

How Many Neurons Does it Take to Make a Mind?

Braden A. W. Brinkman
University of Illinois at Urbana-Champaign
(Dated: December 13, 2009)

Abstract

The human race has made great strides in understanding the world around us, but the understanding of why it is we understand eludes us! We know that thought originates in the brain, we know that the brain is a network of billions of cells called neurons, we know that neurons transmit electric signals between one another and we even know *how* they transmit these signals, but we still do not know how such features combine to form *a mind*. Understanding how the mind, or consciousness, emerges from a network of neurons has become one of the great problems of modern science.

In this essay I will tell a story (not *the* story!) of our attempts to understand our own minds. I begin our exposition, as physicists tend to do, with the basic “fundamental” element of our system, the neuron. We then increase our cast and investigate how a network of neurons computes and learns, and finally, I end our story with the tale of a most peculiar hypothesis, that a mere network of neurons is not enough - quantum mechanics must play a role in how the brain computes...

I. AN INTRODUCTION

A foolish, or perhaps wise, person once asked, “How many grains of sand does it take to make a pile?”. The most sensible answer would seem to be there is no definite answer. A “pile” is not a well-defined entity: we might say of piles of sand what Associate Justice Potter Stewart once said of a certain controversial kind of entertainment: “perhaps I could never succeed in intelligibly [defining it]. But I know it when I see it...”. Hence the question is in some sense unintelligible, but it underscores an important point: a pile of sand is different from a grain of sand. There are features of a “pile” which distinguish it from a grain - features that no single grain possesses but arise due to the cooperative interactions between the grains. For example, when piles get too large, layers of sand may suddenly fall to the bottom of the pile in an avalanche; similarly, flat layers of sand vibrated vertically may develop localized, mobile, oscillating peaks. It is as if there is some conspiracy being orchestrated by the sand grains!

We need not stop here. We may ask what happens when we treat our collective systems as the individual pieces of an even larger collective system. If we add piles of sand together we get a larger pile of sand - that is, perhaps, not so interesting, so instead we take a great leap and turn our attention to a much more interesting (and complicated) system - *the brain*. The brain is a complex network of neurons, which are themselves complex systems composed of proteins, ions and lipids arranged in structures that enable signal propagation and computation. One of the - if not *the* - most striking emergent features of a network of neurons is the mind. How does a collection of neurons give rise to something like consciousness? Can we even understand this, or even intelligibly define what consciousness is, for that matter?

Whether or not these questions can be answered, we can only know by trying. Much effort has already been put into understanding the brain, so I can only offer a glimpse of it here. I begin this glimpse with the brain’s building blocks, the neurons - what we know about them, how they work - and build them up into the networks that our brains consist of. This will be the primary focus of our attention - how do neural networks conspire to exhibit emergent behaviour like memory, computation, or learning, and how might they give rise to consciousness? Are networks enough, or are we missing an essential ingredient to the puzzle? We end our glimpse with the tale of a controversial hypothesis, that quantum mechanics plays a part in the story of how our minds emerge.

II. A NEURON

A. The Real Deal

Although many kinds of cells play a role in the functioning of the brain, the basic computational unit is the neuron (Fig. 1), a cell that specializes in information transfer. Our first order approximation of the cell consists of four main parts: the cell body, dendrites, the axon and a synapse.

The cell body of a neuron is much like the cell body of most other cells, consisting of

