

Analyzing biological and artificial neural networks: challenges with opportunities for synergy?

David GT Barrett^{1,3}, Ari S Morcos^{1,3,4} and Jakob H Macke²



Deep neural networks (DNNs) transform stimuli across multiple processing stages to produce representations that can be used to solve complex tasks, such as object recognition in images. However, a full understanding of how they achieve this remains elusive. The complexity of biological neural networks substantially exceeds the complexity of DNNs, making it even more challenging to understand the representations they learn. Thus, both machine learning and computational neuroscience are faced with a shared challenge: how can we analyze their representations in order to understand how they solve complex tasks? We review how data-analysis concepts and techniques developed by computational neuroscientists can be useful for analyzing representations in DNNs, and in turn, how recently developed techniques for analysis of DNNs can be useful for understanding representations in biological neural networks. We explore opportunities for synergy between the two fields, such as the use of DNNs as *in silico* model systems for neuroscience, and how this synergy can lead to new hypotheses about the operating principles of biological neural networks.

Addresses

¹ DeepMind, London, UK

² Computational Neuroengineering, Department of Electrical and Computer Engineering, Technical University of Munich, Germany

³ Equal contribution, alphabetical order.

⁴ Current address: Facebook AI Research (FAIR), Menlo Park, CA, USA.

Current Opinion in Neurobiology 2019, 55:55–64

This review comes from a themed issue on **Machine learning, big data, and neuroscience**

Edited by **Maneesh Sahani** and **Jonathan Pillow**

<https://doi.org/10.1016/j.conb.2019.01.007>

0959-4388/© 2019 Published by Elsevier Ltd.

Introduction

Neuroscience is in the midst of a technological transformation, enabling us to investigate the structure and function of neural circuits at unprecedented scale and resolution. Electrophysiological technologies [1] and imaging tools [2] have made it possible to record the activity of hundreds of neurons simultaneously, and optogenetic techniques enable targeted perturbations of neural activity [3,4]. These advances hold the promise of

providing fundamental insights into how populations of neurons collectively perform computations. However, it has also become increasingly clear that interpreting the complex data generated by these modern experimental techniques, and distilling a deeper understanding of neural computation is a challenging problem which requires powerful analysis tools [5].

In parallel, the field of machine learning is undergoing a transformation, driven by advances in deep learning. This has led to a large increase in the performance and widespread use of DNNs across numerous diverse problem domains such as object recognition [6,7], automated language translation [8], game-play [9,10] and scientific applications [11]. Deep networks consist of large numbers of linearly connected nonlinear units whose parameters are tuned using numerical optimization. Neuroscience and cognitive science were influential in the early development of DNNs [12] and convolutional neural networks (CNNs), widely used in computer vision [13,14,6], were inspired by canonical properties of the ventral visual stream.

Even though we have full access to DNNs which allows us to measure complete connectivity and complete activation patterns, it has nonetheless been challenging to develop a theoretical understanding of how and why they work. One reason that it is difficult to understand DNNs is that they usually contain millions of parameters. For example, ‘AlexNet’, which is well known for having demonstrated the potential of CNNs, contains 8 layers and a total of 60 million parameters [6]. Modern state of the art networks are often much larger. We still do not fully understand how and why DNNs can generalize so well without overfitting [15,16], nor do we fully understand how invariant representations arise in these multi-layer networks [17,18].

Therefore, both neuroscience and deep learning face a similar challenge: how do neural networks, consisting of large numbers of interconnected elements, transform representations of stimuli across multiple processing stages so as to implement a wide range of complex computations and behaviours, such as object recognition? What data-analysis techniques are most useful in this endeavor? How can we characterize and analyze representations in high-dimensional spaces? These common challenges open up opportunities for synergy in analysis across neuroscience and machine learning [19]. In this review, we explore statistical tools from both of these

disciplines that provide insight into neural representation and computation by analyzing activity-measurements of both single neurons and neural populations.

Receptive fields: what stimulus-features do single neurons represent?

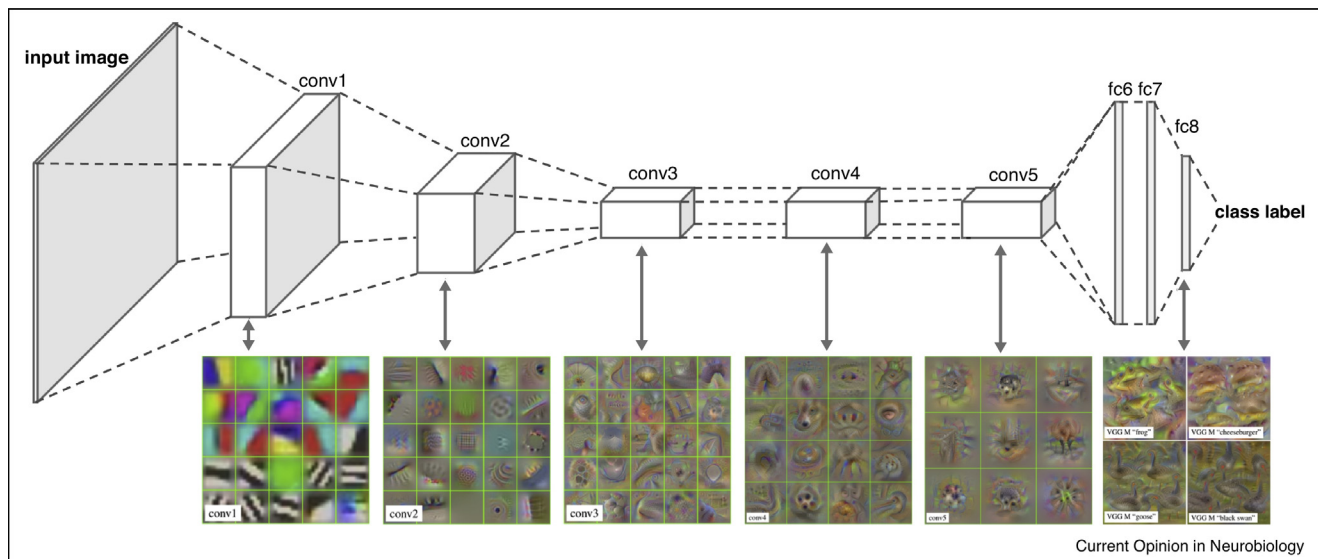
Receptive field analysis is the canonical method in neuroscience for analyzing single neuron activity [21,19]. The receptive field of a neuron usually refers to the region of stimulus space that causes a neural firing response. More generally, it is often characterized as the ‘preferred stimulus’ for a neuron — the stimulus that elicits a maximal response. Perhaps the most famous receptive field analysis was the Nobel prize winning work of Hubel and Wiesel, who discovered that the receptive fields of simple cells in the visual cortex have a spatially localized edge-like structure [22]. Our canonical understanding of visual processing has been largely informed by receptive field analysis. According to this perspective, the receptive fields of neurons along the ventral stream of the visual processing pathway become progressively larger [23], increasingly complex [24,25] and increasingly invariant to changes in the input-statistics [26], culminating in concept cells that are tuned to individual objects, but largely invariant to their visual appearance [27].

