Received: 16 April 2020 | Revised: 15 August 2020 | Accepted: 26 August 2020

DOI: 10.1111/ejn.14963

EDITORIAL

EJN European Journal of Neuroscience FENS WILEY

An evolving perspective on the dynamic brain: Notes from the Brain Conference on *Dynamics of the brain: Temporal aspects of computation*

1 | INTRODUCTION

It is inescapable that we exist in a world that changes; so too must the brain perform its computational feats of perception, motor control, learning, memory, and speech online as the world and the brain's own state dynamically evolve. Recent theories have emphasized that neural computation might directly exploit dynamic principles as a powerful means of processing inputs, exerting control over action, and regulating and updating internal state, rather than by the maintenance of static regimes of activity that attempt to counteract the inevitable process of time. A dynamic perspective on the brain, with a focus on the computational role of transient patterns of neural activity, has been energized by recent advances in recording technologies and analysis methodologies, which have revealed a diversity of patterned neural activity across multiple brain regions and over multiple timescales. In conjunction, methodological advances in the analysis, interpretation, and perturbation of dynamic brain activity have yielded both fresh insight and novel questions regarding the computational nature of transient patterns of neural activity, and the regulation and control of the dynamic brain.

Last summer, neuroscientists from around the globe gathered in Denmark for the Brain Conference on *Dynamics of the brain: temporal aspects of computation*, sponsored by FENS and the Lundberg Foundation and chaired by Gilles Laurent and Ila Fiete. The goal of this meeting was to discuss recent experimental findings and novel theoretical ideas on the role of dynamic neural activity in the computational repertoire of the brain, and to identify promising directions for future research. Below, we survey the research presented at the meeting, covering the state-of-the-art in the isolation and interpretation of dynamic brain activity across a range of model systems and in the support of varied behaviors. We have arranged this report thematically, to focus on the broad concepts that emerged during the meeting relevant for understanding the dynamics of neural computation, although many, if not all, speakers touched on several of these concepts in the course of presenting their research. Along the way, we highlight questions and considerations that arose as future directions for the dynamic perspective on the brain.

Abbreviations: BPTT, Backpropagation through time; ERP, event-related potential; LSTM, Long Short-Term Memory network; MD, mediodorsal thalamus; MEC, medial entorhinal cortex; PFC, prefrontal cortex; PG, diencephalic preglomerular complex; STG, stomatogastric ganglion; SWR, sharp wave ripple; VS, ventral striatum.

Langdon and Chaudhuri contributed equally to this work.

2 | SPACE: A FUNDAMENTAL SUBSTRATE FOR DYNAMIC NEURAL COMPUTATION

The key dimensions for representing a dynamic variable are space and time—accordingly, several talks focused on the neural representation of these quantities in the brain across a diverse range of behavioral tasks. Among these was an in-depth survey presented by Edvard Moser of the spatiotemporal properties of neural activity in the medial entorhinal cortex (MEC), a brain region which is, jointly with the hippocampus, thought to play an important role in spatial navigation (Moser, Moser, & McNaughton, 2017; Strange, Witter, Lein, & Moser, 2014). A subset of cells in the MEC fire according to a hexagonal grid pattern, displaying high activity at spatial locations that repeat at a characteristic spatial scale for each sub-region ("module") of MEC (Hafting, Fyhn, Molden, Moser, & Moser, 2005; Stensola et al., 2012). Influential "bump" attractor models posit that local circuit interactions are responsible for the establishment of these grid-like patterns of activity, with external input able to move the network state continuously among stable modes (McNaughton, Battaglia, Jensen, Moser, & Moser, 2006). Consistent with these models, Moser demonstrated that activity in MEC networks is indeed low dimensional, with structured pairwise correlations between recorded neurons that are invariant across running and epochs of short-wave and REM sleep, establishing that intrinsic neural activity is constrained (at least

in part) by local interactions rather than solely by external input (Gardner, Lu, Wernle, Moser, & Moser, 2019). Furthermore, Moser showed recent work demonstrating that MEC neural activity is not only constrained by a spatial grid structure, but also displays temporal structure—large multi-unit recordings in MEC during sensory-deprived spontaneous locomotion (i.e., in the dark) display robust sequences of activity that are highly stereotyped, directional and consistently repeated on the timescale of tens of seconds, raising a challenge for attractor models of grid cell activity to account for this dynamic activity sustained intrinsically by the network across these long timescales.

Further extending our understanding of space in the brain, Michael Yartsev presented a series of findings on spatial representation in the flying Egyptian fruit bat, an animal in which three-dimensional space is the ethologically relevant domain for brain computation and behavioral control. These bats have sophisticated spatial navigation skills and are able to both travel hundreds of kilometers to a remembered location and then precisely move around that location to forage. In his talk, Yartsev focused on the ways in which bats' representation of 3D space must go beyond simply extending the strategies and representations that we know from 2D navigation. For example, a strategy in which neurons are sharply tuned to particular locations may be much less efficient in 3D, where there are many more locations to cover than in 2D. To highlight the characteristic features of spatial representation in these bats, Yartsev showed two distinct modes of their movement through the air: these bats either perform large-scale commutes over tens or even hundreds of kilometers, or they perform local foraging at their destination by flying in and around trees. In both these modes, bats move in a very restricted domain of the huge movement space afforded by three dimensions and, because they cannot stop in midair, appear to plan a structured trajectory well before movement execution. These observations suggest that neurons might encode these stereotyped patterns with which bats traverse 3D space. To test whether the neural representation of space reflects the flying patterns observed in behavior, Yartsev presented neural data from the hippocampus of fruit bats, recorded wirelessly while they were flying around a room with multiple foraging sites. Apart from finding a large number of three-dimensional place cells, they found that a sizable fraction of the neurons they examined showed spatial tuning for the characteristic movement patterns. Finally, elevating the floor of the room revealed that these hippocampal neurons underwent remapping but maintained their spatial selectivity for particular flight trajectories. Yartsev observed that these hippocampal cells provide a complementary encoding of space to place cells, and that combining these representations can result in much better positional accuracy in three dimensions.

André Longtin presented an inherently dynamical solution to the widespread problem of how an animal might build a spatial map of a static environment, using electrosensation in the weakly electric fish as his model. The particular problem he addressed is one of converting information that is naturally represented in egocentric coordinates as an animal explores a space, such as encounters with landmarks, to an allocentric map where the representation does not depend on the location of the individual. In these fish, a region of the thalamus (the diencephalic preglomerular complex; PG) acts as a bottleneck in this transform, receiving egocentric sensory and motor information from the optic tectum, and feeding it forward to higher areas known to be important for learning (allocentric) spatial representations. Neural recordings during object encounters reveal that while PG neurons do indeed receive information about object encounters in egocentric coordinates, individual neurons respond broadly to object encounters across the entire body, thus ignoring this spatial information. However, these neurons are strongly adapting and their responses thus reflect the time since the last object encounter. Longtin presented a model that showed how the fish might combine this temporal information with an estimate of its velocity to calculate the distance between objects in space, as needed to construct an allocentric map, and then demonstrated how a map constructed from these PG signals is consistent with the distribution of spatial and temporal behavioral errors displayed by these fish during navigation to a food target from different initial locations (Wallach, Harvey-Girard, Jun, Longtin, & Maler, 2018).

Complementary to the representation of space is the representation of head direction. Neurons in the anterodorsal thalamic nucleus and the postsubiculum of the rodent show head directional tuning, encoding an animal's heading direction with respect to the external environment. Adrien Peyrache showed that the representation of head direction is coherent across areas (thalamus and cortex) and across waking and sleep, consistent with an underlying attractor representation (also see Ila Fiete's talk, described below). Significantly, his analyses suggest that thalamocortical coordination in the head direction system is brain state independent. Peyrache next presented work further examining thalamocortical coordination, both in the head direction system and more generally. In particular, replays during SWRs during sleep have been linked to memory consolidation, but the role of the head direction system during these replays is unclear. Peyrache showed that activity in the head direction system is precisely coupled to SWRs, reliably entering a particular set of stable states right before SWR onset. This coupling was homogeneous and specific to head direction neurons, with other neurons and nuclei in the thalamus showing a different pattern of couplings (Viejo & Peyrache, 2019). Peyrache ended by linking differences in coupling properties of thalamic neurons to both the intrinsic properties of the neurons and their functional role in setting cortical state, suggesting a general organizational principle for thalamic responses.

