

Are we quantum computers, or merely clever robots?

Matthew P. A. Fisher
*Department of Physics,
University of California,
Santa Barbara, CA 93106, USA
mpaf@kitp.ucsb.edu*

*Talk presented at the Conference on 90 years of Quantum Mechanics,
NTU, Singapore on 24 January, 2017*

I would like to tell you about a new topic in “applied” quantum mechanics that has been my obsession for the past four years, summarized succinctly in my intentionally provocative title “Are we quantum computers, or merely clever robots?”

My story starts not with quantum mechanics, but with lithium—the 3rd simplest atom but also a psychiatric drug remarkable in its efficacy against serious mood disorders. Lithium—the Li^+ ion alone—is the drug of choice in “tempering mania and bipolar disorder”, employed by millions of people towards this end in the United States since 1974. I, myself, have taken lithium on and off for many years, and have always been fascinated by how a single element can be such an effective drug. Four years ago I set out with determination to uncover the biochemical mechanism behind lithium’s potency.

While there are many medications for fighting depression, bipolar disorder and other debilitating psychiatric disorders, all of these drugs are complex molecules with 10–20 atoms. So lithium’s simplicity was its attraction to me as a physicist. The joke that a physicist, when asked to understand a cow, would begin with the “spherical cow approximation” is, for lithium, no joke at all: Since the Li^+ ion has 2 electrons in the 1s atomic orbital it is “exactly” a sphere. For lithium the “spherical drug approximation” is NOT an approximation. I reasoned that if we cannot understand lithium’s biochemical mode of operation on cognition, there would be little hope to determine the mechanism behind the efficacy of any of the complex molecular psychiatric pharmaceuticals.

Lithium has two stable isotopes, lithium-7 and lithium-6, as I soon learned. The natural abundance of lithium is 92% Li-7 and 8% Li-6. So when one takes lithium as a pharmaceutical one is predominantly taking lithium-7.

I had an idea: Why not purchase isotope modified lithium, the two most readily available isotope compositions being 99% Li-7 and 95% Li-6, and feed these purifications to two groups of rats, with a 3rd group receiving pharmacy lithium (Li-N) and a 4th control group of rats fed no lithium at all, to see if there were any isotopic differential behavioral manifestations. At face value this sounds like a crazy idea, since it is the number of electrons of an atom/ion that determines its chemical properties. One would not expect that the number of neutrons in an atoms nucleus would play any significant role in biochemistry. On the other hand, this rat experiment did not sound so impossibly difficult. Fortuitously, I soon “discovered” that this very experiment had been performed some 30 years ago,¹ with truly remarkable results! This 1986 paper was responsible for my switch from quantum physics into a new (then non-existent) field of “quantum neuroscience”.

In the 1986 paper, the experimenters divided female rats into four groups and for ten days gave the groups either purified Li-7, purified Li-6, pharmacy Li-N or no lithium. After the rats were impregnated and during the 20-day gestation period the groups continued to receive these lithium dosages. After birth of the pups, the experimenters observed the mothering behavior of the different groups looking at various behaviors such as nest building, nursing, grooming of pups, grooming of self, state of alertness, and so on.

The results reported for the control group were all “average.” This is a little worrying; I don’t know whether this experiment was done “blind” (without knowledge that these were the control rats). Still, for the rats taking either pharmacy lithium (92% Li-7) or 99% purified Li-7, the experimenters reported that nest building was absent, nursing was infrequent/short duration, grooming of pups was infrequent and the state of alertness was “low”. For those of you who have either taken lithium or know someone that has, this is what might have been expected. When a person takes high doses of pharmacy lithium, their cognitive abilities become slowed down, with subdued behavior. Incredibly, for the mother rats fed Li-6 their nest building was reported to be excessive, nursing was very frequent/long duration, grooming and retrieval of pups was excessive—so these were the cleanest and safest rat pups in the world! State of alertness, instead of being “average” or “low”, was reported to be “very high”. Apparently, not only did Li-6 and Li-7 have different effects on the mood of the mother rats, but the effects had opposite sign.

