

PAPER • OPEN ACCESS

## Quantum computation with machine-learning-controlled quantum stuff

To cite this article: Lucien Hardy and Adam G M Lewis 2021 *Mach. Learn.: Sci. Technol.* **2** 015008

View the [article online](#) for updates and enhancements.

# Quantum computation with machine-learning-controlled quantum stuff

**OPEN ACCESS**RECEIVED  
24 February 2020REVISED  
7 August 2020ACCEPTED FOR PUBLICATION  
24 August 2020PUBLISHED  
1 December 2020

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Lucien Hardy and Adam G M Lewis

Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario N2L 2Y5, Canada

E-mail: [alewis@perimeterinstitute.ca](mailto:alewis@perimeterinstitute.ca)**Keywords:** quantum computing, quantum machine learning, quantum tomography, quantum information theory, quantum foundations**Abstract**

We formulate the control over quantum matter, so as to perform arbitrary quantum computation, as an optimization problem. We then provide a schematic machine learning algorithm for its solution. Imagine a long strip of ‘quantum stuff’, endowed with certain assumed physical properties, and equipped with regularly spaced wires to provide input settings and to read off outcomes. After showing how the corresponding map from settings to outcomes can be construed as a quantum circuit, we provide a machine learning framework to tomographically ‘learn’ which settings implement the members of a universal gate set. To that end, we devise a loss function measuring how badly a proposed encoding has failed to implement a given circuit, and prove the existence of ‘tomographically complete’ circuit sets: should a given encoding minimize the loss function for each member of such a set, it also will for an arbitrary circuit. At optimum, arbitrary quantum gates, and thus arbitrary quantum programs, can be implemented using the stuff.

**1. Introduction**

Imagine we have some ‘stuff’ with quantum properties. Can we use it as a quantum computer? Join us in picturing a long strip of Plasticine-like stuff, whose unknown properties are accessible only via regularly spaced setting and outcome wires. By appropriate choices of settings, followed by correct interpretations of outcomes, is it possible to implement quantum computation?

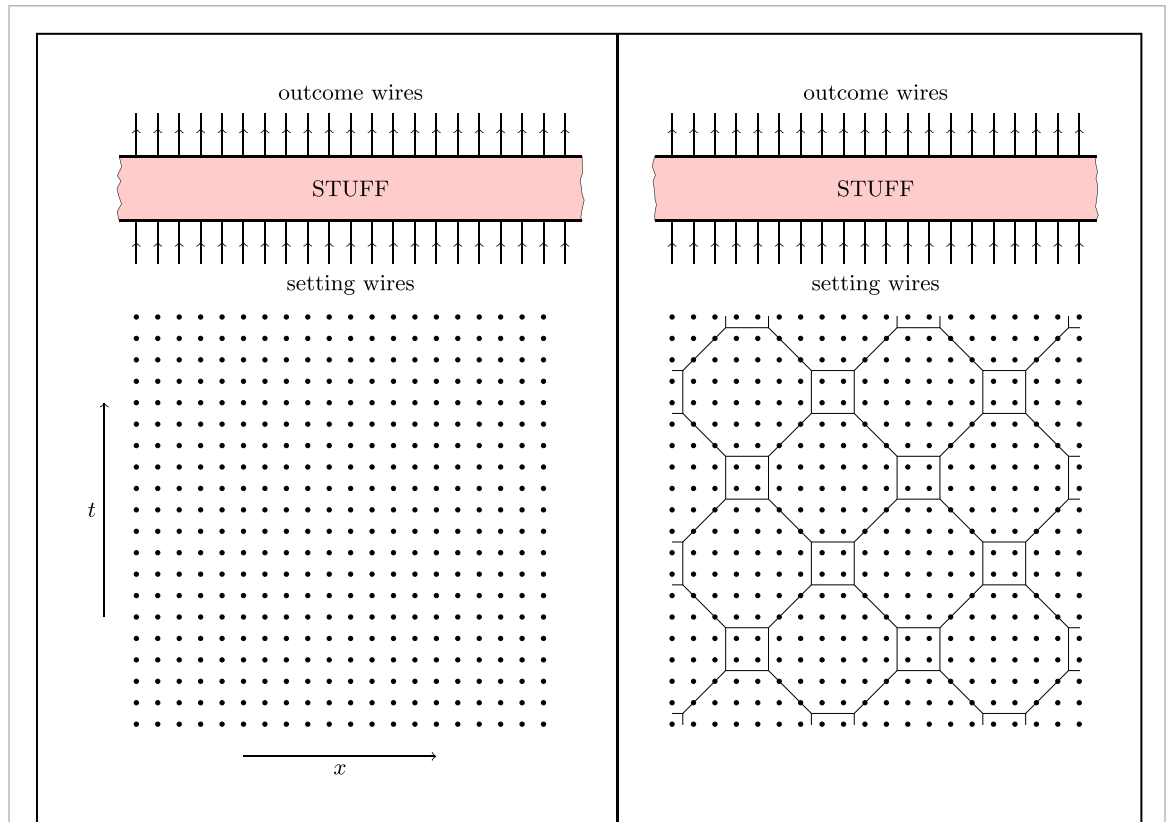
The answer is clearly yes in at least some cases. Most obviously, it is clearly yes if the stuff secretly contains matter known in advance to be capable of quantum computation. Suppose, for example, that the setting wires were to manipulate the local magnetic field and the timing of measurements, with results transmitted by the outcome wires. Any given quantum program could then be implemented, given knowledge of the details.

We consider, as a somewhat lopsided means to achieve quantum computation, the determination of these details. Normally, one would first decide upon a fixed such correspondence between logical and physical operations – Boolean logic to transistor states, say – and then engineer a computer to respect it. Here, we discuss the reverse task, that of mapping arbitrary quantum computation onto the fixed physics of initially uncharacterized matter: stuff.

Our iterative approach to this problem begins with a hypothesized map between physical and logical operations. For example, we might guess that feeding first 2 and then 3 to the first input wire implements the quantum CNOT gate. This guess is likely to be wrong, but, given a few assumptions about the stuff’s internal dynamics, we will illustrate how to determine just *how* wrong.

In principle, one could then converge towards a correct guess by gradient descent, applied to a bespoke set of neural networks. In practice the optimization is likely to be quite difficult, and to depend upon additional assumptions about the stuff than we impose here. In this work, we formulate the problem and a schematic solution in general terms, to be iterated upon later by the study of additional assumptions, algorithmic improvements, and empirical feedback.

The essential problem we are concerned with, the determination of the unknown quantum dynamics of a black box system (in our case, the arbitrary length strip of stuff), might be viewed as a slight generalization of so-called ‘quantum process tomography’ [1–6]. We are aware of two major differences. First, we assume no ability to interact with the quantum stuff except via the classical settings and outcomes on the wires. In



**Figure 1.** On the left we see a section of a length of stuff with input and output wires placed at regular intervals. A time-gated sequence of inputs is fed into the input wires and a similar time-gated sequence of outputs is read off the wires. We represent the input/output at position  $x$  and time  $t$  by a dot as shown. On the right we see the same figure with the dots divided up into octagons and squares. Points lying on the boundary between octagons are assigned to the upper octagon.

particular, we do not presume advance knowledge of how to prepare or to measure states. This echoes previous work on randomized gate benchmarking [7–10], self-consistent quantum process tomography [11, 12], and especially operational process tomography [13].

Second, we will never explicitly represent the full dynamics of the stuff, but only the ability to map chosen quantum circuits onto it. One might presume the weights of any successfully optimized neural networks would need to form at least a partial such representations, but we make no effort to characterize it. This situates our work within the rapidly developing field of ‘quantum machine learning’ [14, 15], and, more specifically, within the subfield concerned with machine-learning-based quantum tomography [16–21].


