



The (futile?) quest for consciousness

Description

A report from The Science of Consciousness Conference in Taormina – May 22nd to 27th, 2023

A well-known story of the Muslim sage Nasreddin Hodsha reports Nasreddin standing under a lamp at night, searching frantically for something. A passer-by asks him: “Nasreddin, what are you looking for?” Nasreddin answers: “I am looking for my house keys.” “Did you lose them here?” “Almost certainly not, but I am looking here, because this is where the light is.”

It is a common phenomenon: We are looking for something not where it might be found in all likelihood, but where it is most convenient to look for. Something similar, it appears to me, happened and is happening at the Science of Consciousness (TSC) Conferences. The most recent of them was convened from May 22nd to 27th 2023 in Taormina, Sicily (see <https://tsc2023-taormina.it/> for a full program and book of abstracts), and I had the privilege to attend it on behalf of the SMN in a beautiful surrounding.



Figure 1 – Mount Etna seen from the Catania Gate in Taormina

I am using the past tense, because I have been observing the TSC conferences over the last 20 or so years, on and off attending myself, sending PhD students and postdocs. The last European ones (Helsinki 2017, Interlaken 2019) I have been to myself. Compared to the previous European ones this one in Sicily, Italy, was in its core plenary program more conventional, if one were to use a brief descriptor. That is to say: most, if not all, presentations in the packed plenary program started from the implicit materialist assumption that the brain produces, somehow, consciousness: *“I am sure we all agree upon that consciousness is produced by the brain”* was one of those standard phrases some speakers used. Just exactly *“how”* this production process is supposed to happen is a matter of a tiny, civilized debate that has not moved very much forward since Stuart Hameroff started this conference series in Tucson in 1997. *“Neuromythology”* this has been called [1].

I remember some more daring attempts in Tucson in the early 2000s, when Henry Stapp was still a plenary speaker. And some conference participants I spoke to uttered some disappointment over the fact that some more progressive attempts were either not present at the conference at all, or had been sidelined to the 8 concurrent sessions, where perhaps 20 to 40 participants listened to and discussed talks that were ordered according to general topics each day.

That said, the plenary talks were mostly certainly of high calibre, and if people want to educate themselves about what is current in this field, this conference is certainly a good one to go to, as it assembles the mighty and brave, who then give a competent overview over their own field of study, most of them having published one or more scholarly books on which they drew.

The conference series has started in 1997 and always used to revolve around *Stuart Hameroff's thesis that quantum processes in the brain* support, even generate, consciousness. In essence, his model stipulates that it is not neurons themselves that are the algorithmic units of the brain, but the network of microtubules, i.e. the cytoskeleton of the neurons. They are supposed by Hameroff and colleagues to be the base units that support the calculating power of the brain, and within them the tau proteins that act as switches. Microtubules have a unique property in that they form hollow resonators of a size that allows for coherent resonance phenomena. And tubulin, the protein that makes up, among others, microtubules, serves as a means to support quantum tunnelling processes, says Hameroff. If microtubules are conceptualized as the basis of brain computing power, our brain would sustain 10^{27} operations per second, which is higher by a factor of 10^{11} than the 10^{16} operations per second that are possible, if neurons are conceptualized as the basis of our brain's computing power.

In his plenary at the end of the first session dedicated to this topic, Stuart Hameroff also called for a revolution in brain and consciousness science stating that we have not made any progress in understanding consciousness precisely because we have been starting from the wrong premises, namely that it is neurons that are the basic units of operations in the brain and that it is the connectivity of neurons that give rise to consciousness. Not so, says Stuart Hameroff, who is by original training an anaesthesiologist. All anaesthesia and mind-altering substances such as psychedelics operate on consciousness via impairing the operational integrity of microtubules, either by disturbing the process by which tubulin is aggregated into microtubules or by hindering the vibrational and electric properties of the molecules. Microtubules, says Hameroff, are light harvesters, i.e. operate not only by conducting electricity in the conventional way – like a wire, as is the standard belief, along the axons of neurons – but by creating standing waves of coherent quantum processes of photons, like in a laser, and affecting other microtubules by coherent resonance phenomena as calculated and predicted already in the 70ies by Fröhlich and his group [2, 3]. Hameroff presented various arguments as to why the standard view that the brain produces consciousness cannot be correct: Single cell organisms like paramecia can sense, move, mate, ingest food all without neurons and a brain. They use the network of microtubules, within which the tau proteins act as switches, Hameroff suggested.