What could conceivably account for this unexpected difference between the two lithium isotopes? While there is a mass difference between the two isotopes, when lithium ions are absorbed in water a very tight shell of four water molecules surrounds the ion with a 2nd shell containing roughly 12 water molecules. These hydration shells are much heavier than either lithium ion, so it seems very unlikely that the lithium isotope mass difference could account for the results of the rat experiment. But with different numbers of neutrons in their atomic nucleus the nuclear spin properties of the two lithium isotopes are very different, as described below. Taken together, these considerations raise the remarkable possibility that nuclear spin processing might be operational in the brain. If present, this processing

would be quantum processing since the nuclear spin is quantized. Might the brain have evolved to enable cognitive quantum processing? Are we, in fact, quantum computers?

There were two books published in 1989 that gave a bad precedent to those of us who want to suggest that quantum processing is operational in the brain. One, by Deepak Chopra, titled *Quantum Healing*, is viewed as pseudoscientific by the scientific establishment—the noted evolutionary biologist Richard Dawkins writes that Chopra uses “quantum jargon as plausible sounding hocus pocus”. Roger Penrose, who’s a serious mathematical physicist, wrote another book in 1989, *The Emperor’s New Mind*, where he argued that classical physics was inadequate to explain consciousness, and somehow the brain was harnessing quantum mechanics, perhaps via wavefunction collapse mediated by quantum gravity in intracellular microtubules. Patricia Churchland, who is a well-known neurophilosopher, writes that, “Pixie dust in the synapses is about as explanatorily powerful as quantum coherence in the microtubules.” That is pretty damning criticism. But what is more remarkable is Patricia Churchland’schutzpah, making such strong statements with little if any knowledge of quantum mechanics. But quantum physicists (myself included) were also duly skeptical, scoffing (and sometimes laughing) at Penrose’s arguments. And not without good reason.

Biological processes are slow, for example protein folding or neuron spike firing rates which occur at frequencies of order 10^4 Hz or 100 Hz, respectively. Converting these frequencies, f , to a temperature scale $T^* = hf/k_B$ where k_B is the Boltzmann constant, gives a “quantum-to-classical” crossover temperature, T^* , above which quantum effects are washed out and the process behaves classically. Only for $T < T^*$ will quantum coherent processes be manifest. For $f = 10^4$ Hz the temperature $T^* = 4 \times 10^{-7}$ K, many orders of magnitude below body temperature, $T_{\text{body}} = 310$ K. Biological processes should then be thermal; we are simply too “hot” for quantum coherence effects to be manifest.

But there is a loophole! The argument above assumes thermal equilibrium. In the worldwide efforts to build a quantum computer in the lab, the challenge is to protect the qubits from thermalizing with the environment; indeed, the qubit must be isolated from the thermal bath or quantum computation is impossible. Quantum computation requires isolation.

If there is quantum processing operational in the brain, it is going to require degrees of freedom (neural qubits) which are isolated. So we can ask, what degrees of freedom, if any, are isolated from the wet environment in biology? There is only one answer: nuclear spins. Protons and neutrons have nuclear spin $1/2$ and the nuclei of many atoms have a nuclear spin which can be $1/2$, 1 , $3/2$, and so on. For a given nucleus it is possible to use NMR to measure a nuclear spin decoherence time, the time that it takes for the nuclear spin to quantum entangle with its environment. For example, the decoherence time of the sodium nucleus when a sodium ion is floating in water is roughly $1/10$ second-long on microscopic timescales but not long on human timescales. These decoherence times vary between the elements.

For a Li-7 ion solvated in water the nuclear spin decoherence time is about 10 seconds. Remarkably a solvated Li-6 ion has a nuclear spin decoherence time of five minutes! That is a long time. Perhaps longer than my own aging memory. The Li-6 nucleus is very isolated, and it was Li-6 which was purported to amplify the cognitive processing of the rats. Putting these clues together convinced me that maybe I should take seriously the proposition that quantum processing with nuclear spins might be operative in the brain.

Some physicists think consciousness is too mysterious to be understood. Ed Witten, the famous String theorist, states that, “Biologists and perhaps physicists will understand much better how the brain works. But why something that we call consciousness goes with those workings, I think that will remain mysterious...”. Close-minded? As I like to joke: As physicists we know an enormous amount about next to nothing! I firmly believe that consciousness is a *real phenomena* and therefore *requires a scientific understanding* that we should seek to discern.

Other physicists believe that quantum mechanics can be ruled out as a requirement for consciousness. A good personal friend of mine, T. Senthil from MIT, is quoted as saying that, “The general assumption has been that of course there is no quantum information processing that’s possible in the brain. He (Fisher) makes the case that there’s precisely one loophole. So the next step is to see if that loophole can be closed.”

