The Quantum Computing Delusion

Dan V. Nicolau, Jr.[1]

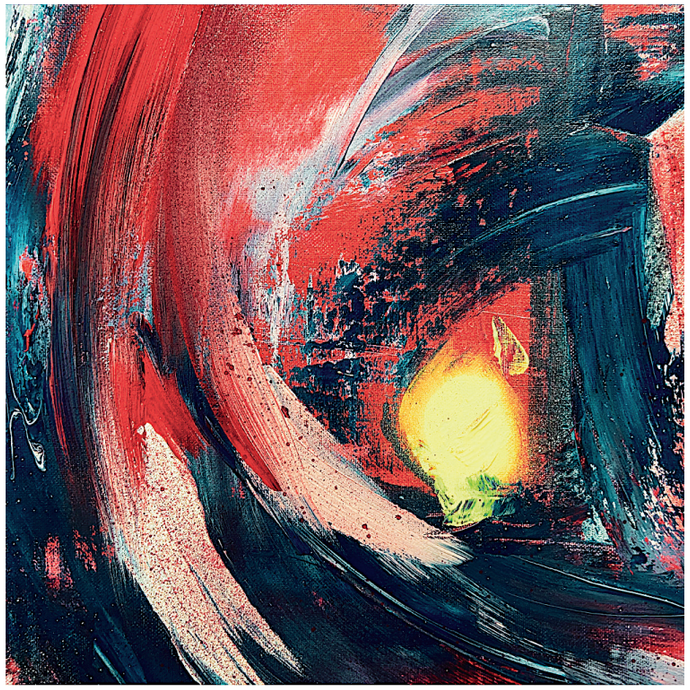

Fig. 1. Lauren Anne Marie Carter,[2] *Éblouie par la lumière quantique* (*Dazzled by the quantum light*). Acrylic on canvas, September 2020.

In the days before Christmas last year — it was a dark and rainy London evening, you know exactly what I mean — I grudgingly accepted an invitation to a house party. After a few glasses of mulled wine, I started talking to a friend of a friend who works at a tech firm that you've heard of. He's smart, young, educated and previously started and then sold his own successful startup. He told me about his work, which sounds interesting, and asked me about mine. I took my phone out and showed him a microscopic image, reprinted on the next page, of one of my lab experiments. It shows a handful of white "snakes" — fluorescent proteins extracted from rabbit muscle, about a millionth of a millimetre across — wandering aimlessly on a similarly protein decorated surface. Unseen to the microscopist, but highlighted in blue by the image analysis algorithm, are "railway tracks" guiding the nano-snakes' movements. I explained that the point of the experiment was to see if, by specially arranging these tracks, the mindless motor proteins could nonetheless generate simple arithmetic sums (2+5, 2+9, 5+9 and 2+5+9, in this case). We wanted to show, basically, that biological — but nonetheless lifeless, in the usual sense of the word — matter was capable of being harnessed to perform basic mathematical operations. My interlocutor, who was clearly intelligent but a layman in respect of bionanotechnology, mathematics or the dark depths of theoretical physics, suddenly became animated: "it's the future of quantum computing!"

What causes otherwise sane, educated, reasonable (even dull!) people to firmly hold and publicly extol such extraordinarily firm views on technologies and discoveries that they know virtually nothing about and that, in many cases, are not even (or not yet) real?

Observe that this is different in kind — not just in quantity — from the many irrational beliefs that we all hold because we grew up with them. When people believed the Earth was flat or that God made it in seven days, that was irrational and out of touch with basic observations about the world around them, but they had been told those things all their lives. We might find some people's beliefs of this type *strange*, *silly* or *stupid* but we do not usually think of them as *insane*. Holding beliefs about the future of a currently non-existent technology in an area of human endeavour that one is completely ignorant of feels different.

Although in common usage, the word 'delusional' is used pejoratively, in medicine it has a very specific meaning. A delusion is an irrational belief, firmly maintained in spite of compelling and voluminous evidence to the contrary *and* not shared by other members of the patient's culture. The second phrase is there precisely to delineate between the two kinds of irrational beliefs described above, to excuse, as it were, irrational beliefs that we grow up with or are surrounded by. A lot of French people appear, interestingly, to believe that the French nation is superior to all others, and if they look, they'll find plenty of (mostly French) people who agree with them, but if you think that you are Napoleon Bonaparte, or even that he is alive, that's a different situation.

