1

Introduction

Wolf Singer, Terrence J. Sejnowski, and Pasko Rakic

As science works to address any number of complex problems, a certain measure of humility must accompany its quest. Viewed over time, it is clear that myriad intricacies are often undervalued, as our collective wisdom and collaborative efforts have failed to resolve any number of issues. Although ultimate answers may be rare, this should not undercut the process of discovery or diminish the measurable progress that has been, or is currently being, made. It simply puts into context a truism: Science is an iterative process. As knowledge expands, each step forward requires us to test the concepts and ideas that emerge. To do this may require us to develop new methods or tools, which in turn may lead us to uncover completely new aspects of the problem that had hitherto escaped attention, thus bringing us back to a point where we need to evaluate, again, where things stand.

So it is, and has been, with our quest to understand the cerebral cortex.

Three decades ago, two of us (Pasco Rakic and Wolf Singer) chaired a Dahlem Workshop in Berlin on the neurobiology of neocortex. This gathering brought together forty distinguished neuroscientists from comparative and evolutionary biology, developmental neurobiology, neuroanatomy, neurophysiology, and behavioral neuroscience for an in-depth discussion of the cerebral cortex and an assessment of current research. The motivation behind this Dahlem Workshop was the realization that although research had advanced system by system and yielded an immense amount of data, the underlying rules and principles were defined for, and understood in, separate research areas, thus complicating communication and cross-disciplinary research. What was clearly lacking was an overarching theory of cortical organization—one that could account for general principles within particular areas as well as for cooperative interactions between cortical regions and cross-system generalities. From the numerous peer reviews of the results (Rakic and Singer 1988[1]), this book captured the conceptual understanding of the time and stimulated future research in developmental, cellular, functional, and cognitive neuroscience.

Years later, at an annual meeting of the Society for Neuroscience, we started to reflect on how the field had changed since that Berlin meeting: What seminal

[1] Although out of press, this publication is available freely online at https://esforum.de/publications/Neocortex_1/chaps/01_Neurocortex_Rakic_and_Singer.pdf (accessed 31.3.19).

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

discoveries had actually been made? Which questions remained unanswered, and what might be needed to address these now? Our discussions led us to explore whether it might be worthwhile to convene another group of experts to assess where things currently stand, in an effort to position research with the conceptual means to move ever forward. Marked by the emergence of completely new disciplines, several key areas demonstrated the extent to which research had expanded dramatically over the past three decades:

- Progress in genetics and molecular biology had revolutionized neuroscientific approaches in virtually all domains, from investigations of development all the way to studies of psychiatric conditions.
- The transfection of neurons and glial cells with genetically encoded marker molecules and the development of transgenic animal models had permitted comprehensive analyses of the brain's connectome, massive parallel recording of neuronal activity at the cellular level, as well as cell-specific interference with neuronal activity.
- The advent of noninvasive imaging technologies and methods to stimulate selected regions of the human brain had boosted the field of cognitive neuroscience.
- The availability of powerful and affordable computational resources now allow us to address the large data sets that were produced through advanced electrophysiological and optical recording methods.
- Last, but not least, the rapidly growing field of computational neuroscience enables us, for the first time, to test the validity of theories and concepts through simulation experiments that are able to cope, although still in a rudimentary way, with the mind-boggling complexity and dynamics of neuronal interactions.

This progress convinced us of the necessity for a new collaboration, yet to do justice to these novel developments, the scope of expertise needed to be broadened. We found a willing partner in Terry Sejnowski, who worked with us to develop a proposal for a forum that would explore the extent to which existing data could be embedded in unifying conceptual frameworks of the neocortex.

