Brain-like functional specialization emerges spontaneously in deep neural networks

Katharina Dobs, Julio Martinez, Alexander J. E. Kell and Nancy Kanwisher

Science advances (2022)

The human brain contains multiple regions with distinct, often highly specialized functions, from recognizing faces to understanding language to thinking about what others are thinking. However, it remains unclear why the cortex exhibits this high degree of functional specialization in the first place. Here, we consider the case of face perception using artificial neural networks to test the hypothesis that functional segregation of face recognition in the brain reflects a computational optimization for the broader problem of visual recognition of faces and other visual categories. We find that networks trained on object recognition perform poorly on face recognition and vice versa and that networks optimized for both tasks spontaneously segregate themselves into separate systems for faces and objects. We then show functional segregation to varying degrees for other visual categories, revealing a widespread tendency for optimization (without built-in task-specific inductive biases) to lead to functional specialization in machines and, we conjecture, also brains.

Link :DOI: 10.1126/sciadv

Rotational dynamics in motor cortex are consistent with a feedback controller

Hari Teja Kalidind, Kevin P Cross, Timothy P Lillicrap , Mohsen Omrani, Egidio Falotico, Philip N Sabes and Stephen H Scott

elife (2021)

Recent studies have identified rotational dynamics in motor cortex (MC), which many assume arise from intrinsic connections in MC. However, behavioral and neurophysiological studies suggest that MC behaves like a feedback controller where continuous sensory feedback and interactions with other brain areas contribute substantially to MC processing. We investigated these apparently conflicting theories by building recurrent neural networks that controlled a model arm and received sensory feedback from the limb. Networks were trained to counteract perturbations to the limb and to reach toward spatial targets. Network activities and sensory feedback signals to the network exhibited rotational structure even when the recurrent connections were removed. Furthermore, neural recordings in monkeys performing similar tasks also exhibited rotational structure not only in MC but also in somatosensory cortex. Our results argue that rotational structure may also reflect dynamics throughout the voluntary motor system involved in online control of motor actions.

Link : https://doi.org/10.7554/eLife.67256

Synapse-type-specific competitive Hebbian learning forms functional recurrent networks

Samuel Eckmann, Julijana Gjorgjieva

bioRxiv (2022)

Cortical networks exhibit complex stimulus-response patterns. Previous work has identified the balance between excitatory and inhibitory currents as a central component of cortical computations, but has not considered how the required synaptic connectivity emerges from biologically plausible plasticity rules. Using theory and modeling, we demonstrate how a wide range of cortical response properties can arise from Hebbian learning that is stabilized by the synapse-type-specific competition for synaptic resources. In fully plastic recurrent circuits, this competition enables the development and decorrelation of inhibition-balanced receptive fields. Networks develop an assembly structure with stronger connections between similarly tuned neurons and exhibit response normalization and surround suppression. These results demonstrate how neurons can self-organize into functional circuits and provide a foundational understanding of plasticity in recurrent networks.

Link : https://www.biorxiv.org/content/10.1101/2022.03.11.483899v1

Dimensionality reduction of calcium-imaged neuronal population activity

Tze Hui Koh, William E. Bishop, Takashi Kawashima, Brian B. Jeon, Ranjani Srinivasan, Sandra J. Kuhlman, Misha B. Ahrens, Steven M. Chase and Byron M. Yu

bioRxiv (2022)

Calcium imaging has been widely adopted for its ability to record from large neuronal populations. To summarize the time course of neural activity, dimensionality reduction methods, which have been applied extensively to population spiking activity, may be particularly useful. However, it is unclear if the dimensionality reduction methods applied to spiking activity are appropriate for calcium imaging. We thus carried out a systematic study of design choices based on standard dimensionality reduction methods. We also developed a novel method to perform deconvolution and dimensionality reduction simultaneously (termed CILDS). CILDS most accurately recovered the single-trial, low-dimensional time courses from calcium imaging that would have been recovered from spiking activity. CILDS also outperformed the other methods on calcium imaging recordings from larval zebrafish and mice.
More broadly, this study represents a foundation for summarizing calcium imaging recordings of large neuronal populations using dimensionality reduction in diverse experimental settings

Link : https://www.biorxiv.org/content/10.1101/2022.03.11.480682v1.full.pdf

Neurons learn by predicting future activity

Luczak, Artur, Bruce L. McNaughton and Yoshimasa Kubo.