That is a nice way to frame things: Neural qubit isolation being the single loophole. Let me ask, then, can we close this quantum loophole? My approach was to make a list of some necessary, but surely not sufficient, conditions that have to be satisfied by biology and biochemistry, in order for evolution to have learned how to quantum process with nuclear spins cognitively. Foremost, you need a common biological element with a very isolated nuclear spin—to serve as a neural qubit. Moreover, you need a mechanism for quantum entangling pairs of nuclear spins, a mechanism for quantum processing with these nuclear spins and a method for quantum-biochemical transduction. Can these conditions be fulfilled in biology? To answer that question I had to learn some biology, chemistry, biochemistry, and neuroscience. It’s actually a lot of fun as a physicist to learn something.

My approach can perhaps be framed more intuitively (albeit less scientifically) if one imagines God coming down to earth and telling us, “Look, I’m clever, I’ve designed chemistry and biochemistry and evolution such that quantum processing with nuclear spins has become operational in your brain. Reverse engineer it”.

If you feel strongly that this simply cannot be true, maintaining that “I don’t believe quantum processing is possible in the brain at all” then you will surely be uninterested in any attempt to reverse engineering this non-existent phenomena. (Sorry about that.) But for me, to get off the ground, it was necessary that I suspend disbelief. Indeed, I am going to work under the assumption that neural quantum processing **is** present. And then seek to reverse engineer. If I’m not successful then probably my assumption is just not true. But if I am successful, well...

Let me briefly describe the progress I've made in this reverse engineering attempt.¹ Firstly, we need to identify a common biological element with a very isolated nuclear spin to serve as our "neural qubit", i.e. what we need, in effect, is an atom with a nuclear spin of $1/2$. Nuclear spins of 1 , $3/2$ or higher have electric quadrupole moments that couple very strongly to electric fields coming from the electric dipole moments of water and decohere rather rapidly (milli-seconds). However, a nuclear spin $1/2$ is only sensitive to magnetic fields. (Li-6, while having a nuclear spin of 1 , has an anomalously weak interaction with electric fields and is sometimes called an "honorary spin $1/2$ ") In solid-state physics we think of magnetic fields as being "weaker" than electric fields. This is also true for atoms/ions in solution since the fluctuating magnetic fields inside water (predominantly from the magnetic moments of the hydrogen nuclear spins) lead to much longer nuclear spin decoherence times for spin $1/2$ nuclei.

So which atom or atoms can serve as our putative neural qubit? I have always been intimidated by biology because there are so many molecules, but, fortunately, there aren't that many atoms that are of primary importance. The important biochemical elements—the elements which take part in the chemistry—all have a single letter: C, H, N, O, P, S. The bio-electrical ions, Na^+ , K^+ , Cl^- , Mg^{2+} and Ca^{2+} , are the other main players. (I should say that calcium is a very special ion in biology and in neuroscience in particular; it's sometimes called an electrical to chemical converter, and I should probably put calcium in both categories). So which of these biological elements have nuclear spin $1/2$? I didn't know, so I had to look it up.

Of course the proton (H^+) has nuclear spin $1/2$, but besides the proton *there is only one common biological element which has nuclear spin $1/2$* —phosphorus. Phosphorus has one stable isotope, phosphorus-31, and the ^{31}P nuclei has spin $1/2$. So ^{31}P is *the only possible neural qubit*. My first attempt to close the quantum loophole has thus been unsuccessful. As I like to joke, rather than closing the loophole, I went through the loophole!

And once through the loophole, one encounters another loophole: The requirement that the phosphorus atom be bonded in a very small molecule that can tumble rapidly to get motional narrowing and long decoherence times. Fortunately, in biology the phosphorus atom is essentially always bonded to four oxygen atoms in a tetrahedral arrangement that is called the phosphate ion, $(\text{PO}_4)^{3-}$. Having forgotten virtually all of what little biology I learned in high school, I did, nevertheless, remember the words adenosine-triphosphate (ATP), adenosine-diphosphate (ADP) and adenosine-monophosphate (AMP). These molecules consist of three, two or a single phosphate ion bonded to the organic molecule adenosine. Indeed, phosphate ions occur throughout all cells in all biological systems. Moreover, and most fortunately for my reverse engineering attempt, the predominant oxygen isotope has zero nuclear spin, thereby forming a protective cage for the central ^{31}P nuclear spin in the phosphate ion.