In section 2 to follow, we introduce the essential ideas required to construe operations upon quantum stuff as a user-defined circuit mapped onto spacetime. These ideas are then refined and formulated precisely in section 4. The iterative process of determining which operations correspond to which circuit, or ‘bootstrap tomography’, is formulated in section 5. Finally, in section 7, we sketch out a machine learning framework to realize this tomographic process.

## 2. Quantum stuff and computation

Picture a strip length- $L$  strip of stuff. Fix two wires at each site  $x_l = l\Delta X$ ,  $l = (0, 1, 2, \dots)$ , one to send in classical settings, and another to read off classical outcomes. Binning the duration of these interactions into discrete time intervals centred at times  $t_n = n\Delta T$ ,  $n = (0, 1, 2, \dots)$  defines a discrete set of spacetime points,  $\mathbf{x} = (x, t) = (l\Delta x, n\Delta t)$ , each labelling a setting-outcome pair. We refer this structure of setting-outcome pairs at discrete events as the *computational lattice*.

The physical dynamics of the stuff might be viewed as a map from sequences of settings to probability distributions over sequences of outcomes. To use the stuff as a quantum computer, we need to know the pertinent aspects of that mapping, construed as one between known mathematical objects. This knowledge can be usefully decomposed into that of two functions: an *encoder* to translate a given program – the ‘logical input’ – into physical settings, and a *decoder* to translate probability distributions over physical outcomes into mathematical objects, the ‘logical output’. The encoder/decoder pair has functioned correctly if and only

**Table 1.** A set of gates that is universal with respect to the given lattice.  $H$  stands for the Hadamard gate,  $P$  stands for the phase gate which acts as  $\begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$ , and  $R$  stands for the  $\frac{\pi}{8}$  gate on the left which acts as  $\begin{pmatrix} 1 & 0 \\ 0 & e^{\frac{\pi}{8}} \end{pmatrix}$ . We divide the set into three ‘subsets’ to emphasize the distinction from a traditional UGS. Subset 1 consists of the unitaries in a UGS. Subset 2 is formed from gates that move information without operating upon it. Subset 3 consists of projective measurements with a defined result. Note that universal sets could easily be devised from POVMs mixing unitary and ‘measurement’ behaviour, and which thus would not respect this classification.



if the emitted logical output is, accounting for the probabilistic nature of the physical outcomes, indeed that dictated by the logical input.

Suppose, for example, we wish to perform the computation  $q_{out} = 2q_{in}$ . Then the encoder must map from the logical input ‘ $2 \times q_{in}$ ’ to some sequence of physical settings. The decoder, in turn, must map the corresponding sequence of physical outcomes into the logical output  $q_{out}$ . The encoder/decoder pair is correct, in this case, when the emitted  $q_{out}$  approaches  $2q_{in}$ , having smoothed out any probabilistic fluctuations.

One might reasonably expect the discovery of a correct encoder/decoder pair to be a fairly daunting task. However, in this case and, as we will see, in general, checking whether a given encoder/decoder pair is correct is quite simple. Furthermore, one can easily construct a smooth measure, or *loss function*, of how far from correct a given encoder/decoder pair is: the RMS error between  $q_{out}$  and  $2q_{in}$ , for example, or some monotonic function of it<sup>1</sup>. A correct encoder/decoder pair will minimize this loss function.

Later, we will detail a means to accomplish this functional minimization automatically on a classical computer, via neural network parameterizations of the encoder and decoder. It is again to be emphasized that the protocol we present is intended as schematic, with the practical solution of the minimization problem left for further refinement in light of additional assumptions placed upon the stuff.

### 3. Gates, circuits, and tessellations

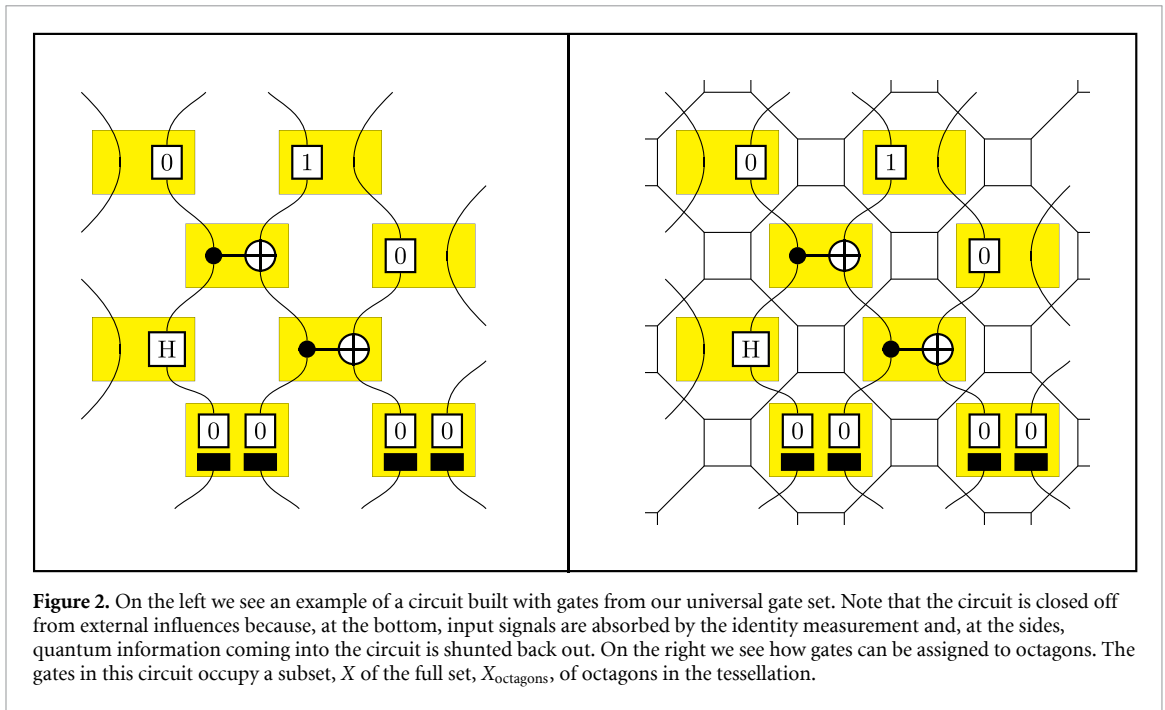
We would like to implement arbitrary quantum computations, not just multiplication by 2. It is well known that arbitrary quantum computations can be decomposed into elementary quantum gates forming a ‘universal gate set’ (UGS). It would thus be sufficient for our encoder/decoder pair to correctly implement each member of some UGS.

Quantum gates, however, map between quantum states, not classical information. By assumption, we do not have direct access to the internal quantum dynamics of the stuff. We will thus instead concern ourselves with quantum *circuits*, wiringtogether of quantum gates that are entirely characterized by classical settings and outcomes.

Any quantum computation may be expressed as a quantum circuit, and doing so further refines the formulation of our task. The logical input thus becomes the classical settings and gate labels defining a particular quantum circuit. The logical output is that circuit’s theoretical outcome distribution. The encoder/decoder pair has functioned correctly if the emitted outcome distribution matches the theoretical one. The loss function can be any of various distance metrics between those distributions.

Later we will demonstrate that, given a few physical assumptions about the stuff, there exist ‘tomographically complete’ sets of quantum circuits. If an encoder/decoder pair, acting upon stuff with the assumed properties, correctly implements all the circuits in a tomographically complete circuit set, it correctly implements all the gates in a UGS. The combination of encoder, stuff, and decoder in that case forms a universal computer. The degree by which it fails to do so can be taken, most primitively, as the summed losses of all the circuits in the set.