As the reader may be aware, major changes in 2006 impacted the Dahlem Workshops, and the institution no longer exists. Its guiding spirit, philosophy, and approach, however, continue to flourish in Frankfurt under the auspices of the Ernst Strüngmann Forum. (For an overview of this transition, see Singer 2016:475–476). Briefly, the Ernst Strüngmann Forum creates an environment that ensures open discourse and encourages divergent ideas. Long-established perspectives are questioned and disciplinary idiosyncrasies exposed. Consensus is never a goal. Instead, topics are examined from multiple perspectives: existing gaps in knowledge are exposed, key questions formulated, and ways of filling such gaps (through future research) are proposed. From April 8–13, 2018, the 27th Ernst Strüngmann Forum was convened in Frankfurt, Germany, to which 48 experts from diverse areas in neuroscience participated.

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

Even a week-long brainstorming encounter of this kind is unable to do complete justice to the state-of-the-art research that has unfolded over three decades, much let alone provide a comprehensive summary. Far more time and effort would be needed just to review the immense amount of data that has accumulated in virtually every domain of research into the cerebral cortex. What could be perceived as a "shortcoming," however, actually gives way to an important insight: In 1987, at the Dahlem Workshop, participants were by and large aware of the developments in the various disciplines and were able to understand the concepts and terminologies used in these fields. In 2018, at the Ernst Strüngmann Forum, transdisciplinary dialogue proved much more difficult: a plethora of abbreviations characterize the language of geneticists and molecular biologists, and the mathematical descriptions of complex dynamics and the highly differentiated taxonomies used in cognitive psychology posed substantial challenges to everyone.

At Dahlem, theories on cortical processing were still dominated by behaviorist concepts, which viewed the brain primarily as a stimulus-response machine. Accordingly, emphasis was placed on serial processing in feedforward architectures. The assumption was that detailed analysis of single-cell responses across the processing hierarchy, all the way up to executive centers, should ultimately permit comprehensive understanding of the system. Hence, the field was mainly interested in describing the gradual transformation of neuronal response properties from sensory surfaces across the hierarchy of cortical processing levels to executive organs. Common concepts for the investigation of sensory processes were feature-selective receptive fields, filter operations to reduce signal-to-noise ratios and redundancies, columns as functional units, maps for the orderly arrangement of neighborhood relations, representations of cognitive objects by responses of individual neurons, and (on the executive side) motor response fields, command neurons, and population vectors. As all information was assumed to be encoded in the discharge rate of neurons, the gold standard was the single-unit recording. Signals reflecting the temporal coordination of population activity, such as field potentials and EEG, were considered too coarse, and it was felt that they provided scant additional information. With a few notable exceptions (see below), these concepts are implemented in the architecture of perceptrons and Hopfield networks as well as their recent extension in deep learning networks. Because of the astonishing performance of these artificial systems in admittedly restricted domains, and because the architecture of these artificial neuronal networks shares similarities with some of the organizational features of the cerebral cortex, one might assume that we now possess valid and explicit models of brain function and hence are close to understanding how the cortex works.

At the Ernst Strüngmann Forum, it became clear that this optimistic view is not warranted; many of the concepts favored during the Dahlem Workshop needed to be abandoned or substantially modified due to novel insights that had since been gained. Importantly, we realized that we are probably further

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

away from a comprehensive understanding of the functions of the cerebral cortex than we imagined thirty years ago. As always in the empirical sciences, technological advances go hand in hand with conceptual developments. In addition to the still valid approach of feedforward processing, the comprehensive study of connectomics (both at the level of intracortical microcircuitry and inter-areal connections) forced us to consider

- functional implications of recurrent coupling within and between cortical areas,
- the immense density of information exchanged among processing streams,
- flat and often reversed hierarchy of putative interactions, and
- distributedness captured by graph theoretical terms such as rich club or small world networks.

These anatomical features are reflected by functional features that could, in part, have been discovered already by single-cell recordings at the time of the Dahlem Workshop. One of them is the concept of an invariant feature-selective receptive field. When feature-selective neurons were exposed to complex patterns, in particular in awake-performing animals, it became obvious how their responses are strongly sensitive to context, behavioral state, and top-down influences resulting from predictions, expectancies, and attention. It was recognized, however, that neuronal responses were variable and not always canonical, in particular in behaving animals, but this variability was attributed to noise fluctuations. The experimenter averaged over trials to extract the "essential" information, as the brain was supposed to average across a population of similar neurons. The new structural data also challenged the concept of columns as a functional unit. They suggest, at least outside input layer four, that horizontal coupling is reciprocal and continuous, even across boundaries between areas. Finally, the flat hierarchy and dense interconnectivity make it appear highly unlikely that areas operate in isolation and only serve as links in a serial processing stream.

