TEMPORAL AND SPATIAL COORDINATION OF THE HIPPOCAMPUS, RETROSPLENIAL CORTEX, AND ANTERIOR THALAMUS

A Thesis
Presented to the Faculty of the Graduate School
of Cornell University
In Partial Fulfillment of the Requirements for the Degree of
Master of Science

by
Rachel Swanson
February 2016

ABSTRACT

The ability to learn about an environment is fundamental to the survival of an organism. The hippocampus is a highly studied component of what is now known to be a distributed, brain-wide network; the coordinated activity of which supports this cognitive function. This thesis first examines the functional contributions made by a less studied node in this network, the retrosplenial cortex, which receives direct sensory and hippocampal input. More broadly, the integration of distributed information processing throughout the brain is necessary. It has been suggested that this can be achieved via the temporal coordination, or coherence, between populations of neurons. The next section of this thesis examines oscillatory co-activity across time in the hippocampus, retrosplenial cortex, and anterior thalamus during a task previously illustrated to involve these regions (Smith 2011). Distinct oscillatory features and single cell action potentials are found to uniquely correlate to brain and behavioral states.

BIOGRAPHICAL SKETCH

Rachel is a born Canadian and raised Californian who went to Rutgers University for her undergraduate studies, majoring in Biochemistry, Psychology, and Public Health. She is fascinated by the self-coordination of the brain and thinks building flexible AI will do a lot more than generate better netflix recommendations.

ACKNOWLEDGEMENTS

Huge thank you to my Masters Thesis adviser, Dr. David Smith, for guiding me through what was initially complete bewilderment at the breadth and depth of current human understanding of the brain. Without his patient and keen guidance, I would not be where I am today. David has a rare mix of qualities - he is both a brilliant observer of the physical world and of the interpersonal world. Thank you for caring as much about your students' happiness and well-being as you care about your own. Thank you also to Dave Bulkin, who took me in when I was a fledgling coder and taught me that hard work was the best way to get anything done. Thank you of course to the rest of the lab, especially the undergraduate students with whom I had the honor of working, Anna and Betsy. Importantly, thank you also to Barb, who captured my intellectual fantasies in my graduate student infancy and gave me a framework by which to approach thinking critically about my own intuitions. Lastly, thank you to Thom, to David, and to the broader Cornell Psychology community, with whom I felt a true sense of belonging and community during my tenure in the department.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

The hippocampus (HPC) is one of the most highly studied regions of the mammalian brain. This is in part due to its role in human episodic memory and Alzheimer's disease (Scoville 1957; Squire 2009), the discovery of long term potentiation (Lomo 1966), and its multi-synaptic removal from primary sensory transducers and motor-effectors. Paradoxically, a predominant feature of the HPC is the presence of pyramidal cells whose firing rates have clear external correlates utilizing sensory and motile information - namely *place cells* (O'Keefe 1971) - providing a tractable link between higher order cognition and single cell physiology. Bilateral damage to the hippocampus in humans results in severe anterograde and temporally graded retrograde amnesia (Scoville 2000), and manifests most prominently as a deficit in spatial navigation and contextual learning (*for review see*: Smith 2008) in mice, rats, bats, cats, rabbits, chinchillas, guinea pigs, and monkeys (Geva-Sagiv 2015). It has been proposed that the network-level algorithms facilitating navigation in mammals ranging from mice to humans have been co-opted over evolutionary timescales to subserve episodic memory - i.e. navigation in mnemonic space (Buzsaki 2005). Notably, the association area of the cortex (including the HPC and related circuitry) has shown the largest expansion and differentiation during mammalian evolution. As such, the study of animal navigation and contextual learning, i.e. the learning of continuously present features of an environment, is an excellent paradigm by which to study a function essential for mammalian viability - memory of when and where.

Extended Hippocampal Circuitry: anatomy and disconnection studies

The hippocampus does not accomplish these integral cognitive feats alone, however. Extensive anatomical, lesion, and electrophysiological studies have identified the entorhinal cortex, medial prefrontal cortex, retrosplenial cortex, medial mamillary and supramamillary nuclei of the hypothalamus, anterior nuclei of the thalamus, amygdala, inferior colliculus, and several brainstem nuclei as members of an extended circuit by which memories are encoded, stored, and recalled (Aggleton 1999). For the purposes of this thesis, I will center discussion on these features in the hippocampus (HPC), retrosplenial cortex (RSC), and anterior thalamic nuclei (ATN); the primary foci of the following experimental chapters. The ATN, HPC, and RSC are densely interconnected, via both direct and indirect pathways. The ATN is a primary subcortical projection target of the subiculum of the hippocampal formation, directly through the fornix and indirectly through the mamillary bodies. The ATN projects back to the HPC indirectly through the RSC. The RSC is situated in the posterior-most region of the dorsal midline cortex and provides substantial visuospatial sensory information directly to the HPC (Burwell 1998), as it is uniquely positioned between the hippocampal formation and a network of dorso-medial cortical areas shown to be involved in spatial memory and visual processing, including the parietal and cingulate cortices (VanGroen 1990). Lesions of the fornix (Vann 2009), a white matter tract projecting from the hippocampus to the mamillary bodies and anterior thalamus, the anterior thalamus (Jenkins 2002), and the hippocampus (Jenkins 2006), result in a decrease in IEG expression in the the hippocampus, subicular complex, and the entorhinal and retrosplenial cortices. Furthermore, lesions of both the RSC and ATN result in profoundly impaired contextual and spatial memory acquisition (e.g. O'Mara 2013; Law 2012; Waburton 1997; Waburton 1999; Danker 2010; Keene 2008),

perhaps through the disruption of HPC activity (Cooper 2001). Humans with RSC and ATN damage exhibit anterograde amnesia similar to that observed in patients with hippocampal damage, a symptom uniquely caused by damage in a small subset of the regions associated with the hippocampus. Korsakoff's syndrome, a chronic memory disorder resulting from alcohol misuse, is characterized by the atrophying of the mamillary bodies and anterior thalamic nuclei (Harding 2000). Human patients with retrosplenial lesions show a marked deficit in episodic processing, and a decreased ability to use navigational cues to move through a known environment (Aggleton 2010; Valenstein 1987; Sculpizio 2013).

Extended Hippocampal Circuitry: electrophysiology

Do the HPC, ATN, and RSC make distinct functional contributions to mnemonic processes, and can these contributions be experimentally dissociated beyond those differences observed behaviorally? In short, yes and yes. Distinguishing single-cell physiological features of this circuit are head direction cells, grid cells, place cells, and cells that fire in response to complex conjunctions of location and cue. These cells are not present in equal numbers throughout the hippocampal circuit, and it is thought that coordination across each of these regions enables the eventual formation of hippocampal place cells. Head direction cells, as the name implies, show an increase in firing rate selectively when an animal's head is pointed in a particular direction in allocentric space, independent of gaze direction (Taube 1990, 2003; Robertson 1999). Head direction cells are found in the lateral mammillary nucleus, the anterior thalamic nucleus, the pre and post subiculum, (Taube 1995) and to a small degree in the retrosplenial cortex (Cho 2001). Grid cells, found in the entorhinal cortex one synapse upstream of the hippocampus (e.g. Moser 2014), are cells that fire (a colloquialism often used to indicate a

statistically significant increase in the rate of action potentials generated by a cell) in multiple spatial locations, forming interlocking equilateral triangles along the surface of an enclosure. Finally, place cells, cells that selectively increase their firing rate in a particular location are found in the hippocampus (O'Keefe) and retrosplenial cortex (Cho 2001; Smith 2011). Importantly, cells that fire in response to complex conjunctions of location and object have also been found in the HPC and RSC (Komoroski 2009; Smith 2012; Wood 2000). The firing properties of these cells are of particular interest, because they are viewed as being at 'the top' of this hierarchy, and it has been suggested by Eichenbaum and colleagues that the joint encoding of context and object by a single neuron are the single cell constituents of episodic memory - namely a memory of an object situated in a context (Komoroski 2009). Furthermore, it has been shown very recently in one of the first high-profile retrosplenial electrophysiological studies that the spatial firing of most RSC neurons reflects conjunctions between an animals position and two or three frames of reference (Nitz 2015), as predicted by many previous physiological and lesion studies, suggesting that the RSC serves as a site of convergence of the many frames of reference (e.g. egocentric and allocentric) relevant to an animal navigating through space.

Consolidation

The above single cell physiology studies begin to paint a picture of the differential contributions made by each of these regions, but how do these properties relate to memory? Importantly, each time an animal is introduced to a novel environment, a different subset of hippocampal pyramidal cells increase their firing rate in response to novel locations. Once an animal is re-introduced to a learned context, the same HPC cells fire in the same locations (*for review see* Smith 2008). It is thought that the intrinsic circuitry of the HPC lends itself to the

rapid encoding of novelty, forming functionally-connected circuits or ‘schemas’, allowing an animal to learn (e.g. McKenzie 2014). During sleep and quiet wakefulness, these distinct hippocampal network states are reactivated in tandem with the reactivation of the cortical cells active during original learning. It is thought that the hippocampal-neocortical dialogue is essential for the stabilization of memories, and that semantisized information is stored in a distributed fashion throughout the neocortex, biasing the processing of incoming sensory information. This can be thought of as ‘schemas’ or ‘contexts’. These are the general principles of systems consolidation theory, different versions of which debate the eventual independence of memories from the hippocampus. Although these theories have been present for decades, only recently has the technology been available to test these principles directly. For example, the direct reactivation of an ensemble of retrosplenial cortical cells active during fear conditioning resulted in a fear conditioned response soon after learning, independent of hippocampal activity. This coherent ‘neocortical memory context’ in the retrosplenial cortex suggests both redundancy in the encoding of contextual memories (Cowansage 2014) and illustrates the importance of the retrosplenial cortex in encoding context. RSC-hippocampal interactions continue to be important for learning in subjects that have been fully trained, or are considered ‘experts, as the temporally coordinated activity of these regions facilitates the updating of existing memories (Tse 2007; Tse 2011).

The subsequent experimental chapters serve to further examine the functional relationships and contributions made by each of these regions. The first study attempts to further disentangle the role the HPC and RSC play in the encoding of cues and contextual information. The second study examines coordinated population level activity in the HPC, RSC, and ATN,

measured via the *local field potential* signal, which has lent great insight into the coordination and exchange of information amongst distributed brain regions.

CHAPTER 2

RSC LESIONS DO NOT IMPAIR CONDITIONAL DISCRIMINATION LEARNING

INTRODUCTION

The RSC and HPC have both been demonstrated as essential in the the formation of episodic memories, operationalized for the purposes of studying these functions in more detail in animals as the encoding of items and contextual information, as discussed above. As learning occurs, both RSC and HPC neurons consistently fire differentially in response to environmental events, including spatial (place fields) and non-spatial (odor cue/item) information (Vann 2009; Bucci 2008). Moreover, much evidence suggests that individual rat HPC neurons develop conjunctive, or 'item-place', firing selectivity (e.g. Wood 2000; McKenzie 2014). This means that a single HPC neuron will encode information about both odor cue and place or context, 'binding' the two. Importantly, these complex item-place responses develop later in learning, from simple place fields (Komorowski 2009). These conjunctive HPC responses are thought of as the building-blocks of episodic memory representations. Does the HPC develop these responses independently or are other brain regions involved? Prior work in our lab suggests RSC involvement is critical. First, RSC neurons were found to fire in response to cues that predict reward (Gabriel 1993). Second, 44% of RSC neurons responded this way on the first day of training (Smith 2012). Lastly, fornix lesions, which partially disconnect the RSC from HPC, did not disrupt RSC selective neuronal responses to cues that predicted reinforcement (Smith 2004). Thus, the RSC encodes behaviorally relevant cues early in training, later encoded by HPC item-place responses, and does not need HPC input to do so. Our hypothesis is that the RSC encoding

of cue significance provides critical input for the development of hippocampal conjunctive representations. In order to investigate this, we will induce lesions in the RSC and then train subjects in a conditional discrimination task, which demands the encoding of both context and items, developed previously (*see methods for additional detail;* Komorowski 2009).

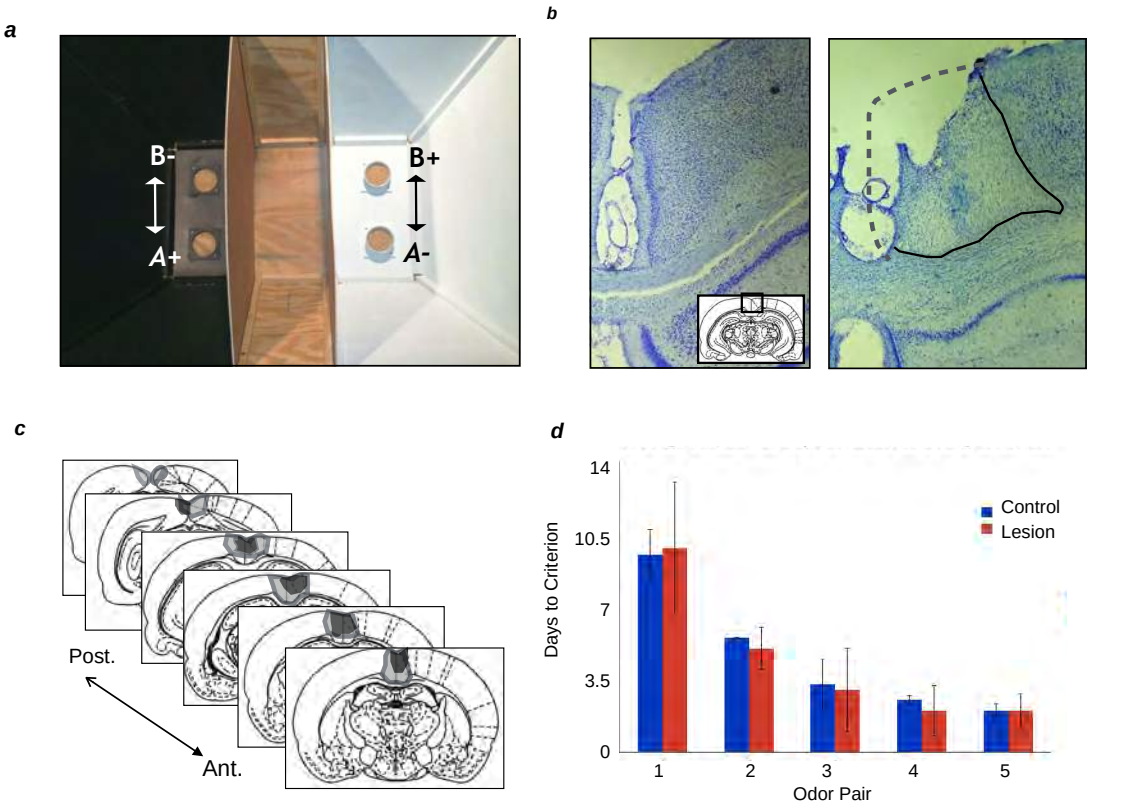

Percent Lesion	1748	1769	1782	1802
Left Hemis.	42.5	45.71	32.18	30.21
Right Hemis.	45.98	30.47	26.71	52.3
Total	44.24	38.09	29.45	41.255
Total in sx range	51.87	44.68	34.52	48.23

AP (mm)	ML (mm)	DV (mm from surface of skull)	Vol (µL)
-2.2	+- 0.5	-3.0	0.35
-3.9	+- 0.5	-3.0	0.35
-5.5	+- 0.5	-3.5	0.20
-5.5	+- 1.0	-2.8	0.35
-6.7	+- 1.1	-2.8	0.35
-8.0	+- 1.3	-2.8	0.30

FIGURE 1 | **Experimental paradigm and learning results**
a The conditional discrimination testing chamber has two contexts distinguishable by odor and color. Olfactory pair valence is reversed in the opposite context, and multiple olfactory pairs (1 - 5) are presented to the animal. ***b*** Example histological section demonstrating typical extent of exocitotic lesion. **c** Lesion extent along the posterior to anterior axis; histologically determined lesion extent of experimental animals is tabulated below. **d** Above: Average learning curve across control (n=4) and lesion (n=4) animals. There is no significant difference in learning between groups, and all animals can successfully learn the task. Below: stereotaxic coordinates for lesion surgeries.

RESULTS

All control and lesion animals learned this complex discrimination task successfully. Learning on the first odor pair took an average of 9 days, as can be seen in ***figure 1d***, however with the introduction of each new odor pair, animals became progressively faster at successfully discriminating the rewarded cue based on the conjunction of context (white or black) and odor (A or B). Because our pre-determined criterion was 2 days, it seems as if performance is reaching an asymptote, however on the last odor pair animals performed at established criterion. **Figure 2a** demonstrates a representative learning curve for one animal across sessions, which was similar across all animals. As can be seen from this learning curve, another performance metric that could be assessed is the drop in performance with the introduction of each new odor pair. It is possible that although animals with RSC lesions perform at similar levels on average, their ability to generalize a learned rule was impaired. This was not the case, however, as can be seen in ***figure 2b,*** which demonstrates clearly that both lesion and control groups exhibited a similar drop in performance with the introduction of each new odor pair., improving significantly over time. A final possibility is that, due to the variability in the extent of the lesions produced (***figure 1c)***, behavioral differences were being obfuscated by bifurcating our data. We then plotted percent lesion against percent change in performance from odor pair 1 to 2, when the biggest drop in performance occurred, and did not find a significant correlation between lesion size and performance (***figure 2c***). In all metrics assessed, we did not find a significant difference between animals with RSC lesions and animals without lesions.

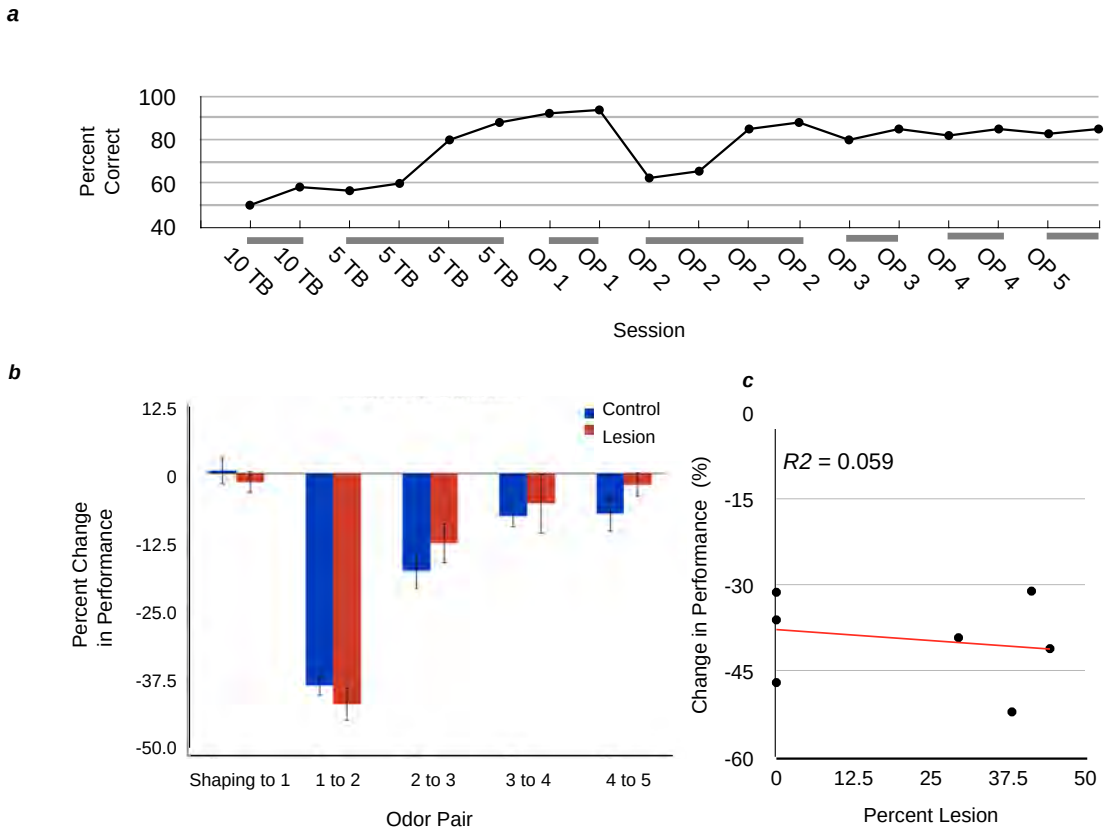

FIGURE 2 | **Further inspection reveals no deficit in learning of complex discrimination task in animals with retrosplenial cortical lesions**

a Representative learning curve of an animal with resplenial lesions. Breaks in teh gray bar along the x axis signal changes in odor pairs. Rats gradually learn the conditional discrimination rule during the first odor pair and then quickly learn to adapt this strategy to new odor pairs. ***b*** Percent change in performance on the first day of each new odor pair. No significant difference between control and lesion groups. **c** Lesion extent is not significantly correlated with percent drop in performance from odor pair 1 to 2.

