transmission of Zika and Dengue viruses to humans using defensive symbionts. By putting *Wolbachia* into mosquitoes, these insects might become less competent vectors of pathogenic viruses. Although in its infancy, the application of defensive symbionts is an exciting area ripe for future development.

What remains to be discovered? A lot, but particularly the evolutionary history and consequences of defensive symbioses. Microbes can evolve quickly — it has been shown that bacteria inside the guts of *C. elegans* (Figure 1D) can evolve within days to protect hosts. The flexibility and speed with which defensive symbionts evolve are not well understood. Nevertheless, there is potential for symbionts that defend longer-lived hosts (for example, mammals and trees) to evolve and counter infection faster than hosts themselves, thereby 'taking over' the coevolutionary race with parasites. Defensive symbionts might also be more effective and metabolically cheaper than immune systems. Given all these benefits, why fight the battle yourself? Recruit a defensive symbiont to do the job for you.

Where can I find out more?

Arnold, A.E., Mejia, L.C., Kyllo, D., Rojas, E.I., Maynard, Z., Robbins, N., and Herre, E.A. (2003). Fungal endophytes limit pathogen damage in a tropical tree. Proc. Natl. Acad. Sci. USA *10*, 15649–15654.

Bates, K.A., Clare, F.C., O'Hanlan, S., Bosch, J., Brookes, L., Hopkins, K., McLaughlin, E.J., Daniel, O., Garner, T.W.J., Fisher, M.C., *et al*. (2018). Amphibian chytridiomycosis outbreak dynamics are linked with host skin bacterial community structure. Nat. Commun. *9*, 693.

Ford, S.A., and King, K.C. (2016). Harnessing the power of defensive microbes: evolutionary implications in nature and disease control. PLoS Pathogens *12*, e1005465.

Hopkins, S.R., Wojdak, J.M., and Belden, L.K. (2017). Defensive symbionts mediate host-parasite interactions at multiple scales. Trends Parasitol. *33*, 53–64.

Jaenike, J., Unkless, R., Cockburn, S.N., Boelio, L.M., and Perlman, S.J. (2010). Adaptation via symbiosis: recent spread of a *Drosophila* defensive symbiont. Science *329*, 212–215.

Palmer, R.M., Stanton, M.L., Young, T.P., Goheen, J.R., Pringle, R.M., and Karban, R. (2008). Breakdown of ant–plant mutualism follows the loss of large herbivores from an African Savannah. *Science 319*, 192–195.

Vorburger, C., and Perlman, S.J. (2019). The role of defensive symbionts in host-parasite coevolution. Biol. Rev. *93*, 1747–1764.

Department of Zoology, University of Oxford,
Oxford OX1 3PS, UK.
E-mail: kayla.king@zoo.ox.ac.uk

Primer

Electroencephalography

Andrea Biasiucci[1], Benedetta Franceschiello[1,2] and Micah M. Murray[1,2,3,4,*]

Electroencephalography (EEG) is the non-invasive measurement of the brain's electric fields. Electrodes placed on the scalp record voltage potentials resulting from current flow in and around neurons. EEG is nearly a century old: this long history has afforded EEG a rich and diverse spectrum of applications. On the one hand, foundations of EEG in clinical diagnostics have dovetailed more recently into brain-triggered neurorehabilitation treatments. On the other hand, EEG has not only been a workhorse for providing brain correlates of constructs in the field of experimental psychology, but has also been used as a true neuroimaging method with more recent extensions in translational as well as computational neuroscience. The versatility and accessibility of the technique, in combination with advances in signal processing, allow for this 'old dog' to still deliver new tricks and innovations.

Electroencephalography's biophysics and measurement

We introduce the uninitiated reader to electrophysiology in general and EEG in particular through a set of analogies intended to provide a clear illustration of which neuronal activity results in measurable signals at the scalp surface (Figure 1). Imagine you are a journalist equipped with a hand-held microphone, which will here be analogous to a recording electrode. You are reporting from a soccer match. If you are standing next to the coach, you can interview her and comprehend her voice despite the noise throughout the stadium. This is akin to recording action potentials of individual neurons. If you are in the press box, you will not be able to record the ongoing conversations between the coach and players on the field, but instead can capture the general commentary of other reporters inside the press box as well as the hum of the audience outside. This is akin to recording local field potentials, where there are contributions of both proximal and relatively distal events. Finally, from your hotel balcony, having lost your press credentials, you may nonetheless be able to hear the joyful cry in unison of the team's supporters from within the stadium when a goal is scored. This is analogous to EEG recordings.

