numbers can be set to automatically iterate with each session).

Throughout an experimental session, the experimenter's monitor displays three main components: (1) a panel showing a stream of exactly what the participant is viewing on their monitor, with overlaid real-time gaze or touch positions if desired, (2) text panels displaying real-time information of various kinds, such as messages from hardware, or summary information about participant performance, and (3) text dialog boxes, Boolean check boxes, and sliders. The interactive components in (3) can be used to modify experimental variables in real-time.

All components of both the initialization and experimenter view can be customized using Unity's editor and external configuration files.

2.2.8. Realtime I/O

USE handles real-time communications with other software and hardware via SerialPortController and UDPPortController objects, which include methods for setting up ports, reading incoming data from system buffers, storing it in USE-specific buffers for use by other methods, and clearing these buffers. For instance, we communicate with a python script controlling our eyetracker via UDP, and with the SyncBox via serial. One special case of serial data are event codes, which are handled by a specific EventCodeController and sent to the SyncBox. Typically, codes are prepared a frame in advance, and sent as quickly as possible after the new frame onset. However, Unity's internal limitations mean that the latency between frame onset and event codes is less stable than might be desired, and for millisecond precision event code timestamps are adjusted with automatized scripts in offline post-processing (see 5.3).

2.3. The USE SyncBox – timekeeping and communication

The USE SyncBox is used to channel communication between experimental hardware during an experimental session, and to generate a single, highly accurate timing record that enables the time-alignment of the various data streams produced by these different pieces of hardware, as well as the physical onset of frames on an experimental display. It consists of a commercial Arduino Mega2560 r3 board with custom firmware and a custom shield connecting it to a number of I/O ports. Fig. 4 shows how the box might be used in a typical experiment.

In the laboratories we have used the SyncBox, it has been used to send event codes to different neural acquisition devices (Neuralynx, BrainAmp, and Neuroscan systems), simple timing pulses to Tobii eyetrackers, and commands to control fluid and pellet dispensers used to reward non-human primates for performance on tasks. It has also been used to receive signals from the light sensors described in 4.3.1, as well as from a custom joystick. Other I/O capabilities can be added by modifying the firmware, and, if necessary, building an adapter to modify the layout of the event code lines. The SyncBox is thus a multi-purpose I/O device that can be used in a variety of experimental settings and quickly adapted between setups.

The two-way connection to the experimental computer is over a USB serial port. There are eight digital I/O channels connected to BNC jacks, five analog input channels also connected to BNC jacks, and 16 digital output channels connected to a 34-pin rectangular connector. Two of the analog inputs are dedicated to receiving data from light sensors attached to the participant monitor, which are fed through a pre-amplifier circuit which performs amplification, low-pass filtering, and DC subtraction on the light sensor signals (see the Syncbox Manual). In the example experiment (4.1), one digital channel was used to output a timing pulse every second that was received by a Tobii

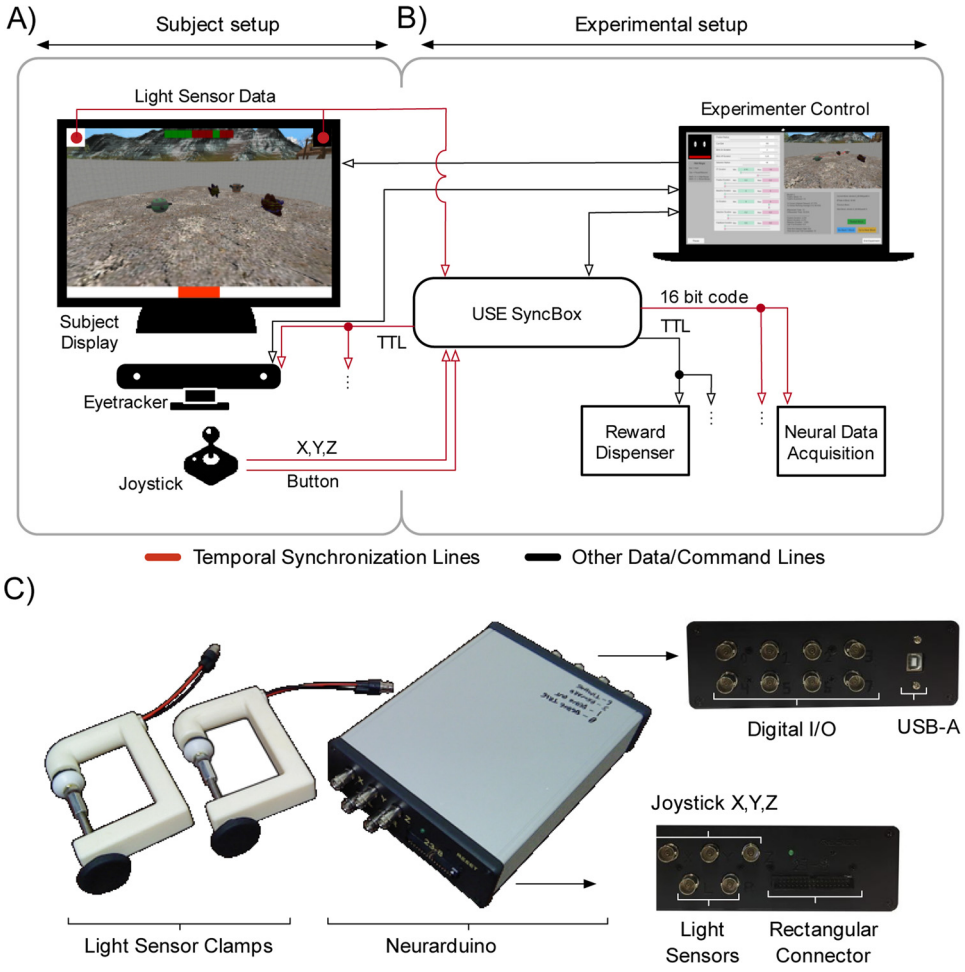

Fig. 4. Example Connectivity diagram showing full capabilities of USE. A typical setup employing the full USE suite. **(A)** The participant setup (left) are those components of the experiment visible to participants, including the monitor, joystick, eyetracker, and light receptors. The virtual environment is displayed on the monitor, participants move through it using the joystick, and the eyetracker records their gaze behaviour. The monitor has patches in the top corners that change between black and white to indicate frame changes, which are picked up by the light receptors. The monitor is controlled directly by the experimental computer over its input line. Joystick output and photo-diode signals are communicated to the USE SyncBox. **(B)** The experimental setup includes the experimental computer, USE SyncBox, neural data acquisition device, and reward dispenser. The experimental computer runs the experiment, controls the monitor, and sends commands to the USE SyncBox. The USE SyncBox can forward event codes for time synchronization to the neural acquisition device, TTL pulses as needed to control a reward dispenser, or regular TTL pulses used for time synchronization to the eyetracker (TTL pulse). The experimental computer is directly connected to the participant display and eyetracker (1), all other communication is funneled through the USE SyncBox, whether from control computer to the peripheral device (2) or from the peripherals to the control computer (3). Lines denote communication from one component to another, with the arrowhead showing the direction of signals. Red is used to highlight those lines that allow post-hoc temporal alignment. **(C)** Photos of a USE SyncBox, its ports, and the custom clamp used to hold light receptors in place over participant monitors. All files and instructions needed to create these are available on the USE website.

TX300 eyetracker, while the rectangular port was used to send event codes to a Neuralynx Digital SX Lynx SDSata acquisition system. Full details of all hardware are available on the USE website.

Full firmware details and code are also available on the USE website. In brief, the box runs an *interrupt-driven loop* every 0.1 ms, and a *host communication loop* that runs as quickly as possible, but without guaranteed timing. The host loop controls all communication over the serial port with the experimental computer, The interrupt-driven loop reads from and writes to all other inputs/outputs and maintains information about scheduled events.

The host computer can issue commands to the box over the serial port to control most aspects of its functioning, including event timing and the interval at which data is reported. The example experiment (4.1) received data with a sampling interval of 3.3 ms, and faster reports are possible. Details of the box's timing capabilities are found in 5.1.

2.3.1. Frame detection using light sensors

To precisely track frame onsets and detect skipped frames (where an expected frame is not rendered) or stuck frames (where a single frame persists for two or more frames), we place the box's light sensors over two small patches at the top corners of the experimental display (see Fig. 4). One patch changes from black to white every frame (*timing patch*), while the other encodes a binary sequence (*coding patch*), specifically the 24-character sequence consisting of the 3-digit binary representations of the numbers 0-7. Deviations from the expected timing and sequence of blacks and whites can then be detected, and skipped or stuck frames identified (5.2). Light sensors were housed in custom 3D-printed clamps that fit a wide range of monitors (.STL files for 3D printing available on our website).

2.3.2. Time synchronization with external devices

There are three ways of syncing the time streams of external devices with the USE SyncBox, and thereby the rest of the experimental equipment. First, devices can receive up to 16-bit event codes sent on command over the USE SyncBox's rectangular port (e.g. the Neural Data Acquisition box in Fig. 4). Second, they could receive regular pulses sent over one of the USE SyncBox's single digital ports (e.g. the eyetracker in Fig. 4). Finally, they could themselves send data to the USE SyncBox (e.g. the joystick in Fig. 4). Custom adapters may need to be built for any of these purposes, and we provide schematics of adapters that alter the rectangular port's output for Neuroscan and BrainAmp acquisition systems on the USE website. Connecting an external device's output to the USE SyncBox may also require modifying the USE SyncBox's firmware.

If devices are not connected to the USE SyncBox in this manner, their time synchronization will be limited by Unity's update cycle. Thus, a standard consumer-grade joystick or keyboard connected over USB will have an unavoidable jitter of up to a frame (16.7 ms on a standard monitor), as well as any delays introduced by the ports themselves.

2.4. Analysis pipeline

2.4.1. Time-alignment of data files

The various data files produced during an experiment are produced by devices using different clocks, which all need to be aligned to a single timestream. In the current USE version, this happens offline. Generally speaking, experimental hardware receives timing pulses or event codes from, or sends digital signals to, the USE SyncBox. Alignment is then a simple matter of assigning USE SyncBox timestamps to each matching code or pulse, either received or generated, in the other hardware's data.