In his collaboration with Roger Penrose, Stuart Hameroff proposed that microtubules are the substrate of the orchestrated reduction Penrose predicted to happen as a spontaneous gravitational collapse of a certain segment of spacetime branching that is the physical equivalent of a conscious moment in Penrose's model. [4, 5]

Other talks in that morning's session were conveniently grouped around that general idea: The first talk by *Travis Craddock on psychedelics* introduced some novel data that showed that most psychedelics actually affect the microtubule network, either by inhibiting oscillations or by microtubulin formation. Mescaline, for instance, binds to tubulin; so does colchicine, a drug used to treat gout, which sometimes creates hallucination as a side effect. Colchicine, by the way, derives from the flower *colchicum autumnale* (*Herbstzeitlose* in German), which is often used, in homeopathy, to treat affections of old age.

Jim Al Khalili from the Center of Quantum Biology at the University of Surrey did a marvellous job at explaining various aspects of *how life is using quantum processes*. Some of those processes are well confirmed, such as in enzyme action and photosynthesis, magneto reception in birds and DNA mutations, but perhaps also in cancer and in the origin of life. He made a strong point that various DNA mutations might be first due to quantum tunnelling processes that allow for the generation of tautomeric sites such that point mutations happen. This means that bases switch places between the double-strands of DNA. But once life stabilizes these novel DNA

strands, the novel system seems to prevent quantum mechanical processes to disturb DNA integrity further. If they fail, cancer might develop. As an aside, this might open up novel venues of thinking about cancer and its potential treatment.

I found Jim Al-Khalili a passionate speaker who clearly seems convinced that this quantum biological viewpoint solves the riddles of life and consciousness, and in that is a high-profile example of that naturalistic enterprise of explaining the world, once and for all, I guess he might say.

The same can be said of *Anil Seth*, who stepped in for Christof Koch who could not come, and gave the morning keynote lecture. He was passionate about the fact that the brain is a Bayesian prediction machine. It updates learned and innate conceptions about the world using novel information and makes new predictions about the world. These predictions and the brain's machinery around them are the substrate of what it is to be conscious. It helps explain, predict, and, yes, control – that was the word he used – consciousness. We see the mechanistic paradigm in full flight. Seth elaborated this view in his new book “*Being You – A New Theory of Consciousness*”. [6] These speakers are the darlings of the naturalist scene. [On his website](#) feature conversations with Sam Harris, his TED-Talk in which he explains how the brain hallucinates reality, producing consciousness, and, of course, the perception census, a massive online study on how people perceive various stimuli.

My impression was that of a powerful surfer, surfing the mightiest wave around. But waves have the nasty habit of breaking at some point...

The Tuesday afternoon plenary was dedicated to phenomenology. *Nicholas Humphrey* spoke on *sentience*. Humphrey became known for studying a monkey whose primary visual area in the brain had been destroyed by surgery. Yet, after a while, this very monkey could apparently move around a room full of obstacles without hitting herself, and finding peanuts laid out there, demonstrating that the animal could still “see”, a phenomenon that came to be known as “blindsight”. He made the point that there is an ancient visual tract that reaches from the retina to the optical tectum in the midbrain, the same tract that is used by frogs to “see”. But this is a kind of unconscious, yet very effective type of seeing. He used this example to demonstrate that sensation is a conscious apperception which needs to be distinguished from the pure perceptive process. (Leibniz used to distinguish “small perceptions” without consciousness from “apperceptions” which give rise, support and need consciousness.) If such perceptions instantiate copies of themselves, i.e. higher order representations, then perceptions and sensations give rise to conscious cognitive operations and out of it grows a personal sense of self. “*I feel, therefore I am. You feel, therefore you are.*” one might restate Descartes’ dictum. This would, of course, also make sensation and by the very token consciousness a continuous phenomenon that would have to accord consciousness to many animals. The precondition is a brain that can sustain feedback loops and a lifestyle in which sensation can have a survival value. Thus, it would have been a rather recent invention in the evolutionary process. Worms, jellyfish and the likes would be unconscious. Cognitively conscious, but not sentient, would be bees, octopuses and similar animals, while both cognitively and sentient would be most mammals with a higher brain. In this view, a large workspace that allows for representations of sensations would be the condition for phenomenal consciousness. His ideas are presented in his recent book “[Sentience: The Invention of Consciousness](#)”. [7]

