Particularly interesting is the methodology used for averaging of multicolor data, which was handled in two different ways: (i) alignment of all of the proteins with respect to a reference protein, which was imaged in each multicolor sample (Fig. 1), and (ii) merging of the multicolor data into a single channel before alignment and classification, which may be necessary when one of the proteins of interest is asymmetrically distributed with respect to the reference. This aspect of the work is beautifully demonstrated in the reconstruction of the HsSAS-6 protein, which extends away from the Cep152 ring: the reconstructed image captures not only its asymmetric localization but also its average orientation with respect to the axis of the procentriole.

Software development is a central component of both works. The particle alignment and fusion program from Heydarian et al.[7] is called datafusion2d and is written for use with MATLAB, and the collection of algorithms used by Sieben et al.[8] has been packaged into a MATLAB application named SPARTAN. Notably, the authors provide both software packages in an open-source manner, available via publically accessible repositories, thus helping to ensure the transparency and reproducibility of the methods and facilitating input from future contributors.

The studies highlighted here represent early steps in the use of fluorescence imaging for structural biology at the nanoscale, but what can we expect in the future? Fluorescence nanoscopy continues to advance at a rapid pace[1], driven by new fluorophores, labeling methods, instrumentation, and imaging and data analysis concepts, all of which will lead to higher-quality image data for use in structure determination. New approaches, such as the direct alignment and averaging of 3D fluorescence images (rather than 2D), could simplify the averaging pipeline and yield reconstructions of greater complexity. These developments offer the potential for particle analysis that goes beyond static structures, extending to the detection of multiple conformational states among a population of particles, and the visualization of flexible components of protein complexes. Although the labels used for fluorescence imaging introduce a spatial offset between the measured signal and the structure of interest, smaller labels in the form of nanobodies, DNA aptamers[10], and genetically encoded tags allow this offset to be reduced to a few nanometers. As these studies illustrate, the combination of SPA with fluorescence imaging has great potential to complement existing structural biology methods, by providing high-resolution structures with protein specificity, thereby establishing a new tool for the study of biology at the nanoscale. ❐

Mark Bates
Department of Nanobiophotonics, Max Planck Institute for Biophysical Chemistry, Göttingen, Germany.
e-mail: mark.bates@mpibpc.mpg.de

Published online: 1 October 2018
https://doi.org/10.1038/s41592-018-0151-7

References

1. Sahl, S. J., Hell, S. W. & Jakobs, S. *Nat. Rev. Mol. Cell Biol.* **18**, 685–701 (2017).
2. Frank, J. *Nat. Protoc.* **12**, 209–212 (2017).
3. Bharat, T. A. M., Russo, C. J., Löwe, J., Passmore, L. A. & Scheres, S. H. W. *Structure* **23**, 1743–1753 (2015).
4. Mennella, V. et al. *Nat. Cell Biol.* **14**, 1159–1168 (2012).
5. Szymborska, A. et al. *Science* **341**, 655–658 (2013).
6. Broeken, J. et al. *Methods Appl. Fluoresc.* **3**, 014003 (2015).
7. Heydarian, H. et al. *Nat. Methods* https://doi.org/10.1038/s41592-018-0136-6 (2018).
8. Sieben, S., Banterle, N., Douglass, K. M., Gönczy, P. & Manley, S. *Nat. Methods* https://doi.org/10.1038/s41592-018-0140-x (2018).
9. Bates, M., Huang, B. & Zhuang, X. *Curr. Opin. Chem. Biol.* **12**, 505–514 (2008).
10. Strauss, S. et al. *Nat. Methods* **15**, 685–688 (2018).

Competing interests

The author declares no competing interests.

NEUROSCIENCE

Deep learning reaches the motor system

A new article by Pandarinath et al. describes an artificial neural network model that captures some key aspects of the activity of populations of neurons in the primary motor cortex.