<sup>1</sup>The loss function will, strictly speaking, be a functional of the encoder/decoder pair. These will eventually be represented as neural networks, however, parameterized a finite set of weights. The loss will be a plain old function of these.



We want the stuff not only to implement an arbitrary computation, but to do so via a restricted number of setting/outcome retrievals. To achieve this, we will require the encoder’s implementation of each individual gate to occupy only a finite volume of spacetime. Specifically, we group together events in the computational lattice  $\mathbf{x} = (l\Delta x, x\Delta t)$ , depicted in the left panel of figure 1, into a *tesselation* of octagons and squares, depicted in the right panel of figure 1.

Each octagon, or *tessel*, will be used to implement one of a universal set of two qubit gates. The qubits for these gates will, under this implementation, be input at the lower slanted edges, and output at the upper slanted edges. Octagons are convenient because they adjoin to form causal-diamond-like structures. The squares appear by geometric necessity and will implement a fixed ‘do nothing’ gate, different from the identity. We will eventually view the encoder and decoder as neural networks to be trained by a machine learning algorithm. One failure mode of this algorithm will occur if the tessellation encloses too few events per gate to implement the UGS properly. In that case, we can uniformly scale up the tessellation and try again.

We label each octagon by its midpoint,  $\mathbf{x}$ . Let  $X_{\text{octagons}}$  be the set of octagon positions we consider in some given tessellation over the length  $L$  and some time duration,  $T$ .

The two qubit gates situated on the octagons will form a regular lattice which we will call the *gate lattice*. A universal gate set (UGS) must include enough gates to do universal quantum computation *with respect to the gate lattice*. An example of a universal set of gates, complete with respect to this lattice is given in table 1. Each of the gates in the table is a two-qubit gate. The first subset of gates are the usual gates included in a UGS adapted for the lattice.

Because we assume control over the stuff only through the classical settings, we need to include two more subsets of gates not usually mentioned as part of a UGS. The second subset are the identity and swap gates which allow us to transport qubits around the circuit. The third subset of gates are the preparation and measurement gates. These gates are non-unitary. The first gate in subset 3 performs the identity measurement on the incoming qubits and then prepares two qubits, each in the 0 state. The second gate in subset 3 projects the left qubit onto the 0 basis while leaving the right qubit unchanged. We have similar notation for the other gates. Unlike the other gates in the table, these measurement gates have outcomes associated with them.

It is necessary to include the second and third subsets of gates in our UGS because we need to train our stuff to implement them the same as any other gate. They do not come for free.

We can connect gates together to form what we will call a *fragment* (denoted by  $\mathcal{F}$ ). Whenever we build a fragment, there must be some gates having one or two open inputs. We will call such gates *initial gates*. A circuit,  $\mathcal{C}$ , is a special case of a fragment for which all initial gates are of the type that ignore any open inputs. The gates in the table provide two ways to do this. First, the preparation gate in subset 3 simply absorbs incoming quantum systems so they do not affect probabilities for the circuit. Second, any gates that include an identity map can be used to shunt quantum information coming through an open input back out through an open output so probabilities for the circuit are not affected. The circuit in figure 2 contains examples of

both types. We can calculate a probability for a circuit using the rules of quantum theory. A fragment that is not a circuit is subject to outside influences and so will not necessarily have a probability associated with it.

We will choose some particular UGS, call it  $G$ , to proceed. It does not have to be the one described here, and there might UGS's that are more suited to this project. However, the gate set must have elements that enable us to close off in the manner just described. There must also be some gates with outcomes, so we can read off the results of the computation.

#### 4. Building circuits

In this section, we describe how sequences of operations upon the stuff may be arranged in such a way to permit comparison with a given quantum circuit.

Thus, within each octagon of the tessellation, we will *encode* (by inputting an appropriate sequence of signals into the setting wires attached to the stuff) and *decode* (by selecting on an appropriate sequence of output signals from the outcome wires attached to the stuff) in an attempt to implement a putative element of the UGS,  $G$ . Initially we do not know what the appropriate encoding and decoding are. Thus we start with some initial choice and then, through the machine learning algorithm, train until we settle on encodings and decodings that minimize the loss function.

Let the putative encoding (decoding) for gate  $g \in G$  be denoted  $E^n[\mathbf{x}, g]$  ( $D^n[\mathbf{x}, g]$ ) for the  $n$ th iterative step in the training for the octagon at  $\mathbf{x}$ . We write  $Y^n[\mathbf{x}, g] = (E^n[\mathbf{x}, g], D^n[\mathbf{x}, g])$ . We are, then, admitting the possibility that the same gate might require different encodings and decodings in different regions of spacetime. One could imagine eventually adjusting the machine learning algorithm to be introduced to exploit any assumed homogeneity, but we will not explore this issue further. At the  $n$ th step we will have a particular encoding and decoding scheme for every gate for every octagon:

$$\{(\mathbf{x}, Y^n[\mathbf{x}, g]) : \text{for all } \mathbf{x} \in X_{\text{octagons}}, \text{ for all } g \in G\}. \quad (1)$$

This specifies an encoding for every element of  $G$  for every octagon. We will attempt to implement various quantum circuits using this encoding. Then, using the machine learning algorithm trained upon the resulting empirical information, we will iterate the encoding and decoding scheme, obtaining a new one to be used during the  $(n+1)$ th step.

We consider a fragment,  $\mathcal{F}$  (as described in section 3), made from the gates,  $g$ , in our UGS with locations assigned to the positions of some of the octagons in some octagon square tessellation. The fragment  $\mathcal{F}$  is specified by

$$\mathcal{F} = \{(\mathbf{x}, g) : \text{for all } \mathbf{x} \in X \subseteq X_{\text{octagons}}\}$$

where  $X$  is the set of octagon positions in the tessellation which contain a gate. A circuit,  $\mathcal{C}$ , is a special case of a fragment in which every initial gate is of the type that ignores any open inputs into it. The attempted implementation of circuit  $\mathcal{C}$  during the  $n$ th step is now given by

$$Y^n[\mathcal{C}] = \{(\mathbf{x}, Y^n[\mathbf{x}, g]) : \text{for all } \mathbf{x} \in X\} \quad (2)$$

in the octagons. Actually, this is not entirely sufficient, since we must also specify what happens operationally in regions of spacetime outside the given octagons (tessels). We will return to this point at the end of this section. Each gate  $g$  in  $\mathcal{C}$  maps quantum information to classical outcomes with various probabilities. The full circuit, however, fixes this internal quantum information (by including state preparation and measurement 'gates', for example), and is thus characterized by a single probability  $p_{\mathcal{C}}$  for a specific set of *classical* outcomes to be observed. Quantum mechanics can be formulated [22] as an assignment of a  $p_{\mathcal{C}}$  to every possible circuit  $\mathcal{C}$ .

In the next section, we will show that under certain assumptions, the converse is also true. That is, if each circuit within a *tomographically complete circuit set* indeed occurs with its predicted  $p_{\mathcal{C}}$ , the underlying *operational* map is indeed that specified by the relevant sequence of gates.

In order to certify that the computation within the stuff is indeed quantum mechanical, we thus seek encoder/decoder pairs that perform this same assignment via the stuff. That is, the theoretical map, from circuit settings and gate labels  $\mathcal{C}$  to outcome probabilities  $p_{\mathcal{C}}$ , is also the operational one, from the tessellation of  $\mathcal{C}$  and a given  $Y^n[\mathcal{C}]$  to the *observed* outcome probability  $p_{\mathcal{C}}^n$ .