DISCUSSION

Given the current state of the literature, and the sheer size of the RSC (in the adult rat, the posterior cingulate cortex and RSC are not differentiated, and the full extent of the RSC comprises approximately 1/10 of the cortex), we are surprised at this null result. Furthermore, in a recent study in our lab (Adam Miller, not yet published), RSC lesions induced using the same surgical method were found to cause impairments in a delayed spatial alternation task with respect to controls with similar lesion extents. Importantly, this task differs from our task in that it is inherently spatial and does not involve manipulations of context. A possibility is that although much evidence supports the notion that the RSC is involved in early encoding of cues, perhaps this information is only utilized by the hippocampus in more complex tasks involving spatial manipulation of said cues. In the conditional discrimination task, space with respect to cues explicitly does not matter, as the location of the cups are manipulated from trial to trial. Lastly, as discussed previously, the RSC is at the end of a long stream of brain regions along the dorsal midline of the cortex that are known to be involved in the processing of visual information (called the dorsal stream). Keeping this, and the privileged monosynaptic connectivity between the HPC and the olfactory bulb, it is possible that an impairment would have been seen had we used visual instead of olfactory cues to distinguish rewarded vs non-rewarded item pairs. This may be difficult, however, because one would have to ensure that the rat could not *smell or* see a difference between the two sets of cues (as we use to pairs of the same odors as the task currently stands). This could be controlled for, however, and it is suggested that this second study be conducted.

METHODS

Subjects and surgical procedures

Surgical subjects were 8 male Long Evans rats (Charles River Laaboratories, Wilmington, MA) weighing 300-350 g at the time of surgery. Rats were selected in pairs and received either a retrosplenial lesion or went through the same surgical procedures, but received injections of saline instead of an excitatory NMDA neurotoxin. Each pair of rats went through all training and then the next pair of rats was selected, underwent surgery, and received training. Four control and four lesion rats were trained total.

In the retrosplenial lesion group, an excitatory NMDA neurotoxin (was induced along the entire rostral-caudal axis of the RSC. These lesions were produced by infusing 0.3 µl of a 100 mM solution of N-methyl D-aspartate (NMDA; dissolved in 0.9% saline) into the RSP. The solution was infused at a rate of 0.1 µL/min through a 1µL Hamilton syringe (Hamilton Company, Reno, NV). During infusion, the syringe was lifted ~0.1 mm to aid diffusion. After each injection, the syringe was left in place for 3 min before being slowly removed. The stereotaxic coordinates and volumes injected are specified in ***figure 1d bottom***. Control rats went through all of the same surgical procedures, differing only in the content of what was injected. The surgical procedure itself took on average of six hours per rat. Throughout this procedure, rats were anesthetized, resting on a heating pad, and monitored by the experimentor. Injection cannula were built by hand using hypodermic glass tubing placed in 23 ga cannula so the injector could be clamped to a stereotaxic manipulator during surgery. 12 holes total were drilled through the rat skull after shifting the skin above, and the specified volume of liquid was injected slowly over the course of one minute. We then left the injector in the tissue for 3 minutes

before moving on to the next site in order to ensure all fluid had been injected. Upon the completion of surgery, rats received stitches to close the wound and were given an antibiotic (5mg/kg Baytril) and an analgesic (5 mg/kg ketoprofen). As the rats recovered from the anaesthetic, a prophylaxis against seizures was administered (diazepam; 0.2cc; 10 mg/ml, i.p.; Sabex, Boucherville, Quebec). The same surgical procedures were used for the Sham rats except that no damage was done to the skull or brain. A follow-up injection of ketoprofin was given from 8 to 12 hours later, from the end of surgery. All procedures complied with the guidelines established by the Cornell University Animal Care and Use Committee. After one week of recovery from surgery, the rats were placed on a restricted feeding regiment (80-85% of free feeding weight) and then began training.

Training Procedure

Animals were trained in a conditional discrimination task very similar to that used by Komorowski (2009) and McKenzie (2014); this task was developed in the Eighenbaum lab. The task itself is visualized in ***figure 1a.*** The conditional discrimination apparatus is one box with two dividers, separating the enclosure into three areas: the 'black context', the inter-trial waiting area, and the 'white' context. The white and black contexts differ in texture, color, and ambient odor. Two ramekins were used containing odors A and B. The training made use of a well-known digging task used to study olfactory memory (Eichenbaum 1998) in which rats are trained to dig in cups of odorized bedding material to retrieve buried 45 mg sucrose pellets (Bioserve, Frenchtown, NJ). In each A/B pair of odors, one cup was rewarded and one was not, and the cup valence was reversed in each context for the entire duration of training with that odor pair (eg context black: A+ B-, context white: A- B+). At the beginning of a session, the rat was placed in

the inter-trial waiting location. The experimenter then baited the odor pair (A or B) identified on a training list, placed the ramekin pair down at the same time in the appropriate context, either A or B, and lifted the waiting chamber divider. The rat then approached the cups. Whichever cup was approached first counted as the rat's selection, and rats were not allowed corrections. The ordering of contexts and arrangement of the ramekins within a context was arranged pseudorandomly. 40 trials were performed in one session, using one of three separate lists to ensure rats did not learn a list. The lesion and control pair first received one day of acclimation, during which they were allowed in all regions of the conditional discrimination box to forage for sucrose pellets. They then were trained in ten trial blocks (ie ten black context in a row, ten white context in a row) for two days in order to facilitate learning. Rats then were trained on lists of 5 trial-blocks until they reached an 80% for two days in a row performance criterion. Then rats were switched to 1 trial-blocks with the same odor pair, during which behavior dropped minimally. After reach the same behavioral criterion, the odor pair was then changed to C/D, and rats went through the same training procedure, but were presented odor pairs in only one trial blocks for all subsequent odor pairs, including C/D. Upon completion, rats were euthanized and histology was performed to determine lesion extent, ***figure 1d.***

Histology

Brains were sections at 40 micrometers, mounted on slides and stained with cresyl violet. Each section was inspected under a microscope, and the extent of the visible lesion was traced onto the corresponding anatomy coronal slice drawing digitally. Pixels within the outline region were then automatically counted and recorded.

CHAPTER 3

COORDINATION OF THE RETROSPLENIAL CORTEX, HIPPOCAMPUS, AND ANTERIOR THALAMUS VIA NEURONAL OSCILLATIONS

INTRODUCTION

Oscillations in the brain

Discussion thus far has focussed on behavioral studies and single cell physiology. Another signal that can be measured from the surface of the cortex or extracellularly in the brain is the local field potential (LFP, or EEG if recorded non-invasively). This is an aggregate signal that reflects net voltage fluctuation in the local vicinity of a recording electrode. This signal cannot be easily decomposed into its respective ionic current sources, unlike a sequence of action potentials being emitted from a single cell, however it is still of great use, as it is an excellent proxy measure of the collective activity of many neurons in a local region (Einevole 2013). Originally proposed by Donald Hebb in 1949, co-active neurons, or cell assemblies, are thought to serve as the syllables from which neural ‘words’ and ‘sentences’ are constructed (Buzsaki 2006). In other words, single cells do not act in isolation to convey information. The temporally coordinated activity of transient groups of cells within and across brain regions, sometimes referred to as ‘networks’, is an essential feature of brain computation (Billeh 2014; Buzsaki 2005). Hence, a signal that can measure this population-level activity with great temporal precision, the LFP, provides information relevant to local and global processing (particularly when current experimental methodology restricts the number of single cells that can be simultaneously recorded from). This

is supported by decades of research that have correlated fluctuations in power in frequencies of the spectrally decomposed LFP signal* with various internal and external behavioral states. Low frequency delta rhythms (~1-4 Hz) markedly distinguish sleep from waking (Steriade 2001), higher frequency theta band oscillations (~6-12 Hz) distinguish active exploration from quiet wakefulness, and even higher gamma band frequencies distinguish attentive from inattentive states (~30-100 Hz) (Buzsaki 2006). Furthermore, different brain areas (including layers within a region - highly varied in their cell profiles, cytoarchitecture, and connectivity) have distinct *spectral profiles*, or peaks in the power of the decomposed LFP signal, that reliably co-vary with the aforementioned internally and externally driven variables.

Oscillatory synchrony as a measure of functional connectivity

The brain is a highly distributed system with no central command, performing multiple computations in parallel*. The way by which this distributed information is coordinated within and across regions in the service of adaptive behavior is a central question in systems and computational neuroscience. It has been proposed that transient oscillations organize local neuronal activity, and if these oscillations occur in synchrony in two disparate brain areas - even over long distances - this facilitates the exchange of information by organizing spike-timing on similar timescales and therefore increasing the probability of coincident excitation of

**The LFP is a single time-series that can be decomposed into its spectral components using standard signal processing techniques. This results in a family of sinusoidal oscillations at different frequencies, with power at each frequency band varying with respect to one another and time. It is important to note that this imposes an oscillatory structure on raw LFP data, and ignores non-oscillatory activity - this is beyond the scope of this thesis, but see He 2010: The temporal structures and functional significance of scale-free brain activity*

**the computation analogy has its weaknesses, but serves the purposes of this thesis.*

downstream neurons (Fell 2011). This is often referred to as ‘functional connectivity’ (Buzsaki 1994, Fries 2007), which is distinguished from ‘structural connectivity’. Notably, the anatomical connectivity of the brain cannot change on the timescales necessary for behavior; the relationship between structural and functional connectivity is an active area of research (Sporns 2010). This synchronous activity can be coordinated by the coupling of the phases of two oscillations, the coupling of the amplitude of one oscillation with the phase of another, or by correlated changes in amplitude (Siegel 2012). A benefit of all of these approaches is they allow the observation of dynamic coupling in *time*, which can then be correlated with other variables of interest.

Spectral Features of the Hippocampus and related circuitry

Theta (~6 - 12 Hz)

It is well known that theta rhythmicity is a dominant oscillatory feature of the hippocampal formation during behavior, playing a fundamental role in the coordination of local circuit computation via the temporal organization of local neuronal ensembles (Buzsaki 2006). Furthermore, theta range oscillations and theta-modulation of neuronal firing rates have been documented in a broad range of brain regions, including the entorhinal cortex, perirhinal cortex, medial mamillary and supramamillary nuclei of the hypothalamus, amygdala, inferior colliculus, and several brainstem nuclei (Buzsaki 2002). Many of these regions share monosynaptic connectivity with the hippocampus. Coherent changes in theta oscillatory activity and neuronal spiking have been correlated with spatial learning (e.g. DeCoteau 2007; van Wingerden 2010; Lisman 2013). For example the pacing of prefrontal hippocampal neurons by hippocampal theta and the increase of coherent theta oscillations in the mPFC and HPC at periods of behavioral significance (e.g. the choice point of a maze) or during memory retrieval have been documented

in numerous studies (e.g. O'Neill 2013, Siapas 2005). Theta coherence has also been documented between the ATN and HPC (Albo 2003) and the HPC and RSC in rats (Gabriel 2004; Young 2009), and between the HPC and RSC more recently in humans, finding transient increases in theta synchrony only during episodic memory retrieval (Foster 2014).

Gamma (low: ~30 - 60 Hz, high: ~ 70 - 100 Hz)

Both theta and gamma rhythms have been studied extensively in relation to memory processes in the rodent and primate brain. Gamma rhythms in the HPC formation are important locally for cell assembly synchrony as well as globally for communication between disparate brain regions during various aspects of hpc dependent memory. (Montgomery and Buzsaki, 2007; Tort et al., 2009). Gamma can be further subdivided into low and high gamma, which occur preferentially during the peak and trough of the theta rhythm. A recent study by Schomberg and colleagues has demonstrated that this millisecond timescale peak to trough variation in gamma amplitude reflects the selective modulation of inputs to the HPC from two of its major upstream afferents - the entorhinal cortex and area CA3 of the HPC (2014). This study is of key importance, in that it demonstrates precisely how a frequently observed brain rhythm supports communication between regions - via temporal segregation of inputs.

Beta (~20-40 Hz)

A less documented spectral feature of the extended hippocampal circuit is the beta rhythm. Notably, although the observed frequency bounds of this region overlap with low gamma, a number of studies have identified these frequency bands as indicating functionally distinct

underlying processing. A review of these studies reveals two seemingly contradictory themes - 1) beta power transiently increases in response to reward receipt or anticipation and occurs in bursts in a novel environment, vs. 2) beta power increases with learning (which implies a decrease in novelty). In a human MEG study, a sharp increase in beta power was observed over the frontal lobe, approximately 100 ms after viewing a cue predicting reward. A concurrent decrease in theta power was additionally seen. Notably, the power of this burst in beta activity increased as a function of contextual novelty (Bunzeck 2011). Note, a different study in rats (van Windergerden 2010) did not find a decrease in theta-band activity, but instead an increase in theta-band phase locking of orbitofrontal neurons during reward expectancy. A recent study published in Nature from the Moser lab has also identified a transient increase in 20-40 Hz power following presentation of a cue prior to receiving a reward in the HPC and lateral but not medial entorhinal cortex. This burst in beta power was coherent between the HPC and mEC and increased as a function of learning (Igarashi 2014). Berke et al (2008) recorded in mouse hippocampi during the exploration of a novel environment and found transient bursts in 23-30 Hz frequency oscillations in addition to the presence of a strong and continuous theta rhythm. Finally, Howe et al (2011) found a similar increase in beta power over the course of learning in the striatum, and a concurrent decrease in gamma power, both occurring as transient bursts following reward receipt. A computational model by Kopell et al (2011) suggests a role for beta-frequency dynamics in short-term memory, and in a review by Engel & Fries (2010), it was suggested that beta plays a role in ‘maintaining the current cognitive state’.

Present Study

In the present study, we performed exploratory analyses examining the temporal dynamics of HPC, RSC, and ATN interaction via simultaneous recording of LFP and action potentials in all three regions. We then sought to characterize these dynamics. No studies to our knowledge have looked at HPC, RSC, ATN joint-dynamics, and at most have examined HPC-RSC or HPC-ATN oscillatory coupling in the theta range. This data was recorded by Dr. David Smith, and has resulted in the publication of three separate papers examining spiking activity in the RSC and HPC (Gill 2011; Smith 2012). LFP and LFP-spike relationships have not previously been analyzed.

RESULTS

Theta range oscillations dominate the spectra of all three regions

Given the intimate connectivity of the RSC and ATN with the HPC, and predominance of the theta rhythm in the extended hippocampal circuit, the presence of a spectral peak in the 6-12 Hz range was expected, and found, in all three regions (*see spectral density plots in figure 1c, these are averaged across learning stage and tetrodes within a region*). Spectral density plots are made for each rat in order to facilitate visual comparison of the similarities and differences between rats. A theta peak is clear in all four rats; subsequent analyses look at averages across rats. The secondary peak around 18 Hz is likely a harmonic of the theta rhythm, and is a bi-product of this type of signal decomposition (harmonics of decreasing power will be seen every integer multiple of our peak signals; Masimore 2004). Raw (unfiltered) traces simultaneously recorded in each region and taken during rat mobility are plotted in *fig 1b,* demonstrating a strong theta rhythmicity that is easily visible by eye. An 8 Hz sine wave is overlaid in the lower left of this figure for visual comparison.

Theta and beta range oscillations are differentially modulated by learning and task epoch

Average spectrograms are plotted across rats in ***figure 2***, separated into categories of interest defined a priori: brain region (HPC, ATN, and RSC) and learning stage (Pre-Learning and Asymptote). Pre-learning through asymptote stages of learning occur on the same maze in the same room, however the pre-learning task differs from the learning task in that all arms of the plus maze have an equal probability of reward (*see: Methods).* Each spectrogram is aligned to

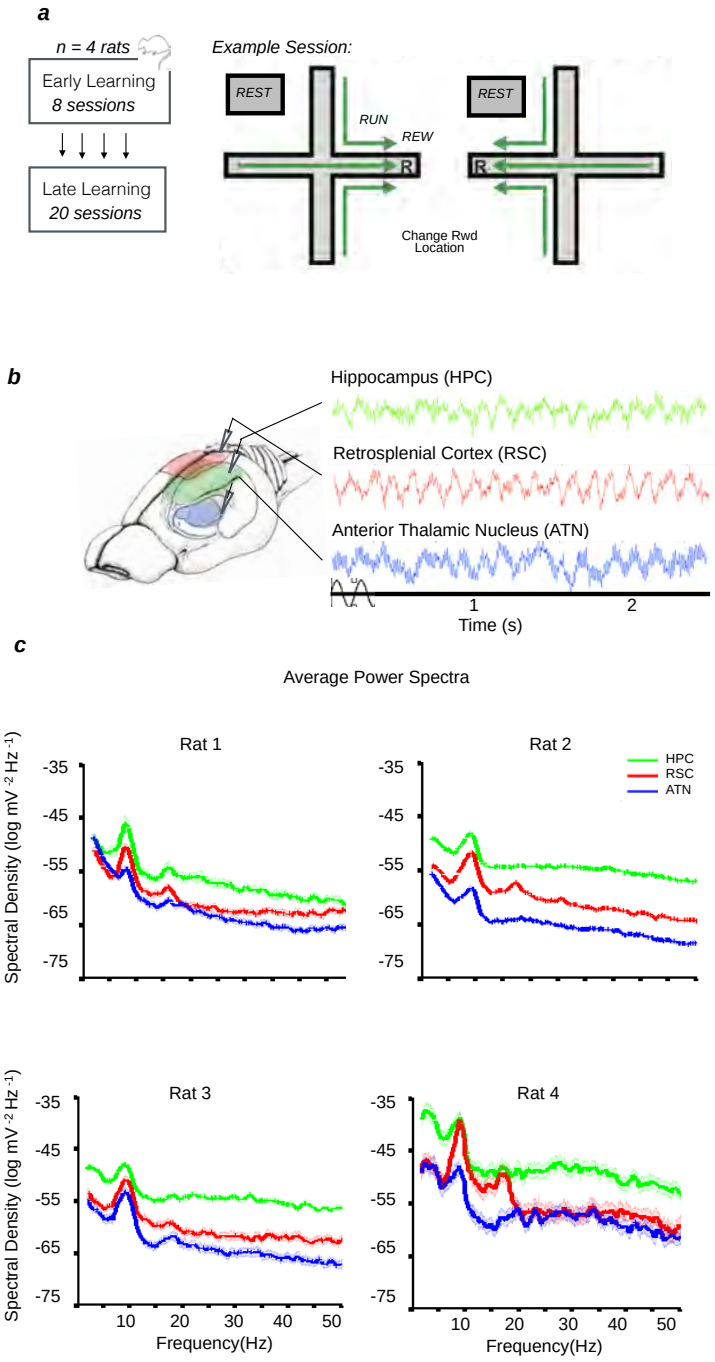

FIGURE 1 | **Overview of task and broad spectral features of LFP in each rat**
a Four rats with high quality simultaneously recordings in the HPC, RSC, and ATN comprise the data used for this study. These recordings were obtained during distinct versions of the plus maze task (pre learning and late learning) ***b*** The plus maze task employed in (Smith 2011); animals employ a win-stay strategy over the course of learning. ***c*** Average spectral density plots across all recordings sessions and all tetrodes in each of 4 rats. Plots show prominent a theta peak (~ 4 - 12 Hz).

receipt of reward at time 0 (indicated by a black line), and is further subdivided into Run, Rew, and Rest epochs of the behavioral task. A black shaded box indicates times of behavioral variability (this is when the rat is being picked up and placed on the inter-trial waiting platform prior to the onset of the Rest epoch). Underneath each spectrogram, the progression of theta (red) and beta (green) power (normalized between 0 and 1), can be seen for ease of visual comparison, and at the bottom of each column, it can be seen that the average running speed across animals is greatest during run, drops to zero during the reward period, and remains low during the rest period. There are minimal differences in average run time between pre-learning (PL) and asymptote (AS).

A few trends are visually observable: 1) theta (6 - 12hz) power is strong in all regions throughout each epoch, differing little as a function of learning, 2) Surprisingly, beta activity (from 20 to 40 Hz) can be seen in each region during pretraining, although much more so in the HPC and ATN. Notably, this beta activity seems to be equally high in Rew and Rest periods during PT. 3) In well trained animals, beta activity appears to be more selective for the reward period in the HPC and ATN, but is no longer present in the RSC.

