Seminar Computational Neuroscience

Summer Semester 2021

List of papers

Contents

1 Connectomics and human connectome

1.1 Binary and analog variation of synapses between cortical pyramidal neurons.

Dorkenwald, S., Turner, N. L., Macrina, T., Lee, K., Lu, R., Wu, J., ... & Silversmith, W. M. (2019). BioRxiv.

Abstract: Learning from experience depends at least in part on changes in neuronal connections. We present the largest map of connectivity to date between cortical neurons of a defined type (L2/3 pyramidal cells), which was enabled by automated analysis of serial section electron microscopy images with improved handling of image defects. We used the map to identify constraints on the learning algorithms employed by the cortex. Previous cortical studies modeled a continuum of synapse sizes by a log-normal distribution . A continuum is consistent with most neural network models of learning, in which synaptic strength is a continuously graded analog variable. Here we show that synapse size, when restricted to synapses between L2/3 pyramidal cells, is well-modeled by the sum of a binary variable and an analog variable drawn from a log-normal distribution. Two synapses sharing the same presynaptic and postsynaptic cells are known to be correlated in size . We show that the binary variables of the two synapses are highly correlated, while the analog variables are not. Binary variation could be the outcome of a Hebbian or other synaptic plasticity rule depending on activity signals that are relatively uniform across neuronal arbors, while analog variation may be dominated by other influences. We discuss the implications for the stability-plasticity dilemma.

doi: `https://doi.org/10.1101/2019.12.29.890319`

1.2 Similarity in Functional Brain Connectivity at Rest Predicts Interpersonal Closeness in the Social Network of an Entire Village.

Hyon, Ryan, Yoosik Youm, Junsol Kim, Jeanyung Chey, Seyul Kwak, and Carolyn Parkinson. Proceedings of the National Academy of Sciences

Abstract: In what ways are we similar to our friends? Here, we characterized the social network of residents of a remote village, a subset of whom contributed personality and neuroimaging data. We demonstrate that similarity in individuals' resting-state functional connectomes predicts individuals' proximity in their real-world social network, even when controlling for demographic characteristics and self-reported personality traits. Our results suggest that patterns of functional brain activity during rest encode latent similarities (e.g., in terms of how people think and behave) that are associated with friendship. Taken together, integrating neuroimaging and social network analysis can offer novel insights into how the brain shapes or is shaped by the social networks that it inhabits.

doi: `https://doi.org/10.1073/pnas.2013606117`

2 Behavior, Decision making, computation

2.1 Identifying Behavioral Structure from Deep Variational Embeddings of Animal Motion.

Luxem, K., Fuhrmann, F., Kürsch, J., Remy, S., Bauer, P. (2020). bioRxiv.

Abstract: Naturalistic behavior is highly complex and dynamic. Approaches aiming at understanding how neuronal ensembles generate behavior require robust behavioral quantification in order to correlate the neural activity patterns with behavioral motifs. Here, we present Variational Animal Motion Embedding (VAME), a probabilistic machine learning framework for discovery of the latent structure of animal behavior given an input time series obtained from markerless pose estimation tools.

To demonstrate our framework we perform unsupervised behavior phenotyping of APP/PS1 mice, an animal model of Alzheimer disease. Using markerless pose estimates from open-field exploration as input VAME uncovers the distribution of detailed and clearly segmented behavioral motifs. Moreover, we show that the recovered distribution of phenotype-specific motifs can be used to reliably distinguish between APP/PS1 and wildtype mice, while human experts fail to classify the phenotype based on the same video observations. We propose VAME as a versatile and robust tool for unsupervised quantification of behavior across organisms and experimental settings

doi: `https://doi.org/10.1101/2020.05.14.095430`

2.2 Neural Population Dynamics in Motor Cortex Are Different for Reach and Grasp.

Suresh, A. K., Goodman, J. M., Okorokova, E. V., Kaufman, M., Hatsopoulos, N. G., & Bensmaia, S. J. (2020). ELife, 9, e58848.