What is the ^{31}P nuclear spin decoherence time when the phosphate ion is free floating in water? Disappointingly only one second. This comparatively short

decoherence time (e.g. relative to five minutes for Li-6) is because at our body's $\text{pH} = 7$ there are free protons floating around that can bind to the negatively charged phosphate ion, and the proton and phosphorus nuclear spins interact with one another. The effect of the indirect magnetic dipole-dipole interaction on the ^{31}P nuclear spin “averages out” provided the phosphate ion is tumbling freely in water. But the “direct” nuclear spin coupling (mediated by electrons via the hyperfine interaction) is $\text{SU}(2)$ invariant and does not. While this exchange interaction is quite small (a frequency of 100Hz or so) it does entangle the phosphorus and proton nuclear spins, and when the proton jumps on and off the phosphate ion (in chemical equilibrium) leads to decoherence of the ^{31}P nuclear spin with a disappointingly short (1 second) time. And that is not very good news for my reverse engineering effort.

But wait. Might there not be another biological cation (positively charged) that can, in some situations, out-compete the proton in binding to the negatively charged phosphate ion? The Na^+ , K^+ , Mg^{2+} and Ca^{2+} ions being the promising candidates. As I learned from the cellular biology books Na^+ and K^+ do not take part in any chemistry and while magnesium can bind to phosphate ions this only occurs in very basic conditions ($\text{pH} = 10\text{--}11$) rarely present in biology. But, at $\text{pH} = 7.4$ present in the mammalian extracellular fluid, calcium ions can outcompete the proton in binding to phosphate ions—rescuing my reverse engineering effort.

The story is quite interesting. In 1975 A.S. Posner while examining the X-ray crystal structure of bone mineral (the crystal hydroxyapatite) noticed that within the unit cell were two “structural units” with chemical formula $\text{Ca}_9(\text{PO}_4)_6$ (subsequently named “Posner clusters”). Although Posner suggested that these clusters might be important in bone growth, they were not believed to be stable as free floating molecules in solution. But 35 years later in 2010, a cryogenic electron microscopy study of in vitro bone growth identified one-nanometer molecules free-floating in simulated body fluid—and proposed these to be Posner-clusters.⁴ If indeed stable in solution, such clusters should rightly be re-named “Posner *molecules*”—the name I have since adopted.

Since both calcium and oxygen have zero nuclear spin, it is only the six ^{31}P nuclear spins that are present in Posner molecules. As we shall see, Posner molecules can then provide an almost ideal setting for quantum information storage and processing of ^{31}P nuclear spins.

How long is the ^{31}P nuclear spin decoherence time when a single Posner molecule is tumbling in water? The dominant decoherence mechanism comes from the spatially and temporally fluctuating magnetic dipolar fields from the protons in the nearby water molecules. But these magnetic fields are fluctuating very rapidly so one would expect a “motional narrowing on steroids”. A very rough estimate gives a spin coherence times of a million seconds. Maybe it's only a thousand seconds, maybe it's ten thousand, maybe it's 10^8 seconds. But this spin decoherence time should be accessible with ^{31}P NMR and such experiments are currently on the agenda.

Another critical requirement for the feasibility of neural nuclear spin quantum processing is a bio-chemical mechanism for quantum entangling two phosphorus nuclear spins. How might this happen? When ATP is catalyzed to AMP (which releases an electron volt of energy) two of the three phosphate ions are cleaved from ATP. These two bonded phosphate ions, which share a central oxygen, is essentially a phosphate dimer—usually called pyrophosphate. The two ^{31}P nuclear spins in pyrophosphate are analogs of the two proton spins in the hydrogen molecule H_2 . And as in H_2 these two spins can be entangled in either a total spin singlet (like the para state of H_2) or a total spin triplet (like the ortho state of H_2). In either case the two ^{31}P nuclear spins will be quantum entangled.

Moreover, there is another biochemical reaction in which pyrophosphate is cleaved into two separate phosphate ions. And one expects the two ^{31}P nuclear spins to remain entangled even when separated, thereby creating non-local quantum entanglement — Einstein’s “spooky action at a distance”. Thus, whenever energy is required within (or without) the cell a flurry of such chemical reactions would release swarms of free, but pairwise entangled, phosphate ions. Presumably this is going on throughout your brain. Or not. But it is going on throughout my brain, since I can’t stop thinking about it!