To quantify these changes in theta and beta power, six 2-way repeated measures ANOVAS were conducted total; 2 per brain region corresponding to theta and beta. Two factors and their interaction were assessed for theta/beta in the HPC/RSC/ATN: learning stage (pre and post) and training epoch within a session (RUN (1), REW (2), AND REST (3)). Because violations of sphericity were observed, a Huynd-Feldt correction was applied. Calculated epsilon values used in these corrections are included in the statistics below (eps). Any value less than 1 indicates a violation in sphericity, and epsilon can be thought of as a ‘correction factor’ to

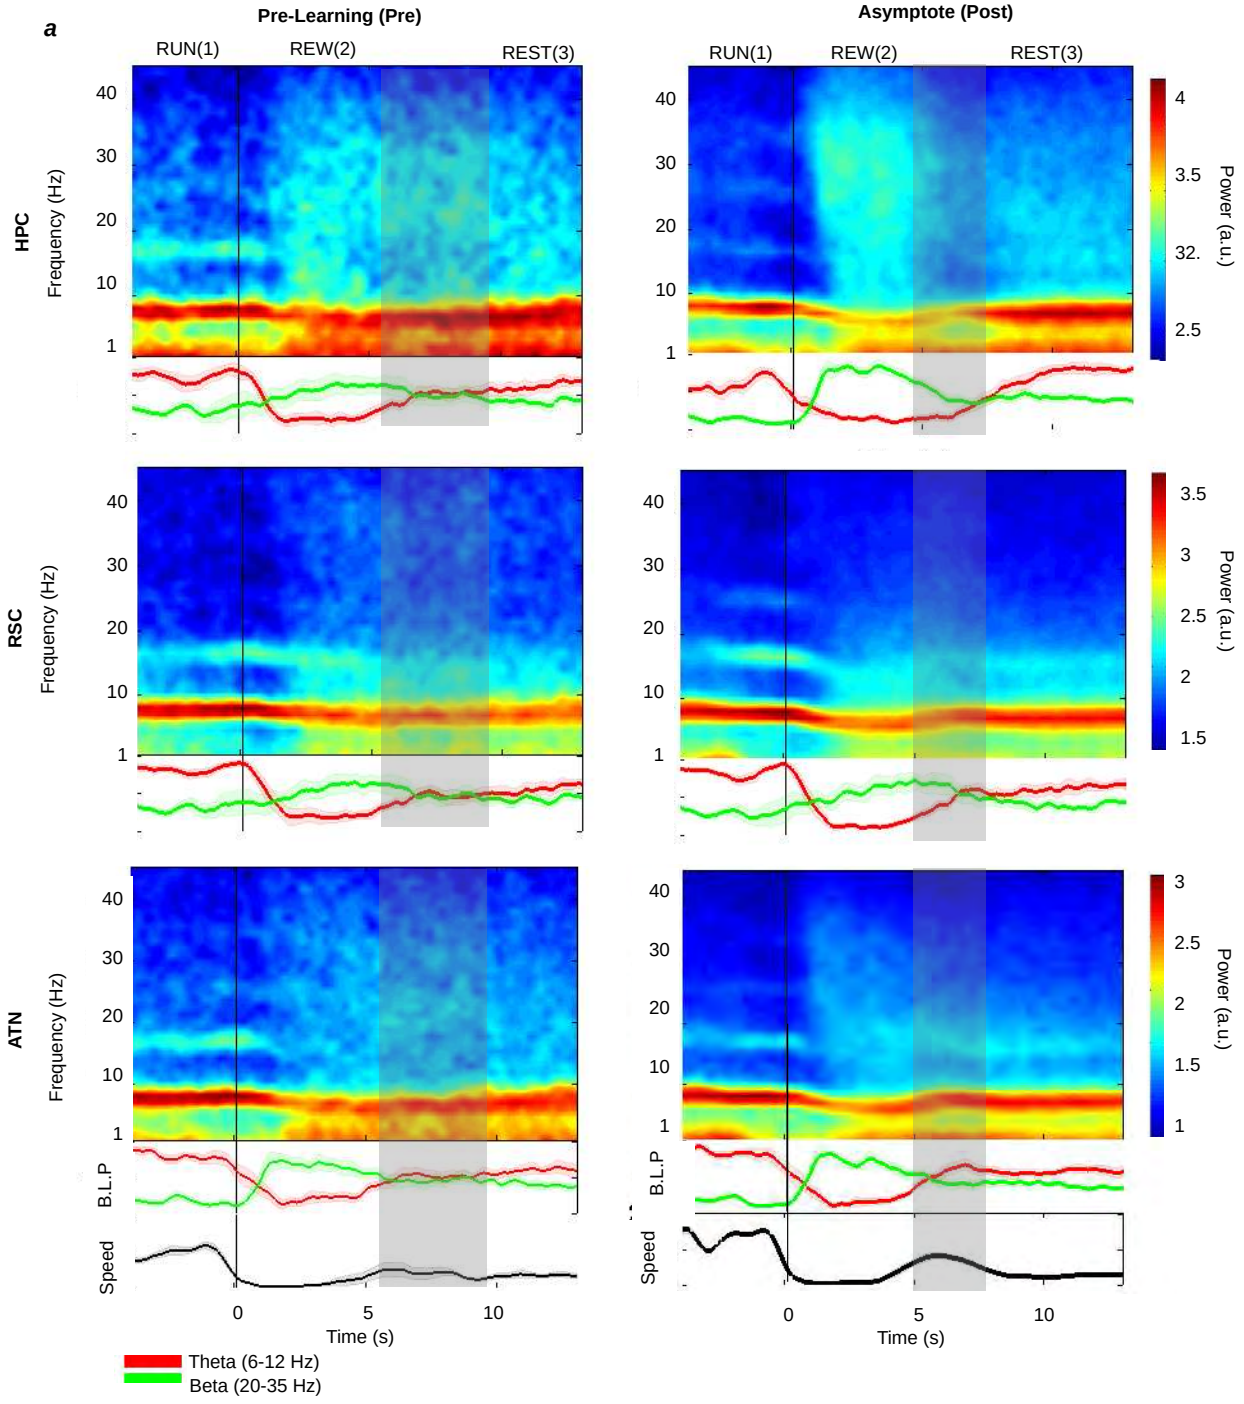

FIGURE 2 | **Beta and theta are prominent oscillatory features of the extended HPC circuit** Spectrogram plots above compare power of frequencies 1 through 40 in each region across stages of the behavioral task. Running speed is plotted below with shaded SEM (calculated across rats). Each region shows and persistent theta rhtyhm. Additionally the HPC and ATN show transient beta that is uniformly higher during pre-learning, and is more temporally localized to reward at asymptote.

account for these features of the data in order to avoid violating assumptions made by the ANOVA analysis method. Refer to methods for additional details. ***Figure 3c*** illustrates the results of these tests. In short, across all regions theta power did not change significantly as a function of learning stage, but did change as a function of epoch (***HPC theta***: F(2,6)=8.29, p=.033, eps =.7721; ***RSC theta:*** F(2,6)=13.02, p=.016, eps=.7324; ***ATN theta***: F(2,6)=1.96,p=.240,eps=.5993), showing a significant decrease in theta at reward receipt. There was no significant interaction effect between learning stage and epoch. Retrosplenial beta power changes showed no significant modulation by learning stage or epoch, however beta power changes in the HPC and ATN had significant main effects and interaction effects (***HPC beta***: learning F(1,3)=14.61, p=.032, epoch F(2,6)=22.92, p=.022, interaction F(2,6)=18.48,p=.003; ***ATN beta***: learning *ns,* epoch F(2,6) = 25.54, p=.002, eps-.8684, interaction F(2,6)=2.87, p=.174, eps=.6171). Most notably, during pretraining, an increase in beta was seen from RUN to REW epochs, but beta power remained high during the subsequent REST epoch. Furthermore, there was a significant *decline* in beta power from PL to ASYMP learning during the RUN period. This decline in beta did not remain, however, and beta powers rebounded to the same levels they were at in PL upon receipt of reward. However, beta power levels declined again during ASYMP REST. These significance trends underscore what can be visually identified in the spectrograms - beta power is high across epochs during PT, but increases only during reward receipt when animals are well-trained.

Periodic beta bursts and novelty

Is beta occurring for sustained intervals, as is the case with theta during active locomotion, or is it occurring in periodic bursts, as has been suggested by previous studies (Howe 2011)? ***Figure 3a*** illustrates that, along the lines of previous studies distinguishing this frequency from low gamma, beta is occurring in isolatable 'bursts' of activity, the duration of which is approximately log-normally distributed, with a peak duration of ~150ms. Notably, although the previous spectrograms were interpreted as a sustained increase in power, it is clear that this increase is instead driven by a higher probability of beta bursting activity in the HPC and ATN. ***Figure 3b*** highlights these transient bursts with spectrogram plots from individual sessions. Note the complete absence of beta activity during the trial period in the upper right spectrogram. The upper two spectrograms are recorded in the HPC during PL and ASYMP; the lower two are recorded in the ATN and RSC of the same ASYMP session.

A bigger picture: strong correlation between frequency of dominant network oscillations and running speed

Although this set of findings is corroborated by recent empiral studies exploring the role of beta in reward anticipation and contextual novelty, an alternative explanation that must be explored is the influence of running speed on changes in power in various frequencies. It has recently been documented that a medial septal glutamatergic circuit directly controls the frequency of theta oscillations in the hippocampus via an increase in firing rate of medial septal VGluT2+ cells, resulting in a nearly linear increase in the peak frequency of theta as a function of increasing

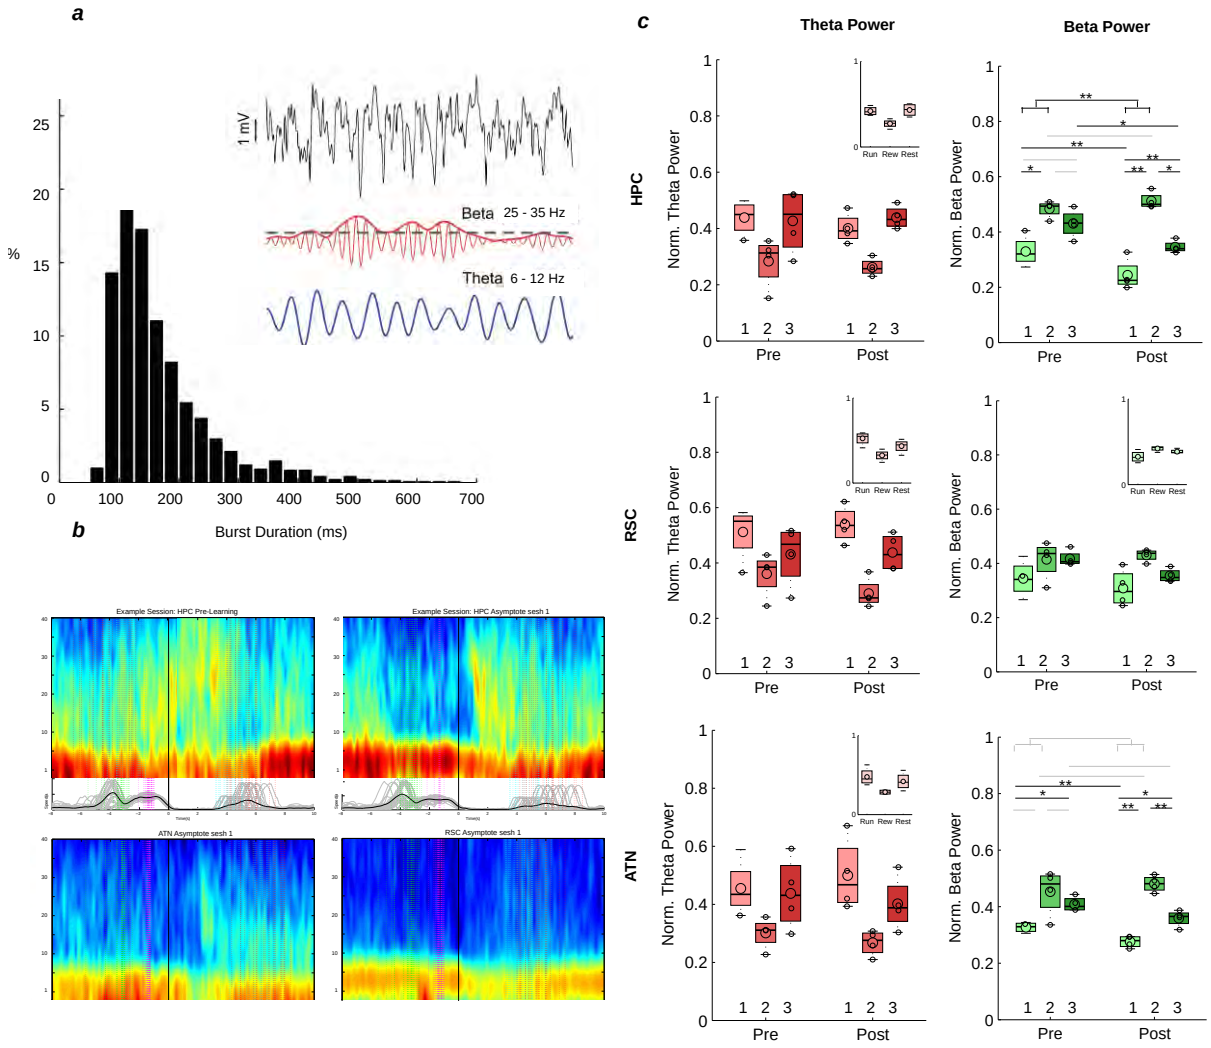

FIGURE 3| **Beta bursts lead to significant increases in beta power in the HPC and ATN during select epochs**

a Distribution of duration of transient increases in beta power in ms; identified as periods > 2SD above mean beta for a session, separated by at least one cycle. Duration of beta bursts is approximately log-normally distributed. Inset shows raw LFP signal filtered in the beta and theta frequency ranges. ***b*** Above left: example session; beta bursts occur during reward, run, and iti periods during pre-learning. Above right: strong selective decrease in beta power during trial, but rebounding of beta burst upon receipt of reward. Lower left and right: ATN and RSC spectrograms respectively for the above mentioned asymptote session. **c.** Box plots showing results of six 2-way repeated measures ANOVAS assessing the significance of either theta or beta power changes as a function of 2 factors: Learning stage (pre, post) and epoch (run (1), reward (2), and rest (3). This was performed independently for each brain region. Results reported in text; in sum, if inset is present, significant effect of epoch only was found. Upon the discovery of a significant interaction in ATN and HPC between learning stage and epoch, post-hoc tests via paired t-tests of a-priori interest were performed (* p<.05; ** p<.001).

running speed (Fuhrman, 2015). It is well known that the medial septum projects to nearly all brain regions in the extended hippocampal circuit, each of which exhibit strong theta rhythmicity. Furthermore, hippocampal gamma frequency has been shown to increase as a function of running speed, in tandem with theta (Ahmed 2012). In the present study, running speed is confounded with epoch, as is illustrated by ***figure 4d***, which shows minimal overlap between the distribution of running speeds across all sessions (including both learning stages) for each of the three epochs. Therefore any change in beta or theta power as a function of epoch could just as reasonably be attributed to changes in the running speed of the animal. Worth noting, however, is the fact that beta power is significantly higher during the REW period, even for speed-matched REW and REST epochs (beta power was pulled only during running speeds occurring during overlap in the distributions). F**igure 4a** and **b** illustrate a clear and opposite correlation between theta power (top), beta power (bottom) and speed, with theta power increasing as a function of running speed, and beta power decreasing. Note, in figure 4a and b scatterplots on the right, on the top of the plot is the distribution of running speeds (i.e. the distribution of our x-axis values). Lower speeds more highly sampled, perhaps leading to variability seen in plot. To the right of the plot is the distribution of power at each frequency band along the y axis. Because power is normalized between 0 and 1, means are approximately equal. The color of each histogram corresponds to the colors in the figure, green: HPC, red: RSC, blue: ATN. This allows us to see, for example, that the distribution of band-limited power in the theta (top plot) and beta (bottom plot) frequencies are more widely distributed in the HPC than in the RSC.

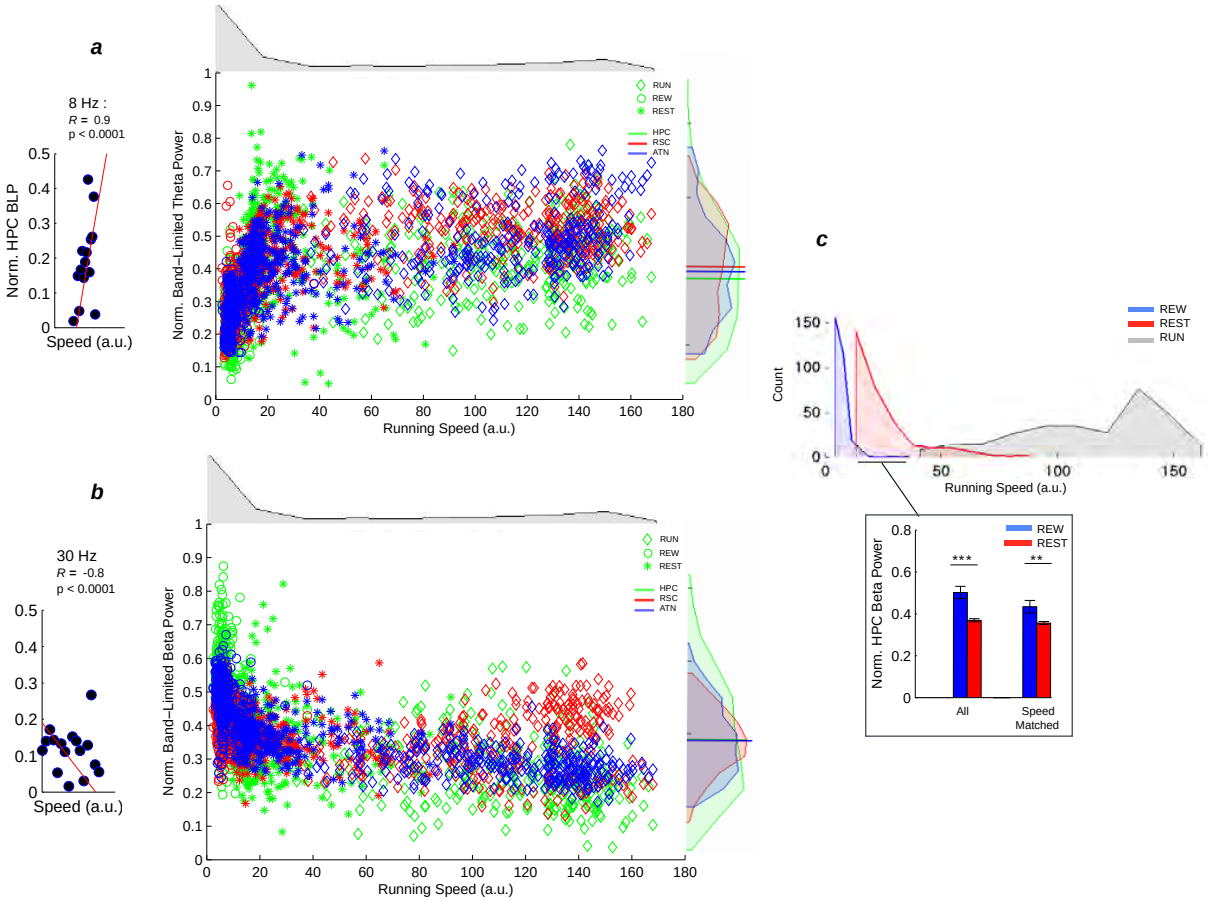

FIGURE 4 | **Beta and theta are strongly correlated with changes in running speed**
a Pearson's correlation and linear regression line fitted against theta power from an example session. Theta power increases significantly as a function of running speed. Right: scatterplot of theta power (normalized between 0 and 1) across all rats and all learning stages, plotted against increasing running speed. Theta appears to increase with running speed initially and then becomes highly variable. This is seen in all regions (blue = ATN; red = RSC; green = HPC). Symbol shape signal epoch. Top of plot is distribution of running speeds. Lower speeds more highly sampled, perhaps leading to variability seen in plot. Right of plot is distribution of power. Because power is normalized between 0 and 1, means are approximately equal. ***b*** Beta shows reverse trend, decreasing in power with increasing running speed. ***c*** Distribution of running speed times in each epoch. Distributions minimally overlap. Inset: beta power is significantly higher during reward than rest period, even if only speed matched trials are used ($p<.001$, $p<.01$, t-test matched for number of samples.