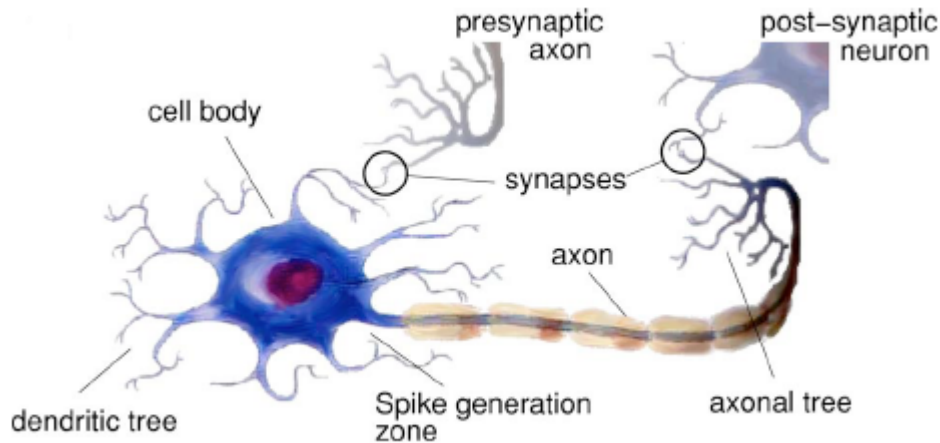


FIG. 1: Simplified picture of a neuron, from Ref. 1

the nucleus and other sub-cellular structures. The other three pieces of our simplified cell are unique to neurons. The axon is a long “cable” along which the electrical signals that constitute our thoughts travel. The axon extends from the cell body and ends in the synapse, which is the junction between the transmitting and receiving neurons. There are two types of neurons, chemical and electrical, both with the job of transmitting signals, but for different purposes and by different means. It is important to note that the “strength” of information transfer is not constant, which has led to suspicion that understanding changes in synaptic properties may be key to understanding learning and memory. Finally, dendrites are “trees” that grow from the cell body and branch into smaller projections which receive signals from the synapses of other neurons.

The flow of information through a neuron is a one-way street - it flows into the dendrites to the cell body and if the sum of all incoming signals is strong enough, the neuron generates a signal that is sent through the axon to the synapse and onto the next cell². Signals are generated as spikes in voltage that propagate along the axon. Information is encoded in the rate at which spikes are fired, thus much effort has been put into understanding how the timing of different spike trains can encode various kinds of information, such as in the ocular nervous pathways¹.

There is of course far more to neurons than we have space to discuss; what I have given is the basic, most essential properties. Learning further details is left as an exercise to the reader. The information provided in this subsection has been adapted from Ref. 3, which should be consulted for such further details.

B. Mathematical Models

The first step to testing our understanding of neurons is to devise a mathematical model of them and see if it reproduces observed behaviours. How we devise the model will depend on what level of description we want. The obvious thing to do is model the neuron as exactly as

we can to mimic actual cell behaviour. The price we pay for realism is a computational one: a real neuron has hundreds or thousands of dendrites, which makes computing their properties quite a challenge! Though this obstacle can be overcome, our ultimate goal is to understand how neurons behave in a network, and running a simulation of a network of a realistic number of these realistic neurons is more likely to contribute to global warming and the heat death of our computers than our understanding of the brain. To overcome this obstacle we must adopt the tried-and-true physicist strategy of stripping away the complicating pieces of our system of interest while leaving (what we hope are) the essential features intact. Reducing the neuron to just a few compartments we are still able to observe many of the interesting features of real neurons. The Hodgkin-Huxley model is a one-compartment model for signal propagation along the axon, and its agreement with measurements has earned it the title of one of the most beautiful applications of physics to biology⁴.

If we want to know how neurons compute, more conceptual, abstract models that do not depend strongly on the actual biophysical details are of great use; often one treats the neuron almost entirely as a black box that receives an input signal and generates an output signal to be sent on to the next neuron. This lets us focus on the connectivity of the neurons without getting lost in the details of the biophysical mechanisms behind signal propagation.

A more detailed summary of the various levels of neuron models is found in Ref. 4. As explained therein, to really understand how to tackle the problem of understanding the brain, it is likely we will have to understand how the various levels of modelling interact with one another. This is an important step for our understanding of the brain, but it is not enough. Though the computational properties of single neurons are interesting in themselves they are not enough to understand consciousness - it is not thought to be a coincidence that the smarter inhabitants of our planet tend to have more brain cells! Of course, it is not simply the number of neurons in our brains that matters, but how they are connected - i.e., the important question to ask is not “How many neurons does it take to make a mind”, but rather, “How does a network of neurons make a mind?”⁶.

