



The Case for Conscious Experience Being in Individual Neurons

Jonathan Edwards¹, Pavel G Somov

¹ University College London, University of London

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.

Abstract

The idea that individual nerve cells might have conscious experiences has been around ever since cells were identified in the seventeenth century, but in the era of modern neuroscience the case for individual human neuronal experience has received little attention. A series of arguments will be presented suggesting that all the human conscious experiences that we talk about are events in individual neurons, not global to the brain or organism. We conclude that cellular consciousness is the only plausible way to explain 'our' experiences within current physics and biology, however implausible it might at first seem. This implies that our experiences are multiple in space as well as time, consistent with the neuroscience watchword that there is 'no single place where everything comes together in the brain'. The central argument is that events of experience must involve rich integration of information and individual neurons are the only places in brains where integration of information occurs. Any more global 'binding' is neither required nor physically possible. The detailed nature of events of integration of signals in a neuron's dendrites remains uncertain but recent developments provide some candidates.

Introduction

Around the year 2000, Edwards and Sevush independently came to the conclusion that our conscious experiences must be events in individual neurons, not spread across neuronal networks or brain regions (Edwards, 2005; Sevush, 2006). Somov (2012) has come to a similar position. Our routes to this unexpected conclusion proved almost identical, centred on the realisation that neurobiology provided no other possible option. The main objection to the idea is that it conflicts with our intuitive sense of self. Conscious events in individual neurons would make our subjectivity massively multiple. Yet it appears this is how neuroscience shows things must be.

The philosopher David Chalmers recently won a 1998 bet with neuroscientist Christof Koch that, 25 years later, even the general form of the neural correlates of consciousness – the biophysical events in brain that most closely underlie our conscious experiences – would remain unknown (Horgan, 2023). Ironically, if what follows below is roughly correct, at least that general form may have come to light within about 5 years¹ – as the direct electromagnetic coupling between a pattern of synaptic potentials and an indivisible dynamic unit, possibly an acoustic or phonon-based mode (in simple terms a sound wave), within a neuronal dendritic tree – occurring in each of many cells in parallel (Edwards, 2020).

The question that dominates almost all attempts to understand the neural correlates of consciousness is how events of experience come to include rich patterns that are in some way unified, ‘bound’ or combined into single events. Theories of mechanisms are driven by a desire to find some elusive non-local ‘binding’ of signals, perhaps through some form of ‘entanglement’ or ‘resonance’. The strange thing is that conventional neuroscience already answers this question – just in a way that clashes with intuition.

A brain contains several billion neurons whose structure allows each to receive signals from about 10,000 other neurons and to generate their own signal in response, again passed to another 10,000 neurons (Glynn, 2003). The brain is not so much talking to itself as members of a vast colony talking to each other. To generate complex meanings, signals have to function together, many neurobiologists taking the view that they combine more the way words do in sentences (Barlow, 1972) than pixels on a screen. As Buzsaki (2011) has emphasised, these complex meanings can only occur where signals operate together in a ‘reader’ or receiver structure. That structure is a neuron and the form in which the signals operate together is a pattern of electrical potentials generated at synapses laid out along branches of the neuron’s input dendrites – its ‘sensorium’ perhaps. It might be argued that the pattern does not actually matter – that the signals just add up somehow – but if not, we are left with combined meaning being manifest nowhere! Signal transmission also depends on action potentials in output axon fibres and on chemical neurotransmitters in synaptic clefts, but it is only as a pattern of synaptic potentials that the signals *act together*. The way potentials act together on the neuron’s dendritic tree is still not well understood because of complexity (Spruston, 2008). Some signals may combine piecemeal, branch by branch, but this is the one place in the brain where at least a good number of signals can be ‘experienced’ by a cell component to which they could have complex combined meaning. (As discussed later, recent studies (Aru et al., 2020) suggest that partition of different sorts of information to dendritic subzones may still lead to a single complex pattern-based computation.) The implication is that this is where human experiences occur and the search is on to find out what that component might be, in terms of up-to-date condensed matter physics.

What follows is an exploration of the case for the individual neuron approach. This has changed little in twenty years but

plausible options for mechanisms have become clearer (Edwards, 2005, 2006, 2008, 2013, 2020, 2021a, 2021b, 2023; Sevush 2006, 2012; Somov, 2012). Two complementary questions stand out: Why should brains, as wholes, *not* be conscious? And why does it make more sense for neurons to be conscious? Because many find the idea hard to access, we open with a dialogue between the authors – a biomedical scientist (JE) and a psychologist (PS), who both happen to consider themselves "neural colonies" – to try to clarify how the theory cashes out, in familiar cultural terms.

Dialogue

PS: The single neuron theory raises many questions but, before getting into details, many people will want to understand in broad terms how the theory can fit with the way we normally think about ourselves. Perhaps the bottom line is this: Is a neuron a what or a who?

