

Using Quantum Computing to Infer Dynamic Behaviors of Biological and Artificial Neural Networks

Gabriel A. Silva

Department of Bioengineering and Department of Neurosciences
Center for Engineered Natural Intelligence and Kavli Institute for Brain and Mind
University of California San Diego, La Jolla CA 92093 USA

Email: gsilva@ucsd.edu

Abstract. The exploration of new problem classes for quantum computation is an active area of research. An essentially completely unexplored topic is the use of quantum algorithms and computing to explore and ask questions *about* the functional dynamics of neural networks. This is a component of the still-nascent topic of applying quantum computing to the modeling and simulations of biological and artificial neural networks. In this work, we show how a carefully constructed set of conditions can use two foundational quantum algorithms, Grover and Deutsch-Jozsa, in such a way that the output measurements admit an interpretation that guarantees we can infer if a simple representation of a neural network (which applies to both biological and artificial networks) after some period of time has the potential to continue sustaining dynamic activity. Or whether the dynamics are guaranteed to stop either through 'epileptic' dynamics or quiescence.

Contents

1	Introduction	2
2	Problem Set Up	3
2.1	Preliminaries and Problem Structure	4
2.2	Functional Interpretations of $f(V)$ Outcomes	6
3	Determining membership in S_n and S_n^c using Grover's algorithm	7
4	Applying Deutsch-Jozsa to Determine if a Network Can Sustain Activity	8
4.1	Evaluating S_n and S_n^c with a Two-Part U_f	8
4.2	Constructing U_f to evaluate $f(S_n)$	9
5	Measurement Interpretations of U_f	10
6	Discussion	11
6.1	Quantum versus Classical Considerations for Studying Neural Networks	11
6.2	Practical Considerations and Applications of This and Future Related Work	12

1 Introduction

Quantum computing, at least theoretically at present, has the potential to revolutionize areas where classical computers show limitations, particularly in cryptography and the simulation of complex physical and chemical systems, including quantum mechanics itself. Notably, quantum algorithms like Shor’s for integer factorization could significantly disrupt traditional cryptographic techniques, presenting a mix of challenges and opportunities in the realm of data security [1]. While research in quantum computing has historically focused on these well-established topics, the exploration of new problem classes that are particularly suitable for quantum computation is an active area of research, to include the application of known quantum algorithms and the development of new ones. In fact, Google Quantum and X Prize just announced a new competition specifically to promote the development of practical, real-world quantum computing algorithms and applications (<https://www.xprize.org/prizes/qc-apps>).

A still-nascent topic is the application of quantum computing to the modeling and simulation of biological and artificial neural networks. In particular, a completely unexplored topic is the use of quantum algorithms and computing to explore and ask questions *about* the functional dynamics of neural networks. We propose that prioritizing research in these areas is important, as it can potentially significantly advance our understanding of the brain and mind and the increasingly rapid development of artificial intelligence.

To be sure, the intersection of quantum computing, neuroscience, and related topics has generated considerable interest already, leading to quite a few books, technical papers, and popular articles. Some of what has been written is speculative but scientifically thoughtful (see, for example, [2], [3], [4], [5], [6], [7]), and a lot is not. Probably most well-known are the arguments by Hammaroff, Penrose, and others suggesting that microtubules in neurons act as quantum computing elements in fundamental ways [8], [9]. While still controversial, recent experimental results suggest that neurons might indeed leverage properties of quantum effects in how they compute [10], a remarkable result. Our focus here, however, is different. We are interested in the practical application of quantum algorithms for probing and understanding neural dynamics rather than exploring the possibility of quantum computational mechanisms within the brain itself.

Despite the growing interest in these topics, rigorous mathematical and quantitative arguments applying quantum algorithms for solving meaningful neuroscience and artificial neural network questions remain in their infancy.

There are two broad areas in which quantum computing may contribute to the scientific study of neural networks. The first is large-scale simulations of neural dynamics across scales of spatial and temporal organization, bounded and informed by the known physiology (in the case of the brain) and known models (for artificial intelligence and machine learning). The second is carefully chosen and structured problems about neural dynamics that make careful and clever use of quantum algorithms.

Sufficiently large-scale simulations would allow observing, experimenting, and iterating numerical experiments under a wide range of parameter and model conditions. If, as neuroscientists suspect, complex emergent cognitive properties are partly due to sufficiently large interactions among foundational physiological and biological components and processes across temporal and spatial scales of organization, the need to carry out very large iterative simulations may be critical to understanding the dynamics that give rise to cognitive properties. Simulations of this kind would support understanding emergent effects that depend on the scale of the computational space, to the extent it can be computed.

However, open-ended large-scale simulations alone will not be sufficient for understanding how the brain, or artificial intelligence for that matter, works. In effect, open-ended large-scale simulations in isolation are what led to the significant challenges and missed targets faced by the highly publicized and hugely funded Blue Brain Project [11]. Arriving at such an understanding necessitates carefully chosen and defined

problems *about* the neural network dynamics being simulated. This is critical. Observing neural dynamics in action — for example, the firing patterns of large numbers of neurons — in isolation and without context, by itself, cannot reveal the underlying algorithms that are operating on those dynamics or why they exist. It is no different than attempting to understand the brain from a systems perspective by studying a single participating protein or ion channel in isolation. Quantum computing potentially has a unique role to play in deciphering and understanding the contributions that scale has on functional network dynamics and behaviors.