Major arguments for an extension and reinterpretation of classical concepts came from experiments in which researchers recorded from more than one neuron at a time. It soon became clear that the fluctuations of neuronal responsiveness were correlated. Some maintain that these correlations reflect noise, hence the term "noise fluctuations." Others, however, observe that correlated firing contained information as it depended on stimulus configuration and behavioral context. Parallel recordings from electrode arrays have also revealed a puzzling but well-coordinated dynamics of cell populations. It was observed that individual neurons can engage in oscillatory patterning of their responses and that these temporally structured responses could synchronize with amazing precision in the millisecond range, depending on stimulus configurations, central states, and top-down signals. Furthermore, these oscillations are organized as traveling waves across the cortex and are both generated

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

spontaneously and induced by stimuli. After the discovery of these coordinated population dynamics in the cerebral cortex, very similar oscillatory phenomena and traveling waves were observed in another structure sharing essential features of recurrency and connected with the cerebral cortex: the hippocampus (Muller et al. 2018). These observations led to a renaissance of interest in dynamics and in recording methods able to capture spatially and temporally coordinated (synchronized) activity of local cell populations with multiunit activity (MUA), intracortical local field potentials (LFPs), electrocorticography recordings from cortical surface electrodes and, at a still coarser spatial and temporal scale, of EEG, MEG, and fMRI signals, respectively. Together with massive parallel recordings of single-cell activity, these approaches revealed a surprising degree of temporal coordination of distributed neuronal activity, both within and across cortical areas, including the nesting of oscillatory activity across distinct frequency bands. Finally, measurements of coherence allowed identification of stimulus and task-dependent formation and dissolution of widespread functional networks and to track the flexible routing of communication between cortical areas. Although the oscillatory patterning of EEG signals in distinct frequency bands was well established at the time of the Dahlem Workshop, and although it was known that these coarse signals reflect synchronized activity, these dynamic signatures of cortical processes were not considered in a functional context: they were merely taken as a state variable correlated with changes in sleep stages and arousal levels. One likely reason is that in the 1980s, most cortical physiology focused on the visual system, and it was thought that processing of (stationary) visual patterns required no computations in the temporal domain. Since then, however, increased research has been devoted to the auditory system, speech recognition, short-term memory, motor control, and spatial navigation, and interest in dynamic processes has increased. A role of precisely timed neuronal activity has also been recognized when it became clear that mechanisms of use-dependent synaptic plasticity were exquisitely sensitive to precise timing relations between pre- and postsynaptic activity, both during development and adult learning. In parallel, computational models became more dynamic, especially those that analyze the computational potential of recurrently coupled networks.

At this Ernst Strüngmann Forum there appeared to be a broad consensus that neuronal information processing capitalizes on the spatial as well as the temporal dimensions of the brain: not only the frequency but also the timing of discharges are informative. However, we are still at the very beginning of our attempts to explore the puzzling complexity of the dynamics that emerge from delay-coupled neuronal networks and to figure out whether and, if so, how the brain actually uses the exceedingly high-dimensional state space provided by these dynamics for computation and the storage of information. One possibility is that the brain exploits these dynamics to define relations that comply with the time-sensitive learning rules for the processing of temporally structured stimuli (the processing of sequences and language) as well as for the

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

realization of generative functions such as are required in predictive coding. In this context, it was noted as surprising that theories on cortical functions took so long to incorporate concepts of pattern generation and dynamic routing, as these had been present in the fields studying pattern generators in invertebrates, lower vertebrates, and insects.