Alignment with data produced internally by the host computer, and with the physical state of monitor frames, is performed in a separate script. The frame-by-frame data stored during the experiment (3.2.7) includes the putative current status (black or white) of each of the two patches under the light sensors connected to the USE SyncBox, as reported by Unity. By analyzing the sensor voltages over the course of the experiment, the physical state of these patches on each frame is determined (see 5.2). Any skipped or stuck frames can then be identified by comparing the states of the clock and signal sensor, and the physical onsets of each frame, as opposed to Unity's estimated onset times, are determined. The frame data can then be modified to reflect the actual status of every single frame in the experiment. At this point, time alignment of all data is complete, as all datastreams are referenced to the USE SyncBox's timestamps, and can thus be directly compared with each other.

2.4.2. Experimental session reconstruction

During a typical experimental session, the frame-by-frame position, rotation, scale, and other properties of interest for each object in the scene is recorded in a *FrameData* file (4.2.6). This enables the complete reconstruction of the experimental session. During reconstruction, Unity's physics engine is ignored, and instead each object is directly assigned its properties as recorded. Each replayed frame contains a perfect re-presentation of the three-dimensional scene, including the camera's position, and thus the image on the monitor is identical to that originally seen by the participant. A video of the entire experimental session, or simply particular moments of interest, can be generated offline, with no need to record during the session itself. In the current version of the suite, this frame-by-frame reconstruction requires custom code modifications for each experiment, a process that we intend to streamline in future version updates.

Other suites have replay capabilities (e.g. Jangraw et al., 2014), which USE extends in several ways. First, gaze positions, mouse or touchscreen clicks, or data recorded by other equipment during the session can be overlaid on the screen. Also, since all skipped and stuck frames are known (see 4.3.1), replays are perfect representations of what was actually displayed. Individual frames can be exported at full or reduced size and analyzed at will, for example using saliency estimation algorithms (e.g. Borji et al., 2014). Finally, gaze targets can be determined, as described in the following section.

2.4.3. Gaze target determination

We smooth and classify gaze data using a modified version of an existing algorithm that provides superior processing of active gaze data (Andersson et al., 2016; Corrigan et al., 2017; Larsson et al., 2013, 2014; Nyström and Holmqvist, 2010), adapted to handle noisier data by using estimated angular acceleration (Engbert and Kliegl, 2003; Engbert and Mergenthaler, 2006) and robust statistics (Leys et al., 2013; Wilcox, 2012). Supplementary Methods 1 provides details of the algorithm. The results of this classification are detailed in 5.4 below.

After smoothing and classification, gaze targets can be determined (Fig. 5). The logic here is something like the inverse of gaze target determination in typical static tasks, in which an AOI is specified around an object of interest, and gaze points that land within this AOI are treated as landing on the object. In USE, any object whose two-dimensional silhouette on the screen lies within a specified distance from each gaze point is treated as a potential target of that point. For any frame of interest, the corresponding gaze points are identified, and a ShotgunRaycast method is used to determine which objects have been foveated (Fig. 5A). This defines a conical section filled with multiple raycasts in Unity's three-dimensional worldspace, whose smaller end is a circle centered on the gaze point, placed exactly on the camera (i.e. the surface of the screen). The larger end is defined such that its silhouette on the camera is identical in size and location to the smaller end. Thus, if any of the rays intersect with an object, its two-dimensional silhouette on the screen lies within the circle defined by the smaller end, i.e. is within the desired degrees of visual angle of the gaze point. The function returns a complete list of all objects hit by the rays making up the conical section. The density of sampling and the radius of the circle surrounding the gaze point are experimenter-defined.

Whether a gaze target is reliably detected depends on the shape of its associated mask, or *collider*, which defines the surfaces that raycasts can hit (Fig. 5B). A mesh collider that perfectly matches the shape of the object will also result in perfect detection, but for shapes defined by a high number of polygons, mesh colliders are computationally expensive. Simpler shapes (usually spheres) are computationally cheap but involve a tradeoff in detection accuracy. Various intermediate collider types can be used to define shapes with greater or lesser degrees of fidelity to the precise object shape. Which one experimenters use will depend on their particular needs. For most of the 3D objects used in our studies (Fig. 2C; Watson et al., in press), the high fidelity afforded by a mesh collider does not come at a high enough computational cost to warrant concern.

After replaying the session using shotgun raycasts for each gaze point, the experimenter then has the sample-by-sample specification of all objects falling within the desired degrees of visual angle from each gaze point, which can be further analyzed according to gaze type (e.g. fixations, smooth pursuits), or objects, as required by the experiment (see 5.4).

3. Results

3.1. USE SyncBox timing

The USE SyncBox enables digital signals to be generated with low latency and high precision (on the order of 0.1 ms for both). Specifically, digital timing pulses are generated with a clock stability of 100 ppm and jitter between pulses of approximately +/- 0.02 ms. When commanded to send a single digital pulse or event code, the delay between trigger and output is 0.01 to 0.11 ms, limited by the box's scheduling interval of 0.1 ms. Analog signals are digitized at 1 ms intervals with sample timing known to 0.1 ms accuracy. Light sensor signals pass through a pre-amplifier before digitization, which causes an additional delay of 0.33 ms.

Full details of the tests supporting these timing specifications are found in the Supplementary Material 2.

3.2. Frame detection

To quantify the performance of frame detection and frame onset determination for a monitor with 60 Hz frame rate, signals were recorded from the clock and signal light sensors over approximately 1.5 h. The data for each sensor was analyzed by identifying peaks and troughs in the first derivative of the voltage trace, where intervals peak-to-trough were classified as black frames, while trough-to-peaks were classified as white. The median duration of black or white periods in the clock signal was 16.7 ms. In the entire test period, there were exactly 32 periods (of a total of 336,722, ˜0.0001%) whose duration was not within 3.3 ms (one reporting interval) of 16.7, 33.3, 50.0, or 66.7 ms. Thus, effectively all white and black periods have a duration that corresponds to an integer multiple of the expected duration of a single

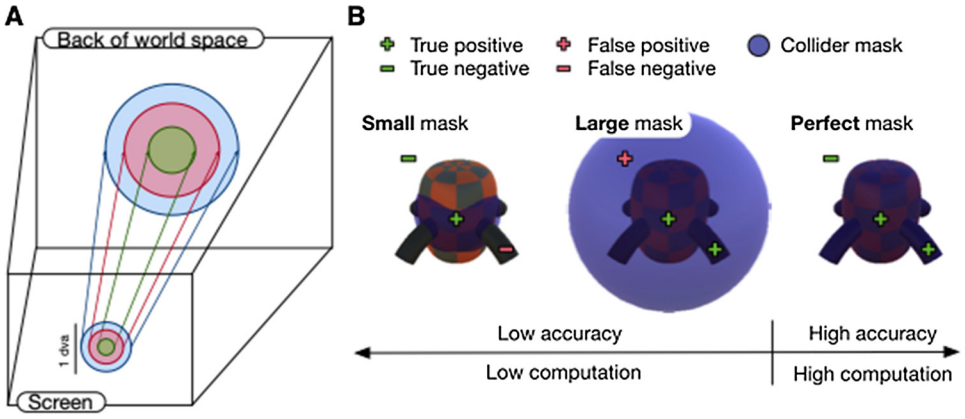

Fig. 5. ShotgunRaycast allows gaze target determination in a 3D world. **(A)** Moving through a 3D world means that the size and shape of the silhouette of an object on the screen constantly changes, making it a challenge to determine gaze targets. The ShotgunRaycast solves this challenge. It defines a conic section, beginning with a circle on the camera and ending with a larger circle a long distance into the world space, which has the same center on the screen and whose projection onto the screen subtends exactly the same angle. This has the effect of finding any objects whose silhouette on the screen lies either completely or partially within the circle defining the smaller end of the conic section. Experimentalists define the density of sampling, and radius of the circle in degrees of visual angle. **(B)** Detection of an object depends on the shape of the collider associated with it. Simple colliders (e.g. spheres; left, middle) are computationally inexpensive. However, they are inaccurate, either because raycasts miss the object - resulting in false negatives (left) - or raycasts hit the outside of an object - resulting in false positives (middle). Alternatively, mesh colliders that perfectly match the shape of an object are perfectly accurate, but are computationally expensive when the object is defined by a high number of polygons.

frame on a 60 Hz monitor. Any frame where this multiple is greater than 1 constitutes a skipped or a stuck frame, which can be identified by investigating the corresponding *coding patch's* data.

3.3. Frame onset to event code timing

A critical question for many neuroscientists will be *what is the relationship between the physical onset of a frame and timing signals sent on that frame?* We tested this by connecting an oscilloscope to (a) light sensor placed on the host computer's monitor, and (b) one line of the USE SyncBox's rectangular port. A USE script was written to alternate the patch of the monitor below the light sensor between black and white every frame, and to send an event code as rapidly as possible at the start of each frame.

Timing was referenced to event code output. The duration between frame onset and event code onset was stable but inconsistent. During any given stable interval, it was approximately constant with a jitter of about +/- 1 ms, but when there were disruptions to frame refreshes, and the patch did not flicker from black to white for two or more frames, the offset reset itself. Stable frame offset values were approximately 0–10 ms. This means that offsets between frame onset and event code onset are stable over periods of tens of seconds or longer, but are not known *a priori* or repeatable between stable periods. There is no Unity-controllable way in which this can be improved on, due to the imprecision of Unity's internal estimates of frame onset.

In summary, the timing of USE is typically precise, but contains unpredictable (non-systematic) glitches, which makes it necessary to run offline scripts to achieve sub-millisecond synchronization with frame onsets. We provide post-processing scripts that implement the accurate sub-millisecond time synchronization of frame onsets to event codes and other experimental hardware (see 5.2).

3.4. Gaze classification

Gaze data from participants in the example experiment were classified using the algorithm described in 4.4.3. Table 1 shows that the velocities, amplitudes, frequencies and durations of each classified gaze period are as in previous psychophysical studies (Andersson et al., 2016; Nyström and Holmqvist, 2010; Otero-Millan et al., 2008). Furthermore, saccades exhibit the classic *main sequence* (Fig. 6a), in which velocity and amplitude are linearly related, with an inflection point occurring between 10–15° (Inchingolo and Spanio, 1985).