It is important to emphasize that EEG can detect only a portion of all the varieties of electrical activity going on in the brain, and does so despite the co-occurrence of other kinds of both physiologic electrical activity (such as cardiac, ocular, and other muscular activity) and environmental noise (such as computer screens and other electric equipment, power lines). To be clear, EEG does not measure action potentials, but rather postsynaptic potentials. Action potentials are the rapid current flow from the soma along the axon, resulting from the depolarization (making more positive) of a neuron from its typical –70 mV resting potential to –55 mV. By contrast, postsynaptic potentials result from relatively slower currents subsequent to neurotransmitter release at the axon's terminal boutons.

In part, the anatomical geometry of individual pyramidal neurons (and their orderly columnar arrangement in most cortical structures) facilitates the measurement of EEG in the following way. An excitatory postsynaptic potential at an apical dendrite will result locally in an intracellular current source or positivity (and an extracellular current sink or negativity). At the soma, there will be an intracellular current sink and extracellular current source. These source–sink configurations are also known as current dipoles. They are the main source of potentials measured by EEG. In order to be measurable at the scalp surface, neuronal populations need to be active essentially simultaneously. This allows for the summation of currents which then conduct in an isotropic fashion, independently of their frequency spectra, throughout the brain volume and in turn through the vasculature,

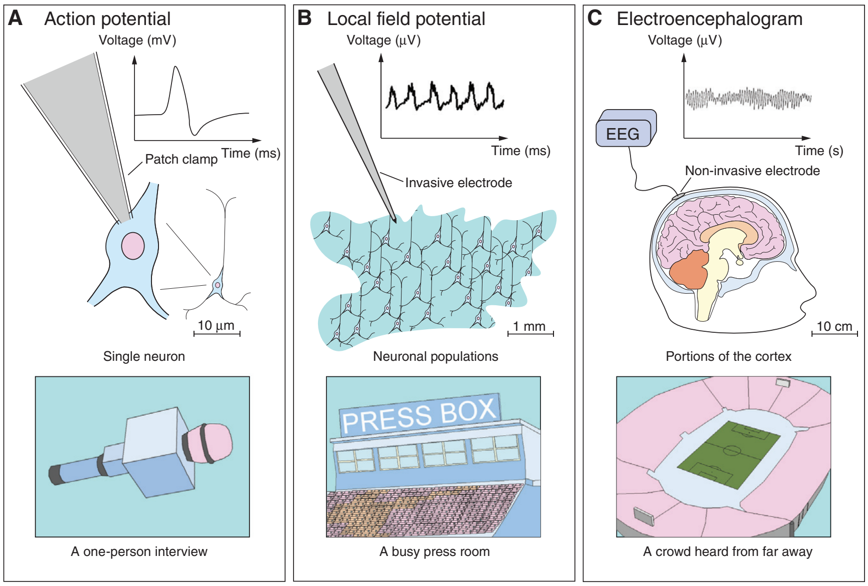

Current Biology

Figure 1. Recording the electrical activity of the brain.
The activity of single neurons, of populations of neurons and of portions of the cortex can be measured directly by means of electrical signals. (A) The patch-clamp technique is used to record action potentials from individual neurons (single-unit recordings). Here, a pyramidal neuron is schematized. Action potentials are short-duration (1 ms), high-amplitude (~100 mV) pulses. (B) Electrode microarrays can be inserted in brain tissue to record the activity of populations of neurons. Depending on filter settings, one can isolate multi-unit activity (action potentials) or lower-frequency local field potentials (LFPs); the latter of which have spectral characteristics similar to scalp-recorded EEG. (C) In electroencephalography (EEG), macroscopic electrodes placed on the surface of the scalp measure the electrical activity of large portions of the brain. EEG oscillations vary according to the synchronized or desynchronized activity of underlying neuronal populations. In addition, the signal recorded on the scalp is attenuated and distorted by volume conduction through various intracranial media as well as the scalp. By analogy, single-unit activity is akin to a one-on-one interview with a coach. LFPs are like capturing commentators on a football match in a press box. EEG is like hearing the roar of the crowd from outside the stadium.

cerebral spinal fluid, dura matter, skull, muscles, fat and skin to the EEG electrodes.