Once liberated and free floating the ^{31}P nuclear spins in these phosphate ions will decohere in a second. But if ample calcium is present (as in the extracellular fluid) the phosphate ions can be enveloped forming well protected clouds of Posner molecules, $\text{Ca}_9(\text{PO}_4)_6$. Moreover, the ^{31}P nuclear spins within the Posner molecules would inherit the pair-wise entanglement, thereby creating inter-molecular nuclear spin entanglement between multiple Posner molecules. (Such inter-molecular nuclear spin entanglement might be most helpful in re-vamping the defunct efforts to make a liquid state NMR quantum computer.) And once “safely” ensconced inside the Posner molecules the ^{31}P nuclear spins can remain coherent for very long times (perhaps 10^6 seconds).

As I am frequently asked, even if this is actually going on inside the brain and plays some cognitive role, wouldn’t one’s thinking be deleteriously affected by applying big magnetic fields to one’s head? After all, nuclear spins are affected by magnetic fields—they start precessing. Perhaps you have heard about avian magnetoneavigation, where birds apparently measure both the strength and direction of the Earth’s magnetic field to help navigate. Well, it turns out that we humans can also detect magnetic fields with our head, although much larger fields in the Tesla range. If you have ever been wheeled into the bore of a large magnet (several Tesla field) for an MRI scan, you might remember that this is done slowly. Otherwise, if moved in quickly, you will feel dizzy. But (hopefully) only if the magnet is turned on! Despite not *being* birdbrains, we can, as birds do, measure magnetic fields with our brains.

Would these magnetic fields not cause decoherence of the ^{31}P nuclear spins inside the Posner molecules in our heads? No, not quantum decoherence, since the applied magnetic fields are classical. But, our nuclear spins would precess—

and maybe make us feel dizzy? By contrast, if free radicals–molecules with a free electron spin–are close to a (not so rapidly rotating) Posner molecule, the magnetic dipolar interactions between the electron and ^{31}P nuclear spins could cause them both to spin flip–thereby quantum entangling and decohering the phosphorus spins (and, erasing quantum information). Perhaps this is one of the reasons why “free radicals” are so bad for our health.

Since magnetic fields will inevitably cause some deleterious decoherence of any quantum information carried by the ^{31}P nuclear spins–even when bonded in Posner molecules–one might ask whether there are ways to “hide” or “embed” some quantum information within the $2^6 = 64$ nuclear spin states in the Posner molecule. In fact, there is. Quantum chemistry calculations reveal that Posner molecules have a 3-fold rotational symmetry axis so that the 64 nuclear spins states can each be chosen to acquire a phase factor $\exp(i2\pi\tau/3)$ under a 120 degree rotation.⁵ The parameter τ , which can take on one of three values ($\tau = 0, \pm 1$), is a “pseudospin” (qutrit, not qubit) that can be associated with each Posner molecule. And this “pseudospin” is very isolated from the environment, with potentially extremely long (days, weeks, months, . . .) decoherence times.

Moreover, because the phosphorus nuclei are indistinguishable Fermions this pseudospin dictates the allowed rotational angular momentum states of the Posner molecule, $L = \tau + 3 \times (\text{integer})$, in units of \hbar . The most general spin/rotation wave function for a Posner molecule is then a linear combination of these three pseudospin sectors. And remarkably for the two sectors with $\tau = \pm 1$ the allowed rotational angular momentum states, not coming in $\pm L$ pairs, do not admit a non-rotating (real) wavefunction–*one cannot make a wavefunction that is not rotating*.

Why is that important? As I have already mentioned, in bone crystal there are two adjacent Posner molecules within each unit cell. Moreover, quantum chemistry calculations show that two Posner molecules when brought together (in vacuum) can chemically bind to one another.⁵ Once bound the two Posner molecules (A and B, say) will clearly have no relative rotational motion, which implies that they can only bind if/when their respective pseudospins have opposite signs, $\tau_A = -\tau_B$. Remarkably, Posner-pair bonding implements a projective quantum measurement onto a “pseudospin-singlet”, $\tau_A + \tau_B = 0$.