In this work, we show how a carefully constructed problem can leverage two foundational quantum algorithms, Grover and Deutsch-Jozsa, in such a way that the output measurements from the quantum computations admit an interpretation that guarantees we can infer a particular behavior about the network dynamics. Specifically, we address the following question: Given a simple representation of a neural network (which applies to both biological and artificial networks) after some period of time of the evolution of its dynamics, does the network have the potential to continue sustaining dynamic activity? Or are the dynamics guaranteed to stop in one of two ways: either through ‘epileptic’ type saturated dynamics or quiescence (i.e., all activity dying away)? We show that by structuring (asking) the problem in a particular way, the implementation of the Grover and Deutsch-Jozsa algorithms arrives at a solution much more efficiently than would be possible by strictly classical methods. As we argue in the Discussion section when one considers the combinatorial size of the computational space of the real brain, for example, these types of questions become inaccessible to classical methods.

Our intent here is to explore an interesting proof of concept problem and to motivate the potential of quantum computation to neural networks as an important research direction. We introduce a number of novel results, including the technical structure of how the problem was set up, and, to our knowledge, one of the only few examples of a practical application of Deutsch-Jozsa. The first application of this kind to the study of neural network dynamics. We have intentionally chosen a tractable and applied problem that has real-world implications and consequences for neural networks, solved in the simplest and most intuitive possible way using two of the best-known foundational quantum algorithms. Given the bespoke nature of setting up, mathematically proving (when possible), and solving quantum computing problems, the results we discuss here are an exercise in thinking about how neural network dynamics questions need to be structured within a quantum computational framework by leveraging deep knowledge about the physiology and models of neural networks. While the problem as structured and the quantum computing approaches we discuss are simple, we have attempted to be mathematically rigorous. The primary focus of this paper is the theoretical arguments necessary to set up and solve the problem. We do not explicitly suggest the design of quantum circuits, although we do discuss throughout the paper the various requirements of circuits that might implement the necessary logic flows and unitary functions.

2 Problem Set Up

Assume we have a (neurobiological or artificial) neural network whose dynamics has evolved for some period of time T_o . We want to evaluate a specifically constructed instance of the Deutsch-Jozsa algorithm to determine if a network is in a state that has the potential to inherently (i.e., in the absence of external stimulation or inputs) sustain continued dynamic activity or if it cannot. We will show that we can map the functional condition that the network cannot sustain activity (for two distinct cases) to an output of the Deutsch-Jozsa algorithm associated with the result of a constant function. Conversely, we can interpret a none-constant (but not necessarily balanced) output of the algorithm as a network in a state that will continue sustaining dynamic activity for some future period of time of at least an additional time step beyond T_o .

The Deutsch-Jozsa algorithm determines whether a given function, which takes a binary input and

returns a binary output, is either 'constant' or 'balanced'. A 'constant' function is one that returns the same output — either all 0's or all 1's — for any input, reflecting a uniform behavior across its domain. In contrast, a 'balanced' function returns 0's for exactly half of the possible inputs and 1's for the other half.

2.1 Preliminaries and Problem Structure

Define a graph model of a neural network as $G = (V, E)$, where V is the vertex set of the graph G , and E is the set of edges. A network is modeled by G in the sense that V and E may have additional features or properties specific to the network being modeled. A node in the network is a vertex with such additional features or properties.

For every vertex $v_i \in V$, at some time T_o , i.e., a 'snapshot' of the dynamically evolved state of G at T_o , we can define a summation and activation function that determines if all the weighted inputs into v_i result in an activation, i.e., firing, event at T_o . There are no restrictions on the dynamics prior to T_o , i.e., between $0 \leq t \leq T_o$, or any stimuli that may have caused or influenced those dynamics. The only consideration is the instantaneous state of the dynamics of the network at T_o . There are many ways to explicitly construct such a function within the theory of artificial neural networks and computational and theoretical neurobiology. For example, our lab has developed a geometric latency-dependent physiologically accurate model that we have shown is implemented as a computational optimization rule in individual neurons and the connectome of the worm *C. elegans* [12] [13] [14] [15] [16].

In the most general case, we can define this function for all $v_i \in V$ by

$$f(V) = \begin{cases} \text{if } \Sigma_r \geq \Sigma_T, v_i = 1 \\ \text{if } \Sigma_r < \Sigma_T, v_i = 0 \end{cases} \quad (1)$$

Σ_r represents a 'running summation'. It reflects the totality of weighted inputs into v_i at a given point in time. Σ_T is a 'threshold constant'. It reflects the processes that model the internal dynamics of v_i that ultimately determine if that node will fire or not. For example, in the framework developed by our group, such internal dynamics are due to the interplay between the geometry of the edges between nodes, the latencies of discrete signaling events as they travel on those edges before they arrive at downstream nodes, refractory periods, and weights, to define a particular form of a neurobiological activation function. Many other representations and models exist for both neurobiological networks and artificial neural networks.

In the simplest case, $f(V)$ is a linear summation of weighted inputs overcoming a threshold constant resulting in a deterministic firing event. A form of an 'integrate and fire' neuron. For simplicity, we assume this latter form throughout the rest of the paper but refer the reader to the broader literature for a more comprehensive discussion.

At T_o every $v_i \in V$ will have associated with it a value $\Sigma_r \in \mathbb{R}$ that can be normalized between some defined range

$$\epsilon \leq \Sigma_r \leq (\epsilon + \Delta)$$

For example, $0 \leq \Sigma_r \leq 1$. This range can be approximated by an n -bit string sufficient to encode some desired degree of precision:

$$\Sigma_r \in \{0, 1\}^n$$

Note that we do not (need to) distinguish with different notation the continuous version of Σ_r from its n -bit string representation. Thus, $f(V)$ is a function that compares an n -bit input $\Sigma_r \in \{0, 1\}^n$, the running summation for each vertex $v_i \in G(V, E)$, to a constant value threshold $\Sigma_T \in \{0, 1\}^n$. It evaluates this for each v_i at some time T_o and outputs either 0 or 1, indicating the vertex v_i fires (activates) when $f(V) = 1$ and does not when $f(V) = 0$.

