Published in final edited form as:

Curr Opin Neurobiol. 2016 October ; 40: 14–22. doi:10.1016/j.conb.2016.05.005.

Itinerancy between attractor states in neural systems

Paul Miller

Volen National Center for Complex Systems, Brandeis University, Waltham MA 02454-9110, U. S. A

Abstract

Converging evidence from neural, perceptual and simulated data suggests that discrete attractor states form within neural circuits through learning and development. External stimuli may bias neural activity to one attractor state or cause activity to transition between several discrete states. Evidence for such transitions, whose timing can vary across trials, is best accrued through analyses that avoid any trial-averaging of data. One such method, hidden Markov modeling, has been effective in this context, revealing state transitions in many neural circuits during many tasks. Concurrently, modeling efforts have revealed computational benefits of stimulus processing via transitions between attractor states. This review describes the current state of the field, with comments on how its perceived limitations have been addressed.

Graphical abstract

Introduction

Attractor states come in many forms, the simplest of which is a point attractor. For example, the firing rate of a neuron responding to a constant stimulus is likely to approach a constant value after a transient period. The constant final rate is a point attractor, because if the rate were temporarily altered (say by a brief pulse of excitation) it would return to its prior constant value. Importantly, at a more detailed level of description—that of the membrane potential, rather than the firing rate—the point attractor is revealed to be a cyclic attractor as the neuron spikes at regular intervals following a periodic voltage trace. If perturbed from the periodic trace, the membrane potential would return to the same orbit, albeit with a phase

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Conflict of interest statement

Nothing declared.

shift. So the notion of a point attractor can depend on the level of description: a point attractor at one level can be comprised of other types of attractor at a more detailed level. In this review we focus on systems possessing many point attractors when described at the level of neural firing rates. Systems with multiple point attractors are important because the system's state—*i. e.* which point attractor it dwells at—becomes history-dependent, so the system can represent memories and respond to stimuli in a sequence-specific manner.

Models of neural circuits containing multiple point attractor states have a long history in theoretical neuroscience [1–8]. These circuits can learn discrete input patterns and later reproduce them upon presentation of altered or incomplete versions of the original pattern. In these circuits different sets of neurons—whose coordinated activities comprise the neural representations of the original stimuli—can maintain their activity in the absence of direct external stimulus. Evidence from many short-term memory tasks indicates that neural circuits are indeed capable of producing such stimulus representations that persist in the absence of external stimuli [9,10].

When neural activity transitions between different attractor states, the term attractor-state itinerancy is used. The transitions can be prompted either by transient external stimuli or due to internal processes. A trivial example of the latter is loss of persistent activity representing short-term memory during a delay due to internal noise fluctuations [10–13].

A more interesting example is the bistable percept, models of which invoke negative feedback such as firing rate adaptation to weaken an active cluster of neurons until an inactive cluster takes over (Fig. 1) [14,15]. Noise causes the timing of transitions between states to be variable, leading to an important observation: the average firing rate across trials, being a weighted mean of the activities corresponding to the two percepts, would not reveal any switching between percepts, which is the most relevant aspect of the activity. Trial-specific analyses, such as hidden Markov modeling (HMM) [16–19], are essential for uncovering such switches between states when the transition times are variable.

Evidence for multiple attractor states and itinerancy

Existence of discrete attractor states

Both perceptual reports and electrophysiological measurements provide evidence for multiple attractor states [20–24]. By its definition, a point attractor requires that in the presence of a fixed stimulus, if neural activity is perturbed a little, it returns or is attracted back to its prior pattern. However, given the difficulty of such perturbations, the presence of attractors has been inferred by responses to gradual variation of the stimulus [20]. As a stimulus is parametrically altered, the neural response or percept changes by small degrees until a transition point is reached, at which a small change in stimulus causes a large change in neural response and percept as the prior attractor state destabilizes and a new one is reached [21]. The observation of hysteresis, meaning that the transition point depends on recent history, strongly suggests the presence of discrete attractor states [22]. This type of evidence for attractor states has been found for simple auditory percepts [23], for face recognition [24], and for spatial environments [22] across a number of species.

Evidence from neural activity for state transitions in the absence of stimulus change

HMM is a method for ensemble data analysis that can reveal any reliable structure underlying distinct states of neural activity as well as any pattern in the order that states are visited by a network. A description of the operation of HMM is presented in Figure 2. Hidden Markov models of multiple simultaneously recorded *in vivo* spike trains have revealed transitions between relatively stable discrete activity states in premotor and prefrontal cortex of monkey performing a spatial short-term memory task [17,25,26] or freely viewing a scene [27], and in gustatory cortex of rats processing tastants [28]. The timing of such state transitions both correlates with ongoing behavior [27,29] and is predictive of an ensuing response [30].

Single spike trains from monkeys have been used to provide evidence for state transitions that do not appear in trial-averaged data. During motor planning, within-trial spike counts taken from anterior cingulate cortex over a time-window near the mid-point of a delay yielded a bimodal distribution across trials [31]. A more sophisticated analysis of lateral intraparietal cortex data taken during a decision-making task used a latent variable model of single spike trains to uncover state transitions [32].

Switching between attractor states is necessary for an internally generated switch between motor output patterns [33]. Transitions between discrete motor states have been revealed by HMM for the crawling of *c. elegans* [34], the sequences of syllables in bird songs [35], and lever-pressing rates in rats [36]. Since the observed motor patterns are produced by underlying neural activity, these studies lend indirect evidence to transitions between attractor states within neural circuits.