Aaron P. Batista and James J. DiCarlo

When can we declare that we *understand* something? One pragmatically minded definition is that we understand a phenomenon when we can duplicate it. Thus, the airfoil design of an airplane wing reveals a better understanding of bird flight than did the feathered wings of mythology's Icarus. This intuition animates an important direction in systems neuroscience: the use of artificial neural networks to duplicate the input–output relationship of real networks of neurons, with the hope that such networks will inform on internal neurobiological mechanisms. Even with different physical implementations (in this case a computer versus cortical tissue), if we can mimic the functionality of a population of motor cortex neurons, then we can claim that we have begun to understand how those neurons work. This form of understanding by model building is a tale of successive approximations: over time we build models that capture more and more details of a system's function, and are increasingly constrained by the biological system's internal structure. This approach, which has been influential in efforts to obtain a mechanistic understanding of the visual system[1,2], has now begun to bear fruit as a tool for investigating the brain's motor system[3].

In their groundbreaking article, Pandarinath et al.[3] bring the modern computational technique of artificial neural networks to bear on the problem of understanding the function of the primary motor cortex. The role of the brain's motor system is to convert neural representations of desired goals into a pattern of muscle contractions that enable the organism to attain the goal (Fig. 1). There is a growing appreciation that the temporal coordination (that is, dynamics) of muscle activity is a key feature of the computations performed by the motor cortex[4]. As muscle coordination involves precise temporal sequencing, it stands to reason that the primary motor cortex expresses temporal dynamics[5]. Pandarinath et al. have designed a neural network architecture that can mimic the

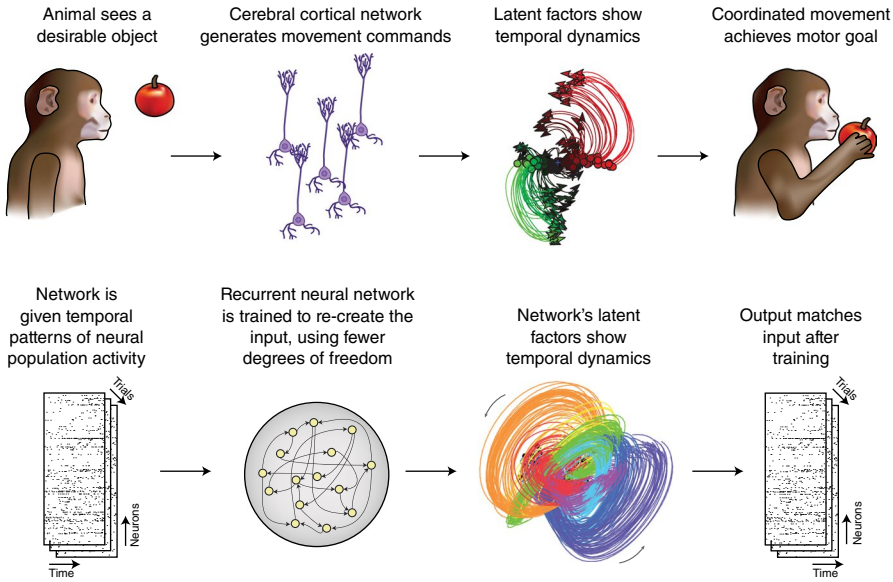

Fig. 1 | An artificial neural network (bottom) can capture the dynamical structure present in neural population activity (top). Credit: Kim Caesar/Springer Nature

temporal dynamics observed in the motor cortex[3]. But how do the primary motor and lateral premotor cortexes generate dynamics from a static input?

To gain a toehold on this mechanistic question, the authors implemented an autoencoder, a special type of artificial neural network. The premise of an autoencoder network is deceptively simple: it seeks to reproduce its input. Although this could be done trivially with one-to-one mapping, the cleverness of an autoencoder is that it passes the data through a 'bottleneck'—an intermediate stage that is of lower dimensionality than the input (and the output). This constraint forces the network's hidden layers to discover the deep structure present in the data. As evidence that the underlying structure has been identified, the autoencoder network can generate synthetic data that closely match real data.