After determining fixation targets using ShotgunRaycast (4.4.3), fixations to both target and distractor objects were used to generate a heatmap of normalized fixation locations on objects, showing that participants tended to focus on object centres, though there were also a smaller number of fixations to their peripheries (Fig. 6b), as one might expect in such a task (see Supplementary Materials 1 for an explanation of how normalized fixation locations were determined). Finally, participants began to preferentially fixate targets over distractors later in the block, showing that their gaze behaviour reflected their rule-learning (Fig. 2e), as has been shown in previous static studies of categorization learning (cf. Blair et al., 2009; Rehder and Hoffman, 2005; McColeman et al., 2014).

4. Discussion

4.1. Overview

The *Unified Suite for Experiments* is a complete integrative suite for the development, control, and analysis of active tasks, simplifying a number of crucial challenges facing researchers interested in more realistic stimuli and possibilities of control. Its State/Level architecture supports nested hierarchies of control that enable the generation of tasks of any degree of complexity, the USE SyncBox enables reliable

Table 1
Summary statistics (means and standard errors) for gaze periods classified as fixations, smooth pursuits, and saccades, averaged across all subjects performing the example experiment. The mean velocity of saccades and peak velocities of fixations and smooth pursuits are not shown, as these are irrelevant to their characterization and tend to be misleading.

	Fixation	Smooth Pursuit	Saccade
Mean Velocity (deg/s)	10.1 ± 0.5 SE	16.6 ± 0.7 SE	—
Peak Velocity (deg/s)	—	—	210.0 ± 3.6 SE
Mean Amplitude (deg)	0.390 ± 0.009 SE	2.5 ± 0.07 SE	5.7 ± 0.1 SE
Rate (/s)	2.6 ± 0.07 SE	1.4 ± 0.05 SE	2.4 ± 0.07 SE
Duration (s)	0.35 ± 0.08 SE	0.32 ± 0.01 SE	0.04 ± 0.0006 SE

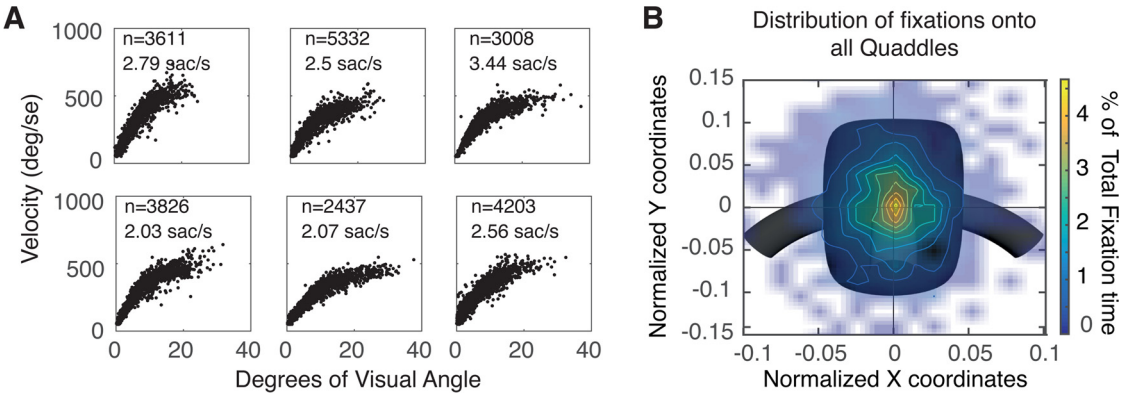

Fig. 6. Gaze behavior during example experiment. **(A)** Saccadic main sequence of six representative participants performing the context-dependent feature learning task. Saccade rate ranged from ˜2-3.5/s. An inflection point in the main sequence occurs around 10–15 deg. **(B)** Density of fixations onto a standardized object for one participant (see methods). Fixation centers were normalized for change in object size with depth and mapped onto a standard object located 4 world units away from the camera. Details of the normalization process are in Supplemental Materials 3. Color indicates the relative amount of time that participants spent fixating each point. Participants were most likely to fixate close to object centers.

timekeeping and communication among experimental devices, and the offline analysis scripts enable precise time-alignment of data files and reconstruction of gaze targets.

4.2. Unique features of USE

Other active experimental design suites have been published (Brookes et al., 2018; Doucet et al., 2016; Jangraw et al., 2014), as well as suites that leverage active tasks for artificial learning agents (Beattie et al., 2016; Leibo et al., 2018). These offer their own advantages, and we recommend that any researchers interested in active tasks review them carefully. There are, however, several unique features of USE that may make it particularly attractive for various purposes. These include:

- ***The State/Level architecture.*** In our experience, a specialized framework for the development of flexible experimental hierarchies greatly speeds up experimental development, and makes the structure of control more apparent.
- ***The USE SyncBox.*** A generic timekeeping and communication device is a powerful tool for neuroscience research, and enables the rapid extension of USE to any new experimental hardware without a corresponding loss of temporal precision or control.
- ***Physical frame onset detection.*** USE's ability to determine precise frame timing light sensors over the code and signal patch, and to thereby identify which specific frames were skipped or stuck during the experiment, enables the perfect reconstruction of the experimental session without relying on the game engine's estimates of frame onset times, which do not account for skipped and stuck frames and, in our experience can sometimes be as much as 20 ms off the actual onset.
- ***Shotgun raycasts.*** So far as we are aware, all published work using eyetracking in 3D scenes either uses simple areas of interest, which do not change with regard to object shape, or single raycasts, which do not account for eyetracker imprecision or the spatial extent of the fovea. Shotgun raycasts, on the other hand, enable precise control over the degrees of visual angle surrounding gaze points in which objects are designated as targeted, without assuming anything about the underlying object shapes.
- ***Translational capabilities.*** USE enables the rapid generation and testing of slightly different versions of the same experiment, suitable for research using different populations of humans and non-humans, as well as artificial learning agents of all types.
- ***Unity online resources.*** Unity is one of the most popular game engines in the world, and as such has many resources available for users of all skill levels, including tutorials, help forums, and (paid) support from Unity employees.

These advantages make USE a powerful tool for active experimentation.

Furthermore, the modular, generic and algorithmic principles underlying USE imply that it is possible for interested researchers to take any of these features that are currently unique to USE, and incorporate them into other development protocols. We hope to encourage cross-fertilization of this kind, where the best development tools and concepts from different laboratories influence those in others. For example, the gaze classification algorithm we describe in 3.4.3 was adapted from that used by Doucet et al.ös (2016) laboratory (Corrigan et al., 2017), for which we are grateful.

4.3. Possible extensions

USE is robust and full-featured enough to be of use to any researchers interested in active tasks. However, there are many ways in which its ease-of-use, flexibility, and power could be improved. We are currently considering several of these, including:

- Creating a generic event-driven data collection system for objects in the 3D environment that automatically records their location, rotation, scale, and other alterable properties whenever these properties change. This would be combined with a generic replayer system that could read in such data files and reconstruct each frame scene, without the necessity of coding a custom replayer for each study.
- Specific wrappers for communication with a large variety of experimental hardware and software (other eyetrackers, response devices, neural acquisition devices, etc).
- Real-time analysis of light sensor data in the USE SyncBox with enough precision to enable signal output to be time-locked to frame onsets, thereby enabling real-time feedback-loops such as those used in Brain-Machine Interface contexts.
- Sub-millisecond access to system-level monitor flip commands, as is available in static experimental design suites (Brainard, 1997; Kleiner et al., 2007; Peirce, 2007, 2008; Pelli, 1997), but not in Unity. This would enable more accurate online estimation of frame onset times, and thus greatly improve frame-locked temporal synchronization in cases when using light sensors is not feasible, as is the case for, e.g., most stereoscopic goggles.
- Sub-millisecond access to inputs from serial, UDP, and TCP/IP ports, which would enable much more precisely-timed communication to/from external devices, without requiring them to be connected to

the USE SyncBox. (This would still include the unavoidable delays associated with these ports, but their precision and accuracy would be much higher than in the current suite.)

USE is under active development, and we intend to update the files available for download as components are modified or created.

4.4. Concluding remarks

Interest in naturalistic, complex and active tasks is burgeoning, and will continue to do so. We very much look forward to seeing the effects that these tasks have on theory and method in the neurosciences, cognitive sciences, and in artificial intelligence research. We hope that the set of software and hardware components we have produced will play a small part in this process by facilitating temporally precise, well-controlled experiments in more complex and naturalistic settings.

Acknowledgements

This work was supported by Grant MOP 102482 from the Canadian Institutes of Health Research (to T.W.) and by the Natural Sciences and Engineering Research Council of Canada, Brain in Action CREATE-IRTG program (to M.R.W. and T.W.).

Appendix A. Supplementary data

Supplementary material related to this article can be found, in the online version, at doi:https://doi.org/10.1016/j.jneumeth.2019.108374.

References

Andersson, R., Larsson, L., Holmqvist, K., Stridh, M., Nyström, M., 2016. One algorithm to rule them all? An evaluation and discussion of ten eye movement event-detection algorithms. Behav. Res. Methods. https://doi.org/10.3758/s13428-016-0738-9.

Asaad, W.F., Eskandar, E.N., 2008a. A flexible software tool for temporally-precise behavioral control in Matlab. J. Neurosci. Methods 174, 245–258.

Asaad, W.F., Eskandar, E.N., 2008b. Achieving behavioral control with millisecond resolution in a high-level programming environment. J. Neurosci. Methods 173 (2), 235–240. https://doi.org/10.1016/j.jneumeth.2008.06.003.

Asaad, W.F., Santhanam, N., McClellan, S., Freedman, D.J., 2013. High-performance execution of psychophysical tasks with complex visual stimuli in MATLAB. J. Neurophysiol. 109 (1), 249–260. https://doi.org/10.1152/jn.00527.2012.

Barlow, S., 2015. Her Story. Sam Barlow., London.

Beattie, C., Leibo, J.Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., et al., 2016. DeepMind Lab. *arXiv*, *abs/1612.03801*. doi:arXiv:1612.03801v2. .

Bennett, A.J., Perkins, C.M., Tenpas, P.D., Reinebach, A.L., Pierre, P.J., 2016. Moving evidence into practice: cost analysis and assessment of macaques' sustained behavioral engagement with videogames and foraging devices. Am. J. Primatol. 78 (12), 1250–1264. https://doi.org/10.1002/ajp.22579.

Blair, M.R., Watson, M.R., Walshe, R.C., Maj, F., 2009. Extremely selective attention: eye-tracking studies of dynamic allocation of attention to stimulus features in categorization. J. Exp. Psychol. Learn. Mem. Cogn. 35, 1196–1206.