More generally, it is an open question how neural dynamics could be spatially invariant, low dimensional and robust on the one hand and yet allow for flexible representations to be used in different contexts. Ila Fiete began her talk with the search for low-dimensional structure in neural activity, presenting work identifying the low-dimensional representation of latent variables in the brain. She introduced a new method to extract smooth geometric structures (so-called "manifolds") from neural population activity recordings using insights from topological analysis. When applied to the head direction system, this method revealed a one-dimensional ring structure within the data, with the angular position on the ring encoding the animal's heading direction, allowing for excellent unsupervised decoding of heading direction. This ring appeared to be an attractor, in that activity states that diverged from the ring flowed back on to it. Furthermore, the low-dimensional attractor was preserved across wake and REM sleep, suggesting a rigid, invariant representational space (Chaudhuri, Gerçek, Pandey, Peyrache, & Fiete, 2019). She then used correlations between neurons across different environments and sleep recording data to show that the grid cell representation was similarly low dimensional (in this case, 2D; Trettel, Trimper, Hwaun, Fiete, & Colgin, 2019) and invariant across state. In the second half of the talk, Fiete presented work showing how such low-dimensional rigid representations could nevertheless be used to flexibly encode higher-dimensional cognitive variables. This work was driven by recent observations of spatial-coding-like signals (particularly grid cell-like responses) in a number of abstract tasks and contexts, suggesting that invariant low-dimensional representations might be reused across contexts. The proposed coding scheme used multiple 2D (or low-dimensional) grid cell modules to represent a higher-dimensional variable, with each module encoding a 2D projection of the higher-dimensional variable (Klukas, Lewis, & Fiete, 2020). Fiete showed how this modular scheme has several advantages over simply building a higher-dimensional grid representation. It is efficient, with the modular structure providing a representational capacity that grows exponentially with the number of modules, providing enough coding states to encode a high-dimensional variable. Moreover, the same architecture can be reused to encode variables of any dimension without reconfiguring the whole circuit. Thus, the grid cell representation is both rigid and flexible, able to represent an arbitrary continuous high-dimensional variable and update the representation by integrating an input signal encoding changes.

3 | IMPLICIT AND EXPLICIT REPRESENTATIONS OF TIME

Dynamical systems provide an implicit code for time, yet it remains an open question to what extent temporal features of an environment or task are also explicitly represented by the brain. A number of talks interrogated the neural processing of explicitly temporal tasks to ask how neurons represent time for use as a feature in prediction, decision-making and action. Joseph Paton addressed this question in the context of a time-based decision, in which the delay between consecutive tones indicates which of two alternatives should be selected in order to obtain reward. In this task, neural activity in the (dorsal) striatum of rats organizes into a sequential representation that tiles the relevant temporal interval. This dynamic representation is flexible, in that subpopulations are "tuned" to a particular moment relative to the full interval (Mello, Soares, & Paton, 2015), and functional, in that the rate of progression through the neural sequence correlates with the likelihood of choosing the long duration option (Gouvêa et al., 2015). The implication that this sequence acts as a clock—that is, as a direct representation of elapsed time—was confirmed through a series of experiments in which thermal manipulation was used to directly modulate the progression of the striatal temporal representation, showing how cooling and heating the striatum (slowing and speeding the sequence respectively) produced a bidirectional and dose-dependent effect on choice behavior during the task. Interestingly, this manipulation did not alter the properties of movement itself, suggesting these striatal dynamics are indeed tracking time as a decision variable for action, at least somewhat independent of the representation and control of movement itself. But how then does this dynamic representation of elapsing time relate to the representation of movement, with which striatal activity has also long been associated? By controlling the spatial location of the trial initiation port, but not changing the locations of the "short" and "long" response ports, Paton and colleagues were able to demonstrate that the striatal representation of elapsed time is composed of separable movement-dependent and -independent subspaces. That is, a distinct fraction of the striatal population was insensitive to spatial location at trial initiation (forming an allocentric temporal representation unresponsive to the position of the mouse) and explained a relatively large fraction of the overall variance in these recordings. Moreover, these movement-dependent and -independent temporal representations were spatially distinct within the architecture of the striatum, suggesting a dorsal/egocentric to ventral/allocentric organization of neural representation within this brain structure.

Considering instead the representation of elapsed time subsequent to a decision, and during the anticipation of an upcoming reward, Angela Langdon presented a new model for the dynamic regulation of reward prediction by learned temporal expectations. This model proposes that the reward learning circuitry centering on the midbrain dopamine system separately learns both the amount and timing of an upcoming reward, rather than an aggregate value prediction as posited by many classic reinforcement learning models.

Looking at the impact of lesions of the ventral striatum (VS) on both dopamine prediction error signals and anticipatory behavior during an odor-guided choice task in rats showed that the temporal specificity of reward predictions, and thus temporally precise dopaminergic reward prediction errors, is critically dependent on an intact VS (Takahashi, Langdon, Niv, & Schoenbaum, 2016). However, dopamine prediction error signals to changes in the amount of reward were unaffected, consistent with a neural separation between these two dimensions of reward prediction. Furthermore, by decoding reward-related activity in neurons recorded from the VS during the same task, Langdon demonstrated that the neural representation of reward predictions in this region dynamically varies with time and is segregated into distinct subspaces that reflect the hidden block-wise structure of the task. This suggests that latent structure of multiple types, including, but not limited to, temporal structure, is learned from experience during a task, and used to dynamically regulate the neural representations that support reward-guided behaviors.

In a thorough behavioral study of learned temporal expectations in a speeded response task in humans, Matthias Grabenhorst asked how probability is represented over time in the brain? Many influential models have suggested that humans and other animals predict the timing of events by computing the hazard rate: the conditional probability that an event is about to happen, given that it has not yet occurred. Using an elegant task design, in which the probability of the target event occurring at a particular moment in time was exponentially distributed or "flipped-exponentially" distributed (i.e., events became more likely at longer delays), Grabenhorst was able to demonstrate that the distribution of reaction times to the onset of the target reflect the reciprocal of the probability density of events in time, rather than the hazard rate, and that temporal uncertainty, which is usually assumed to monotonically increase in time, was also dynamically modulated by this learned probability distribution. This result was replicated across visual, auditory, and somatosensory modalities, suggesting the reciprocal probability density of events in time is a fundamental, and domain general, computation in the brain (Grabenhorst, Michalareas, Maloney, & Poeppel, 2019).

4 | NEW DIMENSIONS IN OLFACTORY SPACE

Research probing the nature of neural representations was not confined to the physical dimensions of space and time; several speakers focused on the neural representation of abstract spaces. In particular, work presented probed the organizational and dynamic structure of sensory representations in piriform cortex, a cortical structure in mammals dedicated to olfaction. Odor representations in piriform cortex are generally thought to be highly decorrelated across distinct odors and optimized for high discriminability, which would seem to require representations that are stable across time. However, work presented at the meeting complicated this picture, pointing both to shared structure in piriform representations as well as changes over time in the representation of olfactory space.

Bob Datta began by asking whether odor representations in piriform cortex might actually reflect the shared chemical structure of the odorants? The presented results suggest they do: while responses in layer 2 of mouse piriform cortex are highly decorrelated, as is ideal for discrimination but poor for classification, responses in layer 3 are organized to reflect certain structures in odor chemical space and edit out others. This odor chemical representation is actively reshaped by recurrent local circuits in cortex, which integrate both inputs from the olfactory bulb and recent odor experience, in order to produce an odor representation finely balanced between the demands of discrimination and generalization (Pashkovski et al., 2020).