If two chemically bonded Posner molecules subsequently unbind and drift apart, they will remain *entangled in a pseudospin-singlet!* This would be ideal for possible pseudospin quantum processing. Imagine a network of pseudospin entangled Posner molecules. The pseudospin entanglement, induced originally by Posner pair binding, will modify the subsequent pair binding rates even between Posner molecules that have never been bonded before. By way of example, consider two pairs of pseudospin entangled Posner molecules (A and A' with $\tau_A + \tau_{A'} = 0$) and another entangled pair (B and B' with $\tau_B + \tau_{B'} = 0$) with A and B well separated from A' and B'. Due to the pseudospin quantum entanglement, if A and B are subsequently brought together the probability that they will chemically bind together will depend on whether or not A' and B' have bonded together. The joint binding probability is

larger than the probability that each pair binds separately. Effectively, non-local pseudospin quantum entanglement can modulate quantum correlations between the pairwise binding of multiple Posner molecules, even when the pairs are well separated—perhaps in different neurons!

While pseudospin entanglement could in principle enable complex quantum processing with multiple Posner molecules, this processing will be for nought unless there is a mechanism for quantum-to-biochemical transduction. Fortunately, there is. When two Posner molecules are bonded together their collective rotation will be much slower than the rotations of the separate unbonded molecules. And with slower rotations it will be easier (in a slightly acidic environment) for the bonded pair to melt than for individual Posner molecules to melt, thereby disassociating the Posner-pair into their calcium and phosphate ionic constituents.

Quantum to biochemical transduction can occur when these released calcium ions trigger biochemical processes in the brain. In fact, after calcium ions flow into a presynaptic neuron (after an action potential reaches the synaptic region) and bind to the small vesicles that store neurotransmitters, a bio-chemical cascade will be triggered that eventually releases the neurotransmitters into the synaptic cleft. These neurotransmitters can diffuse across the synaptic cleft, bind to neuroreceptors on the postsynaptic neuron, and trigger the postsynaptic neuron to fire. Thus, the release of calcium ions when a bonded pair of Posner molecules disassociates provides an ideal mechanism for transducing pseudospin entanglement into neuronal firing.

Due to the spatial non-locality that correlates the pair-wise disassociation of Posner molecules, the quantum entangled web of pseudospins can induce non-local modulations of firing across the familiar “classical” network of inter-connected neurons. This corresponds to a quantum to classical “readout” for the putative hybrid quantum/classical computers within our brains!

How to determine if any of this is actually operational in the brain? After all, the story I have uncovered consists of many, many “pieces” which must precisely fit together to make a complete “picture”—much as in a child’s picture puzzle. Multiple experiments will be required to determine the precise shapes of each piece, and to see if they neatly fit together. This will surely be easier *in vitro* (that is, in test tubes) than in real brains (*in vivo*).

Hopefully, if quantum processing is actually operational in the brain, it should be possible to co-opt the biological “machinery” to duplicate the processing in a test tube—by analogy with genetic engineering. The recipe for the simplest “synthetic quantum brain” would be as follows: Fill a test tube with simulated body fluid (minus the phosphate ions), add pyrophosphate, add the enzyme pyrophosphatase to create free phosphates that can bind with calcium ions to form pseudospin entangled Posner molecules, add calcium-indicators (molecules that fluoresce when calcium ions bind), pour half of the fluid into a second test-tube, drop a bit of hydrochloric acid into each test tube to encourage the Posner molecules to disassociate, and detect the emitted fluorescence from the calcium-indicators.

M. P. A. Fisher

If any *cross-correlations are detected* in the light emitted from the two test tubes it would be truly revolutionary, indicating the presence of nonlocal quantum entanglement between chemical reactions, perhaps paving the way for a liquid state quantum computer and for a quantum theory of the brain.

Let me finish with a quote from Isaac Newton: “I can calculate the motion of the heavenly bodies, but not the madness of men”. Now almost 300 years since Newton’s death, can we not do better?

References

1. “Quantum Cognition: The possibility of processing with nuclear spins in the brain”, M. P. A. Fisher, *Annals of Physics* **362**, 593 (2015).
2. “Aberrant Parenting and Delayed Offspring Development in Rats Exposed to Lithium”, J. A. Sechzer *et al.*, *Biological Psychiatry* **21**, 1258 (1986).
3. “Atomic structure of amorphous calcium phosphate deposits”, F. Betts *et al.*, *Proc. Nat. Acad. Sci.* **72**, 2088 (1975).
4. “The role of prenucleation clusters in surface-induced calcium phosphate crystallization”, A. Dey *et al.*, *Nature Materials* **9**, 1010 (2010).
5. To be published, M. Swift, M. P. A. Fisher and C. G. Van de Walle (2017).