For realistic neural networks — both biological and artificial — the cardinality of V , $|V|$, could be in the billions. (See Section 6 below for a discussion about the dimensionality and computational space of biological neural networks.) But the set of Σ_r values will be much smaller. Given a realistic (i.e., functionally necessary) n -bit precision to represent Σ_r , many nodes would be expected to have the same values of Σ_r . While some other values of Σ_r may not be repeated and, therefore, unique.

Mathematically, we can construct an ordered list \mathbb{V} (in the set-theoretic sense) such that its elements map a value of Σ_r for every v_i . Informally, we can write this as

$$\mathbb{V} : v_i \rightarrow \Sigma_r \text{ for every } v_i \in V \quad (2)$$

It is also possible that only a subset of the available set of total numbers representing all possible values of n -bit strings that Σ_r can take on will actually be realized by the network. In other words, we expect that there will be some n -bit strings that, for some particular instance of an evolution of the network dynamics, do not show up as a value of Σ_r .

Let the complete set of all possible n -bit strings of length n be

$$\mathbb{S}_n := \{s_i | s_i \in \{0, 1\}^n\} \text{ for all } n\text{-bit strings of length } n \quad (3)$$

Then the following definitions follow.

$$S_n := \{s_i | s_i \in \{0, 1\}^n \wedge s_i = \Sigma_r \in \mathbb{V}\} \quad (4)$$

is the subset $S_n \subset \mathbb{S}_n$ who's elements include all n -bit string values necessary to represent all values of Σ_r .

$$S_n^c := \{s_i | s_i \in \{0, 1\}^n \wedge s_i \neq \Sigma_r \in \mathbb{V}\} \quad (5)$$

is the subset $S_n^c \subset \mathbb{S}_n$ who's elements include all n -bit string values *not* necessary to represent values of Σ_r .

It is obvious that S_n^c is the complement of S_n and that

$$S_n \cup S_n^c = \mathbb{S}_n$$

Note how there exists a surjective mapping from $\mathbb{V} \rightarrow \mathbb{S}_n$ (Figure 1), and

$$|\mathbb{V}| = |V| \gg |\mathbb{S}_n| \geq |S_n|$$

For example, while $|V|$ could be in the billions, the binary representation of values of Σ_r to two significant digits (for example, $0 \leq \Sigma_r \leq 1$) would only necessitate 7 bits, which can encode up to 128 values. Two significant digits are required to encode 101 values, i.e., 0 to 0.99 plus the value 1.0.

Equations 2, 3, 4, and 5 motivate the technical details of how we use Grover's algorithm to determine set membership in S_n given \mathbb{V} and how we construct the unitary function to be evaluated by the Deutsch-Jozsa algorithm.




Figure 1: There exists a surjective mapping $\mathbb{V} \rightarrow \mathbb{S}_n$. For $v_i \in G(V, E)$, its corresponding value Σ_r at time T_o maps to one of n -bit string values $s_i \in \mathbb{S}_n = \{0, 1\}^n$.

2.2 Functional Interpretations of $f(V)$ Outcomes

The goal is to evaluate a specifically constructed instance of the Deutsch-Jozsa algorithm to determine if a network is in a state that has the potential to inherently (i.e., in the absence of external stimulation or inputs) sustain continued dynamic activity or whether it cannot.

An inability of a network to sustain activity occurs under two conditions. In the first case, the network becomes 'epileptic', which in our context here implies that all vertices simultaneously fire at T_o , thereby ensuring that they continue to do so for times $t > T_o$. This essentially results in a perpetually saturated constant and global firing state for the entire network. (It is possible that epileptic dynamics may lead to quiescence or maybe in an epileptic state only for some period of time. A discussion about the various conditions that produce these dynamics is beyond the scope of this paper.) A network in this state cannot encode or process information. Formally, this will occur when

$$f(V) = 1 \ . \forall v_i \in G(V, E) \quad (6)$$

In the second case, the network is not capable of sustaining dynamic activity, with the activity eventually dying away, resulting in a quiescent network. There are a number of conditions that can result in this (see, for example, Fig 3. and accompanying text in [12]). This will occur when

$$f(V) = 0 \ . \forall v_i \in G(V, E) \quad (7)$$

In both these cases, $F(V)$ is a constant function that takes as inputs the n -bit strings in S_n and outputs a binary value from $\{0, 1\}$.

When $f(V)$ is not constant, it ensures that the network can sustain activity on its own for at least $T_o + 1$ time steps, i.e., the next time step in its evolution. We will show that we can distinguish between the two states represented by Equations 6 and 7 by running the algorithm twice. Note that we are not claiming nor do we require $f(V)$ to be balanced (in the language of Deutsch-Jozsa). Just that it is not constant. This is consistent with the unitary function requirement in Deutsch-Jozsa needing to take on the general form $f(\{0, 1\}^n) \rightarrow \{0, 1\}$.

We now have a set of conditions that map the structure of Deutsch-Jozsa to an interpretation associated with neural network dynamics. We next need to consider how to evaluate whether $f(V)$ is constant or not

constant for all v_i simultaneously. This will determine whether the network can sustain recurrent activity. Any classical approach would necessitate evaluating and checking each instance of $F(V)$ sequentially for all v_i in the network or the use of an equivalent amount of computational resources.