JE: A neuron is a what, in that we can think of it as an object with size and shape. But the idea is that some neurons, maybe thousands or millions of them, must also each be a 'me', in the sense that *'there is something it is like to be in a world'* for a me. That may not quite make a neuron a who, or Jo Edwards, but maybe not so far off, at least *for how things seem to the neuron*. It is not a human being – that is the colony. A trickier question is what a person is.

My view is that a person is best seen as a self-story, told over a period of years, partly in real time, partly with memory re-plays, and future plans, by a colony of cells to a colony of cells. We know different cells play different roles in bringing in, collating, sifting and presenting the self-story – in eyes and ears and brains. Cells in arms and legs play a role in sensing space and so on. What people find harder to make sense of is that neuroscience tells us that the story is also being told to *many* cells, within the brain, presumably because making use of the story is a complicated business needing a team effort.

The philosopher Daniel Dennett (1991) originally proposed much the same in terms of a story being passed around the brain in "multiple drafts". Neuroscience can give a simple and specific account of that – each draft is the story as presented to a cell. Some cells may only get hum-drum versions, with rather few savouring the full sense of being me but judging by what we know of the cerebral cortex there are likely thousands or millions of these. However much we may have a sense of being a single me getting the story, neuroscience makes that very unlikely. It is often said of a brain 'there is no single place where everything comes together'. In that sense Descartes (1649) was wrong. But it makes no sense for there to be nowhere everything comes together because we experience it doing so. And there is no point in having signals that do not signal to anything. The answer must be that everything comes together in lots of places – as inputs to many neurons. Each enjoys a sense of being 'Jo' or 'Pavel'.

PS: "Some people may say that a "who" doesn't have to have a story to be seen as a person. Some might say that a "who" is just a capacity for experience, a locus of subjectivity, a space that is aware, an animate subject (as opposed to an inanimate object). In other words, a "who" then is just a living and sentient organism that might or might not have a self-story: a self-maintaining, self-serving, self-preserving, autopoietic living system that operates on a primitive "me/not-me," "self/other" duality. An amoeba would be an example of this kind of "who" without an apparent self-story of personhood. With all this in mind, would you be willing to say that a neuron is a subject? That a neuron has subjective experiences?

Would you be willing to say that a neuron is sentient? That a single neuron is aware? In other words, are you willing to experientially anthropomorphize a single neuron?"

JE: You have summarised the key idea. I doubt we even need to call a neuron an organism or worry if it is 'autopoietic'. I doubt that has to do with being a subject, to be honest. But yes, the neuron is what can listen to a story and is a 'who' in that sense. In theory, it might be moved from one colony to another and enjoy different stories. The person idea, on the other hand, I think needs its own story, and also implies uniqueness, and if there are neurons that are who's or me's we can be pretty sure there are at least thousands within one head.

PS: "You say that a person is an ongoing self-story being told by a colony of cells to a colony of cells. Some might say that a story of personhood, to qualify as a true story of self, has to be self-told, self-generated. Not just told to you but told by you. Is it ever the case, based on your understanding, that a neuron auto-generates a story of self? Or is the story of self always told to/given to a neuron by another neuron (or a group of neurons)?"

JE: I think this is where our concept of self is flawed. The arguments can get technical but are simple enough and I think important. If the telling is a process of the sort we generally understand, it will be an influence of some A on some B. We do not have any way of understanding a direct influence of an A on A. Certainly, an A can influence a B and that B can influence A later. An element of that may be relevant, but nerve cell As are connected to thousands of Bs, each of which is influenced by thousands more As so it is very unlikely that A will just get its own influence back. Moreover, there is nothing gained by something telling itself about itself. It would already know. And the story of me is not a story about a single nerve cell so, no, I cannot see how that can work. This also leads to the double-edged role of sequence in brain networks which I will come back to.

PS: "You seem to suggest that at any given time, inside your skull there may co-exist various kinds of self-stories of "Jo Edwards," some are experientially parochial/provincial, while others are more comprehensive and integrated. All these co-occurring self-stories are being experienced as a sense of "me, Jo." And all these self-stories are drafts, in that they are continuously revised and updated, correct? You also seem to suggest that there is a certain continuity to each story because a particular draft is correlated with a given neuron's place in the overall system and its inputs which are largely fixed/stable over time. If so, would it be then accurate to say that any given neuron has an evolving sense of self-story and personhood? Or is it the kind of situation where any given neuron might be presented with an entirely different self-story (in content and scope) from one moment to another?"

