Available online at www.sciencedirect.com

ScienceDirect

Current Opinion in Neurobiology

Statistical learning and representational drift: A dynamic substrate for memories

Jens-Bastian Eppler[1], Matthias Kaschube[2] and Simon Rumpel[3]

In many brain areas, neurons exhibit continuous changes in their tuning properties over days, even when supporting stable percepts and behaviors—a phenomenon termed representational drift. How do neuronal circuits maintain stable function when their constituent elements are in constant flux? Here, we review recent theoretical and experimental work on interconnected levels, ranging from perpetual changes in synapses driving drifts in tuning of individual neurons to emergent stability at the population level, preserving similarities of activity patterns associated to specific percepts or behaviors. We propose that statistical learning, beyond its well-established roles during development and adaptation to new contexts, is also essential under steady behavioral and environmental conditions to safeguard the stability of representational similarities. We discuss implications for learning, memory, and forgetting. This framework reconciles the apparent paradox between unstable neural activity and stable perception, suggesting that representations are maintained through dynamic processes rather than static neural codes.

Addresses
[1] Centre de Recerca Matemàtica, Edifici C, Campus Bellaterra, 08193 Bellaterra, Spain
[2] Frankfurt Institute for Advanced Studies and Institute of Computer Science, Goethe University Frankfurt, Ruth-Moufang-Straße 1, 60438 Frankfurt, Germany
[3] Institute of Physiology, Focus Program Translational Neurosciences, University Medical Center of the Johannes Gutenberg University Mainz, Duesbergweg 6, 55128 Mainz, Germany

Corresponding author: Rumpel, Simon (sirumpel@uni-mainz.de)

Current Opinion in Neurobiology 2025, **94:**103107

This review comes from a themed issue on **Systems Neuroscience 2025**

Edited by **Robert Froemke**, **Máté Lengyel**, **József Fiser** and **Livia de Hoz**

For complete overview of the section, please refer the article collection - Systems Neuroscience 2025

Available online 2 September 2025

https://doi.org/10.1016/j.conb.2025.103107

Representational drift

Behavior and sensory experiences can remain remarkably consistent over extended periods, from weeks to a lifetime. When we hear a song, specific neuronal populations in the auditory cortex become active. One might assume that hearing the same song days later would reactivate the same neurons. However, recent research suggests that the tuning of neurons linked to stable behaviors or sensory representations is not fixed but instead changes over hours to days. This phenomenon, known as representational drift, has been observed in various brain regions, including the motor [1], barrel [2], parietal [3], olfactory [4], visual [5], and auditory cortices [6], as well as the hippocampus [7–9].

The discovery of representational drift has sparked discussions about its functional implications, as changes in neuronal tuning appear to contradict the idea of a stable percept or memory [10]. Traditionally, the activity pattern associated with a particular sensory stimulus is thought to constitute its neuronal representation, forming the basis of perception. However, representational drift suggests that the same stimulus can have multiple representations over time. This, in turn, implies that activation of different neuronal populations may generate the same percept. This raises a fundamental question: What connects these different neuronal representations, or activity patterns, to a consistent and stable percept?

Representation of similarities

An influential concept from cognitive science may offer insight into this puzzle: The representation of abstract cognitive entities, such as percepts or memories, can be understood as a *representation of similarities* [11,12]. In this framework, it is not the particular properties of a sensory stimulus that are primarily encoded in the corresponding activity pattern. Various activity patterns rather encode the relatedness between the associated representational entities by a larger or lesser degree of similarity. Recent developments have expanded on this idea, proposing that neural representations form maps that encode relationships between stimuli [13–15].

Depending on the level of hierarchical processing, these relationships may capture either fundamental perceptual features or more abstract semantic aspects. Notably, this coding scheme is degenerate: different pairs of activity patterns can encode the same degree of relatedness as long as they exhibit the same level of similarity. Consequently, representational maps convey the underlying relational structure of neural information independently of the specific activity patterns [16,17].

To better differentiate between these two levels of representation, we will focus on 'activity patterns' as the classical neuronal representation of a sensory stimulus at a given time point, while 'representational similarity' will be used to describe the higher-level correlate of the representation of cognitive entities.

How can the relational structure of a representational map remain stable despite continuous changes in the tuning of individual neurons? This principle can be illustrated with a simple analogy (Figure 1). Imagine a circle drawn on a three-dimensional object (or manifold), such as a chess pawn (Figure 1a). As the object rotates, the three-dimensional coordinates of the points on the circle shift continuously, much like drifting neuronal activity patterns (Figure 1b). However, the geometric relationships between these points—their relative distances and the overall circular structure—remain unchanged (Figure 1c).

The relevance of this analogy becomes evident when examining longitudinal recordings of neuronal population activity in the mouse auditory cortex, previously published [6]. Over multiple days, individual neurons exhibit changing responses to pure tones of different frequencies (Figure 1d). However, the overall pairwise similarities between population responses remain structurally consistent, preserving the perceptual relatedness between stimuli (Figure 1e).

The common experience of a stimulus consistently evoking the same percept suggests that its representational similarity and the structure of the representational map in which it is embedded remains stable. Indeed, recent studies have reported the persistence of representational similarities across various brain regions. For instance, Gallego et al. [18] demonstrated that although neuronal activity patterns in the macaque motor cortex continuously changed, the representational similarity between population patterns of different movement directions remained intact. Similarly, in the mouse visual cortex, the internal structure of population activity elicited by naturalistic movies was preserved over multiple days [5,19]. Comparable findings have been made in the hippocampus. Rubin et al. [20] and Sheintuch et al. [21] showed that while the specific neurons encoding a task changed over time, the abstract representational structure reflecting alternating runs between rewarded ends of a linear track remained stable across both time and experimental subjects, appearing as a consistent pattern of similarities within neural state space. The structure of a representational map appears not only to be resilient to tuning changes of individual neurons, but even has the capability to recover after a transient disturbance induced by the targeted ablation of individual neurons [22].