However, this is a special case of Deutsch-Jozsa because we need to evaluate S_n as inputs and not the full domain \mathbb{S}_n . The Deutsch-Jozsa algorithm assumes that the oracle function being evaluated is defined over the entire domain $\{0, 1\}^n$. It cannot evaluate a defined subset of $\{0, 1\}^n$. In other words, the algorithm's outputs — and the interpretation of its outputs to determine if the function is constant or balanced — only make sense when it evaluates $f(\mathbb{S}_n = \{0, 1\}^n)$. Thus, it cannot explicitly evaluate $f(S_n)$. The algorithm's results would not reflect the behavior of the defined subset of interest, in this case, S_n . To get around this, we will extend S_n through a specific construction that, when evaluated by Deutsch-Jozsa, allows us to indirectly interpret the behavior of the network.

3 Determining membership in S_n and S_n^c using Grover's algorithm

First, however, we must determine which members of \mathbb{S}_n are in S_n . Given how we set up the problem, this requires searching for each n -bit string $s_i \in \mathbb{S}_n$ in the list \mathbb{V} of Σ_r values that resulted from the dynamic evolution of the network up to time T_o , the time at which we are evaluating the network. The resultant set is S_n . The collection of elements in \mathbb{S}_n that are not in \mathbb{V} make up S_n^c . As introduced above, recall that we expect some members of \mathbb{S}_n to be present in \mathbb{V} only once, some to be repeated, and some not to be present at all. We could potentially use an implementation of Grover's algorithm to construct S_n much more efficiently than a purely classical approach that would necessitate sequentially checking each n -bit string element of \mathbb{S}_n in a huge list \mathbb{V} .

Recall that the list \mathbb{V} contains the values of the running summation Σ_r for each vertex v_i at the observation time T_o (equation 2). Each element of \mathbb{V} will be one of the n -bit string values in \mathbb{S}_n . Label each of the members in \mathbb{V} by $\{\nu_1, \nu_2, \dots, \nu_M\}$. For each n -bit number $s_i \in \mathbb{S}_n$, we construct an oracle function U_f that identifies if s_i is in \mathbb{V} . In other words, we check if there exists a ν_i in \mathbb{V} that equals some value of s_i in \mathbb{S}_n . Formally, the condition we are checking for is

$$\exists \nu_i \in \mathbb{V} \text{ such that } \nu_i = s_i \text{ for some } s_i \in \mathbb{S}_n \quad (8)$$

We want U_f to perform the operation

$$U_f |\nu_i\rangle = (-1)^{f(\nu_i)} |\nu_i\rangle \quad (9)$$

$|\nu_i\rangle$ encodes a state corresponding to an element $\nu_i \in \mathbb{V}$. The function $f(\nu_i)$ evaluates to 1 iff $\nu_i = s_i$ for one of the elements $s_i \in \mathbb{S}_n$, and 0 otherwise.

In practice, the construction of U_f would necessitate a programmable oracle function that can be re-coded for each n -bit string s_i . This might involve the use of conditional quantum gates that implement the logical statement 'if the current state of $\nu_i = s_i$ apply a phase flip; otherwise, do nothing'.

The standard diffusion operator D in Grover's algorithm, in this case, would be defined as

$$D = 2 |\sigma\rangle \langle \sigma| - I \quad (10)$$

for a uniform superposition of all $s_i \in \mathbb{S}_n$ given by

$$|\sigma\rangle = \frac{1}{\sqrt{N}} \sum_{s_i=0}^{N-1} |s_i\rangle \quad (11)$$

I is the $N \times N$ identity matrix. Equation 11 defines the search space over which the algorithm will search. D effectively inverts the probability amplitudes about the average amplitude for a given s_i , thereby amplifying the states present in \mathbb{V} marked by U_f .

To check if each $s_i \in \mathbb{S}_n$ is in \mathbb{V} , we need to run the algorithm iteratively 2^n times. Continuing the example from above, for a practically reasonable precision of Σ_r we would expect this to be at most order of magnitude 10^2 times, given a two-digit precision requires 128 values.

We anticipate the number of Grover iterations needed for amplification of a specific s_i to be approximately

$$\frac{\pi}{4} \sqrt{\frac{N}{\mu}} \quad (12)$$

where μ here is the number of occurrences of a particular s_i in \mathbb{V} , i.e., the number of ν_i that equal a particular s_i . As would be expected, the more occurrences that are present the more efficient the algorithm operates. Optimizing the number of iterations however is non-trivial, and may benefit from other information or patterns in how $G(V, E)$ is constructed that may provide further structure and knowledge about \mathbb{V} . In the worst case we would anticipate running the algorithm $\frac{\pi}{4}N$ times, which would yield a quadratic speedup over any classical search method which would still be significant given the size of V , i.e., $|V|$ and related considerations about the dimensionality of the computational space (see the discussion in Section 6).

If searching for a specific s_i returns s_i after measuring, we assume it is present in \mathbb{V} . If instead the algorithm returns a different value $s_j \neq s_i$ in \mathbb{S}_n we conditionally assume s_i is not in \mathbb{V} . We say conditionally here because, in practice, it is important to consider and take into account the inherent probabilistic nature of quantum mechanics and quantum algorithms. There is of course a non-zero probability that $s_j = s_i$ when in fact s_i is not in \mathbb{V} . In practice, some form of verification may be needed, such as re-running the algorithm for the same value of s_i more than once.

Finally, once we have identified S_n , determining S_n^c is trivial.

4 Applying Deutsch-Jozsa to Determine if a Network Can Sustain Activity

With S_n and S_n^c identified, we are now ready to make use of the Deutsch-Jozsa algorithm to determine if the network can sustain inherent recurrent activity at times beyond T_o or if it cannot. However, as discussed in detail above, the set we need to evaluate is S_n , not the full domain $\mathbb{S}_n = \{0, 1\}^n$.