How can we systematically explore the relationship between running speed and power across all frequency bands? This is done in ***figure 5***. The findings in this figure have been documented in the HPC, but have not yet been documented in the RSC or ATN. In all three brain regions, spectral activity ranging from 2 to 200 Hz demonstrates a highly consistent relationship with running speed, and this relationship is conserved across two unique spatial tasks (PL and ASYMP). This figure has a number of components, which I will elaborate on in-text for the purposes of discussion. For each frequency band, a pearson's correlation coefficient was computed (as in the previous figure), meaning a normalized measure of the linear covariance between running speed and power in the specified frequency band was estimated. If power in that frequency increases as a function of running speed, as is the case with theta, then the coefficient will be positive. The coefficient will be negative in the reverse scenario, as is the case with beta. The length of the bar plots on the left of each spectrogram represent the magnitude of the correlation coefficient, and the color (red or black) indicates the significance of this coefficient (significant, $p<.05$, or n.s.). This will not capture many of the possible relationships that could exist between running speed and frequency modulation, but will capture a strong trend visible in these graphs - the incremental increase or decrease of peak frequency across running speeds. Using this method, there are three to four ranges of frequencies that can be distinguished by their relationship (+ or -, signif or not) with running speed. The boundaries of these regions correspond well with the boundaries of the dominant spectral components of the hippocampal region, including delta (weakly negatively correlated with increasing running speed), theta (strongly positively correlated), beta/slow gamma (strongly negatively correlated), and high gamma (positively correlated). Various mechanisms by which theta and gamma oscillations can

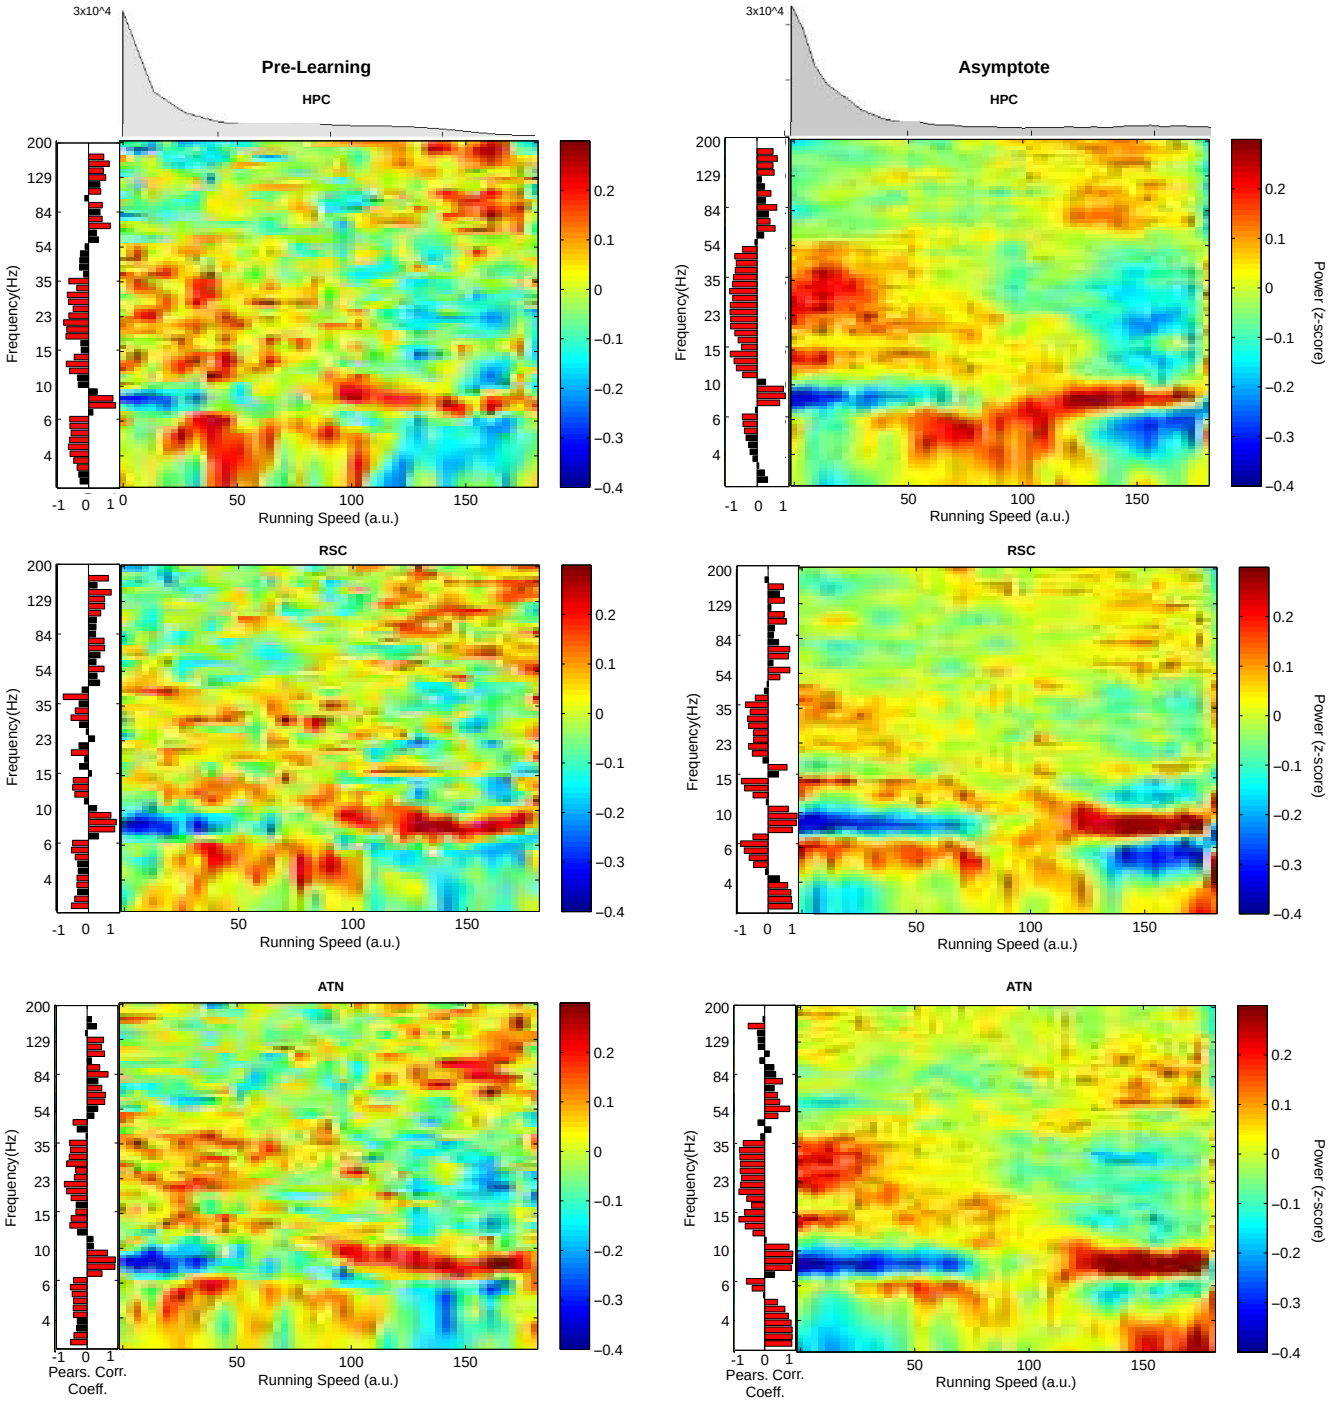

FIGURE 5 | **The major oscillatory components of the HPC, ATN, and RSC (theta, beta/low gamma, high gamma) are strongly correlated with running speed**

In both Pre-Learning (left) and Asymptote performance (right), increases in HPC, RSC, and ATN theta and high gamma are significantly correlated with increases in running speed (designed by bar plot on left of each figure (red: p<.05, black: n.s.; person's r). The reverse is true for beta, which is strongest during low running speeds and significantly decreases with increased running speed. This relationship hold true in two spatially similar but cognitively distinct tasks, suggesting the beta burst observed is driven by decreases in running speed of the animal. Top of figure: dist. of running speeds. Power (zscored).

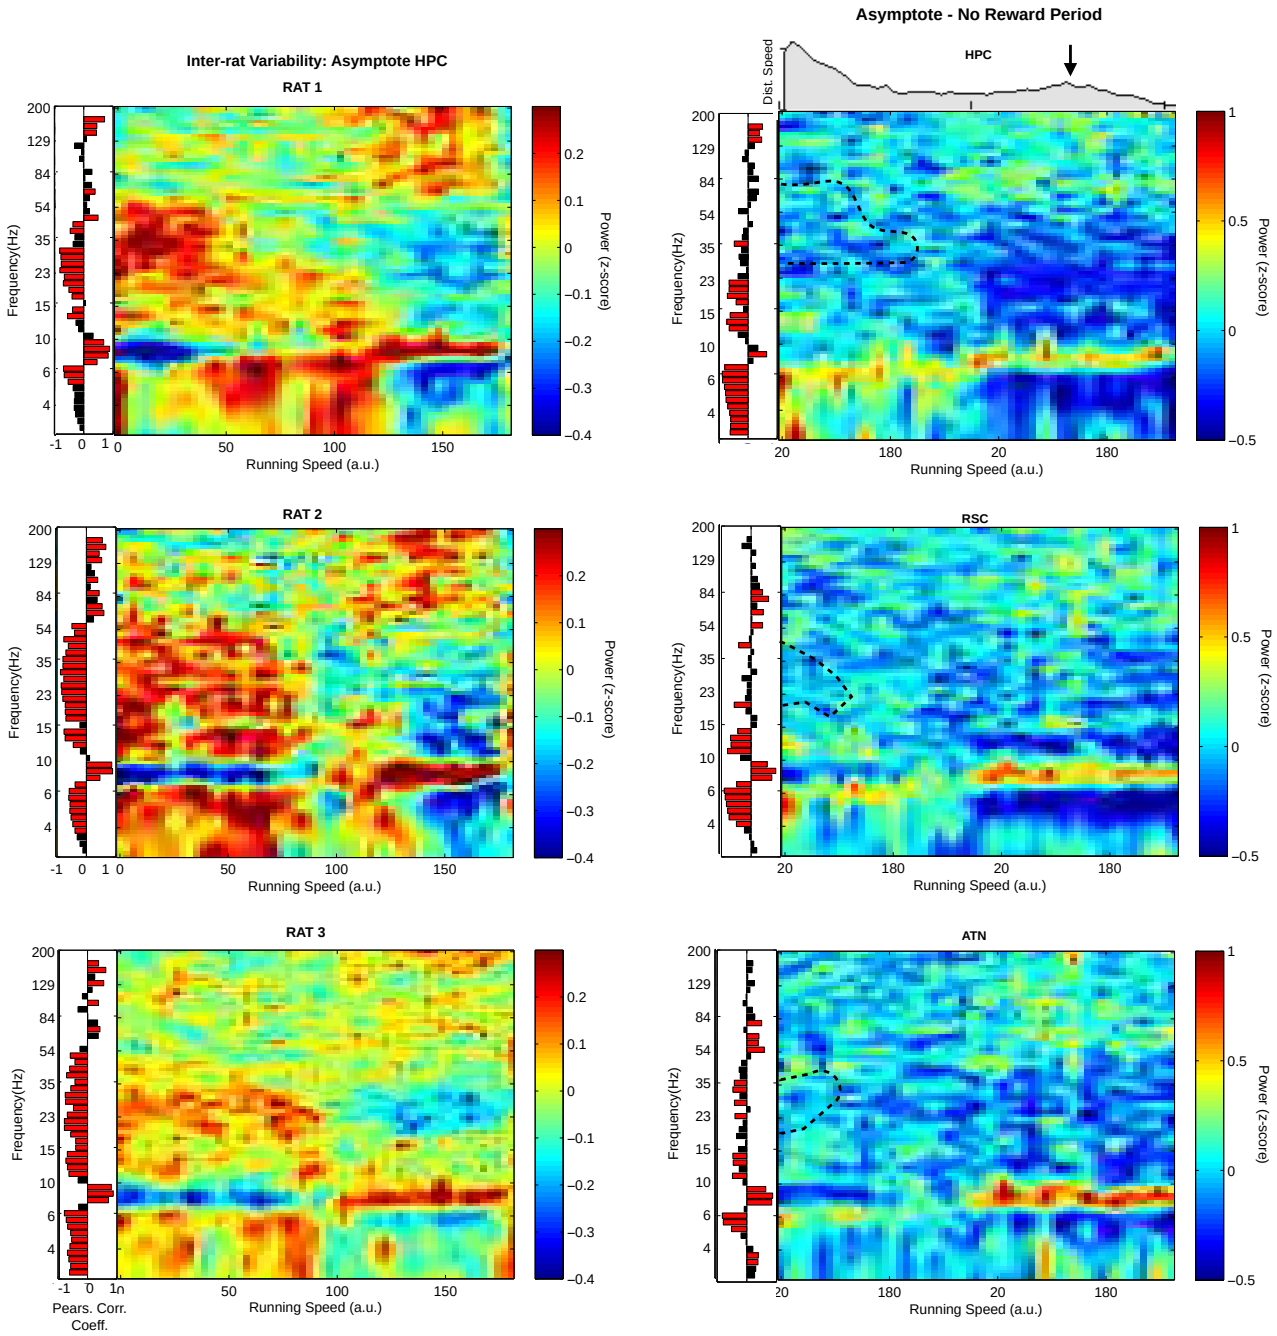

FIGURE 6 | **Correlation between running speed and power seen in all rats. Removal of reward beta may cause decrease in observed low-speed beta power.**
a Intimate relationship between running speed and distinct frequency bands conserved across rats, although specific features, such as boundaries of co-varying correlaten coefficients vary from rat to rat.
b Removing reward period results in a visible decrease in presence of low-running speed beta power, even during speeds shared by respective running speed distributions of each epoch. Black dotted line outlines location of high beta power seen in previous plot. Only viewing right end of this increase in power because be removing reward periods, removed almost all lowest running speed events. Arrow indicates notable peak in higher speeds

be coordinated as a function of the local circuitry of a region have been proposed, which suggests that the presence of the same family of oscillations seen in the HPC as in the RSC and ATN, where theta frequency dynamics were already known to be present, is perhaps not surprising (Kopell 2010). These results have strong implications for a number of studies investigating the cognitive correlates of various oscillations. In many studies, when an animal receives a reward - it slows down. These data suggest that a peak in beta activity, could be related to novelty or reward processing, the local dynamics of the region being recorded from which are controlled by animal movement, or both.

Although we cannot dissociate these effects in our present data, we can look to see if removing the reward period entirely results in the subsequent loss of low-running-speed beta activity. For running speed 1-20, we have too few samples to reliably plot cross-frequency power, however for running speed of 20 and above, which partially overlap with our REW distribution, we can make similar plots. These can be seen in **figure 6b.** There are many caveats associated with these plots, including the large loss in number of speed-coherence samples due to subtraction of this period and the minimal overlap between the REW and REST distributions, however these plots suggest that the removal of the REW period results in a loss of higher running speed beta power that was present in the previous plot (black dotted outlines). In order to ensure these cross-frequency running speed relationships are a general feature of these brain regions and are not driven by a particular rat, averages for 3 of 4 rats (restricted by size on page) were plotted in **figure 6a.** Interestingly, the boundaries of frequency bands with co-varying correlation coefficients are narrow for some rats (beta in rat 1) and wider for other rats (beta in rat 2). It is likely that the boundaries of distinct states of a network, measured as various

covarying properties of a set of frequencies, change not only from state to state, but also brain region to brain region and animal to animal. Despite this, hoppocampal theta is remarkably consistent in frequency range and modulation by running speed in all three regions and across animals.

Coherent theta and beta oscillations suggest strong functional connectivity between all three regions in the theta frequency range, and preferential communication between the HPC and ATN in the beta range

Thus far we have examined the shared and unique spectral properties of the HPC, RSC, and ATN in relation to task demands and running speed. However, we have not yet examined whether these oscillations are occurring in phase, i.e. if they are coherent. As discussed in the introduction of this chapter, transient synchronization of oscillations has been shown to coordinate neuronal ensembles in distributed brain regions, facilitating mutually-influential co-activation (Fries 2007). In order to investigate the functional connectivity of the HPC-ATN-RSC inter-regional circuit during reward-directed learning, we calculated two measures of oscillatory synchrony; coherence, estimated by normalizing the cross correlation of two signals by their respective power spectra (Mitra, **chronux.org**, *see methods)*, and phase-locking values (PLV) (Lachaux, 1999), a measure of the consistency of the phase relationship between two signals, not weighted by power. We used both methods, because although coherence is ubiquitously used to describe synchronous oscillations in continuous or point process data, it has been shown that by normalizing the cross spectra of two signals by their respective power, it is possible to have a high level of coherence even with a low level of phase-syncronization (Srinath 2014).

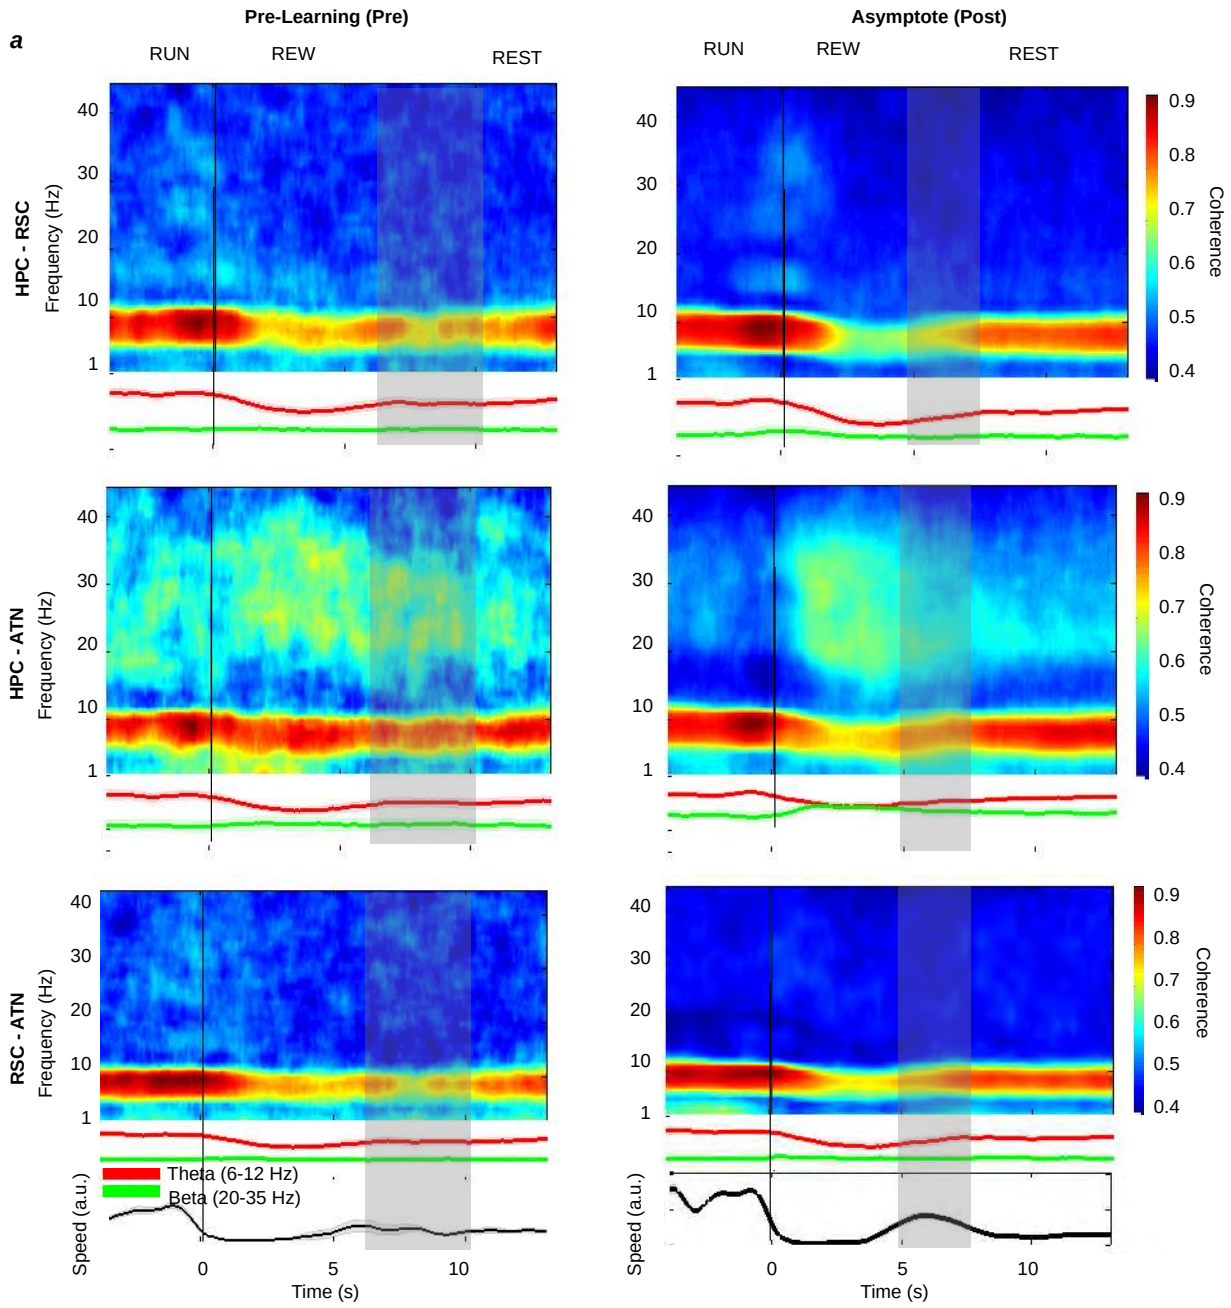

FIGURE 7| **Theta coherence coordinates activity of all three regions; beta coherence specific to the HPC and ATN**

Coherence plots are centered around reward receipt (t=0) in order to facilitate detection of rewarded-related oscillatory activity. Black bars indicate time during which rat is being picked up. Average coherence plotted for each pair of brain regions, separated by learning stage. All three regions exhibit strong theta coherence peaking at reward receipt and immediately decreasing with decreased running speed. Additionally, the beta bursts seen in HPC and ATN spectrograms are coherent during all epochs in pre training and only reward receipt in post training. Given timing of beta bursts discussed previously, this suggests that if beta is present in both regions, it will be coherent.

Because the mechanism we are interested in is facilitative coordination of activity via time windows of joint excitability, i.e. phase synchronization between two brain regions, PLV (also referred to more aptly as Inter Site Phase Clustering (ISPC)) is complementary and equally appropriate.

As can be seen in ***figure 7,*** the predominant spectral features of each region (namely theta and beta/low gamma) are strongly coherent. Theta-range coherence is seen between all three pairs of regions in both PL and ASYMP task stages, however beta coherence is only seen between the HPC and ATN. Following a pattern similar to that seen during changes in beta power in both the HPC and ATN, pre-learning beta coherence is present in all stages of the task, however at asymptote, it decreases during the run period, peaks at the reward, and remains high during rest.