Borji, A., Sihite, D.N., Itti, L., 2014. What/Where to look next? Modeling top-down visual attention in complex interactive environments. IEEE Trans. Syst. Man Cybern. Syst. 44 (5), 523–538. https://doi.org/10.1109/TSMC.2013.2279715.

Bouchard, S., Guitard, T., Bernier, F., Robillard, G., 2011. Virtual reality and the training of military personnel to cope with acute stressors. In: In: Brahnam, S., Jain, L.C. (Eds.), Virtual Reality in Psychotherapy, Rehabilitation, and Assessment, vol. 6. Springer-Verlag, Inc, New York, NY, pp. 109–128.

Brainard, D.H., 1997. The Psychophysics Toolbox. Spat. Vis. 10 (4), 433–436.

Brookes, J., Warburton, M., Alghadier, M., Mon-Williams, M.A., Mushtaq, F., 2018. Studying human behaviour with virtual reality: the Unity Experiment Framework. biorxiv. https://doi.org/10.1101/459339.

Chaytor, N., Schmitter-Edgecombe, M., 2003. The ecological validity of neuropsychological tests: a review of the literature on everyday cognitive skills. Neuropsychol. Rev. 13 (4), 181–197.

Order, Colossal, 2015. Cities: Skylines. Paradox Interactive., Tampere, Finland.

Corrigan, B.W., Gulli, R.A., Doucet, G., Martinez-Trujillo, J.C., 2017. Characterizing eye movement behaviors and kinematics of non-human primates during virtual navigation tasks. J. Vis. 17 (12), 15. https://doi.org/10.1167/17.12.15.

Doucet, G., Gulli, R.A., Martinez-Trujillo, J.C., 2016. Cross-species 3D virtual reality toolbox for visual and cognitive experiments. J. Neurosci. Methods 266, 84–93. https://doi.org/10.1016/j.jneumeth.2016.03.009.

Engbert, R., Kliegl, R., 2003. Microsaccades uncover the orientation of covert attention. Vision Res. 43 (9), 1035–1045. https://doi.org/10.1016/S0042-6989(03)00084-1.

Engbert, R., Mergenthaler, K., 2006. Microsaccades are triggered by low retinal image slip. Proc. Natl. Acad. Sci. U.S.A. 103 (18), 7192–7197.

Hwang, J., Mitz, A.R., Murray, E.A., 2019. NIMH MonkeyLogic: behavioral control and data acquisition in MATLAB. J. Neurosci. Methods 323, 13–21. https://doi.org/10.1016/j.jneumeth.2019.05.002.

Inchingolo, P., Spanio, M., 1985. On the identification and analysis of saccadic eye movements—a quantitative study of the processing procedures. IEEE Trans. Biomed. Eng. 32 (9), 683–695. https://doi.org/10.1109/TBME.1985.325586.

Jangraw, D.C., Johri, A., Gribetz, M., Sajda, P., 2014. NEDE: an open-source scripting suite for developing experiments in 3D virtual environments. J. Neurosci. Methods 235, 245–251. https://doi.org/10.1016/j.jneumeth.2014.06.033.

Kingstone, A., Smilek, D., Eastwood, J.D., 2008. Cognitive Ethology: a new approach for studying human cognition. Br. J. Psychol. 99 (3), 317–340. https://doi.org/10.1348/000712607X251243.

Kleiner, M., Brainard, D., Pelli, D.G., 2007. What's new in Psychtoolbox-3? *Perception*, *36*, ECVP Abstract Supplement.

Kruschke, J.K., 1992. ALCOVE: an exemplar-based connectionist model of category learning. Psychol. Rev. 99 (1), 22–44.

Larsson, L., Nyström, M., Stridh, M., 2013. Detection of saccades and postsaccadic oscillations in the presence of smooth pursuit. IEEE Trans. Biomed. Eng. 60 (9), 2484–2493. https://doi.org/10.1109/TBME.2013.2258918.

Larsson, L., Nyström, M., Stridh, M., 2014. Discrimination of fixations and smooth pursuit movements in high-speed eye-tracking data. Conf. Proc. IEEE Eng. Med. Biol. Soc. 3797–3800. https://doi.org/10.1109/EMBC.2014.6944450.

Leibo, J.Z., d'Autume, C.D.M., Zoran, D., Amos, D., Beattie, C., Anderson, K., et al., 2018. Psychlab: a psychology laboratory for deep reinforcement learning agents. arXiv 1801.08116v2.

Leys, C., Ley, C., Klein, O., Bernard, P., Licata, L., 2013. Detecting outliers: do not use standard deviation around the mean, use absolute deviation around the median. J. Exp. Soc. Psychol. 49 (4), 764–766. https://doi.org/10.1016/j.jesp.2013.03.013.

Liao, H., Qu, Z., 2013. Virtual experiment system for electrician training based on Kinect and Unity3D. Proceedings from International Conference on Mechatronic Sciences, Electric Engineering and Computer.

McColeman, C.M., Barnes, J.I., Chen, L., Meier, K.M., Walshe, R.C., Blair, M.R., 2014. Learning-induced changes in attentional allocation during categorization: a sizable catalog of attention change as measured by eye movements. PLoS One 9 (1), e83302. https://doi.org/10.1371/journal.pone.0083302.

Niantic, 2016. Pokémon GO. Niantic., San Francisco.

Nintendo EDP, 2016. Super Mario Run. Nintendo, Kyoto.

Nyström, M., Holmqvist, K., 2010. An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data. Behav. Res. Methods 42 (1), 188–204. https://doi.org/10.3758/BRM.42.1.188.

Mayer, R.E., 2018. Computer games in education. Annu. Rev. Psychol.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., et al., 2015. Human-level control through deep reinforcement learning. Nature 518 (7540), 529–533. https://doi.org/10.1038/nature14236.

Oemisch, M., Watson, M.R., Womelsdorf, T., Schubö, A., 2017. Changes of attention during value-based reversal learning are tracked by N2pc and feedback-related negativity. Front. Hum. Neurosci. 11, 7. https://doi.org/10.3389/fnhum.2017.00540.

OpenAI, 2018. Gym. Retrieved November 20, 2018 From gym.openai.com.

Otero-Millan, J., Troncoso, X.G., Macknik, S.L., Serrano-Pedraza, I., Martinez-Conde, S., 2008. Saccades and microsaccades during visual fixation, exploration, and search: foundations for a common saccadic generator. J. Vis. 8 (14). https://doi.org/10.1167/8.14.21. 21.1-2118.

Peirce, J.W., 2007. PsychoPy–Psychophysics software in Python. J. Neurosci. Methods 162 (1-2), 8–13. https://doi.org/10.1016/j.jneumeth.2006.11.017.

Peirce, J.W., 2008. Generating stimuli for neuroscience using PsychoPy. Front. Neuroinform. 2, 10. https://doi.org/10.3389/neuro.11.010.2008.

Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat. Vis. 10 (4), 437–442.

Rehder, B., Hoffman, A.B., 2005. Eyetracking and selective attention in category learning. Cogn. Psychol. 51 (1), 1–41.

Risko, E.F., Laidlaw, K., Freeth, M., Foulsham, T., Kingstone, A., 2012. Social attention with real versus reel stimuli: toward an empirical approach to concerns about ecological validity. Front. Hum. Neurosci. 6, 143. https://doi.org/10.3389/fnhum.2012.00143.

Schmuckler, M.A., 2001. What is ecological validity? A dimensional analysis. Infancy 2 (4), 419–436. https://doi.org/10.1207/S15327078IN0204_02.

Shanley, P., 2017. Why aren't video games as respected as movies? Retrieved September 6, 2018 from. . https://www.hollywoodreporter.com/heat-vision/why-arent-video-games-as-respected-as-movies-1067314.

Slater, M., Wilbur, S., 1997. A framework for immersive virtual environments (FIVE): speculations on the role of presence in virtual environments. Presence Teleoperators Virtual Environ. 6 (6), 603–616.

Squad, 2015. Kerbal Space Program. Squad., Mexico City.

The NDP Group, 2009. More americans play video games than go out to the movies. Retrieved September 10, 2018 from. https://www.npd.com/wps/portal/npd/us/news/press-releases/pr_090520/.

Unity Technologies, 2017a. Company Facts. Retrieved July 17, 2017, from. https://unity3d.com/public-relations.

Unity Technologies, 2017b. Learn With Unity. Retrieved July 17, 2017, from. https://unity3d.com/learn.

Wagner, F., Schmuki, R., Wagner, T., Wolstenholme, P., 2006. Modeling Software With Finite State Machines — a Practical Approach. Auerbach Publications, Boca Raton.

Wang, J.X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J.Z., et al., 2018. Prefrontal cortex as a meta-reinforcement learning system. Nat. Neurosci. 21

(6), 860–868. https://doi.org/10.1038/s41593-018-0147-8.
Watson, M.R., Voloh, B., Naghizadeh, M., Womelsdorf, T. (In press). Quaddles: a multi-dimensional 3-D object set with parametrically controlled and customizable features. Behav. Res. Methods.
Watson, M.R., Voloh, B., Naghizadeh, M., Womelsdorf, T., 2018. Hierarchical Processing Revealed by Choice and Information-Sampling in a Dynamic Feature-Based Context-Dependent Learning Task. Manuscript in preparation. .
Wilcox, R.R., 2012. Introduction to Robust Estimation and Hypothesis Testing, 2nd ed. Elsevier/Academic Press, Amsterdam.
Witmer, B.G., Singer, M.J., 1998. Measuring presence in virtual environments: a presence questionnaire. Presence Teleoperators Virtual Environ. 7 (3), 225–240.
Youngblut, C., 2007. What a Decade of Experiments Reveals About Factors That Influence the Sense of Presence: Latest Findings. Institute for Defense Analysis, Alexandria, Virginia.

Glossary

Active task: Experimental tasks which involve some combination of realistic, usually moving, stimuli, continuous opportunities for action, ecologically valid tasks, complex behaviours, etc. Here, they are contrasted with **static tasks (see below)**

Arduino: A multi-purpose generic micro-processor, here used to control inter-device communication and time synchronization.

Raycast: A game-engine method that sends a vector between two points in a virtual three-dimensional environment, and returns the first object in that environment it hits. Often used to determine if a character in a game can see or shoot another character.