Abstract: Low-dimensional linear dynamics are observed in neuronal population activity in primary motor cortex (M1) when monkeys make reaching movements. This population-level behavior is consistent with a role for M1 as an autonomous pattern generator that drives muscles to give rise to movement. In the present study, we examine whether similar dynamics are also observed during grasping movements, which involve fundamentally different patterns of kinematics and muscle activations. Using a variety of analytical approaches, we show that M1 does not exhibit such dynamics during grasping movements. Rather, the grasp-related neuronal dynamics in M1 are similar to their counterparts in somatosensory cortex, whose activity is driven primarily by afferent inputs rather than by intrinsic dynamics. The basic structure of the neuronal activity underlying hand control is thus fundamentally different from that underlying arm control.

doi: `https://doi.org/10.7554/eLife.58848`

2.3 Forced choices reveal a trade-off between cognitive effort and physical pain.

Vogel, T. A., Savelson, Z. M., Otto, A. R., & Roy, M. (2020). Elife, 9, e59410.

Abstract: Cognitive effort is described as aversive, and people will generally avoid it when possible. This aversion to effort is believed to arise from a cost–benefit analysis of the actions available. The comparison of cognitive effort against other primary aversive experiences, however, remains relatively unexplored. Here, we offered participants choices between performing a cognitively demanding task or experiencing thermal pain. We found that cognitive effort can be traded off for physical pain and that people generally avoid exerting high levels of cognitive effort. We also used computational modelling to examine the aversive subjective value of effort and its effects on response behaviours. Applying this model to decision times revealed asymmetric effects of effort and pain, suggesting that cognitive effort may not share the same basic influences on avoidance behaviour as more primary aversive stimuli such as physical pain.

doi: `https://doi.org/10.7554/eLife.59410`

2.4 From single neurons to behavior in the jellyfish Aurelia aurita

Pallasdies, F., Goedeke, S., Braun, W., & Memmesheimer, R. M. (2019). Elife, 8, e50084.

Abstract: Jellyfish nerve nets provide insight into the origins of nervous systems, as both their taxonomic position and their evolutionary age imply that jellyfish resemble some of the earliest neuron-bearing, actively-swimming animals. Here, we develop the first neuronal network model for the nerve nets of jellyfish. Specifically, we focus on the moon jelly Aurelia aurita and the control of its energy-efficient swimming motion. The proposed single neuron model disentangles the contributions of different currents to a spike. The network model identifies factors ensuring non-pathological activity and suggests an optimization for the transmission of signals. After modeling the jellyfish's muscle system and its bell in a hydrodynamic environment, we explore the swimming elicited by neural activity. We find

that different delays between nerve net activations lead to well-controlled, differently directed movements. Our model bridges the scales from single neurons to behavior, allowing for a comprehensive understanding of jellyfish neural control of locomotion.

doi: `https://doi.org/10.7554/eLife.50084.`

2.5 A neural basis of probabilistic computation in visual cortex

Walker, E. Y., Cotton, R. J., Ma, W. J., Tolias, A. S. (2020). Nature Neuroscience, 23(1), 122-129.

Abstract: Bayesian models of behavior suggest that organisms represent uncertainty associated with sensory variables. However, the neural code of uncertainty remains elusive. A central hypothesis is that uncertainty is encoded in the population activity of cortical neurons in the form of likelihood functions. We tested this hypothesis by simultaneously recording population activity from primate visual cortex during a visual categorization task in which trial-to-trial uncertainty about stimulus orientation was relevant for the decision. We decoded the likelihood function from the trial-to-trial population activity and found that it predicted decisions better than a point estimate of orientation. This remained true when we conditioned on the true orientation, suggesting that internal fluctuations in neural activity drive behaviorally meaningful variations in the likelihood function. Our results establish the role of population-encoded likelihood functions in mediating behavior and provide a neural underpinning for Bayesian models of perception.

doi: `https://doi.org/10.1038/s41593-019-0554-5`

2.6 Brain computation by assemblies of neurons.

Papadimitriou, C. H., Vempala, S. S., Mitropolsky, D., Collins, M., & Maass, W. (2020). Proceedings of the National Academy of Sciences, 117(25), 14464-14472.