Since attractor states are produced by the internal structure of a neural circuit it is likely that spontaneous activity, in the absence of a stimulus to drive activation of a particular state, would randomly sample the network's activity states [37–41]. This feature combined with synaptic plasticity, can stabilize discrete attractor states [42]. In support of the hypothesis that spontaneous activity samples attractor states, multi-electrode data from awake ferrets shows that spontaneous activity *in vivo* approaches activity evoked by natural stimuli over the course of development [43], with statistics more compatible with wandering between multiple states than with fluctuations about a single state [44]. Recently, direct evidence from HMM demonstrates that spontaneous activity in rodents can exhibit rapid transitions between discrete states [45].

Whole-brain data from fMRI or MEG [46,47], and large-scale simulations of the brain based on tractography studies [47,48], provide evidence for rapid transitions between a few discrete functional connectivity states in the resting state of humans. These transitions and the underlying discrete states would not be revealed by time-averaged data [47–49].

State transitions without point attractors

The wealth of evidence for state transitions does not necessarily imply stochastic switching between point attractor states. Indeed, in EEG measurements, state transitions are observed as changes in the amplitude, phase, or frequency of oscillatory patterns [50,51]. While a transition between point-attractor states in one circuit could trigger a change in oscillation in

another (via *e. g.* neuromodulator release [52]), it is also possible for a circuit possessing multiple oscillatory (cyclic) attractor states [53,54] to transition stochastically between the states.

Similar state transitions can also arise in deterministic systems, either through chaos [55,56] in which case states have an oscillatory structure [57], or through heteroclinic sequences [58–61] in which case states are weakly unstable fixed points. When comparing these deterministic dynamical models with point-attractor state itinerancy, both the statistics of state-durations across trials and the dynamics of neural activity within a state should differ. Extensions of HMM can extract such internal state dynamics [62] to help determine which dynamical system is present.

Computational models with attractor-state itinerancy

Network multistability can arise from distinct bistable units or as an emergent property of connected monostable units

Multistability can arise from recurrent feedback in the circuit when the firing rate of single neurons varies continuously with input (*i. e.*, when single neurons taken in isolation show neither jumps in firing rate nor hysteresis so are monostable). In such cases, the representations of stimuli are distributed, meaning each pattern requires coordinated firing of many neurons and each neuron can take part in many patterns, typically with independent groups of other cells [6,8]. Chaotic networks comprised of standard monostable single units can be trained to produce a discrete number of stable network states whose activity can be switched by transient inputs [33,63,64]. Randomly connected oscillators can also produce multiple stable states, distinguished by phase relationships between oscillators, as well as amplitude and frequency of oscillation [65].

In another approach, several groups have focused on the behavior of clusters of neurons, also called assemblies [66], within which preferentially strong connections [67–69] generate a single bistable unit. Such bistable units could arise from a Hebbian strengthening of connections among coactive excitatory cells, eventually leading to similarly responsive cells being preferentially connected to each other [42,70], as observed empirically [71].

In models, the level of cross-inhibition between bistable units determines the number of coactive units. If only a single unit can be active due to strong cross-inhibition, *i. e.*, in a winner-takes-all circuit, information is stored very sparsely and the number of stable states equals the number of units. If multiple units can be active at once due to relatively weak cross-inhibition, there is a combinatorial explosion in the number of potentially stable states, suggesting great computational benefits. The information capacity of such circuits is demonstrated in their ability to dissociate the number, duration and amplitude of repeated identical stimuli [72] while the richness of their dynamics is revealed in their ability to produce chaotic transients while also containing many stable states [73].

Attractor-state transition models for decision-making

Decision-making models with attractor-states representing the final choice yield many of the neural dynamics observed in perceptual choice tasks [74–77]. In a variant of such models,

the initial “undecided” state can itself be an attractor state, in which case decisions are made by noise-induced transitions [78,79].

Evidence for a delayed rapid transition from one discrete attractor state to another during decision making has been accrued during taste processing (where the decision is to expel or ingest a tastant) [28,30], in a somatosensory task [80], and in a detection of motion direction task [32]. However, an alternative analysis of the latter task revealed ramping as the typical dynamics during stimulus acquisition, with rapid transitions corresponding only to changes of mind [81].

Surprisingly, noise-induced transitions between point attractors can outperform the optimal method of integration for deciding between noisy stimuli [79], if biological constraints such as the limited range of neural firing rates and the presence of within-circuit noise are taken into account. In such circuits a particularly strong fluctuation can kick the activity out of an “undecided” state into one of the decision states, such a transition being more likely to the decision state with greater underlying activation. The advantage of such attractor-based methods is that internal noise, which limits the usefulness of perfect integration [82] (it accumulates as Brownian motion in standard integration models of decision making), can actually boost performance by attractor transitions, generating a decision within an appropriate time window, in a method resembling stochastic resonance [78].

Bayesian sampling with multiple attractor states

In tasks with a constant stimulus, noise-driven transitions between attractor states can account for the distribution of durations of individual competing percepts [14,15,83]. Moreover, point attractor models of perceptual rivalry are compatible with the hypothesis that percepts represent the outcome of a Bayesian sampling process if the fraction of time spent in an attractor state is equal to the posterior probability of that percept corresponding to the stimulus [84]. Importantly, synaptic plasticity in circuits with multiple attractor states can provide a basis for Bayesian computations [85–87].