Carl Schoonover & Andrew Fink proposed that piriform cortex is not primarily involved in odor identification per se and that instead it serves as a fast learning system for encoding regularities in the olfactory environment—a short-term scratchpad for recent experiences. By testing the stability of odor responses to a panel of odors presented at varying intervals across days and weeks, they found that odor representations in long-term recordings from populations of neurons in piriform cortex are stable over short periods (i.e., days) but are profoundly reorganized on a timescale of weeks. Schoonover and Fink hypothesized that this representational drift arises from the continuous encoding of odor memory traces, causing continuous overwriting of older ones. In this view, ongoing odor experience, for example arising from sampling odors in the home cage, will overwrite old memory traces, thereby causing the representations of infrequently encountered odors to change. Two further observations support this interpretation: daily experience with a set of odorants dramatically reduced the instability of their corresponding odor representations, although this experience-induced stabilization lasted only so long as the animal continued to have regular experience with the stimuli. Critically, if daily odor presentation was halted, these odor neural representations became labile once again.

The question of which dimensions of dynamic neural representation are invariant, and which might be flexibly controlled is a fundamental one that recurred through the meeting. While space and time are stable dimensions of the external environment, one feat of representation and computation in the brain is to perform abstraction over dimensions, as in the case of an egocentric to allocentric map transformation, or in the construction of an abstract space for odor representation that generalizes over specific chemical signatures.

Which dimensions of neural representations are flexible in different brain areas, tasks and model systems, and which are invariant, yielding an irreducible feature of experience in all domains? Future efforts to answer this question will yield valuable insight into the fundamental organization of dynamic neural representations in the brain.

5 | BIOPHYSICAL AND ENVIRONMENTAL CONSTRAINTS AND OPPORTUNITIES

A natural dynamic constraint on neural activity is smoothness, which appears in two guises. First, many variables encoded in the brain (or latent variables) are continuous over time and do not change dramatically on very short timescales. Second, the tuning curves of neurons are smooth, meaning that the responses to nearby stimulus values are similar. Exploiting these constraints, Jonathan Pillow discussed methods for identifying low-dimensional latent dynamical structure from neural data (Wu, Pashkovski, Datta, & Pillow, 2018; Wu, Roy, Keeley, & Pillow, 2017). Pillow showed how to formalize these smoothness constraints by using appropriate prior distributions over both the latent variable and neural tuning curves, and combined these with a statistical model for spike generation and an inference method to estimate the latent variables. The resulting method has the appealing feature of allowing both the underlying latent dynamics and the tuning curves to be nonlinear (unlike most previous methods), and is able to extract complicated low-dimensional structure from data. Pillow showed how this method could be used to extract latent manifolds from neural responses in both the hippocampus and the piriform cortex, recovering the underlying spatial map from hippocampal responses and a 2D odor representation from piriform cortex (where standard methods like principal components analysis performed poorly).

The requirement of smoothness is also likely to shape the particular computational solutions used for a task. An example of this was provided by Mark Churchland, who pointed out that the dominant signals in motor cortex do not seem to reflect either kinematic parameters of the movement or correspond to muscle activity in a simple way. Instead, he argued that they corresponded to a dynamical system set up to drive muscular activity and that such a dynamical system required smoothness, either as a consequence of fundamental biophysical constraints on what neurons can do, or from the need to make trajectories robust to noise-induced perturbations. For a dynamical system, smoothness requires that similar patterns of activity lead to future trajectories (or outcomes) that are also similar. Churchland formalized this requirement by defining a measure of trajectory "tangling" that is high when nearby states have very different derivatives (i.e., lead to different outcomes), and predicted that neural trajectories in motor cortex should have low tangling. Indeed motor cortex shows much less tangling than either muscle responses or sensory cortex and a number of features of motor cortex responses can be predicted from the requirement of low tangling (Russo et al., 2018).

Internal biophysical constraints are not the only restrictions on neural computation; the environment imposes its own dynamic constraints on the brain as well. Among these is the fundamental learning problem of credit assignment, in which an organism must learn what events or actions in a dynamic and multi-dimensional environment produced an outcome, even if the precipitating event is no longer present. This problem is exacerbated by the stark mismatch in the timescales on which neural activity is typically observed to evolve (milliseconds to seconds) and the sometimes extended delay between events or actions and their associated outcome (which can arrive minutes, hours or even days after the precipitating event). Robert Gütig took this fundamental conundrum and showed a model in which a spiking neural network can solve the problem of spatiotemporal credit assignment when features (and their associated spiking neural responses) are fast, but feedback is relatively slow. He introduced the concept of "aggregate-label learning" to train a neural network to emit a discrete number of spikes that matches a feedback signal that is proportional to the number of times a patterned cue was present. Algorithmically, this learning rule relies on the insight that while spike counts do not provide a finite, continuous gradient along which to adjust synaptic efficacies during learning, one can substitute the voltage required to elicit the next spike (Gütig, 2016). This solution produced a spiking neural network that responded to the occurrence of various temporally patterned inputs, embedded in noise and temporally divorced from the feedback signal used for training, with the appropriate number of spikes. In a novel extension, Gütig then showed how this same learning mechanism could be used in a "self-supervised" fashion, to train a spiking network to accurately identify spatially and temporally extended regularities directly from signals from the environment without explicit feedback.

As Wolfgang Maass pointed out, the architectures and algorithms for learning in artificial neural networks far outstrip the capabilities of our models for biological learning. There exist a number of dynamical processes in the brain that likely play an important role in allowing the brain to do complex learning over time. Maass' talk focused on two powerful ideas from artificial neural networks that allow efficient temporal computing—Long Short-Term Memory networks (LSTMs) and Backpropagation through time (BPTT)—and showed how biophysical features of neurons may afford the brain similar capabilities. With a network of LSTMs, the individual units in the neural network are not biologically plausible neurons but abstract nodes possessing several regulatory gates. LSTMs have been important in the construction of artificial

networks that can easily store information over time and learn long-term dependencies. Maass showed how adding a population of neurons with adapting thresholds to a spiking network could allow a biologically constrained network to show performance nearly on par with artificial LSTMs (Bellec, Salaj, Subramoney, Legenstein, & Maass, 2018). Next, BPTT is the current gold standard for training artificial recurrent neural networks. In this algorithm, error signals are propagated backwards through time, so that a neuron's synaptic weights can be modified depending upon its effect on a (much) later outcome. Maass showed how a combination of local eligibility traces at synapses (for which there are several candidate biological mechanisms) and top-down feedback signals (such as might arise from neuromodulation) could combine to provide a neural algorithm that is similar to BPTT (Bellec et al., 2018, 2019). In summing up, Maass argued that the temporal computing capabilities of the brain dramatically improve when one accounts for slow temporal processes and urged a more thorough accounting for neuronal biophysics on longer timescales in our models of the brain.

A theme that emerged from these talks is that biophysical and environmental constraints need not be simply understood negatively, as barriers that organisms must overcome. They can also serve as priors, such as smoothness, which can be used to make data analysis techniques more specific and powerful. Furthermore, these constraints can also serve as resources for neural computation, as in the case of slow synaptic and cellular timescales, that provide robust mechanisms to ensure stability and control over dynamic activity in the brain.

6 | BRAIN STATES: AT THE INTERSECTION OF INTERNAL NEURAL DYNAMICS AND THE EXTERNAL WORLD

Neural dynamics have a very strong internal component, reflecting the role of both local circuit influences and the modulation of global brain state by different behavioral drives, such as hunger and sleep. In a fascinating demonstration of the complex relationship between behavior, neural state, and the environment, Jennifer Li and Drew Robson presented behavioral and whole-brain imaging results from freely moving larval zebrafish, in which the animals switch between hunting (exploitative) and exploratory behavioral states (Marques, Li, Schaak, Robson & Li 2020). These states shape numerous aspects of behavior, affecting locomotor strategy, hunting probability, hunting accuracy and so on, as well as both coarse and fine motor movements. Intriguingly, Li & Robson showed that these behavioral states are themselves at least partially independent of both hunger and the presence of prey: for example, even after an unsuccessful hunting bout, fish will switch into the exploration state and ignore prey. At the neural level, they found that global brain state oscillates along an axis in principal component space that reflects dorsal raphe neural activity. Identifying the neurons in the dorsal raphe that were most correlated with the transition into the exploitation (i.e., hunting) state led them to a model of zebrafish brain state alternation involving a distributed network of trigger signals that feed into a generalized trigger signal from the dorsal raphe. This dorsal raphe trigger signal initiates the transition to the exploitation state, with time dependence well modeled by a stochastic nonlinear oscillator, consisting of a short impulsive rise phase and long relaxation phase, and with the duration of the exploitation state set by the amplitude of the trigger signal. They ended by arguing that this dynamic behavioral state transition mechanism reflects an ancient and evolutionarily conserved system with parallels to serotonergic neuromodulation in *C. elegans*.