We can then define one of several loss functions measuring how closely  $Y^n[\mathcal{C}]$  indeed implements some  $\mathcal{C}$ ,

$$L[Y^n[\mathcal{C}], \mathcal{C}] = \text{err}(p_{\mathcal{C}}^n, p_{\mathcal{C}}), \quad (3)$$

where  $\text{err}(x, y)$  is some convenient positive function with a global minimum at  $x = y$ , for example  $\text{err} = |x^2 - y^2|$ .

Now suppose we have fixed a set  $\mathcal{S}$  of circuits that we wish to simultaneously implement, for example because they form a tomographically complete circuit set. A loss function over the full set can then be written as

$$\mathcal{L}[Y^n, \mathcal{S}] = \sum_{\mathcal{C} \in \mathcal{S}} L[Y^n[\mathcal{C}], \mathcal{C}]. \quad (4)$$

Since the internal quantum information of the stuff might depend upon occurrences outside the spacetime region we have identified with our circuit, for this to work in practice we must also specify what happens at the spacetime points,  $\bar{X}$ , that are not in the octagons labelled by  $\mathbf{x} \in X$ . This includes those in the squares and those that occur elsewhere.

One strategy out of potentially many is to choose a ‘null’ encoding  $E_0$  at each ‘external’ spacetime point, and to simply ignore signals on the outcome wires there, so that we need not specify any decoding. The null encoding might also be iterated (‘trained’) in order to, for example, appropriately ‘zero out’ quantum information in the external regions; in that case it would be denoted  $E_0^n$ . Similar considerations apply to the squares. We do not need to concern ourselves with the encoding for spacetime points in the future of the circuit,  $\mathcal{C}$ , via the causal assumption that influences cannot travel backwards in time.

## 5. Bootstrap tomography

Now let us discuss the construction of so-called ‘tomographically complete circuit sets’.

If the loss function,  $\mathcal{L}[Y^n, \mathcal{S}]$ , is minimized when summed over a big enough set of test circuits, we would like it to be the case that any circuit gives the correct probabilities (to within some small error). In this section we state a theorem (proven in appendix A) that if the empirical probabilities,  $p_C^n$ , are exactly equal to the ideal probabilities,  $p_C$  (so the loss function (4) is minimized), for a certain set of circuits,  $\mathcal{S}_{\text{tom}}$ , then this is true for all circuits. The circuits in  $\mathcal{S}_{\text{tom}}$  have the property that they are bounded in size. We conjecture that a robust version of this theorem also holds - namely that the loss function (4) over  $\mathcal{S}_{\text{tom}}$  bounds the loss function for a fraction of  $1 - \kappa$  of circuits for any  $\kappa > 0$ .

We will need the important notion of a *bounded fragment (or circuit)*. This is one that fits inside a box of some constant size,  $\Delta L$  and  $\Delta T$ , where this box size does not increase in size as we increase  $L$  and  $T$ .

It is shown in appendix A that we can associate a vector,  $\mathbf{r}_X^n[\mathcal{F}]$ , with any fragment in region  $X$ . This vector linearly relates the given fragment to a tomographically complete set of fragments for the given region. In the case of quantum theory, the vector  $\mathbf{r}_X[\mathcal{F}]$  is linearly related to the superoperator associated with the fragment. The vectors,  $\mathbf{r}_X^n[\mathcal{F}]$ , are used to calculate the probabilities for circuits.

We can determine the vectors,  $\mathbf{r}_X[\mathcal{F}]$ , by doing tomography on a set of circuits  $\mathcal{F} \cup \bar{\mathcal{F}}$  for different  $\bar{\mathcal{F}}$ . We say we have *fragment tomography boundedness* if we can do tomography on a bounded set of fragments (pertaining to  $X$ ) by means of a bounded set of circuits.

In appendix A we define a *composition tomograph*,  $\mathbf{\Lambda}_X^{X_1, X_2, \dots}$  which tells us how to combine  $\mathbf{r}$  vectors pertaining to non-overlapping (though possibly adjacent) regions  $X_1, X_2, \dots$  to obtain the tomographic information pertaining to the region  $X = X_1 \cup X_2 \cup \dots$ .

We say we have *composition tomography boundedness* if the composition tomograph for a composite region formed from any number of regions can be determined from the composition tomographs for composites that fit inside bounded boxes. In other words, we can do the calculation for any circuit from calculations pertaining to smaller bounded parts of that circuit.

In appendix A the following theorem is proven.

**Theorem 1.** If we have fragment tomography boundedness and composition tomography boundedness, then there exists a bounded set of circuits  $\mathcal{S}_{\text{tom}}$  such that if, at iteration  $\tilde{n}$ , we have

$$p_C^{\tilde{n}} = p_C \quad \text{for all } C \in \mathcal{S}_{\text{tom}}$$

then  $p_C^{\tilde{n}} = p_C$  for any  $C \in \mathcal{S}_{\text{circuits}}$ .

$\mathcal{S}_{\text{tom}}$  is then called a ‘tomographically complete circuit set’.

We conjecture (but do not prove) that a robust version of this theorem holds. For this we need the notion of a *locally determined measure* over the set of all circuits of some given size. We consider ‘rectangular’ circuits formed from templates having  $N$  octagons in the horizontal (spatial) direction and  $T$  octagons in the vertical (temporal) direction. For each position in this template we choose gates randomly with some position



dependent weights. For positions on the boundary these weights are non-zero only for gates that close the circuit appropriately. For positions on the interior, we demand that all the weights are non-zero. In this way we generate the complete set of circuits for the given template, each having some probability which provides the locally determined measure.

**Conjecture 1.** If we have fragment tomography boundedness and composition tomography boundedness, then there exists a bounded set of circuits  $\mathcal{S}_\kappa^{\text{tom}}$  such that for all  $N$  and  $T$  and for any  $\kappa$  such that  $0 < \kappa \leq 1$ , at iteration  $\tilde{n}$  we have

$$\text{err}(p_{\mathcal{C}}^{\tilde{n}}, p_{\mathcal{C}}) \leq \mathcal{L}[Y^{\tilde{n}}, \mathcal{S}_\kappa^{\text{tom}}]f(N, T)$$

for a fraction  $1 - \kappa$  of circuits,  $\mathcal{C}$ , with respect to a locally determined measure over the set circuits that fit inside a rectangular template of dimensions  $N$  and  $T$ . Here  $f(N, T)$  is a finite order polynomial function of  $N$  and  $T$ .

The motivation for this conjecture is that the tomography process will fix the parameters in the gates to some error and this will bound the error for an arbitrary circuit (as quantified by  $\kappa$ ,  $\epsilon_\kappa$ , and  $f(N, T)$ ). If we enlarge the set,  $\mathcal{S}_\kappa^{\text{tom}}$ , we can expect to get a better bound on the error since we then collect more information. Conjecture 1 only guarantees that the fraction  $1 - \kappa$  of circuits agree with the ideal case to within some error. One possible problem is that it remains possible then that the interesting circuits (i.e. the ones used to do non-trivial quantum computation) fall into the set of circuits whose error is not bounded. This problem deserves more attention.

The causaloid framework [23–26] is used to prove Theorem 1. This framework was originally developed as a framework for modelling indefinite causal structure in the context of Quantum Gravity. These theorems mean that we can rely on the quantum stuff to implement an arbitrary circuit as long as the measured probabilities for circuits in  $\mathcal{S}_{\text{tom}}$  are close enough to the ideal probabilities calculated from Quantum Theory. This is good because it would not be practical to measure the probabilities for all  $\mathcal{C} \in \mathcal{S}_{\text{circuits}}$  since the rank of this set grows very rapidly with  $L$  and  $T$ .