Nature Machine Intelligence (2022):

Understanding how the brain learns may lead to machines with human-like intellectual capacities. It was previously proposed that the brain may operate on the principle of predictive coding. However, it is still not well understood how a predictive system could be implemented in the brain. Here we demonstrate that the ability of a single neuron to predict its future activity may provide an effective learning mechanism. Interestingly, this predictive learning rule can be derived from a metabolic principle, whereby neurons need to minimize their own synaptic activity (cost) while maximizing their impact on local blood supply by recruiting other neurons. We show how this mathematically derived learning rule can provide a theoretical connection between diverse types of brain-inspired algorithm, thus offering a step towards the development of a general theory of neuronal learning. We tested this predictive learning rule in neural network simulations and in data recorded from awake animals. Our results also suggest that spontaneous brain activity provides 'training data' for neurons to learn to predict cortical dynamics. Thus, the ability of a single neuron to minimize surprise—that is, the difference between actual and expected activity—could be an important missing element to understand computation in the brain.

Link : https://www.nature.com/articles/s42256-021-00430-y.pdf

The Geometry of Representational Drift in Natural and Artificial Neural Networks

Kyle Aitken, Marina Garrett, Shawn Olsen, Stefan Mihalas

bioRxiv (2021)

Neurons in sensory areas encode/represent stimuli. Surprisingly, recent studies have suggested that, even during persistent performance, these representations are not stable and change over the course of days and weeks. We examine stimulus representations from fluorescence recordings across hundreds of neurons in the visual cortex using in vivo two-photon calcium imaging and we corroborate previous studies finding that such representations change as experimental trials are repeated across days. This phenomenon has been termed "representational drift". In this study we geometrically characterize the properties of representational drift in the primary visual cortex of mice in two open datasets from the Allen Institute and propose a potential mechanism behind such drift. We observe representational drift both for passively presented stimuli, as well as for stimuli which are behaviorally relevant. Across experiments, the drift most often occurs along directions that have the most variance, leading to a significant turnover in the neurons used for a given representation. Interestingly, despite this significant change due to drift, linear classifiers trained to distinguish neuronal representations show little to no degradation in performance across days. The features we observe in the neural data are similar to properties of artificial neural networks where representations are updated by continual learning in the presence of dropout, i.e. a random masking of nodes/weights, but not other types of noise. Therefore, we conclude that a potential reason for the representational drift in biological networks is driven by an underlying dropout-like noise while continuously learning and that such a mechanism may be computational advantageous for the brain in the same way it is for artificial neural networks, e.g. preventing overfitting.

Link: https://doi.org/10.1101/2021.12.13.472494

Does the brain care about averages? A simple test

A. Tlaie, K. A. Shapcott, P. Tiesinga, M. L. Schölvinck, and M. N. Havenith
bioRxiv (2021)

Trial-averaged metrics, e.g. in the form of tuning curves and population response vectors, are a basic and widely accepted way of characterizing neuronal activity. But how relevant are such trial-averaged responses to neuronal computation itself? Here we present a simple test to estimate whether average responses reflect aspects of neuronal activity that contribute to neuronal processing in a specific context. The test probes two assumptions inherent in the usage of average neuronal metrics:

1. Reliability: Neuronal responses repeat consistently enough across single stimulus instances that the average response template they relate to remains recognizable to downstream regions.
2. Behavioural relevance: If a single-trial response is more similar to the average template, this should make it easier for the animal to identify the correct stimulus or action.