In recent years, receptive field analysis has also become a canonical method for analyzing response properties in artificial neural networks [28–31,20**]. For example, a receptive field analysis method known as activation maximization [20**,32,33] can be used to synthesize images which maximally activate units in artificial neural networks (Figure 1) [20**]. Similar to the ventral pathway, the receptive fields of units in CNNs become progressively larger [31] and increasingly complex. Earlier units having edge-like receptive fields and later units respond to more complex textures [28,20**]. The units of the final layer have class-specific receptive fields (Figure 1) which correspond to specific object categories, akin to concept-cells [27,34,20**].

Ablations: how relevant are single neurons to the overall computation?

Historically, neuroscience has focused on single neurons with clearly defined tuning properties, in part, due to the technical constraints involved in recording neural activity. However, a number of recent studies have demonstrated that neurons with clearly defined tuning, such as visual cortex simple cells, are relatively uncommon and that neurons with ambiguous tuning properties often contain substantial amounts of task-relevant information, especially when analyzed as populations [36–41]. This leads to the question: how important are selective, clearly tuned neurons?

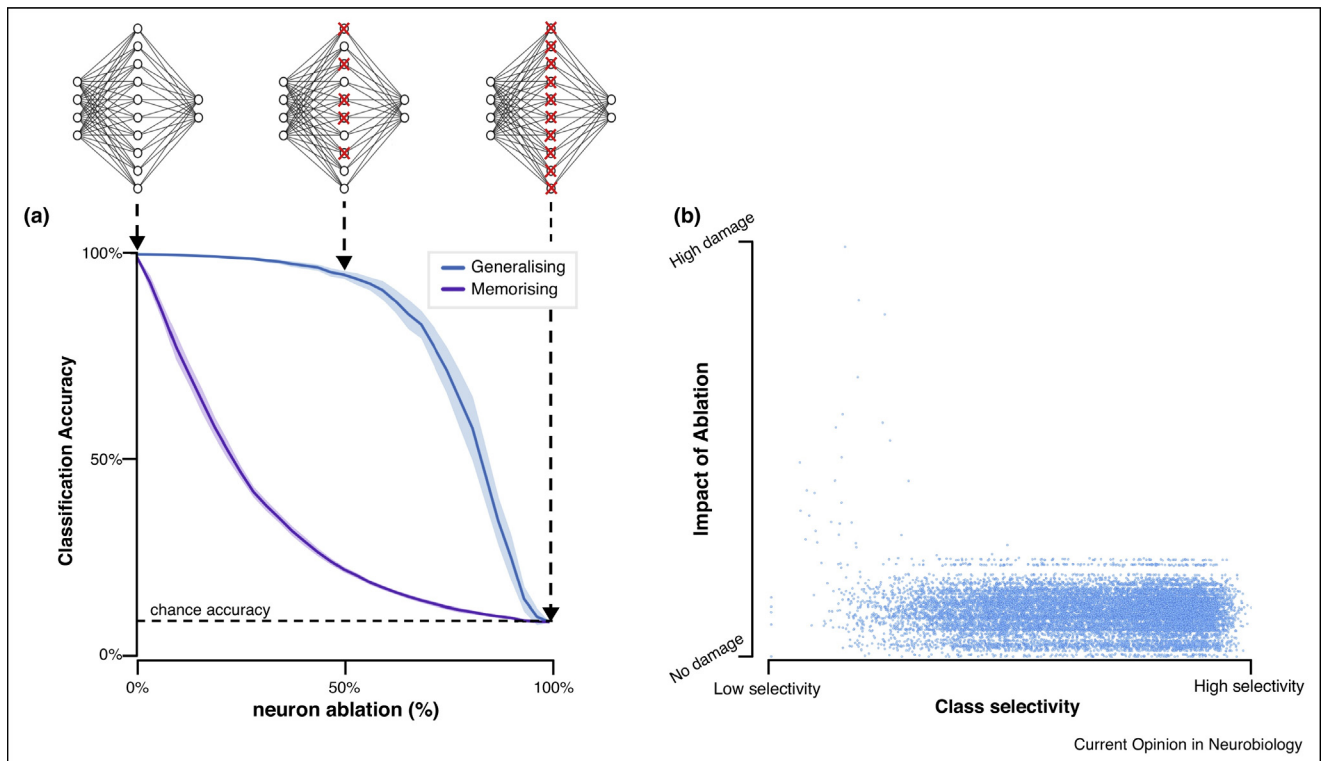
Figure 1



Receptive field analysis of neural filters from a DNN known as VGG-M. A schematic of the VGG-M network (top) illustrates how an input image is progressively processed through banks of convolutional (conv) filters, followed by a series of fully connected (fc) layers before finally producing a class label as output. Receptive fields (bottom) for each layer are calculated using activation maximisation [20**]. The receptive fields become increasingly complex, with earlier units having edge-like structure, intermediate layers responding to complex textures and the final layer responding to object-like images. It is not always possible to find a semantically meaningful description of a receptive field, as can be seen for the intermediate layers.

Adapted by permission from Springer Nature: Springer International Journal of Computer Vision, Visualizing Deep Convolutional Neural Networks Using Natural Pre-images, Aravindh Mahendran and Andrea Vedaldi, 2016.