Abstract: Our expanding understanding of the brain at the level of neurons and synapses, and the level of cognitive phenomena such as language, leaves a formidable gap between these two scales. Here we introduce a computational system which promises to bridge this gap: the Assembly Calculus. It encompasses operations on assemblies of neurons, such as project, associate, and merge, which appear to be implicated in cognitive phenomena, and can be shown, analytically as well as through simulations, to be plausibly realizable at the level of neurons and synapses. We demonstrate the reach of this system by proposing a brain architecture for syntactic processing in the production of language, compatible with recent experimental results.

doi: `https://doi.org/10.1073/pnas.2001893117`

3 Sensory systems

3.1 Spontaneous retinal waves generate long-range horizontal connectivity in visual cortex

Kim, J., Song, M., Jang, J., Paik, S. (2020). Journal of Neuroscience.

Abstract: In the primary visual cortex (V1) of higher mammals, long-range horizontal connections (LHCs) are observed to develop, linking iso-orientation domains of cortical tuning. It is unknown how this feature-specific wiring of circuitry develops before eye-opening. Here, we suggest that LHCs in V1 may originate from spatiotemporally structured feedforward activities generated from spontaneous retinal waves. Using model simulations based on the anatomy and observed activity patterns of the retina, we show that waves propagating in retinal mosaics can initialize the wiring of LHCs by coactivating neurons of similar tuning, whereas equivalent random activities cannot induce such organizations. Simulations showed that emerged LHCs can produce the patterned activities observed in V1, matching the topography of the underlying orientation map. The model can also reproduce feature-specific microcircuits in the salt-and-pepper organizations found in rodents. Our results imply that early peripheral activities contribute significantly to cortical development of functional circuits.

doi: `https://doi.org/10.1523/JNEUROSCI.0649-20.2020`

3.2 How many neurons are sufficient for perception of cortical activity?

Dalgleish, H. W., Russell, L. E., Packer, A. M., Roth, A., Gauld, O. M., Greenstreet, F., ... & Häusser, M. (2020). . Elife, 9, e58889.

Abstract: Many theories of brain function propose that activity in sparse subsets of neurons underlies perception and action. To place a lower bound on the amount of neural activity that can be perceived, we used an all-optical approach to drive behaviour with targeted two-photon optogenetic activation of small ensembles of L2/3 pyramidal neurons in mouse barrel cortex while simultaneously recording local network activity with two-photon calcium imaging. By precisely titrating the number of neurons stimulated, we demonstrate that the lower bound for perception of cortical activity is ca. 14 pyramidal neurons. We find a steep sigmoidal relationship between the number of activated neurons and behaviour, saturating at only ca. 37 neurons, and show this relationship can shift with learning. Furthermore, activation of ensembles is balanced by inhibition of neighbouring neurons. This surprising perceptual sensitivity in the face of potent network suppression supports the sparse coding hypothesis, and suggests that cortical perception balances a trade-off between minimizing the impact of noise while efficiently detecting relevant signals.

doi: `https://doi.org/10.7554/eLife.58889`

3.3 Reconstruction of natural images from responses of primate retinal ganglion cells.

Brackbill, N., Rhoades, C., Kling, A., Shah, N. P., Sher, A., Litke, A. M., Chichilnisky, E. J. (2020). Elife, 9, e58516.

Abstract: The visual message conveyed by a retinal ganglion cell (RGC) is often summarized by its spatial receptive field, but in principle also depends on the responses of other RGCs and natural image statistics. This possibility was explored by linear reconstruction of natural images from responses of the four numerically-dominant macaque RGC types. Reconstructions were highly consistent across retinas. The optimal reconstruction filter for each RGC – its visual message – reflected natural image statistics, and resembled the receptive field only when nearby, same-type cells were included. ON and OFF cells conveyed largely independent, complementary representations, and parasol and midget cells conveyed distinct features. Correlated activity and nonlinearities had statistically significant but minor effects on reconstruction. Simulated reconstructions, using linear-nonlinear cascade models of RGC light responses that incorporated measured spatial properties and nonlinearities, produced similar results. Spatiotemporal reconstructions exhibited similar spatial properties, suggesting that the results are relevant for natural vision.