4.1 Evaluating S_n and S_n^c with a Two-Part U_f

One way to work around this requirement is to extend S_n into \mathbb{S}_n by taking advantage of the fact that $S_n \cup S_n^c = \mathbb{S}_n$. We need a construction that extends S_n in such a way that when we evaluate U_f , we can interpret the result in a way that gives us information about the behavior of the network — specifically, the evaluation of Equation 1 over S_n — indirectly inferred from the output of the algorithm.

To achieve this, we define a 'two-part' oracle function U_f that operates on the state $|s_i\rangle|-\rangle$, where $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ as is standard. Recall that Equation 1 defines a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. Then, for $s_i \in S_n$

$$\text{If } f(s_i) = 1 \text{ then } U_f |s_i\rangle|-\rangle = -|s_i\rangle|-\rangle \quad (13a)$$

In this case U_f flips the phase of $|s_i\rangle$ when the ancillary qubit is in the $|-\rangle$ state.

$$\text{If } f(s_i) = 0 \text{ then } U_f |s_i\rangle|-\rangle = |s_i\rangle|-\rangle \quad (13b)$$

Here, U_f does not alter the state.

For $s_i \in S_n^c$, instead of evaluating $f(s_i)$, like we did for $s_i \in S_n$, U_f will set each $s_i = 0$ or $s_i = 1$. From a functional perspective, similar to Equation 13,

$$\text{If } s_i = 1 \forall s_i \in S_n^c \text{ then } U_f |s_i\rangle |-\rangle = -|s_i\rangle |-\rangle \quad (14a)$$

Analogous to Equation 13a, U_f flips the phase of $|s_i\rangle$. But, as in Equation 13b,

$$\text{If } f(s_i) = 0 \text{ then } U_f |s_i\rangle |-\rangle = |s_i\rangle |-\rangle \quad (14b)$$

From a practical perspective, implementing U_f would initialize all $|0\rangle^n$ n -qubit registers in superposition and the ancillary qubit in the $|-\rangle$ state:

$$|\psi_o\rangle = \frac{1}{\sqrt{2^n}} \sum_{s_i=0}^{2^n-1} |s_i\rangle \otimes |-\rangle$$

We then apply U_f as defined in Equations 13 and 14 by first checking the membership of s_i in either S_n or S_n^c , followed by applying $f(s_i)$ when appropriate or 'hardcoding' the output to either 0 or 1. For the case where $s_i \in S_n^c = 0$ U_f applies the identify matrix and leaves the state unchanged.

Checking for membership in S_n versus S_n^c using a set of quantum circuits is a nuanced process, even though we know the membership of s_i in each. Because we know *a priori* set membership in S_n and S_n^c before the Deutsch-Jozsa step, conceptually, this could be achieved by constructing a series of quantum circuits that 'mark' $s_i \in S_n^c$ so that instead of being evaluated by the function given by Equation 13 they are evaluated by the operation in Equation 14.

Because $\mathbb{S}_n = \{0, 1\}^n \equiv |0\rangle^{\otimes n}$ is in superposition, the operation to mark the members of S_n^c must be carried out simultaneously across this superposition. This could be achieved using a set of controlled quantum circuits, each designed to recognize and selectively apply a phase flip to the states corresponding to $s_i \in S_n^c$, using additional ancillary qubits. Functionally, this would necessitate many controlled circuits within U_f where each circuit is designed to recognize and act on one value known to be in the set S_n^c . The marking operations effectively occur simultaneously because the quantum computation happens in parallel across the superposition. Using a phase flip to carry out the marking operation might be a computationally feasible approach because it is a reversible operation that does not collapse the superposition. It subtly alters the state as part of the operation U_f . The other part of U_f then evaluates $f(S_n)$, the combined U_f operating on the superposition across the entire domain \mathbb{S}_n . We discuss a form of $f(S_n)$ in the next section.

4.2 Constructing U_f to evaluate $f(S_n)$

We next need an explicit description to construct the part of U_f that evaluates $f(S_n)$ as defined by Equation 1 as a logical quantum circuit.

We first need to encode each $s_i \in S_n$ that corresponds to values of $\sum_r \in \mathbb{V}$ and the n -bit constant \sum_T into quantum states. Conceptually, we can assign n -qubits to encode each n -bit value and then construct $f : \{0, 1\}^n \rightarrow \{0, 1\}$ to implement Equation 1 in such a way that it compares the n -qubit encoded values of \sum_r and \sum_T . The comparison output reflects the function output, either 0 or 1, corresponding to a firing (discrete signaling) event outcome for each v_i in the network.

To achieve this, we can build quantum gates that implement a binary subtraction $\sum_r - \sum_T$. If the result of this subtraction is positive or zero then $\sum_r \geq \sum_T$ and $f(S_n) = 1$. If the subtraction is negative then $\sum_r < \sum_T$ and $f(S_n) = 0$.

The output value of $f(S_n)$ is stored in an ancillary qubit, which has an assigned state $|1\rangle$ for positive values of the binary subtraction, and $|0\rangle$ when it is negative. This can be achieved by a quantum circuit that conditionally flips the ancillary qubit based on the outcome of the subtraction. Measurement of the that qubit then collapses to $|1\rangle$ or $|0\rangle$ appropriately.