Figure 2



Impact of neural ablations on DNN accuracy. (a) By ablating neurons and measuring DNN classification accuracy, the impact of neural ablations on DNN performance can be evaluated. A network with good generalization ability (cyan) is much more robust to ablation than networks which overfit the data using memorisation strategies (purple). (b) The relationship between the class selectivity of a neuron and the importance of a neuron for classification can be evaluated by measuring the impact of neuron ablation on classification accuracy as a function of the neuron class selectivity index. Most neurons have little impact on accuracy, and neuron class selectivity does not correlate strongly with the impact of neuron ablation.

Figures reproduced with modification from [35**].

Ablation analysis allow us to answer this question by silencing neurons and measuring the impact of this intervention on network output. However, it is difficult to perform single-neuron ablation analysis in biological neural networks, due to technical constraints [42], although this may change in the coming years due to the advent of optogenetics and targeted two-photon stimulation [3,4]. In theoretical neuroscience single-neuron ablation analysis is possible, and has allowed us to understand the impact of ablations on neural computation in model systems of both intact and damaged biological networks [43].

Ablation analysis is also possible in DNNs, as we can perfectly characterize the activity of every neuron in response to any arbitrary ablation. For example, ablation was recently used to evaluate the relationship between unit selectivity and unit importance as well as the relationship between network robustness and generalization in CNNs performing image classification tasks [35**]. By measuring how network accuracy dropped as increasing

numbers of neurons were deleted and training on both normal datasets (which allow generalizable solutions) and datasets in which image labels were randomized (which destroys structure and forces memorization), it was found that networks which learn generalizable solutions (i.e. those solutions which generalize to images never seen during training) were more robust to ablations than those which simply memorize the training data (and are therefore unable to generalize to unseen images; Figure 2a). Leveraging the ability to simultaneously characterize unit selectivity and unit importance, the impact of ablation on network accuracy and on the class selectivity of each unit was calculated. Perhaps surprisingly, there was little relationship between class selectivity and importance, suggesting that units critical to network computation were just as likely to exhibit ambiguous tuning as they were to be clearly tuned (Figure 2b). If the maximum drop of accuracy for *any* class is calculated rather than the drop in overall accuracy [44], a significant but relatively weak (~ -0.22) correlation between class selectivity and maximum class accuracy drop was observed — some

highly selective units were critical for individual classes, whereas many were unimportant.

These results raise an intriguing question: If selectivity is not predictive of unit importance, what is? This question has been extensively studied in the deep learning literature, primarily with the aim of ‘pruning’ neural networks to generate smaller networks which can run on resource limited devices [45,46]. One of the best signals of feature map importance has been found to be the summed absolute value of weights, suggesting that the magnitude of a feature map may be a deciding factor for importance [46]. Pruning remains an active area of research, and a better understanding of the factors which influence unit importance may prove highly insightful. This is particularly true for recurrent neural networks (RNNs) [47]: in RNNs, the contribution of a neuron can be nonlinearly amplified or reduced through recurrent interactions, which can make it more challenging to decipher which neurons ultimately contribute to a computation.

Dimensionality reduction: how can we characterize distributed representations?

If neural networks employ distributed representations [37,39,41], what analysis methods are available for us to make progress towards understanding these representations? A powerful approach which has been applied in neuroscience is to search for low-dimensional structure in large populations of neurons [5,48]. The use of low-dimensional codes embedded in high dimensional representations has been hypothesized to be an important organizing principle of neural computation in the brain [49,43,50], and so, dimensionality reduction methods are particularly suitable for characterizing neural activity.

Dimensionality reduction methods attempt to identify the neural signals that are maximally correlated with sensory stimuli [51], behavioural observations [52], neural activity in other brain areas [53] and instantaneous and temporal covariation within a brain area [54–56]. Multiple studies have found that neural dynamics and information about stimuli and behaviour are indeed concentrated in low-dimensional subspaces or manifolds [38,5,57,58,59] and are often distributed across many neurons [60,61]. For example, it has been found that dynamics in the motor cortex can be organized into low-dimensional manifolds within a high-dimensional distributed representation [62,61].

Do these manifolds only reflect correlations in the data, or do they impose constraints on what can be learned? To investigate this question, [57] used multi-electrode recordings to control a brain-machine interface. By changing the decoder for the brain-machine interface, they demonstrated that animals can learn to adapt to new decoders whenever the required activity patterns remained in the underlying manifold, but could not adapt

otherwise. This result suggests that low-dimensional neural manifolds represent functional structure and are not simply a byproduct of statistical correlations. This also has implications for future data-analysis: if the dynamics underlying neural data are low-dimensional and distributed across neurons, a small number of measurements is sufficient for decoding them [63,48].

Substantial redundancies in representations have also been reported for DNNs. [64] observed that most parameters in DNNs can be predicted from a small number (~5%) of parameters without significant reduction in network performance. [65] reported that the intrinsic dimensionality of neural networks can be orders of magnitude lower than the number of parameters. Moreover, ablation studies like the ones described in the previous section have found that more than 85% of parameters can be removed with minimal performance loss. In addition to linear dimensionality-reduction techniques, nonlinear dimensionality reduction techniques, such as t-SNE [66] have also been used extensively to analyze DNNs (e.g. [9]).

Relatively little is known about the relationship between DNN performance and the dimensionality of stimulus representations. Recent theoretical studies using approaches from statistical physics [67**] and computational neuroscience [68] are providing frameworks for studying the geometrical properties of object-related manifolds in neural networks: How high-dimensional are representations? Do they correspond to smooth or highly irregular manifolds? How do these geometrical properties impact the capacity and robustness of neural networks? These advances are providing opportunities for development and application of theory-driven data analysis techniques.

Cross-correlation: how can we compare representations across networks?

Just as we may need to go beyond single neuron analyses toward neural population analysis, we often need to go beyond single population analysis so that we can compare representations across multiple populations of neurons. Such analysis is necessary whenever we need to answer questions about the transformation of representations across layers, or across time. Also, looking towards opportunities for synergy between computational neuroscience and machine learning, it may be useful to compare representations in biological networks directly to representations in artificial networks.