The quintessential technology involved in recording EEG involves the combination of electrodes, composed of conductive materials, and operational amplifiers (Figure 2). Typically, the resistive contact between the electrode and skin is improved with electrolytic gels or salts. Another, more recent approach has been so-called ‘dry’ electrodes that capitalize on innovations in material sciences as well as electronics to minimize the necessity of preparing participants’ scalps and therefore to reduce the setup time. It is likely that these innovations will be paralleled by progress in printable electrode technologies as well as continued development of polymers and wearable devices. In the interim, contemporary EEG electrode caps are tolerated well by people of all ages from premature neonates to the elderly. Innovations in amplifier design have allowed for faster sampling rates and increased numbers of simultaneously recorded channels. Standard commercial EEG systems, which are also approved for clinical use, can readily acquire data from at least 128 channels, with greater than 10 kHz sampling rate at all channels, and with 24 bit resolution at each amplifier.

Such a system typically costs less than $60,000 and will have a lifetime of at least 10 years. In terms of cost-effectiveness, the reader might bear in mind that a 3 Tesla MRI scanner costs on the order of $2–3 million, and standard magnetoencephalography (MEG) equipment carries a similar price tag. A further practical benefit of EEG is that it is easily combined with other brain mapping and imaging methods — MRI, MEG, functional near-infrared spectroscopy, non-invasive brain stimulation, and so on — as well as with neuropharmacological, physiological, and interventional regimes. Systematic improvements in the portability of EEG systems allow for recordings in real-world environments including at the bedside as well as in classrooms and on athletic fields. What is more, modern systems allow for precisely synchronized recordings from multiple individuals; so-called EEG hyperscanning. Lastly, advances in signal processing permit on-line analyses as well as neurofeedback that can in turn be used for brain–machine interfaces as well as control of stimulus delivery to augment perception and/or performance by waiting for the ‘optimal’ state of the subject.

A brief history

Hans Berger, credited as the discoverer of the human EEG in the late 1920s, devised the technique in order to provide a ‘window onto the brain’ (Figure 2). He recorded signals that fluctuated rhythmically when the eyes were shut, but which became far less rhythmic and of generally smaller amplitude when the eyes were open. The scientific community originally dismissed Berger and his EEG. This may in part have been because Berger was something of a loner and also because of his convictions regarding telepathic phenomena. Some dismissed the scalp EEG as a cardiac or muscle artifact. Others contended that brain activity should not become less rhythmic and of generally smaller amplitude when the eyes were opened (today this phenomenon is recognized as ‘alpha blocking’). Still others contended that the rhythmic fluctuations that Berger measured were simply too slow to reflect actual neural activity, which at the time was assumed to be limited to action potentials.

It was only when the British physiologists Edgar Adrian and Bryan Matthews replicated Berger’s observations in 1934 that EEG was accepted as a non-invasive measure of the electric fields of the brain (it should also be noted that Adrian and Matthews fully credited Berger with the discovery of scalp EEG). Unfortunately, Berger’s research was halted by