III. A NETWORK

Trying to understand the properties of interconnected neurons by a detailed description of neurons would be about as useful as trying to understand water by solving the Schrodinger equation; accordingly, anyone who wants to understand the collective properties of neurons in networks will take the abstract approach. Abstracted neurons in a network are treated as nodes given labels to indicate their present state, and the abstracted connections between the nodes are lines with labels to indicate the strength and direction of information flow between two neurons.

Of the first models of neural networks was developed by John Hopfield, with the goal of modeling a possible mechanism for memory storage and recall⁵. In his model each neuron (out of a total of N) is assigned a state V_i which is either V_i^0 or V_i^1 , typically taken to be 0 and 1, respectively. A state of 0 represents a non-firing neuron whereas a state of 1 represents a neuron firing at the maximum firing rate. The strength of the connection from a neuron

j to i is denoted T_{ij} . The total input signal into the neuron i is then

$$\mathcal{H}_i = \sum_j T_{ij} V_j + I_i, \quad (1)$$

where I_i is a possible external input to the neuron, taken to be 0 here. In real neurons it is known that a neuron will only fire an output signal if the total signal input to it is greater than some threshold voltage. If we take the threshold for each neuron to be U_i , then the state of each neuron is given according to the rule

$$V_i = \begin{cases} 1, & \mathcal{H}_i > U_i \\ 0, & \mathcal{H}_i < U_i. \end{cases}$$

The state of the network is given as a list $\{V\}$ of the state of each neuron. The initial state of the network is initial randomized, and the state of each neuron is updated randomly and independently with a mean attempt rate W . The independence of updates results in the neurons firing out of sync, just as observed in real neuron networks.

To model information storage and recall in the network we define a set of desired “memories” (i.e., states) $\{V^s\}$ for $s = 1, \dots, n < N$. The memories are stored by means of the strength of connections between the neurons. If we choose

$$T_{ij} = \sum_s (2V_i^s - 1)(2V_j^s - 1), \quad (2)$$

with $T_{ii} \equiv 0$, then the memory states are fixed points of the system. That is, if we initialize the system in one of the states V^s and update the system according to the above algorithm, the state will not flow away from the initial state. If we start the system with a state nearby one of the memory states, we expect the network to eventually flow to the nearby memory state. Of course, like my father with names, memory recall isn’t perfect. The network will sometimes evolve to a state that isn’t quite the correct one. Numerical results found that approximately $n = 0.15N$ simultaneous memories could be stored before errors in recall became significant. (I am unsure if my father’s difficulties in recall are due to an increased number of memories over time or a decreased number of neurons!). Further difficulties with recall arose due to memories that were too similar and often became merged, and many memory states are only metastable.

The memory recall/computation procedure exhibited by this model is due to the collective behaviour of the neurons in the network, driven by the nonlinear interactions of the input/output relationships between the neurons - no single neuron in the network could have produced the same results. Because the result of the algorithm proves to be only weakly dependent on the details of the model, adding further neurological details should retain the basic behaviour of the model⁵.

This model is a schematic picture of how a network might perform memory storage and recall that arise from collective properties. A physical interpretation is that memories can be encoded by adjusting the connection strength between neurons. Should we expect this model to be anything more than schematic? Could it describe a real neural network? The

answer is a definite “maybe”. To see why consider the reason why the Hopfield net evolved to a stable state instead of wandering around all possible states for eternity. The reason is that there is an energy (Lyapunov) function