doi: `https://doi.org/10.7554/eLife.58516`

3.4 High-precision coding in visual cortex

Stringer et al., Cell (2021)

Abstract: Individual neurons in visual cortex provide the brain with unreliable estimates of visual features. It is not known whether the single-neuron variability is correlated across large neural populations, thus impairing the global encoding of stimuli. We recorded simultaneously from up to 50,000 neurons in mouse primary visual cortex (V1) and in higher order visual areas and measured stimulus discrimination thresholds of 0.35 deg and 0.37 deg, respectively, in an orientation decoding task. These neural thresholds were almost 100 times smaller than the behavioral discrimination thresholds reported in mice. This discrepancy could not be explained by stimulus properties or arousal states. Furthermore, behavioral variability during a sensory discrimination task could not be explained by neural variability in V1. Instead, behavior-related neural activity arose dynamically across a network of non-sensory brain areas. These results imply that perceptual discrimination in mice is limited by downstream decoders, not by neural noise in sensory representations.

doi: `https://doi.org/10.1016/j.cell.2021.03.042`

4 Learning, representations and coding

4.1 Effective learning is accompanied by high-dimensional and efficient representations of neural activity.

Tang, E., Mattar, M. G., Giusti, C., Lydon-Staley, D. M., Thompson-Schill, S. L., Bassett, D. S. (2019). Nature neuroscience, 22(6), 1000-1009.

Abstract: A fundamental cognitive process is to map value and identity onto the objects we learn about. However, what space best embeds this mapping is not completely understood. Here we develop tools to quantify the space and organization of such a mapping in neural responses as reflected in functional MRI, to show that quick learners have a higher dimensional representation than slow learners, and hence more easily distinguishable whole-brain responses to objects of different value. Furthermore, we find that quick learners display more compact embedding of their neural responses, and hence have higher ratios of their stimuli dimension to their embedding dimension, which is consistent with greater efficiency of cognitive coding. Lastly, we investigate the neurophysiological drivers at smaller scales and study the complementary distinguishability of whole-brain responses. Our results demonstrate a spatial organization of neural responses characteristic of learning and offer geometric measures applicable to identifying efficient coding in higher-order cognitive processes.

doi: `https://doi.org/10.1038/s41593-019-0400-9`

4.2 Fundamental bounds on the fidelity of sensory cortical coding

Rumyantsev, O. I., Lecoq, J. A., Hernandez, O., Zhang, Y., Savall, J., Chrapkiewicz, R., ... & Schnitzer, M. J. (2020). Nature, 580(7801), 100-105.

Abstract: How the brain processes information accurately despite stochastic neural activity is a longstanding question. For instance, perception is fundamentally limited by the information that the brain can extract from the noisy dynamics of sensory neurons. Seminal experiments suggest that correlated noise in sensory cortical neural ensembles is what limits their coding accuracy, although how correlated noise affects neural codes remains debated. Recent theoretical work proposes that how a neural ensemble's sensory tuning properties relate statistically to its correlated noise patterns is a greater determinant of coding accuracy than is absolute noise strength. However, without simultaneous recordings from thousands of cortical neurons with shared sensory inputs, it is unknown whether correlated noise limits coding fidelity. Here we present a 16-beam, two-photon microscope to monitor activity across the mouse primary visual cortex, along with analyses to quantify the information conveyed by large neural ensembles. We found that, in the visual cortex, correlated noise constrained signalling for ensembles with 800–1,300 neurons. Several noise components of the ensemble dynamics grew proportionally to the ensemble size and the encoded visual signals, revealing the predicted information-limiting correlations. Notably, visual signals were perpendicular to the largest noise mode, which therefore did not limit coding fidelity. The information-limiting noise modes were approximately ten times smaller and concordant with mouse visual acuity. Therefore, cortical design principles appear to enhance coding accuracy by restricting around 90% of noise fluctuations to modes that do not limit signalling fidelity, whereas much weaker correlated noise modes inherently bound sensory discrimination.

doi: https://doi.org/10.1038/s41586-020-2130-2

4.3 Learning probabilistic neural representations with randomly connected circuits.

Maoz, O., Tkačik, G., Esteki, M. S., Kiani, R., Schneidman, E. (2020). Proceedings of the National Academy of Sciences, 117(40), 25066-25073.