Stanislas Dehaene presented results on the triggering of conscious perception, a quite different but also seemingly global brain state. He introduced an appealingly simple method to decompose a cognitive task into a sequence of operations by testing for stability in the underlying neural representations. The method (King & Dehaene, 2014) proceeds by training a classifier (such as a support vector machine) to decode aspects of the stimulus from neural data at one moment in time, and then asks how this decoder generalizes to other points in time. If the representation is sequential, decoding performance should be high around the training point and low elsewhere (with the falloff determined by the timescale of the sequence). By contrast, if the representation is sustained, then the classifier should generalize well. Thus, how decoding generalizes across time may illuminate the temporal organization of mental representations. Dehaene applied this decoding approach to a masking task, where a picture is briefly flashed followed by a mask. The delay between the target and the mask affects whether the target is seen subliminally or consciously, and he asked what aspects of brain responses are correlated with conscious visibility. Their results showed evidence for early gradual unconscious evidence accumulation in visual areas, that seemed sequential, followed by an all-or-none transition to a distributed metastable state that is sustained over time, involves prefrontal cortex (PFC), and is correlated with conscious perception (van Vugt et al., 2018). Interestingly, the early unconscious transient could be used to partially predict whether the stimulus would be consciously seen or not. Dehaene argued that this was evidence for a "global workspace" picture of consciousness, where many segregated unconscious processors exist in parallel and the transition to consciousness reflects the global availability of a piece of information.

Many cognitive tasks require the ability to flexibly switch between different brain states in different contexts. In a set of detailed studies of frontothalamic interactions in mice during

a context-guided choice task, Michael Halassa showed how the thalamus is a critical node for the rapid reconfiguration of task-relevant dynamic representations in the PFC. In the task, a "rule" cue prior to each trial communicated which of a spatially conflicting visual or auditory cue should be chosen in order to receive a reward (i.e., "attend vision" or "attend auditory"). Rule-selective sequences in populations of PFC neurons coded for the context during the delay before the choice cues were presented (Schmitt et al., 2017). Training with an additional set of contextual cues demonstrated that these sequences do indeed represent the rule and not simply the contextual cue itself (Rikhye, Gilra, & Halassa, 2018). Interestingly, these PFC sequences are not a purely local phenomenon: bilateral optogenetic inhibition of the mediodorsal thalamus (MD) specifically during the delay period diminished rule maintenance in PFC, suggesting MD coordinates with PFC to sustain rule representation during the task. Furthermore, MD also displays context-selective activity during the delay period, and based on inactivating PFC inputs to MD, appears to be computed from PFC responses that lack context selectivity. This contextual representation then feeds back to cortex to exert two processes: amplification of context-relevant PFC inputs and suppression of context-irrelevant ones. In that manner, PFC input–output patterns are configured in a context-appropriate manner (Rikhye et al., 2018).

Moving to a larger-scale picture of the relationship between internal states, external inputs and information processing in the brain, Wolf Singer outlined a biophysically grounded general theory of cortical computation. In the first part of his talk, he contrasted two strategies by which neurons could encode relationships between features. One is a feedforward architecture, where units respond to specific conjunctions of features (a so-called "labeled-line" code). This strategy is simple but computationally inefficient, requiring a very large number of neurons to encode the possible feature combinations. Moreover, it has trouble encoding relationships between features that are separated in time and with novel combinations, and does not account for the large number of lateral and feedback projections in the cortex. A second strategy encodes relationships dynamically, exploiting the natural tendency of cortical networks to oscillate. Singer suggested that this tendency, when combined with recurrent connections endowed with Hebbian learning, allows cortical columns coding for related features to transiently synchronize, converting related features into temporal correlations and thus binding them together. The temporal patterning coordinates the timing of spikes, allowing for the operation of learning rules. Such assembly codes coexist with the feedforward labeled-line codes, with synchronous patterns better able to drive the selection of conjunctive features in further layers. In the second part of his talk, Singer moved beyond assembly formation, highlighting that the framework described above does not account for high-dimensional and asynchronous activity patterns. He suggested that cortex acts as a high-dimensional coding space that is able to store prior information about stimuli, integrate these priors with input signals, and rapidly represent the resulting computations in an easy-to-read-out format for future classification and action selection (Singer & Lazar, 2016). The lower-dimensional synchronized assemblies described above represent the readout of this Bayesian computation. Thus, resting state activity exhibits a high-dimensional correlation structure, reflecting stimulus priors stored in synaptic weight distributions. Stimuli that match prior expectations (i.e., predicted stimuli) induce low-dimensional synchronized sub-states. These readout patterns are easily separable by downstream circuits. Moreover, they persist for some time in cortical activity, exhibiting fading memory (as in reservoir computing ideas of neural computation) and allowing for the encoding of sequences. An intriguing feature of this proposed framework is that the firing rates of neurons and their finer-timescale synchrony code for different aspects of stimuli, with firing rate signaling surprise and salience (e.g., a mismatch between sensory evidence and predictions) while synchrony (perhaps in the gamma band) signifies a match with prior expectation. Singer presented evidence supporting a number of predictions of this framework and ended with a call for the development of new mathematics to analyze high-dimensional dynamically evolving activity vectors.

The talks presented investigated internal brain states across scales, systems and levels of abstraction as they ranged from zebrafish hunting to human consciousness, and from the specifics of rule-switching in the mouse to a general theory of cortical computation. A common theme that emerged is the utility of using dynamics to study internal states, with two natural questions being how a given internal state shapes finer-timescale neural dynamics and what the dynamics of inter-state transitions are. Dynamics may thus offer an integrative perspective on brain states, how they are formed and how they evolve, along with new ways to identify and define them across organisms.

7 | SEQUENCES: A GENERAL MOTIF FOR DYNAMIC NEURAL COMPUTATION

The idea that computation in the brain uses transient sequences has a long history, ranging from stereotyped motor trajectories as seen in central pattern generators (Marder & Bucher, 2001) to more abstract and flexible sequences in the context of navigation and decision-making tasks (Buzsáki & Tingley, 2018; Harvey, Coen, & Tank, 2012). Of late this idea has regained computational prominence, with sequences of neural activity observed in a variety of

model systems, brain areas and behavioral computations. A number of talks exemplified this resurgence of interest in neural sequences and both Gilles Laurent and Michael Long showed extensive lists of the variety of brain areas and task domains in which neural sequences have now been found, including hippocampal replay and preplay, bird song encoding in the HVC, olfactory neural trajectories, behavioral choice sequences in parietal and frontal cortex, basal ganglia dynamics and motor cortex dynamics. In the context of this renewed interest, several talks focused on general purpose mechanisms for sequence generation and learning, inspired by the idea that sequences of neural activity might act as a temporal scaffolding, with neural representations or motor commands inheriting temporal structure through binding to the appropriate stage in the sequence.

Sequences in the mammalian hippocampus, such as place cell trajectories, can be activated in response to behavior or can be internally generated in the absence of the corresponding behavior, either as replays during sharp wave ripples (SWRs) or as theta sequences. These observations suggest that the hippocampus may act as a general-purpose sequence generator. Claudia Clopath synthesized a number of empirical features of hippocampal sequences to construct a model showing how CA3 neurons could form abstract sequences or a "temporal backbone," which could then be used to flexibly and rapidly learn desired spiking sequences in a downstream area (such as CA1) by binding them to the appropriate moment in the abstract sequence (Nicola & Clopath, 2019). Intriguingly, rather than these abstract sequences being learned at the behavioral timescale, in the model the default timescale at which these sequences evolve is set by the intrinsic theta rhythm—a pronounced component of hippocampal neural activity—allowing them to be learned using rapid Hebbian learning. Adding a second oscillatory input with a slightly different timescale to the sequence neurons (putatively from the medial septum) caused oscillatory interference between the input and the intrinsic theta oscillations, which served to dilate the timescale of the neural representation and produce activity that varied at the appropriate behavioral timescale. The model suggested that during sharp-wave ripples (henceforth SWRs), when replay is seen, the external input drops (as is true for medial septal input) revealing the rapid intrinsic timescale.