Because the coherence values so closely mirror coordinated increases and decreases in power across regions, it was possible that we would not obtain similarly high results with PLV. This was not the case, however, as seen in ***figure 8***, leading to the conclusion that has been landed upon in studies examining other regions in this circuit - theta synchronization coordinates long-distance information exchange. The selective beta synchronization between the the HPC and ATN was unexpected, however, and has not been previously illustrated. The only known literature we are aware of regarding hippocampal beta synchronization is the earlier discussed recent Moser paper detailing LEC and HPC beta coherence in response to cues predicting reward. Another set of 2-way repeated measures anovas (6 performed between coherence values and 6 performed between ISPC values) resulted in the following: sustained high theta coherence across all periods with no significant changes by learning stage or epoch, with the exception of a

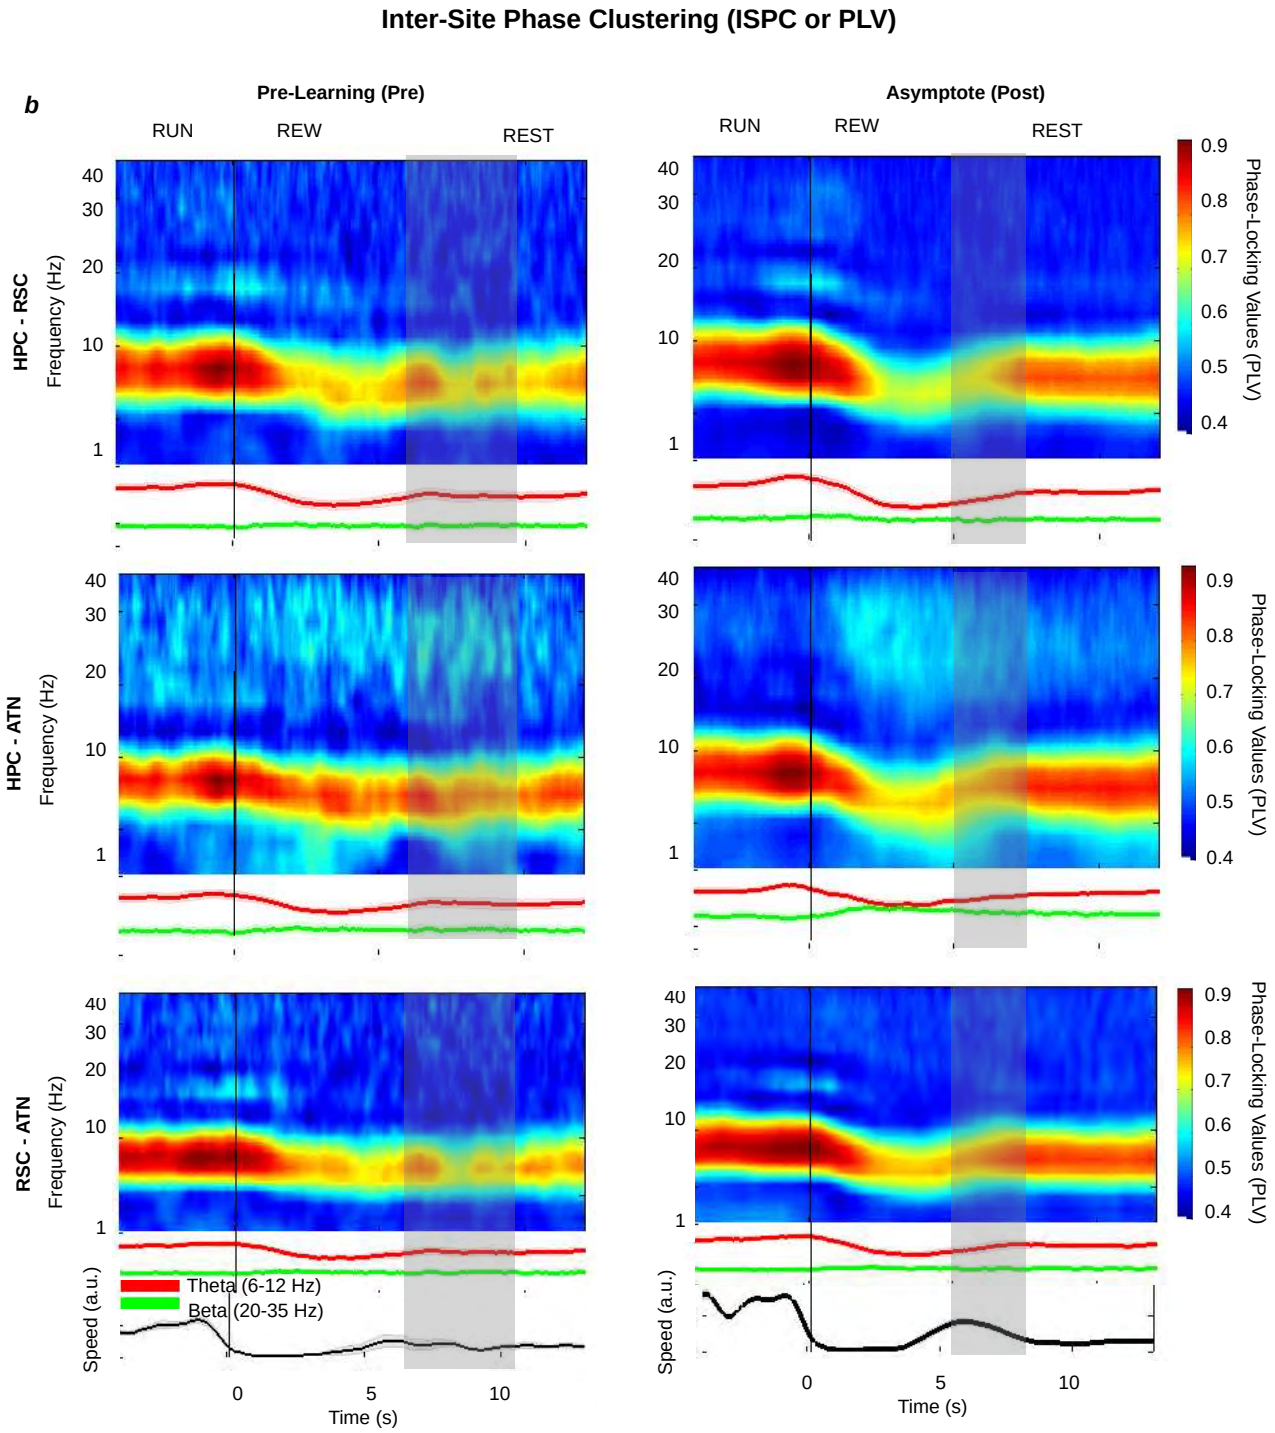

FIGURE 8 | **ISPC values, a measure of phase locking independent of signal amplitude, show same theta (all pairs of regions) and beta (exclusive to HPC-ATN) phase-coupling.**
ISPC or phase locking values (PLV) measure the consistency in phase differences between two regions over time. Densely clustered phase-lags lead to a high PLV value, widely varying phase lags lead to a low PLV value. PLV results mirror coherence results.

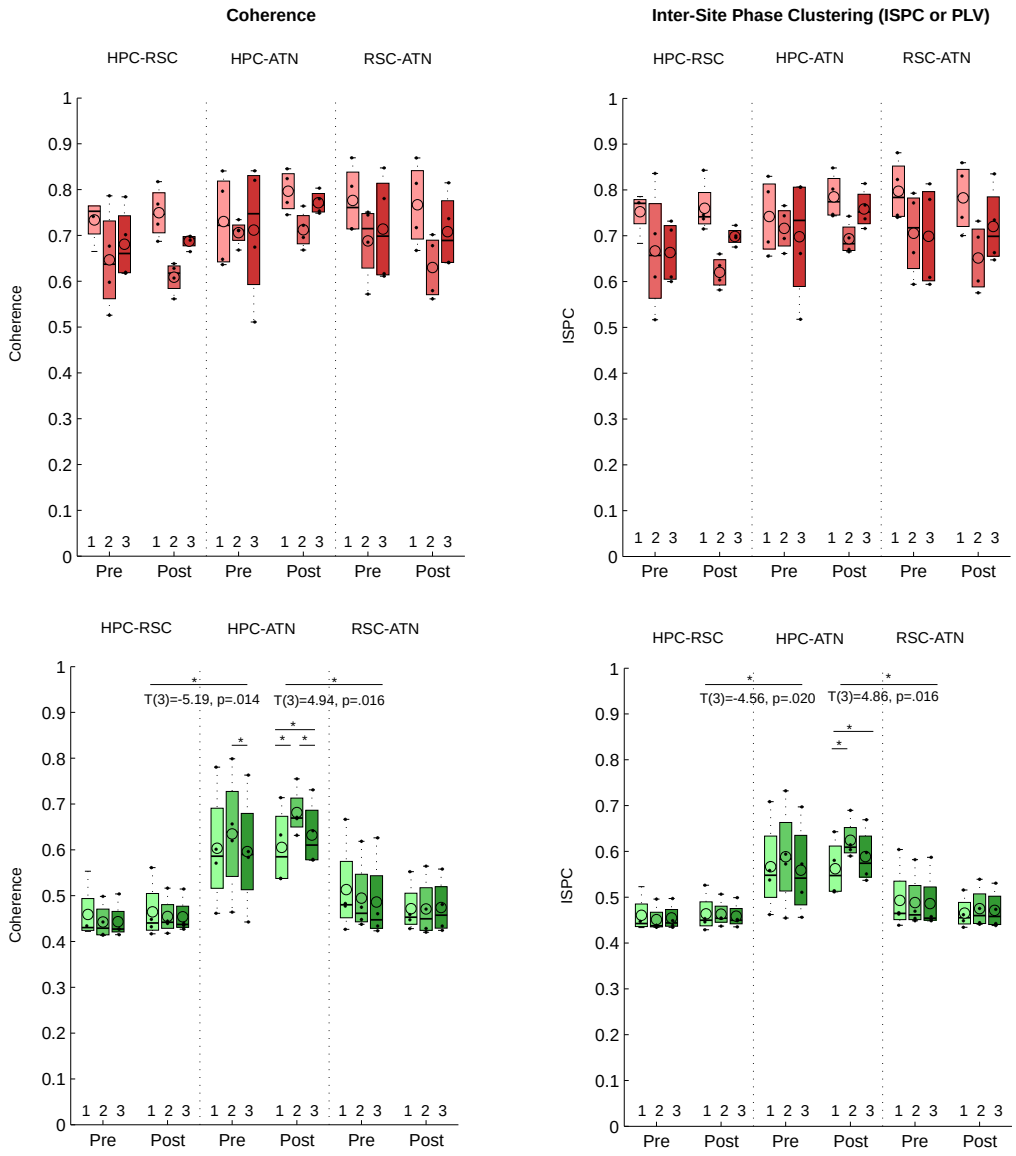

FIGURE 9 | **HPC-ATN beta coherence and ispc values are significantly increased during reward** High sustained theta coherence and ispc values to not show a significant differences in group means as a function of learning or epoch. Decrease in theta during reward period is seen, but not significant. This is consistent with result finding low theta power during low running speeds. Average beta coherence and ispc in the ATN-HPC is significantly higher than HPC-RSC beta coherence and RSC-ATN coherene. Results also indicate a significant interaction effect and main effect of epoch. HPC-ATN show significant increase in beta coherence upon reward receipt in asymptote but not pre-learning trials, and coherence shows a significant decrease in coherence. Overall, theta coherence is dominant in all regions and beta coherence is unique to the HPC-ATN circuit and increases upon receipt of reward (see discussion for statistics), * p<.05, **p<.001

significant main effect of epoch found via both ISPC and Coherence between the RSC and ATN (Coherence RSC-ATN theta: $F(2,6)=12.55$, $p = .012$, eps = .85, PLV (or ISPC) theta: $F(2,6)=11.22$, $p = .009$). Beta coherence between the HPC and ATN, on the other hand, was significantly higher averaged across epoch and learning stage than HPC-RSC and RSC-ATN synchronization, in both Coherence scores and PLV's, results ***figure 9)***. Furthermore, a significant interaction effect was found between learning stage and epoch between the HPC and ATN in beta in both coherence and ISPC (PLV HPC-ATN interaction: $F(2,6) = 8.12$, $p = .0623$; Coherence HPC-ATN interaction: $F(2,6) = 7.26$, $p = .042$, eps = .7587). Upon closer inspection, in both PL and ASYMP training, beta coherence is significantly higher during reward receipt than during the RUN period.

Coherence as a function of running speed

Using PLV values, we employed the same strategy as discussed previously in order to look at the variation of coherence between regions as a function of running speed, ***figure 9.*** Due to the known correlation between coherence and regional amplitude, we used PLV values for these calculations. In all regions, the frequency of peak theta coherence increased by a small degree as a function running speed, with higher theta coherence at higher running speeds. Beta coherence, on the other hand, did not show the same linear trend seen in power. During pre-training, beta coherence between HPC and ATN appears to peak at around 100a.u. This is very close to the average running speed of the animal at the choice point of the plus maze. A week negative correlation is seen between running speed and beta synchrony in the HPC - ATN during ASYMP trials. It is possible that with a greater range of running speeds, stronger speed-dependence of

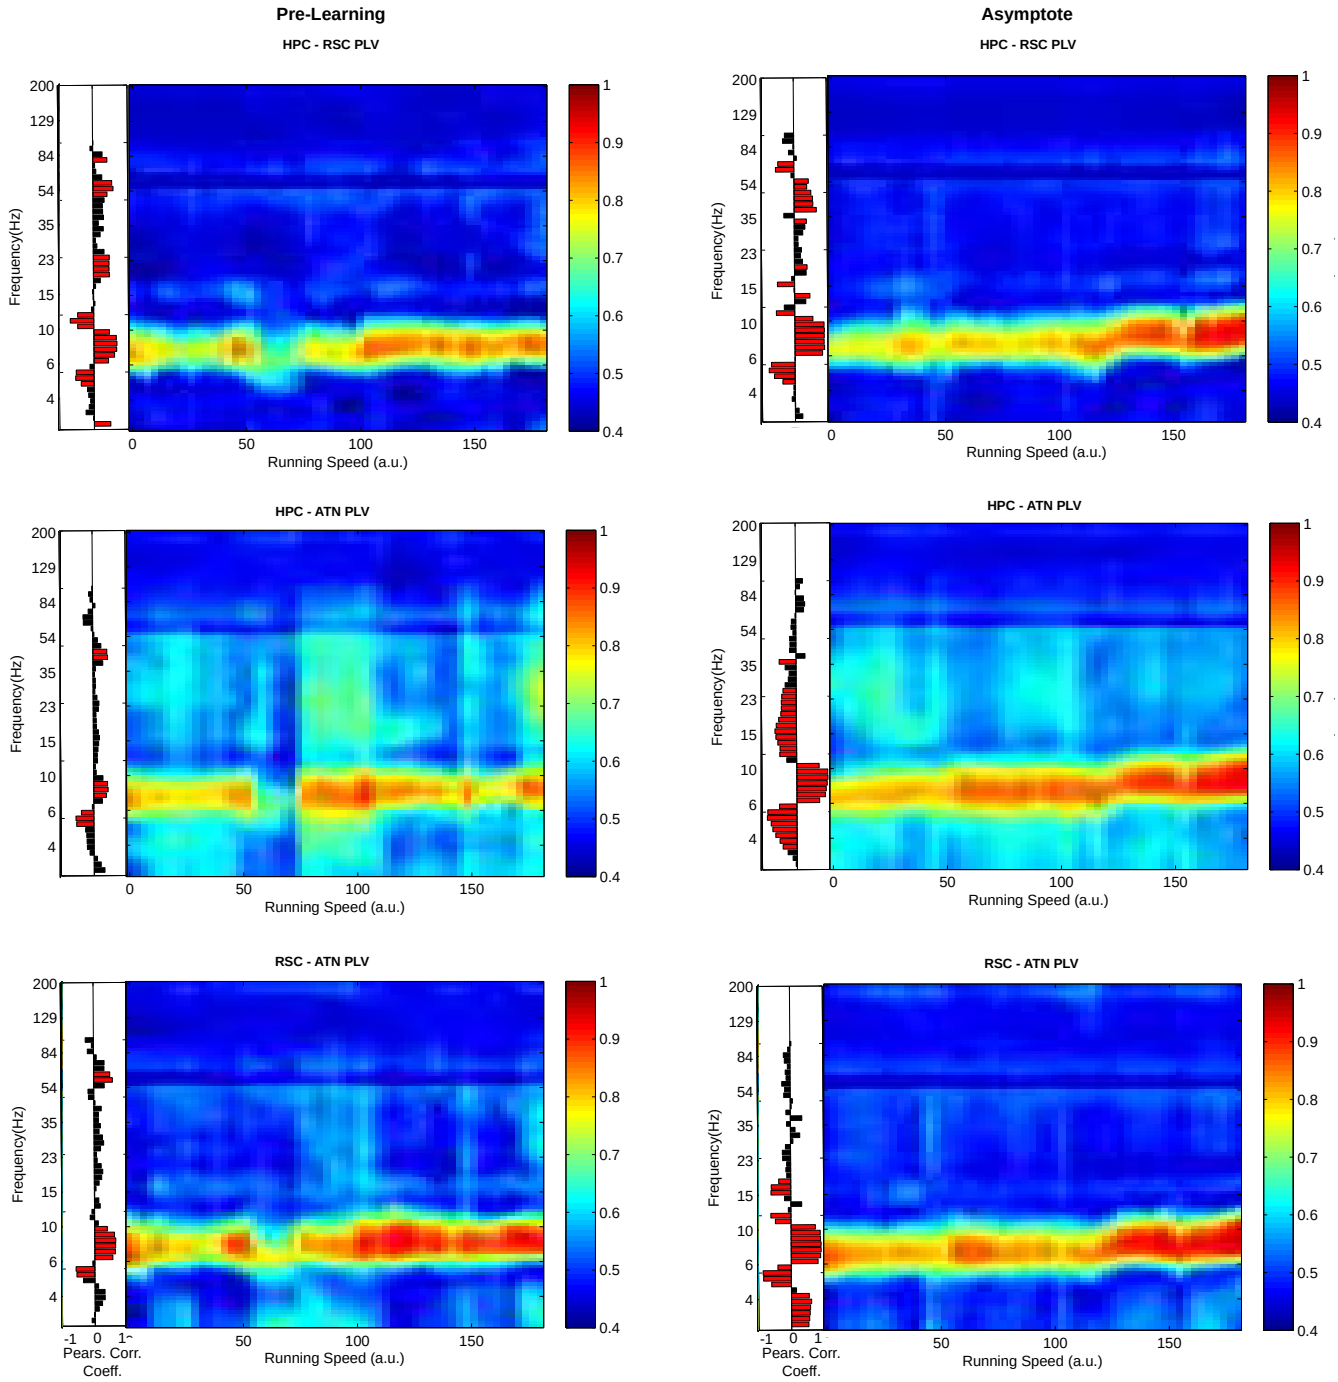

FIGURE 10 | **Theta coherence but not beta coherence show reliable correlations with running speed.**

Theta coherence increases in intensity and frequency with increased running speed. Beta coherence is not linearly modulated by running speed and seems to peak in PT at approximately 100 a.u. and around 100 and also lower running speeds at asymptote.

correlation scores would emerge.

Coherence as a function of location on the maze

Given the lack of linear speed dependency of the beta coherence signal, we thought about yet a third way to cut the data that would allow us to look at changes in oscillatory coherence between structures as a function of location on the maze. The ATN, RSC, and HPC are all directly involved in navigation, thus although observing our data as a function of epoch serves to segregate reward form maze run times, it is possible that averaging across trials obfuscates the spatial complexity of the task, which had been previously highlighted as being reliably encoded by hippocampal and retrosplenial cells in the form of place cells, cells with conjunctive response properties, and, especially in the retrosplenial cortex, cue responsive cells (Smith 2012). **Figure 11** demonstrates a clear difference between theta and beta coherence, and between pre learning and asymptote levels and *locations* of coherence. The upper half of each quadrant represents z-scored HPC-RSC, HPC-ATN, and RSC-ATN theta and beta coherence values during both pre learning (upper half of quadrant) and during asymptote performance (lower half of quadrant). During PT in all brain region pairs, theta and beta coherence are uniformly high across the entire plus-maze. Remarkably, when this task is well learned, relative theta coherence in all three regions is only seen at reward locations and on the stems of the maze leading to reward. In our average spectrograms, this results in a constant blur of hippocampal theta coherence, implying chronic connectivity, however it is clear from these findings that theta coherence becomes remarkably spatially selective. Note - the supposed increase in theta coherence at the site of reward directly contradicts our decreased theta at low running speeds finding. However, given

the method used to make these plots and the variability of rat location, it is likely that this increase is prominent because the rat is spending most of his time at the reward site during a trial (~half), and this site is a far more consistent location than the high-speed trajectories traversed on the maze. A marked spatial change in beta coherence also occurs that differs from the change in theta coherence. Beta coherence is initially high across all regions of the maze, but then substantially decreases between the HPC-RSC and RSC-ATN, and shows a similar decrease in the HPC-ATN, BUT an increase at the E and W reward locations, and potentially as the rat approaches the choice point of the maze (see ***figure 11d***). Since each of these arms was visited the same number of times (brain region were recorded from simultaneously), a difference in number of visits cannot account for these differences. Because running speed is comparable across pre-learning and asymptote sessions, taken together these findings suggest that the beta increase observed during reward and decrease in beta during trial running is highly likely to be related to processing of the reward cue.