For the two tomographic boundedness properties to hold requires in each case that (i) that a mathematical prerequisite holds and then (ii) that the physics of the stuff accords. The mathematical prerequisite is that the properties hold for ideal circuits constructed on the given lattice from the given universal gate set. To check this requires a mathematical calculation. Since we have a universal gate set, it is immediately clear that this mathematical prerequisite holds for fragment tomography boundedness (as we can use the UGS to construct a tomographically complete set of fragments that are bounded). We conjecture in appendix A that the mathematical prerequisite holds for composition tomography boundedness for any UGS. If the mathematical prerequisites hold, then we can consider whether the physical properties accord. This will, most likely, be settled through working with the stuff. However, we can always cook up situations in which the stuff fails to have the properties. For example, the properties will fail if the stuff has hidden signalling between far separated locations. For example, the bit of the stuff at  $\mathbf{x}$  might send a radio frequency signal to the bit of stuff at some  $\mathbf{x}'$  in the deep future outside any bounding box where this radio signal cannot tomographically probed by circuits that live in a bounding box.

While it would be satisfying to prove Conjecture 1 mathematically, it is possibly more useful to test it empirically. The conjecture does, in any case, necessarily involve assuming the boundedness properties which are, themselves, in need of empirical investigation. We test the conjecture empirically by determining the extent to which minimizing the loss function on various bounded sets of circuits allows us to reproduce the probabilities for sets of larger circuits.

## 6. Random circuit sampling

In the preceding sections we have illustrated that minimization of the loss function (4) by an implementation  $Y^m[\mathcal{C}]$ , acting upon the stuff with respect to a tomographically complete set of quantum circuits  $\mathcal{S}_{\text{tom}}$ , certifies that each individual gate in a UGS has also been correctly implemented by  $Y^m[\mathcal{C}]$ . The loss function (4) is inconvenient to use directly, however, because (a) computing a probability for every circuit in  $\mathcal{S}_{\text{tom}}$  for every training iteration is likely to be expensive and (b) the differing contributions to the overall gradients by the (possibly many) terms in the sum over  $\mathcal{S}_{\text{tom}}$  are expected to confound gradient descent optimization.

We will instead operate upon randomly constructed circuits, with each training iteration acting upon either a single example or perhaps a small minibatch of them. Random circuits can be constructed in various ways. For example, we could start by randomly assigning gates from the UGS to a few positions, resulting in a fragment. Next, we can identify the locations of open wires, and close the circuit by assigning random preparation measurement gates to each. Alternatively, one could begin with the preparation gates, randomly



add gates from the UGS capable of receiving inputs from those already present, and at some also-random point close the construction with a measurement gate, repeating the process from scratch if it has failed to deliver a circuit. Any such random constructions will be controlled by some parameters which dictate the average size of the circuits.

Consider a set (minibatch),  $\mathcal{S}_{\text{rand}}$ , of circuits generated randomly by some such technique. We make the following random sampling assumption

**Assump: Random circuit sampling.** There exists a bounded set of circuits  $\mathcal{S}_{\kappa}^{\text{rand}}$  such that for all  $N$  and  $T$  and for any  $\kappa$  such that  $0 < \kappa \leq 1$ , at iteration  $\tilde{n}$  we have

$$\text{err}(p_{\mathcal{C}}^{\tilde{n}}, p_{\mathcal{C}}) \leq \mathcal{L}[Y^{\tilde{n}}, \mathcal{S}_{\kappa}^{\text{rand}}]g(N, T) \quad (5)$$

for a fraction  $1 - \kappa$  of circuits,  $\mathcal{C}$ , with respect to a locally determined measure over the set circuits that fit inside a rectangular template of dimensions  $N$  and  $T$ . Here  $g(N, T)$  is a finite order polynomial function of  $N$  and  $T$ .

Whether such a sampling process actually exists is a physical hypothesis about the stuff. We motivate this hypothesis by illustrating a random sampling process upon a simple numerical circuit model in appendix B. Since a qualifying  $\mathcal{S}_{\text{rand}}$  is also a robust  $\mathcal{S}_{\text{tom}}$  in the sense of Conjecture 1, this model motivates Conjecture 1 as well.

A practical loss function is thus

$$\mathcal{L}_{\text{rand}}[Y^n, \mathcal{S}_{\text{tom}}] = \sum_{\mathcal{S}_{\text{rand}} \subset \mathcal{S}_{\text{tom}}} \sum_{\mathcal{C} \in \mathcal{S}_{\text{rand}}} L[Y^n[\mathcal{C}], \mathcal{C}], \quad (6)$$

where  $\mathcal{S}_{\text{rand}}$  is a randomly constructed subset of  $\mathcal{S}_{\text{tom}}$ . Since  $\mathcal{S}_{\text{rand}}$  would canonically contain only one or a few elements, the interior sum is much simpler than that of (4). In addition, by virtue of the random circuit sampling assumption, the outer sum can be treated by individually optimizing each of its summands. In other words, repeatedly optimizing  $L[Y^n[\mathcal{C}], \mathcal{C}]$  with respect to randomly constructed  $\mathcal{C}$  from  $\mathcal{S}_{\text{tom}}$  also optimizes (6) and, by assumption, (4).

## 7. Machine learning proposal

### 7.1. Neural network implementation of encoder/decoder

We have repeatedly alluded to our intention to view the encoder/decoder pair as neural networks to be trained by a machine learning algorithm. In this section we present a schematic realization of this process.

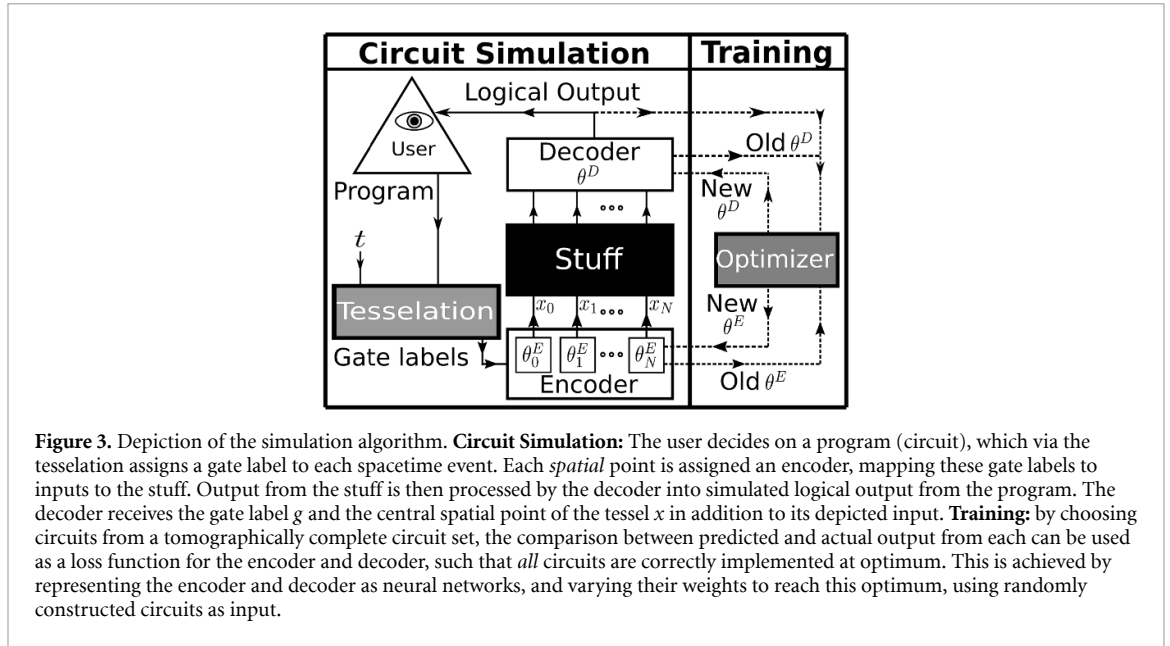
Recall that in the preceding sections we have formulated the problem of translating arbitrary quantum programs into operations upon the stuff as *bootstrap tomography*: the functional minimization of the loss function (4) (in practice (6)) obtained by comparing the observed and predicted outcome distributions with respect to the encoder/decoder pair at training iteration  $n$ ,  $E^n[\mathbf{x}, g]$  and  $D^n[\mathbf{x}, g]$ . Here  $g = g(t, x)$  is the gate label assigned to the point  $\mathbf{x} = (t, x)$  by the tessellation of the given circuit.