The basic idea of quantum mechanics — that efforts at quantum computing are built on — is: to each possible state of a physical system, *i.e.* every way it *could* be, we can assign a number called an *amplitude*. Amplitudes are conceptually similar to probabilities, in the sense that a higher amplitude is associated with a higher likelihood of finding the world in *that* state. But they are different from amplitudes in that they can be negative and can even be imaginary numbers. This means that while probabilities can only add up, amplitudes can cancel out ("interfere destructively") as well.

A quantum computer is composed of qubits, which are like bits in a normal computer except that they can simultaneously be in both the 0 and 1 state, as with Schroedinger's cat being both alive and dead. So, a quantum computer with, say, 200 qubits (something in reach of near-future technology) could hold more "computational states" than the number of atoms in the known Universe.

What does this mean for computation? Regular computers, while very fast, are sequential — one operation at a time. That works very well for some things, like handling spreadsheets, playing chess or finding paths through cities using GPS: situations where we understand the rules of the game, in some sense. But it fails at many of the computational tasks humanity cares about, for example the discovery of new drugs, understanding the economy and planning under uncertainty and constraint. For these and myriad of other problems, the computer is forced, it would seem (though we can't yet formally prove it), to search for the proverbial needle in the haystack by checking each strand of hay, one at a time. For a drug "made" of a — modest, by biochemical standards — 200 molecular parts, each of which could be in one of two states, finding the needle might involve looking at a number of hay strands greater, once again, than the number of atoms in the Universe.

It's basically for this reason that so much hope is placed on quantum computing. In principle, since a quantum computer with N qubits can be in all 2^N states at the same time, using the drug discovery example above, it could search through the astronomically large solution space all in one go, finding the molecular configuration or configurations that successfully kill the pathogen (or whatever the aim of the drug is). Since virtually every area of human endeavour — and even logical reasoning itself — is in some sense founded on ultimately computational problems of this type, it follows, as I've argued before, that such number-crunching power would make us indistinguishable from demigods (at least as far as intelligence goes).

Of course, for a computer to be actually useful, it mustn't just compute: we also need to get the answer out. To do that, we need a way to turn the amplitudes of the quantum computer — once it has finished its work — into real probabilities corresponding to the answers of the original problem. The rules for how amplitudes of a quantum system (not just a computer) convert to probabilities are well known and are amongst the most fundamental laws of physics as we understand them. In the case of a quantum computer that 'naively' looked at all the strands of hay, pulling the probabilities out would simply give us a random answer, like listening to a bunch of people playing different tunes on different instruments all at once. Obviously, that would not be useful.

Of course, we can do better than that. We can try to exploit how amplitudes interact with one another by setting up something like an orchestra: a (hopefully) clever pattern of interference, so that most or all of the wrong answers to the problem cancel out (interfere destructively), while for the right answer or answers, the amplitudes all reinforce each other.

The challenge, naturally, is setting up such an orchestra that would work for all problems of interest and, crucially, *without* knowing what the right answer is in advance, or even if one exists! And, ay, there's the rub. In the tug-of-war between the enormous state-storing power of the qubits and the limitations necessarily imposed by any structured orchestration of their interactions, the quantum computer loses some of the "speed-up" conferred on it by its ability to store an exponentially large set of strands of hay. The question is, how much?

The answer to this question is not definitely known, and of course depends on the difficulty of the computational problem we're trying to solve, but all the theoretical results we have point to an answer that goes like: "*for most problems, the quantum computer loses most of its power.*" Although I've smudged a lot of technical details in the discussion above, a specific result, called Grover's Theorem, says that — at best — all that quantum computers (as currently conceived) can do when searching through a large, disorganised haystack (a database) is run squared-times faster than a regular computer. In other words, if a garden variety computer needs T seconds to do a job, the ideal quantum computer, if it existed, could do it in $\sqrt{T}$ seconds.

Is that a big deal? For easy computational problems, of the type regular computers can do quickly, we don't need quantum computers (or other alternative computing technologies), so the answer is 'no'. On the other hand, for the problems we believe to be fundamentally 'hard', requiring unreasonable (exponential) time on a regular computer, the quantum computer, in its ideal embodiment, would turn an exponential into a slightly smaller exponential. The image I have in mind when I think about that is of the authorities in Bangalore in the 1980s, who, faced with exponentially growing traffic, razed the city's glorious, green roundabouts to enlarge intersections. It bought them about 6 months.