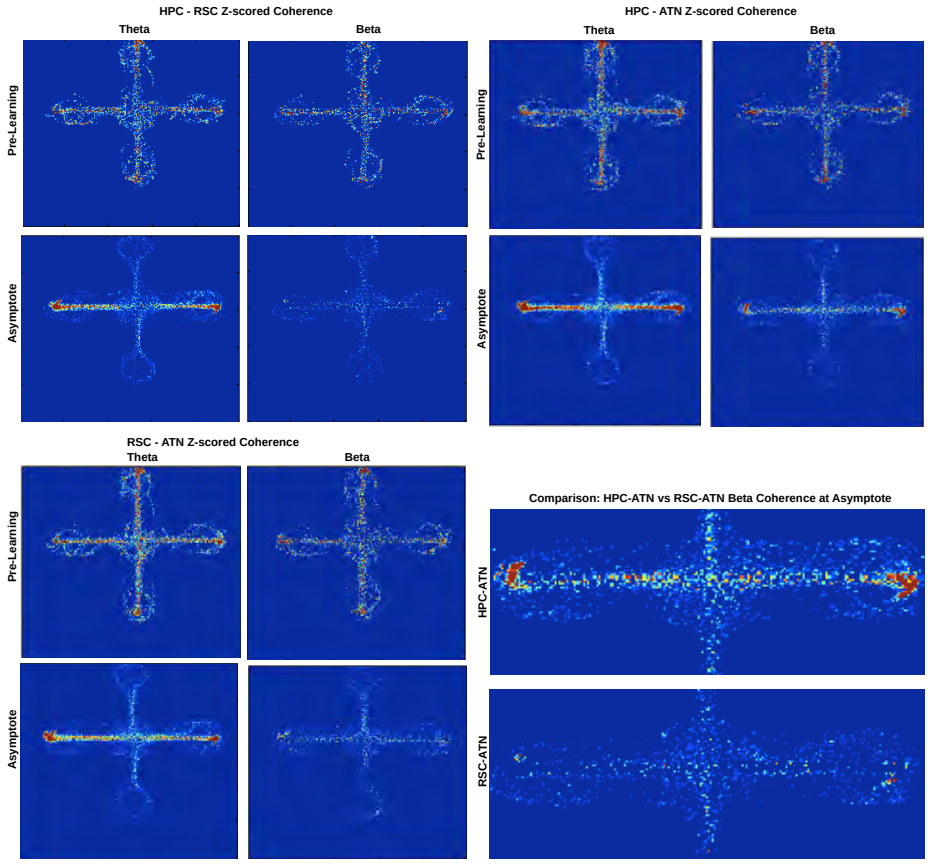

FIGURE 11 | **Theta and beta z-scored coherence is modulated by learning stage and location.**
Plots represent zscored coherence average across sessions and rats separately for both theta and beta range coherence, by learning stage, and by region pair (H-R upper left, H-A upper right, R-A lower left). During initial exposure to the plus maze, both theta and beta coherence can be seen on all arms of the maze. Following learning, theta coherence occurs preferentially on the rewarded arms of the maze, showing a completely different pattern in spatial

Phase Delay between HPC, RSC, and ATN

Lastly, in computing phase locking values, it is also possible to compute a measure of the difference between oscillatory signals of the same frequency in different brain areas at every point in time. The ISPC (or PLV) is a measure of the variance of these values, and can be computed as an average (assuming the phase lag between regions in a particular frequency band is consistent), or as a sliding average, making only weak stationarity assumptions. In **figure 12**, the phase difference itself it plotted in three different ways in order to visually illustrate this method. **Figure 12a** consists of three rose plots, one showing the average phase lag between the HPC (green), RSC (red) , and ATN (blue). An average phase lag of the difference seen between the HPC and RSC, assuming a 9 Hz theta cycle, is approximately 31 ms. This means on average, retrosplenial theta lags behind hippocampal theta by 31 ms. This is very similar to the time lag found between mPFC and HPC theta ~28ms, and is supported by the time constant of monosynaptic conduction delays. ATN theta is much closer to hippocampal theta, with an approximate average delay of only 6ms. This is much shorter than would be expected if the HPC was the soul driver of ATN theta, and suggests theta being recorded in the ATN is either the result of volume conduction from the HPC (less likely due to the small delay; volume conduction is instantaneous), or due to the HPC and ATN being driven by the same theta pacemaker. As mentioned previously, the medial septal nucleus has been shown to drive HPC theta and cause many of the features documented in this thesis and elsewhere regarding changes in the power and peak frequency of theta as a function of running speed. It is possible, then, that ATN theta is also being driven by this region. The 2nd rose plot in this figure demonstrates that on average

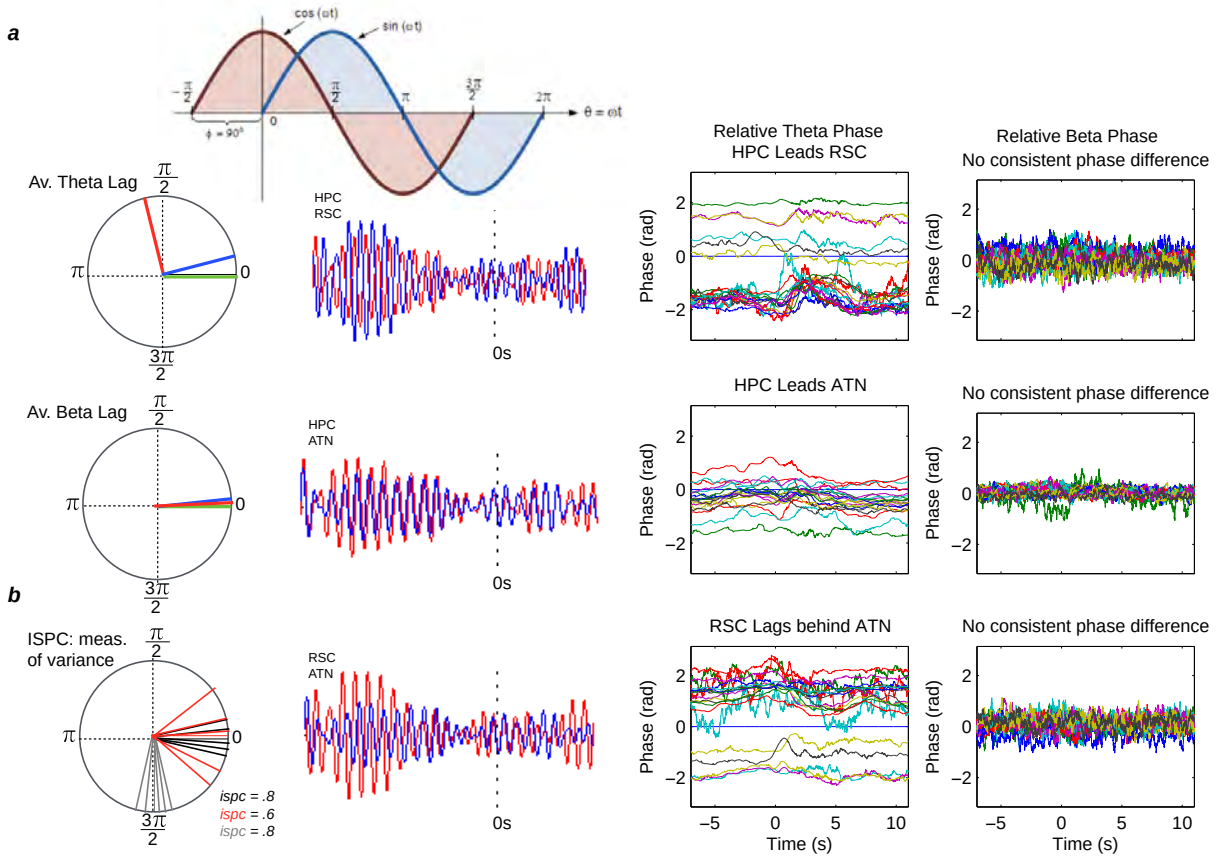

FIGURE 12 | **Theta and beta average phase delays**
a Average theta (top) and beta (bottom) phase delays averaged across all sessions between HPC (green), RSC (red), and ATN (blue).***c*** Example theta filetred traces from the labeled pairs of brain regoins, demonstrating visually what a phase delay looks like. ***b*** How can a zero phase-lag lead to high PLV/ISPC values? ISPC is a measure of the variance in phase differences between two signals. It does not weight values based on proximity to 0 or 180 degrees, and does not differentiate between positive and negative lags. ***c*** Average phase lag between regions across time for different trials is very consistent.

beta conduction delays are approximately 0. This means that the instantaneous beta coherence documented between the HPC and ATN that becomes highly reward selective is occurring at a very similar time. We cannot rule out the possibility of volume conduction from the hippocampus, although if this were the case we would also see zero phase lag in theta. Thus, we propose that another circuit, perhaps cholinergic projections from the basal ganglia known to optimally form cell assemblies in the beta frequency range and to be associated with reward and attention may serve to bias the HPC and ATN circuits towards resonation at beta frequencies (Tingley 2015) or perhaps directly pace this circuitry, facilitating communication between these regions. **Figure 12b** shows filtered theta traces overlaid from each region, demonstrating visually what different phase differences look like. **Figure 12c** (note these are radians so -2 and 2 are very similar is a plot of the phase lags found over the same time window used consistently throughout this thesis for each trial, illustrating the remarkable consistency of the theta phase lag between regions.

Spike-Phase Coherence

All previously discussed analyses have considered the temporal dynamics of local field potentials between regions, i.e. LFP-LFP questions. Another signal that was recorded, however, are action potentials of many single cells in each region. Thus, we can ask a parallel question as follows: does the phase at which action potentials are discharged from single cells show tuning for particular phases of our spectrally decomposed oscillatory frequencies? Are cells firing only at a particular phase of one frequency range, perhaps one that is physiologically relevant and prevalent in all of our regions, such as theta frequency ranges? Does this differ from region to region? Or from functionally-defined cell type to cell type?

The following analysis focused on a subset of these questions. For the purposes of this thesis, we have selected the theta frequency range to analyze spike-phase tuning, keeping in mind that this analysis could be expanded to the full frequency spectrum. Spike-phase questions are of broad interest, because if spike-phase coherence is found, it further supports an underlying assumption in this thesis, being that coherent oscillations organize the spiking of local ensembles of neurons in time in order to facilitate information exchange between regions. The first question we asked was 'is the spiking of individual units phase tuned to local theta?'. This is a reasonable question, given the predominance of theta in each region. Furthermore, previous studies have found a subset of theta phase-tuned cells in many regions that exhibit strong theta oscillations during active behavior (e.g. Wilson 2005). As can be seen in ***figure 13***, a subset of cells in each region were in fact found to be theta phase-tuned to local theta. Notably, although only a cell with no phase coherence is illustrated in this figure in the ATN, many theta phase-tuned cells were found. As can be inferred, because theta oscillations are coherent between regions, if a cell was phase-tuned to local theta, it was also found to be phase tuned, or 'coherent' to HPC theta. For this thesis, a measure called mean resultant length was calculated for each session, which is an average measure of the magnitude and direction of phase tuning of each cell for a specified frequency and time. In this case, the frequency is theta and we are looking at the total duration of the session. This takes on a value between 0 and 1, where 0 is random spiking with respect to LFP phase, and 1 is spiking at exactly the same phase for all spikes. In other words, we have one value per cell per session, indicating 'how phase tuned' the cell is. We can then calculate this value for every cell to get a sense of trends in the theta phase tuning of the entire population of cells recorded in each region, across sessions and rats. We can see in ***figure 14*** that the HPC,

RSC, and ATN cell populations show different magnitudes and phases of HPC theta tuning. Hippocampal cells show the greatest degree of phase tuning to HPC theta, however a trend is still observed in RSC cells that is anti phase to HPC population tuning, mirroring what we saw in our single cell examples. We see less phase clustering in our ATN cells, however we have a much lower cell sample in the ATN, and this could be influencing our observations. In sum, theta phase

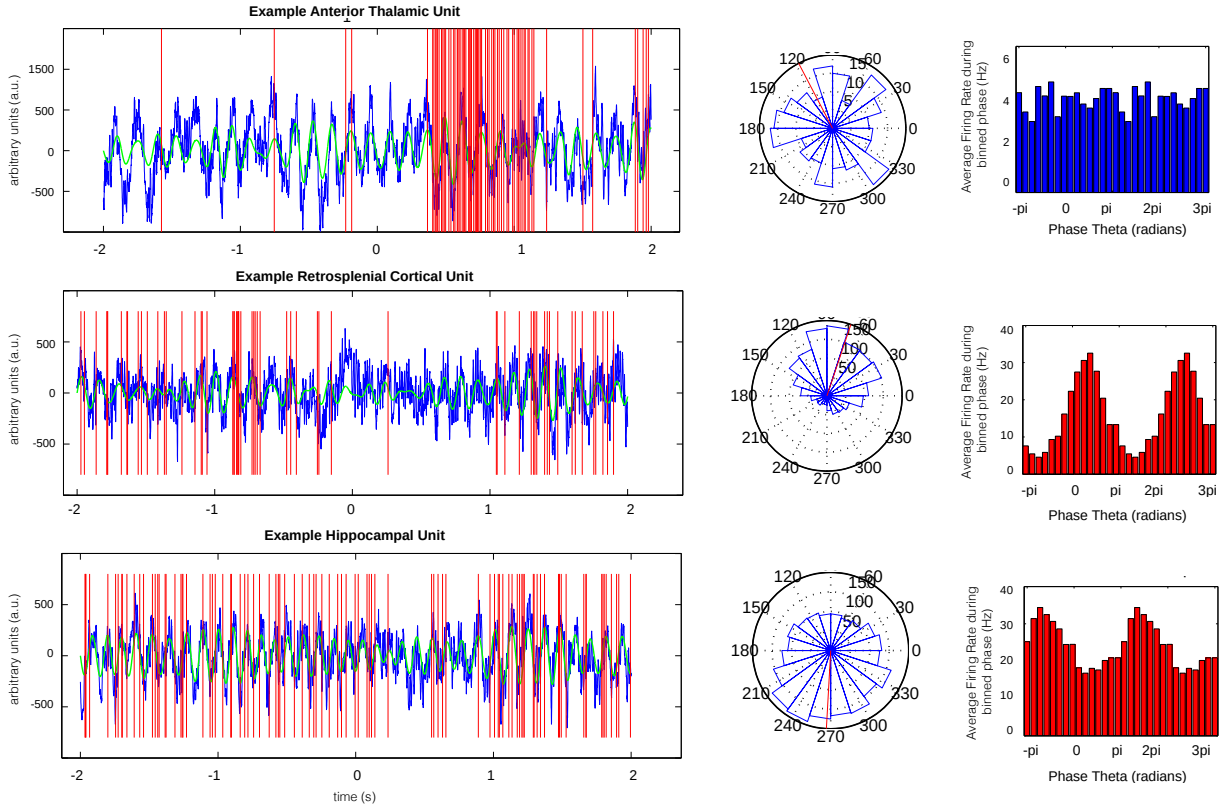

FIGURE 13 | **Theta Phase Tuning of Randomly Selected Examplar Cell in the ATN, RSC, and HPC**

In this plot, each row corresponds to data recorded from a single cell in the ATN, RSC, and HPC, and the corresponding LFP recorded in the same area at that time. ***a*** In these three vertically stacked plots, a selection of raw LFP trace (blue) can be seen with the theta-filtered signal overlaid in green. Red lines denote spikes of a single cell. These are centered around reward receipt, and as can be seen each of these cells have distinct temporal profiles. ***b*** These are circular histograms (rose plots) denoting the distribution of phases of the respective theta signal at which spikes for each cell occur. The red line indicates the mean direction of tuning. ***c*** These are linear histograms plotting the same information as in ***b***, but repeated for two cycles to allow us to see the periodicity of firing more easily. As can be seen, the anterior thalamic example cell is not significantly phase tuned to theta in that range, however the RSC and HPC units are tuned to their respective thetas, and furthermore occur anti phase to one

tuning to HPC theta is seen a subset of cells in each of the three regions we recorded from, the HPC, RSC, and ATN.

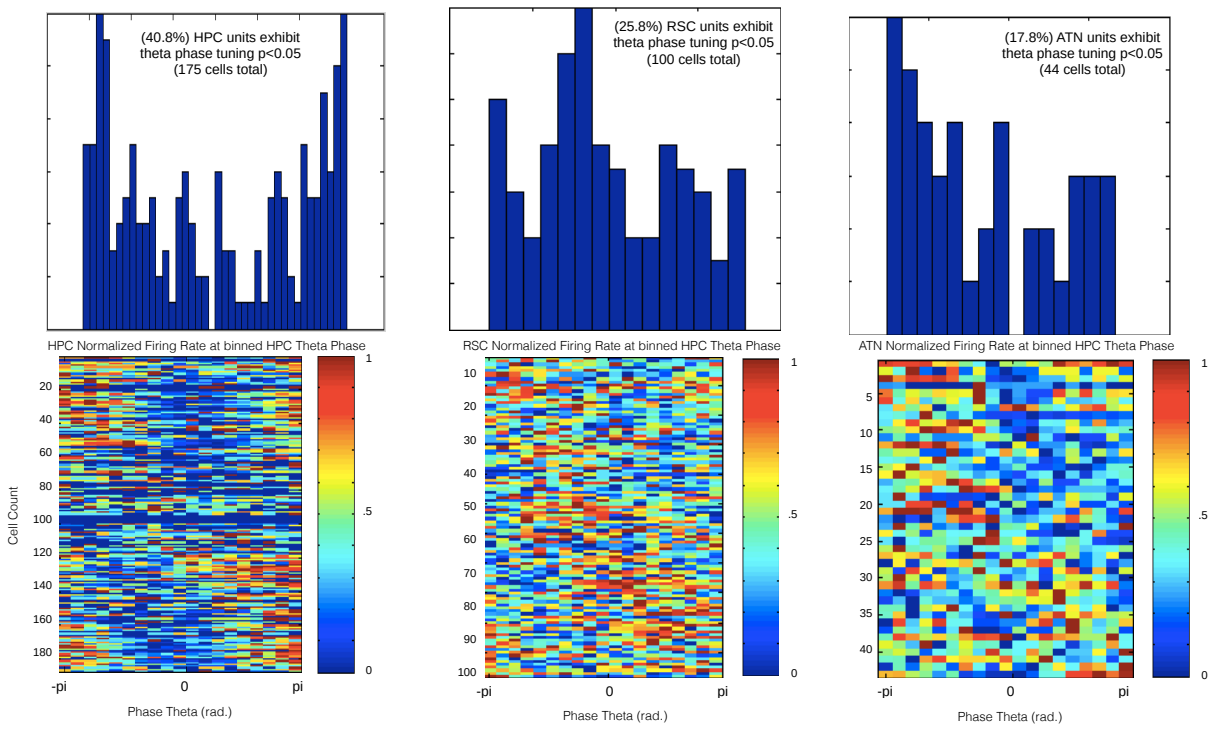

FIGURE 14 | **Average preferred phase of all recorded units in each region to representative HPC theta-frequency oscillation**

a The first column of plots represent the average phase tuning of all HPC cells to HPC theta, from -pi to pi (top), and the changes in normalized firing rate for each cell (0 through 185) with respect to theta phase (bottom). The top plots are essentially collapsed versions of the bottom plots. ***b and c*** Same as *a*, except in the RSC and ATN. All cells are plotted collected from all rats across all sessions. HPC units seem to show a greater population-level phase preference than the RSC or ATN, but that could be due to a difference in the number of cells recorded. All regions show some degree of population-level phase tuning to HPC theta. HPC and RSC average phase tuned properties appear to be anti phase to one another. Notably, the phase difference between peak average population preferred phase in each region approximately mirrors the theta lags seen in figure 12.

DISCUSSION

Theta is known to dominate the hippocampal local field potential signal during awake behavior. Over the decades, this same high amplitude frequency has been documented during similar behavior across distributed regions throughout the brain, all of which have been shown to work in tandem with the hippocampus in the service of learning and memory (*see introduction)*. Theta coherence is the primary frequency within which functional connectivity between these regions, such as the medial prefrontal cortex, entorhinal cortex, and hippocampus, has been found. It comes as no surprise then, that two regions known to be intimately involved with HPC function also exhibit strong theta coherence during nearly all waking behavior. Despite expecting this, theta coherence has not been shown between the ATN and HPC, and has not been characterized between all three regions at one time.

The first set of observations made involve the predominance of theta oscillations in all three regions, and the coherence in theta between these regions. Expectedly, theta coherence is modulated by running speed. Unexpectedly, however, it seems power in each region, ranging from 4 to 200 Hz, is also modulated by running speed. These findings have been reported in gamma frequencies and in theta frequencies in the hippocampus (Ahmed 2012), but have not been systematically studied in other brain regions to our knowledge. Although it is known that the phase of lower frequency oscillations often modulates the amplitude of higher frequency oscillations, the presence of the same dominant frequencies was not a given, nor were the consistencies across regions in increases and decreases in correlation in each frequency band. As mentioned previously, it has been shown recently that the the medial septum directly paces

hippocampal theta via the increasing firing rate of local glutamatergic cells as a function of increased running speed (Fuhrman, 2015). This results in an upward shift in the peak frequency of theta in the HPC. It is also known that the medial septum projects to the RSC and ATN. It is strongly believed that the medial septum helps to coordinate regions distributed throughout the brain via the pacing of local theta. Given the observed larger phase lag between the RSC and HPC however, it is also possible that the HPC itself is driving RSC theta during active behavior. It has also been illustrated that RSC independently produces theta (Talk 2004). We cannot make these determinations given the current findings.

The second set of observations involve beta. First, an increase of beta power at the time of reward and during novel exploration occurs in the HPC and ATN, but not RSC. Secondly, this increase in beta power occurs in 'bursts', perhaps serving to synchronize transient neuronal ensembles. Lastly, this whenever these beta bursts are simultaneously present in the HPC and ATN, they are coherent. However, given the strong dependency of frequency of dominant spectral power on running speed discussed above, and the observation that beta power is highest during low running speeds, we cannot conclusively dissociate two possible explanations for the observed bursts in beta power during reward, as both reward/novelty and running speed changes with or without reward have resulted in these observations in the HPC. In the present study, running speed is confounded with epoch, in that almost all low running speeds occur during the reward period. In support of the theory that these beta bursts are a 'novelty' read out of ATN and HPC activity, however, is the observation that beta power is high throughout the pre training period, during higher running speeds. This observation, in addition to the remarkable change in the spatial selectivity of theta and beta coherence between all region pairs, explored in the last

figure, suggests that beta coherence between the HPC and ATN could be related to some sort of 'this is new' or 'be alert' signal processing. This is of course speculation at this stage and with this data, and both possibilities are supported.