$$E = -\frac{1}{2} \sum_{i,j} T_{ij} V_i V_j - \sum_i (I_i - U_i) V_i,$$

which is (at least on average) negative for any change in V_i , thus E tends towards a minimum as the system updates. Any physicist who has taken a course in statistical mechanics should immediately recognize this as the Ising model Hamiltonian repackaged and repainted for marketing to biophysicists, with $V_i = \sigma_i$, $T_{ij} = J_{ij}$, $I_i - U_i = h_i$, and modified so that each neuron exhibits its own local “field”, h_i , instead of an overall global one. The appearance of the Ising model is not a coincidence, and is instead due in part to our treating neurons as two-state systems in the network. The deep reason for this is the Ising model is the maximum-entropy (in the information theoretic sense) model describing pairwise correlated binary bits - which is just what the neurons are in the Hopfield net. In this information-theoretic view one measures (or hopes to) the probability distribution of the network state $\{\sigma\}$

$$P(\sigma_1, \dots, \sigma_N) = Z^{-1} \exp \left[\sum_i h_i \sigma_i + \frac{1}{2} \sum_{i \neq j} J_{ij} \sigma_i \sigma_j \right], \quad (3)$$

with the parameters h_i and J_{ij} chosen so that the averages $\langle \sigma_i \rangle$ and $\langle \sigma_i \sigma_j \rangle$ match those measured experimentally, with Z^{-1} the usual normalization factor⁷.

How well does treating the network as an Ising model work? The data in Fig. (2) appear in Ref. 7 comparing this model to measurements on a small network of $N = 10$ neurons. The figure of merit is the ratio of the information entropy differences $r = (S_1 - S_2)/(S_1 - S_N)$, where S_k is the information entropy assuming there exist k -wise correlations among neurons. If this ratio is close to 1 it implies the pairwise correlations capture most of the behaviour of the network, which is indeed seen in the data.

Measurements of 250 (sub-)networks of 10 neurons yield an average ratio of $\langle r \rangle \approx 0.9$, indicating the pairwise correlated neurons reproduced most of the measurements made on the network; hence, the need for further correlations is of limited necessity. It would thus seem that the Ising model does a fairly accurate job of capturing the behaviour (or at least the correlations) of an actual network of neurons - at least for small numbers of neurons.

What about larger networks? Do pairwise interactions still capture the essential behaviour of the network as we increase the number of neurons? This sounds suspiciously like the thermodynamic limit in thermodynamics, so one might hope that our understanding of the Ising model in this limit will carry over to our understanding of neural networks. For weak couplings J_{ij} , which are observed to be independent of N , the “multi-information” $I_N = S_1 - S_N$ is proportional to $N(N - 1)$. As seen in Fig. (3), this is experimentally seen to hold up to $N = 15$. The I_N for independent neurons behaves similarly, and there is a crossover in extrapolations of the curves at $N \approx 200$; this implies that $I_N = S_1$ and this point, and hence S_N , the information entropy of the N correlated system, is zero. This