These findings highlight that when considering the term "representational drift", it is important to note that while the tuning of individual neurons indeed drifts, the representational similarity—defined as the different degrees of similarities of the activity patterns associated with the corresponding representational entities—appears to remain stable.

Drift in synaptic connections

Despite the abundance of studies documenting representational drift, the mechanisms driving these changes are less well understood. Ongoing changes occur not only in neuronal tuning but also in the network structure in which neurons are embedded. Synapses—both excitatory and inhibitory—undergo continuous modifications, including the formation and elimination of connections as well as changes in synaptic strength. The steady-state dynamics of synaptic changes have been shown to shape fundamental statistical properties of neuronal circuits, such as the size distribution of synaptic connections or their overall density [23,24]. These dynamics have been observed to varying degrees across different brain regions, even under stable environmental and behavioral conditions [25,26]. While multiple processes may contribute to the tuning changes associated with representational drift, including changes in intrinsic excitability [27,28], alterations in a neuron's synaptic inputs appear to be a key factor.

Synaptic changes fall into several categories: Synaptic modifications can be driven by experience, as seen in classical Hebbian forms of activity-dependent plasticity, which play a crucial role in learning and memory [29]. Other mechanisms include internal state-dependent synaptic plasticity [30] and homeostatic synaptic plasticity, which serves to stabilize neuronal activity by counteracting excessive change and maintaining functional network dynamics [31]. Growing evidence suggests that synaptic changes can also occur through stochastic, activity-independent processes. These changes stem from the biochemical properties of synaptic structures and the inherent randomness of molecular processes at synapses, such as receptor trafficking, protein turnover, and the probabilistic nature of neurotransmitter release [26]. Notably, such stochastic synaptic changes persist even in the absence of neuronal activity [23,32,33].