The problem of automatically varying a function to optimize a loss function is the central concern of *machine learning*, whose subfield of *deep learning* [27, 28] has become explosively popular in the past decade or so. Deep learning becomes especially useful relative to other approaches when the function to be optimized represents, or depends upon, a complicated probability distribution, as one might expect the encoder/decoder pair to.

A deep learning optimization, or ‘training’, begins by representing the target function as a linear composition of smooth *activation functions* called a *neural network*. Different *network structures* are defined by different arrangements of activation functions, with ‘deep’ learning being somewhat vaguely defined by its focus upon network structures formed of ‘many’ successive layers. The activation functions themselves are parameterized by their *weights*  $\theta$  so that different functions decompose into a given network structure by different choices of weights.

Every function theoretically has some neural network representation [28]. While this in itself is not especially impressive, the linearity and smoothness of a neural network’s activation functions permit functional derivatives with respect to the full network to be expressed as sums of partial derivatives of the weights, essentially via the chain rule. The gradient of the full function with respect to some loss can then be efficiently descended by descending those of the individual activation functions, a process known as *backpropagation* [28–30].

Precisely how backpropagation is best applied in a given situation is somewhat problem-dependent. We will save detailed consideration of this matter in future studies, when we implement bootstrap tomography on a small, classically simulated spin chain.




---

**Algorithm 1** Schematized optimization of the encoders and decoder.
 

---

- 1: **procedure** Optimize( $\{\theta_x^E\}, \theta^D$ )
  - 2: **for** number of training iterations **do**
  - 3:  $Circuit \leftarrow$  a randomly generated circuit.
  - 4:  $LogicalOutput \leftarrow$  a batch of  $m$  calls to SIMULATECIRCUIT( $\{\theta_x^E\}, \theta^D, circuit$ ).
  - 5:  $PredictedOutput \leftarrow m$  predicted circuit outcomes.
  - 6:  $Loss \leftarrow$  distance metric between  $LogicalOutput$  and  $PredictedOutput$  per (3).
  - 7: Vary  $\{\theta_x^E\}$  and  $\theta^D$  to minimize  $Loss$ , as dictated by the chosen optimization strategy.
  - 8: **end for**
  - 9: **end procedure**
- 

For now we will instead depict the problem at the more abstract level depicted in figure 3. We will denote the neural network representation of the decoder as  $\mathcal{D}[\theta^D; \mathbf{x}, g]$ , so that

$$D^n[\mathbf{x}, g] = \mathcal{D}[\theta^D; \mathbf{x}, g]. \quad (7)$$

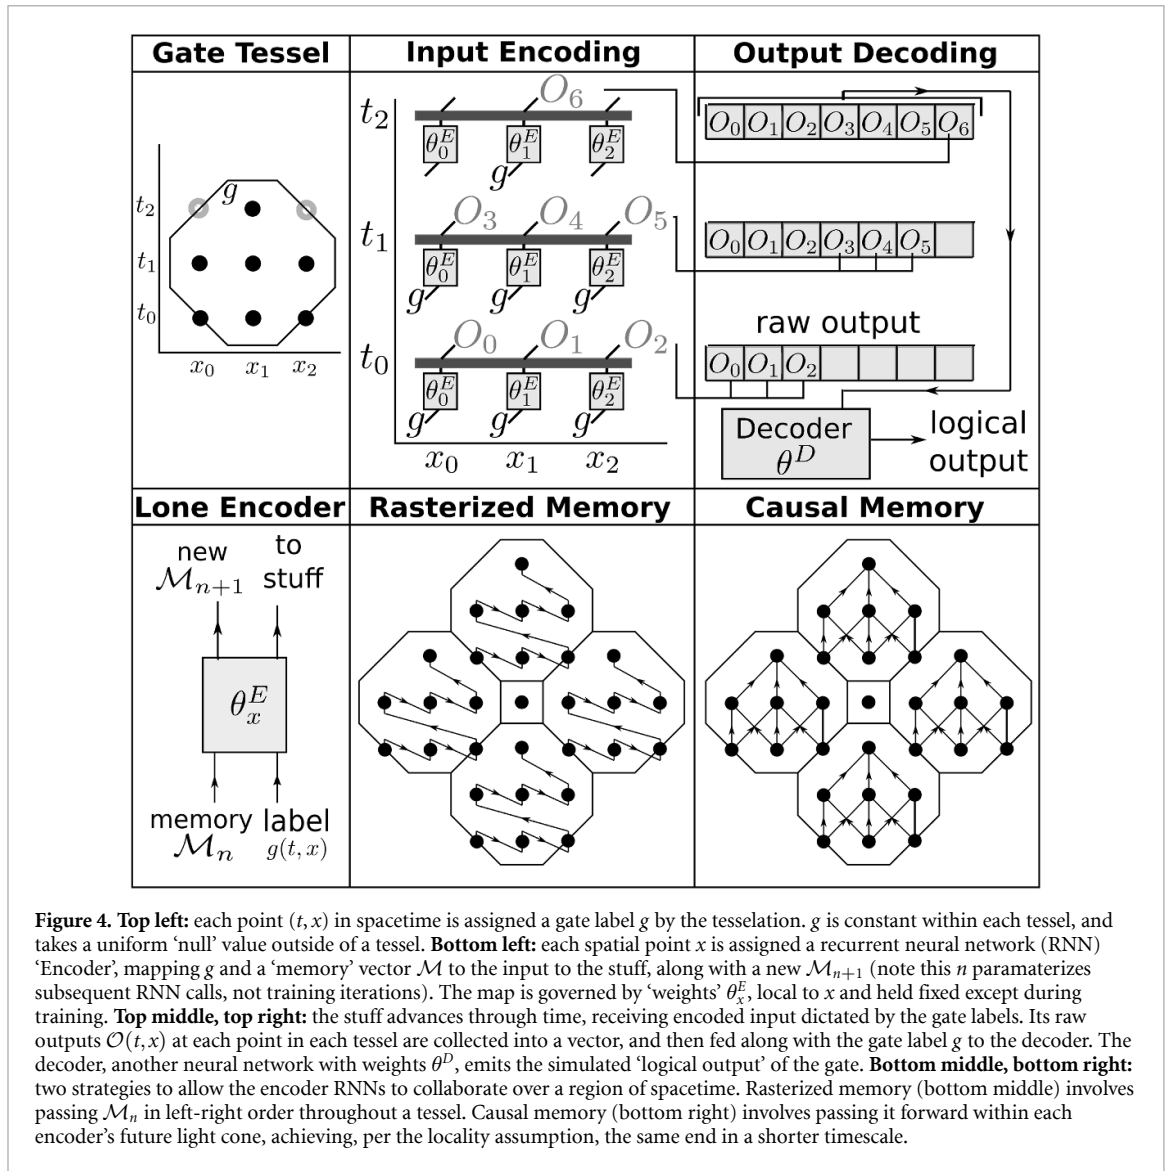
The network representation  $\mathcal{D}$  itself is fixed for all  $n$ . Training updates are instead implemented by varying the weights  $\theta^D$  between training iterations; they are otherwise fixed.