In the rodent, when animals are active, hippocampal sequences are thought to be structured by the phase of the theta oscillations: activity at the early phase is thought to correspond to an animal's current location, and activity at the later phase thought to correspond to future plans. This patterning by the theta oscillation has a counterpart at the level of the gamma oscillation, with so-called "fast" gamma rhythms reflecting periods of high CA1 coupling to the medial entorhinal cortex, possibly important for representing current location, and "slow" gamma reflecting high CA1 coupling to CA3, possibly linked to retrieving sequences and planning trajectories.

Laura Colgin and Matthew Wilson both presented compelling evidence for the role of hippocampal sequences and frequency-based patterning in learned spatial behaviors. Laura Colgin showed results from a delayed match-to-sample spatial memory task in which a rat had to learn and remember the location of a reward across trials. She used Bayesian decoding of simultaneously recorded place cell ensembles to look at how place cell sequences developed across learning and whether these sequences were abnormal when animals failed to remember. As the animal learned the location of the reward, place cell sequences that predicted paths toward the reward developed. These sequences predicted longer paths in correct as opposed to error trials. Over the course of learning, replay of trajectories during SWRs also developed a bias to terminate at the goal location across correct but not error trials. Finally, preliminary data suggested that slow gamma power increased during the sample phase of error trials, suggesting that slow gamma rhythms may interfere with memory encoding.

Matt Wilson showed results from a navigation task where a mouse had to run along an H-shaped maze. In one arm of the maze, the mouse had to turn in an experimenter-determined direction. It then had to remember the direction in which it turned and turn in the same direction at the end of the next arm in order to receive reward. Thus, the task had a component where the mouse had to learn and encode the location of the reward and a second component where it had to retrieve the memory to get the reward. Wilson paired this task with theta-phase locked optogenetic stimulation of parvalbumin-positive interneurons neurons in area CA1. Intriguingly, activating inhibitory neurons did not lead to performance deficits. Instead, for the right combination of theta phase and task state, inhibiting CA1 lead to enhanced performance. Stimulation at the theta trough during the memory retrieval phase enhanced performance dramatically (by about 15% on a 65% baseline), and stimulation at the theta peak during memory encoding also enhanced performance. These results suggest that the hippocampus shifts between encoding and retrieval on every theta cycle, with the encoding phase potentially driven by increased coupling to entorhinal cortex and the retrieval phase displaying increased coupling to CA3. During the encoding phase of the task, the animal's current location is important but where it is going is not; the converse is true during the retrieval phase of the task. Given these two competing task demands, suppressing the task irrelevant component improves performance (Siegle & Wilson, 2014). Wilson then turned to the replay of sequential activity during SWRs. Such replay is thought to be important for reward learning, and Wilson asked if the relationship between reward learning and SWRs may be different between

sleep and quiet wakefulness. Using analyses of simultaneous recordings of neurons in the hippocampus and in the ventral tegmental area (which contains reward modulated dopamine neurons), he showed that while reward-related VTA neurons were coordinated to hippocampal replays during quiet wakefulness, the relationship was much weaker during sleep, and reward-related VTA neurons actually reduced their firing (Gomperts, Kloosterman, & Wilson, 2015). Thus, during sleep information seemed to be replayed but not reinforced in the reward system, suggesting that replays have different roles in different states.

In another model system known for producing neural sequences, Michale Fee posed the general question of how the songbird brain learns vocal behaviors consisting of a complex sequence of motor gestures. While there are reinforcement learning models that address this question, they require a good representation of the underlying state space in which learning should be performed. In the songbird, Fee argued that the premotor nucleus HVC acts as a simple sequence generation circuit, generating a sparse representation of time that provides an appropriate state space for song learning. Appropriate connections from HVC to a downstream motor area, RA, could then drive motor commands at the right times. Recordings of HVC neurons in young birds show how these sequences might emerge over the course of development, starting from a single parent sequence (protosyllable). Over time, this sequence starts to split into two, with neurons initially participating in both before becoming selective for one sequence or the other (Okubo, Mackevicius, Payne, Lynch, & Fee, 2015). This splitting continues, generating sequences selective for each syllable in the bird's adult song. Fee's proposal thus uses unsupervised learning to construct an inherently dynamical latent space that can then be used as a substrate for reinforcement learning (Mackevicius & Fee, 2018).

Michael Long began with the observation that despite widespread noise in the brain, dynamic sequences of neural activity can be surprisingly precise. This is true not just for responses to external sensory stimuli but, at least in the songbird, for internally generated sequences of activity. He considered a set of candidate models that might allow for such sequences. Learning on a set of randomly chosen initial synapses yielded activity sequences that either did not propagate through the network or were not sparse. A synfire chain model (Abeles, Prut, Bergman, & Vaadia, 1994) yielded sequences that were composed of discrete steps unlike the continuous sequences observed in the data. Finally, a "polychronization" model (Izhikevich, 2006) with a spread of synaptic delays allowed for sparse continuous sequences that propagated through the population. In a demonstration of a close theory–experiment loop, Long then looked for the source of these synaptic delays and used a combination of tracing studies, whole cell recordings and calcium imaging to argue that conduction delays from local axons showed the right distribution of timescales and are sufficient to account for the predicted delays. Thus, local conduction delays, which are often ignored in models of interacting neurons, may play an important dynamical role. A lively discussion followed, in which Wolf Singer pointed out that the myelination properties of axons change during learning (Sampaio-Baptista & Johansen-Berg, 2017), and Eve Marder noted that conduction velocity can change for bursts or due to changes in brain temperature (which can be caused by, e.g., the presence of an opposite-sex conspecific). Thus, we ended by discussing the exciting idea that conduction delays, amongt other cellular and biophysical variables that go beyond simple rate and spike dynamics, might be hitherto underexplored dynamical variables.

Taking a more abstract system-independent perspective, Giulio Bondanelli addressed coding with transient trajectories, which are closely linked to sequences. Classical population coding typically assumes that unchanging stimuli are encoded by the steady-states or time-averaged firing rates of neurons, but neural responses exhibit strong temporal dynamics even when stimuli do not change. Moreover, stimulus decoding is sometimes better during transient phases than when dynamics have converged to a fixed point (Mazor & Laurent, 2005). In a typical linear dynamical system, responses will decay away monotonically in the absence of a stimulus, making them poor candidates for coding with transients. Building on ideas from a class of linear systems called "non-normal," Bondanelli presented a framework for encoding multiple stimuli in strongly amplified transient trajectories by choosing the connectivity matrix to be the sum of appropriate low-rank pieces (Bondanelli & Ostojic, 2020), and showed that it could explain various observed features of auditory cortical neural data, such as non-monotonic transient activity at stimulus offset and better discriminability during the offset transient phase (Bondanelli, Deneux, Bathellier, & Ostojic, 2019).

While the various ideas regarding computing with neural sequences presented at the meeting were compelling, it remains an open question how general the proposed mechanisms of sequence generation and computing with transient activity are across the brain. Working at the interface of theory and experiment in the songbird vocal learning circuit or on hippocampal circuitry in spatial tasks in rodents affords a level of anatomical and neurological detail that lends credence to these theories, but potentially restricts their relevance for understanding dynamic sequence-like activity in other brain areas or model systems. Is it perhaps the case that different areas have evolved different dynamic and circuit solutions to produce similar patterns of sequential activity? These questions remain to be answered, but the convergent picture arising from hippocampal circuits in mice and vocal song circuits in songbirds provides a promising place to start.