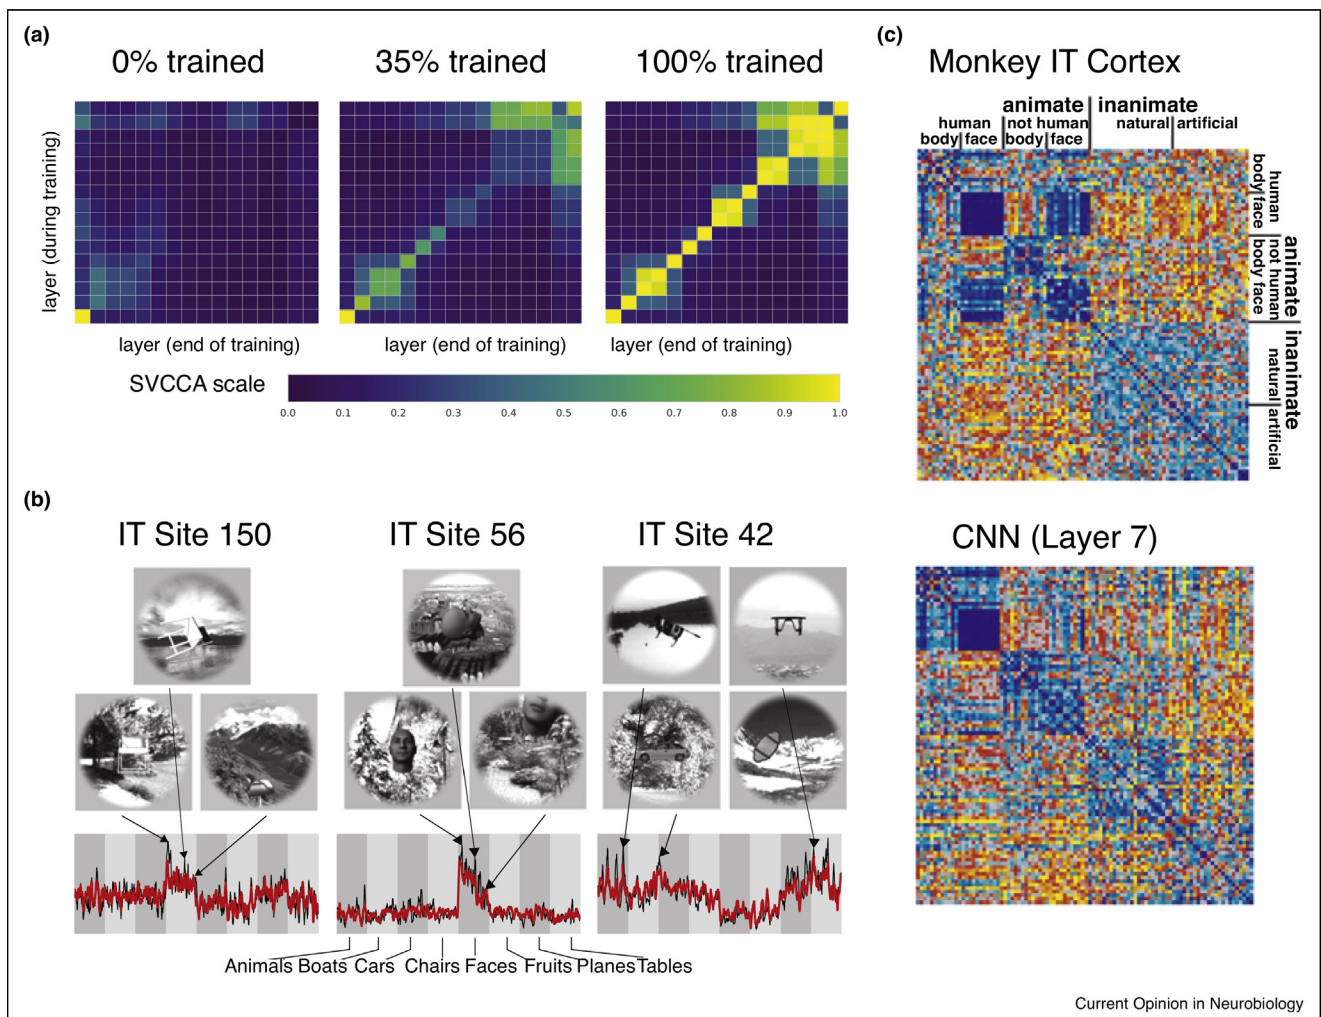
Canonical Correlation Analysis (CCA) is a suitable method for comparing representations across networks [72,53,73,69**,74]. CCA exposes directions which capture correlations between two data-sets (which may contain different numbers of neurons), just as Principal Component Analysis exposes directions in a data-set that

maximize variance. Raghu *et al.* [69**] developed an extension of CCA to compare activation-vectors of different neural layers. This analysis revealed low-dimensional structure in distributed representations, consistent with previous dimensionality reduction analysis. Going further, they used this approach to compare different ‘snapshots’ of a network during training and they found that the neural representations in the lower layers of DNNs stabilize much earlier than representations in higher layers (Figure 3a). Based on this observation, they proposed a new, computationally efficient training-approach in which lower layers are sequentially ‘frozen’ to focus computation-time on higher layers. Finally, they

used their CCA-based algorithm to compare two different network architectures trained on the same task (a ResNet and a ConvNet). They found that their representations were similar in early layers, but not in late layers. Li *et al.* [75*] used a related approach to investigate whether DNNs trained on the same task but with a different initial random seed can achieve similar representations. They reported that different networks have units which span overlapping subspaces, even when individual neurons differ.

Representational similarity analysis [76], which was originally proposed to compare representation between

Figure 3



Statistical methods for comparing representations across networks: **(a)** Learning dynamics in a DNN analyzed using Singular Vector Canonical Correlation Analysis (SVCCA): Each entry in the matrix describes how similar, during training, the representation in each layer is with each other layer after training. One can see that some layers are highly similar to their neighbouring layers, and that some layers converge quickly to their final values. **(b)** Representational similarity analysis (RSA) evaluates the (dis-)similarities that two networks assign to different pairs of inputs. The example shows a comparison between recordings in IT and a deep neural network. **(c)** Linear regressions can be used to align individual activation vectors between networks. The example shows 3 activation vectors measured in IT cortex (black), and reconstructions from those of a the top layer of a CNN (red). Panel a was reproduced with modification from [69**], b was reproduced with modification from [70], originally published under a CC-BY attribution license: <https://creativecommons.org/licenses/by/4.0/legalcode>, and c from [71].

different neural imaging modalities, can also be used to compare networks. In this approach, each network (or layer of a network) is characterized by a ‘representational similarity matrix’ (Figure 3b), which describes which pairs of stimuli a network considers to be similar or dissimilar. These pairwise similarities can be calculated from correlation matrices of the corresponding network activations.