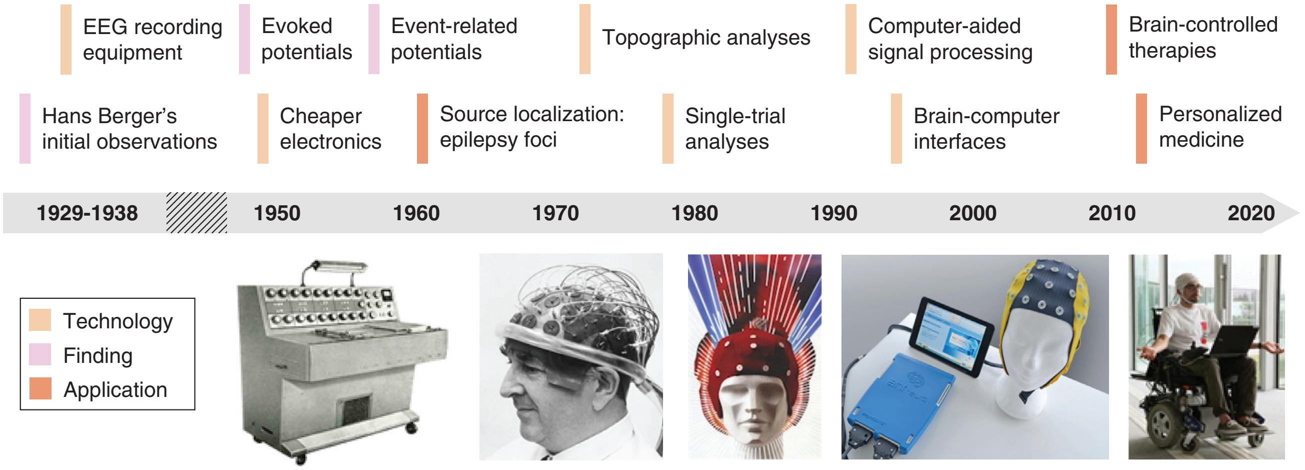

Current Biology

Figure 2. The evolution of human EEG technology and its applications.
Timeline of the evolution of electroencephalography in terms of its technology, exemplar findings, and applications. The images, from left to right, represent: a 9-channel Nihon Kohden ME-91D Electroencephalograph (from 1959: this was Japan's first 9 channel EEG, courtesy of Nihon Kohden Corporation); a picture of EEG pioneer Derek Fender (© Bettmann, Getty Images); a photo of one of the iconic designs of EEG caps from the1980s ("In all the world there's only one", courtesy of Electro-Cap International, Inc.); a modern portable waveguard™ original EEG cap and eego™ sports EEG amplifier (courtesy of ANT Neuro b.v.); and an EEG-controlled wheelchair (courtesy of José del R. Millán and EPFL).

the Nazi regime in Germany, and he ultimately committed suicide in 1941. Nonetheless, Berger's legacy has been the introduction of a technique that rapidly (if not immediately) took a firm foothold not only in clinical applications, but also in research domains from neurophysiology to computer science.

What is done with EEG (and what perhaps should not be)

Arguably, the versatility and accessibility of EEG have been a double-edged sword in terms of how EEG has been used and viewed by the scientific community. On the one hand, because EEG is a direct and real-time measure of the brain's neural activity, it is possible to characterize the integrity of specific neurophysiologic pathways, states of consciousness/sleep, as well as the precise temporal dynamics of brain (dys)function. Moreover, when this exquisite temporal resolution is combined with estimations of the underlying sources, it becomes feasible to also characterize brain networks, their connectivity, as well as the extent to which specific activity can be considered serial *versus* parallel during the course of a given function.

On the other hand, because EEG is a measure of electrical potentials — voltage describes the work needed to move electrical charges between two locations without accelerating them — it is absolutely critical for users of EEG (and electrophysiology more generally) to fully understand the analytical and interpretational consequences of measurements between an active and reference site. Further compounding this concern is that the neural activity generating EEG conducts throughout the brain volume, meaning that one cannot assume that a signal recorded at one point on the scalp has its source directly below. The situation is analogous in many regards to a geographical surveyor. On the one hand, they must measure altitude against an arbitrarily defined scale: sea-level is a reference that is relative rather than absolute. On the other hand, measuring only the highest peak provides little information about the terrain.

Spectral decomposition

If one foregoes the above points for the time being, it is possible to decompose the EEG into a series of sine waves to generate a frequency spectrum of the data. These waves can be characterized at each point in time by their amplitude, or power in the case of a rectified signal, and phase, which is where the sine wave is positioned in a cycle, measured either in degrees from 0 to 360 or in radians from 0 to 2. The frequency bands are often *a priori* delimited. For example, many studies divide the EEG into delta (δ: ~0.2–3.5 Hz), theta (θ: ~4–7.5 Hz), alpha (α: ~8–13 Hz), beta (β: ~14–30 Hz), gamma (γ: ~30–90 Hz) and (very) high frequencies (>90 Hz). These subdivisions have subsequently been shown to have empirical bases. One pitfall, however, has been that frequency bands were effectively interpreted as themselves directly mediating a specific brain process in a 1-to-1 fashion, despite evidence that oscillatory activity is nested with itself and exhibits complex amplitude–phase relationships across frequencies. And in addition to changing its amplitude and phase, oscillatory activity can also accelerate and decelerate.