JE: Yes, I think under normal circumstances each neuron will enjoy a continuous and constantly updating story without dislocations. On the other hand, we know that sometimes in mental illness, or more simply in dreams, a self-story can become dislocated. In a dream we may seem to jump back in time to take some exams we sat twenty years before. A person suffering from severe depression may feel that they witness a story that is not theirs; they have depersonalisation. I am suggesting that what feels that here is each of many neurons. I strongly suspect that within a brain there is only one story at a time, although different groups of neurons may get it in different detail. The exception to this would be if there is major damage to connections, as in the people treated by corpus callosum division – so-called split-brain cases. Even

there, the two sides mostly seem to keep each other updated.

People might argue: If the brain is telling itself one coherent story, why can we not consider the brain as a single ‘consumer of’, or ‘listener to’ that story? This is where we need to get into the details of the argument against brain-wide consciousness.

Why Looking for Brain-wide Consciousness is a Wild Goose Chase

1. The Brain as ‘Hologram’

Ramon y Cajal's 1890s pictures of separate neurons² established for neuroscience the Neuron Doctrine, which effectively says all interesting events in brains occur separately in individual cells; everything else is branching telephone wires. In the mid-twentieth century, however, it was found that functions like memory (Lashley, 1950) and responses to odours (Freeman, 1987) could not be pin-pointed, and seemed ‘distributed’. Although never widely accepted, Pribram's (1991) idea of a holographic (or ‘holonomic’) brain was typical of a view that brain functions, including consciousness, are spread out as some sort of unified wavelike event. This was fuelled by electroencephalogram (EEG) traces, although these are not the sort of waves to do that. Pribram made use of the idea that in holograms *all input information can be found at every place sampled*. The irony is that the Neuron Doctrine, with each cell sending signals in branches to about 10,000 other cells and receiving about 10,000 inputs, means that a distributed function with all input information present at many sites is already dealt with. It already is a ‘hologram’.

There is only a problem with the Neuron Doctrine if we want all unifications of information in each cell to be *further* unified into one grand event. The idea that somehow binding must occur *between* neurons, in addition to signals meeting at dendrites, seems to be based on confusion of two sorts of unification (also discussed by Bayne, 2000). The so-called ‘phenomenal binding problem’ is the problem of how to get many signals, encoding the rich pattern of an experience, together. Enough information needs to be unified into each event of experience. That gets confused with the idea that all experience must be unified into a single event at any one time in a brain – a single ‘me’ event. We do not have evidence for the latter. We have an intuition that it is like that, but science often shows intuitions are wrong.

This confusion may have led to the idea that we should be wanting to bind together *output spikes* of neurons scattered across the cortex – which make up part of the EEG. However, these signals do not begin to have integrated meaning until they converge with other signals at receiving dendritic trees (Buzsaki, 2011). Moreover, if signals did bind across the cortex, we hit the flip side of the binding problem: the boundary problem. We know that, at any one time, our experiences do not include all the information our brains are processing. Vast amounts of visual, interoceptive and other stuff are processed subconsciously. Patterns of converging signals in dendritic trees will be selective. It is much harder to see how patterns of firing across the cortex can be selectively bound, even if there was *something for which they could be bound* which, as far as we know, there is not. There is no ‘person’ hovering over the brain-space watching EEG patterns dance. There is talk of things like ‘circuits’ but we shall see that doesn't work.

A lot of attention has been paid to synchronous oscillations in neuronal potentials and in particular gamma frequencies, that sometimes, but not always, correlate with conscious activity. Yet, as von der Malsburg (1994) pointed out when raising a role for synchrony in binding, the synchrony must operate where the signals arrive synchronously at dendrites. Elsewhere it has no causal significance.

In simple terms, why would brains have massively complicated selective connection pathways if computation could be done by some general buzzing effect? It makes no scientific sense. But it seems that it is hard for people to let go of the idea of some global ‘me-ness’.

2. Grey Matter Geometry is Unsited to Representation.

It is said to be hard to see how to get an image of a sunset out of a grey soggy lump. It sure is! Rather than being a regular array, the brain is like a termite nest of tubes criss-crossing in all directions between cells dotted about with no clear geometric regularity. Worse still, the nest is folded and puckered to fit the skull. The cell dendrites and axons are not in neat rows but all amongst each other like spaghetti in a tin. All signals in the brain are similar so we pretty much have to assume meaning is encoded in spatial relations. Spatial relations of neurons across the cortex are just too muddled to encode a precise visual image. People seem to forget. Some look to an information theory, which is independent of physical structure. Yet, in information terms, the answer is clear – the only places where information is combined are the individual dendritic trees. Within a dendritic tree, the structure is beautifully organised geometrically, with synapses laid out in regular arrays along linear branches.