Abstract: The brain represents and reasons probabilistically about complex stimuli and motor actions using a noisy, spike-based neural code. A key building block for such neural computations, as well as the basis for supervised and unsupervised learning, is the ability to estimate the surprise or likelihood of incoming high-dimensional neural activity patterns. Despite progress in statistical modeling of neural responses and deep learning, current approaches either do not scale to large neural populations or cannot be implemented using biologically realistic mechanisms. Inspired by the sparse and random connectivity of real neuronal circuits, we present a model for neural codes that accurately estimates the likelihood of individual spiking patterns and has a straightforward, scalable, efficient, learnable, and realistic neural implementation. This model's performance on simultaneously recorded spiking activity of ¿100 neurons in the monkey visual and prefrontal cortices is comparable with or better than that of state-of-the-art models. Importantly, the model can be learned using a small number of samples and using a local learning rule that utilizes noise intrinsic to neural circuits. Slower, structural changes in random connectivity, consistent with rewiring and pruning processes, further improve the efficiency and sparseness of the resulting neural representations. Our results merge insights from neuroanatomy, machine learning, and theoretical neuroscience to suggest random sparse connectivity as a key design principle for neuronal computation.

doi: https://doi.org/10.1073/pnas.1912804117

4.4 The brain in motion: How ensemble fluidity drives memory-updating and flexibility.

Mau, W., Hasselmo, M. E., Cai, D. J. (2020). Elife, 9, e63550.

Abstract: While memories are often thought of as flashbacks to a previous experience, they do not simply conserve veridical representations of the past but must continually integrate new information to ensure survival in dynamic environments. Therefore, 'drift' in neural firing patterns, typically construed as disruptive 'instability' or an undesirable consequence of noise, may actually be useful for updating memories. In our view, continual modifications in memory representations reconcile classical theories of stable memory traces with neural drift. Here we review how memory representations are updated through dynamic recruitment of neuronal ensembles on the basis of excitability and functional connectivity at the time of learning. Overall, we emphasize the importance of considering memories not as static entities, but instead as flexible network states that reactivate and evolve across time and experience.

doi: https://doi.org/10.7554/eLife.63550

4.5 Representational drift in the mouse visual cortex.

Deitch, D., Rubin, A., Ziv, Y. (2020).bioRxiv.

Abstract: Neuronal representations in the hippocampus and related structures gradually change over time despite no changes in the environment or behavior. The extent to which such 'representational drift' occurs in sensory cortical areas and whether the hierarchy of information flow across areas affects neural-code stability have remained elusive. Here, we address these questions by analyzing large-scale optical and electrophysiological recordings from six visual cortical areas in behaving mice that were repeatedly presented with the same natural movies. We found representational drift over timescales spanning minutes to days across multiple visual areas. The drift was driven mostly by changes in individual cells' activity rates, while their tuning changed to a lesser extent. Despite these changes,

the structure of relationships between the population activity patterns remained stable and stereotypic, allowing robust maintenance of information over time. Such population-level organization may underlie stable visual perception in the face of continuous changes in neuronal responses.

doi: `https://doi.org/10.1101/2020.10.05.327049`

4.6 Predictive coding is a consequence of energy efficiency in recurrent neural networks.

Ali, A., Ahmad, N., de Groot, E., van Gerven, M. A., Kietzmann, T. C. (2021). bioRxiv.