8 | SOCIAL INTERACTION AND COMMUNICATION: A DYNAMIC LOOP

In a more naturalistic setting, sensory processing and behavioral control must proceed within the context of a dynamic interaction with social partners as well as the environment. A number of talks emphasized this important aspect of neural computation, focusing on communication and social interaction in songbirds, *Drosophila* and mice. In an elegant analogy to the development of language specificity during human development, Sarah Woolley presented work showing how the neural representation of song vocalizations varies across the auditory processing hierarchy of juvenile songbirds. In thalamoreceptive layers of auditory cortex the neural representation of a particular song is highly similar across individual birds, but differs greatly across birds in the deep output and secondary layers. Like language learning, the development of this song representation depends critically on the experience of hearing the tutor's song. Perhaps surprisingly, juveniles exposed to a cross-species tutor, and thus a song with different characteristic auditory features ("syllables"), learned to produce a song with the tutor's species-specific syllables remarkably well. Woolley then demonstrated the highly adaptive nature of the neural representation of song in the auditory cortex, which failed to display selective neural responses for own-species vocalizations in higher auditory areas in the cross-species tutored birds. Rather, auditory neural activity in birds that acquired a cross-species song showed selectivity for the acoustic features of the song "language" they had learned through experience (Moore & Woolley, 2019).

Successful communication occurs in a dynamic environment in which the brain must continuously process incoming information and modulate behavioral output online as the interactive setting evolves. Mala Murthy presented her lab's work on dynamic communication, using courtship in *Drosophila* as a model system. Successful courtship is promoted by the male production of a song—produced by wing vibrations—while the female arbitrates the mating decision. Male song structure and intensity depend on the interactions between the male and the female; rather than repetitively executing stereotyped wing-vibration sequences, the male continuously modulates song production according to the social interaction (Coen et al., 2014). What then are the auditory features of the male song that female (and male) brains respond to? Murthy identified pC2 neurons in the Drosophila brain of both sexes that act as auditory pulse feature detectors, demonstrating a common brain response to this property of the male song. However, the relationship between this neural response and behavior diverged for females and males: females slowed down when the pC2 neurons responded to a particular pulse rate, while males sped up and also sang. This dynamic feedback loop between the courting individuals through their common pC2 neurons allows both males and females to detect and modulate both locomotor activity and song production in an inherently social behavior (Deutsch, Clemens, Thiberge, Guan, & Murthy, 2019). Additionally, Murthy presented work mapping auditory activity throughout the entire central brain of *Drosophila* (Pacheco, Thiberge, Pnevmatikakis, & Murthy, 2019). The discovery of widespread and diverse auditory responses in nearly every brain region of both males and females suggests that courtship song has a strong modulatory impact on a variety of sensory and motor processes. Murthy concluded with a call for more sophisticated tools to map behavior at the highest resolution (Calhoun, Pillow, & Murthy, 2019; Pereira et al., 2019), in order to more precisely map the role of internal states on these dynamic acoustic courtship behaviors.

The social consequences of many ethologically relevant behaviors raise the question of whether socially relevant cues are represented differently in the brain from those that are non-social, and, as a consequence, whether distinctly social neural representations are implicated in impairments of social behaviors. Tal Tamir showed neurons in the PFC of mice that responded preferentially to socially relevant olfactory cues, such as male or female odors, over non-social (food) olfactory cues. At the population level, neural activity in PFC showed a similar pattern at baseline, but followed distinct low-dimensional trajectories for social and non-social stimuli both during and after stimulus presentation. Interestingly, in neural populations recorded from Cntnap2 mice (a genetic model of autism), the separation between social and non-social representations in the PFC was greatly reduced. In a demonstration of brain dynamics over a relatively longer timescale, Tamir then showed that the separation between social and non-social neural representations in the PFC increased with experience over consecutive days in the wild-type mice, a refinement of dynamic neural representation that failed to occur in the autism-model mice (Levy et al., 2019).

Together, these talks highlighted the importance of the social dimension of neural representation, a behavioral setting that remains relatively underexplored. Insight into neural computation will naturally benefit from a richer and more detailed understanding of the dynamic, interactive setting in which much of behavior takes place. Dynamical systems approaches to multi-agent systems suggest new frameworks for integrating empirical data from multiple brains and behavior and provide paradigmatic examples from artificial intelligence of competitive (such as generative adversarial networks) or cooperative computation (such as distributed decision-making systems) in an interactive social setting.

9 | THE DEVELOPMENT OF NEURAL DYNAMICS: LEARNING AND ORGANIZATION ACROSS MULTIPLE TIMESCALES

While the focus on dynamic neural activity has typically been at the timescale of seconds and minutes during a task, a major piece of the puzzle is uncovering how the brain is dynamically reorganized across the days, months, and years of development to support neural computation. Several speakers took this interesting perspective, probing how patterns of brain activity evolve through early life, when brain circuits are both substantially and relatively rapidly reshaped as an animal or human acquires new experiences and novel abilities. Julijana Gjorgjieva asked how spontaneous brain activity might refine functionally specialized neural circuits in early development, even before any sensory experience has been acquired. She presented a biophysically realistic model of synaptic plasticity to show how spontaneous activity can establish remarkably precise fine-scale structure in the spatial organization of dendritic synapses (Kirchner & Gjorgjieva, 2019). Using a burst-timing-dependent plasticity rule based on the action of neurotrophic factors in which postsynaptic calcium spread induces spatial competition, this model demonstrates how functional synaptic clustering emerges in response to spontaneous waves of activity in the developing retina. Gjorgjieva proposed that the critical ingredients of spontaneous activity and synaptic plasticity are already present in the early developing brain, allowing networks of neurons to wire themselves to the finely structured circuits observed in adulthood.

At the opposite end of the spatial scale, Shruti Naik showed how the macroscopic brain signal in scalp EEG of very young infants evoked by unfamiliar face stimuli evolves through development as they acquire sophistication in their ability to recognize faces. While averaging single trial responses reveals stereotyped features of the face-evoked event-related potential (ERP) by 12 weeks of age (tracking a developmental milestone in early visual areas), individual trial responses are highly variable, raising the question of how such dynamic variability in brain activity can support reliable face recognition. By quantifying this across-trial variability in face-evoked EEG activity of individuals between 2 and 6 months of age, Naik demonstrated how the distribution of the latency of single-trial ERP-like events becomes gradually more concentrated around the time of the mean ERP component—a quenching of variability around these stereotyped patterns of brain activity—in accordance with developmental age. This suggests that the stabilization of single-trial dynamics around the large-scale activity patterns typically measured by the grand-average ERP is a critical stage of the maturating infant brain.

10 | EVOLVING FLEXIBLE CONTROL OF THE DYNAMIC BRAIN

Over an even longer timescale, organisms have evolved a complex set of mechanisms by which to control the dynamic patterns of neural activity that support the various behaviors they perform. A major theme was the control of dynamic neural activity, asking the fundamental question of how complex patterns of activity are reliably reproduced by an organism despite sometimes wild variation in sensory input and the broader environment.

Eve Marder asked how finely-tuned do the parameters that control intrinsic properties and synaptic efficacies of neurons need to be for "good enough" circuit activity? In other words, how variable can brains be and still produce successful behavior? She focused on the stomatogastric ganglion (STG) neurons of wild-type crabs (i.e., crabs that have evolved to be successful in their natural habitat, not a laboratory) responsible for producing the pyloric rhythm in this organism. She pointed to the variability in this three-neuron circuit across individuals: cell morphology is highly variable and wiring is inefficient and "tortuous." And yet the pyloric rhythm is highly stereotyped despite a two- to sixfold variation in circuit parameters for different individuals, suggesting degenerate mechanisms by which this macroscopic dynamic pattern can be achieved in the circuit. Indeed, by generating families of models with different conductance densities, she showed how distinct circuit mechanisms are able to achieve largely identical oscillatory patterns of activity (Gutierrez, O'Leary, & Marder, 2013; Prinz, Bucher, & Marder, 2004). These distinct model circuits reveal their difference in their response to perturbation: a prediction borne out in the individual response of STG circuits to perturbations arising from temperature, pH and chemical manipulations. For instance, increasing temperature will ultimately disrupt the pyloric rhythm for all individuals, but the temperature at which this occurs, and the dynamic patterns of activity produced as the pyloric rhythm fails were highly variable across individual STG preparations. This intriguing demonstration of dynamic degeneracy in what is a relatively small and well-characterized circuit raises the interesting and important question of how neural circuits maintain stability and resistance to perturbation in an environment in which unexpected changes in conditions are bound to occur.