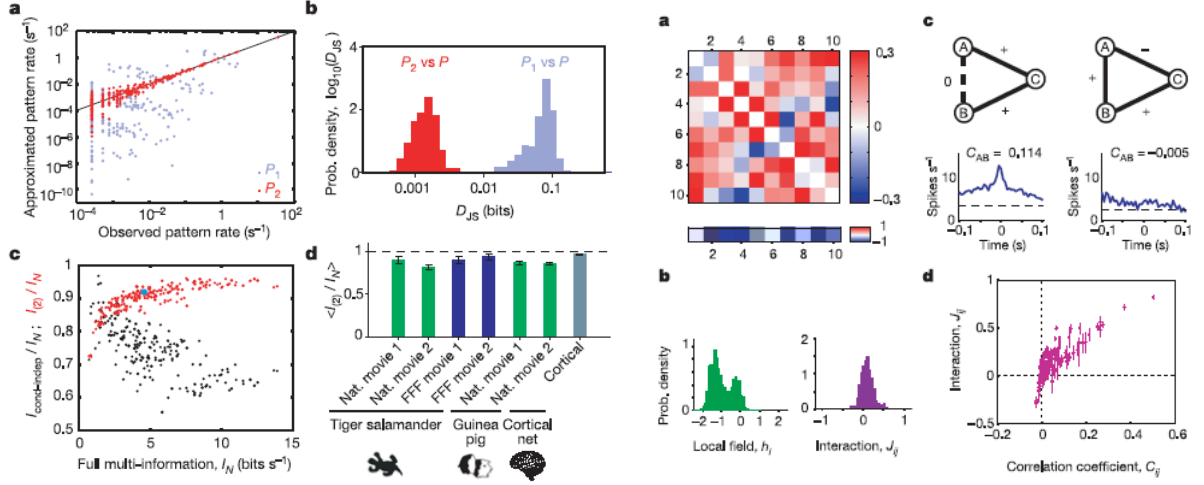


FIG. 2: Left: Data testing the pairwise model's applicability. a) shows the model predictions (red dots) match well with the observed signal pulse rates (equality shown by solid black line). b) compares the pairwise model with the independent neuron model, and c) plots the ratio r for each model. d) plots the ratio r for networks from different sources; Right: Testing pairwise correlations. a) plots values of J_{ij} and h_i beneath, with the corresponding histogram of values in b). c) indicates that strong interactions with a mutual neuron can correlate weakly interacting neurons while frustrate interactions can lead to low correlation. d) shows J_{ij} versus the cross-correlation coefficient C_{ij} . The figures with original captions can be found in Ref. 7.

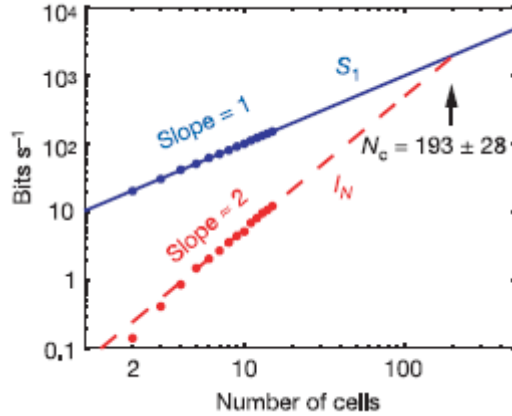


FIG. 3: Testing large N behaviour. From Ref. 7.

cannot actually occur, otherwise there could be no variation in firing patterns of the cells, while variation is still observed experimentally. However, collections of roughly about 175 ganglia in retinas do appear to be the largest size of correlated cell clusters, indicating

the $N_c \approx 200$ limit may be a genuine critical size of networks above which the behaviour qualitatively changes.

So, it would appear that the Ising model (and perhaps the Hopfield model, by extension) has considerable promise in describing the dynamics of actual neural networks, simplifying a potentially N -wise correlated system to a pairwise correlated system! Or has it? The careful reader will notice that the experiments described so far have been for relatively small numbers of neurons, about 15 at most. Anything more than that is an extrapolation or not a direct measurement of the probability distribution of the correlations. As diligent scientists we ought ask if the behaviours observed for small numbers of cells really do extrapolate to large numbers. The authors of Ref. 8 argue that the analysis presented above work well in a perturbative regime in which the experiments were done, but extrapolations beyond that regime are meaningless. So, the potential reduction in the complexity modeling neural networks, while appropriate for relatively small systems, is still unproven for larger networks. The next chapter on pairwise neural networks is yet to be written - current technological capabilities limit us to measuring the state of only about 100 neurons at a time, far from the 100 billion or so neurons actually present in the brain. Until we can do experiments in these non-perturbative regimes and are able to compare them with the predictions of the pairwise models we cannot be able to say if they really capture the essential properties of neural networks.

Thus far I have focused on modelling memory and recall in networks and presented the experimental results supporting the use of Ising models for networks, and the emergent properties that arise due to these models and algorithms. There is of course much more to the modeling of the brain - I have said nothing of models of learning rules, which model how synaptic connection strengths change depending on how often one neuron causes another neuron to fire. This is known as Hebb's rule: the connections between neurons which fire together tend to grow stronger, while connections between neurons that do not fire together tend to weaken. Learning is associated with "training" the network such that certain neural pathways are re-enforced. This should ring some bells - in the Hopfield model memories were encoded in the synaptic connection strengths, an idea very much in the spirit of Hebb's rule. Models of learning ("synaptic plasticity") typically introduce a system of differential equations which relate the changes in the strength of synaptic couplings to changes in the firing rate through that synapse³.