State Machine (also Finite State Machine): A way of conceptualizing and implementing control in software, such that at any one moment the software is in one, and only one, state. In **hierarchical state machines**, as used in the present software suite, these can be organized into different levels, such that each level can only be in one state, but a state can pass control to a lower level.

Static task: Experimental tasks like those traditionally used in the cognitive neurosciences. Simple, usually stationary, stimuli, limited opportunities for action, simple behaviours, etc. Here, they are contrasted with **active tasks (see above)**.

Unity: One of the most popular video game engines. Freely available.

Video game engine: A software development kit designed to handle many of the common issues involved in creating video games, such as interfacing with controllers, simulating physical collisions and lighting, etc.

Appendix E. MAD saccade: statistically robust saccade threshold estimation

Journal of Eye Movement Research
x(y):z

MAD saccade: statistically robust saccade threshold estimation

Benjamin Voloh
Vanderbilt University
Nashville, USA

Marcus R. Watson
York University
Toronto, Canada

Seth König
Vanderbilt University
Nashville, USA

Thilo Womelsdorf
Vanderbilt University
Nashville, USA

Saccade detection is a critical step in the analysis of gaze data. A common method for saccade detection is to use a simple threshold for velocity or acceleration values, which is typically estimated from the data using the mean and standard deviation. However, this method has the downside of being influenced by the very signal it is trying to detect, the outlying velocities or accelerations that occur during saccades. We propose instead to use the median absolute deviation (MAD), a robust estimator of the standard deviation that is not influenced by outliers. We modify an algorithm proposed by Nyström and colleagues, and quantify saccade detection performance in both simulated and human data. Our modified algorithm shows a significant and marked improvement in saccade detection, showing both more true positives and fewer false negatives. We conclude that robust estimators can be widely adopted in other common, automatic gaze classification algorithms due to their ease of implementation.

Keywords: Saccades, median absolute deviation, MAD, eye tracking, free viewing

Introduction

The analysis of gaze behaviour yields critical insights into processes underlying visual attention, perception, and executive control, as well as their mechanistic underpinnings (Bahill, Clark, & Stark, 1975; Engbert & Kliegl, 2003; Otero-Millan, Troncoso, Macknik, Serrano-Pedraza, & Martinez-Conde, 2008). One prominent component of gaze behaviour are saccades, ballistic eye movements that rapidly re-orient the eye and thus the image impinging on the retina. Because of their ballistic nature, saccades can be differentiated from other gaze events – such as fixations or smooth pursuits – by examining the eye velocity and/or acceleration (Chau, 2011; A. Duchowski et al., 2002; König & Buffalo, 2014; Larsson, Nystrom, & Stridh, 2013; Nyström & Holmqvist, 2010). A common approach is to apply a threshold in the velocity (e.g. (Nyström & Holmqvist, 2010)) and/or acceleration profile (e.g. (A. Duchowski et al., 2002)). Typically, threshold crossings mark the presence of saccades; further analysis can then demarcate more precise onsets and offsets (Larsson et al., 2013; Nyström & Holmqvist, 2010). Although other, computationally more elaborate methods exist that can outperform thresholding algorithms (Komogortsev, Gobert, Jayarathna, Do Hyong Koh, & Gowda, 2010; König & Buffalo, 2014; Salvucci & Goldberg, 2000), thresholds remain popular because they are simple to implement, particularly in experiments with rigorously defined, simple gaze behaviour.

A critical and common step in algorithmic saccade detection is thus the choice of the threshold. However, variability in saccadic profiles, the presence of other gaze events such as fixations or smooth pursuits (Larsson et al., 2013), measurement noise (Dai, Selesnick, Rizzo, Rucker, & Hudson, 2016; Holmqvist, Nyström, & Mulvey, 2012), or sampling frequency (Mack, Belfanti, & Schwarz, 2017), all make it difficult to reliably detect saccades algorithmically. Thus, saccadic detection can be improved if the threshold is estimated robustly from the data itself, and may even adapt to changing conditions (A. T. Duchowski, 2003; Engbert & Kliegl, 2003; Nyström & Holmqvist, 2010).

Recently, Nyström and colleagues developed an algorithm that iteratively calculates an adaptive threshold, which has the benefit of setting the lowest possible threshold given changes in background noise and fixation characteristics (Nyström & Holmqvist, 2010). This algorithm performs better than nine other recent algorithms on data measured while participants viewed static stimuli (Andersson, Larsson, Holmqvist, Stridh, & Nyström, 2017). It iteratively calculates the threshold as a function of the mean and standard deviation. However, because these quantities are highly biased by the presence of outliers, the outliers may be undetected, a phenomenon called "masking" (Wilcox, 2012). In the velocity or acceleration domain, saccades are by definition outliers, as they take up a far smaller number of gaze points than fixational intervals, and they have much higher velocities and accelerations. This implies that the saccade detection threshold is modulated by the very signal it is trying to detect (Fig 1). A robust estimation of the saccade threshold would ideally be independent of saccades.

The purpose of this work is to formally compare saccade detection using statistically robust and non-robust threshold estimation. We have further set out two goals for threshold estimation; first, it should be robust and insensitive to variations of maximum saccade amplitudes; and second, that it can be efficiently implemented. To this end, we modify the algorithm by Nyström and colleagues by using the median and median absolute deviation (MAD), which are robust to outliers in the data (Leys, Ley, Klein, Bernard, & Licata, 2013; Wilcox, 2012). The median absolute deviation is a robust estimate of the standard deviation, and can take into account non-normal distributions via a scale factor (Leys et al., 2013; Wilcox, 2012). Previous gaze research has used this measure to detect outliers and clean data (Rütsche, Baumann, Jiang, & Mojon, 2006), while other research has recommended using a median filtering to reduce the influence of noise in general (Daye & Optican, 2014; Liston, Krukowski, & Stone, 2013). We propose using MAD as a threshold estimator in and of itself. We find that robust threshold estimation leads to improved saccade detection, particularly at higher noise levels, and is robust to varying threshold confidence levels. It also outperforms the original algorithm as proposed by Nyström and colleagues. We conclude that threshold estimation based on MAD can be widely and easily applied in other saccade detection algorithms where the threshold should be robustly estimated (e.g. (Larsson et al., 2013)).

Methods

All analyses were performed in MATLAB 2015b (Mathworks) using custom code. Implementation code for the robust estimation and the saccade simulation code is available on reasonable request to the corresponding author.

Saccade Simulation

We simulated saccades to create ground truth scan paths in order to objectively compare algorithms. Simulations were based on the procedure proposed by (Dai et al., 2016). Two dimensional saccades were generated using a parametrical model of saccades that reproduces the saccadic main sequence. The parametric model generates a saccade waveform from the sum of a soft ramp function and a shifted negated soft ramp function (Equation 2. (Dai et al., 2016)). We used the following parameter values for Equation 2 selected from uniform distributions ranging between these listed values : η= 0.45-0.65, c = 4.5-7.5, and τ = 2-6. Importantly, τ represents saccade amplitude. We used saccade amplitudes ranging from 2-6 degrees of visual angle (dva), representing ranges that are prevalent in experiments with static images in humans.

Adaptive Algorithm for Saccade Detection

Nyström and colleagues (2010) proposed an innovative method to determine the saccade threshold in an adaptive, iterative way (Fig 1A) (Nyström & Holmqvist, 2010) . It first determines a global threshold for saccade detection, and then a local threshold for onset/offset estimation. We obtained code for the adaptive algorithm from the personal website of Marcus Nyström (http://www.humlab.lu.se/en/person/MarcusNystrom/ , link: "*Source code for the algorithm described in Nyström, M., & Holmqvist, K. (2010). An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data. Behavior research methods, 42(1), 188-204.*"). This code was slightly modified from its published version. We refer interested readers to the original publication for the full details of the algorithm, but describe here the relevant details for the adaptive threshold calculation.

The determination of the saccade peak velocity threshold θ_{PT} can be broken down into the following steps (Figure 1A). First, θ_{PT} is set to an initial value in the range 100-300 deg/s. Second, for all velocity samples x lower than θ_{PT}, a new threshold is calculated as:

$$\theta_{PT} = \mu + \lambda * \sigma \quad (1)$$

where μ and σ is the mean and standard deviation over samples *x,* the parameter λ (lambda) is a scale factor equal to 6. This procedure is then repeated until the error between iterations is less than 1 deg/s. The saccade onset/offset threshold θ_{ST} has a similar form:

$$\theta_{ST} = \mu + 3 * \sigma \quad (2)$$

In our tests, we compare two versions of this original algorithm. First, we use the algorithm as presented on the website. However, in the current implementation (1.0), the algorithm does not (re-)calculate the threshold over all remaining data samples at each iteration. Instead, for each putative inter-saccadic interval (i.e. between threshold crossings), a number of samples are removed at the start and end of the inter-saccadic interval, defined as the minimum fixation duration (40 ms) * sampling frequency (500 Hz) / 6, which comes out to 3 samples removed at the flanks of each inter-saccadic intervals. In our simulations, this amounts to the removal of ~1% of the data. Thus, a second version of the algorithm does *not* excise any parts of the data, which is the algorithm as originally proposed in the publication. We refer to the first version – the *adaptive threshold* algorithm with excised data - as "AT-excise", and the second one using all data as "AT".

Robust estimation of mean and deviation

To get a robust estimate of the central tendency and variability of the data, we instead propose to use the median and median absolute deviation (MAD) (Leys et al., 2013; Wilcox, 2012). In this framework, we treat saccades as outliers to be detected (Engbert & Kliegl, 2003). A robust measure of variability σ' is defined as:

$$\sigma' = b * MAD(x) \quad (3)$$

$$MAD(x) = median(\,|x_i - median(x)|) \quad (4)$$

The factor b is equal to 1.4826 assuming the underlying distribution (i.e. ignoring outliers) is normal (Leys et al., 2013; Wilcox, 2012). Thus, we proposed to calculate the peak velocity and saccade duration thresholds as:

$$\theta_{PT} = \mu' + \lambda' * \sigma' \quad (5)$$

$$\theta_{ST} = \mu' + 3 * \sigma' \quad (6)$$

where μ' is the median. We refer to this algorithm as "AT-MAD", and also compare it to one where we excise data (as described above), referred to as "AT-MAD-excise".