Another pitfall is that the decomposition process itself assumes that the EEG is exclusively a mixture of sinusoids, whereas non-sinusoidal signals as well as arrhythmic activity are also intrinsic features of EEG. These pitfalls underscore how our understanding of the biophysical mechanisms generating EEG is partial and incomplete. Despite these pitfalls, frequency decompositions are an effective means for quantifying and distinguishing between two main classes of brain activity: evoked and induced. Evoked activity is both time-locked and phase-locked to an event, such as a visual stimulus. Evoked activity can be observed reliably in response to single events, but is often of relatively small magnitude. Consequently, most studies of evoked activity perform signal averaging, taking many trials of the same or similar event to increase the signal-to-noise ratio (after all, the brain is never silent). This procedure yields the event-related potential (ERP) as an output. By contrast, induced activity is neither phase-locked nor time-locked to an event.

The reference problem and what to do about it

ERPs have been and remain a major workhorse across both clinical and research applications of EEG. Like any technique, EEG/ERPs can and do result in misuse and misinterpretation. Some of this can be remedied through a better understanding of the biophysics of EEG measurements. The measurement of voltage requires a reference site, but there is no ideal reference, because nowhere on the scalp or body is there a perfectly electrically neutral locus. Voltage time series will therefore change their shape when a different reference is used. Just so, the variance and by extension statistical results also change with a change in the reference site. The consequence is clear; one obtains divergent results with different choices of reference site. The quagmire is also clear; there is no objective way of deciding which results ought to be taken as the ground truth, and one cannot simply replicate others' potentially erroneous decisions.

It is at this point where one might be tempted to abandon altogether the use of EEG/ERPs in favor of techniques that do not suffer from this reference problem (this is indeed a major advantage of magnetoencephalography, despite its price tag and practical inconveniences). A more constructive tactic is to return to the biophysics of EEG measures and the benefit of multi-channel recordings. By instead focusing on reference-independent and global measures of the electric field at the scalp, it becomes evident how EEG/ERP can objectively and quantitatively speak to core processes relevant for both within-subject and between-subject designs in the full spectrum of research. These processes include whether effects are driven by modulations in response strength (gain), modulations in the network of active brain regions, modulations in the timing or duration of cerebral processes, or any combination thereof (and moreover as a function of time). What is more, such measures also address current trends for 'big data' by providing clear ways of homogenizing datasets and extracting relevant features across different laboratories and acquisition parameters.

Reference-independence is a spatial, rather than temporal, notion. In the case of EEG/ERPs, this notion manifests in the fact that the shape of the electric field at the scalp (the topography) is independent of the reference in the same manner that the mountains surveyed in the geographer analogy above do not change shape when sea-level changes. Put alternatively, while the reference will influence the specific value ascribed to a given location (whether microvolts in the case of electrodes or altitude in the case of mountains), the spatial gradients are immutable across the global terrain (scalp in the case of EEG/ERPs). Moreover, and because biophysical laws have established that changes in topography are indicative of changes in the configuration of the underlying sources in the brain, quantitatively measuring topographic features equips users of EEG/ERPs with the ability to analyse whether and when networks in the brain are changing. It thus follows logically that increased spatial sampling provides better characterization of EEG/ERP topography.

Pioneers of EEG/ERP, including individuals such as Herbert Vaughan Jr., Derek Fender, and Dietrich Lehmann, already recognized during the 1960s (and perhaps earlier) the importance and value of EEG analyses based on spatial features. This led to breakthroughs not only in EEG technologies (Figure 2), but also in interpretations of EEG signals. For example, Vaughan Jr. and his colleagues, including figures like Joseph Arezzo, Daniel Javitt and Charles Schroeder, focused in large part on understanding the neurophysiological underpinnings of EEG/ERP signals with a particular emphasis on comparative studies across non-human primates and humans. They achieved this by conducting topographic and source estimation analyses of scalp-recorded data alongside intracranial measurements of multi-unit activity and local field potentials measured simultaneously from all cortical layers. Contemporaneously, Lehmann and his colleagues, including figures like Daniel Brandeis, Thomas König, Christoph Michel, Wolfgang Skrandies, and Werner Strik, described and quantified how EEG and ERPs are structured, rather than chaotic, both in their spatial and temporal properties. Lehmann named these structures 'functional microstates' and posited that they constitute the monads or 'atoms of thought'. This proposition has been borne out across domains in clinical and basic research.