Recall of sequences via attractor-state itinerancy

Models of word recall, within which each word is stored as a particular attractor state, have explained statistical features of recall patterns from word-lists [88–91]. Within these models, transitions are facilitated by adaptation—as proposed by others in the more general framework of latching models [92–96]—combined with oscillations in external drive. The sequence of recall or the sequence of words within a grammar [97] can be related to the connectivity between neurons comprising the corresponding attractor states, or the number of units shared between different states. The connectivity necessary for recall of sequences within winner-takes-all circuits can be achieved by spike-timing dependent plasticity [98–100] or other bidirectional Hebbian plasticity rules [61].

Are attractor-state models contradicted by physiological data?

Variability of neural activity in multi-attractor systems

One argument against taking point-attractor models seriously is they appear to preclude the obvious variability in neural activity apparent at all spatial and temporal scales. Since the attractor state is a fixed point of the dynamics, from a mathematical viewpoint once in an attractor state neural activity is fixed and static. Therefore one has to be careful about the extent to which an attractor description is valid [101]. In particular, as mentioned in the opening paragraph, the level of description is important—a high-level variable (such as mean firing rate of a group of neurons) can be static, so in an attractor state, even when the microscopic constituents of that high-level quantity vary incessantly.

From an empirical viewpoint, when HMM produces a better description of spike trains than does the PSTH (*e. g.*, as measured by log likelihood [28]) one should take the attractor state model seriously. From a modeling perspective, one can separate time-scales, averaging over the very fast time-scale of single spikes (a few ms), while taking slowly varying properties as parameters for a model of the dynamics of the intermediate time-scale (tens of ms). If the intermediate time-scale model possesses attractor states then much can be learned from an attractor-state formalism, even if these states vary slowly and disappear over time. Indeed, as external stimuli, neuromodulators, and internal facilitating or adaptive processes cause the attractor landscape to shift, biasing neural activity toward new states, so the computational repertoire of such systems is revealed.

Taken in this light, attractor-state models do not preclude variability, even in a high-level variable such as the mean firing rate of a group of cells. Within an attractor state, variability of neural firing rate arises both from the slow drift of parameters and from inherent noise. Greater changes in neural activity arise from the switches between attractor states, which provide slow timescales for activity variation, that can span orders of magnitude [102]. Finally, attractor models can explain the observed reduction in trial-to-trial variability of neural activity upon stimulus presentation [103] when an external stimulus forces the activity into a particular state, following a period of itinerancy between many states during spontaneous activity [45,104–106].

Firing-rate jumps between discrete attractor states can be small

Early implementations of discrete attractor models suggested the firing rates of neurons comprising a unit would jump between zero and a maximal level as the unit switched between “on” and “off” values in different attractor states. However, differences in neural firing rate of much more than 10Hz are rarely seen in persistent activity or during the non-transient period of a constant stimulus [107,108]. Thus discrete attractor models appeared to be at odds with the data.

In most models the firing rate of the “on” state is limited only by saturating feedback. This can be kept in check via synapses dominated by NMDA receptors with a slow time constant [109] or when there is partial overlap of neurons representing the states [8]. Alternative models demonstrate that low firing rates of cells in the “on” state are possible if the network’s stable states do not depend on the two stable fixed points of excitatory cells with

strong feedback [110–112]. In one case, high levels of noise shift activity away from deterministic fixed points [110], while in the other case, the inhibition-stabilized (IS) regime is used [111,112].

The IS network is intriguing (Fig. 3), as strong negative feedback stabilizes activity of excitatory cells at what would otherwise be an unstable fixed point (like a pencil balancing on its tip). The IS network is deemed “paradoxical” in that a reduction of external excitatory drive to inhibitory cells leads to an increase in their activity due to a resulting increase of within-circuit excitatory input. The excitatory cells stabilize at a higher firing rate with increased inhibition because of the concurrent destabilizing positive excitatory feedback. Of note, if units of excitatory and inhibitory cells in an IS circuit are weakly connected with each other, they can generate multiple stable “on” states of arbitrarily low firing rates (Fig. 3) [111]. Moreover, unlike other attractor models that rely on cross-inhibition, inhibition can rise in phase with excitation in the IS regime (Fig. 1b), as observed in many experiments [113–115].

Discussion

Attractor states are inevitably present in any coupled system, because the value of one variable constrains to some extent the possible values of other variables. An attractor state simply means that there is a restricted set of combined values of all the variables, toward which the system returns following any small deviation from that restricted set. Even in chaotic neural systems, the combinations of firing rates that the system moves through over time are restricted to a small set of the possible combinations, so chaotic systems possess strange attractors. In oscillators, the phase relationship between different variables is typically fixed and returned to following a deviation, so oscillators are also attractors, namely cyclic attractors. (See [53] for a more general review with examples of attractors in neural circuits.)

Given that the observable ongoing dynamics of neural activity is more compatible with other types of attractor state, a focus on point attractors may seem artificial. However, whenever neural activity persists in one distinguishable pattern for a time much longer than the switching time to the next pattern, a model of switching between point attractors is a valuable starting point for analysis of the system. Even if the distinct patterns are oscillatory in nature, point attractors can be used in a model, in which the transitions are induced by an additional oscillatory term rather than by noise alone.