We apply this test to a large publicly available data set featuring electrophysiological recordings from 42 cortical areas in behaving mice. In this data set, we show that single-trial responses were less correlated to the average response template than one would expect if they simply represented discrete versions of the template, down-sampled to a finite number of spikes. Moreover, single-trial responses were barely stimulus-specific – they could not be clearly assigned to the average response template of one stimulus. Most importantly, better-matched single-trial responses did not predict accurate behaviour for any of the recorded cortical areas. We conclude that in this data set, average responses do not seem particularly relevant to neuronal computation in a majority of brain areas, and we encourage other researchers to apply similar tests when using trial-averaged neuronal metrics.

Link : https://www.biorxiv.org/content/10.1101/2021.11.28.469673v1

The information theory of developmental pruning: Optimizing global network architectures using local synaptic rules

Scholl, Carolin, Michael E. Rule, and Matthias H. Hennig.
PLoS computational biology 17.10 (2021)

During development, biological neural networks produce more synapses and neurons than needed. Many of these synapses and neurons are later removed in a process known as neural pruning. Why networks should initially be over-populated, and the processes that determine which synapses and neurons are ultimately pruned, remains unclear. We study the mechanisms and significance of neural pruning in model neural networks. In a deep Boltzmann machine model of sensory encoding, we find that (1) synaptic pruning is necessary to learn efficient network architectures that retain computationally-relevant connections, (2) pruning by synaptic weight alone does not optimize network size and (3) pruning based on a locally-available measure of importance based on Fisher information allows the network to identify structurally important vs. unimportant connections and neurons. This locally-available measure of importance has a biological interpretation in terms of the correlations between presynaptic and postsynaptic neurons, and implies an efficient activity-driven pruning rule. Overall, we show how local activity-dependent synaptic pruning can solve the global problem of optimizing a network architecture. We relate these findings to biology as follows: (I) Synaptic over-production is necessary for activity-dependent connectivity optimization. (II) In networks that have more neurons than needed, cells compete for activity, and only the most important and selective neurons are retained. (III) Cells may also be pruned due to a loss of synapses on their axons. This occurs when the information they convey is not relevant to the target population.

Link : https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1009458

Self-healing codes: How stable neural populations can track continually reconfiguring neural representations

Rule, Michael E., and Timothy O'Leary.

Proceedings of the National Academy of Sciences 119.7 (2022)

As an adaptive system, the brain must retain a faithful representation of the world while continuously integrating new information. Recent experiments have measured population activity in cortical and hippocampal circuits over many days and found that patterns of neural activity associated with fixed behavioral variables and percepts change dramatically over time. Such "representational drift" raises the question of how malleable population codes can interact coherently with stable long-term representations that are found in other circuits and with relatively rigid topographic mappings of peripheral sensory and motor signals. We explore how known plasticity mechanisms can allow single neurons to reliably read out an evolving population code without external error feedback. We find that interactions between Hebbian learning and single-cell homeostasis can exploit redundancy in a distributed population code to compensate for gradual changes in tuning. Recurrent feedback of partially stabilized readouts could allow a pool of readout cells to further correct inconsistencies introduced by representational drift. This shows how relatively simple, known mechanisms can stabilize neural tuning in the short term and provides a plausible explanation for how plastic neural codes remain integrated with consolidated, long-term representations.

https://doi.org/10.1073/pnas.2106692119

The impact of sparsity in low-rank recurrent neural networks

Herbert, Elisabeth and Ostojic,Srdjan.

(2022) *bioRxiv*.