Algorithm Comparison

To compare the performance of different algorithms, we used event-based comparison to match true and reported saccades (Dai et al., 2016; Warby et al., 2014) . A match is logged if there is sufficient sample overlap (>20%) between a true and reported saccade. All matched saccades are true positives (TP), unmatched true saccades are false negatives (FN), and unmatched reported saccades are false positives (FP). From these, we further calculate the precision (=TP/TP+FP) and recall (=TP/(TP+FN)). These are used to compute the F1 score (=2*precision*recall/(precision+recall)), an aggregate performance measure.

To determine the timing characteristics of onsets and offsets, we take the difference in onsets/offsets between the (matched) true and reported saccades. Onset/offsets in the true reference were defined as the first point where the velocity dropped below 5 deg/s in the noise-free simulation. The onset/offset lags were defined as the average of the difference within a simulation. The onset/offset jitters were defined as the standard deviation of the differences within a simulation.

To determine the difference in performance, we perform a pair-wise t-test by taking the difference in F1 scores between algorithms for each simulation. P-values were multiple-comparison corrected using the Bonferroni procedure.

Experiment and Data

The York University Office of Research Ethics approved the present study as confirming to the standards of the Canadian Tri-Council Research Ethics guidelines (Certificate # 2016-214) We analyzed gaze data from performance during a feature-based attention task (n=3). Participants had to choose between one of two objects presented on a monitor screen (Watson, Voloh, Naghizadeh, & Womelsdorf, 2017). Over a block of trials, they learned to choose the rewarded target. When they reached performance criterion, the block would switch and feature-reward associations had to be re-learned.

At the beginning of each trial, participants were presented with a blank white screen and a fixation point that they needed to fixate for one second before they could begin a trial. Upon fixation, this central point would turn from blinking red to solid black, before disappearing. The blank white screen persisted for another 600ms before the subjects were presented with a 3D rendered arena with two

objects placed at random positions. If participants broke this fixation too early, then the trial would be aborted.

Subjects had to fixate one of the two objects for 300 ms before receiving a cue that let them know that they were able to choose that object (the cue consisted of a translucent dot superimposed on stimulus). Participants were then able to choose that object by pressing down the spacebar and continuing to fixate the object for another 100 ms. Auditory feedback in the form of a low or high pitched beep, and visual feedback in the form of a red dot superimposed on the stimulus or a yellow dot superimposed on the stimulus for 300 ms, were used to indicate incorrect and correct responses respectively. After feedback, the objects disappeared, and the inter-trial interval consisted of the arena for 800ms. Participants had 30 seconds to make a response in a trial, otherwise the trial would be aborted, and they would be presented with instructions letting them know to respond faster.

The experiment was run using custom code for the Unity3D game engine. Gaze data was collected using the Tobii TX300 (sampling frequency, 300 Hz), in a similar, low light environment. Participants were seated 50-60cm away from the monitor. Although they were seated, they were otherwise unrestrained.
We randomly selected 60 seconds of data from each participant at least 10 minutes into the session. Noise levels for each participant were defined as RMS of the x- and y-gaze positions during the inter-saccadic (i.e. fixational) periods (Holmqvist et al., 2012). Manual classification was done with a custom GUI that had four displays; x-position, y-position, velocity, and (x,y) gaze.

Results

To illustrate the strength of robust statistics, we show an example simulation with moderate noise where AT failed to detect the saccade (Fig 1C) but AT-MAD succeeded in doing so (Fig 1D; see also Fig 5E-F). In this example, the relatively high velocity values (Fig. 1B) push the threshold higher than its initial starting value (Fig. 1C). After many iterations, the threshold is too high to detect this saccade. However, because robust threshold estimation is relatively insensitive to outlier values, the AT-MAD algorithm successfully exits after one iteration, and is thus able to detect the saccade (Fig. 1D).

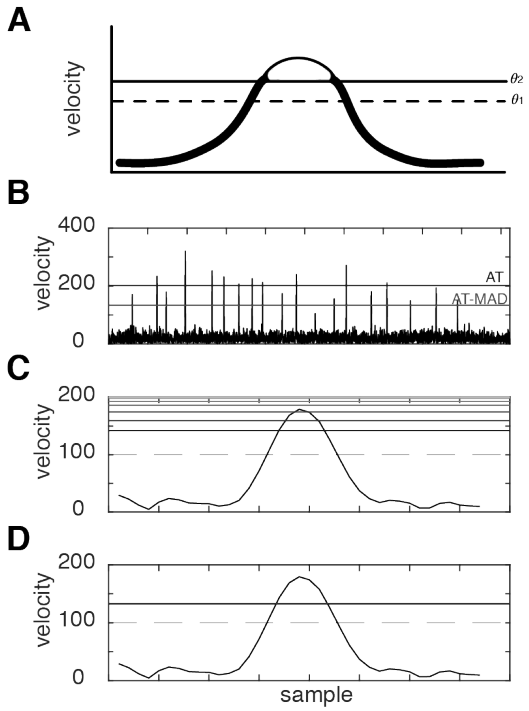

Fig. 1. Example where adaptive thresholding detects saccades in a data-driven manner, but can fail with non-robust statistics
(A) Schematic of AT algorithm. The threshold for detection is initialized at θ_{PT1}. All points below this (thick line) are then used to calculate θ_{PT2}. This algorithm proceeds until it converges to a solution. (B) Velocity of simulated scanpath with 20 low amplitude saccades. Horizontal lines depict the threshold as determined by the AT algorithm (black) and AT-MAD algorithm (grey). The AT-MAD algorithm finds a lower bound, than AT though still well above the background noise. (C-D) Example saccade that was not detected by the AT algorithm (C) but was detected by the AT-MAD algorithm (D) (see Fig 5E,F for examples from human data). The initial threshold is depicted as a dotted line. Sold horizontal lines represent the threshold on successive iterations, with darker (lighter) lines showing earlier (later) iterations. Notice that it increases beyond the initial threshold, but the AT-MAD algorithm successfully stops iterating, whereas AT does not.

We compared the detection performance (F1 score) of four different versions of the algorithm as a function of noise level (Figure 2A,B). We found that using the AT algorithm as originally proposed had the lowest performance. Performance improved if some data flanking the saccades was removed (AT-excise), as in the version published on Nyström's website, suggesting that the original algorithm remains sensitive to the relatively high velocities just below threshold, on the flanks of saccades. However, AT-MAD and AT-MAD-excise had similar performance over all noise levels, showcasing the insensitivity of detected thresholds to outliers when using robust statistics. Robust threshold estimation consistently and significantly improved F1 scores by ~0.02-0.1 for noise levels

equal to or greater than 0.4 (Fig 2B; $p<0.05$, multiple comparison corrected). If some data in the inter-saccadic intervals was excised (AT-excise vs AT-MAD-excise), the improvement peaked at 0.056, whereas the improvement was greatest (0.096) if all data was considered (AT vs AT-MAD). Improvements could be traced to both a higher true positive rate (Fig 2C), and a lower false negative rate (Fig 2D).

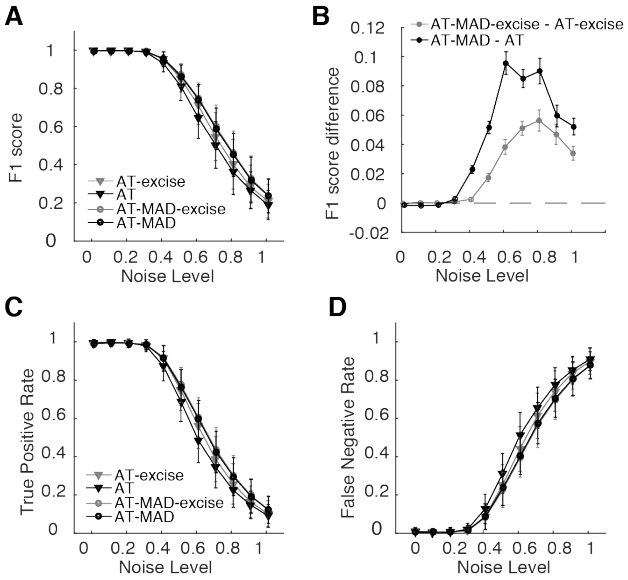

Fig 2. Robust estimation of threshold parameters results in improved detection performance
(A) Mean and standard deviation of F1 score as a function of noise for two versions of the original algorithm (triangles, AT and AT-excise) and two versions using the robust estimator (circles, AT-MAD and AT-MAD-excise). Versions using MAD consistently outperform (compare circles vs triangles). Using the original data, excising some data improves detection performance (compare grey and blacks lines, triangles), whereas it makes no difference for versions using the MAD estimator (grey and black lines, circles). (B) Mean and standard deviation of the pairwise difference in F1 score comparing AT vs AT-MAD (black), and AT-excise and AT-MAD-excise (grey). Filled circles represent statistically different score ($p<0.05$, multiple comparison corrected). Threshold estimation using MAD shows significantly improved performance for noise levels at 0.4 and above, particularly for the algorithm as originally proposed (black). (C-D) Mean and standard deviation of the true positive rate (C), and false negative rate (D) for the four different versions of the algorithm. Improvements in (A-B) can be traced both to a higher true positive rate (C) and a lower false negative rate (D).

While robust threshold estimation results in better saccade detection, it may do so at the expense of ill-characterized saccade durations. Thus, we compared onset and offset lags between the original and MAD versions of the algorithm. We found that onset and offset lags were comparable for all versions of the algorithms, increasing with noise (Fig 3A-B). All algorithms had a jitter of ~2ms across all noise levels, although the variability in jitter increased with noise (Fig 3C-D). Thus, based on the simulation results, the MAD algorithm consistently and more reliably detects saccades at higher noise levels, but shows similar saccade durations.