Once the network has been trained (deep learning techniques are used—'deep' because the network is a cascade of multiple stages, and 'learning' because the network's parameters are adjusted through successive approximations), its hidden structure can be examined. Herein lies the main result of the study: the activity of the network's units demonstrates temporal dynamics that are reminiscent of those exhibited by populations of real neurons. That is to say, the same computational structure arises both from real neurons and from a collection of model neurons and connections that only coarsely approximate the actual physical substrate. It is easy to imagine that principles of neural computation have been captured, although it remains challenging to describe those principles in human-interpretable terms.

This deep structure discovered by the network reinforces an emerging understanding in systems neuroscience, which is that the recorded activity of neurons is a readout of an underlying structure of lower dimensionality[6]. Just as weather can provide information about climate, even though the two are separate concepts, from observations of neural activity one can infer the properties of the latent factors that are the essence of network computation. Think of the fluctuating pixel intensities on a television screen, which are different from the narrative action that drives those pixels. It is by recording from populations of neurons at once, and building computational models, that neuroscientists are able to identify underlying 'latent factors'. Evidence that the latent-factor view captures bedrock aspects of neural information processing is the fact that it is difficult for animals to learn how to drive their population neural activity in a manner that violates the latent structure present in the activity of a population of neurons[7]. One outcome of this latent-factors perspective is that scientists are beginning to be able to understand the neural basis of behavior with single-trial resolution: each action is a unique event, and much of interest is lost when we average over repeats of nominally identical (but actually different) actions (see, e.g., ref. [8]). The network built here can capture these idiosyncratic individual-trial aspects of neural activity.

Pandarinath et al. named their network with the clever double-acronym LFADS, in part for its discovery of latent factors via the construction of a dynamical system. Currently, those in the field should see LFADS as a tool—as the authors urge—rather than a literal model of the brain. A future application of LFADS could be to improve the performance of brain–computer interface systems by developing a richer, more flexible controller from the neural instantiation of motor intentions to the movement of external devices. Beyond this important potential application, it is impossible to resist the deep connections between the architecture of real neurons and the powerful computational tools provided by simulated networks of neurons. And, now that deep learning has begun to be applied to the motor system, the field might rapidly catch up to visual neuroscience, where artificial neural networks have advanced enough that they are actually providing mechanistic insight into the neural bases of perception, cognition and, perhaps soon enough now, action. ❐

Aaron P. Batista[1]* and James J. DiCarlo[2]
[1]*Department of Bioengineering, University of Pittsburgh, Pittsburgh, PA, USA.* [2] *Department of Brain and Cognitive Sciences and McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA, USA.*
**e-mail: aaron.batista@pitt.edu*

Published online: 1 October 2018
https://doi.org/10.1038/s41592-018-0152-6

References

1. Yamins, D. L. K. et al. *Proc. Natl. Acad. Sci. USA* **111**, 8619–8624 (2014).
2. Yamins, D. L, Hong, H., Cadieu, C. & DiCarlo, J. J. in *Advances in Neural Information Processing Systems 26* (eds Burges, C. J. C. et al.) 3093–3101 (NIPS Foundation, La Jolla, CA, 2013).
3. Pandarinath, C. et al. *Nat. Methods*. https://doi.org/10.1038/s41592-018-0109-9 (2018).
4. Churchland, M. M. et al. *Nature* **487**, 51–56 (2012).
5. Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. *Neuron* **68**, 387–400 (2010).
6. Cunningham, J. P. & Yu, B. M. *Nat. Neurosci.* **17**, 1500–1509 (2014).
7. Sadtler, P. T. et al. *Nature* **512**, 423–426 (2014).
8. Peixoto, D. et al. *bioRxiv* Preprint at https://www.biorxiv.org/content/early/2018/03/17/283960 (2018).

Competing interests

The authors declare no competing interests