Many others during the 1960s and 1970s focused their efforts on using EEG/ERPs to identify correlates of perceptual and psychological constructs. These efforts gave rise to a plethora of prototypical ERP response components and a veritable alphabet soup of monikers. Individuals such as W. Grey Walter, Samuel Sutton, Emanuel Donchin, Steven Hillyard, Risto Näätänen and their colleagues deserve particular recognition for their (ongoing) contributions to the field. From a purely signal processing standpoint, an ERP component is defined by two features: its latency relative to the stimulus or event and its topography. More often than not, however, ERP components have been defined by their latency and amplitude, in particular polarity, at a given scalp location (*versus* a pre-selected reference) and are in turn interpreted as reflective of a given sensory, perceptual, or motor process in a nearly one-to-one fashion similar to the spectral decomposition described above.

The old dog still teaching us new tricks

EEG's analytic potential and range of applications have yet to be exploited fully (Figure 3). For a time, EEG stood alone as a non-invasive technique for measuring brain function. With

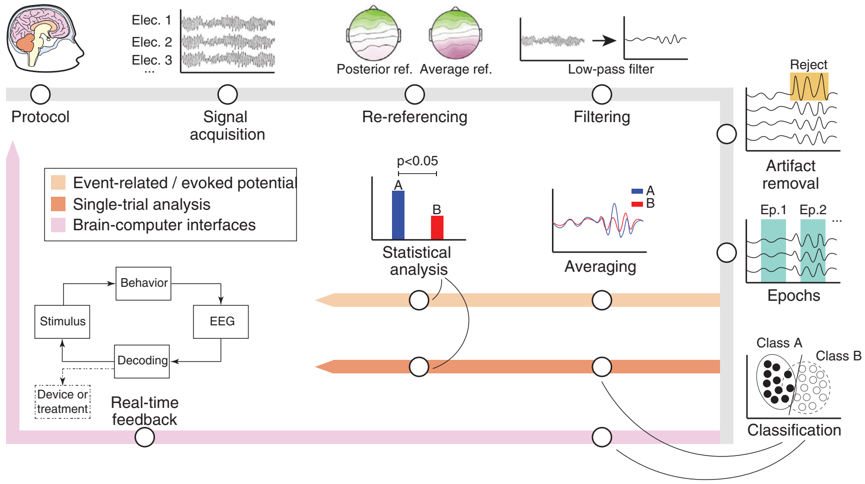

Current Biology

Figure 3. How to process and analyse EEG.

Here, we present three ‘flows’ depicting some of the EEG processing and analysis steps. The gray part of the diagram is shared among all analyses, and is commonly referred to as ‘pre-processing’. During this phase the signal is acquired, and each channel’s recording is referenced to shared voltage. Then, signals are filtered in the frequency domain to exclude potential non-biological noise. The signal is inspected for artifacts due, for example, to muscular activity and relevant epochs of data are extracted for subsequent processing. At this point, epochs of different ERPs are grouped and averaged, obtaining sets of time series that are compared statistically to assess significant differences. Single-trial analysis treats epochs differently: instead of averaging, after grouping epochs by condition it involves training a classifier — basically an algorithm able to discern the condition that generated the signal from a set of features of the signal itself — and providing a statistical measure of how well the classifier performs. Brain–computer interfaces use the same classification scheme as single-trial analyses, but instead of quantifying classifier performance they provide feedback on the current decision of the machine to the user in real time.

the advent of techniques like MEG in 1968, positron emission tomography in 1975, transcranial magnetic stimulation in 1985, and functional magnetic resonance in 1990, EEG no longer held a monopoly. As is often the case, one or another method became more fashionable alongside thematic trends, such as functional localization or connectomics. In a market-driven manner, such competition among methods has often led to innovation, and EEG is no exception. With a renewed interest by the scientific community in brain dynamics as well as in real-world applications, EEG is experiencing a renaissance and is regaining its position as a pre-eminent neuroscience technique.

EEG as a neuroimaging technique

Let there be no confusion: EEG is both a brain mapping as well as brain imaging tool. The loci of intracranial sources can be reasonably well estimated as a function of time based on surface recordings at the scalp. This electromagnetic inverse problem can be solved, and simulations as well as empirical findings have demonstrated localization errors below 1 centimeter. But the solution to the inverse problem is both ill-poised and non-unique, in part because the brain and its coverings behave like a volume conductor, meaning that any electrode on the surface detects activity from distal portions of the brain to one degree or another.