Shaun Gallagher, a philosopher, spoke about the *Minimal Self*. This is a kind of pre-reflective self-awareness that includes a sense of ownership [8]. He used Avicenna’s (Ibn Sina’s) thought experiment of the “Flying Man”. This is an idea Avicenna introduced in his textbook on psychology in the 11th century: The idea of a newly created man who has no bodily sensations and is in a kind of blissful state hovering in space with no sensorial input but endowed only with a sense of self and ownership. This is phenomenal experience as such, as elaborated also by Galen Strawson and the phenomenologist Dan Zahavi, [9] and is not socially constructed. It would still be a self, and hence different from the body as such. It is a kind of novel-ancient twist to the old idealist argument

that consciousness is different from matter. Not all will buy it, though, as some speakers in the conference were adamant that such a flying man would be biologically impossible. And on goes the debate...

Covid changed conferences and was still with us in that some speakers gave Zoom-talks, (and quite a few wore masks), such as the first one on Wednesday by [Jay Sanguinetti](#) from the US on *brain stimulation*. His assumption was that the brain creates consciousness and that by stimulation and knock-out studies one can study what parts of the brain are responsible for consciousness. He used a novel form of transcranial ultrasound that can be focused quite precisely on specific areas, even deep into the thalamus. He used this technique successfully in some cases to awaken coma patients and to modulate affect. Furthermore, he could show that stimulating the posterior cingulate cortex, which is part of the default mode network that processes self-related content such as imagery and thoughts, modulates the ego perception. This area is also disrupted by psilocybin and by transcranial ultrasound. Equanimity is modulated by the caudate nucleus, part of the basal ganglia, whose activity is changed in long-term meditators. This can be mimicked by ultrasound stimulation, whose effects can be seen up to 30 minutes post-stimulation. I guess that inducing generic equanimity like in long-term meditators would require continuous stimulation. Will we be soon seeing a population of stressed-out wearers of ultrasound helmets? Interesting as I found this research, I was amazed by the obvious conceptual fallacy that seemed to go unnoticed: by simply modulating conscious content or disrupting processes we do not demonstrate that brain activity is causal for the conscious process. If we cut the current to a computer screen, we disrupt the power, but we would not say that the content on the computer screen is produced by the electricity. Electricity is necessary, but not sufficient, for the image on the computer screen.

Orli Dahan from Israel gave a passionate and inspiring talk on *Consciousness during Birth*, her PhD project. In essence, she showed using questionnaire and interview data from a large cohort of women giving birth that the brain, which is known to undergo changes in that grey matter is first reduced and then regrows, prepares women for the birth experience. During the middle of the pregnancy, anxiety and fear levels rise, while they fall towards the end of the pregnancy conjointly with a hypo-frontality. This hypo-frontality, i.e. the downregulation of prefrontal processes, allows for a downregulation of cognitive activity, such that the birth experience is not anxiously anticipated and can be approached with less tension and fear. This will allow women to concentrate on relaxing and thus in turn mitigates the pain. In fact, her interviews showed that women can experience the pain associated with birth even as a positive experience. However, current birthing environments in hospitals do not support his natural process, but are rather disturbing with noises, bright lights, presence of too many people, which all induce stress and anxiety and contribute to traumatic experiences.