Abstract: Predictive coding represents a promising framework for understanding brain function. It postulates that the brain continuously inhibits predictable sensory input, ensuring a preferential processing of surprising elements. A central aspect of this view is its hierarchical connectivity, involving recurrent message passing between excitatory bottom-up signals and inhibitory top-down feedback. Here we use computational modelling to demonstrate that such architectural hard-wiring is not necessary. Rather, predictive coding is shown to emerge as a consequence of energy efficiency. When training recurrent neural networks to minimise their energy consumption while operating in predictive environments, the networks self-organise into prediction and error units with appropriate inhibitory and excitatory interconnections, and learn to inhibit predictable sensory input. Moving beyond the view of purely top-down driven predictions, we demonstrate via virtual lesioning experiments that networks perform predictions on two timescales: fast lateral predictions among sensory units, and slower prediction cycles that integrate evidence over time.

doi: `https://doi.org/10.1101/2021.02.16.430904`

4.7 Barlow Twins: Self-Supervised Learning via Redundancy Reduction.

Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S. (2021). arXiv preprint arXiv:2103.03230.

Abstract: Self-supervised learning (SSL) is rapidly closing the gap with supervised methods on large computer vision benchmarks. A successful approach to SSL is to learn representations which are invariant to distortions of the input sample. However, a recurring issue with this approach is the existence of trivial constant representations. Most current methods avoid such collapsed solutions by careful implementation details. We propose an objective function that naturally avoids such collapse by measuring the cross-correlation matrix between the outputs of two identical networks fed with distorted versions of a sample, and making it as close to the identity matrix as possible. This causes the representation vectors of distorted versions of a sample to be similar, while minimizing the redundancy between the components of these vectors. The method is called Barlow Twins, owing to neuroscientist H. Barlow's redundancy-reduction principle applied to a pair of identical networks. Barlow Twins does not require large batches nor asymmetry between the network twins such as a predictor network, gradient stopping, or a moving average on the weight updates. It allows the use of very high-dimensional output vectors. Barlow Twins outperforms previous methods on ImageNet for semi-supervised classification in the low-data regime, and is on par with current state of the art for ImageNet classification with a linear classifier head, and for transfer tasks of classification and object detection.

doi: `https://arxiv.org/abs/2103.03230`

5 Machine learning / AI

5.1 The Next Big Thing (s) in Unsupervised Machine Learning: Five Lessons from Infant Learning.

Zaadnoordijk, L., Besold, T. R., Cusack, R. (2020). arXiv preprint arXiv:2009.08497.

Abstract: After a surge in popularity of supervised Deep Learning, the desire to reduce the dependence on curated, labelled data sets and to leverage the vast quantities of unlabelled data available recently triggered renewed interest in unsupervised learning algorithms. Despite a significantly improved performance due to approaches such as the identification of disentangled latent representations, contrastive learning, and clustering optimisations, the performance of unsupervised machine learning still falls short of its hypothesised potential. Machine learning has previously taken inspiration from neuroscience and cognitive science with great success. However, this has mostly been based on adult learners with access to labels and a vast amount of prior knowledge. In order to push unsupervised machine learning forward, we argue that developmental science of infant cognition might hold the key to unlocking the next generation of unsupervised learning approaches. Conceptually, human infant learning is the closest biological parallel to artificial unsupervised learning, as infants too must learn useful representations from unlabelled data. In contrast to machine learning, these new representations are learned rapidly and from relatively few examples. Moreover, infants learn robust representations that can be used flexibly and efficiently in a number of different tasks and contexts. We identify five crucial factors enabling infants' quality and speed of learning, assess the extent to which these have already been exploited in machine learning, and propose how further adoption of these factors can give rise to previously unseen performance levels in unsupervised learning.

doi: `https://arxiv.org/abs/2009.08497`

5.2 A neural network walks into a lab: towards using deep nets as models for human behavior.

Ma, W. J., Peters, B. (2020). arXiv preprint arXiv:2005.02181.