Mathematical solutions can be improved by incorporating biophysically based constraints; for example, only the brain and its grey matter produce EEG, and EEG records only ohmic currents. Because MEG does not suffer from either the abovementioned reference problem nor from volume conduction, some have claimed that MEG has higher spatial resolution. But it should not be overlooked that EEG is not only sensitive to both radially and tangentially orientated dipolar fields, but is also capable of detecting activity from both superficial as well as deep sources. By contrast, MEG detects only tangential fields from superficial sources. These aspects make it challenging to directly compare localization accuracy. Nonetheless, when data from the same number of sensors were compared, EEG outperformed MEG (though this remains a subject of debate that will surely produce continued innovations).

Computational neuroscience

While other branches of neuroscience have already greatly benefitted from computer-based simulations to provide explanations of neural activity beyond what could be observed in classical electrophysiology, only in recent years have computational models demonstrated their potential for EEG. Just like virtual reality in experimental psychology, computational models allow experimenters to simulate situations and interactions that are very hard to test and observe in the real world. The first computational models in neuroscience aimed to describe the firing activity of different types of neurons based on their spiking activity. Simulated neurons precisely replicated the neural activity at a single-cell level, providing biologically meaningful insights compatible with intracranial recordings. In addition, computational models were used to characterize networks of neurons rather than individual units, thus enabling the identification of a signal across a neuronal population.

In contrast with these approaches, other branches of computational models dealt with the description of behavioural or perceptual mechanisms, providing a macroscopic and qualitative analysis of phenomena. These models are very helpful in investigating the relationships across areas and the causality of a single phenomenon or experimental task. State of the art methods in computational models described how population dynamics could determine specific features of the electrical activity recorded by EEG. In particular, these methods addressed how network modifications, for example representing a neurological accident, could explain the variation of the recorded EEG signal. The use of qualitative models of the brain, embedding features of behaviour and perception, together with population network dynamics, represents one of the most innovative and promising developments in EEG and neuroscience.

Brain decoding and real-world neuroscience

The growing interest in understanding variability in data at the group and individual levels is a major driving force behind the development of single-trial analyses, thereby providing more direct links between brain activity and behavior. Before the introduction of these techniques, ERPs were extracted from the 'noisy' background of the ongoing EEG through signal averaging, thus requiring stimuli be repetitively presented under the same conditions. In practice, it is often infeasible to maintain the participant and their environment under stable conditions throughout the duration of an experiment that would allow for sufficient numbers of trials for all conditions of interest. Just so, it is unreasonable to think that the state of the participant remains unchanged throughout an experiment. For example, there is recognized contribution of pre-stimulus state on stimulus processing and perception.

The advent of machine learning techniques in the treatment of biophysical signals allows experimenters and clinicians to determine which of one-among-many possible response features had been elicited or generated on each trial. In unsupervised methods, the classification emerges from clustering of observed features. In supervised methods, the classifier is trained *a priori* on an independent dataset (or portion thereof). A discriminant function is then used in order to classify a 'test' dataset group membership is unknown. A key aspect of machine learning is therefore 'the' feature or features of interest. These can be identified either based on the performance of the classifier or based on their neurobiological significance, though these two extremes need not be mutually exclusive. This aspect is perhaps best illustrated in the case of brain–computer interfaces (BCIs), which refer to those systems that can be trained to recognize the brain signatures associated to specific tasks and decode the current mental task of a user in real-time.

BCI technology was initially developed to provide a communication and control channel for severely impaired patients to control telepresence robots, spellers, and wheelchairs by thoughts. More recently, brain-triggered therapies for stroke rehabilitation have been tested in clinical studies. These investigations consistently showed that long-lasting recovery of motor function can be achieved also years after the initial accident, and that the modulation of EEG motor rhythms plays a crucial role in the recovery process. The capacity to decode mental states in real-time and modify the feedback to the subject accordingly opens unprecedented opportunities in neuroscience. The possibility to move away from trial-based studies makes the quest for a "neuroscience of everyday-life" more actual than ever.