Neural population dynamics are often highly coordinated, allowing task-related computations to be understood as neural trajectories through low-dimensional subspaces. How the network connectivity and input structure give rise to such activity can be investigated with the aid of low-rank recurrent neural networks, a recently-developed class of computational models which offer a rich theoretical framework linking the underlying connectivity structure to emergent low-dimensional dynamics. This framework has so far relied on the assumption of all-to-all connectivity, yet cortical networks are known to be highly sparse. Here we investigate the dynamics of low-rank recurrent networks in which the connections are randomly sparsified, which makes the network connectivity formally full-rank. We first analyse the impact of sparsity on the eigenvalue spectrum of low-rank connectivity matrices, and use this to examine the implications for the dynamics. We find that in the presence of sparsity, the eigenspectral in the complex plane consist of a continuous bulk and isolated outliers, a form analogous to the eigenspectra of connectivity matrices composed of a low-rank and a full-rank random component. This analogy allows us to characterise distinct dynamical regimes of the sparsified low-rank network as a function of key network parameters. Altogether, we find that the low-dimensional dynamics induced by low-rank connectivity structure are preserved even at high levels of sparsity, and can therefore support rich and robust computations even in networks sparsified to a biologically-realistic extent.

Link : https://www.biorxiv.org/content/10.1101/2022.03.31.486515v1.full.pdf

Detecting and correcting false transients in calcium imaging

Jeffrey L. Gauthier , Sue Ann Koay , Edward H. Nieh , David W. Tank ,Jonathan W. Pillow and Adam S. Charles
Nature Methods (2022)

Population recordings of calcium activity are a major source of insight into neural function. Large datasets require automated processing, but this can introduce errors that are difficult to detect. Here we show that popular time course-estimation algorithms often contain substantial misattribution errors affecting 10–20% of transients. Misattribution, in which fluorescence is ascribed to the wrong cell, arises when overlapping cells and processes are imperfectly defined or not identified. To diagnose misattribution, we develop metrics and visualization tools for evaluating large datasets. To correct time courses, we introduce a robust estimator that explicitly accounts for contaminating signals. In one hippocampal dataset, removing contamination reduced the number of place cells by 15%, and 19% of place fields shifted by over 10 cm. Our methods are compatible with other cell-finding techniques, empowering users to diagnose and correct a potentially widespread problem that could alter scientific conclusions
Link : https://rdcu.be/cJ0uF

Stable memory and computation in randomly rewiring neural networks

Acker, D., Paradis, S., and Miller, P
J Neurophysiol 122

Our brains must maintain a representation of the world over a period of time much longer than the typical lifetime of the biological components producing that representation. For example, recent research suggests that dendritic spines in the adult mouse hippocampus are transient with an average lifetime of ~10 days. If this is true, and if turnover is equally likely for all spines, ~95% of excitatory synapses onto a particular neuron will turn over within 30 days; however, a neuron’s receptive field can be relatively stable over this period. Here, we use computational modeling to ask how memories can persist in neural circuits such as the hippocampus and visual cortex in the face of synapse turnover. We demonstrate that Hebbian plasticity during replay of presynaptic activity patterns can integrate newly formed synapses into pre-existing memories. Furthermore, we find that Hebbian plasticity during replay is sufficient to stabilize the receptive fields of hippocampal place cells in a model of the grid-cell-to-place-cell transformation in CA1 and of orientation-selective cells in a model of the center-surround-to-simple-cell transformation in V1. Together, these data suggest that a simple plasticity rule, correlative Hebbian plasticity of synaptic strengths, is sufficient to preserve neural representations in the face of synapse turnover, even in the absence of activity-dependent structural plasticity

Link : https://doi.org/10.1152/jn.00534.2018

Representational drift in the mouse visual cortex

Deitch, D., Rubin, A., and Ziv, Y.
Curr Biol 31

Recent studies have shown that neuronal representations gradually change over time despite no changes in the stimulus, environment, or behavior. However, such representational drift has been assumed to be a property of high-level brain structures, whereas earlier circuits, such as sensory cortices, have been assumed to stably encode information over time. Here, we analyzed large-scale optical and electrophysiological recordings from six visual cortical areas in behaving mice that were repeatedly presented with the same natural movies. Contrary to the prevailing notion, we found representational drift over timescales spanning minutes to days across multiple visual areas, cortical layers, and cell types. Notably, neural-code stability did not reflect the hierarchy of information flow across areas. Although individual neurons showed time-dependent changes in their coding properties, the structure of the relationships between population activity patterns remained stable and stereotypic. Such population-level organization may underlie stable visual perception despite continuous changes in neuronal responses.