Just as it is possible to compare representations in a DNN with another DNN, it is possible to compare DNN representations directly to biological neural representations. This is particularly interesting, as it opens up the possibility of directly using artificial neural networks as model systems for biological neural networks. This approach was recently adopted by Yamins *et al.* who used linear regression to identify correspondences between activation vectors of biological neural activity measurements and DNN activity measurements [71,23]. By predicting each ‘recorded’ activation vector from a linear combination of DNN-activation vectors from different layers of a DNN, they found that V4 activity could be best reconstructed from intermediate layers, and IT activity from top layers of the DNN (Figure 3c). This result is broadly consistent with receptive field analysis in artificial and biological neural networks which indicates that there are similarities between visual processing in a DNN and the ventral stream (Figure 1).

Challenges and opportunities

Neuroscientists and machine-learning researchers face common conceptual challenges in understanding computations in multi-layered networks. Consequently, there is an opportunity for synergy between the disciplines, to redeploy DNN analysis methods to understand biological networks and *visa-versa*. To what extent is synergy possible, and what challenges need to be overcome?

The first challenge is that DNNs allow full experimental access, whereas this is not possible for biological neural networks. Consequently, many DNN analysis methods can exploit information that is unavailable to neuroscientists. For instance, the activity of all units in a network in response to arbitrary stimuli can be simultaneously measured; the complete weight matrix (‘the connectome’) is known; and precise ablation and perturbation experiments can be performed. Moreover, the full behavioural history of the network (including every stimulus it has ever seen) is known, as is the learning process which determined the weights. Finally, it is usually possible to take gradients with respect to model parameters in DNNs and use these gradients for analysis.

The second challenge is that, despite both conceptual and empirical similarities [71,70,77*,25,23] between biological and artificial neural networks at the computational and algorithmic level (see [78*,79] for differences), there

are manifest differences at the mechanistic level. Neural networks used in deep learning are not biologically plausible because, for instance, they rely on the backpropagation algorithm (though see [80*,81]). Also, they do not produce spike-based representations. The constraints and demands faced by artificial and biological networks are also very different. For instance, brains need to be incredibly power efficient, whereas DNNs must be small enough to fit into computer memory. Whether and when DNNs and biological neural networks use similar representations and algorithms remains an open question. Consequently, analysis methods that may be informative for DNNs may not be appropriate for biological networks.

A third challenge is that we do not yet fully understand the solutions that DNNs learn, despite having full experimental access in DNNs, and despite them having simpler neural machinery. Since biological networks are substantially more complicated than DNNs, we should expect that it will be even more challenging to understand computation in the brain.

Notwithstanding these challenges, there are ample opportunities for synergy. First of all, several analysis methods for DNNs *can* be applied to biological systems without modification. For example, dimensionality reduction techniques such as CCA only require access to activity recordings. CCA could be used to study consistencies in activity patterns in the same neurons over time, across different layers, regions or animals, or to study the similarity of representations across subjects or brain areas. Moreover, a variety of dimensionality reduction algorithms which are specifically suited to biological neural network data have been developed, such as methods that are robust to missing data [82], methods that allow multiple recordings to be combined [56], and methods that are well matched to the statistics of spiking-noise [55,83] and to non-stationarities [58,59*].

Second, there is potential for overcoming some limitations that prevent the direct use of DNN analysis methods in neuroscience. For example, for some algorithms which depend on access to gradients, alternative ‘black-box’ variants are being developed which do not require such access and which might enable future application to biological systems, such as recent work on adversarial examples [84*]. It will be an important avenue for future work to adapt such methods to the statistical properties of neural activity measurements, and in particular to the fact that typically, we only have access to sparse, noisy and limited, non-stationary measurements in biological neural networks.

Third, DNNs can serve as idealized, *in-silico* model systems, which can allow researchers to rapidly develop, test and apply new data-analysis techniques on models that can solve complex sensory processing tasks [67**,77*].

Fourth, deep learning is providing new tools for the development of flexible, efficient and powerful neural analysis tools. For example, many Bayesian inference methods were too computationally expensive previously for large-scale data analysis. New approaches for training DNNs to perform Bayesian inference [85,86] are now opening up new avenues to develop efficient Bayesian inference methods for complex, mechanistic models in neuroscience (e.g. [83,87,88]).

Fifth, an important component of analysis and measurement in machine learning is the use of benchmark datasets to empirically compare algorithm performance. This approach may also be useful in neuroscience where it can often be difficult to determine, or even quantify, progress. Publicly accessible, large-scale, standardized data-sets are becoming available in neuroscience [89], which may enable the development of neuroscience benchmarks and challenges, for example, to predict the response-properties of neurons along the visual hierarchy, or for comparing representations between artificial networks and the brain [90]. These approaches might be useful in comparing, selecting, and ruling out competing models.

Finally, the fact that it has been difficult to understand DNNs — despite full experimental access and the use of simple neural units — serves as a reminder that better experimental tools for probing the activity and connectivity of neural circuits are necessary, but not sufficient. Rather, to understand computations in biological neural networks, we additionally require powerful methods for data-analysis, and ultimately we will require quantitative theories explaining these phenomena. Here again, another opportunity for synergy arises between the disciplines. Since DNNs have demonstrated that it is possible for neural systems to support a wide range of complex behaviours, the theoretical insights and understanding that has been developed for DNNs may be directly useful for informing new theoretical neuroscience. For instance, the observation that many of the features of DNNs can arise from very simple underlying principles such as numerical optimization of a loss-function suggests that many features of biological systems may also be understood from similar underlying optimization principles [91].

The co-incidental revolutions underway in neuroscience and machine learning have opened up a wide array of questions and challenges that have been long beyond our reach. ‘Out of adversity comes opportunity’ (B. Franklin), and where there are shared challenges there are opportunities for synergy. At their beginnings, the study of biological and artificial neural networks often confronted these challenges together, and although our disciplines have drifted and diverged in many ways, the time seems to be right, now, to return to this inter-disciplinary collaboration, in theory and in analysis.

Conflict of interest statement

The authors declare no conflict of interest.