A method of analysis like HMM can reveal abrupt state transitions in systems without point attractors, so long as the state-specific features of neural activity are included in the analysis. When the states are distinguished by significant coordinated changes in firing rate across many neurons, as is the case when the states are distinct point attractors, then the spike counts in small time-windows are the appropriate variables for HMM analysis. However, when states are distinguished by oscillation frequency, then the amplitudes of different frequency components of the Fourier transform of small time-windows are the appropriate variables for HMM analysis. Whether neural activity patterns are epitomized by punctate jumps between relatively stable patterns, or by a gradual meandering of firing rates without

any semblance of stationariness, should be revealed in the coming years as larger datasets of simultaneously recorded neurons are analyzed. HMM will be a valuable tool in this enterprise.

Conclusions

Circuits with multiple discrete attractor states have many computational benefits, are easily formed by known mechanisms of synaptic plasticity, and are suggested by both behavioral and electrophysiological data when trials are analyzed separately from each other. Therefore single-trial analysis methods, such as hidden Markov modeling, are to be recommended for any investigation of the dynamics of neural activity, with a view to extracting any underlying attractor-state itinerancy.

Acknowledgments

I am grateful to NIDCD for funding under grant DC09945. I am grateful for useful feedback on a draft of the manuscript by Jonathan Cannon and Tony Ng.

References

1. Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA*. 1982; 79:2554–2558. [PubMed: 6953413]
2. Hopfield JJ. Neurons with graded response have collective computational properties like those of two-state neurons. *Proc Natl Acad Sci USA*. 1984; 81:3088–3092. [PubMed: 6587342]
3. Cohen MA, Grossberg S. Absolute Stability of Global Pattern Formation and Parallel Memory Storage by Competitive Neural Networks. *IEEE transactions on Systems, Man, and Cybernetics*. 1983; SMC-113:815–826.
4. Grossberg S. Some nonlinear networks capable of learning a spatial pattern of arbitrary complexity. *Proceedings of the National Academy of Science USA*. 1968; 59:368–372.
5. Treves A. Graded-response neurons and information encodings in autoassociative memories. *Phys Rev A*. 1990; 42:2418–2430. [PubMed: 9904294]
6. Battaglia FP, Treves A. Stable and rapid recurrent processing in realistic autoassociative memories. *Neural Comput*. 1998; 10:431–450. [PubMed: 9472489]
7. Amit DJ, Brunel N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex*. 1997; 7:237–252. [PubMed: 9143444]
8. Curti E, Mongillo G, La Camera G, Amit DJ. Mean field and capacity in realistic networks of spiking neurons storing sparsely coded random memories. *Neural Comput*. 2004; 16:2597–2637. [PubMed: 15516275]
9. Funahashi S, Bruce CJ, Goldman-Rakic PS. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol*. 1989; 61:331–349. [PubMed: 2918358]
10. Fuster JM, Jervey JP. Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. *Science*. 1981; 212:952–955. [PubMed: 7233192]
- 11*. Rolls ET, Deco G. Stochastic cortical neurodynamics underlying the memory and cognitive changes in aging. *Neurobiol Learn Mem*. 2015; 118:150–161. The reduction of concentrations of dopamine, acetylcholine and noradrenaline with aging conspire to reduce the stability of attractor states and thus decrease the duration of short-term memories. [PubMed: 25536108]
12. Miller P. Stabilization of memory states by stochastic facilitating synapses. *Journal of Mathematical Neuroscience*. 2013
13. Miller P, Wang XJ. Stability of discrete memory states to stochastic fluctuations in neuronal systems. *Chaos*. 2006; 16:026110. [PubMed: 16822042]
14. Moreno-Bote R, Rinzel J, Rubin N. Noise-induced alternations in an attractor network model of perceptual bistability. *J Neurophysiol*. 2007; 98:1125–1139. [PubMed: 17615138]