Link: https://doi.org/10.1016/j.cub.2021.07.062

Self-organized reactivation maintains and reinforces memories despite synaptic turnover

Michael Jan Fauth, Mark CW van Rossum

e Life 8

Long-term memories are believed to be stored in the synapses of cortical neuronal networks. However, recent experiments report continuous creation and removal of cortical synapses, which raises the question how memories can survive on such a variable substrate. Here, we study the formation and retention of associative memory in a computational model based on Hebbian cell assemblies in the presence of both synaptic and structural plasticity. During rest periods, such as may occur during sleep, the assemblies reactivate spontaneously, reinforcing memories against ongoing synapse removal and replacement. Brief daily reactivations during rest-periods suffice to not only maintain the assemblies, but even strengthen them, and improve pattern completion, consistent with offline memory gains observed experimentally. While the connectivity inside memory representations is strengthened during rest phases, connections in the rest of the network decay and vanish thus reconciling apparently conflicting hypotheses of the influence of sleep on cortical connectivity.

Link : https://doi.org/10.7554/eLife.43717

The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single Stimulus-Tuned Attractor Account for Patterns of Noise Variability

Guillaume Hennequin Yashar Ahmadian Daniel B. Rubin Máté Lengyel Kenneth D. Miller

In: Neuron 98

Correlated variability in cortical activity is ubiquitously quenched following stimulus onset, in a stimulus-dependent manner. These modulations have been attributed to circuit dynamics involving either multiple stable states ("attractors") or chaotic activity. Here we show that a qualitatively different dynamical regime, involving fluctuations about a single, stimulus-driven attractor in a loosely balanced excitatory-inhibitory network (the stochastic "stabilized supralinear network"), best explains these modulations. Given the supralinear input/output functions of cortical neurons, increased stimulus drive strengthens effective network connectivity. This shifts the balance from interactions that amplify variability to suppressive inhibitory feedback, quenching correlated variability around more strongly driven steady states. Comparing to previously published and original data analyses, we show that this mechanism, unlike previous proposals, uniquely accounts for the spatial patterns and fast temporal dynamics of variability suppression. Specifying the cortical operating regime is key to understanding the computations underlying perception.

Link : https://doi.org/10.1016/j.neuron.2018.04.017

Drifting assemblies for persistent memory: Neuron transitions and unsupervised compensation
Yaroslav Felipe Kalle Kossio, Sven Goedeke,Christian Klos and Raoul-Martin Memmesheimer
PNAS

Change is ubiquitous in living beings. In particular, the connectome and neural representations can change. Nevertheless, behaviors and memories often persist over long times. In a standard model, associative memories are represented by assemblies of strongly interconnected neurons. For faithful storage these assemblies are assumed to consist of the same neurons over time. Here we propose a contrasting memory model with complete temporal remodeling of assemblies, based on experimentally observed changes of synapses and neural representations. The assemblies drift freely as noisy autonomous network activity and spontaneous synaptic turnover induce neuron exchange. The gradual exchange allows activity-dependent and homeostatic plasticity to conserve the representational structure and keep inputs, outputs, and assemblies consistent. This leads to persistent memory. Our findings explain recent experimental results on temporal evolution of fear memory representations and suggest that memory systems need to be understood in their completeness as individual parts may constantly change.

Link : https://doi.org/10.1073/pnas.2023832118

Dynamic and reversible remapping of network representations in an unchanging environment
Isabel I.C. Low,Alex H. Williams ,Malcolm G. Campbell ,Scott W. Linderman and Lisa M. Giocomo
Neuron 109

Neurons in the medial entorhinal cortex alter their firing properties in response to environmental changes. This flexibility in neural coding is hypothesized to support navigation and memory by dividing sensory experience into unique episodes. However, it is unknown how the entorhinal circuit as a whole transitions between different representations when sensory information is not delineated into discrete contexts. Here we describe rapid and reversible transitions between multiple spatial maps of an unchanging task and environment. These remapping events were synchronized across hundreds of neurons, differentially affected navigational cell types, and correlated with changes in running speed. Despite widespread changes in spatial coding, remapping comprised a translation along a single dimension in population-level activity space, enabling simple decoding strategies. These findings provoke reconsideration of how the medial entorhinal cortex dynamically represents space and suggest a remarkable capacity of cortical circuits to rapidly and substantially reorganize their neural representations.