Despite the probable degeneracy in specific brain circuits and mechanism, Gilles Laurent pointed to the remarkable prevalence of sequential neural activity across phylogenetically distinct organisms operating in hugely different environments. He proposed that this prevalence owes directly to the fact that the physical and biological world is dominated by correlations over many timescales; brains

have adapted over evolution toward common solutions for controlling behavior in such a world. He then traced three cases of transient neural dynamics across distinct model systems to highlight a fundamental motif of stereotyped and low-dimensional sequential trajectories that move away from, and then return to, a resting state (i.e., a fixed point in the state space). In the olfactory system of the arthropod, he showed how the population response of projection neurons of the antennal lobe to an odor stimulus traces out a dynamic trajectory that is highly stereotyped, and evolves on a low-dimensional manifold toward an odor-specific fixed point (Mazor & Laurent, 2005; Wehr & Laurent, 1996). In natural conditions, however, odors are brief, and this transient activity does not evolve rapidly enough to reach this fixed point. Interestingly, downstream Kenyon cells only decode the odor during the transient phase of the projection neuron population response, confirming that this dynamic pattern of activity is indeed the critical representation of the odor. Asking next how the dynamic response of a population of neurons might be controlled, Laurent introduced the chromatophore system of the cuttlefish. A pattern of chromatophores provide camouflage for the animal, and each are controlled by muscles to expand and contract, blanching the macroscopic pattern after a threat from the environment. By tracking the state of tens of thousands of chromatophores following a blanching event, Laurent and colleagues demonstrated that the global chromatophore state follows a stereotyped trajectory away from, and then back to the resting state. This stereotypy arises despite the high dimensionality of the pattern itself, consistent with the existence of a low-dimensional motor control representation that orchestrates this enormously complex spatiotemporal pattern (Reiter et al., 2018). Finally, he presented neural data from the dorsal cortex of turtles in response to the electrical stimulation of single neurons, demonstrating surprisingly reliable sequences of activity that propagate through tens of neurons in the local cortical circuit after even single spikes are elicited in individual pyramidal neurons (Hemberger, Shein-Idelson, Pammer, & Laurent, 2019). That this is possible suggests a topology of excitation that effectively primes certain patterns of sequential activity to flow through the cortical circuit.

11 | SUMMARY

Ultimately, the diversity of research presented at this Brain Conference revealed that we are at an exciting juncture in the study of the brain: theoretical and empirical progress has afforded a view of the building blocks of dynamic computation in neural systems. The research presented provided a compelling picture of how transient neural activity can represent space, time and various features of a task, be it an abstract experimental manipulation or a naturalistic feature in a social setting such as communication or mating. The frontier is now to press forward in our understanding of how dynamic computation in the brain is flexibly controlled, and what mechanisms allow the reliable propagation of transient patterns of activity across different circuit architectures and in different environmental conditions. One of the key concepts for debate that emerged was the tension between invariance and flexibility: neural dynamics are naturally constrained by specific local connections, brain architecture, and biophysical mechanisms, yet time and again throughout the meeting, researchers presented patterns of neural population activity that displayed surprisingly similar dynamics despite different model systems, different stimuli, and different tasks and global behavioral states. Continuing the search to find the fundamental mechanisms by which dynamic neural activity is flexibly controlled to produce these diverse behaviors will be an exciting next chapter for the field.

ACKNOWLEDGEMENTS

We thank Ila Fiete and Gilles Laurent for inviting us to participate in the meeting and write this meeting report, as well as all the speakers for helpful suggestions on the text. A.J.L. was supported by awards R01DA042065 and R01DA050647 from the National Institute on Drug Abuse.

CONFLICTS OF INTEREST

The authors declare no competing interests.

PEER REVIEW

The peer review history for this article is available at https://publons.com/publon/10.1111/ejn.14963

DATA AVAILABILITY STATEMENT

No primary data.

Angela J. Langdon[1]
Rishidev Chaudhuri[2]

[1]*Princeton Neuroscience Institute & Department of Psychology, Princeton University, Princeton, NJ 08544, USA*
[2]*Center for Neuroscience, Department of Mathematics, Department of Neurobiology, Physiology & Behavior, University of California, Davis, Davis, CA 95616, USA*

Correspondence
Angela J. Langdon, Princeton Neuroscience Institute & Department of Psychology, Princeton University, Princeton, NJ, USA.
Email: alangdon@princeton.edu

Rishidev Chaudhuri, Center for Neuroscience, Department of Mathematics and Department of Neurobiology, Physiology & Behavior, University of California, Davis, Davis, CA, USA.
Email: rchaudhuri@ucdavis.edu

ORCID

Angela J. Langdon https://orcid.org/0000-0003-4742-6976
Rishidev Chaudhuri https://orcid.org/0000-0003-2301-3261

REFERENCES

Abeles, M., Prut, Y., Bergman, H., & Vaadia, E. (1994). Synchronization in neuronal transmission and its importance for information processing. *Progress in Brain Research*, *102*, 395–404.

Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., & Maass, W. (2018). Long short-term memory and learning-to-learn in networks of spiking neurons. *Advances in Neural Information Processing Systems*, *31*, 787–797.

Bellec G., Scherr F., Subramoney A., Hajek E., Salaj D., Legenstein R., Maass W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, *11*, (1), http://doi.org/10.1038/s41467-020-17236-y.

Bondanelli, G., Deneux, T., Bathellier, B., & Ostojic, S. (2019). Population coding and network dynamics during OFF responses in auditory cortex. *BioRxiv*, 810655.

Bondanelli, G., & Ostojic, S. (2020). Coding with transient trajectories in recurrent neural networks. *PLOS Computational Biology*, *16*(2), e1007655. https://doi.org/10.1371/journal.pcbi.1007655

Buzsáki, G., & Tingley, D. (2018). Space and time: The hippocampus as a sequence generator. *Trends in Cognitive Sciences*, *22*, 853–869. https://doi.org/10.1016/j.tics.2018.07.006

Calhoun, A. J., Pillow, J. W., & Murthy, M. (2019). Unsupervised identification of the internal states that shape natural behavior. *Nature Neuroscience*, *22*, 2040–2049. https://doi.org/10.1038/s41593-019-0533-x

Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., & Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, *22*, 1512–1520. https://doi.org/10.1038/s41593-019-0460-x

Coen, P., Clemens, J., Weinstein, A. J., Pacheco, D. A., Deng, Y., & Murthy, M. (2014). Dynamic sensory cues shape song structure in Drosophila. *Nature*, *507*, 233–237. https://doi.org/10.1038/nature13131

Deutsch, D., Clemens, J., Thiberge, S. Y., Guan, G., & Murthy, M. (2019). Shared song detector neurons in drosophila male and female brains drive sex-specific behaviors. *Current Biology*, *29*, 3200–3215.e5. https://doi.org/10.1016/j.cub.2019.08.008

Gardner, R. J., Lu, L., Wernle, T., Moser, M.-B., & Moser, E. I. (2019). Correlation structure of grid cells is preserved during sleep. *Nature Neuroscience*, *22*, 598–608. https://doi.org/10.1038/s41593-019-0360-0

Gomperts, S. N., Kloosterman, F., & Wilson, M. A. (2015). VTA neurons coordinate with the hippocampal reactivation of spatial experience. *Elife*, *4*, e05360. https://doi.org/10.7554/eLife.05360

Gouvêa, T. S., Monteiro, T., Motiwala, A., Soares, S., Machens, C., & Paton, J. J. (2015). Striatal dynamics explain duration judgments. *Elife*, *4*, e11386.

Grabenhorst, M., Michalareas, G., Maloney, L. T., & Poeppel, D. (2019). The anticipation of events in time. *Nature Communications*, *10*, 5802. https://doi.org/10.1038/s41467-019-13849-0

Gutierrez, G. J., O'Leary, T., & Marder, E. (2013). Multiple mechanisms switch an electrically coupled, synaptically inhibited neuron between competing rhythmic oscillators. *Neuron*, *77*, 845–858. https://doi.org/10.1016/j.neuron.2013.01.016

Gütig, R. (2016). Spiking neurons can discover predictive features by aggregate-label learning. *Science*, *351*, aab4113. https://doi.org/10.1126/science.aab4113

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*, 801–806. https://doi.org/10.1038/nature03721

Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, *484*, 62–68. https://doi.org/10.1038/nature10918

Hemberger, M., Shein-Idelson, M., Pammer, L., & Laurent, G. (2019). Reliable sequential activation of neural assemblies by single pyramidal cells in a three-layered cortex. *Neuron*, *104*, 353–369.e5. https://doi.org/10.1016/j.neuron.2019.07.017

Izhikevich, E. M. (2006). Polychronization: Computation with spikes. *Neural Computation*, *18*, 245–282. https://doi.org/10.1162/089976606775093882

King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. *Trends in Cognitive Sciences*, *18*, 203–210. https://doi.org/10.1016/j.tics.2014.01.002

Kirchner, J. H., & Gjorgjieva, J. (2019). A unifying framework for synaptic organization on cortical dendrites. *BioRxiv*, 771907.