In short, as far as we can tell, for most problems of interest, quantum computers would offer only a modest speed-up on their electronic counterparts. None of this is controversial. Grover's result is almost 30 years old and quantum computing research goes back to the 1980, so we have had plenty of time, brain power and money to think about ways around these limitations.

Friends of quantum computing point out that there are important computational problems for which quantum computers could make something currently intractable, tractable. In particular, they point to Internet cryptography, much of which is based on the presumed difficulty of the problem of factorising numbers into primes (*e.g.* finding out that 143 can be made by multiplying 11 and 13). This is indeed the case: a quantum computing 'orchestra' scheme called Shor's algorithm can rapidly find prime factors. If quantum computers can successfully overcome the formidable technological challenges they face (but see below), they may render much of existing Internet security systems vulnerable.

I can't resist the temptation of comparing that to someone defending a delusion of being constantly followed by the CIA by pointing to an occasional black van passing by.

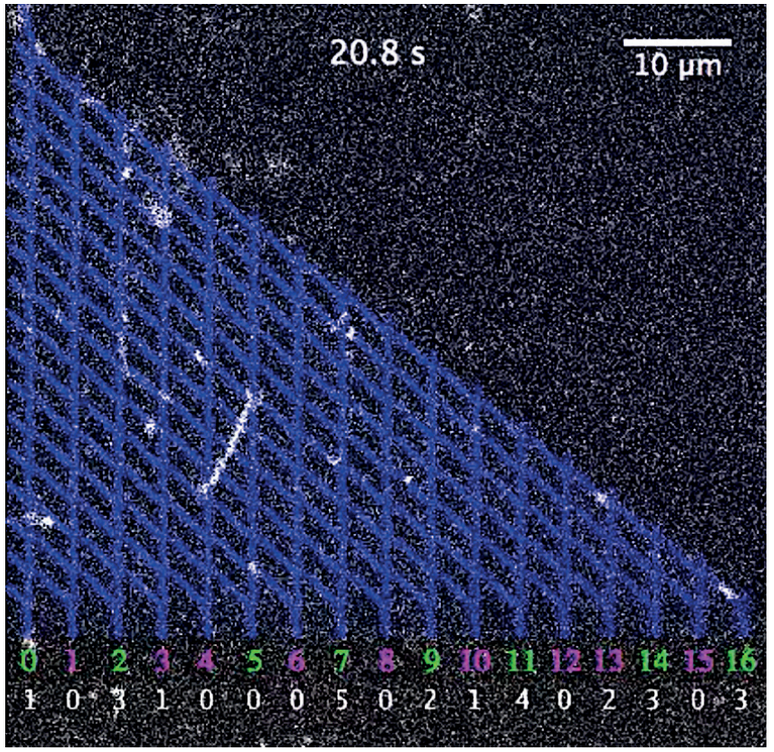

Fig. 2. Fluorescent proteins extracted from rabbit muscle, about a millionth of a millimetre across.

Firstly, the zeal of the quantum computing community for breaking Internet passwords is itself suspect. If a bunch of materials engineers told you about a cutting edge nanotechnology and, on questioning about its potential uses, all they could offer was that it could be used to easily smash every window in the world, most people would raise an eyebrow. Less pejoratively, the use of factorisation for online security is largely a historical accident, since factorisation is not believed to be in the class of 'exponentially hard' problems. We already have myriad security systems based on the purported difficulty of those harder problems, all of which could replace factorisation as the *de rigueur* security technology, rendering the speed advantage of quantum computers essentially insignificant.

Note that none of this has to do with questions of whether quantum computers can scale in practice. There are, to be sure, serious engineering challenges facing their development. These include, *inter alia*, 'decoherence,' increasing the density of signalling and wiring — which is hard to do without degrading the system's stability — and temperature control. There is also the quite real possibility, pointed out by Simon Levin, one of the fathers of computational complexity theory, that we may not understand the laws of quantum physics as well as we think we do, potentially building quantum computers on foundations of theoretical sand.

But that's not the point. If there was a reasonable expectation that quantum computers could, in principle, offer exponential speed-up over regular computers for even a few key computational problems of practical importance, the excitement around the technology would be more than justified and investment in overcoming those engineering problems would be warranted. Unfortunately, based on what we have learned over the past 40 years, that is simply not the case.