Figure 1

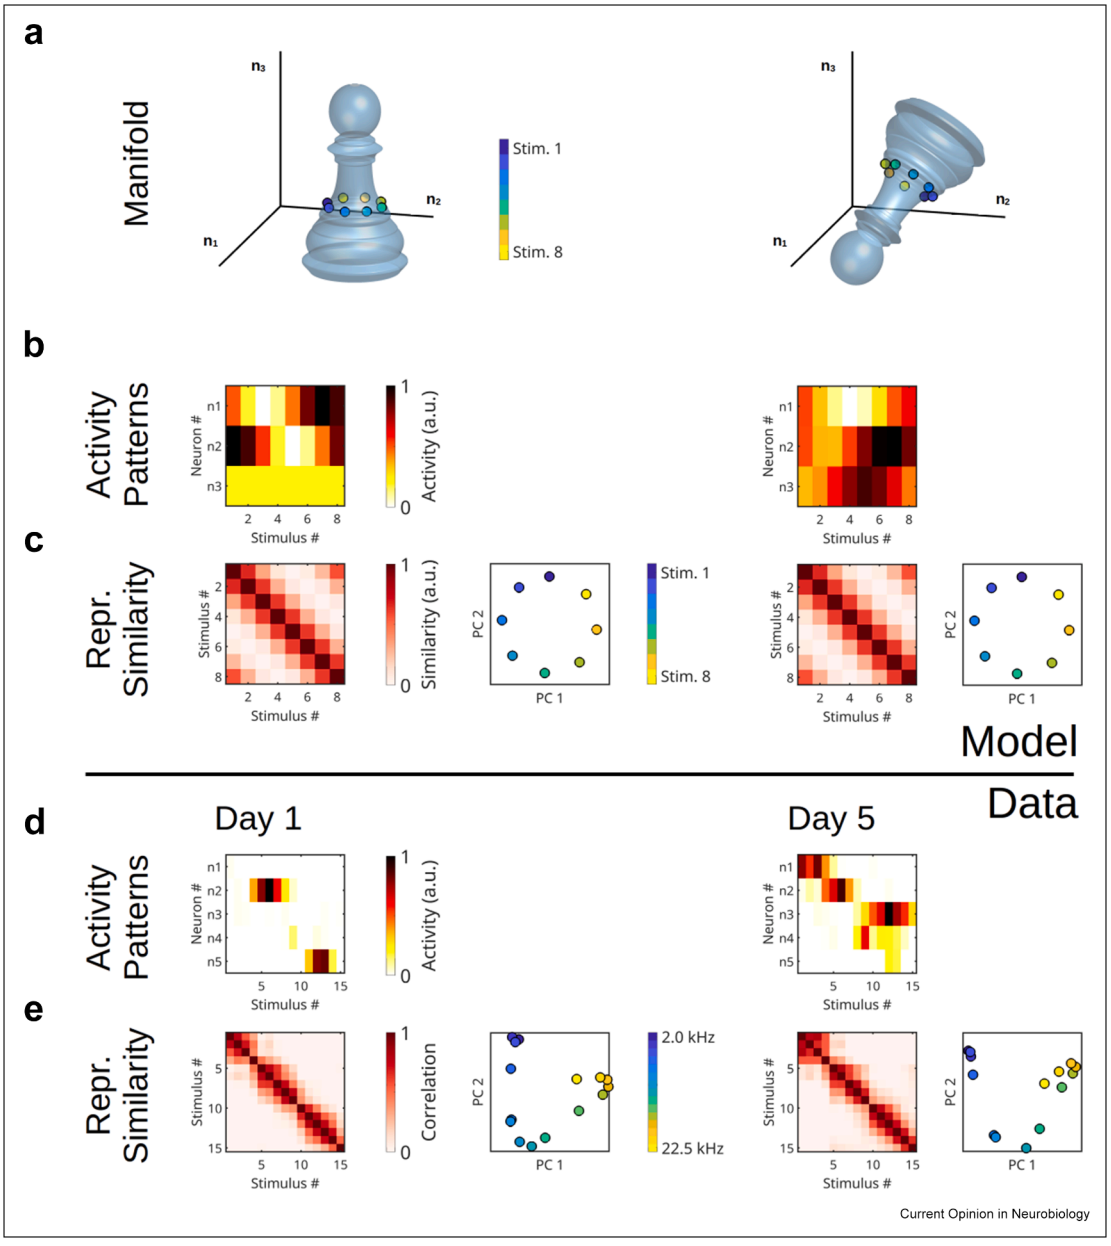

Representational similarity remains stable despite changing activity patterns. a) A rotating 3D manifold (illustrated by a chess pawn) contains eight points representing different stimuli arranged in a circular pattern. While the rotation changes the spatial coordinates of each point, their relative positions and distances remain constant. In this illustration, the three spatial dimensions represent activities of three hypothetical neurons n_i. **b)** The 3D coordinates of these points change during rotation. **c)** Despite coordinate changes, the distance matrix between points and their projection onto the first two principal components remains preserved. **d)** Similar dynamics become apparent in recordings from mouse auditory cortex [6], where the responses to a set of 15 pure tones drift over time, as shown for 5 example neurons. **e)** Despite these changes in tuning, the representational similarity remains largely stable on a population level, as shown by correlation matrices and principal component projections (analysis considering the 1000 most active neurons recorded from 12 animals).

Changes in synapses may underlie representational drift

While some studies have explored how neural circuits can sustain stable activity patterns despite continuous synaptic perturbations [34–36], recent theoretical work has increasingly focused on capturing the phenomenon of representational drift and clarifying the relationship between synaptic changes and neuronal tuning shifts.

Several recent studies have explored learning-dependent synaptic changes as a potential driver of representational drift. Devalle et al. [37] made a notable contribution by showing that continuous Hebbian learning alone can account for key experimental observations of drift. Their model reconciles seemingly contradictory findings by demonstrating that re-exposure to a stimulus can either accelerate drift [38] or promote stabilization [4], depending on the degree of Hebbian learning associated with other stimuli in the interim. Complementary work by Ratzon et al. [39] proposed that representational drift emerges during continual learning of a single stimulus as a form of implicit regularization: after rapid initial learning and slower consolidation, network activity becomes sparser as the network shifts between multiple equally effective solutions within a degenerate solution space. Similarly, Qin et al. [40] showed that ongoing Hebbian-like learning can preserve a stable task readout in a network of neurons showing tuning drift. Both of these latter studies focused primarily on the representation of a single stimulus. Notably, while all three models highlight the importance of learning-dependent mechanisms in driving representational drift, they also incorporate stochastic elements.

These studies align with a broader perspective that interprets representational drift as a natural outcome of continuous learning in neural circuits. Driscoll et al. [41] articulated this view, suggesting that ongoing changes in neural representations reflect adaptive processes rather than mere instability. Supporting this idea, Duncker et al. [42] demonstrated that continual learning in recurrent networks drives persistent shifts in neural activity patterns while preserving task performance. Such ongoing learning processes may also serve a computational role: Johnston & Fusi [43] showed that sequentially learning multiple tasks can give rise to abstract representations that enhance generalization.

In addition to studies focusing solely on continuous, learning-dependent changes, several others have explicitly examined how stochastic network changes interact with synaptic plasticity driven by learning. Theoretical work has proposed that effective learning and memory require an optimal balance between these two forms of synaptic changes [44,45]. Interestingly, the magnitude of stochastic changes has been suggested to be at least comparable to that of learning-induced changes [45]. Recent experimental work provided evidence for such a balance: Bauer et al. [46] showed that sensory deprivation in the mouse visual cortex altered the direction but not the magnitude of tuning drift in neurons. This observation was interpreted as Hebbian plasticity not being the main driver of representational drift, but rather steering its direction to match the statistics of the environment. Consistent findings were made in the mouse auditory cortex, showing that fear learning merely induced biases in the ongoing dynamics of sensory representations [6]. New analyses of this dataset by Eppler et al. [47] revealed that the long-term dynamics of signal and noise correlations in the cortex could only be reproduced in a model incorporating balanced Hebbian and stochastic forms of synaptic changes.

Together, a picture emerges in which, alongside Hebbian plasticity, stochastic changes in synapses may represent a fundamental yet underexplored property of the brain. The degree of stochasticity within the neuronal network parallels the concept of temperature in physical systems, which similarly reflects the degree of stochasticity [48]. In the future, it will be interesting to test to what extent this level of stochasticity can be regulated in a physiological context.

Statistical learning for stable representational similarities

Given recent findings suggesting that representational drift under stable conditions reflects a balance between Hebbian and stochastic synaptic changes, the following questions arise: What might drive ongoing Hebbian plasticity? And could this continuous form of learning help maintain representational similarities as individual neuronal responses change over the course of hours to days?