3. Brain-wide Fields are not ‘Bound’.

Another idea that hopes to ‘unify’ signals across the cortex is that they form a unified electromagnetic (EM) field – even if, given the remarks above, this seems unpromising. The potentials that form the EM field might seem unified but *in themselves*, they are neither unified nor non-unified, neither together nor not together. If we are looking for unification of lots of signals to form a rich pattern, that ‘unification’ is only meaningful in a scientific theory if it is a unification of *an effect on something*. As William James (1893) pointed out, any other idea of unification has no effect on a theory’s predictions because it *doesn’t do anything*. In an embrace, potentials on two sides of one brain are no more unified than on the same side of two brains. Consciousness is not going to be a field. It is going to be *something being affected by a field*. The field is almost certainly EM, so the question is what is experiencing that field? We need some sort of ‘aerial’, but we want more than a radio or phone aerial, which just picks up changes in one potential over time. We want aerials that can sense spatial patterns – many potentials at a time. Neuronal dendrites are really just that, with neat geometry. But an aerial only picks up potentials in its immediate locality, in this case, one dendritic tree. And nothing at a larger scale in the brain comes near to being such a sensing device.

4. The Catch22 of Network Sequence

What makes a network of nerves a network is a relation of sequence between cells. Cell A sends to B and B to C. For re-

entrant circuits, C sends back to A – more sequence. But if you want a group of neurons to experience a pattern as a group, you want them informed in parallel, not in sequence. James Blackmon (2016) has pointed out that brain events occurring in sequence cannot also be fused into a single event. Being in sequence is our clearest way of defining *separate* events. Network, or sequential, relation is exactly the wrong thing for binding, which is why people keep looking for an extra ‘binding’ relation that cuts across connectivity because they know, deep down, that a network is no use. But if a group of nerves is going to be a single subject and it is not because they are a network in sequence what distinguishes nerves in one brain from those in another?

5. The Catch22 of Entanglement.

Some have hoped that quantum entanglement of signals might provide unification. There is another hitch, however. If two events are entangled in quantum theory, it means that whatever information you get from one, the information you get from the other is the same, or correlated. But the whole point of having many bits of information is that they are independent. In other words, if *signals*, such as those shown on EEG, were entangled, information value would be lost – a worse-than-useless situation.

There are, nevertheless, potential advantages in whatever *receives* a rich set of signals having internally correlated or entangled components. It could increase its energy content and power output. It could lock it into a geometric framework (as for bird brain navigation). At a quantum level, it can indeed support a *unified* rich distributed interaction in a well-defined domain of space and time. As indicated later, such correlation may be plausible in a neuron but there is no known receiver unit of this sort at a larger scale.

6. Multiple subjectivity was never a problem.

People ask why, if there is more than one ‘me’ subject in a brain, we have no sense of being multiple. We wouldn’t. The brain does not represent its method of representation. It has ways of representing its own timing but no means to represent other internal relations, with no rulers or ways to use them. Nothing would signify ‘another me over there’. An absence of any sense of multiplicity is no surprise. No subject is aware of the subjectivity of another in any other context, anyway. There is also no need to worry about ‘conflict’ if no cell is aware of what any other cell ‘decides’. All experiencing cells can get updated with a team decision and accept it is what ‘I’ decided.

7. The fallacy of the central control agent.

It may be assumed that, at whatever scale, there must be a single conscious ‘control agent’ in a mind, to prevent chaos. But from a logistic point of view, we almost certainly need *lots* of agents. The reason is that a control agent, if it means anything, is something that says ‘OK’ or ‘No way’, or, with some range, ‘up a bit’ or ‘way down’. When we ask human beings questions, there are all sorts of answers: ‘it’s grandmother’ or ‘too hot’ or ‘maybe next Thursday’. That requires team effort because one unit cannot produce all these outputs. As an analogy, get 160,000 Marines in black to stand in formation each with a white banner. The commander asks: “Who is president?”. Each marine has learnt a specific role

and up go banners in a pattern showing the face of Joe Biden to an overhead plane. Conflict does not arise, but no one control agent can do this. Each marine hears the question and says to himself 'For this one, I leave my banner down' or 'For this one, it's up'. The output is not just a majority vote, like in an election, it is something much more sophisticated.

How does that seem to a psychologist, Pavel?

PS: Ever since Freud, at least, I see psychology as aware of 'multiplicity' in mental agency. I think this fits together well. OK, so we bite the bullet and think of consciousness as belonging to individual cells. Exactly where do we locate consciousness in the neuron itself? What could be the information processing mechanism within a neuron?"

JE: In two sentences, I see consciousness as the relation of synaptic potentials to some global cytoskeletal structural unit in dendrites. The information processing is the quantum-level event of a field interacting with a coupled field excitation. I will try to make that intelligible.

Why Consciousness in Individual Neurons Makes More Sense

1. Function

The only places in the brain where rich combinations of signals are integrated, or form a unified causal influence, are in individual neuron dendritic trees. As von der Malsburg (1994), Triesman (1998) and Buzsaki (2011) imply in different ways, this must be where signals are 'bound' in terms of 'feature binding'. If a brain is to respond to a red square and a blue triangle and not a blue square and a red triangle it has to organise the way these signals *converge* onto computational events and the computational events are in individual neuron dendritic trees.