Lastly, there is a strong correlation between running speed and the frequency and power of various frequencies, i.e. as running speed increases, the power and frequency of theta and high gamma increase, and the power and frequency of beta/slow gamma decrease. This anti-correlation in oscillatory regimes could serve an adaptive ecological function. It has been shown that theta cycles serve to organize place cell assemblies in time (Buzsaki 2005), and that the preferred theta phase of these assemblies remains consistent, despite increasing frequency of peak theta-power during running. This means the 'place code' of the HPC is invariant with respect to location and phase of theta, arguably serving as a more reliable read-out of 'where am I?' to downstream regions. Notably, whisking and other peripheral motor patterns are also theta-synchronized. An interesting experimental question could address this same line of reasoning, except exploring beta instead of theta oscillations. How do the timing of beta oscillations (which peak in power during reward receipt and novelty) relate to dopaminergic activity in the basal forebrain? When we see a decrease in running speed in a non-attentive animal (perhaps during quiet wakefulness), do we still see this bursting in beta power? In sum, although the dissociation of 'novelty/reward' and 'running speed' is experimentally feasible, it is important to recognize that the beta oscillatory regime observed in the HPC and ATN, running speed, and attention are all ecologically and physiologically related.

In conclusion, the present study found evidence of sustained theta synchronization between all three regions, the HPC, RSC, and ATN, during movement. Furthermore, theta power,

and thus theta coherence, decreased in tandem with a decrease in locomotion at the site of reward receipt. During this same time, beta power was anti-correlated with theta power in that it increased, although notably only in the ATN and HPC. These increases were not sustained, and instead occurred in transient ‘beta bursts’, which when occurring were always coherent between the HPC and ATN. Theta phase delays between regions matched expected phase delays given their monosynaptic connectivity (~30ms), however there was no phase delay between the HPC and ATN in the beta range. These changes in beta power co-occurring in both regions instantaneously could be driven by some outside region (perhaps the basal forebrain), or it is possible each region is intrinsically oscillating at this frequency. Lastly, a subset of neurons recorded in the HPC, ATN, and RSC were found to release action potentials timed to HPC theta. Notably, cells identified as ‘place’ or ‘reward’ sensitive did not show any difference in the number of tuned cells with respect to cells not identified as being selective for particular task events.

METHODS

Selection of subjects and sessions

Four of the original 16 subjects were used in this study. These subjects were selected because they had ‘good quality’ simultaneous LFP recordings in the HPC, ATN, and RSC in at least one pre-learning and one asymptote session. Good quality was defined as recordings that had less than 1% of the LFP data clipped (due to gain settings during data acquisition), minimal 60 Hz line noise, at least one recorded cell present per session, and power spectral density plots that did not appear to have wide peaks at 60 Hz and its harmonics. For some subjects, sessions recorded across learning (between pre-learning and asymptote) were available, however we did not have enough mid-learning sessions to analyze changes across learning, and instead chose to analyze pre and post learning effects.

Subjects and surgical procedures (modified description from prior studies)

The subjects were 4 food restricted (80–85% of free feeding weight) adult male Long-Evans rats (Simonsen, CA). Movable stereotrode recording electrodes, fabricated by twisting together 2 25 lm lacquer coated tungsten wires (McNaughton et al., 1983), were stereotaxically positioned just above the CA1 field of the hippocampus (2.5–4.5 mm posterior to bregma, 2.5 mm lateral, and 1.7 mm ventral to the brain surface) and the RSC (3.5–4.5 mm posterior to bregma, 0.5 mm lateral, and 0.3 mm ventral). RSC electrodes were implanted in the granular retrosplenial area b (Rgb), also referred to as the posterior cingulate cortex or Brodman’s Area 29c. The rats were anesthetized with sodium pentobarbital (40 mg/kg). They were also given atropine sulfate (0.2

mg/kg) to prevent respiratory congestion, an antibiotic (5 mg/kg Baytril), and an analgesic (5 mg/kg ketoprofen). All procedures complied with guidelines established by the University of Washington Animal Care and Use Committee.

Behavioral Training (modified description from prior studies)

Each rat began training with random foraging sessions, termed ‘pre-learning’. Rats started each trial on a randomly designated arm of a plus maze and foraged for reward (2 drops of chocolate milk) on a different randomly designated arm. The rats were given two blocks of 10 training trials, separated by 30 s of darkness. The training procedures and behavioral requirements did not differ between the two blocks of trials. Trials began when the rats were placed on the maze facing outward at the end of an arm. The rats were allowed to search the maze until they located and consumed the reward. The rats were then placed on a 5.5 by 25-cm platform adjacent to the maze for an inter-trial delay period, during which time the experimenter pretended to bait each of the four arms of the maze, but actually baited only the appropriate arm for the upcoming trial. The duration of the delay varied from trial to trial and among experimenters, but typically lasted 18–25 s. “Run” is defined as the time period from trial start to .5s prior to reward receipt. “Rew” is defined as the time from reward receipt to .5s prior to the rat being picked up to be moved to the inter-trial delay platform, “Rest” is defined as the period while the rat was on the ITD platform.

Rats then began training in the same room on the same maze, but with task manipulation. Rats were trained to retrieve rewards (two drops of chocolate milk) from the east arm of a plus maze during the first half of each training session (15 trials total) and from the west arm during

the second half of the session (15 trials total). Training continued until subjects reached a pre-determined behavioral criterion of at least 75% correct choices on two consecutive sessions. After achieving this criterion, the rats were given 2–10 postcriterial training sessions for the collection of additional neuronal data during asymptotic performance. These asymptotic performance sessions were used for all analyses and are designated ‘Asymptote’. The maze itself occupied a circular area, enclosed by curtains,with objects placed around the perimeter to serve as visual cues [for details, see Yeshenko et al. (2004)].

Data Collection

Neuronal spike data and video data were collected throughout learning with the Cheetah Data Acquisition System (Neuralynx, Bozeman, MT). For this study, we examined data that were collected as part of a study investigating spatial and event-related responses of hippocampal neurons while the rats ran the trials (Smith and Mizumori, 2006b). In this study instead of examining spiking activity, we examined simultaneously recorded local field potential signal in the hippocampus (HPC), retrosplenial cortex (RSC) and anterior thalamus (ATN). Neuronal LFP and video data were collected with the Cheetah Data Acquisition System (Neuralynx, Bozeman, MT). Before training, the recording probes were lowered in 40-micrometer increments until isolatable units were encountered. The electrodes were also advanced to obtain new units between asymptotic performance sessions to improve the cell yield. Local field potentials (LFPs) were sampled continuously at 2 kHz, band-pass filtered at 1–500 Hz, and stored to disk. The rat’s position and direction of travel were monitored by digitized video (sampled at 30 Hz)

of an LED array attached to the rat's head. Event times were flagged manually, and confirmed via inspection of x,y position coordinates, obtained from the video data.

Data Preparation

All LFP, spike, event, and video records were collected, transferred into MATLAB, and stored in the form of one structure per session, labeled by subject number, date, and training stage.

Analysis

LFP preparation and spectra

The LFP is related to the aggregate level of ensemble activity measured as fluctuations in intracellular voltage. Quantifying LFP activity is done by computing a power spectrum. The power spectrum is a measure of the variance in a signal as a function of frequency and is related to the time domain signal through the Fourier transform of the auto-correlation function.

For all analyses, one 'representative' stereotrode was selected in each session from each region. Stereotrodes were selected by maximizing the number of spikes recorded on a given stereotrode in a given session, ensuring high stereotrode quality, while minimizing the degree of LFP data clipping (sessions had already been selected on the basis of at least one stereotrode from each region with less than 1% of the data clipped). LFP were not down sampled from the original 2016 Hz sampling rate.

Average LFP power spectral density plots were computed for each region for each rat using the chronux *mspectrumc* function, giving us a spectral density estimate for each brain

region for each rat using the multi-taper method. This method is an extension of the short-time FFT, designed to increase the signal-to-noise ratio of the frequency representation by applying several tapers that have varying temporal (and thus spectral) characteristics to each data segment, resulting in several independent power spectra estimates (independent due to the orthogonality of the tapers), which are then averaged together. We set our time-bandwidth product (TW) to 3 and used 5 tapers. Output was converted to dB for plotting, 95% confidence intervals are plotted around each average spectrum, computed across sessions within rats.

LFP spectrograms were computed using the MATLAB *spectrogram* function for an epoch of -8 to +13s surrounding trial start (selected on the basis of the average trial and iti dwell-time). This time period includes each of our 3 behavioral epochs of interest, RUN, REWARD, and REST. Spectrograms were computed for each trial and then averaged together across trials, across sessions of the same type (either pre-learning or asymptote), and across rats within session type, producing two spectrograms per recording location (one for PL, one for ASYMP). Spectrograms were computed using the following parameters: a sliding window width of 1 s (this assumes data are stationary within each window - a *weak stationarity assumption*), every ΔT = 100 ms (meaning there is a 90% overlap in our sliding time window, introducing autocorrelation within each frequency band). We computed two different sets of spectrograms, the first set on a linear scale from 1 to 40 Hz and the second set on a log scale from 3 to 200 Hz.

Identification of beta bursts

High-amplitude β-oscillations were identified by first band-pass filtering all single-session LFP traces in the β-range, using the MATLAB *firl* and *filtfilt* functions. The absolute value of the amplitude of the band-pass filtered signal was then converted to z-scores for each

session separately. Peaks were identified where the normalized trace crossed a threshold of >2.5 SD above the session mean. Peaks were considered members of individual bursts if consecutive threshold crossings were separated by at least one oscillation cycle. The duration of beta peaks was then combined across all sessions across rats to generate one distribution describing 'β-burstiness'.

Significance Testing (spectrograms, coherence, and PLV)

Significance testing of power changes within the beta (defined as 25 to 35 Hz) and theta (defined as 6 to 12 Hz) frequency bands was done by first extracting the average theta and beta time course from each region (HPC, RSC, ATN) for each training session (8 pre-learning sessions; 20 asymptote sessions), normalizing the maximum power value within each session to 1 separately for each frequency band, and then computing one average theta and beta time course for each rat across sessions of the same type (PL or ASYMP). Run, reward, and rest time periods were then selected for each rat and all values within those time periods were averaged together. This eventually led to 3 x 3 x 2 x 2 values per rat: average run power, average reward power and average rest power (RUN, REW, REST), for each brain region (HPC, RSC, ATN), within each learning stage (PL, ASYMP) for both beta and theta (THETA, BETA). We were interested in 6 questions: how theta power changes across epochs and across learning, and how beta power changes across epochs and across learning, separately for each brain region. In order to test this, we conducted 6 2-way repeated measures anovas using the MATLAB *anovan* function, with learning epoch (run (1), reward (2), and rest (3)) as one factor, and learning stage (pre-learning (pre) and asymptote (asymp)) as our second factor. In order to perform a repeated measures anova using this matlab function, it is necessary to define 'subject' as a random effect. We had

n=4 values per group (one value for each rat). Repeated measures are particularly vulnerable to violations of sphericity, which can positively bias the F-statistic. In order to account for this, we adjusted our degrees of freedom for each factor and the interaction between factors using each respective Huynh-Feldt epsilon value (custom MATLAB script, designated *eps* in all reporting of statistics). Following the finding of a significant main effect of epoch, we re-grouped our data and performed a one-way repeated measures ANOVA. Following a significant main effect of learning, it was not necessary to do post-hoc testing, because factor learning has only two levels. Following a significant interaction effect, 9 pairwise comparisons and one complex contrast of a priori interest were assessed using paired t-tests. Importantly, this means separate error terms were used for each comparison, instead of a pooled error term, effectively choosing precision over power. This is suggested when violations of sphericity are present. Notably, how to perform follow-up testing with repeated measures anovas is controversial, with contention predominantly centering around whether or not to correct for multiple comparisons. The results presented in this thesis do not make corrections. If Bonferroni corrections or FDR corrections are made, changes in LFP power results are the same, changes in coherence are not; the increase in beta coherence at reward during asymptote between the ATN and HPC drops just below significance at a p=.06. Significance testing for coherence and ISPC was done in precisely the same way, except average coherence or ISPC instead of average power values were used.

Phase-locking analyses

LFP-LFP phase locking analyses were computed using two methods: the multi-taper method via chronux (Mitra and Pesaran 1999; Pesaran et al 2002) and inter-site phase clustering, or phase locking values (Lachaux 1999; custom MATLAB software).

Coherence

The chronux toolbox estimates coherence using the multi-taper method, which refers to the way the power spectral density of each signal is calculated prior to taking the cross correlation between signals and normalizing by their respective power spectra (this subsequent value is referred to as *coherence* and is a measure of the phase-relationship between two signals). For a brief introduction to the multi taper method, see *LFP preparation and spectra* section above. Following calculation of the spectra of two signals, at each frequency band, coherence gives us an estimate of how 'coherent' two signals are, or how correlated the changes in phase are of the two signals. It is also possible to compute coherence across time using a sliding time window, allowing us to visualize changes in our coherence estimates. In these analyses, we computed coherence across time using the same window length and overlap as we did for our spectrogram calculations (1s window, 90% overlap between windows). The tapers are parameterized by their length in time, T, and their bandwidth in frequency, W. We set our time-bandwidth product (TW) to 3 and used 5 tapers. We calculated coherence between all pairs of regions, HPC-RSC, HPC-ATN, and RSC-ATN, for each trial for each session, and then averaged across trials, sessions, and rats, separately for learning stage and each aforementioned pair, in order to obtain estimates across the same behavioral time window (-8 to 13s surrounding reward) for which spectrograms were computed.

Intersite-phase clustering/Phase-locking values

An alternative measure of the phase consistency between two signals is referred to as intersite-phase clustering, sometimes called phase-locking values. Note, there are several terms in the literature used to describe phase-based connectivity, including phase-locking value/

statistic/factor, phase synchronization, and phase coherence. Arguably, these are suboptimal terms, partly because the same terms are used to indicate different methods. In this description, I will use ISPC, as it is a concise description of the method (clustering in polar space of phase angle differences between two LFP signals). All data were collected and averaged the same way as above, however instead of using a sliding time-window to obtain overlapping power spectral density estimates, I convolved each signal with a family of wavelets prior to computing ISPC, with center frequencies ranging from 3 to 45 Hz, spaced logarithmically, and a length of three cycles. Wavelets are constructed by windowing a sine wave of a desired frequency with a gaussian. The standard deviation of the gaussian determines the width of the wavelet, and the frequency of the sine wave determines the 'center frequency' or the frequency one is interested in extracting from an LFP signal via convolution. It is also necessary to add an amplitude scaling factor to the gaussian if interested in extracting power information, however we are only interested in extracting phase information. A benefit of wavelet analysis is the user can specify the central frequencies. Additionally, it is easily possible to vary the width of the wavelet as a function of frequency, so that the stationarity assumption, which holds only where the wavelet has not tapered to 0, can change across frequencies. By defining a wavelet width of 3 cycles, I am stating that the duration of the stationarity assumption changes across frequencies bands, meaning at higher frequencies I will have better temporal resolution and lower frequency resolution. This is preferred, given the spectral properties of the LFP. The result of these convolutions is the wavelet transform of the LFP, or WTLFP (f,t) a matrix of complex numbers whose absolute values (or length) and arguments (or angles) represent the amplitude and phase of the LFP at frequency *f* against time *t*. In sum, the WTLFP represents fluctuations in frequency

band-limited power and phase over time.

The ISPC is calculated by convolving two LFP signals (in our case HPC and ATN, HPC and RSC, or RSC and ATN) with the same family of wavelets. Then, for each frequency band the time series of phase values of length n from one region is subtracted from the time series of phase values of length n from another region, resulting in a vector of length n-1. These differences represent the change in lag between two signals, and can itself be informative if determining the average phase lag (windowed or not) if it is desired. The ISPC is measure of the dispersion of the angles of these differences. It is calculated by windowing the above data, thus obtaining a series of angles (which can be conceptualized on a unit circle as in figure 12), and then calculating the mean resultant length of these vectors. This value has a magnitude and an angle. The angle is simply the average angle of all vectors, and the magnitude is what we call the ISPC. The ISPC is 1 if two signals are perfectly coherent (ie the phase difference is always exactly the same), and the ISPC value is 0 if the phase angle differences are evenly dispersed across the unit circle. Coherence estimates a similar thing in a different way, thus calculating both of these measures and confirming they are the same is of value. All averaging across trials, sessions, and rats was done in the same way.

Power/Coherence as a function of running speed

Plots of power at each running speed represent averages across all rats, separated by learning stage and brain region. Spectrograms were calculated as described previously for every trial, except frequencies between 3 and 200 Hz, logarithmically spaced, were used. These were matched with running speed during that trial. Spectrogram/running speed matched from all trials, sessions of the designated type, and rats were concatenated, and this matrix was resorted in

ascending order by running speed. The large matrix was then smoothed and plotted. Pearson's correlation coefficients were calculated between speed and each band-limited power time series. Calculation of ISPC values changing as a function of running speed was done in the same fashion.

Coherence as a function of location

In order to calculate changes in theta and beta coherence between each region pair as a function of location, we used our previously calculated coherence measures, extracted theta and beta coherence time series, and then computed animal location at each point during these time series using the recorded video data. This was done separately for each trial, when coherence was z scored, and then all trials, sessions, and rats were averaged together for each learning stage (pre learning and asymptote) and coherence pair). Because we were interested in relative changes in coherence, we z scored coherence values. Colorscaling on all plots is the same, starting at 0, therefore preventing visualization of decreases in coherence in order to make increases more easily visible.

Phase-Spike Coupling Z-scores (PSCz); Mean Resultant Length

All spike-phase analyses (i.e. the assessment of spike phase-locking to a local field potential time series of interest) were done by calculating the mean resultant length (very common; e.g. *see O'Neill 2013*) and assessed for significance using a time-shuffled resampling procedure described in detail under the title *Phase-Amplitude Coupling z score (PACz)* in Michael X Cohen's fantastic book; Analyzing Neural Time Series Data (includes MATLAB code). This method was adopted for spikes instead of for power-phase lfp-lfp coupling, so I will refer to it as

PSCz (phase-spike coupling z score). Lastly, this average phase-locking measure was calculated for each unit across the entire session. This is because PSCz (and MRL) are sensitive to the number of spikes used to calculate the measure, and have a positive bias. This means with less spikes the increasing likelihood of finding significant phase-locking. Some papers have employed a 'spike-thinning' method in order to keep the number of spikes consistent (O'Neill 2013), however I did not implement this and instead chose to maximize my spike count by finding a single value across the entire session, as mentioned. I also only looked at phase tuning of spikes to theta (6-12 Hz). This same analysis could be done for any band-pass filtered LFP range (such as Beta).

The mean resultant length is calculated as follows. Within one session at a time, field potentials from the region of interest are first band-pass filtered in the theta range by constructing a fir1 filter, applying the filter using the function filtfilt, and then extracting instantaneous phase using the hilbert MATLAB functions (see previous methods descriptions). Then, for each unit recorded in a session, each unit spike was assigned a phase value based on its simultaneous field potential sample. The mean resultant length is then calculated across the session for this unit by summing each phase value for said unit vector and dividing by the total number of spikes. The resulting value will be a number between 0 and 1, where 0 means there is no phase-tuning of the spikes, and 1 means there is perfect phase tuning. Statistical significance of phase-tuning was assessed by generating a null hypothesis distribution, with the null hypothesis being 'expected MRL value when timing of spikes ablated'. This was done by randomly shuffling the spike time series 1000 times, and calculating the same MRL value at each iteration. A null hypothesis distribution was then generated with these MRL values expected under the null hypothesis. The

actual MRL value was then translated into a z-score based on the distribution of null hypothesis MRL values. Notably, using a z-score is sensible in this case because the null hypothesis distribution is a distribution of means of means, thus it is normally distributed. An alternate way to test significance of phase-locking is by using the Rayleigh test for circular uniformity, which is ideal only for unimodal distributions (used in O'Neill 2013).

CHAPTER 4

CONCLUSIONS AND FUTURE DIRECTIONS

A yet to be conquered frontier across all levels of neuroscience is a rich, mathematically detailed account of the processes by which an organism learns, stores, and retrieves information. This is not only a fascinating problem from a computational perspective, likely yielding a set of algorithms which will revolutionize everything from artificial intelligence to how we retrieve information from the internet, but it is also a problem of great romance. To understand episodic memory is to understand a small piece of how we construct a coherent story about ourselves.

The null result of the first study highlighted the usefulness - and shortcomings - of utilizing lesion methods to approach the study of a brain region function. The second study, very much a foil to the first, highlighted the importance of an experimental design catered to a specific set of a priori hypotheses, and in turn the shortcomings of exploratory data analysis (due to our inability to dissociate novelty from low running speed). Despite this, incredibly rich dynamics were seen in the three interconnected brain regions, which underscored the potential usefulness of the 'coherence as a proxy of information exchange' approach to analyzing local field potential data.

Summary of Findings

Importantly, a series of novel findings were made in this series of studies. In study one, we found that the RSC was not necessary when rats are solving this conditional discrimination task, which has previously been shown by Howard Eichenbaum's lab to be HPC and mPFC -

dependent. This came as a surprise, given the monosynaptic connectivity the RSC shares with the HPC and mPFC, and given the previously discussed putative functional role of this region. However, as discussed in the chapter, the ventral hippocampus receives privileged input directly from the olfactory bulb, which means it is very likely there is a route by which olfactory information (in this case our item) can reach the hippocampus, allowing the formation of item and context. A suggested future direction following the lesion study in chapter 2 is to repeat the study with visual instead of olfactory cues. A future step likely to be of use is also the inclusion of temporary lesions at different stages of learning instead of permanent lesions prior to learning.