Interestingly, statistical learning operates within similar timescales as representational drift, ranging from minutes [49] to hours and days [50] and even weeks [51]. Statistical learning—the ability to extract regularities from the environment through experience—has been widely documented across various domains. Initially demonstrated in infant language acquisition [52], where infants identify word boundaries from continuous speech based on transitional probabilities, statistical learning has since been shown to occur across visual [53], tactile, and auditory modalities [54]. This learning process unfolds automatically, often unconsciously, typically without explicit reinforcement, and updates internal models based on environmental statistics in ways that can be formally described through Bayesian inference [55]. The domain-general nature of statistical learning is clear in both its broad applicability [56] and its neural implementation [57], with theoretical frameworks offering robust mathematical foundations for understanding these learning processes [58]. However, statistical learning is traditionally thought to modify neuronal activity patterns during development or when needed to adapt to new environmental conditions. But could it also play a role in maintaining stability under consistent environmental conditions?

We propose that continuous statistical learning may also serve as a major driver of the Hebbian component of continuous synaptic changes. By counteracting stochastic synaptic fluctuations, along with homeostatic

mechanisms [31], statistical learning functions as a stabilizing mechanism, preserving stable representational similarities even in unchanging environments. Thus, even when neural circuits experience consistent environmental statistics—whether in the form of stimulus relationships, task structure, or movement patterns—statistical learning may be required to maintain these regularities in the neural code. Through statistical learning, the environment may act as a stable reference frame for neuronal dynamics during drift. This dynamic balance between undirected, random fluctuations and directed, structured learning could allow neuronal circuits to maintain representational similarities while retaining plasticity [45].

If stochastic changes in the network were left unchecked, the overlap between activity patterns evoked by two stimuli would fluctuate unpredictably (Figure 2a), altering representational similarity and, consequently, the overall representational structure. However, by maintaining a consistent number of neurons responding to each stimulus and preserving their overlap, the representational structure remains stable despite drift in responses of individual neurons (Figure 2b). The ongoing process of statistical learning counterbalancing stochastic synaptic changes can be likened to the idiomatic expression of 'herding cats,' describing an almost impossible task requiring continuous effort. In this analogy, statistical learning functions

Figure 2

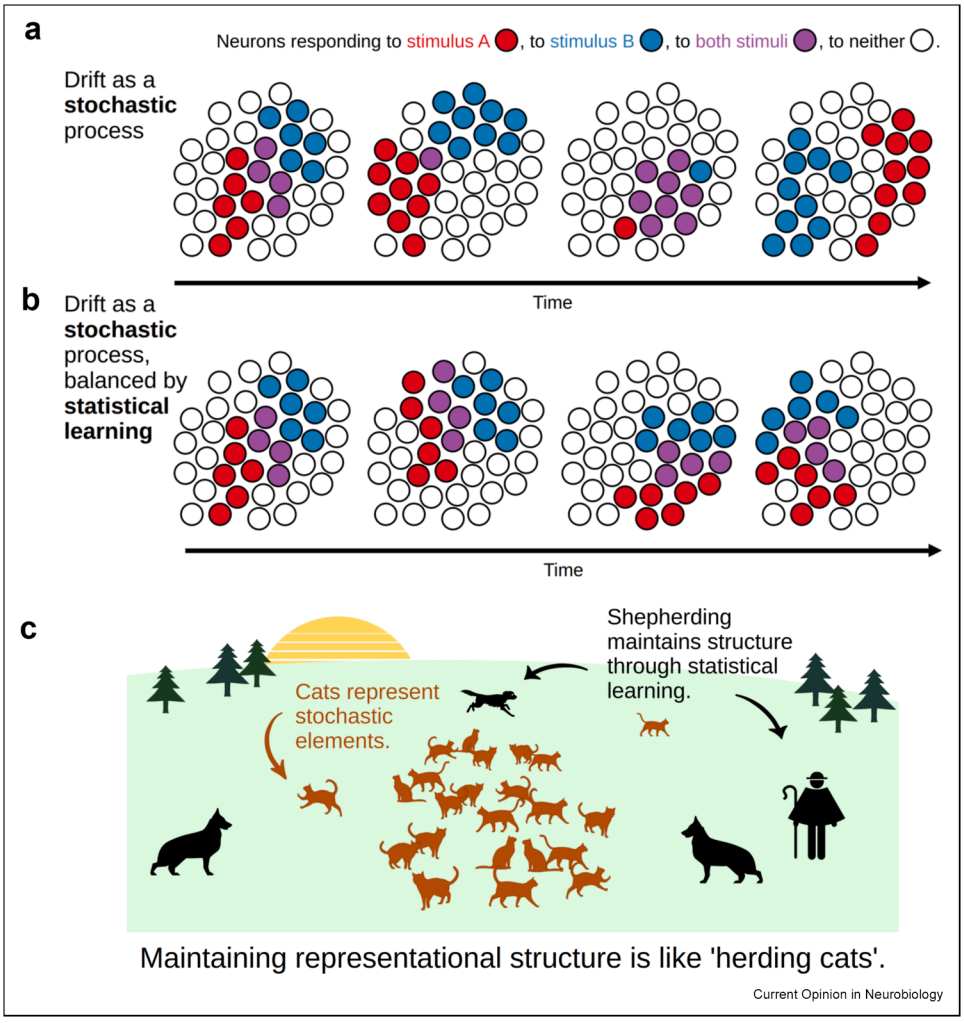

Statistical learning maintains neural representations by balancing stochastic processes. a) Without statistical learning, stochastic changes in the network cause random fluctuations in the overlap between neural populations responding to a pair of stimuli, resulting in unstable representational similarity. **b)** Introducing statistical learning as a stabilizing force may maintain consistent overlap between stimulus-evoked activity patterns and preserve representational similarity. **c)** 'Herding cats' as a metaphoric illustration of this dynamic balance: the cats (brown) represent perpetual stochastic changes in the neuronal network, while the shepherd and herding dogs (black) represent continuous statistical learning essential to maintain order and structure.

as a shepherd with dogs, skillfully bounding the stochastic changes in the neuronal network, represented by the cats, to uphold the representational structure of the environment (Figure 2c).