Another example of collective phenomena in the brain is the synchronized firing of large numbers of neurons. Synchronization is important in vision, movement, memory and epilepsy⁹, so how might one model it? The classic model synchronized oscillators is the Kuramoto model. The Kuramoto model describes a large collection of coupled oscillators, a large fraction of which may spontaneously synchronize if the coupling between them is sufficiently strong. This situation sounds reminiscent of how neurons in a network might behave, and indeed the Kuramoto model has been used to study the synchronization of neurons^{9,10}. The model is described by the differential equation

$$\dot{\theta}_i = \omega_i(t) + \frac{1}{N} \sum_{j=1}^N K_{ij}(t) \sin(\theta_j - \theta_i),$$

where θ_i is the phase of the i th oscillator, ω_i is the natural frequency of the oscillator (the natural firing rate for neurons) and K_{ij} is the coupling between different oscillators. In the usual Kuramoto model the frequency and coupling are fixed in time, but to incorporate learning by changing synaptic couplings one might expect that firing rates and synaptic strengths will change over time, and so this time dependence is included in the model when applied to neurons. This generalization will hopefully provide a more realistic model of synchronization in neural networks, and could in principle be coupled with a model of learning which gives rules for how $K_{ij}(t)$ should change in time.

I shall spend no more space discussing the collective properties of neural networks (the examples I have provided are by no means exhaustive!), not for lack of interest but because there are more interesting questions to spend the remaining space on! The attentive reader will notice that while I have given a few examples of collective phenomena that (purportedly) describe neural activity, none of these phenomena have that hallmark sense of “consciousness” associated with them. They may all be very important to the eventual emergence of a mind from our collection of neurons, but nowhere in these phenomena do we presently see even hints of consciousness just around the corner. If we were to manage to write down simplified models of various neurological functions such as memory, learning, synchronization, and so on, and bundle them into a “Standard model of Theoretical Neurobiology”, should we expect to get consciousness out? If not, what ingredients are we missing? I will spend my remaining space discussing a particular hypothesis: the missing ingredient is quantum mechanics.

IV. A QUANTUM HYPOTHESIS

One would naïvely expect quantum mechanics to play no role in the brain. Our experience with macroscopic quantum effects tells us that temperatures must be very, *very* low (much lower than body temperature!), and sufficiently isolated from the environment to sustain coherent quantum mechanical states. How could the brain ever maintain such a state at body temperature and given its ceaseless interaction with its environment? Most scientists would expect it couldn't. Indeed, the speculation that quantum mechanics is responsible in some way for consciousness often seems to stem not from well reasoned arguments exposing the fallacy of the line of thinking in my previous sentences, but from arguments along the lines of “Consciousness is weird and hard to understand. Quantum mechanics is weird and hard to understand. They must be related.” This is, of course, an insufficient reason to suspect a quantum mechanical origin of consciousness. There are less erroneous conjectures, such as the supposition that if living beings truly possess free will, then we need a non-deterministic theory to describe computation in the brain, and quantum mechanics fits the bill. Some go further than even this, arguing that many aspects of the brain are “non-computational” in the Gödel-Turing sense, and so something beyond even quantum mechanics must come into play, and must be somehow inherently “non-computational”. In particular, Roger Penrose suggests that a true theory of quantum gravity coupled to structures in the brain is required to understand consciousness¹¹. This sounds like too speculative and too fantastical a claim to be taken seriously, but to Penrose's credit he, along with Stuart Hameroff, have put forth

what is perhaps the only concrete proposal of a mechanism for quantum computing in the brain, and so it is worth investigating.