In other words, in functional terms, binding only occurs in individual dendritic trees. That is therefore the obvious place for subjective binding into events of experience should be. Subjective binding anywhere else would not do anything and so has no reason either to have evolved or to be possible to report. The more difficult question is what this binding is based on subjectively.

2. Could an Individual Cell Support a Whole Experience?

It might be surprising that an individual neuron could get enough information to sense a landscape or printed page. It might be the structure with the most complicated informational input in the body but even so, is it enough? In the 1970s neurobiologists such as Horace Barlow (1972) worked out the likely number of signals to expect our experiences to be based on. The calculations drew on how many sensory inputs were available and evidence from experiments about how much information we can retrieve from short-term memory. Barlow came up with the number of 1,000 signals. That might seem few, but some neurobiologists suggest even fewer. Barlow's 1,000 signal argument gives one cell the right sort of complexity. More than one cell starts to look too complex – certainly, more than 100 would be. Our experiences are not going to need more signals than are initially carried in an optic nerve and we know much of what comes in through optic

nerves is ignored at any one time. We are almost certainly not looking for a capacity of more than 50,000 channels – which is what pyramidal cells can have. At any one time, many neurons may integrate less than 1,000 signals, but if it is the pattern of active and null signals at 50,000 sites that matters, a wide range of meanings can be achieved with relatively few active signals.

If experiences did make use of input to many neurons, then there are two questions. Why put even more inputs together when in computational terms the input to one neuron is already surprisingly complex? How does it help to make these inputs ‘together’ if each cell can only respond to its own input and not that of others? If inputs to other cells produce crosstalk, we have a problem like that of entanglement – loss of independent variation.

3. The Reason for Going Quantum.

The reasons for making experiences quantum-level events need not have anything to do with obscure things like entanglement or wave functions. The main reason is that we need experience to be a direct interaction rather than indirect or sequential, mediated, interaction as in Blackmon’s (2016) argument above. We are looking for a single indivisible event rather than something done piecemeal. That is essentially what is meant by quantum-level events; they are not piecemeal or indirect, involving many steps. Any theory that deals with *individual indivisibles*, which human subjects seem to be, really has to operate at the level of such indivisibles in physics – the quantum level. Some of these indivisibles exist only within atoms but more general quantum field theory allows for individual dynamic units that ‘inhabit’ extended matter, which can be at the cellular scale. The simplest example is the phonon-based mode we call a sound wave.

4. What Event of Interaction Could Be an Experience?

Modern physics *individualises* events by making them relations between individual field excitations (photons, electron orbitals) and the universal fields to which they couple. In simple terms, such an event is the history of an individual ‘particle’ being influenced by forces – often called a quantum system, although usage is confusing. It might be the arising of an electron orbital through a relation between a free electron and the electromagnetic field around a nucleus. For an event of experience to be a truly individual *unified* event, as suggested above, it should belong to a field excitation experiencing a pattern of potentials to which it can couple or relate. In brains the potentials that provide information about the world are EM potentials – the signals at synapses that are integrated in dendritic trees. So, the ‘subject’ or ‘soul’ ought to be some mode of field excitation, associated with charge, that ‘inhabits’ the whole tree, or maybe some component structure extending throughout. Again, it is simplest to think of an appropriate excitation here as something like a sound wave.

The likely structures are membrane or cytoskeleton. A sitar may be a helpful analogy. The dendritic membrane is encrusted with knoblike ‘spines’ – much like tuning knobs. It is where the potentials arise but looks messy as a home for a mode of excitation. The microtubular threads of the cytoskeleton are, in contrast, very ordered, looking more likely to be inhabited by a mode that can be affected by all the potentials, just as the sitar’s internal sympathetic strings are affected by all the notes played as fingers move over the fretboard. But as for the sitar, the truly global excitation may be the acoustic

behaviour of the whole ‘instrument’, dependent on the bridge, the hollow body and the air within.

The excitations in the sitar – and maybe the mode in the dendritic tree – are acoustic modes or sound waves. In a structure like a human cell, with molecular dipoles, or the quartz crystal of a crystal radio, the arrangement of electrical charges in the substrate allows an acoustic mode to couple directly to an EM field. There are more complex alternatives, but this gives a general idea. The mode of excitation can, in a unified direct interaction, respond to a complex pattern of electrical potentials in time and space.

There is something more about acoustic and similar modes that may be relevant to our stream of consciousness. Acoustic modes can continue to exist when ‘un-played’. The sitar resting on the shelf has ‘ground state’ modes ready to sing. The crystal radio without its aerial has inactive piezoelectric quartz modes. These modes can both continue to exist and respond to what is going on – something that probably does not apply to simple modes like atomic orbitals which are either there or not. The acoustic mode can take up energy and give it out – and so could contribute to the cell firing – acting in response to its experience, just as the quartz crystal in the radio responds to radio waves by producing music.