This perspective suggests a fundamental connection between representational drift and statistical learning. Conceptualizing continuous statistical learning as a safeguard for stable neuronal representations raises the question of what happens when this mechanism fails. Notably, humans tend to avoid sensory input with low empirical entropy, a measure of information content, which correlates with perceived boredom [59]. As a form of negative affect, boredom likely serves as an adaptive drive helping to continuously sustain statistical learning [60]. Supporting this idea, recent research shows that novel, enriched environments help stabilize hippocampal representations [61].

Functional implications of a dynamic substrate for memories

Traditionally, changes in mature, healthy neuronal circuits were thought to only occur in response to external stimuli. However, the emerging view of a continuously remodeling neural network challenges classical models of memory storage and forgetting, which assume memories are imprinted onto a largely static structure [62]. Future research should explore how representational drift influences long-term memory and forgetting.

Two classical theories offer complementary perspectives on forgetting, both aligning with our framework: trace decay theory suggests memories fade over time [63–65], while interference theory posits that memories degrade when replaced by new information [66,67]. These theories parallel the mechanisms of representational drift—stochastic synaptic changes may drive time-dependent decay, whereas Hebbian plasticity may underlie interference by integrating new experiences [37,38,47].

Rather than a mere failure of memory retention, forgetting may serve an adaptive role, enabling neuronal circuits to clear outdated information and optimize encoding of relevant experiences, ultimately enhancing adaptability to new environments.

Furthermore, we speculate that the interplay between stochastic changes and statistical learning might also play a role in spontaneously generating novel associations between neural representations, a process often linked to creativity [68]. Such generation of new connections is thought to emerge through a combination of variability and selection of viable solutions [69]. In our framework, stochastic changes could serve to spontaneously generate novel combinations of neural activity patterns, potentially providing a neural substrate for creative processes [70].

Summary

Together, we here propose a framework in which the inherent instability of synapses is balanced by synaptic changes induced by statistical learning, helping to preserve the functionally relevant aspects of neural representations by continuously reinforcing the statistical structure of the environment. Understanding this interplay between representational drift and statistical learning may shed light on the basis underlying the intrinsically dynamic nature of neuronal circuits that can be observed even under stable environmental conditions. This fundamental property is likely advantageous for adaptive behavior and continuous learning in changing conditions and superior to an otherwise static system in which plastic adaptations occur solely as a reaction to outside triggers.

Author contributions

Jens-Bastian Eppler: Conceptualization, Writing – Original Draft, Visualization. Matthias Kaschube: Conceptualization, Writing – Reviewing and Editing, Funding acquisition. Simon Rumpel: Conceptualization, Writing – Original Draft, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by research grants Deutsche Forschungsgemeinschaft SPP 2041 Project #347573108, Deutsche Forschungsgemeinschaft/Agence nationale de la recherche Project #431393205, Deutsche Forschungsgemeinschaft DIP "Neurobiology of Forgetting". This research was supported in part by the National Science Foundation Grant No. NSF PHY-1748958 and the Gordon and Betty Moore Foundation Grant No. 2919.02.

The authors would like to thank the members of the Rumpel and Kaschube labs and Noam Ziv and Alex Roxin for helpful discussions on the manuscript.

Data availability

Data used for the illustration in Figure 1 is publicly available from the G-node repository: https://doi.org/10.12751/g-node.trwj8c.

References

Papers of particular interest, published within the period of review, have been highlighted as:

* of special interest
** of outstanding interest

1. Rokni U, Richardson AG, Bizzi E, Seung HS: **Motor learning with unstable neural representations**. *Neuron* 2007, **54**: 653–666, https://doi.org/10.1016/j.neuron.2007.04.030.

2. Margolis DJ, Lütcke H, Schulz K, Haiss F, Weber B, Kügler S, Hasan MT, Herlmchen F: **Reorganization of cortical population activity imaged throughout longterm sensory deprivation**. *Nat Neurosci* 2012, **15**:1539–1546, https://doi.org/10.1038/nn.3240.

3. Driscoll LN, Pettit NL, Minderer M, Chettih SN, Harvey CD: **Dynamic reorganization of neuronal activity patterns in parietal cortex**. *Cell* 2017, **170**:986–999, https://doi.org/10.1016/j.cell.2017.07.021.

4. ** Schoonover CE, Ohashi SN, Axel R, Fink AJP: **Representational drift in primary olfactory cortex**. *Nature* 2021, **594**: 541–546, https://doi.org/10.1038/s41586-021-03628-7.
Highly influential paper coining the term 'representational drift.' Apart from describing drift in the mouse olfactory system, the authors show that repeated stimulus presentation can lead to a slowdown of drift in activity patterns.

5. Deitch D, Rubin A, Ziv Y: **Representational drift in the mouse visual cortex**. *Curr Biol* 2021, **31**:4327–4339, https://doi.org/10.1016/j.cub.2021.07.062.