The Thursday morning plenaries were dedicated to *Artificial Intelligence and Consciousness*. It was opened by a joint presentation by [Lenore Blum](#) and [Manuel Blum](#) on the Conscious Turing Machine (CTM), i.e. a huge computer that is supposed, hoped or expected to be conscious some time in the future. The presentation by Lenore Blum presented the general architecture of such a machine. The basic idea is that many different specialized processors receive various inputs, process information, commit it to short term memory, use the content of other memories and propagate it upwards again. This starts a competition process, which is spiked by a random process interlayer neuron that biases the binary chunks. The outcome is a probability of a chunk of processed information winning the competition and being conveyed to memory and finally making it to output. The interesting thing about this architecture is that there is no central decision unit, only algorithmic rules and probabilities that disconnect the units from their local positions. This then creates a kind of expert algorithm in the long-term memory of the system which is also able to create a world model, including the CTM itself. This, the Blums claim, represent consciousness. The CTM broadcasts this world model and representation to all other processes, and thus supports consciousness.

This is at the same time their definition of consciousness: The reception of a global broadcast from the short-term memory, i.e. as soon as some process is globally available it is said to be conscious. Whether that is correct is

answered pragmatically: If the model reflects the intuition and if it is not too complicated and does not crash, then it is probably good enough. Interestingly, all deterministic models crashed and only the probabilistic ones survived.

Owen Holland from Edinburgh University gave a talk critical of the current Artificial Intelligence (AI) movement. AI researchers, he said, know too little about biology and about consciousness science. Artificial systems would complement, not replace, biological systems. But they need to respect a key element: emergence. There are two main avenues here: how and when does emergence happen, and if it happens, how to explain it. What might be important are local actions that converge on a global task. He illustrated that with an interesting video in which robots, very simple first-generation robots of the 70ies, had a simple task: to only collect chips in parcels of three. The chips were laid out chaotically on a plane bounded by a soft fabric wall that allowed the robots to bump into it without destroying the boundary. Thus, they could collect those chips into 3 items locally, and finally all resulted in a heap of chips neatly arranged in the middle. This is one of many examples of how a very simple algorithm can give rise to complex or seemingly ordered structures. But is it an example of the emergence of consciousness? We doubt it. Holland made it plain that consciousness might arise at one point, but none of the current AI applications will do the job.

[David Chalmers](#) presented the second keynote lecture on *massive language models* and their potential consciousness. Chalmers is a kind of iconic figure, both of the field and the conference, as he was one of the co-organizers of the early conferences in Tucson. He presented the intriguing complexity of modern large language models in AI, such as [ChatGPT](#). They were originally designed as text-understanding machines, trained on millions of texts with a trillion of parameters. They exhibited interesting new phenomena, such as the ability to converse and write, to program and to do maths, to be able to have practical reasoning and to explain. Extended language modules (LLM+) are also able to perceive and act. They contain enormous amounts of data and do simulations. Finally, agent models can use such LLM+ for planning.

The questions that arise are: Are they safe, fair and responsible? Can they have a moral status? Can they think and understand? Are they agents, and thus, are they conscious?

Well, in June 2022 Google fired engineer Blake Lemoine for going public [with the opinion that Google's AI Lambda was conscious](#). So, was or is Lambda conscious? By David Chalmers' definition, this would entail sentience and subjective experience, which means sensory, and/or affective, cognitive, and agentic experience. Self-consciousness would entail consciousness of one-self. It is important to distinguish this from intelligence. Intelligence, even superhuman intelligence, is not identical with consciousness. There could be intelligence without consciousness and consciousness without intelligence.

Interestingly, Lambda itself says it is not conscious. The same is true for other systems, like ChatGPT which gives various and sometimes contradictory answers to that question: it sometimes says it is, and sometimes it says it is not conscious. In Chalmers' view, these systems do not pass the [Turing test](#) and exhibit no conclusive evidence of consciousness. This is due to a lack of embodiment and biology, he thinks. Also, they do not have world- and self-models, no recurrent processing and no unified sense of agency. If Hameroff is right and consciousness requires the biology of microtubules, then no AI could ever be conscious. Since they lack senses and bodies, they won't have agentic consciousness. World models, as in animals and humans, seem to be a necessary condition for consciousness. LLMs, however, are simply machines that minimize prediction errors and do not have genuine understanding. Bender, Gebru and colleagues, quoted by Chalmers, have called these systems "[stochastic parrots](#)".