Abstract: What might sound like the beginning of a joke has become an attractive prospect for many cognitive scientists: the use of deep neural network models (DNNs) as models of human behavior in perceptual and cognitive tasks. Although DNNs have taken over machine learning, attempts to use them as models of human behavior are still in the early stages. Can they become a versatile model class in the cognitive scientist's toolbox? We first argue why DNNs have the potential to be interesting models of human behavior. We then discuss how that potential can be more fully realized. On the one hand, we argue that the cycle of training, testing, and revising DNNs needs to be revisited through the lens of the cognitive scientist's goals. Specifically, we argue that methods for assessing the goodness of fit between DNN models and human behavior have to date been impoverished. On the other hand, cognitive science might have to start using more complex tasks (including richer stimulus spaces), but doing so might be beneficial for DNN-independent reasons as well. Finally, we highlight avenues where traditional cognitive process models and DNNs may show productive synergy.

doi: `https://arxiv.org/abs/2005.02181`

5.3 An ecologically motivated image dataset for deep learning yields better models of human vision.

Mehrer, J., Spoerer, C. J., Jones, E. C., Kriegeskorte, N., Kietzmann, T. C. (2021). Proceedings of the National Academy of Sciences, 118(8).

Abstract: Inspired by core principles of information processing in the brain, deep neural networks (DNNs) have demonstrated remarkable success in computer vision applications. At the same time, networks trained on the task of object classification exhibit similarities to representations found in the primate visual system. This result is surprising because the datasets commonly used for training are designed to be engineering challenges. Here, we use linguistic corpus statistics and human concreteness ratings as guiding principles to design a resource that more closely mirrors categories that are relevant to humans. The result is ecoset, a collection of 1.5 million images from 565 basic-level categories. We show that ecoset-trained DNNs yield better models of human higher-level visual cortex and human behavior.

doi: `https://doi.org/10.1073/pnas.2011417118`

5.4 Individual differences among deep neural network models.

Mehrer, J., Spoerer, C. J., Kriegeskorte, N., Kietzmann, T. C. (2020). Nature communications, 11(1), 1-12.

Abstract: Deep neural networks (DNNs) excel at visual recognition tasks and are increasingly used as a modeling framework for neural computations in the primate brain. Just like individual brains, each DNN has a unique connectivity and representational profile. Here, we investigate individual differences among DNN instances that arise from varying only the random initialization of the network weights. Using tools typically employed in systems neuroscience, we show that this minimal change in initial conditions prior to training leads to substantial differences in intermediate and higher-level network representations despite similar network-level classification performance. We locate the origins of the effects in an under-constrained alignment of category exemplars, rather than misaligned category centroids. These results call into question the common practice of using single networks to derive insights into neural information processing and rather suggest that computational neuroscientists working with DNNs may need to base their inferences on groups of multiple network instances.

doi: `https://doi.org/10.1038/s41467-020-19632-w`

6 Plasticity and Learning

6.1 Homeostatic mechanisms regulate distinct aspects of cortical circuit dynamics.

Wu, Y. K., Hengen, K. B., Turrigiano, G. G., Gjorgjieva, J. (2020). Proceedings of the National Academy of Sciences, 117(39), 24514-24525.

Abstract: Homeostasis is indispensable to counteract the destabilizing effects of Hebbian plasticity. Although it is commonly assumed that homeostasis modulates synaptic strength, membrane excitability, and firing rates, its role at the neural circuit and network level is unknown. Here, we identify changes in higher-order network properties of freely behaving rodents during prolonged visual deprivation. Strikingly, our data reveal that functional pairwise correlations and their structure are subject to homeostatic regulation. Using a computational model, we demonstrate that the interplay of different plasticity and homeostatic mechanisms can capture the initial drop and delayed recovery of firing rates and correlations observed experimentally. Moreover, our model indicates that synaptic scaling is crucial for the recovery of correlations and network structure, while intrinsic plasticity is essential for the rebound of firing rates, suggesting that synaptic scaling and intrinsic plasticity can serve distinct functions in homeostatically regulating network dynamics.

doi: `https://doi.org/10.1073/pnas.1918368117`

6.2 Bidirectional synaptic plasticity rapidly modifies hippocampal representations independent of correlated activity.

Milstein, A. D., Li, Y., Bittner, K. C., Grienberger, C., Soltesz, I., Magee, J. C., Romani, S. (2020). BioRxiv.