6. Aschauer DF, Eppler J-B, Ewig L, Chambers AR, Pokorny C, Kaschube M, Rumpel S: **Learning-induced biases in the ongoing dynamics of sensory representations predict stimulus generalization**. *Cell Rep* 2022, **38**, 110340, https://doi.org/10.1016/j.celrep.2022.110340.

7. Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK: **Neuronal code for extended time in the hippocampus**. *Proc Natl Acad Sci* 2012, **109**:19462–19467, https://doi.org/10.1073/pnas.1502758112.

8. Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, Gamal AE, Schnitzer MJ: **Long-term dynamics of CA1 hippocampal place codes**. *Nat Neurosci* 2013, **16**:264–266, https://doi.org/10.1038/nn.3329.

9. Hainmueller T, Bartos M: **Parallel emergence of stable and dynamic memory engrams in the hippocampus**. *Nature* 2018, **558**:292–296, https://doi.org/10.1038/s41586-018-0191-2.

10. Chambers AR, Rumpel S: **A stable brain from unstable components: emerging concepts and implications for neural computation**. *Neuroscience* 2017, **357**:172–184, https://doi.org/10.1016/j.neuroscience.2017.06.005.

11. Shepard RN, Chipman S: **Second-order isomorphism of internal representations: shapes of states**. *Cogn Psychol* 1970, **1**:1–17, https://doi.org/10.1016/0010-0285(70)90002-2.

12. ** Edelman S: **Representation is representation of similarities**. *Behav Brain Sci* 1998, **21**:449–498, https://doi.org/10.1017/s0140525x98001253.
Landmark paper linking similarity of neuronal activity patterns and perceptual similarity, thereby introducing the concept of representational similarity.

13. Behrens TEJ, Muller TH, Whittington JCR, Mark S, Baram AB, Stachenfeld KL, Kurth-Nelson Z: **What is a cognitive map? Organizing knowledge for flexible behavior**. *Neuron* 2018, **100**:490–509, https://doi.org/10.1016/j.neuron.2018.10.002.

14. Kriegeskorte N, Wei XX: **Neural tuning and representational geometry**. *Nat Rev Neurosci* 2021, **22**:703–718, https://doi.org/10.1038/s41583-021-00502-3.

15. Noda T, Aschauer DF, Chambers AR, Seiler JP, Rumpel S: **Representational maps in the brain: concepts, approaches, and applications**. *Front Cell Neurosci* 2024, **18**, 1366200, https://doi.org/10.3389/fncel.2024.1366200.

16. Kriegeskorte N, Mur M, Bandettini P: **Representational similarity analysis - connecting the branches of systems neuroscience**. *Front Syst Neurosci* 2008, **2**, https://doi.org/10.3389/neuro.06.004.2008.

17. Kriegeskorte N, Mur M, Ruff DA, Kiani R, Bodurka J, Esteky H, Tanaka K, Bandettini PA: **Matching categorical object representations in inferior temporal cortex of man and monkey**. *Neuron* 2008, **60**:1126–1141, https://doi.org/10.1016/j.neuron.2008.10.043.

18. * Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE: **Long-term stability of cortical population dynamics underlying consistent behavior**. *Nat Neurosci* 2020, **23**:260–270, https://doi.org/10.1038/s41593-019-0555-4.
Even though activity patterns in monkey motor cortex change with time, neural representational similarities are remarkably stable.

19. Xia J, Marks TD, Goard MJ, Wessel R: **Stable representation of a naturalistic movie emerges from episodic activity with gain variability**. *Nat Commun* 2021, **12**:5170, https://doi.org/10.1038/s41467-021-25437-2.

20. Rubin A, Sheintuch L, Pinchasof O, Rechavi Y, Geva N, Ziv Y: **Revealing neural correlates of behavior without behavioral measurements**. *Nat Commun* 2019, **10**:4745, https://doi.org/10.1038/s41467-019-12724-2.

21. Sheintuch L, Geva N, Baumer H, Rechavi Y, Rubin A, Ziv Y: **Multiple maps of the same spatial context can stably coexist in the mouse hippocampus**. *Curr Biol* 2020, **30**:1467–1476, https://doi.org/10.1016/j.cub.2020.02.018.

22. * Noda T, Kienle E, Eppler JB, *et al*.: **Homeostasis of a representational map in the neocortex**. *Nat Neurosci* 2025, **28**: 1533–1545. https://doi.org/10.1038/s41593-025-01982-7.
Maintenance of representational similarities in mouse auditory cortex despite drift in neuronal population activity vectors. Representational similarity is actively restructured within days after targeted neuron ablation.

23. ** Yasumatsu N, Matsuzaki M, Miyazaki T, Noguchi J, Kasai H: **Principles of long-term dynamics of dendritic spines**. *J Neurosci Off J Soc Neurosci* 2008, **28**:13592–13608, https://doi.org/10.1523/JNEUROSCI.0603-08.2008.
Early paper demonstrating stochastic changes in synapses independent from neural activity.

24. Loewenstein Y, Kuras A, Rumpel S: **Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo**. *J Neurosci* 2011, **31**: 9481–9488, https://doi.org/10.1523/JNEUROSCI.6130-10.2011.

25. Attardo A, Fitzgerald JE, Schnitzer MJ: **Impermanence of dendritic spines in live adult CA1 hippocampus**. *Nature* 2015, **523**:592–596, https://doi.org/10.1038/nature14467.

