CellPress REVIEWS

Scientific Life

Emerging Opportunities for Advancing Cognitive Neuroscience

Uri Hasson[1,*] and Howard C. Nusbaum[2]

Cognitive neuroscience can be substantially advanced if structured mechanisms are created to increase its social impact and to develop synergies with some currently more distant disciplines that are developing relevant knowledge. We present such opportunities and argue that pursuing these can benefit from establishing a centralized coordinating organizational approach.

As an emergent discipline between cognitive psychology and neuroscience, cognitive neuroscience (CN) aims to explain relationships among neural systems, psychological states or processes, and behavior. At its best, it does so with few linking hypotheses, relying on studies that propel the field forward by adding to existing knowledge or revising prior conceptions. There is little doubt that the field will advance if it continues in this manner.

However, based on our own experience as scientists in the field, as well as our experience with decision processes in funding agencies, we believe that CN research can be substantially strengthened (in terms of conceptual advances, social and scientific impact, breadth and pace) if efforts are made in four key areas (discussed below). We contend that progress in some of these areas requires establishing a transdisciplinary coordinating group, with a different mission from that of extant scientific societies, and that would function outside research funding bodies. Its focus would be on developing infrastructure to shape CN and its outward-looking interactions, rather than on communication of research advances.

Our views on these issues were shaped both by carrying out research in CN and by recent experience at the US National Science Foundation (NSF), where H.N. was Division Director for the Behavioral and Cognitive Sciences Division, and U. H. was Program Director for the Cognitive Neuroscience Program. Our duties brought us into contact with investigators of all career levels in CN and in disciplines with interests relevant to CN. Investigators and scientific organizations have diverse aims within the discipline, such as advancing existing research programs, establishing new collaborations (within or across disciplines), applying new paradigms, or developing new research communities (via new coordination networks, conferences, or workshops). Notably, however, what we advocate here did not naturally emerge from interactions with principal investigators or scientific societies, and we concluded that such initiatives are not optimally suited to *de novo* origination within a single funding agency in the absence of analytic consideration and guidance from the field.

Based on our experience, we believe that development in the following four directions will be particularly important for the advancement of CN, and that pursuing some of these directions will require a centralized coordinating structure.

Open CN to Perspectives and Practices from Scientific Disciplines Where Interactions Are Limited

Currently, several research communities are independently developing knowledge that can advance CN. Some of these communities have a primarily applied (problem-driven) interest in neurobiological data. This includes the engineering community, as well as communities engaged in brain-computer interface research, data science, and machine learning. Based on our observations of synergies that develop among research communities, we believe improved interactions with engineering, and specifically neuro-engineering, will catalyze new approaches for understanding neural mechanisms, as engineers often rely on (and develop) data-driven analyses that can yield novel insights or produce new hypotheses about the neural basis of behavior.

The development of methods for training compact convolutional neural networks that can be applied to relatively small datasets is an example of a relevant engineering solution [1]. Additionally, advances in automatic speech reconstruction/decoding from electrocorticography data show that sophisticated decoding approaches can produce new hypotheses about the brain mechanisms that underlie speech comprehension and production [2]. IEEE's Society for Systems, Man, and Cybernetics exemplifies a community where cutting-edge engineering advances often rely on analysis of neurobiological data, but whose findings are having limited impact on theory development in CN. Examples of other areas are presented in Box 1.

Develop Processes to Strengthen CN's Impact on Socially Relevant Areas and Industry

Insights from CN can have near-term impact on key areas, including cybersecurity (e.g., authentication systems using biometric signatures), clinical practice (e.g., presurgical planning), and areas that are driving the economy, such as machine learning, personalized medicine and Artificial Intelligence. CN investigation methods can be used to advance 'explainable AI', for example by applying research tools akin to those used in CN investigations to understand the organization of artificial networks.

Box 1. Examples of Research Areas that Could Be Better Integrated with Mainstream CN

(i) There are currently no mechanisms for integrating leading-edge pattern classification paradigms into mainstream CN research. The annual International Workshop on Pattern Recognition in Neuroimaging (PRNI) provides an opportunity for interactions, but there are only a few others. One possibility would be to implement machine-learning competitions (as supported by Kaggle) to allow CN researchers to obtain new answers to theoretically motivated questions by capitalizing on advances and know-how in data science.

(ii) Mathematical psychology has been relatively less integrated into CN than other approaches. However, researchers have begun using neurobiological data within hybrid neuro/cognitive models that are designed to inform both functional and neurobiological explanations [7–9]. This framework can directly inform CN but has yet to be integrated into mainstream studies and analysis streams.