The encoder  $E^n[\mathbf{x}, g]$  controls the internal quantum information of the stuff, and thus unlike the decoder must interact with it in real time. It nevertheless needs to share information within tessels in order to track that quantum information's flow. This point will be elaborated upon in the upcoming subsection B. For now we briefly note that we handle this problem by representing  $E^n[\mathbf{x}, g]$  with a 'fleet' of 'recurrent' neural networks, one at each spatial point  $x$ . We denote the encoder weights at some  $x$  as  $\theta_x^E$ , and the full set of weights over all spatial points as  $\{\theta_x^E\}$ . Thus

$$E^n[\mathbf{x}, g] = \mathcal{E}[\{\theta_x^E\}, \mathcal{M}; \mathbf{x}, g]. \quad (8)$$

The network representation  $\mathcal{E}$  is again fixed for all training iterations  $n$ , with variation between iterations instead implemented by manipulations of the weights  $\{\theta_x^E\}$ . The output of the network additionally depends upon a *memory vector*  $\mathcal{M}$ . This vector is passed between different  $\theta_x^E$  within a tessel, allowing encoders following different constant- $x$  'worldlines' to communicate. We will elaborate upon this point in Subsection B. Let us now follow the logic of figure 3 in words. The user selects a program, which is mapped by the tessellation into a field of gate labels  $g(t, x)$ . Each gate label at  $x$  is passed along with the appropriate memory vector  $\mathcal{M}$  to the encoder at that same  $x$ ,  $\mathcal{E}[\theta_x^E, \mathcal{M}; \mathbf{x}, g]$ . This yields a raw input signal, to be sent to the setting wire at  $x$ . Once an entire tessel has been implemented, the corresponding 'raw' outcome settings from the stuff are passed into the decoder, which maps them into 'logical' output.

If the weights  $\theta^D$  and  $\{\theta_x^E\}$  are optimal, as indicated by the loss function (6) reaching its global minimum, the logical output may be interpreted as the correct result of the program. Otherwise, the gradients of the weights in the direction of decreasing (6) can be calculated, as by the 'optimizer' in figure 3,



**Figure 4.** **Top left:** each point  $(t, x)$  in spacetime is assigned a gate label  $g$  by the tessellation.  $g$  is constant within each tessel, and takes a uniform ‘null’ value outside of a tessel. **Bottom left:** each spatial point  $x$  is assigned a recurrent neural network (RNN) ‘Encoder’, mapping  $g$  and a ‘memory’ vector  $\mathcal{M}$  to the input to the stuff, along with a new  $\mathcal{M}_{n+1}$  (note this  $n$  parameterizes subsequent RNN calls, not training iterations). The map is governed by ‘weights’  $\theta_x^E$ , local to  $x$  and held fixed except during training. **Top middle, top right:** the stuff advances through time, receiving encoded input dictated by the gate labels. Its raw outputs  $\mathcal{O}(t, x)$  at each point in each tessel are collected into a vector, and then fed along with the gate label  $g$  to the decoder. The decoder, another neural network with weights  $\theta^D$ , emits the simulated ‘logical output’ of the gate. **Bottom middle, bottom right:** two strategies to allow the encoder RNNs to collaborate over a region of spacetime. Rasterized memory (bottom middle) involves passing  $\mathcal{M}_n$  in left-right order throughout a tessel. Causal memory (bottom right) involves passing it forward within each encoder’s future light cone, achieving, per the locality assumption, the same end in a shorter timescale.

and descended along to obtain new, better optimized weights. In machine learning parlance, this process is called ‘training’ the networks. Note the gradients of course cannot be directly backpropagated through the stuff, but the loss can still be descended per a variety of algorithm genres, such as reinforcement learning. The next iteration operates upon a new randomly constructed circuit, and the process is repeated until a desired convergence threshold is reached. Illustrative pseudocode is provided as Algorithm 1. From a technical perspective, of course, the bulk of the difficulty would reside in the choice of a practical optimization procedure, which in this work we do not consider.

### 7.2. RNN fleet implementation of encoder

As discussed previously, the encoder  $E^n[\mathbf{x}, g]$  and decoder  $D^n[\mathbf{x}, g]$  have differing relationships with the real time behaviour of the stuff, which suggests a particular network implementation for the encoder that we call an ‘RNN fleet’. We will elaborate upon this point here.

Figure 4 depicts several aspects of the behaviour of the encoder/decoder algorithm in spacetime. In the top left panel, we see a single tessel, assigned to a gate labelled  $g$ . The top middle panel depicts this same tessel, implemented as operations upon the encoders and thus upon the stuff. As the stuff proceeds through time, the gate label  $g$  is fed to the encoder at each  $(t, x)$  in the tessel. As we see in the top right panel, the output from the tessel is collected into a vector. Once all of it has been collected, it is sent to the decoder, which produces the logical output of the gate.

Since nothing in this procedure depends upon the logical output directly, the precise time at which the decoder  $D^n[\mathbf{x}, g]$  operates is not especially important, within reason. The encoder  $E^n[\mathbf{x}, g]$ , on the other hand, controls the internal quantum information of the stuff. The order in which its inputs are processed therefore critical. In addition, it must be synchronized within a tessel.

The problem of processing ‘time series’ with a specific temporal ordering occurs repeatedly in machine learning. The prototypical example is machine translation, since correct translations depend upon prior context. A genre of network structures known as ‘recurrent neural networks’, or RNNs [28, 29, 31, 32], are adapted to remember such context. In addition to their weights, which are fixed except during training, such networks maintain a ‘memory’ vector  $\mathcal{M}$ . The output of the RNN depends on the memory as well as upon the weights and the input. But unlike the weights, the memory is modified between successive function calls, allowing it to represent contextual information: *which* contextual information to record is in turn determined by the weights.

We thus implement the encoder  $E^n[\mathbf{x}, g]$  at each separate spatial point  $x$  as an RNN =  $\mathcal{E}[\theta_x^E, \mathcal{M}; \mathbf{x}, g]$ , as depicted in the bottom left panel of figure 4. Each  $\mathcal{E}[\theta_x^E, \mathcal{M}; \mathbf{x}, g]$  follows a particular constant- $x$  worldline through the tessellation, partly motivating the term ‘RNN fleet’.

Additional motivation for the term comes from the need of the various RNNs to be synchronized within a tessell. Thus, the memory vector  $\mathcal{M}_n$  is not simply passed forward along a worldline, but is instead shared within each tessell between networks of different weights. Consequently, we need to synchronize the processing of *several* time series at different spatial points.

The bottom middle and bottom right panels depict two strategies for synchronizing within tessells. We call the first and simplest strategy, in the bottom middle panel of 4, *rasterized memory*. In this paradigm, encoders  $\mathcal{E}[\theta_x^E, \mathcal{M}; \mathbf{x}, g]$  are called sequentially at each timestep from left to right. Starting from a fixed null value, memory is thus passed within the zig-zagging ‘spacelike’ lines depicted in the bottom middle panel of figure 4, which resemble the ‘rasterized’ beam path of a cathode ray tube television. Circuit simulation using the rasterized strategy is illustrated by the pseudocode of Algorithm 2.