Consider first the case for one specific vertex in the network $v_i \in G(V, E)$, with a corresponding entry $\nu_i \in \mathbb{V}$, where ν_i encodes the value of Σ_r at time T_o (Equation 2). Start with $2n + 1$ qubits. The first $q_0 \dots q_{(n-1)}$ qubits represent Σ_r . The next $q_n \dots q_{(2n-1)}$ qubits represent Σ_T . The last q_{2n} qubit is the ancillary qubit that encodes the comparison result. ν_i can be represented by the state $|\nu\rangle = |\nu_{n-1}\nu_{n-2} \dots \nu_o\rangle$. The constant value of the threshold Σ_T can be represented by the state $|c\rangle = |c_{n-1}c_{n-2} \dots c_o\rangle$. The ancillary qubit is initialized to the $|0\rangle$ state.

One simple approach to carrying out the comparison is via a binary subtraction comparison for each pair of qubits between $|\nu\rangle$ and $|c\rangle$ starting from the most significant digit. This tells us if the result of the subtraction is negative or positive since we are interested in determining the sign of the resultant operation, not the actual result. For each pair of bits (ν_j, c_j) , perform a controlled operation that determines if $\nu_j < c_j$. If it does, then it implies that ν_i is less than Σ_T . If the bit values are equal, just move to the next pair until an inequality occurs.

For example, in keeping with our example from above, a reasonable assumption is $n = 7$, i.e., we can represent Σ_r and Σ_T by sufficient precision as two-digit values, which can be encoded by 7-bit numbers. Assume that $|\nu\rangle = |\nu_6\nu_5\nu_4\nu_3\nu_2\nu_1\nu_o\rangle = |1010101\rangle$, representing a decimal value of $\Sigma_r = 0.85$, and that $|c\rangle = |c_6c_5c_4c_3c_2c_1c_o\rangle = |0110010\rangle$, representing a value of $\Sigma_T = 0.5$. Starting from the most significant bit, we compare $|\nu_j\rangle$ to $|c_j\rangle$. In this example the first comparison we do is $1 > 0$, which implies that $|\nu\rangle > |c\rangle$. Because we do not actually need to know the result of the comparison, we can immediately stop here and set the ancillary qubit to $|1\rangle$. A quantum circuit implementation would again likely make use of a series of controlled gates to achieve this.

5 Measurement Interpretations of U_f

When we evaluate U_f , if we set $S_n^c = 0$ and measure the output to be constant, i.e., $|0\rangle^{\otimes n}$, we can interpret this as the network's dynamic activity dying away at $T_o + 1$. The network is quiescent. This is because, in this case, the output can only be constant when $f(S_n) = 0$, given that $S_n^c = 0$. In other words, the network is necessarily quiescent because the output from the Deutsch-Jozsa algorithm will evaluate to a constant function only when $f(s_i) = 0 \forall s_i \in S_n$.

If, on the other hand, we set $S_n^c = 1$ and measure the output to be constant, we can interpret this as the dynamic activity in the network becoming epileptic since it will only return constant when $f(S_n) = 1$.

But if U_f is *not* constant for either case having set $S_n^c = 0$ or $S_n^c = 1$, i.e. the measured output is a non-zero bit string, we interpret that as the network being in a dynamic state that is neither epileptic nor quiescent, and therefore, has at least the *potential* for continued dynamic activity. This is the case because not constant will only occur when U_f evaluates $f(S_n)$ to some mixture of 1's and 0's. In other words, a mixture of activated and non-activated vertices in the network.

As discussed above, in general, what we cannot say is whether U_f , and therefore $f(S_n)$, is balanced. But for our purposes, in the context of evaluating the signaling potential of the network, this does not matter.

6 Discussion

6.1 Quantum versus Classical Considerations for Studying Neural Networks

In this work, we showed how the dynamical signaling potential of neural networks at future times, given the system’s temporally evolved state, can be determined using quantum computing. While solving the problem as we have set it up is not outside the realm of what can be achievable by classical methods, we have shown how leveraging quantum algorithms would solve the problem more efficiently, at least quadratically better for the Grover search part of the process and exponentially for the evaluation of S_n by Deutsch-Jozsa. For example, an average 3 lb adult human brain contains approximately 85 billion neurons, i.e. $|V| = 85$ billion. At this scale, for the Grover search part of the solution, we estimate that it would take Frontier, the fastest existing current supercomputer, on the order of hours to solve the problem. On a laptop, the task would take many hours to days. But at least theoretically, at worst, it would take seconds to minutes to solve on a quantum computer. Still, this is easily solvable using classical architectures.

However, the disparity between quantum and classical temporal and computing resources very quickly increases beyond any conceivable classical capabilities as a network’s structure, size, and complexity increase. For example, in the human brain, each of the 85 billion neurons form between 10,000 to 100,000 synaptic connections, leading to an estimated 10^{16} , 10 quadrillion, synapses [17], [18], [19]. If we consider just 1000 neurons, each acting as statistically independent computational elements capable of firing on their own, if of those just 2% are firing at any given time, the number of possible distinct encodings is $1000!/20! \cdot (1000 - 20)! = 10^{41}$. If the entire number of neurons in the neuronal network of the brain was to be taken into account, the dimensionality of the computational space expressed as the number of distinct encodings would be order of magnitude $10^{10^{11.95}}$, or about 892 billion digits. By contrast, there are about 10^{21} stars in the observable universe. Even more so, the assumption that individual neurons act as independent computational elements where a summation function thresholds to an activation function, the way we constructed $f(V)$ as in Equation 1 and considered the size of the network as $|V|$, is a gross oversimplification. We know that each neuron is a functional network *onto* itself, with individual dendrites and dendritic compartments acting themselves as independent computational elements [20] [21]. [22].