Note that the sound wave, or equivalent, is proposed not as something heard but as *the hearer*: a, suitably ‘immaterial’, dynamic unit that Descartes or Leibniz would have called a ‘soul’. There is no homuncular regress; this is a ‘me’ and the only sort there ever was.

As of now, the existence of suitable modes of excitation in neuron dendrites, to ‘listen’ to the synaptic potentials is speculation, but work by Craddock (2017) suggests that long-range modes of the sort described may well occupy the cytoskeleton. All we need is one that can alter the cell’s electrical output to make the model fit together. Time will tell.

Experimental testing of theories of the neural correlations of consciousness is hard. The judgment on the Chalmers/Koch bet was that we have got nowhere in 25 years. However, in the last few years it may be fair to say that the evidence has shifted progressively towards events *within* individual cells rather than across networks. There is even a Dendritic Integration Theory of Consciousness, and a very interesting one (Aru et al., 2020). To quote: “consciousness relies on the dendritic integration of two anatomically and functionally segregated data streams” [within a neuron]. The theory proposes that patterns of signals in apical and basal dendrites of a layer 5 pyramidal neuron may underlie the comparisons that form the basis of predictive coding, or difference detection, theories of cognition, and that anaesthetics may work by decoupling the two. Aru et al. propose various ways to test their proposals empirically. It is now taken as realistic to talk of complex pattern-based computation within cells – building on classic work from Koch and Segev (2000). Aru and colleagues do not place consciousness within a neuron but they do place the rich functional operations they consider to be conscious there. Which, to us, begs a question.

Formal proof that events of experience occur in normal brains at one scale or another is likely to always be impossible or near impossible simply because the evidence has to depend on reports from a whole functioning system. Nevertheless, the sort of testing that Aru suggests promises to make it clear in such detail that the neural correlates are inside individual cells that the idea that conscious events are at some other scale, already known to make little physical sense, becomes perverse.

To summarise, putting events of conscious experience in individual cells makes much better sense than in brains for several reasons. We know what the signals are – EM potentials – and it would make sense if these could combine in a single event of influence on a quantised mode in the cell structure, something that looks plausible but is yet to be identified.

Interestingly, fundamental events in physics are now neither billiard balls colliding, nor binary logical events in computers. New field excitations come into being in the context of rich field patterns in a sense because that domain of the universe has an *appetite* for them, or *desires them to be*. Our model suggests that within a neuron there are events that are both fundamental interactions and models of the interaction of the whole organism with its environment. That allows 'appetite' at a fundamental level to take on a role as appetite of the whole 'person'. It also suggests that at root we are ruled by desire or emotion – only secondarily by reason.

Getting Used to the Idea

Detail of the biology and information theory is given elsewhere in cited papers. What we want to highlight here are broader implications, including how we get used to the idea, which has been far from trivial, even for those of us who have done so. We return to dialogue.

JE: Why do you think so many people find the single cell approach hard to grasp or accept?

PS: I think there are several key issues. First, it's the matter of size. It is counterintuitive to think of yourself as being the size of a neuron, of yourself (your "entire" sense of self) as being objectively microscopic. The size of the experience (of am-ness or self-ness) is in disproportion to the size of the experiencer – perhaps a kind of neurologically Napoleonic, narcissistic coefficient; we have to be a tad self-aggrandizing to survive the grandeur of the cosmos we find ourselves in. It is yet another kind of Copernican shock, another ego-dethroning.

Second, single cell/single neuron view of consciousness is a threat to the Illusion of Phenomenological Singularity of Being. For various historical, legalistic and cultural reasons, we are programmed to think of ourselves as singular, not plural, as an I, not a We, as an individual, not a commune. The single neuron theory of consciousness is anarchic. The notion that there is more than one of me inside of me directly threatens our sense of agency. If the kingdom or queendom is to run smoothly, there can be only one sovereign on the throne.

Third, we struggle to anthropomorphize a cell. In fact, we dare not. This Fear of Anthropomorphizing is the legacy of scientific/empirical reductionism.

Fourth, Paradigm Obedience. Science is business, another way of getting bananas. Paradigms are revolutions. Revolutions are resisted by the establishment of the jungle. We, modern-day apes who fancy themselves "human," are afraid to de-humanize ourselves by anthropomorphizing the very cells that make us human and humane.

JE: How might the single cell approach make it easier to understand each other socially?