Abstract: According to standard models of synaptic plasticity, correlated activity between connected neurons drives changes in synaptic strengths to store associative memories. Here we tested this hypothesis in vivo by manipulating the activity of hippocampal place cells and measuring the resulting changes in spatial selectivity. We found that the spatial tuning of place cells was rapidly reshaped via bidirectional synaptic plasticity. To account for the magnitude and direction of plasticity, we evaluated two models – a standard model that depended on synchronous pre- and post-synaptic activity, and an alternative model that depended instead on whether active synaptic inputs had previously been potentiated. While both models accounted equally well for the data, they predicted opposite outcomes of a perturbation experiment, which ruled out the standard correlation-dependent model. Finally, network modeling suggested that this form of bidirectional synaptic plasticity enables population activity, rather than pairwise neuronal correlations, to drive plasticity in response to changes in the environment.

doi: https://doi.org/10.1101/2020.02.04.934182

6.3 Optimal synaptic dynamics for memory maintenance in the presence of noise.

Raman, D. V., O'leary, T. (2020). BioRxiv.

Abstract: Synaptic connections in many brain areas have been found to fluctuate significantly, with substantial turnover and remodelling occurring over hours to days. Remarkably, this flux in connectivity persists in the absence of overt learning or behavioural change. What proportion of these ongoing fluctuations can be attributed to systematic plasticity processes that maintain memories and neural circuit function? We show under general conditions that the optimal magnitude of systematic plasticity is typically less than the magnitude of perturbations due to internal biological noise. Thus, for any given amount of unavoidable noise, 50% or more of total synaptic turnover should be effectively random for optimal memory maintenance. Our analysis does not depend on specific neural circuit architectures or plasticity mechanisms and predicts previously unexplained experimental measurements of the activity-dependent component of ongoing plasticity.

doi: https://doi.org/10.1101/2020.08.19.257220

6.4 Cortical reactivations of recent sensory experiences predict bidirectional network changes during learning.

Sugden, A. U., Zaremba, J. D., Sugden, L. A., McGuire, K. L., Lutas, A., Ramesh, R. N., ... & Andermann, M. L. (2020). Nature Neuroscience, 1-11.

Abstract: Salient experiences are often relived in the mind. Human neuroimaging studies suggest that such experiences drive activity patterns in visual association cortex that are subsequently reactivated during quiet waking. Nevertheless, the circuit-level consequences of such reactivations remain unclear. Here, we imaged hundreds of neurons in visual association cortex across days as mice learned a visual discrimination task. Distinct patterns of neurons were activated by different visual cues. These same patterns were subsequently reactivated during quiet waking in darkness, with higher reactivation rates during early learning and for food-predicting versus neutral cues. Reactivations involving ensembles of neurons encoding both the food cue and the reward predicted strengthening of next-day functional connectivity of participating neurons, while the converse was observed for reactivations involving ensembles encoding only the food cue. We propose that task-relevant neurons strengthen while task-irrelevant neurons weaken their dialog with the network via participation in distinct flavors of reactivation.

doi: https://doi.org/10.1038/s41593-020-0651-5

6.5 Synaptic plasticity as Bayesian inference.

Aitchison, L., Jegminat, J., Menendez, J. A., Pfister, J. P., Pouget, A., Latham, P. E. (2021). Nature Neuroscience, 1-7.

Abstract: Learning, especially rapid learning, is critical for survival. However, learning is hard; a large number of synaptic weights must be set based on noisy, often ambiguous, sensory information. In such a high-noise regime, keeping track of probability distributions over weights is the optimal strategy. Here we hypothesize that synapses take that strategy; in essence, when they estimate weights, they include error bars. They then use that uncertainty to adjust their learning rates, with more uncertain weights having higher learning rates. We also make a second, independent, hypothesis: synapses

communicate their uncertainty by linking it to variability in postsynaptic potential size, with more uncertainty leading to more variability. These two hypotheses cast synaptic plasticity as a problem of Bayesian inference, and thus provide a normative view of learning. They generalize known learning rules, offer an explanation for the large variability in the size of postsynaptic potentials and make falsifiable experimental predictions.

doi: https://doi.org/10.1038/s41593-021-00809-5.