26. Kasai H, Ziv NE, Okazaki H, Yagishita S, Toyoizumi T: **Spine dynamics in the brain, mental disorders and artificial neural networks**. *Nat Rev Neurosci* 2021, **22**:407–422, https://doi.org/10.1038/s41583-021-00467-3.

27. Haimerl C, Machens C: **Representational drift without synaptic plasticity** [Poster presented at Cosyne 2023]. In *Poster 1-136*; 2023. Cosyne 2023.

28. Delamare G, Zaki Y, Cai DJ, Clopath C: **Drift of neural ensembles driven by slow fluctuations of intrinsic excitability**. *eLife* 2024, **12**, https://doi.org/10.7554/eLife.88053.

29. Nicoll RA: **A brief history of long-term potentiation**. *Neuron* 2017, **93**:281–290, https://doi.org/10.1016/j.neuron.2016.12.015.

30. Brzosko Z, Mierau SB, O P: **Neuromodulation of spike-timing-dependent plasticity: past, present, and future**. *Neuron* 2019, **103**:563–581, https://doi.org/10.1016/j.neuron.2019.05.041.

31. Rule ME, O'Leary T: **Self-healing codes: how stable neural populations can track continually reconfiguring neural representations**. *Proc Natl Acad Sci U S A* 2022, **119**, e2106692119, https://doi.org/10.1073/pnas.2106692119.

32. Minerbi A, Kahana R, Goldfeld L, Kaufman M, Marom S, Ziv NE: **Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity**. *PLoS Biol* 2009, **7**, e1000136, https://doi.org/10.1371/journal.pbio.1000136.

33. Rubinski A, Ziv NE: **Remodeling and tenacity of inhibitory synapses: relationships with network activity and neighboring excitatory synapses**. *PLoS Comput Biol* 2015, **11**, e1004632, https://doi.org/10.1371/journal.pcbi.1004632.

34. Mongillo G, Rumpel S, Loewenstein Y: **Inhibitory connectivity defines the realm of excitatory plasticity**. *Nat Neurosci* 2018, **21**:1463–1470, https://doi.org/10.1038/s41593-018-0226-x.

35. Acker D, Paradis S, Miller P: **Stable memory and computation in randomly rewiring neural networks**. *J Neuro Physiol* 2019, **122**:66–80, https://doi.org/10.1152/jn.00534.2018.

36. Kossio YFK, Goedeke S, Klos C, Memmesheimer R: **Drifting assemblies for persistent memory: neuron transitions and unsupervised compensation**. *Proc Natl Acad Sci* 2021, **118**, https://doi.org/10.1073/pnas.2023832118.

37. * Devalle F, Zou L, Cecchini G, Roxin A: **Representational drift as the consequence of ongoing memory storage**. *bioRxiv* 2024, https://doi.org/10.1101/2024.06.25.600729.
Ongoing Hebbian learning of stochastically distributed stimuli is sufficient to account for different phenomena of representational drift, e.g. the observed stabilization after repeated stimulus presentation and the different effects of time and experience on a neuron's activity and tuning, respectively.

38. * Geva N, Deitch D, Rubin A, Ziv Y: **Time and experience differentially affect distinct aspects of hippocampal representational drift**. *Neuron* 2023, **111**:2357–2366, https://doi.org/10.1016/j.neuron.2023.05.005.
Place cell tuning in mouse hippocampus changes with experience, whereas elapsed time primarily affects the drop-in and drop-out of individual neurons.

39. Ratzon A, Derdikman D, Barak O: **Representational drift as a result of implicit regularization**. *eLife* 2024, **12**, RP90069, https://doi.org/10.7554/eLife.90069.

40. Qin S, Farashahi S, Lipshutz D, Sengupta AM, Chklovskii DB, Pehlevan C: **Coordinated drift of receptive fields in hebbian/anti-Hebbian network models during noisy representation learning**. *Nat Neurosci* 2023, **26**:339–349, https://doi.org/10.1038/s41593-022-01225-z.

41. Driscoll LN, Duncker L, Harvey CD: **Representational drift: emerging theories for continual learning and experimental future directions**. *Curr Opin Neurobiol* 2022, **76**, 102609, https://doi.org/10.1016/j.conb.2022.102609.

42. Duncker L, Driscoll L, Shenoy KV, Sahani M, Sussillo D: **Organizing recurrent network dynamics by task-computation to enable continual learning**. *Adv Neural Inf Process Syst* 2020.

43. Johnston WJ, Fusi S: **Abstract representations emerge naturally in neural networks trained to perform multiple tasks**. *Nat Commun* 2023, **14**, https://doi.org/10.1038/s41467-023-36583-0.

44. Humble J, Hiratsuka K, Kasai H, Toyoizumi T: **Intrinsic spine dynamics are critical for recurrent network learning in models with and without autism spectrum disorder**. *Front Comput Neurosci* 2019, **13**:38, https://doi.org/10.3389/fncom.2019.00038.

45. * Raman DV, O'Leary T: **Optimal plasticity for memory maintenance during ongoing synaptic change**. *eLife* 2021, **10**, e62912, https://doi.org/10.7554/elife.62912.
Elegant theoretical work demonstrating benefits of stochastic synaptic changes for memory storage during continual learning.

46. Bauer J, Lewin U, Herbert E, Gjorgijeva J, Schoonover CE, Fink AJP, Rose T, Bonhoeffer T, Hübener M: **Sensory experience steers representational drift in mouse visual cortex**. *Nat Commun* 2024, **15**:9153, https://doi.org/10.1038/s41467-024-53326-x.