Acknowledgements

We would like to thank Adam Santoro, Neil Rabinowitz, Lars Buesing and Thomas Reber for insightful discussions and comments on the manuscript, as well as Louise Deason and the DeepMind team for their support. JHM acknowledges support by ‘SFB 1233 Robust Vision’ of the German Research Foundation (DFG), the German Federal Ministry of Education and Research (BMBF, project ‘ADMIMEM’, FKZ 01IS18052 A-D), and the Human Frontier Science Program (RGY0076/2018), and thanks his colleagues at research center caesar and TU Munich for comments on the manuscript.

References and recommended reading

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

- of special interest
 - of outstanding interest
1. Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, Lee AK, Anastassiou CA, Andrei A, Aydn Ç *et al.*: **Fully integrated silicon probes for high-density recording of neural activity.** *Nature* 2017, **551**:232.
 2. Ahrens MB, Li JM, Orger MB, Robson DN, Schier AF, Engert F, Portugues R: **Brain-wide neuronal dynamics during motor adaptation in zebrafish.** *Nature* 2012, **485**:471.
 3. Packer AM, Russell LE, Dalgleish HWP, Häusser M: **Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo.** *Nat. Methods* 2014, **12**:140.
 4. Lerman GM, Gill JV, Rinberg D, Shoham S: **Spatially and temporally precise optical probing of neural activity readout.** *Optics and the Brain*. Optical Society of America; 2018. BTu2C-3.
 5. Cunningham JP, Yu BM: **Dimensionality reduction for large-scale neural recordings.** *Nat. Neurosci.* 2014, **17**:1500.
 6. Krizhevsky A, Sutskever I, Hinton GE: **Imagenet classification with deep convolutional neural networks.** *Advances in Neural Information Processing Systems*. 2012:1097-1105.
 7. Simonyan K, Zisserman A: **Very deep convolutional networks for large-scale image recognition.** *CoRR* 2014. abs/1409.1556.
 8. Wu Y, Schuster M, Chen Z, Le QV, Norouzi M, Macherey W, Krikun M, Cao Y, Gao Q, Macherey K *et al.*: **Google's neural machine translation system: bridging the gap between human and machine translation.** *arXiv preprint arXiv:1609.08144* 2016.
 9. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G *et al.*: **Human-level control through deep reinforcement learning.** *Nature* 2015, **518**:529.
 10. Silver D, Schrittwieser J, Simonyan K, Antonoglou I, Huang A, Guez A, Hubert T, Baker L, Lai M, Bolton A *et al.*: **Mastering the game of go without human knowledge.** *Nature* 2017, **550**:354.
 11. Januszewski M, Kornfeld J, Li PH, Pope A, Blakely T, Lindsey L, Maitin-Shepard J, Tyka M, Denk W, Jain V: **High-precision automated reconstruction of neurons with flood-filling networks.** *Nat. Methods* 2018, **15**:605-610.
 12. Rumelhart DE, McClelland JL: *Parallel Distributed Processing, vol. 1.* Cambridge, MA: MIT Press; 1987.
 13. Fukushima K, Miyake S: **Neocognitron: a self-organizing neural network model for a mechanism of visual pattern recognition.** *Competition and Cooperation in Neural Nets*. Springer; 1982:267-285.
 14. LeCun Y, Boser BE, Denker JS, Henderson D, Howard RE, Hubbard WE, Jackel LD: **Handwritten digit recognition with a back-propagation network.** *Advances in Neural Information Processing Systems*. 1990:396-404.