Although we found evidence of a lack of involvement of the retroslpenial cortex in the aforementioned olfactory item - context conditional discrimination task, it does not mean these regions do not communicate with one another throughout the course of contextual learning in general. Chapter 3 took an approach that differed greatly from the first study, in that it was entirely exploratory. The spike and local field potential (LFP) data in the hippocampus, retrosplenial cortex, and anterior thalamus were previously collected by Dr. David Smith, and were approached with minimal a priori assumptions regarding what would be the predominant spectral and coherence features of the data. To briefly summarize the findings and put them in a broader context, we found strong theta coherence between all three region pairs, and beta coherence between the ATN and HPC specifically during novel experience and reward receipt. Beta coherence between these regions (HPC and ATN) occurred in stereotypes ‘bursts’, similar to what has been found in previous studies, and is likely related to novelty processing and is perhaps timed by the dopaminergic system. It is important to note, however, that in this study running speed and novelty/reward are confounded, and after examining changes in the peak

power and frequency of the LFP at different running speeds, it is clear running speed and power are intimately correlated. This has previously been shown in the HPC for theta and gamma, but has not been shown in the ATN or RSC. The existence of the same spectral relationship to running speed in all three regions suggests a fundamental intrinsic or extrinsic driver of these microcircuits. Furthermore, a subset of neurons in all three regions were found to preferentially phase-lock to the theta-range frequency band. This is expected, as theta phase-locking has been found in many regions that share monosynaptic connectivity with the HPC, including the mPFC (Wilson 2005), and the ATN (Cho 2001).

These data are very rich, and many questions have yet to be inspected. Immediate future steps could be to look at phase-amplitude coupling within regions, in addition to analyzing how oscillations of frequencies other than theta in each region organize spike timing. Of particular interest is whether RSC cells, which were found to strongly respond to reward cues, transiently synchronize their firing to a beta instead of theta rhythm surrounding that time of reward. Very preliminary analyses have found a large subset of cells that spike in time with theta in each region, when averaged across a session. Of course, it is possible (and very very likely) that there is great heterogeneity in the firing rates of cells, and that various cells, due to their cell-type and the intrinsic structure of circuit within which they are embedded, exhibit many different types of firing patterns, occasionally paced by theta or beta. Looking at transient spike-phase coupling is then also of great interest. Lastly, it has been hypothesized for over a decade and illustrated recently that when sharp wave ripples occur in the hippocampus, a particularly transient high frequency burst, the thalamus becomes ‘silent’ and cortex becomes very active (Logothetis 2012). It is thought that these transient, tens of ms long bursts allow privileged communication

between the hippocampus and cortex. This observation strongly supports predictions made by systems consolidation theory, and, given the recordings in thalamus, hippocampus, and cortex, it would be fascinating to see if we can make the same observations in this data. Important to note, however, is that we have very few periods of quiescence, as recordings are done when the animal is actively engaged in a task, not when the animal is sleeping (which is when SPW-R's are most frequently observed).

Thirty years ago, less than two hundred articles mentioned the RSC, and now there are well over one thousand (pubmed). It is Dr. David Smith's great insight that has lead him to this fascinating and important region far before many others. It is my hope that with this growth in interest in higher order cortex, and simultaneous increase in the number and availability of advanced high-density experimental methodologies, that we will begin to *solve* these regions of the brain.

LITERATURE CITED

Aggleton, J. P., & Brown, M. W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. The Behavioral and Brain Sciences, 22(3), 425–489.

Aggleton, J. P. (2010). Understanding retrosplenial amnesia: insights from animal studies. Neuropsychologia, 48(8), 2328–38. doi:10.1016/j.neuropsychologia.2009.09.030

Ahmed, O. J., & Mehta, M. R. (2012). Running Speed Alters the Frequency of Hippocampal Gamma Oscillations. Journal of Neuroscience, 32(21), 7373–7383. doi:10.1523/JNEUROSCI.5110-11.2012.

Albo, Z., Viana Di Prisco, G., & Vertes, R. P. (2003). Anterior thalamic unit discharge profiles and coherence with hippocampal theta rhythm. Thalamus & Related Systems, 2(2), 133–144. doi:http://dx.doi.org/10.1016/S1472-9288(03)00006-2.

Alexander, A. S., & Nitz, D. a. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. Nature Neuroscience, (July), 1–12. doi:10.1038/nn.4058.

Berke, J. D., Hetrick, V., Breck, J., & Greene, R. W. (2008). Transient 23-30 Hz oscillations in mouse hippocampus during exploration of novel environments. Hippocampus, 18(5), 519–529. doi:10.1002/hipo.20435.

Billeh, Y. N., Schaub, M. T., Anastassiou, C. a, Barahona, M., & Koch, C. (2014). Revealing cell assemblies at multiple levels of granularity. Journal of Neuroscience Methods, 236, 92–106. doi: 10.1016/j.jneumeth.2014.08.011.

BG Cooper, SJY Mizumori (2001). Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus.
The Journal of Neuroscience 21 (11), 3986-4001.

Bunzeck, N., Guitart-Masip, M., Dolan, R. J., & Düzel, E. (2011). Contextual novelty modulates the neural dynamics of reward anticipation. Journal of Neuroscience, 31(36), 12816–12822. doi: 10.1523/JNEUROSCI.0461-11.2011.

Buzsaki, G. (2005). Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. Hippocampus, 15(7), 827–840. doi:10.1002/hipo. 20113.

Buzsaki, G. (2006). Rhythms of the Brain. Oxford University Press.

Buzsáki, G. (2002).. Theta Oscillations in the Hippocampus. Neuron, 33(3), 325–340. doi: 10.1016/S0896-6273(02)00586-X.

Buzsaki, G. (2011). Neural Syntax: cell assemblies, synapsembles and readers. Neuron, 68(3), 362–385. doi:10.1016/j.neuron.2010.09.023.

Cohen MX (2014) Analyzing Neural Time Series Data: Theory and Practice Issues in Clinical and Cognitive Neuropsychology.

Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., … Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. Nature, 462, 353–357. doi:10.1038/nature08573.

Cowansage, K. K., Shuman, T., Dillingham, B. C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct Reactivation of a Coherent Neocortical Memory of Context. Neuron, 84(2), 432–441. doi:10.1016/j.neuron.2014.09.022.

Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. Behavioral Neuroscience, 115(1), 3–25. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11256450.

DeCoteau WE, et al. (2007) Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task. Proc Natl Acad Sci USA 104:5644–5649.

Einevoll, G. T., Kayser, C., Logothetis, N. K., & Panzeri, S. (2013). Modelling and analysis of local field potentials for studying the function of cortical circuits. Nature Reviews Neuroscience, 14(11), 770–85. doi:10.1038/nrn3599.

Engel, A. K., & Fries, P. (2010). Beta-band oscillations - signalling the status quo? Current Opinion in Neurobiology, 20(2), 156–165. doi:http://dx.doi.org/10.1016/j.conb.2010.02.015.

Fell, J., & Axmacher, N. (2011). The role of phase synchronization in memory processes. Nature Reviews. Neuroscience, 12(2), 105–18. doi:10.1038/nrn2979.

Florin, E., & Baillet, S. (2015). The brain's resting-state activity is shaped by synchronized cross-frequency coupling of oscillatory neural activity. NeuroImage, 111, 26–35. doi:10.1016/j.neuroimage.2015.01.054.

Foster, B. L., Kaveh, A., Dastjerdi, M., Miller, K. J., & Parvizi, J. (2013). Human retrosplenial cortex displays transient theta phase locking with medial temporal cortex prior to activation during autobiographical memory retrieval. Journal of Neuroscience, 33(25), 10439–46. doi:10.1523/JNEUROSCI.0513-13.2013.

Fuhrmann, F., Justus, D., Sosulina, L., Kaneko, H., Beutel, T., Friedrichs, D., … Remy, S. (2015). Locomotion, Theta Oscillations, and the Speed-Correlated Firing of Hippocampal Neurons Are Controlled by a Medial Septal Glutamatergic Circuit. Neuron, 86(5), 1253–1264. doi:10.1016/j.neuron.2015.05.001.

Gabriel M *Neurobio of Cingulate Cortex.* Boston: Birkhauser (1993).

Geva-Sagiv, M., Las, L., Yovel, Y., & Ulanovsky, N. (2015). Spatial cognition in bats and rats: from sensory acquisition to multiscale maps and navigation. Nat Rev Neurosci, 16(2), 94–108. Retrieved from http://dx.doi.org/10.1038/nrn3888.

Gill, P. R., Mizumori, S. J. Y., & Smith, D. M. (2011). Hippocampal Episode Fields Develop with Learning. Hippocampus, 21(11), 1240–1249. doi:10.1002/hipo.20832.

Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, R. (2009). High-frequency, long-range coupling between prefrontal and visual cortex during attention. Science, 324, 1207–1210. doi: 10.1126/science.1171402.

Harding A, Halliday G, Caine D, Kril J. Degeneration of anterior thalamic nuclei differentiates alcoholics with amnesia. Brain. 2000; 123(Pt 1):141–154.

He, B. J., Zempel, J. M., Snyder, A. Z., & Raichle, M. E. (2010). The temporal structures and functional significance of scale-free brain activity. Neuron, 66(3), 353–369. doi:10.1016/ j.neuron.2010.04.020.

Hindley, E. L., Nelson, A. J. D., Aggleton, J. P., & Vann, S. D. (2014). Dysgranular retrosplenial cortex lesions in rats disrupt cross-modal object recognition. Learning & Memory (Cold Spring Harbor, N.Y.), 21(3), 171–9. doi:10.1101/lm.032516.113.

Howe, M. W., Atallah, H. E., McCool, a., Gibson, D. J., & Graybiel, a. M. (2011). Habit learning is associated with major shifts in frequencies of oscillatory activity and synchronized spike firing in striatum. Proceedings of the National Academy of Sciences, 108(40), 16801–16806. doi: 10.1073/pnas.1113158108.

Hutchison, R. M., Womelsdorf, T., Allen, E. a, Bandettini, P. a, Calhoun, V. D., Corbetta, M., … Chang, C. (2013). Dynamic functional connectivity: promise, issues, and interpretations. NeuroImage, 80, 360–78. doi:10.1016/j.neuroimage.2013.05.079.

Igarashi, K. M., Lu, L., Colgin, L. L., Moser, M.-B., & Moser, E. I. (2014). Coordination of entorhinal-hippocampal ensemble activity during associative learning. Nature. doi:10.1038/ nature13162.

Jenkins, T.A., Amin, E., Brown, M.W., Aggleton, J.P.(2006) Changes in immediate early gene expression in the rat brain after unilateral lesions of the hippocampus. Neuroscience 137:747–759.

Jenkins, T.A., Dias, R., Amin, E., Brown, M.W., Aggleton, J.P.(2002b) Fos imaging reveals that lesions of the anterior thalamic nuclei produce widespread limbic hypoactivity in rats. J. Neurosci. 22:5230–5238.

Katche, C., Dorman, G., Gonzalez, C., Kramar, C.P., Slipczuk, L., Rossato, J.I., Cammarota, M., and Medina, J.H. (2013a). On the role of retrosplenial cortex in long-lasting memory storage. Hippocampus 23, 295-302. doi: 10.1002/hipo.22092.

Keene, C.S., and Bucci, D.J. (2008a). Contributions of the retrosplenial and posterior parietal cortices to cue-specific and contextual fear conditioning. Behav Neurosci 122, 89-97. doi: 10.1037/0735-7044.122.1.89.

Keene, C.S., and Bucci, D.J. (2008b). Involvement of the retrosplenial cortex in processing multiple conditioned stimuli. Behav Neurosci 122, 651-658. doi: 10.1037/0735-7044.122.3.651.

Keene, C.S., and Bucci, D.J. (2008c). Neurotoxic lesions of retrosplenial cortex disrupt signaled and unsignaled contextual fear conditioning. Behav Neurosci 122, 1070-1077. doi: 10.1037/a0012895.

Keene, C.S., and Bucci, D.J. (2009). Damage to the retrosplenial cortex produces specific impairments in spatial working memory. Neurobiol Learn Mem 91, 408-414. doi: 10.1016/j.nlm.2008.10.009.

Komorowski, R. W., Manns, J. R., & Eichenbaum, H. (2009). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. The Journal of Neuroscience, 29(31), 9918–9929. doi:10.1523/JNEUROSCI.1378-09.2009.

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

Lachaux, J. P., Rodriguez, E., Martinerie, J., & Varela, F. J. (1999). Measuring phase synchrony in brain signals. Human Brain Mapping, 8(4), 194–208. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10619414.

Law, L. M., & Smith, D. M. (2012). The anterior thalamus is critical for overcoming interference in a context-dependent odor discrimination task. Behavioral Neuroscience, 126(5), 710–719. doi: 10.1037/a0029698.

Lisman, J. E., & Jensen, O. (2013). The Theta-Gamma Neural Code. Neuron, 77(6), 1002–1016. doi:10.1016/j.neuron.2013.03.007.

Logothetis, N. K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H. C., … Oeltermann, a. (2012). Hippocampal-cortical interaction during periods of subcortical silence. Nature, 491(7425), 547–53. doi:10.1038/nature11618.

Lømo, Terje (1966). "Frequency potentiation of excitatory synaptic activity in the dentate area of the hippocampal formation". Acta Physiologica Scandinavica 68 (Suppl 277.

Long, L. L., Hinman, J. R., Chen, C. M., Escabi, M. a., & Chrobak, J. J. (2014). Theta dynamics in rat: Speed and acceleration across the septotemporal axis. PLoS ONE, 9(5), 1–11. doi: 10.1371/journal.pone.0097987.

Masimore, B., Kakalios, J., & Redish, a. D. (2004). Measuring fundamental frequencies in local field potentials. Journal of Neuroscience Methods, 138(1-2), 97–105. doi:10.1016/j.jneumeth. 2004.03.014.

Maviel, T., Durkin, T.P., Menzaghi, F., and Bontempi, B. (2004). Sites of neocortical reorganization critical for remote spatial memory. Science 305, 96-99. doi: 10.1126/science. 1098180.

McKenzie, S., Frank, A. J., Kinsky, N. R., Porter, B., Rivière, P. D., & Eichenbaum, H. (2014). Hippocampal Representation of Related and Opposing Memories Develop within Distinct, Hierarchically Organized Neural Schemas. Neuron, 83(1), 202–215. doi:http://dx.doi.org/ 10.1016/j.neuron.2014.05.019.

Montgomery SM, Buzsa´ki G (2007) Gamma oscillations dynamically couple hippocampal CA3 and CA1 regions during memory task perfor- mance. Proc Natl Acad Sci U S A 104:14495–14500.

Moser, E. I., Roudi, Y., Witter, M. P., Kentros, C., Bonhoeffer, T., & Moser, M.-B. (2014). Grid cells and cortical representation. Nat Rev Neurosci, 15(7), 466–481. Retrieved from http://dx.doi.org/10.1038/nrn3766.

N. Kopell, C. Borgers, D. Pervoichine, P. Malerba, and A. T. (2010). Gamma and Theta Rhythms in Biophysical Models. doi:10.1007/978-1-4419-0996-1. V. Cutsuridis et al. (eds.), Hippocampal Microcircuits, Springer Series in Computational Neuroscience 5, doi: 10.1007/978-1-4419-0996-1 15.

Nadel, L., Winocur, G., Ryan, L., & Moscovitch, M. (2007). Systems consolidation and hippocampus: two views. Debates in Neuroscience, 1(2-4), 55–66. doi:10.1007/s11559-007-9003-9.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Research, 34(1), 171–175. doi:http://dx.doi.org/10.1016/0006-8993(71)90358-1.

O‘Mara, S. M. (2013). The Anterior Thalamus Provides A Subcortical Circuit Supporting Memory And Spatial Navigation . Frontiers in Systems Neuroscience . doi: 10.3389/fnsys.2013.00045

O’Neill, P.-K., Gordon, J. a, & Sigurdsson, T. (2013). Theta oscillations in the medial prefrontal cortex are modulated by spatial working memory and synchronize with the hippocampus through its ventral subregion. Journal of Neuroscience, 33(35), 14211–24. doi:10.1523/JNEUROSCI.2378-13.2013.

Robertson, R. G., Rolls, E. T., Georges-Francois, P., & Panzeri, S. (1999). Head direction cells in the primate pre-subiculum. Hippocampus, 9(3), 206–219. doi:10.1002/(SICI)1098-1063(1999)9:3<206::AID-HIPO2>3.0.CO;2-H.

Schomburg, E. W., Fernández-Ruiz, A., Mizuseki, K., Berényi, A., Anastassiou, C. A., Koch, C., & Buzsáki, G. (2014). Theta Phase Segregation of Input-Specific Gamma Patterns in Entorhinal-Hippocampal Networks. Neuron, 1–16. doi:10.1016/j.neuron.2014.08.051.

Scoville, W.B., Milner, B.(1957) Loss of recent memory after bilateral hippocampal lesions. J. Neurol. Neurosurg. Psychiatr. 20:11–21.

Siapas, A. G., Lubenov, E. V, & Wilson, M. a. (2005). Prefrontal phase locking to hippocampal theta oscillations. Neuron, 46(1), 141–51. doi:10.1016/j.neuron.2005.02.028.

Squire, L. R. (2009). The Legacy of Patient H.M. for Neuroscience. Neuron, 61(1), 6–9. doi: 10.1016/j.neuron.2008.12.023.

Siegel, M., Donner, T. H., & Engel, A. K. (2012). Spectral fingerprints of large-scale neuronal interactions. Nature Reviews. Neuroscience, 13(2), 121–34. doi:10.1038/nrn3137.

Sirota, A., Csicsvari, J., & Buhl, D. (2003). Communication between neocortex and hippocampus, 100(4).

Smith DM. 2008. The hippocampus, context processing and episodic memory. In: Huston JP, editor. Handbook of Behavioral Neuro- science, Vol. 18, Handbook of Episodic Memory.

Smith, D. M., Barredo, J., & Mizumori, S. J. Y. (2012). Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. Hippocampus, 22(5), 1121–33. doi: 10.1002/hipo.20958.

Smith, D. M., Wakeman, D., Patel, J., & Gabriel, M. (2004). Fornix lesions impair context-related cingulothalamic neuronal patterns and concurrent discrimination learning in rabbits (Oryctolagus cuniculus). Behavioral Neuroscience, 118(6), 1225–39. doi: 10.1037/0735-7044.118.6.1225.

Sporns, O. (2010). Networks of the Brain. MIT Press. Retrieved from https://books.google.com/books?id=v1DBKE7-UrYC.

Srinath, R., & Ray, S. (2014). Effect of amplitude correlations on coherence in the local field potential. Journal of Neurophysiology, 112(4), 741–751. Retrieved from http://jn.physiology.org/content/112/4/741.abstract.

Steriade M, Timofeev I, Grenier F (2001) Natural waking and sleep states: a view from inside neocortical neurons. J Neurophysiol 85: 1969-1985.

Sulpizio, V., Committeri, G., Lambrey, S., Berthoz, A. & Galati, G. Selective role of lingual/parahippocampal gyrus and retrosplenial complex in spatial memory across viewpoint changes relative to the environmental reference frame. Behav. Brain Res. 242, 62–75 (2013).

Talk, A., Kang, E., & Gabriel, M. (2004). Independent generation of theta rhythm in the hippocampus and posterior cingulate cortex. Brain Research, 1015(1-2), 15–24. doi:10.1016/j.brainres.2004.04.051.

Taube, J. S., Muller, R. U., and Ranck, J. B. Jr. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. J. Neurosci. 10, 420–435.

Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. The Journal of Neuroscience, 15(1 Pt 1), 70–86.

Taube, J. S., and Bassett, J. P. (2003). Persistent neural activity in head direction cells. Cereb. Cortex 13, 1162–1172.

Tingley, D., Alexander, a. S., Quinn, L. K., Chiba, a. a., & Nitz, D. a. (2015). Cell Assemblies of the Basal Forebrain. Journal of Neuroscience, 35(7), 2992–3000. doi:10.1523/JNEUROSCI.4432-14.2015.

Valenstein, E. et al. Retrosplenial amnesia. Brain 110, 1631–1646 (1987).

Van Groen, T. & Wyss, J.M. Connections of the retrosplenial granular a cortex in the rat. J. Comp. Neurol. 300, 593–606 (1990).

van der Meer MA, Redish AD (2009) Low and high gamma oscillations in rat ventral striatum have distinct relationships to behavior, reward, and spiking activity on a learned spatial decision task. Front Integr Neurosci 3:9.

Vann, S.D., Brown, M.W., Erichsen, J.T., Aggleton, J.P.(2000b) Using Fos imaging in the rat to reveal the anatomical extent of the disruptive effects of fornix lesions. J. Neurosci. 20:8144–8152.

van Wingerden M, Vinck M, Lankelma J, Pennartz CM (2010) Theta-band phase locking of orbitofrontal neurons during reward expectancy. J Neurosci 30:7078–7087.

Wang XJ (2010) Neurophysiological and computational principles of cortical rhythms in cognition. Physiol Rev 90:1195–1268.

Warburton EC, Baird AL, Aggleton JP. Assessing the magnitude of the allocentric spatial deficit associated with complete loss of the anterior thalamic nuclei in rats. Behavioural Brain Research. 1997; 87(2):223–232.Warburton EC, Aggleton JP. Differential deficits in the Morris water maze following cytotoxic lesions of the anterior thalamus and fornix transection. Behavioural Brain Research. 1999; 98(1):27–38.

Wyss, J.M. & Van Groen, T. Connections between the retrosplenial cortex and the hippocampal formation in the rat: a review. Hippocampus 2, 1–11 (1992).

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron, 27(3), 623–633.

Young, C. K., & McNaughton, N. (2009). Coupling of theta oscillations between anterior and posterior midline cortex and with the hippocampus in freely behaving rats. Cerebral Cortex, 19(1), 24–40. doi:10.1093/cercor/bhn05.