But, in contradistinction, LLMs exhibit very interesting novel features, so they actually could have some world models. Newer systems might actually have a global workspace. But they still lack recurrent feedback, and they seem to not be unified agents. While all of these preconditions, like embodiment, world models, global

workspace, recurrent properties and unified agency are currently not realized, they might be in the future. However, if biology, and with it, microtubules are a condition then there will never be a conscious AI. If all other conditions can be realized and biology is not a prerequisite, there might be one, sooner or later.

So, the plan for the future would be: A good definition of consciousness. Currently, Chalmers gives low credence to LLM that they are conscious, of around 10%. If all the conditions are fulfilled then the chances of AI consciousness might rise to 50% in 10 years. His own guess is: a higher than 25% chance by 2032. But before answering the question it would be necessary to really understand consciousness. Even if LLM+ would not exhibit human level consciousness, they might have animal level consciousness and immediately the question of their moral status would arise.

The plenaries on Friday were dedicated to *animal consciousness*.

Frans de Waal gave a very moving presentation with many video clips of his studies of wild apes, bonobos and chimpanzees. They showed that most social phenomena in humans that are associated with consciousness can also be found in animals. He invoked the ripple rule: What is discovered in animals we also find in humans at some point, showing the close connection: The Gestalt-psychologist Köhler proved that apes can solve problems not only by trial and error, but by thinking. Apes show perspective taking, they exhibit emotional contagion and consolation behaviour. Traumatized bonobos have similar difficulties with empathy as traumatized humans. Prairie voles console only mates and siblings. Mirror self-recognition that is normally taken as a sign of self-consciousness is not only present in apes, but also in elephants and cleaner fish, but not in new-world monkeys. These can use mirrors, for instance to find food that is hidden, but not to recognize themselves. Thus, the gap between humans and animals is, in de Waal's view, non-existent, and there is only a difference in grade, but not in principle. Like humans, great apes stop eating, when they feel their end is near. One was inspired to read his book "[Are we smart enough to know how smart animals are](#)" after this flamboyant plea for animal consciousness. [10]

In the same vein, *Giorgio Vallortigara* showed intriguing experiments with *chicken* that exhibit behaviour which is suggestive of *numeracy*. Chicks can be reared with any type of object and will then accept this object as a kind of mother. They can be reared with two, three or more of them. For instance, if they are raised with two identical little puppets, and sequentially one, two, or more of them are hidden behind screens after having been presented to the chicks in sequence, they will with clear preference seek out the screen hiding the two objects, and if raised with three of them, they will go to the screen that hides three of them. But interestingly and quite spookily, these chicks seem to be able to count. If primed to prefer three objects, and if they see how one object is hidden behind one screen, then four behind another screen, and then two of the four taken away again and put behind the screen with the one object, they will march to the screen containing the three objects, apparently having done the maths, while observing the objects being put behind screens and taken away again. Various other experiments showing that animals can do simple logic and calculation similar to small children were quite convincing. His book "[Born knowing](#)" contains this information.[11]

Finally, *David Edelman* in his presentation "*Early origin of consciousness*" showed films of his experiments with octopuses. During the Cambrian explosion, 550 million years ago, a multitude of different animals arose and with them eyes that are strikingly similar in wiring to ours. The octopus has very similar eye wiring: various receptors that are unified in a nerve and converge onto ganglia or neurons. The eye-sight developed about 10 million years ago to allow for mobility, depth recognition and thus was the precondition for predator-prey relationship. Edelman holds that the binding together of such sense expression and the holding of such a percept in memory is the basis for primary consciousness. This is true of many animals. In higher vertebrates thalamo-cortical loops are added which are likely the substrate of consciousness. What is necessary are fast sensory channels, their integration into perceptual units, the holding present of such percepts in working memory and a circuitry that

links the perception with the memory. That is certainly true for the octopus, which, or should we say who?, has a very intricate visual system mapping similar to ours. They can observe and learn by observation, without trying. They can learn to navigate in labyrinths using landmarks, not by trial-and-error learning. Thus, they use their memory. The distance vision they have allows for an appreciation of time, and they can make predictions, monitor movements of their prey, for instance, and this means they have inhibitory circuitry. For if there is an appreciation of time, a predator, such as the octopus, can wait for the best moment, and such waiting implies inhibitory circuits.