15. Zhang C, Bengio S, Hardt M, Recht B, Vinyals O: **Understanding deep learning requires rethinking generalization**. *International Conference of Learning Representations (ICLR) 2017*. This study demonstrated empirically that commonly used DNNs have sufficient capacity to memorize random data, calling into question commonly used explanations of the generalization ability of deep nets.
16. Dinh L, Pascanu R, Bengio S, Bengio Y: **Sharp minima can generalize for deep nets**. *ICML 2017*.
17. Tishby N, Zaslavsky N: **Deep learning and the information bottleneck principle**. *Information Theory Workshop (ITW), 2015 IEEE*. IEEE; 2015:1-5.
18. Achille A, Soatto S: **Emergence of Invariance and Disentangling in Deep Representations**. *Proceedings of the ICML Workshop on Principled Approaches to Deep Learning 2017*.
19. Victor JD: **Analyzing receptive fields, classification images and functional images: challenges with opportunities for synergy**. *Nat. Neurosci.* 2005, **8**:1651.
20. Mahendran A, Vedaldi A: **Visualizing deep convolutional neural networks using natural pre-images**. *Int. J. Comput. Vis.* 2016, **120**:233-255.
- The authors apply and compare different methods for visualizing neural networks, including activation-maximization and pre-images.
21. Sherrington CS: *The Integrative Action of the Nervous System*. New York: C Scribner's Sons; 1906.
22. Hubel DH, Wiesel TN: **Receptive fields of single neurones in the cat's striate cortex**. *J. Physiol.* 1959, **148**:574-591.
23. Yamins DLK, DiCarlo JJ: **Using goal-driven deep learning models to understand sensory cortex**. *Nat. Neurosci.* 2016, **19**:356.
24. Pasupathy A, Connor CE: **Shape representation in area v4: position-specific tuning for boundary conformation**. *J. Neurophysiol.* 2001, **86**:2505-2519.
25. Güçlü U, van MAJ, Gerven: **Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream**. *J. Neurosci.* 2015, **35**:10005-10014.
26. Rust NC, DiCarlo JJ: **Selectivity and tolerance ("invariance") both increase as visual information propagates from cortical area v4 to it**. *J. Neurosci.* 2010, **30**:12978-12995.
27. Reddy L, Kreiman G, Koch C, Fried I, Quiroga R: **Invariant visual representation by single neurons in the human brain**. *Nature* 2005, **435**.
28. Zeiler MD, Fergus R: **Visualizing and understanding convolutional networks**. *European Conference on Computer Vision*. Springer; 2014:818-833.
29. Yosinski J, Clune J, Nguyen A, Fuchs T, Lipson H: **Understanding neural networks through deep visualization**. *Deep Learning Workshop, International Conference on Machine Learning (ICML) 2015*.
30. Mahendran A, Vedaldi A: **Understanding deep image representations by inverting them**. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition 2015*:5188-5196.
31. Luo W, Li Y, Urtasun R, Zemel R: **Understanding the effective receptive field in deep convolutional neural networks**. *Advances in Neural Information Processing Systems 2016*:4898-4906.
32. Nguyen A, Dosovitskiy A, Yosinski J, Brox T, Clune J: **Synthesizing the preferred inputs for neurons in neural networks via deep generator networks**. *Advances in Neural Information Processing Systems 2016*:3387-3395.
33. Cadena SA, Weis MA, Gatys LA, Bethge M, Ecker AS: **Diverse feature visualizations reveal invariances in early layers of deep neural networks**. *Computer Vision – ECCV 2018*. Cham: Springer International Publishing; 2018, 225-240.
34. Le Q, Ranzato M, Monga R, Devin M, Chen K, Corrado G, Dean J, Ng A: **Building high-level features using large scale unsupervised learning**. *Proceedings of the 29th International Conference on Machine Learning 2012*.
35. Morcos AS, Barrett DGT, Rabinowitz NC, Botvinick M: **On the importance of single directions for generalization**. *International Conference of Learning Representations (ICLR) 2018*. An example of 'in-silico neurophysiology' on CNNs with an intriguing result: by performing ablation-experiments in CNNs, they report that strongly-tuned neurons are not more important for classification performance than ones with complex tuning. Also see[44].
36. Olshausen BA, Field DJ: In *23 Problems in Systems Neuroscience*. Edited by van Hemmen L, Sejnowski TJ. Oxford: Oxford University Press; 2006:182-214. (chapter 10).
37. Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S: **The importance of mixed selectivity in complex cognitive tasks**. *Nature* 2013, **497**:585.
38. Mante V, Sussillo D, Shenoy KV, Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex**. *Nature* 2013, **503**:78-84.
39. Raposo D, Kaufman MT, Churchland AK: **A category-free neural population supports evolving demands during decision-making**. *Nat. Neurosci.* 2014, **17**:1784.
40. Goris RLT, Simoncelli EP, Movshon J: **Origin and function of tuning diversity in macaque visual cortex**. *Neuron* 2015, **88**:819-831.
41. Morcos AS, Harvey CD: **History-dependent variability in population dynamics during evidence accumulation in cortex**. *Nat. Neurosci.* 2016, **19**:1672.
42. Theunissen FE, Miller JP: **Representation of sensory information in the cricket cercal sensory system. ii. information theoretic calculation of system accuracy and optimal tuning-curve widths of four primary interneurons**. *J. Neurophysiol.* 1991, **66**:1690-1703.
43. Barrett DGT, Denève S, Machens CK: **Optimal compensation for neuron loss**. *eLife* 2016, **5**:e12454.
44. Zhou B, Sun Y, Bau D, Torralba A: **Revisiting the importance of individual units in CNNs via ablation**. *arXiv preprint arXiv:1806.02891* 2018.
45. LeCun Y, Denker JS, Solla SA: **Optimal brain damage**. *Advances in Neural Information Processing Systems*. 1990:598-605.
46. Molchanov P, Tyree S, Karras T, Aila T, Kautz J: **Pruning convolutional neural networks for resource efficient inference**. *International Conference of Learning Representations (ICLR) 2017*.
47. Sussillo D, Barak O: **Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks**. *Neural Comput.* 2013, **25**:626-649.
48. Gao P, Trautmann E, Yu BM, Santhanam G, Ryu S, Shenoy K, Ganguli S: **A theory of multineuronal dimensionality, dynamics and measurement**. *bioRxiv* 2017.
49. Gao P, Ganguli S: **On simplicity and complexity in the brave new world of large-scale neuroscience**. *Curr. Opin. Neurobiol.* 2015, **32**:148-155.
50. Mastrogiuseppe F, Ostojic S: **Linking connectivity, dynamics, and computations in low-rank recurrent neural networks**. *Neuron* 2018. 10.1016/j.neuron.2018.07.003.
51. Archer EW, Koster U, Pillow JW, Macke JH: **Low-dimensional models of neural population activity in sensory cortical circuits**. *Advances in Neural Information Processing Systems 2014*:343-351.
52. Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Qi X-L, Romo R, Uchida N, Machens CK: **Demixed principal component analysis of neural population data**. *Elife* 2016, **5**:e10989.
53. Semedo J, Zandvakili A, Kohn A, Machens CK, Byron MY: **Extracting latent structure from multiple interacting neural populations**. *Advances in Neural Information Processing Systems 2014*:2942-2950.
54. Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M: **Gaussian-process factor analysis for low-**