The model goes by the name of “Orchestrated Objective Reduction” (orch OR). It posits that there exists some structure in the brain capable of quantum computation, sufficiently isolated from the rest of the brain that quantum mechanical states remain coherent on timescales relevant to neural processes. What ultimately decoheres the superimposed states and forces them into definite states is presumed to be quantum gravity effects, which is where the non-computability enters (and is also arguably presently the weakest part of the model, as there is no quantum theory of gravity to implement this non-computability, though string theory methods have been employed to try and flesh out this idea¹³).

The key physiological components in the orch OR model are microtubules, often associated with cell structure but which also have roles in other cell functions. Microtubules are tube-like structures made of 13 “strings” made of a dimer called tubulin. The tube shape provides a potential shield from the environment for any quantum mechanical states inside the tube, and the tubulin has two conformations, open and closed, which could provide the “0” and “1” states to be superimposed in a quantum computation. The states can also be viewed as black and white cells in a cellular automata calculation, as shown in Fig. 4. Such a state could in principle compute classically using simple cellular automata rules¹².




FIG. 4: Schematic depiction of an automaton calculation on a microtubule, from Ref. 12.

Penrose and Hameroff argue this still is not enough - though this can perform as a computer, to really establish consciousness one needs to be able to superimpose tubulin states such that the microtubules can begin in some initial state, form superpositions, and then collapse into a final state *which is not computable from the initial state*. This is where Penrose posits that quantum gravity enters: because mass and energy bend spacetime, superpositions of the different tubulin conformations (or positions, etc), represent superpositions of different local spacetimes. When a large enough collection of tubulin dimers are in superpositions they reach a critical mass at which they interact strongly enough through quantum gravity that the superposition collapses into a definite state, but one which cannot be computed. This process is shown schematically in Fig. 5¹².

So, the million-dollar question (or however much a Nobel Prize award gets you these days) is, “Is this a feasible mechanism for quantum computation in the brain?”. The first test to pass is the issue of decoherence. Could the computations as described possibly occur before the environment washes away the quantum superposition? Max Tegmark has done rough, detailed calculations¹³ (rough in the sense that they are order-of-magnitude estimates, but detailed in the sense that you probably couldn’t do them on the back of an envelope,




FIG. 5: Schematic depiction of the quantum orch OR computation. Left to right, tubulin dimers begin to become superimposed, with neighbouring dimers begin to become superimposed until reaching critical mass, and finally the wave function collapses, yielding an uncomputable state. Figure adapted from Ref. 12.

unless it was an unusually large envelope) that suggest colliding ions or water molecules should cause decoherence in about $10^{-20} - 10^{-19}$ seconds, much faster than the proposed timescale on which quantum gravity supposedly acts. He concludes that quantum effects would be too short lived, and our classical treatment of the brain is appropriate. However, the authors of Ref.14 argue that these calculations are at best relevant to a different model that is not orch OR, and at worst incorrect as the decoherence times in Tegmark's paper are predicted to grow with temperature, which conflicts with our intuitions about decoherence times dropping off with temperature. The authors of Ref. 14 establish their own estimates for decoherence times in the orch OR model, coming up with a decoherence time scale of $10^{-5} - 10^{-4}$ seconds, which they argue is slow enough for the quantum mechanical effects to take place. Furthermore, they suggest that the coherent state may be enhanced by Frölich condensation, a pseudo-Bose-condensation in which "incoherent metabolic energy is used to force coherence in much the same way that coherence is induced in lasers"¹⁵. It is proposed that this enhancement will enable the coherent state to persist long enough to be relevant to physiological processes. They further contend that there may be topological states that are stable against decoherence.