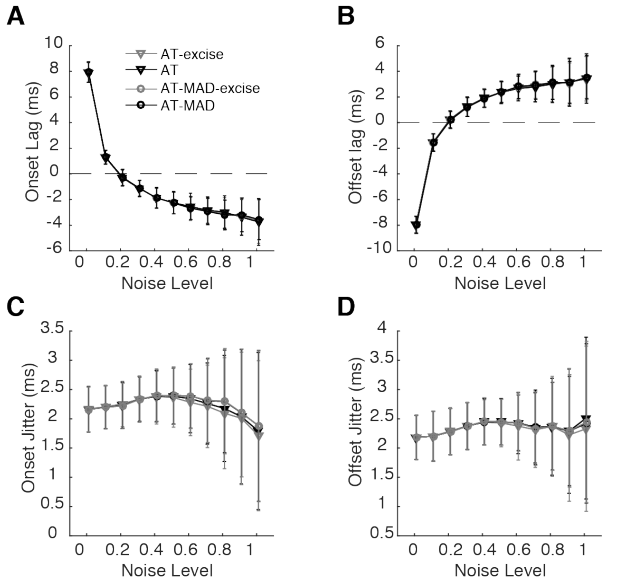

Fig 3. Onset and offset lags are comparable across all algorithms
(A-B) Mean and standard deviation of the onset (A) and offset (B) lags, which is the difference in onset/offset as determined by the algorithm and the reference. All tested algorithms show similar lags. At higher noise levels, the lag decreases (shifted earlier). (C-D) Mean and standard deviation of the onset (C) and offset (D) jitter in lag. Across all tested noise levels, jitter is around 2 ms, but variability in jitter gets higher with increasing noise levels. Jitter at all noise levels is comparable across algorithms.

Experimental questions may necessitate varying levels of confidence in saccade detection. For example, studies comparing saccade characteristics across populations (Bargary et al., 2017) may only wish to analyze well-defined saccades. In this case, experimentalists may wish to have a particularly high detection threshold. This could be manipulated by setting the lambda parameter to higher values. To this end, we performed another set of analyses manipulating the value of lambda (Fig 4). We found that for low noise levels (<0.2), the value of lambda did not distinguish the performance of any four considered algorithms (Fig 4A-B,D-E). Very low lambda (=4, 5) had lower performance, due to a higher number of false positives. Here, AT-excise and AT slightly but significantly outperformed their robust counterparts (Fig 4C, F). However, the benefits of the MAD algorithm start to accrue at noise levels higher than 0.4. At these noise levels, very high lambda value (=10) negatively impact the performance of AT by a

factor of ~0.2 (Fig 4A). At higher noise levels, performance of the AT rapidly declines, to a minimum of ~0.2 at the highest noise level (=1). However, at this same level, AT-MAD achieved a performance of ~0.4. (Fig 4B). Indeed, for noise levels greater than or equal to 0.4, the AT-MAD algorithm had consistently and significantly higher performance than AT, up to a peak of ~0.45 at the highest value of lambda (Fig 4C). The effects were qualitatively similar when comparing AT-excise and AT-MAD-excise (Fig 4D-F), where the performance boost was smaller but still substantial (Fig 4F). Thus, the MAD algorithm allows experimentalists to more robustly define their desired confidence for saccade detection.

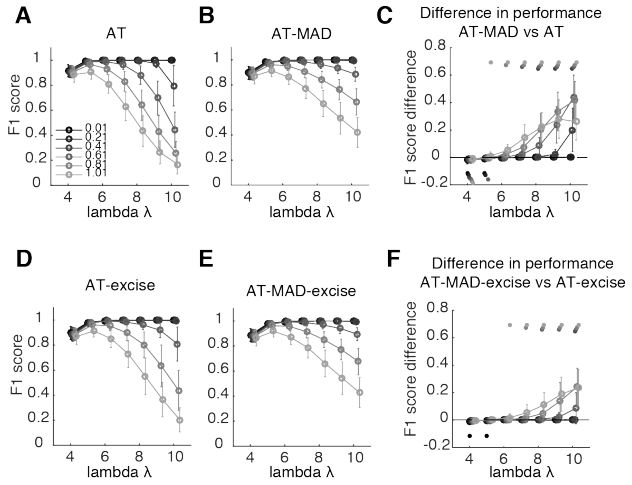

Fig 4. MAD is more robust to changes in threshold confidence level
(A) Mean and standard deviation of the detection performance of the AT algorithm as a function of lambda for various levels of noise. Grayscale indicates the level of noise, with darker (lighter) colors indicating less (more) noise. At low noise levels (<0.2), performance does not depend on the choice of lambda. At moderate noise levels (0.4), performance rapidly decreases at very high lambda values (=10). However, with high noise levels, there is a substantial and rapid decrease in performance with higher lambda values. (B) Same as (A) but for the AT-MAD algorithm. Performance is high for low-moderate noise levels (<0.6). For higher noise levels, performance decreases at a slower rate for higher lambda values. (C) Difference in detection performance between AT-MAD and AT. Grayscale dots represent significant differences at the corresponding noise level. Dots above (below) zero depict significant increases (decreases) ($p<0.05$, multiple comparison corrected). At higher levels of lambda, AT-MAD far outperforms AT. This is especially true for moderate to high-levels of noise. (D-F) Same as (A-C) but for AT-excise (D), AT-MAD-excise (E), and their comparison (F). Results are qualitatively similar as for (A-C). Excising some data points using the original algorithm helps but using MAD still allows for higher lambda values.

Up to this point, results were based on simulated data. A remaining question is how the algorithms would perform on real-world data. We analyzed data from three subjects performing a task with head-unrestrained, free viewing of a static scene. We manually annotated one minute of data randomly sampled from each subject. Noise levels (RMS) for the three subjects were in the range of [0.06 – 0.1] in the x-position, and [0.11 – 0.23] in the y-position, placing it in the lower quarter of simulated noise levels (Table 1). We compared the AT-excise and AT-MAD-excise (which, we note, had the lowest amount of improvement in our simulations, Fig 2). AT-MAD-excise outperformed AT-excise in all cases where lambda was set to 7 or greater. This is due to its higher true positive rate (Fig 5B), generally no false positives (Fig 5C), and lower false negative rate (Fig 5D). In one subject, AT-excise outperformed AT-MAD-excise at a lambda of 6 (Fig 5Ai), due to a substantial number of false positives in AT-MAD-excise (Fig 5Ci). Detection performance of AT-MAD-excise generally stayed stable across all three subjects with increasing values of lambda (Fig 5A), whereas AT-excise showed a steep drop in performance in two subjects with increasing lambda (Fig 5Ai-ii). Generally speaking, AT-excise failed because it would consistently find a higher threshold than AT-MAD-excise, resulting in an inability to detect low-velocity saccades (illustrated in Fig 5E-F). These results complement the simulation results and suggest that AT-MAD-excise outperforms AT-excise on human data by estimating a lower threshold, though one that is still conservative enough to avoid false positives in most cases. Furthermore, they suggest a lambda of 8 can drastically reduce false positives without affecting overall detection performance.

TABLE I
SUMMARY GAZE CHARACTERISTICS DURING STATIC VIEWING

Subject	Y-noise MEAN±STD* (DVA)	Y-noise MEAN±STD* (DVA)	Saccade Amplitude MEAN±STD* (DVA)
1	0.10±0.08	0.23±0.15	5.34±4.41
2	0.06±0.02	0.12±0.06	4.27±3.16
3	0.08±0.05	0.11±0.06	5.06±5.46

*STD=standard deviation

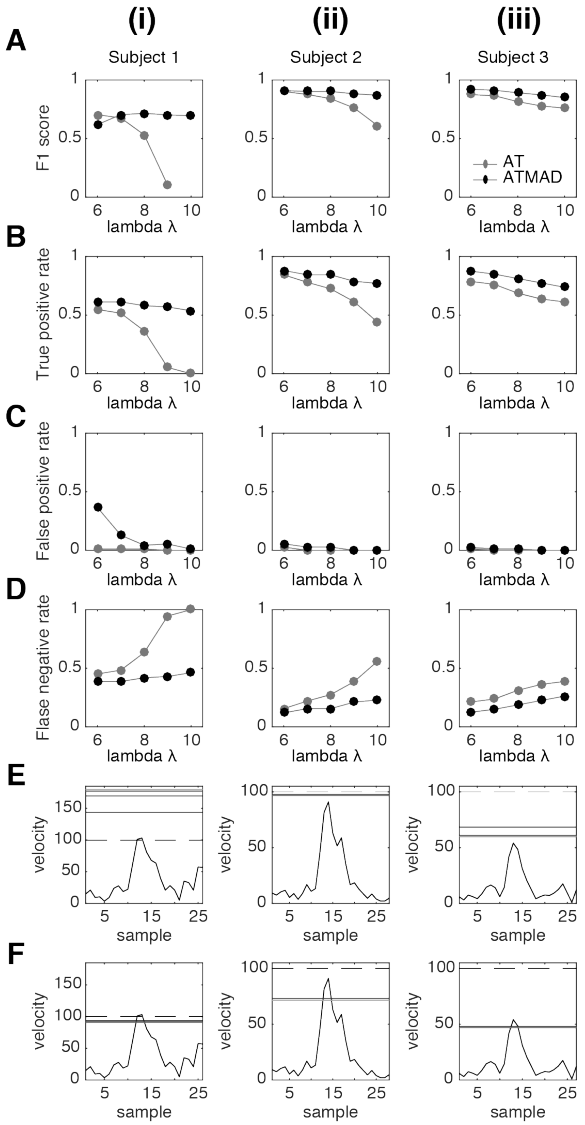

Fig 5. AT-MAD-excise outperforms AT-excise on real world data
(A) F1 score of the AT-excise (grey) and AT-MAD-excise (black) algorithms on one minute of human gaze data for three separate subjects (i-iii). With the exception of Subject 1 at lambda value of 6, AT-MAD-excise outperforms AT-excise. In subject 1 (Ai), the fifth data point is undefined because precision and recall are zero. (B-D) Same as in (A) but depicting the true positive rate (B), false positive rate (C), and false negative rate (D). AT-MAD-excise has a superior true positive and false negative rate to AT-excise. The change in either is linear for AT-MAD-excise, whereas that is not the case for AT-excise. In subject 1 (Ci), AT-MAD-excise is prone to detect a substantial number of false positives for lambda less than 8. (E-F) Example saccades that AT-excise could not detect (E), but AT-MAD-excise could (F). Conventions are the same as in Fig 1.

Discussion

The current work improves on saccade detection by robustly estimating a threshold while mitigating the influence of saccades themselves. The use of a robust estimator allows the estimation of a threshold unbiased by the very signal we are trying to detect (Leys et al., 2013). This allows a more accurate estimation of the background noise levels. We find that the use of MAD, a robust estimator of the standard deviation not biased by outliers, improves the detection of saccades relative to previously published versions of this algorithm (Nyström & Holmqvist, 2010), particularly at higher noise levels. There is both an increase in the number of correctly identified saccades, and a decrease in the number of falsely identified saccades. Moreover, because MAD is not sensitive to outliers in the data, its use allows the experimentalist to confidently define their desired level of confidence.