On the other hand, there are excellent reasons to continue to study quantum computers. For one thing, they are likely to lead to insights into the laws of quantum physics that we may be able to reach by other experimental means. And quantum computer-like systems could

still find an indispensable role in drug discovery, for instance by using quantum computers to simulate quantum molecular processes. Most excitingly, I think, the effort to build and scale quantum computers may teach us about the limits of computing and even, maybe, supply us with a profound new law of physics, which would say that there are some problems for which no computer, of any kind, can find answers in practice, in this world.

So, does the hype around quantum computers fit the medical definition of 'delusion'? Not exactly. For one thing, the *Diagnostic and Statistical Manual of Mental Disorders* tells us that a person cannot be diagnosed as being delusional if the belief in question is one "*ordinarily accepted by other members of the person's culture or subculture*." It's not clear how many believers are needed for a delusional belief at the individual level to escape from the "*folie à…*" diagnostic category, but a cursory Internet search for "quantum computing stock price" makes it clear that this hurdle, wherever it may lie, was passed long ago. When a large number of people come to believe irrational and probably false things based on hearsay, as I suspect is the case of my friend at the Christmas party, that's not considered to be a case of 'clinical' delusion by the psychiatric profession. Instead, we call it something more like 'mass hysteria', which, on my view, is much worse, by dint of lacking the originality and innocence that tend to characterise personal delusions.

[1] Mathematician, physician and engineer, Associate Professor of Computational Mathematics at the Queensland University of Technology and Visiting Professor of Experimental Medicine at the University of Oxford.
[2] Lauren Carter is a London-based painter.

The art of unconventional computing with cellular automata*

Genaro J. Martínez,[1] Andrew Adamatzky,[2] Marcin J. Schroeder[3]

The exploration of unconventional computing in its diverse forms is not only, and not primarily a result of the natural human pursuit for innovation but rather a response to challenges faced by the current information technology. Some of these challenges are not new, e.g. the expected end of applicability of Moores Law or the von Neumann bottleneck in the transfer of data between the CPU (central processing unit) and RAM (random-access memory). However, the bottleneck in the past was just a nuisance, but at present the need for massive processing of synaptic weights in the network for machine learning which requires multiple transfer makes this primary tool of AI (artificial intelligence) inefficient and hopeless in the competition with the natural, biological systems of information processing. An example of another challenge of a very different "down to the earth" type is the high energetic cost of machine learning estimated already as a substantial portion of the energetic needs of the industrial societies which in the near future is expected to become the main consumer of energy. Thus, the question about the frugality of nature in the energetic budget for the human or animal brain is worth billions or trillions of the future dollars.

These and other challenges direct the research towards unconventional forms of computing with the special interest in its natural forms identified in living organisms on the one hand, and in the utilization of new, natural, physical phenomena in information processing.

This brings us to a more general and theoretical question overarching the interests in natural forms of information processing about what constitutes the fundamental distinction between the traditional form of computing originating in the theoretical model of a Turing machine, and unconventional, natural computing. One of the possible answers is that the Turing machine model is based on the principle of a one-way, goal-oriented action initiated and controlled by a pre-defined program, while all natural processes are dynamic, *i.e.* they are based on mutual interactions within the processing system and with its environment.[4]

It is possible to consider a modification of the Turing machine model in which instead of the one-way action of the head on the tape the processing is performed by mutual reading and mutual re-writing of the two interacting central components.[5] This model of symmetric inductive machine remains within the Turing limit of computability as soon as the dynamics of interaction is computable, but nothing makes this computability unavoidable.[6]

The shift of the focus on the dynamic, interaction based forms of information processing can be implemented in the most natural way in the information processing in cellular automata where the art of unconventional computing begins. Unconventional and natural computing[7] has the capacity to handle information at an atomic and molecular level, the first stage. A diversity of scientific fields study and research all these ways on continuous and discrete domains. Lines of research can be found in Table 1.

Table 1: Some ways to unconventional computers.

Quantum computers[8]	**Reaction-diffusion computers[9]**
DNA computers[10]	**Hot ice computers[11]**
Physarum computers[12]	**Collider computers[13]**
Optical computers[14]	**Thermodynamic computers[15]**

In this way, the cellular automata theory conceived by von Neumann in the late 1950s years as a tool of super computation.[16] Von Neumann has been working with primitive and indivisible elements and where this theory offers an inherently and massively computation in parallel. He had discussed that universal Turing machines cannot exploit the process in