15. Shpiro A, Moreno-Bote R, Rubin N, Rinzel J. Balance between noise and adaptation in competition models of perceptual bistability. *J Comput Neurosci.* 2009; 27:37–54. [PubMed: 19125318]
16. Deppisch J, Pawelzik K, Geisel T. Uncovering the synchronization dynamics from correlated neuronal activity quantifies assembly formation. *Biol Cybern.* 1994; 71:387–399. [PubMed: 7993929]
17. Gat I, Tishby N, Abeles M. Hidden markov modeling of simultaneously recorded cells in the associative cortex of behaving monkeys. *Network: Computational and Neural Systems.* 1997; 8:297–322.
18. Radons G, Becker JD, Dulfer B, Kruger J. Analysis, classification, and coding of multielectrode spike trains with hidden Markov models. *Biological Cybernetics.* 1994; 71:359–373. [PubMed: 7948227]
19. Otterpohl JR, Haynes JD, Emmert-Streib F, Vetter G, Pawelzik K. Extracting the dynamics of perceptual switching from 'noisy' behaviour: An application of hidden Markov modelling to pecking data from pigeons. *J Physiol Paris.* 2000; 94:555–567. [PubMed: 11165920]
20. Daelli V, Treves A. Neural attractor dynamics in object recognition. *Exp Brain Res.* 2010; 203:241–248. [PubMed: 20437171]
21. Sigala N, Logothetis NK. Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature.* 2002; 415:318–320. [PubMed: 11797008]
22. Leutgeb JK, Leutgeb S, Treves A, Meyer R, Barnes CA, McNaughton BL, Moser MB, Moser EI. Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron.* 2005; 48:345–358. [PubMed: 16242413]
23. Snowdon CT. Response of nonhuman animals to speech and to species-specific sounds. *Brain Behav Evol.* 1979; 16:409–429. [PubMed: 121704]
24. Rotshtein P, Henson RN, Treves A, Driver J, Dolan RJ. Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nat Neurosci.* 2005; 8:107–113. [PubMed: 15592463]
25. Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Tishby N, Vaadia E. Cortical activity flips among quasi-stationary states. *Proc Natl Acad Sci U S A.* 1995; 92:8616–8620. [PubMed: 7567985]
26. Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E. Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci.* 1996; 16:752–768. [PubMed: 8551358]
27. Rainer G, Miller EK. Neural ensemble states in prefrontal cortex identified using a hidden Markov model with a modified EM algorithm. *Neurocomputing.* 2000; 32–33:961–966.
28. Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB. Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A.* 2007; 104:18772–18777. [PubMed: 18000059]
29. Durstewitz D, Vittoz NM, Floresco SB, Seamans JK. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron.* 2010; 66:438–448. [PubMed: 20471356]
- 30*. Sadacca BF, Mukherjee N, Vladusich T, Li JX, Katz DB, Miller P. The Behavioral Relevance of Cortical Neural Ensemble Responses Emerges Suddenly. *J Neurosci.* 2016; 36:655–669. The authors find that rapid decision-related state transitions in gustatory cortex revealed by HMM of multiple simultaneously recorded spike trains during taste processing are predictive of forthcoming behavior on a trial-by-trial basis. [PubMed: 26791199]
31. Okamoto H, Isomura Y, Takada M, Fukai T. Temporal integration by stochastic recurrent network dynamics with bimodal neurons. *J Neurophysiol.* 2007; 97:3859–3867. [PubMed: 17392417]
- 32*. Latimer KW, Yates JL, Meister ML, Huk AC, Pillow JW. NEURONAL MODELING. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science.* 2015; 349:184–187. Single-trial analysis of individual spike trains shows the trial-averaged ramping is more compatible with abrupt transitions within each trial than gradual integration of evidence. [PubMed: 26160947]

33. Sussillo D, Abbott LF. Generating coherent patterns of activity from chaotic neural networks. *Neuron*. 2009; 63:544–557. [PubMed: 19709635]
34. Lee SH, Kang SH. Characterization of the crawling activity of *Caenorhabditis elegans* using a hidden markov model. *Theory Biosci*. 2015; 134:117–125. [PubMed: 26319806]
35. Katahira K, Suzuki K, Okanoya K, Okada M. Complex sequencing rules of birdsong can be explained by simple hidden Markov processes. *PLoS One*. 2011; 6:e24516. [PubMed: 21915345]
36. Eldar E, Morris G, Niv Y. The effects of motivation on response rate: a hidden semi-Markov model analysis of behavioral dynamics. *J Neurosci Methods*. 2011; 201:251–261. [PubMed: 21782849]
37. Blumenfeld B, Bibitchkov D, Tsodyks M. Neural network model of the primary visual cortex: from functional architecture to lateral connectivity and back. *J Comput Neurosci*. 2006; 20:219–241. [PubMed: 16699843]
38. Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A. Spontaneously emerging cortical representations of visual attributes. *Nature*. 2003; 425:954–956. [PubMed: 14586468]
39. Tsodyks M, Kenet T, Grinvald A, Arieli A. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*. 1999; 286:1943–1946. [PubMed: 10583955]
40. Bressloff PC, Cowan JD, Golubitsky M, Thomas PJ, Wiener MC. What geometric visual hallucinations tell us about the visual cortex. *Neural Comput*. 2002; 14:473–491. [PubMed: 11860679]
41. Ermentrout GB, Cowan JD. A mathematical theory of visual hallucination patterns. *Biol Cybern*. 1979; 34:137–150. [PubMed: 486593]
- 42**. Litwin-Kumar A, Doiron B. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nat Commun*. 2014; 5:5319. From an initially randomly coupled balanced network, spike-timing dependent plasticity produces stimulus-dependent assemblies that are spontaneously reactivated in the absence of stimuli. Ongoing spike-timing dependent and homeostatic plasticity combined with the structured spontaneous activity was necessary for stabilization of the learned connectivity. [PubMed: 25395015]
43. Berkes P, Orban G, Lengyel M, Fiser J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*. 2011; 331:83–87. [PubMed: 21212356]
44. Goldberg JA, Rokni U, Sompolinsky H. Patterns of ongoing activity and the functional architecture of the primary visual cortex. *Neuron*. 2004; 42:489–500. [PubMed: 15134644]
- 45**. Mazzucato L, Fontanini A, La Camera G. Dynamics of multistable states during ongoing and evoked cortical activity. *J Neurosci*. 2015; 35:8214–8231. HMM of multiple simultaneously recorded spike trains in gustatory cortex of rats reveals transitions between discrete activity states during spontaneous activity as well as during activity evoked by a taste stimulus. [PubMed: 26019337]
46. Kringelbach ML, McIntosh AR, Ritter P, Jirsa VK, Deco G. The Rediscovery of Slowness: Exploring the Timing of Cognition. *Trends Cogn Sci*. 2015; 19:616–628. [PubMed: 26412099]
- 47*. Golos M, Jirsa V, Dauce E. Multistability in Large Scale Models of Brain Activity. *PLoS Comput Biol*. 2015; 11:e1004644. When units in a large-scale brain model are connected according to human connectome data derived from diffusive spectrum imaging, the model network so-produced contains multiple attractor states for large ranges of parameters. The resulting dynamics are compatible with functional connectivity data derived from fMRI. [PubMed: 26709852]
48. Hansen EC, Battaglia D, Spiegler A, Deco G, Jirsa VK. Functional connectivity dynamics: modeling the switching behavior of the resting state. *Neuroimage*. 2015; 105:525–535. [PubMed: 25462790]
49. Deco G, Jirsa VK. Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. *J Neurosci*. 2012; 32:3366–3375. [PubMed: 22399758]
50. Freeman WJ. Characterization of state transitions in spatially distributed, chaotic, nonlinear, dynamical systems in cerebral cortex. *Integr Physiol Behav Sci*. 1994; 29:294–306. [PubMed: 7811649]
51. Freeman WJ, Holmes MD. Metastability, instability, and state transition in neocortex. *Neural Netw*. 2005; 18:497–504. [PubMed: 16095879]
52. Osinski BL, Kay LM. Granule cell excitability regulates gamma and beta oscillations in a model of the olfactory bulb dendrodendritic microcircuit. *J Neurophysiol*. 2016 jn 00988 02015.