Conclusions and outlook

EEG has come a long way over the past century. Despite these remarkable achievements, at least two major issues represent pressing challenges. First, the complexity of the neurophysiology across macroscopic, mesoscopic and microscopic levels means that the understanding of the exact generation and functional significance of EEG remains in its infancy and a domain of active research and a certain degree of debate. Second, the status quo in the field is ill-defined. The analysis of the electroencephalogram lacks standardization of the processing steps. While other neuroimaging techniques have consolidated evidence-based workflows, there remains little consensus on a shared pipeline across clinical and research communities. This will be critical for any contribution of EEG to big-data initiatives. The use of a set of unified analysis approaches, when it comes to a transversal and translational technique like human EEG, would greatly facilitate the definition of standards and the translation of findings from research to clinical settings. Such issues notwithstanding, EEG (and electrophysiology more broadly) is an exceptionally powerful neuroscientific tool widely used not only in humans, but also across species and contexts from the laboratory to the real-world.

FURTHER READING

Biasiucci, A., Leeb, R., Iturrate, I., Perdikis, S., Al-Khodairy, A., Corbet, T., Schnider, A., Schmidlin, T., Zhang, H., Bassolino, M., *et al*. (2018). Brain-actuated functional electrical stimulation elicits lasting arm motor recovery after stroke. Nat. Commun. *9*, 2421.

Birbaumer, N., Ghanayim, N., Hinterberger, T., Iversen, I., Kotchoubey, B., Kübler, A., Perelmouter, J., Taub, E., and Flor, H. (1999). A spelling device for the paralysed. Nature *398*, 297–298.

Cohen, M.X. (2017). Where does EEG come from and what does it mean? Trends Neurosci. *40*, 208–218.

Hari, R., and Puce, A. (2017). MEG-EEG Primer (Oxford: Oxford University Press).

Kajikawa, Y., and Schroeder, C.E. (2011). How local is the local field potential? Neuron *72*, 847–858.

Kappenman, E.S., and Luck, S.J. (2011). The Oxford Handbook of Event-Related Potential Components (Oxford: Oxford University Press).

Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S.F., and Baker, C.I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. Nat. Neurosci. *12*, 535–540.

Lefebvre, J., Hutt, A., Knebel, J.-F., Whittingstall, K., and Murray, M.M. (2015). Stimulus statistics shape oscillations in nonlinear recurrent neural networks. J. Neurosci. *35*, 2895–2903.

Lehmann, D. (1987). Principles of spatial analysis. In Handbook of Electroencephalography and Clinical Neurophysiology, *Vol. 1*: Methods of Analysis of Brain Electrical and Magnetic Signals., A.S. Gevins, and R. A, eds. (Amsterdam: Elsevier), pp. 309–354.

Lopes da Silva, F. (2013). EEG and MEG: Relevance to neuroscience. Neuron *80*, 1112–1128.

Michel, C.M., and Murray, M.M. (2012). Towards the utilization of EEG as a brain imaging tool. Neuroimage *61*, 371–385.

Michel, C.M., Murray, M.M., Lantz, G., Gonzalez, S., Spinelli, L., and Grave de Peralta, R. (2004). EEG source imaging. Clin. Neurophysiol. *115*, 2195–2222.

Murray, M.M., Brunet, D., and Michel, C.M. (2008). Topographic ERP analyses: a step-by-step tutorial review. Brain Topogr. *20*, 249–264.

Nunez, P. L., and Srinivasan, R. (2006). Electric Fields of the Brain. (Oxford: Oxford University Press.)

Stone, J.L., and Hughes, J.R. (2013). Early history of electroencephalography and establishment of the American Clinical Neurophysiology Society. J. Clin. Neurophysiol. *30*, 28–44.

Tivadar, R.I., and Murray, M.M. (2019). A primer on electroencephalography and event-related potentials for organizational neuroscience. Organ. Res. Meth. *22*, 69-94.

Vaughan, H.G. (1982). The neural origins of human event-related potentials. Ann. N.Y. Acad. Sci. *388*, 125–138.

[1]The Laboratory for Investigative Neurophysiology (The LINE), Department of Radiology, University Hospital Center and University of Lausanne (CHUV), Lausanne, Switzerland. [2]Ophthalmology Department, University of Lausanne, Fondation Asile des Aveugles, Lausanne, Switzerland. [3]EEG Brain Mapping Core, Center for Biomedical Imaging (CIBM) of Lausanne and Geneva, Lausanne, Switzerland. [4]Department of Hearing and Speech Sciences, Vanderbilt University Medical Center, Nashville, TN 37232, USA.
*E-mail: micah.murray@chuv.ch