Assuming, for now, that decoherence is not a problem, how else might this model be tested? The other main assumption it makes is that the tubulin can switch between its two conformations, and do so fast enough to be relevant to the function of the brain. When the orch OR theory was proposed it was known that tubulin had these two configurations, but the details were unknown. Recently, the authors of Ref. 15 have addressed this issue in light of a new understanding of the biological details of tubulin. The picture of tubulin presented in Ref. 15 is more complicated than in orch OR's concept of it. The conformation of the dimer depends on whether GTP or GDP is bound to one of its monomers. In solution this leads to the open or closed states, but only those dimers with GTP are incorporated into microtubules during assembly. Even if GTP is hydrolyzed to GDP after assembly it does not result in a significant structural change of the dimer. Even if this change were significant enough to be treated as a different state, conversion from GDP back to GTP while in the microtubule is essentially irreversible - the microtubule cannot switch between different states without being completely disassembled and then reassembled, a process that

takes far too long to be used in computations on timescales appropriate to thought processes.

Could there be a revision of the model that salvages it from destruction? The authors of Ref. 15 consider this, postulating that perhaps a different combination of tubulin isomers may work; however, they conclude that the only feasible mechanism which could produce coherent states for long enough, the Frölich condensation, is still not sufficient to maintain coherence long enough to be relevant. The mechanism would require a GTP to GDP conversion rate that is ten orders larger than what is experimentally observed. The authors conclude that there is no foreseeable way to salvage the orch OR model; however consciousness arises, it would seem that Nature does not use quantum mechanics to make it.

V. A CONCLUSION

I am afraid I will have to conclude our story without a proper ending, as it has yet to be written. We are still a long way from understanding consciousness, but we have made great strides in understanding collective phenomena in the mind, which are sure to be essential to the ultimate emergence of the mind from neurons that make up the brain (assuming, of course, that to ask the question “What is consciousness?” is not ultimately unintelligible!). I did not have the space or time to survey all of the observed emergent phenomena in the mind, and so choose to focus on memory and pairwise correlation models as the primary examples, and told the interesting story of how quantum gravity, of all things, might directly enter the formation of the mind. This essay has hopefully convinced the reader (if the reader is not already convinced) that there are tremendously interesting problems in biological systems, especially in the brain. It is with this that I end our tale; I will leave its resolution as an exercise to the reader, the author, and the brilliant scientists already working hard on understanding our own minds.

¹ Rabinovich, M. I., *et al.*, *Rev. Mod. Phys.* **78**, p. 1213 - 1265 (2006)

² This is, strictly speaking, a lie. Some dendrites do transmit information instead of just receiving it, but for the sake of clarity in my exposition certain white lies must be told.

³ Levitan, I. B. and Kaczmarek, L. K., *The Neuron: Cell and Molecular Biology*, Oxford University Press (2002)

⁴ Herz, A. V. M., *et al.*, *Science*, **314**, p. 80-84 (2006)

⁶ Of course, that makes for a less-catchy essay title.

⁵ Hopfield, J. J., *Proc. Nat. Ac. Sci.*, **79** 8, p. 2554-2558 (1982)

⁷ Schneidman, E., *et al.*, *Nature*, **440**, p. 1007-1012 (2006)

⁸ Roudi, Y., Nirenberg, S., Latham, P. E., *Pub. Lib. of Sci.*, **5**, 5, (2009)

⁹ Cumin, D., Unsworth, C. P., *Physica D*, **226**, p. 181-196 (2007)

¹⁰ Lin, C., Lin, M., *Commun. Nonlin. Sci. Numer. Simulat.*, **14**, p. 3258-3260 (2009)

¹¹ Penrose, R., *The Large, the small, and the human mind*, Cambridge University Press (1997)

¹² Hameroff, S. R., <http://www.quantumconsciousness.org/penrose-hameroff/quantumcomputation.html>, accessed Dec. 2009.

¹³ Tegmark, M., *Phys. Rev. E*, **61**, 4, p. 4194-4206 (2000)

¹⁴ Hagan, S., Hameroff, S. R., and Tuszynski, J. A., *Phys. Rev. E*, **65**, 061901 (2002)

¹⁵ McKemmish, L. K. *et al.*, *Phys. Rev. E*, **80**, 021912 (2009)