Lastly, in biological brains, various forms of plasticity — adaptive structural changes in response to dynamic functional changes, which in turn reciprocally influence function — continually adjust the network’s topology, i.e., the formation and retraction of synapses, and the synaptic weights within the network. It is very possible that future artificial neural networks will also have some form of plasticity that allows them to adapt and structurally ‘morph’ in a dynamic way.

So, in fact, the human brain network is a hierarchical (i.e., nested) family of adaptive and dynamic networks across *many* functional scales of organization, from molecular to cellular, individual neurons, networks of neurons, to networks of brain regions. All of it resulting in an astronomically (literally) huge computational space.

Furthermore, there are also about another 86 billion non-neuronal cells in the brain. Of these, about 20% are astrocyte glial cells. Certain sub-types of astrocytes can both listen in and modulate neuronal signaling and information processing [23]. These cells form an independent network onto themselves while at the same time cross-talk with the neuronal network. From a computational perspective, very little is still known about how interactions between neurons and astrocytes result in functional and cognitive outputs. It is one of the most exciting topics in systems and computational neuroscience. As if this was not all enough, the brain’s complexity is not just due to the sheer number of connections and size of its computational space but how these connections and resultant encodings are dynamically reconfigured and modulated. Understanding neuromodulation is an entire research effort in itself.

It is not clear how best to approach this unimaginable degree of combinatorial complexity and make sense of the information being produced. However, it is almost certain that purely classical computational methods alone will not be feasible. It will take significant continued theoretical and hardware advances in quantum computing and significant ingenuity regarding how problems are structured to make any sense of or progress for combinatorial and computational spaces approaching even a fraction of these scales. But for certain key questions and problems quantum computing could hold the key to truly understanding the functioning of the brain and artificial neural networks in ways we do not understand today.

At least for the foreseeable future, it is likely that classical-quantum computing hybrid models may be most successful for simulating and studying the computational space of neural networks. At present, there is a continued ‘leap-frogging’ between problems in combinatorics that quantum computers may be able to solve that are outside of what is possible for classical computers, contrasted with continued advances in classical algorithms that surprise the field and can (theoretically) perform as well as quantum algorithms for certain kinds of problems [24] [21]. And at least at certain scales, current efforts using specialized (classical) neuromorphic hardware are now approaching simulation sizes that are relatively comparable to the human brain [25]. The development of specialized neuromorphic computing specifically for carrying out large-scale simulations of the brain has a very long and rich history [26], [27]. But because, as we argued in the Introduction, simulations alone will not be enough to understand the brain or the deep inner workings of artificial neural networks, appropriately asking the right questions to probe this huge combinatorial and computational space is where we suggest quantum computing has a unique and complementary role to play compared to classical methods alone.

6.2 Practical Considerations and Applications of This and Future Related Work

One potential practical application of the work we present here and possibly related work is the iterative numerical experimentation and testing of different parameter conditions on the future state of an evolved network at the time T_o at which it is queried. The structure and form of synaptic models, internal dynamic models of neurons, membrane potential parameters, and network geometry (in the case of biological neural networks) or variations of activation functions and network architectures (in the case of artificial neural networks) could be systematically investigated without needing to brute force evolve the network dynamics *de novo*. This could be useful as the complexity of the local dynamic rules that govern global network dynamics increase and become more neurobiologically realistic [12] [13]. Beyond the proof of concept solution we developed here, it is interesting to consider the use or development of other quantum algorithms that could more efficiently probe this and related network dynamics problems.

One important consideration is the interplay between computational dimensionality and the physical constraints brain networks are necessarily subject to and (in contrast to artificial neural networks) seem to take advantage of that might provide a new perspective on brain algorithms [12], [13] [28] [14].[15] [16]. Other types of physical considerations may be equally relevant in different ways. This is important because it is almost certain that we have yet to fully discover the right mathematical ideas and descriptions that implement the biological brain’s algorithms.

A particularly important research direction for quantum computing applications to neural networks is the development of quantum algorithms and methods that support computing the time evolution of neural network dynamic models. Successfully doing so in a practical setting at scale is exceedingly challenging both from a theoretical and engineering perspective but of immense consequence. While several algorithmic approaches for the time evolution of quantum computational systems have been discussed [29], [30], it is not at all obvious how to set up and construct the necessary mathematical conditions to compute specific (local) instances of internal neuronal dynamics and then allow them to temporally (globally) evolve on a network.

This effort will likely need to use known quantum algorithms and methods creatively and possibly develop new ones. Any resultant methods would almost certainly have important applications for artificial quantum neural networks and machine learning, as well as neuroscience. As the quantum internet is built, the ability to study, interrogate, probe, and understand quantum networks may necessitate models and computing capabilities that are inherently quantum.

Finally, there is inherent value in thinking about how to structure a question or problem about the dynamic behavior or functionality of neural networks within the context of a quantum computing framework beyond any eventual actual quantum implementation. The theoretical considerations and necessary structure that formulating such problems in a quantum algorithmic context necessitate is an interesting exercise in its own right that provides a unique perspective on the interplay between the dimensionality and computational space of neural networks versus the algorithmic rules that operate on that space. Beyond quantum computing, it is interesting to think about how such a change in perspective may motivate new models and analysis approaches for neural networks that are of relevance to theoretical advances or classical computing implementations that can complement and work alongside quantum computing advances.

Acknowledgments

The author would like to thank Robert Treuhaft, with whom he had some of the earliest conversations related to quantum computing and the brain. And to Eric Traut and Rita J. King for many insightful discussions and feedback.