Klukas M., Lewis M., Fiete I. (2020). Efficient and flexible representation of higher-dimensional cognitive variables with grid cells. *PLOS Computational Biology*, *16*, (4), e1007796http://doi.org/10.1371/journal.pcbi.1007796.

Levy, D. R., Tamir, T., Kaufman, M., Parabucki, A., Weissbrod, A., Schneidman, E., & Yizhar, O. (2019). Dynamics of social representation in the mouse prefrontal cortex. *Nature Neuroscience*, *22*, 2013–2022. https://doi.org/10.1038/s41593-019-0531-z

Mackevicius, E. L., & Fee, M. S. (2018). Building a state space for song learning. *Current Opinion in Neurobiology*, *49*, 59–68. https://doi.org/10.1016/j.conb.2017.12.001

Marder, E., & Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Current Biology*, *11*, R986–R996. https://doi.org/10.1016/S0960-9822(01)00581-4

Marques J.C., Li M., Schaak D., Robson D.N., Li J.M. (2020). Internal state dynamics shape brainwide activity and foraging behaviour. *Nature*, *577*, (7789), 239–243. http://doi.org/10.1038/s41586-019-1858-z.

Mazor, O., & Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*, *48*, 661–673. https://doi.org/10.1016/j.neuron.2005.09.032

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the "cognitive map". *Nature Reviews Neuroscience*, *7*, 663–678. https://doi.org/10.1038/nrn1932

Mello, G. B. M., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Current Biology*, *25*, 1113–1122. https://doi.org/10.1016/j.cub.2015.02.036

Moore, J. M., & Woolley, S. M. N. (2019). Emergent tuning for learned vocalizations in auditory cortex. *Nature Neuroscience*, *22*, 1469–1476. https://doi.org/10.1038/s41593-019-0458-4

Moser, E. I., Moser, M.-B., & McNaughton, B. L. (2017). Spatial representation in the hippocampal formation: A history. *Nature Neuroscience*, *20*, 1448–1464. https://doi.org/10.1038/nn.4653

Nicola, W., & Clopath, C. (2019). A diversity of interneurons and Hebbian plasticity facilitate rapid compressible learning in the hippocampus. *Nature Neuroscience*, *22*, 1168–1181. https://doi.org/10.1038/s41593-019-0415-2

Okubo, T. S., Mackevicius, E. L., Payne, H. L., Lynch, G. F., & Fee, M. S. (2015). Growth and splitting of neural sequences in songbird vocal development. *Nature*, *528*, 352–357. https://doi.org/10.1038/nature15741

Pacheco, D. A., Thiberge, S. Y., Pnevmatikakis, E., & Murthy, M. (2019). Auditory activity is diverse and widespread throughout the central brain of *Drosophila*. *BioRxiv*, 709519.

Pashkovski S.L., Iurilli G., Brann D., Chicharro D., Drummey K., Franks K.M., … Datta S.R. (2020). Structure and flexibility in cortical representations of odour space. *Nature*, *583*, (7815), 253–258. http://doi.org/10.1038/s41586-020-2451-1.

Pereira, T. D., Aldarondo, D. E., Willmore, L., Kislin, M., Wang, S.-S.-H., Murthy, M., & Shaevitz, J. W. (2019). Fast animal pose estimation using deep neural networks. *Nature Methods*, *16*, 117–125. https://doi.org/10.1038/s41592-018-0234-5

Prinz, A. A., Bucher, D., & Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, *7*, 1345–1352. https://doi.org/10.1038/nn1352

Reiter, S., Hülsdunk, P., Woo, T., Lauterbach, M. A., Eberle, J. S., Akay, L. A., … Laurent, G. (2018). Elucidating the control and development of skin patterning in cuttlefish. *Nature*, *562*, 361–366. https://doi.org/10.1038/s41586-018-0591-3

Rikhye, R. V., Gilra, A., & Halassa, M. M. (2018). Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nature Neuroscience*, *21*, 1753–1763. https://doi.org/10.1038/s41593-018-0269-z

Russo, A. A., Bittner, S. R., Perkins, S. M., Seely, J. S., London, B. M., Lara, A. H., … Churchland, M. M. (2018). Motor cortex embeds muscle-like commands in an untangled population response. *Neuron*, *97*, 953–966.e8. https://doi.org/10.1016/j.neuron.2018.01.004

Sampaio-Baptista, C., & Johansen-Berg, H. (2017). White matter plasticity in the adult brain. *Neuron*, *96*, 1239–1251. https://doi.org/10.1016/j.neuron.2017.11.026

Schmitt, L. I., Wimmer, R. D., Nakajima, M., Happ, M., Mofakham, S., & Halassa, M. M. (2017). Thalamic amplification of cortical connectivity sustains attentional control. *Nature*, *545*, 219–223. https://doi.org/10.1038/nature22073

Siegle, J. H., & Wilson, M. A. (2014). Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus. *Elife*, *3*, e03061.

Singer, W., & Lazar, A. (2016). Does the cerebral cortex exploit high-dimensional, non-linear dynamics for information processing? *Frontiers in Computational Neuroscience*, *10*, 99. https://doi.org/10.3389/fncom.2016.00099

Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, *492*, 72–78. https://doi.org/10.1038/nature11649

Strange, B. A., Witter, M. P., Lein, E. S., & Moser, E. I. (2014). Functional organization of the hippocampal longitudinal axis. *Nature Reviews Neuroscience*, *15*, 655–669. https://doi.org/10.1038/nrn3785

Takahashi, Y. K., Langdon, A. J., Niv, Y., & Schoenbaum, G. (2016). Temporal specificity of reward prediction errors signaled by putative dopamine neurons in rat VTA depends on ventral striatum. *Neuron*, *91*, 182–193. https://doi.org/10.1016/j.neuron.2016.05.015

Trettel, S. G., Trimper, J. B., Hwaun, E., Fiete, I. R., & Colgin, L. L. (2019). Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nature Neuroscience*, *22*, 609–617. https://doi.org/10.1038/s41593-019-0359-6

van Vugt, B., Dagnino, B., Vartak, D., Safaai, H., Panzeri, S., Dehaene, S., & Roelfsema, P. R. (2018). The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. *Science*, *360*, 537. https://doi.org/10.1126/science.aar7186

Viejo G., Peyrache A. (2020). Precise coupling of the thalamic head-direction system to hippocampal ripples. *Nature Communications*, *11*, (1), http://doi.org/10.1038/s41467-020-15842-4.

Wallach A., Harvey-Girard E., Jun J.J., Longtin A., & Maler L. (2018). A time-stamp mechanism may provide temporal information necessary for egocentric to allocentric spatial transformations. *eLife*, *7*, http://dx.doi.org/10.7554/elife.36769.

Wehr, M., & Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature*, *384*, 162–166. https://doi.org/10.1038/384162a0

Wu, A., Pashkovski, S., Datta, S. R., & Pillow, J. W. (2018). Learning a latent manifold of odor representations from neural responses in piriform cortex. *Advances in Neural Information Processing Systems*, *31*, 5378–5388.

Wu, A., Roy, N. A., Keeley, S., & Pillow, J. W. (2017). Gaussian process based nonlinear latent structure discovery in multivariate spike train data. *Advances in Neural Information Processing Systems*, *30*, 3496–3505.

How to cite this article: Langdon AJ, Chaudhuri R. An evolving perspective on the dynamic brain: Notes from the Brain Conference on *Dynamics of the brain: Temporal aspects of computation*. *Eur J Neurosci*. 2021;53:3511–3524. https://doi.org/10.1111/ejn.14963