(iii) Research in the neurobiology of language has been slow to capitalize on advances in models developed within the Natural Language Processing and Computational Linguistics communities. These approaches are quite distinct from Neurolingistics and Psycholinguistics. Presently, CN has relatively modest contact with Computational Linguistics, with such efforts driven mainly by the Linguistics community (seen, e.g., in the recent inaugural Workshop on Linguistic and Neuro-Cognitive Resources; LiNCR). Recent publications demonstrate the effectiveness of this integration (e.g., [10,11]).

This includes not only application of data analysis methods developed for studying spatiotemporal brain dynamics, but also application of experimental methodologies, such as evaluating the impact of localized inhibition or excitation, the impact of stimulus degradation, and the effects of memory and attention. This can potentially provide new insights into the internal organization of convolutional neural networks, which can yield further developments in deep learning methods. While there is considerable promise, with economic implications, the impact of mainstream CN research on AI is currently limited, with much of the crosstalk occurring in the space of computer-vision models and reinforcement learning. Consequently, the development of computational models such as deep-learning-based classifiers is impacted by neurobiological levels of analyses that reflect a limited subset of the intellectual approaches employed in CN. And conversely, advances in AI are informing CN in a limited way as those approaches do not share CN's historical focus on gaining insights into specific representational mechanisms (see [3] for critical evaluation, but [4] for a more optimistic assessment).

The machine/statistical learning and engineering communities also stand to gain, because analyzing human brain data presents unique challenges that can (and have already been shown to) catalyze development of new methods in those areas. This includes managing systematic variance across individuals, accounting for long-range connectivity when applying models of dynamical systems (e.g., to neural avalanches), and developing neural networks that learn from sparse data. Providing an organized interface to CN data and methods will also aid in workforce development as Data Science programs do not have an organic way for training their students (often Master's level) to manage neurobiological and physiological datasets, which are skills that will be necessary in the future workforce. Currently, many of these potential connections are very weak, nonexistent, or driven by unnecessarily random opportunities.

Transform the Use of Cyber-infrastructure

The past 20 years have witnessed multiple attempts to use cyber-infrastructure (CI) to facilitate faster analyses and hypothesis testing, as well as reproducible research in CN (examples include Brainlife [12], OpenNeuro [13], and ReproNim [14]). However, the community should begin to consider the potential of CI to enable tackling new, yet-unformulated questions that simply cannot be currently addressed. A recent example showing the utility of such an approach is using massively parallel computing for online decoding of fMRI correlation matrices [5]. Future advances could address other challenges currently out of reach, for instance, real-time source localization of activity from electroencephalography or magnetoencephalography systems for purposes of real-time neurofeedback, or millisecond-level comparisons of brain measurements versus model-predicted activity patterns for purposes of training artificial networks. Some of these efforts might require coordinated construction and time-sharing of large-scale specialized infrastructure that will serve the field as a whole, which is a model that has been successful in other scientific areas. One worry is that without coordination, multiple incompatible approaches to the use of CI will flourish, whereas with coordination, standards can be developed and supported.

Guarantee a Community Standard that Ensures that Much of High-quality CN Research is Openly and Meaningfully Accessible

A separate effort, related to those above but distinct in its mode of development, is the need to assure that CN research products and methods are available without barriers to the clinical, engineering, technological, and other communities mentioned above. Currently there are no open access (OA) journals dedicated to covering the entire field of CN. OA journals relevant to CN investigators are either generic (e.g., *Scientific Reports*, *PLoS One*) or are oriented towards (neuro)biology (e.g., *eLife*) with a narrower readership in cellular, molecular, and systems neuroscience. There is a real need for an OA journal/platform, focused on CN, to serve as a high signal-to-noise nexus for solid high-quality research in the field, and to do so without being overly selective to the point that mainstream reports cannot be published. This journal could also be inclusive with respect to the communities that will contribute to this emerging body of knowledge, opening special issues to the multiple communities identified above. Importantly, such an OA platform should also expose data, analysis methods and workflows to assure reproducibility and other modes of data reuse. Achieving these aims in an international context that satisfies the requirements recently introduced by the European Union's General Data Protection Regulation is already an urgent undertaking.

How Can the Field Advance and Mature Along these Lines?

While some of the above-mentioned efforts could be facilitated by small organized groups of investigators who may obtain funds for national or international research coordination, our experience suggests that developing an organization infrastructure with these goals will produce larger dividends. Independently, we argue that these considerations should not be left to the initiative or discretion of funding-agency personnel. In a nutshell, program directors and administrators at private foundations or governmental agencies have several competing interests and demands on funding, some of which may not fully capture the importance of international coordination and collaboration, or may not sufficiently emphasize industry collaborations. They may therefore be less effective in devising potential strategic plans of actions as compared with groups of experts in these areas. Rather, an internationally oriented organization (or advisory board; see [6]) that can maintain effective connections with public and private funding bodies would be more effective in advancing such efforts. Such an organization would also be in position to identify and advocate for research areas suitable for larger-scale research coordination both in terms of paradigmatic and conceptual synthesis, and basic research. Such efforts may well surpass the usual scale of funded research projects and could involve international coordination among multiple agencies.