- dimensional single-trial analysis of neural population activity.** *J. Neurophysiol.* 2009, **102**:614-635.
55. Macke JH, Büsing L, Cunningham JP, Yu BM, Shenoy KV, Sahani M: **Empirical models of spiking in neural populations.** *Adv. Neural Inf. Process. Syst.* 2012:1350-1358.
56. Nonnenmacher M, Turaga SC, Macke JH: **Extracting low-dimensional dynamics from multiple large-scale neural population recordings by learning to predict correlations.** *Advances in Neural Information Processing Systems 2017*:5706-5716.
57. Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Byron MY, Batista AP: **Neural constraints on learning.** *Nature* 2014, **512**:423.
58. Rabinowitz NC, Goris RL, Cohen M, Simoncelli EP: **Attention stabilizes the shared gain of v4 populations.** *Elife* 2015, **4** e08998.
59. Williams AH, Kim TH, Wang F, Vyas S, Ryu SI, Shenoy KV, Schnitzer M, Kolda TG, Ganguli S: **Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis.** *Neuron* 2018, **98**:1099-1115.
- This study presents a dimensionality-reduction algorithm based on tensor-decompositions, which allows one to separate dynamics on different time-scales.
60. Machens CK, Romo R, Brody CD: **Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex.** *J. Neurosci.* 2010, **30**:350-360.
61. Gallego JA, Perich MG, Miller LE, Solla SA: **Neural manifolds for the control of movement.** *Neuron* 2017, **94**:978-984.
62. Shenoy KV, Sahani M, Churchland MM: **Cortical control of arm movements: a dynamical systems perspective.** *Annu. Rev. Neurosci.* 2013, **36**:337-359.
63. Ganguli S, Sompolinsky H: **Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis.** *Annu. Rev. Neurosci.* 2012, **35**:485-508.
64. Denil M, Shakibi B, Dinh L, Ranzato M, De Freitas N: **Predicting parameters in deep learning.** *Advances in Neural Information Processing Systems 2013*:2148-2156.
65. Li C, Farkhoor H, Liu R, Yosinski J: **Measuring the intrinsic dimension of objective landscapes.** *International Conference of Learning Representations (ICLR)* 2018.
66. van der Maaten L, Hinton G: **Visualizing data using t-sne.** *J. Mach. Learn. Res.* 2008, **9**:2579-2605.
67. Chung SY, Lee DD, Sompolinsky H: **Classification and geometry of general perceptual manifolds.** *Phys. Rev. X* 2018, **8** 031003.
- This study introduces a new theoretical framework for characterizing the geometrical structure of ‘perceptual manifolds’, for example, the space of images corresponding to the same object category, and linking it to the capacity of a classifier. The framework can be applied to both neural activity measurements and deep neural networks.
68. Stringer CS, Pachitariu M, Steinmetz N, Carandini M, Harris KD: **High-dimensional geometry of population responses in visual cortex.** *bioRxiv* 2018.
69. Raghu M, Gilmer J, Yosinski J, Sohl-Dickstein J: **Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability.** *Advances in Neural Information Processing Systems 2017*:6078-6087.
- The authors present an efficient algorithm for identifying correspondences between the representations of two networks, and use it study the dimensionality and convergence of neural networks.
70. Khaligh-Razavi S-M, Kriegeskorte N: **Deep supervised, but not unsupervised, models may explain it cortical representation.** *PLoS Comput. Biol.* 2014, **10** e1003915.
71. Yamins DLK, Hong H, Cadieu CF, Solomon EA, Seibert D, DiCarlo JJ: **Performance-optimized hierarchical models predict neural responses in higher visual cortex.** *Proc. Natl. Acad. Sci. U. S. A.* 2014, **111**:8619-8624.
72. Hotelling H: **Relations between two sets of variates.** *Biometrika* 1936, **28**:321-377.
73. Sussillo D, Churchland MM, Kaufman MT, Shenoy KV: **A neural network that finds a naturalistic solution for the production of muscle activity.** *Nat. Neurosci.* 2015, **18**:1025.
74. Morcos AS, Raghu M, Bengio S: **Insights on representational similarity in neural networks with canonical correlation.** *Advances in Neural Information Processing Systems.* Curran Associates, Inc.; 2018:5732-5741.
75. Li Y, Yosinski J, Clune J, Lipson H, Hopcroft JE: **Convergent learning: do different neural networks learn the same representations?** *International Conference of Learning Representations (ICLR)* 2016.
- The authors introduce methods for studying whether different neural networks trained on the same task converge to similar or different representations.
76. 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.
77. Kriegeskorte N: **Deep neural networks: a new framework for modeling biological vision and brain information processing.** *Annu. Rev. Vis. Sci.* 2015, **1**:417-446.
- This review paper describes CNNs in the context of neuroscience, and summarizes known commonalities and differences between CNNs and the visual system.
78. Geirhos R, Temme CRM, Rauber J, Schütt HH, Bethge M, Wichmann FA: **Generalisation in humans and deep neural networks.** *Advances in Neural Information Processing Systems 2018*:7548-7560.
- By comparing the performance of human observers with those of different DNNs on images corrupted by noise, they find the human visual system to be much more robust to perturbations.
79. Pospisil D, Pasupathy A, Bair W: **Comparing the brains representation of shape to that of a deep convolutional neural network.** In *Proceedings of the 9th EAI International Conference on Bio-inspired Information and Communications Technologies (formerly BIONETICS); ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering): 2016*:516-523.
80. Lillicrap TP, Cownden D, Tweed DB, Akerman CJ: **Random synaptic feedback weights support error backpropagation for deep learning.** *Nat. Commun.* 2016, **7** 13276.
- This study aims to find biologically realistic approaches for optimizing the parameters of DNNs.
81. Guerguiev J, Lillicrap TP, Richards BA: **Towards deep learning with segregated dendrites.** *ELife* 2017, **6**:e22901.
82. Balzano L, Nowak R, Recht B: **Online identification and tracking of subspaces from highly incomplete information.** *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on; IEEE: 2010*:704-711.
83. Gao Y, Archer EW, Paninski L, Cunningham JP: **Linear dynamical neural population models through nonlinear embeddings.** *Advances in Neural Information Processing Systems 2016*:163-171.
84. Brendel W, Rauber J, Bethge M: **Decision-based adversarial attacks: reliable attacks against black-box machine learning models.** *International Conference on Learning Representations 2018.*
- While most algorithms for generating ‘adversarial examples’ require access to the inner workings of a classifier, this study introduces a black-box approach which only requires observations of inputs and outputs.
85. Kingma DP, Welling M: **Auto-encoding variational Bayes.** *2nd International Conference on Learning Representations (ICLR)* 2013.
86. Papamakarios G, Murray I: **Fast epsilon-free inference of simulation models with Bayesian conditional density estimation.** *Advances in Neural Information Processing Systems, vol. 29* 2017.
87. Lueckmann J-M, Goncalves PJ, Bassetto G, Öcal K, Nonnenmacher M, Macke JH: **Flexible statistical inference for mechanistic models of neural dynamics.** *Advances in Neural Information Processing Systems 2017*:1289-1299.

88. Speiser A, Yan J, Archer EW, Buesing L, Turaga SC, Macke JH: **Fast amortized inference of neural activity from calcium imaging data with variational autoencoders.** *Advances in Neural Information Processing Systems* 2017:4024-4034.
89. Hawrylycz M, Anastassiou C, Arkipov A, Berg J, Buice M, Cain N, Gouwens NW, Gratiy S, Iyer R, Lee JH *et al.*: **Inferring cortical function in the mouse visual system through large-scale systems neuroscience.** *Proc. Natl. Acad. Sci. U. S. A.* 2016, **113**:7337-7344.
90. Schrimpf M, Kubilius J, Hong H, Majaj NJ, Rajalingham R, Issa EB, Kar K, Bashivan P, Prescott-Roy J, Schmidt K *et al.*: **Brain-score: Which artificial neural network for object recognition is most brain-like?** *bioRxiv* 2018:407007.
91. Marblestone AH, Wayne G, Kording KP: **Toward an integration of deep learning and neuroscience.** *Front. Comput. Neurosci.* 2016, **10**:94.