PS: Seeing each other as a dynamic, fluid "we" would help us appreciate mood-specific, mode-specific, moment-specific ambassadorship of our communal manifestations. If/when you, the neural colony "Jo" and I, the neural colony "Pavel" meet again for a walk on Hampstead Heath, which of our many proxies, ambassadors, selves will show up? One of the up-and-coming psychological/therapeutic approaches of the day is known as the Internal Family Systems Therapy, which is a kind of neuro-revival of Freud's triune theory of self. It is broader, more humanistic, more neuro-scientifically up-to-date. But it is very "Multimind" (Ornstein, 2015), very Edwards, very Sevush, very "many," very "we." Seeing ourselves in all our parallel-processing, parallel-agenting complexity and reviewing our "theory of others' minds" along the same lines can leverage nothing but tolerance and nuanced understanding.

JE: That makes sense. I tend to be sceptical about us knowing enough about mind to apply theory to clinical practice, but you may be too, and we share the hope that a clearer idea of how brain cells work together may help provide sound grounding for care. 'Multimind' sounds a good ballpark. Many people like Minsky have considered mind as a 'society', but maybe not in terms of subjectivity. Individual neuronal consciousness sometimes gets confused with multiple personalities and that, of course, is wrong. But it is consistent with 'being in two minds' about things, a colony that can be swayed this way or that by members.

PS: Does a neuron have rights (as in human rights)?

JE: The ethics of the neuron model quite often come up. If there are many conscious agents in a brain where does responsibility lie? If many subjects can feel pain, do they have separate rights to relief? And so on. I see responsibilities and rights as concepts we can apply to the colony. What readers see as the implications for cells I think maybe is for them to decide!

PS: When you and I meet or communicate, who are we talking to/communicating with from Time 1 to Time 2? Do we end up engaging with different neurons from a moment to a moment? What accounts for seeming continuity of attitudinal and behavioural vectors?

JE: As I see it, each time we talk to the same whole colony and all members contribute a yes or no to the outcome. But at different times each cell contributes differently. They are part of a co-ordinated colony, so there is continuity, just as the Italians go on being the Italians.

PS: How have these ideas changed your own life?

JE: Chiefly in terms of the excitement of seeing a whole new aspect of understanding ourselves – the way electrical events might connect to meanings and ideas. If we know exactly where that happens, in our cells, we have a chance of making headway, even if it is far from easy.

PS: Does a neuron have a private sense of self, its own sense of "am-ness," "I-ness"? Or is "a sense of self" just another "story" that is input into a given neuron's dendritic tree?

JE: A good question. How much is what a neuron feels its interpretation of input and how much is the interpretation

already in the input? Science has no rules to use here. We have to find them. Some examples are probably answerable. A high-level 'me' cell will sense motion because it gets signals from cells that indicate there is motion, not because it is told of different positions at different times. I am pretty sure a 'me' cell senses it is the only me because it is told that by other cells. My cells no longer sense that. The story must be partly cultural, but I suspect it also comes from some pre-programmed role assignment in the network structure.

PS: Some might say that "self" is "an auto-biographical" memory, a kind of dataset about self. Does a single neuron have a memory (as in "autobiographical" memory)?

JE: Our brain cells know nothing about themselves as brain cells. They carry contributions to the colonial memory, but these can only be made use of through engaging the whole colony, so I see no private or individual neuronal memory in practice.

PS: Would it be accurate to think of all neurons inside our skulls as being on a continuum of autobiographical awareness, where some of the best-connected neurons are nearly totally informed about anything and everything that has been happening to a given anatomical body that they are respectively denizens of, whereas other neurons are only minimally autobiographically informed?

JE: Yes, I would expect cells to have a different range of detail about different aspects according to their positions and roles. Position and role may not be uniformly matched. Dehaene (2014) and Changeux have suggested cells with different 'role levels' may occur in different *proportions* in different areas rather than simply being segregated.

PS: Would it be accurate to metaphorically compare the dendritic tree to a keyboard and real-time conscious experience to a 'desktop' display, with cytoskeleton serving as input processor? What does a neuron output in this metaphor?

JE: Yes, something like that. The output is an action potential down the axon. A response in the cytoskeleton must facilitate this somehow, as the quartz radio crystal activates the speaker.

PS: In speaking to a "you", I speak to a crowd, a colony of neural selves unaware of not being the "only ones" receiving this question. If so, addressing all of the "you" – there – inside this corporeal vehicle – I am making the "you" aware that they have conscious neighbours across the synaptic abyss to left and right, above and below... So, what does it mean to inform all of these mini-you-s of the fact that you'll, that yinzz (as they say in Western PA), that you guys are not alone... Does this break up the monadic loneliness?

JE: I think it reminds us to heed Prospero; that we are such stuff as dreams are made on. It may help us understand that 'body' is of no importance. What matter are patterns in time – stories. I am not sure monads were ever lonely – in theory, they perceive the universe and progress in harmony! Maybe ours are so harmonious they seem like one.