The new evidence on the structural and functional organization of the cerebral cortex suggests that current concepts have to be considerably extended to do justice to the complexity and power of cortical computations. There was consensus that we have to learn to cope with the high-dimensional, nonlinear dynamics of the unimaginably complex interactions among the neurons of cortical networks, and that we will need new tools (e.g., machine learning) to decipher the information content in high-dimensional activity vectors as well as new mathematical instruments to analyze and interpret the trajectories of network states. Concerns were also expressed with respect to the requirement to provide causal evidence for the relations between neuronal activity and behavior. While new methods such as optogenetics and DREADDS permit cell-specific manipulation of neuronal activity, interference with the activity of nodes in a highly interconnected system may have uncontrollable consequences other than those intended. This may force the field to relax the criteria for the establishment of causal relations and in certain cases be satisfied with correlative evidence.

While the new data on connectomics and dynamics has precipitated a shift in concepts and paradigms, which is currently raising more questions than actual answers, the great advances in genetics and molecular biology have dramatically enhanced the resolution of investigations on developmental processes. The basic concepts involved in phylogenetic and ontogenetic development, formulated at the time of the Dahlem Workshop, seem to have passed the test of time. Still, much more is known now about the genetic and molecular networks that determine the birth, division cycles, migration paths, and differentiation steps of stem cells giving rise to excitatory and inhibitory neurons. Among the numerous new insights in the mechanisms determining the fate of precursor cells were the notions that inhibitory interneurons continue to be integrated into cortical circuitry during early postnatal development, that primates possess special mechanisms to increase neuron numbers in supragranular layers, and that genes have been identified that control the overall volume of the neocortex. Since participants at this Forum conduct research on a variety of animal models, the considerable species-specific differences were evident. For example, although radial glial cells in developing rodents are an excellent model to study some aspects of cortical development, the equivalent cells in primates (including humans) have specific genes and molecules as well as possess certain functional capacities that are absent in all subprimate species analyzed thus far. The difference between primary visual cortex in primates and nonprimates is obvious. Likewise, rodents do not even possess some of the cytoarchitectonic and functional areas (e.g., dorsal prefrontal association

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

cortex, Broca and Wernicke areas), which have different neuronal composition and pattern of connections. Thus, the development, anatomy, and function of some human-specific cortical features can only be studied in humans.

In conclusion, and in keeping with the overall nature of science, this Forum was a sincere attempt to understand the major developments that have taken place in neocortex research over the last thirty years, ever with an eye toward the future. It is clear that a large number of methodological breakthroughs in all disciplines of the life sciences drove progress forward, that the analysis of massive new data (especially the big data on connectomics and molecular diversity) is reliant on powerful computational tools, and that a substantial amount of new data has been acquired only through large cooperative efforts, as opposed to research in small groups characteristic of neuroscientific investigation thirty years ago. Equally, however, it is clear that conceptualization lags behind data accumulation. Thus, we posit that the greatest challenge for future endeavors will be to integrate the plethora of facts generated by the highly diverse fields of research into an overarching comprehensive theory on cortical functions. Whether this is at all possible—whether there is even such a thing as a unifying theory of neocortex—remains an open question. Perhaps accumulated knowledge must remain distributed across the community of specialized experts, similar to how functions of the cerebral cortex are distributed. Just as the brain, as a whole, produces intuitively plausible behavior, distributed knowledge might serve to explain a large number of normal and pathological behaviors, ultimately enabling the development of useful tools without meeting the epistemic challenge of having to fit into a unified theory.

Acknowledgments

We wish to express our gratitude to Silke Bernhard, the director of the Dahlem Workshops from 1974 until 1989. We know of no better way to pay tribute to her memory than to continue the dialogue that took root in Berlin some thirty years ago. Equally, we wish to acknowledge the efforts of Andreas and Thomas Strüngmann, who have enabled this unique approach to continue in Frankfurt. Their vision and support of the Ernst Strüngmann Forum is invaluable to basic science, for it offers researchers a much-needed yet rare opportunity to reflect, to reanalyze, to correct, and to propose directions for future research to pursue. On behalf of all participants, we thank you sincerely.

From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3