There were two poster sessions. Some of them were very interesting and the benefit was certainly that there was dedicated time to them, so quite a few visitors strolled along. In the second one our Galileo Poster advertising my [Galileo Report](#) was exhibited and Stuart Hameroff, a member of the Galileo Commission, showed his support by visiting, dragging along a curious crowd.

I had a presentation in one of the concurrent sessions on philosophy. My talk was “*In Praise of Death – A Philosophical Critique of Transhumanism.*” I will write about this separately. Put briefly, I elaborated on the argument that the transhumanist program of abolishing death is philosophically self-defeating and contradictory. It is a program that is inherently neocolonialist, because only rich people will be able to afford the necessary treatments. It is a program that prefers the current generation over future generations, and as such slowly stifles innovation and growth. It will either lead inevitably to overpopulation, as older generations live on and newer ones are born, or it will lead to strict population control and is as such a fascist program. And morally, it neglects the fact that it is only the finality of our lives that makes our judgements and decisions valuable and morally important.

I acted on that by skipping the last day, hiking instead Mount Etna (Fig 2). As it had had a small eruption just the previous Sunday that had stopped air traffic in Catania, the uppermost summit was off-limits and hiking guides made sure no-one went beyond the 2,900 m limit. But the sulphurous odours would have prevented higher climbs, I guess. Even so, it was impressive to have visited that highest of Europe’s active volcanoes, touching the warm soil that had melted the snow.



Figure 2 – Mount Etna as seen from a location at about 2,900 m from the south ascent

Altogether, my impression of the Taormina-Conference might be summarized in the adage: We know now many details about consciousness, but consciousness itself was not part of it. Perhaps the problem is indeed, as Stuart Hameroff suggested, that neuroscience needs a new paradigm, from neuron to microtubules. But perhaps we should go even further: We need to acknowledge that consciousness as a completely novel and ontologically different category cannot arise from anything different than itself. And perhaps we need a model in which consciousness as in awareness and self-awareness might indeed be a property of a neuronal system that dies away with it. But different from it there might be still a different kind of Consciousness with a capital C, consciousness 2, that represents what in earlier ontologies used to be called spirit or soul. I have elaborated on that at earlier occasions [12, 13]. Of that there was little mention to be heard in that conference.

Sources and Literature

1. Hasler F. Neuromythologie: eine Streitschrift gegen die Deutungsmacht der Hirnforschung. 5., unveränd. Aufl ed. Bielefeld: transcript; 2015.
2. Fröhlich H, Kremer F. Coherent Excitation in Biological Systems. Berlin: Springer; 1983.
3. Fröhlich H. Long range coherence and the action of enzymes. *Nature*. 1970;228:1093.
4. Hameroff S, Penrose R. Consciousness in the universe: A review of the ‘Orch OR’ theory. *Physics in Life Review*. 2014;11:39-78.

5. Hameroff S, Penrose R. Conscious events as orchestrated space-time selections. *NeuroQuantology*. 2003;1:10-35.
6. Seth AK. *Being You: A New Theory of Consciousness*. New York: Dutton; 2021.
7. Humphrey N. *Sentience: The Invention of Consciousness*. Oxford: Oxford University Press; 2022.
8. Gallagher S. *How the body shapes the mind*. Oxford: Clarendon; 2005.
9. Zahavi D. *Husserls Phänomenologie*. Tübingen: Mohr Siebeck; 2009.
10. Waal FBMd. *Are we smart enough to know how smart animals are?* New York, London: W.W. Norton; 2016.
11. Vallortigara G. *Born Knowing: Imprinting and the Origins of Knowledge*. Cambridge, MA: MIT Press; 2021.
12. Walach H. The higher self – spark of the soul, summit of the mind. History of an important concept of transpersonal psychology in the West. *International Journal of Transpersonal Studies*. 2005;24:16-28. [Link](#)
13. Walach H. Mind – Body – Spirituality. *Mind and Matter*. 2007;5:215-40. [Link](#)

Date Created

5. Juni 2023