47. Eppler J-B, Lai T, Aschauer D, Rumpel S, Kaschube M: **Representational drift reflects ongoing balancing of stochastic changes by Hebbian learning**. *bioRxiv* 2025, https://doi.org/10.1101/2025.01.05.631363.

48. Ziv NE, Brenner N: **Synaptic tenacity or lack thereof: spontaneous remodeling of synapses**. *Trends Neurosci* 2018, **41**:89–99, https://doi.org/10.1016/j.tins.2017.12.003.

49. Kim R, Seitz A, Feenstra H, Shams L: **Testing assumptions of statistical learning: is it long-term and implicit?** *Neurosci Lett* 2009, **461**:145–149, https://doi.org/10.1016/j.neulet.2009.06.030.

50. Fiser J, Aslin RN: **Statistical learning of higher-order temporal structure from visual shape sequences**. *J Exp Psychol Learn Mem Cognit* 2002, **28**:458, https://doi.org/10.1037//0278-7393.28.3.458.

51. Karuza EA, Emberson LL, Roser ME, Cole D, Anslin RN, Fiser J: **Neural signatures of spatial statistical learning: characterizing the extraction of structure from complex visual scenes**. *J Cognit Neurosci* 2017, **29**:1963–1976, https://doi.org/10.1162/jocn_a_01182.

52. Saffran JR, Aslin RN, Newport EL: **Statistical learning by 8-Month-Old infants**. *Science* 1996, **274**:1926–1928, https://doi.org/10.1126/science.274.5294.1926.

53. Fiser J, Aslin RN: **Unsupervised statistical learning of higher-order spatial structures from visual scenes**. *Psychol Sci* 2001, **12**:499–504, https://doi.org/10.1111/1467-9280.00392.

54. Conway CM, Christiansen MH: **Modality-constrained statistical learning of tactile, visual, and auditory sequences**. *J Exp Psychol Learn Mem Cogn* 2005, **31**:24–39, https://doi.org/10.1037/0278-7393.31.1.24.

55. Orbán G, Fiser J, Aslin RN, Lengyel M: **Bayesian learning of visual chunks by human observers**. *Proc Natl Acad Sci* 2008, **105**:2745–2750, https://doi.org/10.1073/pnas.0708424105.

56. Kirkham NZ, Slemmer JA, Johnson SP: **Visual statistical learning in infancy: evidence for a domain general learning mechanism**. *Cognition* 2002, **83**:B35–B42, https://doi.org/10.1016/s0010-0277(02)00004-5.

57. Turk-Browne NB, Scholl BJ, Chun MM, Johnson MK: **Neural evidence of statistical learning: efficient detection of visual regularities without awareness**. *J Cognit Neurosci* 2009, **21**:1934–1945, https://doi.org/10.1162/jocn.2009.21131.

58. Vapnik V: *The nature of statistical learning theory*. New York: Springer; 1999, https://doi.org/10.1007/978-1-4757-3264-1.

59. Seiler JP, Dan O, Tüscher O, Loewenstein Y, Rumpel S: **Experienced entropy drives choice behavior in a boring decision-making task**. *Sci Rep* 2022, **12**:3162, https://doi.org/10.1038/s41598-022-06861-w.

60. Seiler JP, Dan O: **Boredom and curiosity: the hunger and the appetite for information**. *Front Psychol* 2024, **15**, 1514348, https://doi.org/10.3389/fpsyg.2024.1514348.

61. Elyasaf G, Rubin A, Ziv Y: **Novel off-context experience constrains hippocampal representational drift**. *Curr Biol* 2024, **34**:5769–5773. https://10.1016/j.cub.2024.10.027.

62. Mongillo G, Rumpel S, Loewenstein Y: **Intrinsic volatility of synaptic connections - a challenge to the synaptic trace theory of memory**. *Curr Opin Neurobiol* 2017, **46**:7–13, https://doi.org/10.1016/j.conb.2017.06.006.

63. Brown J: **Some tests of the decay theory of immediate memory**. *Q J Exp Psychol* 1958, **10**:12–21, https://doi.org/10.1080/17470215808416249.

64. Peterson LR, Peterson MJ: **Short-term retention of individual verbal items**. *J Exp Psychol* 1959, **58**:193–198, https://doi.org/10.1080/17470215808416249.

65. Ricker TJ, Vergauwe E, Cowan N: **Decay theory of immediate memory: from Brown (1958) to today (2014)**. *Q J Exp Psychol* 2016, **69**:1969–1995, https://doi.org/10.1080/17470218.2014.914546.

66. Briggs GE: **Acquisition, extinction, and recovery functions in retroactive inhibition**. *J Exp Psychol* 1954, **47**:285–293, https://doi.org/10.1037/h0060251.

67. Robertson EM: **New insights in human memory interference and consolidation**. *Curr Biol* 2012, **22**:R66–R71, https://doi.org/10.1016/j.cub.2011.11.051.

68. Mednick SA: **The associative basis of the creative process**. *Psychol Rev* 1962, **69**:220–232, https://doi.org/10.1037/h0048850.

69. Campbell DT: **Blind variation and selective retention in creative thought as in other knowledge processes**. *Psychol Rev* 1960, **67**:380–400, https://doi.org/10.1037/h0040373.

70. Seiler JP, Eppler J-B, Dan O, Elpelt J, Kaschube M, Rumpel S: **Towards circuit mechanisms of the creative process: describing the functions, mechanisms and neural correlates of creativity**. *PsyArXiv* 2025, https://doi.org/10.31234/osf.io/mpbgy_v1.