References

- [1] Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal of Computation* **26**, 1484–1509 (1997).
- [2] Vierre, E., Geraci, J. & Silva, G. A. *Quantum computing for neuroscience and neuroscience for quantum computing* (In: Convergence: Artificial Intelligence and Quantum Computing, 2nd edition, in press).
- [3] Huang, D., Wang, M., Wang, J. & Yan, J. A survey of quantum computing hybrid applications with brain-computer interfaces. *Cognitive Robotics* <https://doi.org/10.1016/j.cogr.2022.07.002> (2022).
- [4] Zhao, R. & Wang, S. A review of quantum neural networks: Methods, models, dilemma. *arXiv* arXiv:2109.01840 (2021).
- [5] Silva, G. A. Large-scale simulations of the brain may need to wait for quantum computers. *Forbes* <https://www.forbes.com/sites/gabrielasilva/2021/09/02/large-scale-simulations-of-the-brain-may-need-to-wait-for-quantum-computers/?sh=3934b5b07254> (2021).
- [6] Miranda, E. *et al.* An approach to interfacing the brain with quantum computers: Practical steps and caveats. *arXiv* arXiv:2201.00817 (2022).
- [7] Swan, M., Santos, R. P. D. & Lebedev, M. *Quantum Computing for the Brain* (WSPC, 2022).
- [8] Hameroff, S. R. Conscious events as orchestrated space-time selections. *Journal of Consciousness Studies* **3**, 36–53 (1996).
- [9] Hameroff, S. R. Quantum computation in brain microtubules? the penrose-hameroff 'orch or' model of consciousness. *Philosophical Transactions of the Royal Society of Lond* **622**, 1869–1896 (1998).

- [10] Kerskens, C. M. & Perez, D. L. Experimental indications of non-classical brain functions. *Journal of Physics Communications* 105001 (2022).
- [11] Documentary follows implosion of billion-euro brain project. *Nature* <https://www.nature.com/articles/d41586-020-03462-3>.
- [12] Buibas, M. & Silva, G. A. A framework for simulating and estimating the state and functional topology of complex dynamic geometric networks. *Neural Computation* **23**, 183–214 (2010).
- [13] Silva, G. A. The effect of signaling latencies and node refractory states on the dynamics of networks. *Neural Computation* **31**, 2492–2522 (2019).
- [14] Puppo, F., George, V. K. & Silva, G. A. An optimized function-structure design principle underlies efficient signaling dynamics in neurons. *Nature Scientific Reports* **8**, 10460 (2018).
- [15] George, V. K., Puppo, F. & Silva, G. A. Computing temporal sequences associated with dynamic patterns on the c. elegans connectome. *Frontiers in systems neuroscience* doi: 10.3389/fnsys.2021.564124 (2021).
- [16] George, V. K. *et al.* Learning without gradient descent encoded by the dynamics of a neurobiological network. *International Conference on Machine Learning (ICLR) Brain2AI* <https://arxiv.org/abs/2103.08878> (2021).
- [17] Herculano-Houzel, S. The human brain in numbers: A linearly scaled-up primate brain. *Frontiers in Human Neuroscience* doi: 10.3389/neuro.09.031.2009 (2009).
- [18] Azevedo, F. A., Carvalho, L. R., Grinberg, L. T., Marcelo, J. & Ferretti, R. E. Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *Journal of Comparative Anatomy* doi: 10.1002/cne.21974 (2009).
- [19] DeFelipe, J., Marco, P., Busturia, I. & Merchán-Pérez, A. Estimation of the number of synapses in the cerebral cortex: methodological considerations. *Cerebral Cortex* **9**, 722–732 (1999).
- [20] Fisek, M. & Hausser, M. The computing dendrite: From structure to function. *Springer* DOI 10.1007/978-1-4614-8094-5 (2013).
- [21] Tindall, J., Fishman, M., Stoudenmire, E. M. & Sels, D. Efficient tensor network simulation of IBM’s eagle kicked ising experiment. *PRX Quantum* **5**, 010308 (2024).
- [22] Eyal, G. *et al.* Human cortical pyramidal neurons: From spines to spikes via models. *Frontiers in Cellular Neuroscience* doi.org/10.3389/fncel.2018.00181 (2018).
- [23] de Ceglia, R., Ledonne, A., Litvin, D. G. & *et al.* Specialized astrocytes mediate glutamatergic gliotransmission in the CNS. *Nature* **622**, 120–129 (2023).
- [24] Pirnay, N., Ulitzsch, V., Wilde, F., Eisert, J. & Seifert, J.-P. An in-principle super polynomial quantum advantage for approximating combinatorial optimization problems via computational learning memory. *Science Advances* **10**, doi: 10.1126/sciadv.adj5170 (2024).
- [25] World first supercomputer capable of brain-scale simulation being built at Western Sydney University. *Western Sydney University* <https://www.westernsydney.edu.au/icns>.
- [26] Roy, K., Jaiswal, A. & Panda, P. Towards spike-based machine intelligence with neuromorphic computing. *Nature* **575**, 607–617 (2019).

- [27] Schuller, I. K. & Committee), R. S. O. Neuromorphic computing: From materials to systems architecture. *Department of Energy* <https://www.osti.gov/servlets/purl/1283147> (2015).
- [28] Achterberg, J., Akarca, D., Strouse, D., Duncan, J. & Astle, D. E. Spatially embedded recurrent neural networks reveal widespread links between structural and functional neuroscience fundings. *Nature Machine Learning* **5**, 1369–1381 (2023).
- [29] Venegas-Andraca, S. E. Quantum walks: A comprehensive review. *arXiv* arXiv:1201.4780 (2012).
- [30] Kendon, V. Quantum computing using continuous-time evolution. *Interface Focus* **10**, arXiv:20190143.4780 (2020).