53. Miller P. Dynamical systems, attractors, and neural circuits. *F1000 Faculty Reviews*. 2016
54. Weimann JM, Meyrand P, Marder E. Neurons that form multiple pattern generators: identification and multiple activity patterns of gastric/pyloric neurons in the crab stomatogastric system. *J Neurophysiol*. 1991; 65:111–122. [PubMed: 1999725]
55. Tsuda I. Toward an interpretation of dynamic neural activity in terms of chaotic dynamical systems. *Behav Brain Sci*. 2001; 24:793–810. 810–748. discussion. [PubMed: 12239890]
56. Tsuda I. Chaotic itinerancy and its roles in cognitive neurodynamics. *Curr Opin Neurobiol*. 2015; 31:67–71. [PubMed: 25217808]
57. Freeman, WJ.; Barrie, JM. Chaotic oscillations and the genesis of meaning in cerebral cortex. In: Buzsaki, G., editor. *In Temporal Coding in the Brain*. Springer-Verlag; 1994. p. 13–37.
58. Rabinovich MI, Sokolov Y, Kozma R. Robust sequential working memory recall in heterogeneous cognitive networks. *Front Syst Neurosci*. 2014; 8:220. [PubMed: 25452717]
59. Afraimovich V, Gong X, Rabinovich M. Sequential memory: Binding dynamics. *Chaos*. 2015; 25:103118. [PubMed: 26520084]
60. Rabinovich MI, Varona P, Tristan I, Afraimovich VS. Chunking dynamics: heteroclinics in mind. *Front Comput Neurosci*. 2014; 8:22. [PubMed: 24672469]
- 61*. Fonollosa J, Neftci E, Rabinovich M. Learning of Chunking Sequences in Cognition and behavior. *PLoS Comput Biol*. 2015; 11:e1004592. Commencing with a multiple-attractor, winner-takes-all network, bidirectional Hebbian learning in a firing-rate model generates recall of sequences via winnerless competition, meaning previous attractor states lose their stability in a specific direction corresponding to the next item in the sequence to be recalled—i. e., a heteroclinic sequence is produced. A hierarchy of heteroclinic sequences accounts for chunking of items during memory recall. [PubMed: 26584306]
62. Escola S, Fontanini A, Katz D, Paninski L. Hidden Markov models for the stimulus-response relationships of multistate neural systems. *Neural Comput*. 2011; 23:1071–1132. [PubMed: 21299424]
63. Barak O, Sussillo D, Romo R, Tsodyks M, Abbott LF. From fixed points to chaos: three models of delayed discrimination. *Prog Neurobiol*. 2013; 103:214–222. [PubMed: 23438479]
64. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503:78–84. [PubMed: 24201281]
65. Kalitzin S, Koppert M, Petkov G, da Silva FL. Multiple oscillatory states in models of collective neuronal dynamics. *Int J Neural Syst*. 2014; 24:1450020. [PubMed: 25081428]
66. Buzsaki G. Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*. 2010; 68:362–385. [PubMed: 21040841]
67. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol*. 2005; 3:e68. [PubMed: 15737062]
68. Perin R, Berger TK, Markram H. A synaptic organizing principle for cortical neuronal groups. *Proc Natl Acad Sci U S A*. 2011; 108:5419–5424. [PubMed: 21383177]
69. Perin R, Telefont M, Markram H. Computing the size and number of neuronal clusters in local circuits. *Front Neuroanat*. 2013; 7:1. [PubMed: 23423949]
70. Bourjaily MA, Miller P. Excitatory, inhibitory, and structural plasticity produce correlated connectivity in random networks trained to solve paired-stimulus tasks. *Frontiers in Computational Neuroscience*. 2011; 5:37. [PubMed: 21991253]
71. Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöstrom PJ, Mrcic-Flogel TD. Functional specificity of local synaptic connections in neocortical networks. *Nature*. 2011; 473:87–91. [PubMed: 21478872]
72. Miller P. Stimulus number, duration and intensity encoding in randomly connected attractor networks with synaptic depression. *Front Comput Neurosci*. 2013; 7:59. [PubMed: 23675344]
- 73*. Stern M, Sompolinsky H, Abbott LF. Dynamics of random neural networks with bistable units. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2014; 90:062710. Bistable neural units coupled randomly produce rich dynamics, with chaotic activity compatible with multiple attractor states for broad ranges of parameters. [PubMed: 25615132]
74. Wang XJ. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*. 2002; 36:955–968. [PubMed: 12467598]