Link : https://doi.org/10.1016/j.neuron.2021.07.005

Multiple Maps of the Same Spatial Context Can Stably Coexist in the Mouse Hippocampus
Liron Sheintuch , Nitzan Geva , Hadas Baumer,Yoav Rechavi,Alon Rubin and Yaniv Ziv
Curr Biol 30.

Hippocampal place cells selectively fire when an animal traverses a particular location and are considered a neural substrate of spatial memory. Place cells were shown to change their activity patterns (remap) across different spatial contexts but to maintain their spatial tuning in a fixed familiar context. Here, we show that mouse hippocampal neurons can globally remap, forming multiple distinct representations (maps) of the same familiar environment, without any apparent changes in sensory input or behavior. Alternations between maps occurred only across separate visits to the environment, implying switching between distinct stable attractors in the hippocampal network. Importantly, the different maps were spatially informative and persistent over weeks, demonstrating that they can be reliably stored and retrieved from long-term memory. Taken together, our results suggest that a memory of a given spatial context could be associated with multiple distinct neuronal representations, rather than just one.

Link : https://doi.org/10.1016/j.cub.2020.02.018

Cross-attractor repertoire provides new perspective on structure-function relationship in the brain

Zhang, Mengsen, Sun, Yinming and Saggar, Manish
bioRxiv.

The brain is a complex system exhibiting ever-evolving activity patterns without any external inputs or tasks. Such intrinsic dynamics (or lack thereof) are thought to play crucial roles in typical as well as atypical cognitive functioning. Linking the ever-changing intrinsic dynamics to the rather static anatomy is a challenging endeavor. Dynamical systems models are important tools for understanding how structure and function are linked in the brain. Here, we provide a novel modeling framework to examine such structure-function relations. Our deterministic approach complements previous modeling frameworks, which typically focus on noise-driven (or stochastic) dynamics near a single attractor. We examine the overall organizations of and coordination between all putative attractors. Using our approach, we first provide evidence that examining cross-attractor coordination between brain regions could better predict human functional connectivity than examining noise-driven near-attractor dynamics. Further, we observed that structural connections across scales modulate the energy costs of such cross-attractor coordination. Overall, our work provides a systematic framework for characterizing intrinsic brain dynamics as a web of cross-attractor transitions and associated energy costs. The framework may be used to predict transitions and energy costs associated with experimental or clinical interventions.

Link : https://doi.org/10.1101/2020.05.14.097196

Probing the relationship between linear dynamical systems and low-rank recurrent neural network models

Valente, Adrian, Srdjan Ostojic, and Jonathan Pillow.
arXiv preprint (2021).
A large body of work has suggested that neural populations exhibit low-dimensional dynamics during behavior. However, there are a variety of different approaches for modeling low-dimensional neural population activity. One approach involves latent linear dynamical system (LDS) models, in which population activity is described by a projection of low-dimensional latent variables with linear dynamics. A second approach involves low-rank recurrent neural networks (RNNs), in which population activity arises directly from a low-dimensional projection of past activity. Although these two modeling approaches have strong similarities, they arise in different contexts and tend to have different domains of application. Here we examine the precise relationship between latent LDS models and linear low-rank RNNs. When can one model class be converted to the other, and vice versa? We show that latent LDS models can only be converted to RNNs in specific limit cases, due to the non-Markovian property of latent LDS models. Conversely, we show that linear RNNs can be mapped onto LDS models, with latent dimensionality at most twice the rank of the RNN.

https://doi.org/10.48550/arXiv.2110.09804