The goal of the algorithm by Nyström and colleagues is to find the lowest possible threshold that can reliably differentiate saccades from noise and fixations. However, as we have shown, the resulting threshold remains sensitive to the saccades themselves. One reason is that data samples that fall just below threshold (i.e. those flanking detected saccades) can still influence the computed threshold. This concern can be alleviated by excising a number of samples in the inter-saccadic flanks. However, this procedure introduces another user-defined parameter, namely, the duration/number of samples to discard. It is likely the case that the optimal number of samples to discard varies by experimental condition, manipulation, or hard-ware considerations. By considering all of the data, this concern is obviated.

The algorithm we consider operates in the velocity domain, but another alternative is to consider the acceleration domain (A. Duchowski et al., 2002; A. T. Duchowski, 2003; Larsson et al., 2013), particularly when other relatively high velocity eye events, such as smooth pursuits, are present (Larsson et al., 2013). Because the acceleration is a second order derivative, the effect of outliers is amplified. Thus, the use of the MAD estimator would likely improve threshold estimation, especially in combination with an iterative threshold estimation as Nyström and colleagues proposed (Nyström & Holmqvist, 2010). In fact, preliminary tests in our lab on other data where smooth pursuits are prevalent have shown this to be the case (data not shown).

We have shown that the benefits of the MAD estimator mainly accrue at high noise levels. This is particularly important in non-ideal experimental conditions that can lead to noisier data, such as unrestrained viewing, or when working with younger, older, or clinical populations

(Bargary et al., 2017; Rütsche et al., 2006), as well as more in more realistic virtual or game-like settings (Clay, König, & König, 2019). The other benefit of MAD is that it allows experimentalists to define a confidence level (lambda parameter), as required by the experimental questions/equipment. This could be relevant for studies that look at differences in saccade generation across individuals, or populations; such studies may wish to analyze only well-defined saccades (i.e. a higher confidence threshold) (Bargary et al., 2017). Alternatively, lower thresholds may be used to allow the detection of micro-saccades (Engbert & Kliegl, 2003), although this remains to be tested. The use of MAD allows experimentalists to set confidence levels in an unbiased manner.

The current study focuses on the comparison of one published algorithm with a robust alternative. This algorithm has laudable strengths compared to others, such as data-driven (as opposed to user driven) threshold estimation, and a flexible approach that allows setting different thresholds across different experimental subjects, sessions, trials etc. It performs better than nine published algorithms on experiments with static stimuli (Andersson et al., 2017). That said, it remains an open question whether other threshold-based algorithms would benefit from the use of a robust estimator. Because of the ease of implementation, we believe this could be easily tested by interested readers.

This study is related to a method of threshold detection proposed by Engbert and colleagues (Engbert & Kliegl, 2003), which has been widely used in research into microsaccades (Engbert & Kliegl, 2003; Otero-Millan et al., 2008). They use a different, median-based method of estimating the standard deviation. A formal comparison between the two methods is beyond the purview of the paper. However, the present work should be seen as complementary. It provides a formal and rigorous a comparison of robust vs non-robust statistics in the estimation of the standard deviation for the purpose of threshold determination, and suggests that robust estimation may be superior in general.

The use of robust statistics threshold estimation is meant to complement methods that use more sophisticated approaches to determining saccade onsets and offsets (e.g. (Larsson et al., 2013)). We believe that the proposed method is ideal for detecting the presence of saccades, but not necessarily their precise duration. One future direction would be to test thresholding algorithms that use MAD at different sampling frequencies. We expect larger improvements when sampling at lower frequencies, because (1) this data tends to be noisier, (2) peak velocities cannot be reliably recovered (Mack et al., 2017), and (3) outliers exert more influence with less data. However, this remains to be tested. Additionally, robust estimation may be beneficial in non-human animal models, for which head-free eye-tracking is challenging. Indeed, initial testing in our laboratory suggests that robust threshold estimation improves saccade detection in non-human primates in a variety of different tasks (data not shown).

In conclusion, we present here a simple, easily implementable change to a common step in the analysis of saccades, namely, using a robust estimator of the central tendency and deviation to estimate detection thresholds. The simple change leads to improved saccade detection with a published algorithm. The simplicity of this change should encourage further testing and implementation in other thresholding algorithms.

Ethics and Conflict of Interest

The author(s) declare(s) that the contents of the article are in agreement with the ethics described in http://biblio.unibe.ch/portale/elibrary/BOP/jemr/ethics.html and that there is no conflict of interest regarding the publication of this paper.

Acknowledgements

This work was supported by a grant from the Canadian Institute of Health Research CIHR Grant MOP_102482

References

Andersson, R., Larsson, L., Holmqvist, K., Stridh, M., & Nyström, M. (2017). One algorithm to rule them all? An evaluation and discussion of ten eye movement event-detection algorithms. *Behavior Research Methods*, *49*(2), 616–637. https://doi.org/10.3758/s13428-016-0738-9

Bahill, A. T., Clark, M. R., & Stark, L. (1975). The main sequence, a tool for studying human eye movements. *Mathematical Biosciences*, *24*(3–4), 191–204. https://doi.org/10.1016/0025-5564(75)90075-9

Bargary, G., Bosten, J. M., Goodbourn, P. T., Lawrance-Owen, A. J., Hogg, R. E., & Mollon, J. D. (2017). Individual differences in human eye

movements: An oculomotor signature? *Vision Research*, *141*, 157–169. https://doi.org/10.1016/j.visres.2017.03.001

Chau, V. (2011). A flicker change detection task reveals object-in-scene memory across species. *Frontiers in Behavioral Neuroscience*, *5*. https://doi.org/10.3389/fnbeh.2011.00058

Clay, V., König, P., & König, S. U. (2019). *Eye tracking in virtual reality*. https://doi.org/10.16910/jemr.12.1.3

Dai, W., Selesnick, I., Rizzo, J.-R., Rucker, J., & Hudson, T. (2016, December). *A parametric model for saccadic eye movement*. 1–6. https://doi.org/10.1109/SPMB.2016.7846860

Daye, P. M., & Optican, L. M. (2014). Saccade detection using a particle filter. *Journal of Neuroscience Methods*, *235*, 157–168. https://doi.org/10.1016/j.jneumeth.2014.06.020

Duchowski, A., Medlin, E., Cournia, N., Murphy, H., Gramopadhye, A., Nair, S., … Melloy, B. (2002). 3-D eye movement analysis. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 573–591. https://doi.org/10.3758/BF03195486

Duchowski, A. T. (2003). *Eye Tracking Methodology: Theory and Practice*. https://doi.org/10.1007/978-1-4471-3750-4

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, *43*(9), 1035–1045. https://doi.org/10.1016/S0042-6989(03)00084-1

Holmqvist, K., Nyström, M., & Mulvey, F. (2012). *Eye tracker data quality: what it is and how to measure it*. 45. https://doi.org/10.1145/2168556.2168563

Komogortsev, O. V., Gobert, D. V., Jayarathna, S., Do Hyong Koh, & Gowda, S. M. (2010). Standardization of Automated Analyses of Oculomotor Fixation and Saccadic Behaviors. *IEEE Transactions on Biomedical Engineering*, *57*(11), 2635–2645. https://doi.org/10.1109/TBME.2010.2057429

König, S. D., & Buffalo, E. A. (2014). A nonparametric method for detecting fixations and saccades using cluster analysis: Removing the need for arbitrary thresholds. *Journal of Neuroscience Methods*, *227*, 121–131. https://doi.org/10.1016/j.jneumeth.2014.01.032

Larsson, L., Nystrom, M., & Stridh, M. (2013). Detection of saccades and postsaccadic oscillations in the presence of smooth pursuit. *IEEE Transactions on Biomedical Engineering*, *60*(9), 2484–2493. https://doi.org/10.1109/TBME.2013.2258918

Leys, C., Ley, C., Klein, O., Bernard, P., & Licata, L. (2013). Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, *49*(4), 764–766. https://doi.org/10.1016/j.jesp.2013.03.013

Liston, D. B., Krukowski, A. E., & Stone, L. S. (2013). Saccade detection during smooth tracking. *Displays*, *34*(2), 171–176. https://doi.org/10.1016/j.displa.2012.10.002

Mack, D. J., Belfanti, S., & Schwarz, U. (2017). The effect of sampling rate and lowpass filters on saccades – A modeling approach. *Behavior Research Methods*, *49*(6), 2146–2162. https://doi.org/10.3758/s13428-016-0848-4

Nyström, M., & Holmqvist, K. (2010). An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data. *Behavior Research Methods*, *42*(1), 188–204. https://doi.org/10.3758/BRM.42.1.188

Otero-Millan, J., Troncoso, X. G., Macknik, S. L., Serrano-Pedraza, I., & Martinez-Conde, S. (2008). Saccades and microsaccades during visual fixation, exploration, and search: Foundations for a common saccadic generator. *Journal of Vision*, *8*(14), 21–21. https://doi.org/10.1167/8.14.21

Rütsche, A., Baumann, A., Jiang, X., & Mojon, D. S. (2006). Development of visual pursuit in the first 6 years of life. *Graefe's Archive for Clinical and Experimental Ophthalmology*, *244*(11), 1406–1411. https://doi.org/10.1007/s00417-005-0248-4

Salvucci, D. D., & Goldberg, J. H. (2000). *Identifying fixations and saccades in eye-tracking protocols*. 71–78. https://doi.org/10.1145/355017.355028

Warby, S. C., Wendt, S. L., Welinder, P., Munk, E. G. S., Carrillo, O., Sorensen, H. B. D., … Mignot, E. (2014). Sleep-spindle detection: crowdsourcing and evaluating performance of experts, non-experts and automated methods. *Nature Methods*, *11*(4), 385–392. https://doi.org/10.1038/nmeth.2855

Watson, M. R., Voloh, B., Naghizadeh, M., & Womelsdorf, T. (2017). Quaddles: A multidimensional 3D object set with parametrically-controlled and customizable features. *Behavior Research Methods*, *In press*. https://doi.org/10.1101/194381

Wilcox, R. R. (2012). *Introduction to Robust Estimation and Hypothesis Testing*. Amsterdam: Academic Press. nlebk.