Epilogue

There remains an intriguing question. Why is it that the above account of consciousness as a property of individual cells is to some a 'no-brainer' – the obvious way to look at mind – and for others it is at least six impossible things to believe

before breakfast? Are our minds built differently in this respect? Western culture has in the past told us there is one 'me'. Some Eastern cultures view mind very differently, but not as a multi-me either. Yet, according to William James (1893), the polyzoic multi-mind idea has had its adherents for centuries. All that we can perhaps say is that it makes sense to us.

There are clear signs that neuroscience is moving towards an understanding of just how complex events in individual neurons are. In terms of conscious events being in individual cells, the question of "why not?" gets harder and harder to ignore, even if people remain divided about credibility. To end with another quote from Aru et al. (2020) "Perhaps now due to the unprecedented advances in the techniques [] the time is really ripe for figuring out the neural basis of consciousness."

Footnotes

¹ With help from the late Michael Ellis Fisher, who in 2003 pointed out to JE that in modern physics long-range modes of oscillation in condensed matter, such as phonons, can be regarded as indivisible dynamic units.

² See e.g. <https://www.artsy.net/artwork/santiago-ramon-y-cajal-purkinje-cell-of-the-human-cerebellum>

References

- Aru, J., Suzuki, M. & Larkum, M. (2020). *Trends in Cognitive Sciences*, 24(10), 814-825.
- Barlow, H.B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, 1(4), 371-394.
- Bayne, T. (2010). *The Unity of Consciousness*. Oxford: Oxford University Press.
- Blackmon, J. (2016). Hemispherectomies and independently conscious brain regions. *Journal of Cognition and Neuroethics*, 3(4), 1–26.
- Buzsaki, G. (2011). *Rhythms of the Brain*. Oxford: Oxford University Press.
- Craddock, T.J.A., Kurian, P., Preto, J., Sahu, K., Hameroff, S.R., & Tuszynski, J.A. (2017). Anaesthetic alterations of collective terahertz oscillations in tubulin correlate with clinical potency. *Scientific Reports*, 7(1), 9877.
- Dehaene, S. (2014). *Consciousness and the Brain*. London: Viking Books.
- Dennett, D.C. (1991). *Consciousness Explained*. London: Penguin Books.
- Descartes, R. (1649). *Passions of the Soul*. Retrieved from <https://www.earlymoderntexts.com/assets/pdfs/descartes1649part1.pdf>
- Edwards, J.C. (2005). Is consciousness only a property of individual cells? *Journal of Consciousness Studies*, 12(4-5), 60-76.
- Edwards, J.C.W. (2006). *How Many People Are there in My Head, and in Hers?* Exeter: Imprint Academic.
- Edwards, J.C. (2013). EM fields and the meaning of meaning: Response to John Joe McFadden. *Journal of Consciousness Studies*, 20(9-10), 9-10.
- Edwards, J.C. (2020). Quantum-level experience in neural dendrites: An interpretation-neutral model. *Journal of*

Consciousness Studies, 27(11-12), 8-29.

- Edwards, J.C. (2021a). Proximal experience as an essential part of physics. *Journal of Consciousness Studies*, 28(3-4), 76-99.
- Edwards, J.C.W. (2021b). A framework for evolution and consciousness: Panpsychism without tears? *Journal of Consciousness Studies*, 28(11-12), 77-101.
- Edwards, J.C.W. (2023). Modern monads. *Qeios*. <https://doi.org/10.32388/W5X50X>
- Freeman, W.J. (1987). Simulation of chaotic EEG patterns with a dynamic model of the olfactory system. *Biological Cybernetics*, 56, 139–150.
- Glynn, I. (2003). *An Anatomy of Thought*. Oxford: Oxford University Press.
- Horgan, J. (2023). A 25-year-old bet about consciousness has finally been settled. *Scientific American*. Retrieved from <https://www.scientificamerican.com/article/a-25-year-old-bet-about-consciousness-has-finally-been-settled/>
- James, W. (1893). *Principles of Psychology*. New York: Holt & Co.
- Koch, C. and Segev, I. (2000) The role of single neurons in information processing. *Nature Neuroscience* 3, 1171-1177.
- Lashley, K. (1950). In search of the engram. *Society of Experimental Biology Symposium*, 4 454–482.
- Ornstein, R. (2015). *Multimind: A New Way of Looking at Human Behaviour*. San Jose, CA: Malor Books.
- Pribram, K. (1991). *Brain and Perception*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Sevush, S. (2006). Single neuron theory of consciousness. *Journal of Theoretical Biology*, 238(3), 704-725.
- Sevush, S. (2016). *The Single Neuron Theory*. London: Palgrave McMillan.
- Somov, P. (2012). *Neural We*. i-Catching Books.
- Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, 9(3), 206-21.
- Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 353, 1295-1306.
- von der Malsburg, C. (1994). The correlation theory of brain function. In E. Domany, J. L. van Hemmen, & K. Schulten (Eds.), *Models of Neural Networks II*. New York: Springer.