75. Wang XJ. Decision making in recurrent neuronal circuits. *Neuron*. 2008; 60:215–234. [PubMed: 18957215]
76. Wong KF, Huk AC, Shadlen MN, Wang XJ. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Frontiers in Computational Neuroscience*. 2007; 1:6. [PubMed: 18946528]
77. Wong KF, Wang XJ. A recurrent network mechanism of time integration in perceptual decisions. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2006; 26:1314–1328. [PubMed: 16436619]
78. Miller P, Katz DB. Stochastic transitions between neural states in taste processing and decision-making. *J Neurosci*. 2010; 30:2559–2570. [PubMed: 20164341]
79. Miller P, Katz DB. Accuracy and response-time distributions for decision-making: linear perfect integrators versus nonlinear attractor-based neural circuits. *Journal of Computational Neuroscience*. 2013; 35:261–294. [PubMed: 23608921]
80. Ponce-Alvarez A, Nacher V, Luna R, Riehle A, Romo R. Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2012; 32:11956–11969. [PubMed: 22933781]
81. Bollimunta A, Totten D, Ditterich J. Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2012; 32:12684–12701. [PubMed: 22972993]
82. Burak Y, Fiete IR. Fundamental limits on persistent activity in networks of noisy neurons. *Proc Natl Acad Sci U S A*. 2012; 109:17645–17650. [PubMed: 23047704]
83. Moreno-Bote R, Shpiro A, Rinzel J, Rubin N. Alternation rate in perceptual bistability is maximal at and symmetric around equi-dominance. *J Vis*. 2010; 10:1.
84. Moreno-Bote R, Knill DC, Pouget A. Bayesian sampling in visual perception. *Proc Natl Acad Sci U S A*. 2011; 108:12491–12496. [PubMed: 21742982]
85. Bernacchia A. The interplay of plasticity and adaptation in neural circuits: a generative model. *Front Synaptic Neurosci*. 2014; 6:26. [PubMed: 25400577]
86. Nessler B, Pfeiffer M, Buesing L, Maass W. Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Comput Biol*. 2013; 9:e1003037. [PubMed: 23633941]
87. Legenstein R, Maass W. Ensembles of spiking neurons with noise support optimal probabilistic inference in a dynamically changing environment. *PLoS Comput Biol*. 2014; 10:e1003859. [PubMed: 25340749]
- 88*. Katkov M, Romani S, Tsodyks M. Effects of long-term representations on free recall of unrelated words. *Learn Mem*. 2015; 22:101–108. An attractor-transition model of word recall, in which the size of an attractor correlates with the likelihood of recall of any specific word accounts for observed recall patterns. In particular, the observation that recall of easy words decreases the likelihood of later recall of other words is accounted for by the retrieval process returning to the large basin of attraction corresponding to the easy word. [PubMed: 25593296]
89. Recanatesi S, Katkov M, Romani S, Tsodyks M. Neural Network Model of Memory Retrieval. *Front Comput Neurosci*. 2015; 9:149. [PubMed: 26732491]
90. Katkov M, Romani S, Tsodyks M. Word length effect in free recall of randomly assembled word lists. *Front Comput Neurosci*. 2014; 8:129. [PubMed: 25352804]
91. Romani S, Pinkoviezky I, Rubin A, Tsodyks M. Scaling laws of associative memory retrieval. *Neural Comput*. 2013; 25:2523–2544. [PubMed: 23777521]
92. Akrami A, Russo E, Treves A. Lateral thinking, from the Hopfield model to cortical dynamics. *Brain Res*. 2012; 1434:4–16. [PubMed: 21839426]
93. Russo E, Treves A. Cortical free-association dynamics: distinct phases of a latching network. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2012; 85:051920. [PubMed: 23004800]
94. Treves A. Frontal latching networks: a possible neural basis for infinite recursion. *Cogn Neuropsychol*. 2005; 22:276–291. [PubMed: 21038250]
95. Kropff E, Treves A. The storage capacity of Potts models for semantic memory retrieval. *Journal of Statistical Mechanics: Theory and Experiment*. 2005:P08010.