Cognitive Neuroscience can capitalize on such emerging opportunities to improve its science and increase its relevance. However, such efforts would require community-level coordination that is not yet in place.

Disclaimer Statement

No prior or current NSF personnel were consulted with, nor commented on, any version of this manuscript. The opinions expressed are the authors' alone and do not reflect in any way the position of the NSF.

[1]Center for Mind/Brain Sciences (CIMeC), University of Trento, Rovereto, TN, Italy
[2]Department of Psychology, The University of Chicago, Chicago, IL, USA

*Correspondence: uri.hasson@unitn.it (U. Hasson).

https://doi.org/10.1016/j.tics.2019.02.007

References

1. Waytowich, N. *et al.* (2018) Compact convolutional neural networks for classification of asynchronous steady-state visual evoked potentials. *J. Neural Eng.* 15, 066031
2. Angrick, M. *et al.* (2019) Speech synthesis from ECoG using densely connected 3D convolutional neural networks. *J. Neural Eng.* Published online 4 March 2019. http://dx.doi.org/10.1088/1741-2552/ab0c59 Published online 4 March 2019
3. Barsalou, L.W. (2017) What does semantic tiling of the cortex tell us about semantics? *Neuropsychologia*, 105, 18–38
4. Cichy, R.M. and Kaiser, D. (2019) Deep neural networks as scientific models. *Trends Cogn. Sci.* 23, 305–317
5. Wang, Y. *et al.* (2016) Real-time full correlation matrix analysis of fMRI data. In *2016 IEEE International Conference on Big Data (Big Data)*, pp. 1242–1251, IEEE
6. Gunsalus, C.K. *et al.* (2019) Overdue: a US advisory board for research integrity. *Nature*, 566, 173–175
7. Kragel, J.E. *et al.* (2015) Neural activity in the medial temporal lobe reveals the fidelity of mental time travel. *J. Neurosci.* 35, 2914–2926
8. Nunez, M.D. *et al.* (2017) How attention influences perceptual decision making: single-trial EEG correlates of drift-diffusion model parameters. *J. Math. Psychol.* 76, 117–130
9. Turner, B.M. *et al.* (2017) Approaches to analysis in model-based cognitive neuroscience. *J. Math. Psychol.* 76, 65–79
10. Nelson, M.J. *et al.* (2017) Neurophysiological dynamics of phrase-structure building during sentence processing. *Proc. Natl. Acad. Sci. U. S. A.* 114, E3669–E3678
11. Broderick, M.P. *et al.* (2018) Electrophysiological correlates of semantic dissimilarity reflect the comprehension of natural, narrative speech. *Curr. Biol.* 28, 803–809
12. Avesani, P. *et al.* (2019) The open diffusion data derivatives, brain data upcycling via integrated publishing of derivatives and reproducible open cloud services. *PsyArXiv* Published online 21 December 2018. http://dx.doi.org/10.31234/osf.io/y82n7 Published online 21 December 2018
13. Gorgolewski, K. *et al.* (2017) OpenNeuro—a free online platform for sharing and analysis of neuroimaging data. In *Proceedings from the 23rd Annual Meeting of the Organization for Human Brain Mapping (OHBM)*. Vancouver, Canada, p. 1677
14. Kennedy, D.N. *et al.* (2019) Everything matters: the repronim perspective on reproducible neuroimaging. *Front. Neuroinform.* 13, 1

Science & Society

Holding Robots Responsible: The Elements of Machine Morality

Yochanan E. Bigman,[1,*] Adam Waytz,[2] Ron Alterovitz,[3] and Kurt Gray[1]

As robots become more autonomous, people will see them as more responsible for wrongdoing. Moral psychology suggests that judgments of robot responsibility will hinge on perceived situational awareness, intentionality, and free will, plus human likeness and the robot's capacity for harm. We also consider questions of robot rights and moral decision-making.

Advances in robotics mean that humans already share roads, skies, and hospitals with autonomous machines. Soon, it will become commonplace for cars to autonomously maneuver across highways, military drones to autonomously select missile trajectories, and medical robots to autonomously seek out and remove tumors. The actions of these autonomous machines can spell life and death for humans [1], such as when self-driving vehicles kill pedestrians.