96. Lerner I, Bentin S, Shriki O. Integrating the automatic and the controlled: strategies in semantic priming in an attractor network with latching dynamics. *Cogn Sci*. 2014; 38:1562–1603. [PubMed: 24890261]
97. Rolls ET, Deco G. Networks for memory, perception, and decision-making, and beyond to how the syntax for language might be implemented in the brain. *Brain Res*. 2015; 1621:316–334. [PubMed: 25239476]
98. Kappel D, Nessler B, Maass W. STDP installs in Winner-Take-All circuits an online approximation to hidden Markov model learning. *PLoS Comput Biol*. 2014; 10:e1003511. [PubMed: 24675787]
99. Klampfl S, Maass W. Emergence of dynamic memory traces in cortical microcircuit models through STDP. *J Neurosci*. 2013; 33:11515–11529. [PubMed: 23843522]
100. Miller P, Wingfield A. Distinct effects of perceptual quality on auditory word recognition, memory formation and recall in a neural model of sequential memory. *Front Syst Neurosci*. 2010; 4:14. [PubMed: 20631822]
101. Durstewitz D, Seamans JK. Beyond bistability: biophysics and temporal dynamics of working memory. *Neuroscience*. 2006; 139:119–133. [PubMed: 16326020]
- 102**. Schaub MT, Billeh YN, Anastassiou CA, Koch C, Barahona M. Emergence of Slow-Switching Assemblies in Structured Neuronal Networks. *PLoS Comput Biol*. 2015; 11:e1004196. In simulations of hierarchical networks of spiking neurons, noise-driven transitions between attractor states at a low level of the hierarchy provide the fluctuations that can trigger abrupt transitions between attractor states at a higher level, yet on a much slower timescale. [PubMed: 26176664]
103. Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*. 2010; 13:369–378. [PubMed: 20173745]
104. Deco G, Hugues E. Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Comput Biol*. 2012; 8:e1002395. [PubMed: 22479168]
105. Litwin-Kumar A, Doiron B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*. 2012; 15:1498–1505. [PubMed: 23001062]
106. Doiron B, Litwin-Kumar A. Balanced neural architecture and the idling brain. *Front Comput Neurosci*. 2014; 8:56. [PubMed: 24904394]
107. Miyashita Y. Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*. 1988; 335:817–820. [PubMed: 3185711]
108. Warden MR, Miller EK. Task-dependent changes in short-term memory in the prefrontal cortex. *J Neurosci*. 2010; 30:15801–15810. [PubMed: 21106819]
109. Wang XJ. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J Neurosci*. 1999; 19:9587–9603. [PubMed: 10531461]
110. Amit DJ, Treves A. Associative memory neural network with low temporal spiking rates. *Proc Natl Acad Sci U S A*. 1989; 86:7871–7875. [PubMed: 2813363]
111. Latham PE, Nirenberg S. Computing and stability in cortical networks. *Neural Comput*. 2004; 16:1385–1412. [PubMed: 15165395]
112. Golomb D, Rubin N, Sompolinsky H. Willshaw model: Associative memory with sparse coding and low firing rates. *Phys Rev A*. 1990; 41:1843–1854. [PubMed: 9903293]
- 113**. Rubin DB, Van Hooser SD, Miller KD. The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex. *Neuron*. 2015; 85:402–417. Evidence for a switch to the inhibition-stabilized regime with strong inputs is presented via a wealth of neural activity data from primary visual cortex of ferrets backed up by extensive simulations and models. In particular, a shift from supralinear to sublinear summation of spatially separated inputs is accounted for by the models. [PubMed: 25611511]
114. Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL. Paradoxical effects of external modulation of inhibitory interneurons. *J Neurosci*. 1997; 17:4382–4388. [PubMed: 9151754]
115. Anderson JS, Carandini M, Ferster D. Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J Neurophysiol*. 2000; 84:909–926. [PubMed: 10938316]

Highlights

Multiple point-attractor states form in neural circuits with clustered connectivity.

Transitions between attractor states can arise without a change in stimulus.

Transitions and final state reached can be stimulus-history or context dependent.

Variability of the timing of transitions across trials renders the PSTH inadequate.

Hidden Markov modeling reveals such state transitions in multi-unit data.

Figure 1. Noise-induced transitions generate itinerancy between the two attractor states of a bistable system, accounting for perceptual alternation

A) Three separate trials (vertically aligned panels) indicate back and forth switching between states where one unit's activity dominates over the other, but with across-trial variability in the transition times. B) Activity averaged over ten trials, as is common for neural data, obscures all signs of the underlying dynamics. Blue and red curves denote activity of separate excitatory cell-groups.

Figure 2. Hidden Markov modeling (HMM)

A) HMM operates on sets of simultaneously recorded spike trains, from either a single time-series or (as in this case) multiple trials. **B)** HMM assumes a number of discrete states (colored ovals) with a matrix of transition probabilities per unit time between the states (represented by the thickness of arrows). Each state is defined by the probability of emission of a particular pattern in a time-bin (inset histograms). For neural spike trains an emitted pattern is the combination of neurons producing a spike in a small time bin, so a state's emission probabilities are directly related to the firing rate of each neuron in that state. HMM proceeds iteratively in a two-step process: 1) Given a possible model of transition and emission probabilities, it calculates the likelihood the observed spike trains are produced by the model, giving the probability the system is in any given state at any point in time on any trial. 2) Given the fitting of state probabilities to the spike trains, both the emission probabilities of each state and transition probabilities between states are recalculated. Repeated iteration of steps 1) and 2) increases the log likelihood of the match to the data until a local optimum is reached. Typically numerous random initial conditions for the emission and transition probabilities are used in case one particular set of iterations reaches a local optimum that is not the global optimum. **C)** Given the optimized HMM, the posterior probabilities for each state are ascribed to the time-series of spike trains (colored solid lines overlaying the spike trains).

Figure 3. Low firing rates and in-phase inhibition in the inhibition-stabilized regime

A)–C) Two bistable units (one blue, one red) with symmetric cross-inhibition (excitatory connections to inhibitory cells between units) switch between states in which one unit's activity dominates. **D)–F)** With alternative circuit parameters, each excitatory-inhibitory unit is in the inhibition-stabilized (IS) regime, allowing for bistability between states of low firing rates. In the IS regime (E–F), inhibitory input to a cell covaries with excitatory input—notice that cells from the same group are in phase—unlike standard multistable models (B–C) in which cross-inhibition dominates—notice that cells from the same group are out of

phase, because activity of inhibitory cells is determined most by the other unit's excitatory cell. **A), D)** Arrows denote excitatory connections, solid circles denote inhibitory connections. **B–C), E–F)** Blue and red traces denote the two cell-groups. See supplement for code that generated this figure.