The illusion of motion created by such televisions relies upon the time required for the beam to traverse the screen being much shorter than the processing time of the eye. The ability of the rasterized memory strategy to effectively synchronize encoders within a tessell correspondingly depends upon the processing time of the encoders being much shorter than that between simulation timesteps.

The *causal memory* strategy depicted in the bottom right of figure 4 relaxes this assumption. It is based on the assumption that the error incurred by sending encoder input out of sequence falls off with the spacetime distance between the disordered events.

Instead of passing memory sideways between every event in the tessell, we thus pass it forward in spacetime within a fixed ‘light cone’ of predetermined width  $s$ , unless doing so would cross a tessell. Given locality, this should synchronize just as well as the rasterized strategy, without putting impositions upon the relative timescales of the encoder and the simulations. Circuit simulation using the causal strategy is illustrated by the pseudocode of Algorithm 3.

The introduction of a concrete network structure raises the issue of its practicality. In the fully general case of completely inhomogeneous matter, a different set of weights is required for each gate at each spacetime point, which might be expected to fairly quickly exhaust computational resources. This is a symptom of the weak assumptions placed here on the stuff: temporally homogeneous stuff, for example, would require only one set of weights per gate (in the computationally complete set) per spatial point, and spatiotemporal homogeneity only one set of weights per gate. We leave the interesting question of what algorithmic economies can be derived from similar, but weaker, assumptions for future study.

---

**Algorithm 2** Circuit simulation using *rasterized* memory, passed sequentially between each point in each tessell.

---

```

1: procedure SimulateCircuitRasterizedMemory( $\{\theta_x^E\}, \theta^D, circuit$ )
2:  $g(t, x) \leftarrow$  the gate labels assigned to each spacetime point by the tessellation of circuit.
3:  $n_T \leftarrow$  a separate point counter for each tessell  $T$  with fixed  $g$ .
4:  $\mathcal{M}_{n_T}^T \leftarrow$  a separate memory vector for each  $T$ .
5: for  $t$  in the tessellation do
6:   for  $x$  in the tessellation do
7:     Determine which tessell  $T$  the point  $(t, x)$  inhabits.
8:     Retrieve this  $g, n_T$ , and  $\mathcal{M}_{n_T}^T$ .
9:      $InputSignal, \mathcal{M}_{n_T+1}^T \leftarrow$  output mapped from  $g$  and  $\mathcal{M}_{n_T}^T$  by encoder at  $x$  with weights  $\theta_x^E$ .
10:     $RawOutput[n_T] \leftarrow$  the corresponding output from stuff at  $x$ .
11:     $n_T \leftarrow n_T + 1$ .
12:   end for
13:   if this is the last point in a tessell then
14:      $LogicalOutput[T] \leftarrow$  output mapped from  $RawOutput, g$ , and  $x$  by decoder with weights  $\theta^D$ .
15:   end if
16:   end for
17: return  $LogicalOutput$ 
18: end procedure

```

---

---

**Algorithm 3** Circuit simulation using *causal* memory, passed within forward lightcones, confined by tessels.

---

```

1: procedure SimulateCircuitCausalMemory( $\{\theta_x^E\}, \theta^D, \text{circuit}$ , light cone width  $s$ .)
2:  $g(t, x) \leftarrow$  the gate labels assigned to each spacetime point by the tessellation of circuit.
3:  $n_T \leftarrow$  a separate point counter for each tessel  $T$  with fixed  $g$ .
4:  $\mathcal{M}_{t=0, x} \leftarrow$  a separate memory vector for each spatial point  $x$ .
5: for  $t$  in the tessellation do
6: for  $x$  in the tessellation do
7: Retrieve  $g(t, x)$  and  $\mathcal{M}_{t-s, x-s: x+s}$ , the memories between  $x-s$  and  $x+s$ .
8:  $\tilde{\mathcal{M}}_{t, x} \leftarrow$  pointwise multiplication of  $\mathcal{M}_{t-s, x-s: x+s}$  within the same tessel as  $(t, x)$ .
9:  $\text{InputSignal}, M_{t, x} \leftarrow$  output mapped from  $g$  and  $\tilde{\mathcal{M}}_{t, x}$  by encoder at  $x$  with weights  $\theta_x^E$ .
10:  $\text{RawOutput}[n_T] \leftarrow$  the corresponding output from stuff at  $x$ .
11:  $n_T \leftarrow n_T + 1$ .
12: end for
13: if this is the last point in a tessel then
14:  $\text{LogicalOutput}[T] \leftarrow$  output mapped from  $\text{RawOutput}$ ,  $g$ , and  $x$  by decoder with weights  $\theta^D$ .
15: end if
16: end for
17: return  $\text{LogicalOutput}$ 
18: end procedure

```

---

## 8. Final thoughts

In the preceding work, we introduced the concept of quantum stuff, Plasticene-like matter whose unknown internally-quantum dynamics are to be controlled by purely classical settings and monitored by purely classical outcomes. We next mapped quantum gates and circuits onto these settings and outcomes via a tessellated spacetime structure, and presented a loss function whose optimization certifies the correct implementation of these gates. Finally, we proposed a neural-network-based framework to optimize this loss function.

Numerous future directions are possible. Most obvious is the development of a concrete implementation. To this end one might consider adjustments of the framework yielding practical improvements. For example, as currently suggested, members of a tomographically complete circuit set would need to be run very many times in order to converge to their empirical probability, at each step of the iteration. We might consider iterating after each run of the circuit, or after some smaller finite number of runs. One way to proceed would be to consider tomographically complete circuit sets whose members have probability equal to zero. We can then start to run the circuits in the set. Every time we obtain the outcome associated with one of these circuits (which are supposed to have probability zero), we iterate. Even if the probability was nonzero, but less than  $\frac{1}{2}$ , we could consider batches of  $r$  runs of the circuit and iterate if  $s$  runs show a positive result for the circuit. For appropriate choices of  $r$  and  $s$ , the probability would be very small.

The bigger picture here is that we use machine learning implemented on a less powerful computer to train some stuff with unknown properties to implement a much more powerful computer. In the example above we have considered using a classical computer (implementing a neural network) to train quantum stuff to implement a quantum computer. The type of protocol we have described might work in other cases. For example, if we had some classical stuff that was complex in the right kind of way then it might be advantageous to use machine learning to train it to do classical computation that takes advantage of this complexity.

The viewpoint in the present paper is that the settings influence variables in the stuff (perhaps by tuning parameters in the Hamiltonian). The training process consists of finding the right settings to implement quantum computation. We could take a more radical viewpoint in which the training process corresponds to the physical assembly of a quantum computer by automated trial and error. This would be *self assembly quantum computation*. Self assembly is a rich subject. Examples of self assembly can be found at all scales [33]. One example is the self assembly of quantum dots [34]. According to [33], self assembly is ‘the autonomous organization of components into patterns or structures without human intervention’. However, if we were to implement a machine learning approach then we would have a kind of human intervention in the training process. This might lead to a system better suited to quantum computation. We might also consider self assembly classical computing. One might even imagine a sequence of increasingly powerful computers, each assembled in this way by its ancestor.



## Acknowledgment

A G M Lewis is supported by the Tensor Network Initiative at Perimeter Institute. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Research, Innovation and Science.

## Appendix A. Proof of theorem 1

In this appendix we provide definitions of fragment tomography boundedness and composition tomography boundedness and we prove Theorem 1. The idea is to consider tomography on fragments. We consider two types of tomography. First, we have *fragment tomography*, whereby we obtain a mathematical object ( $\mathbf{r}_X[\mathcal{F}]$  below) associated with the fragment. Second, we consider *composition tomography* whereby we obtain the rule for composing these mathematical objects for composite regions such as  $X_1 \cup X_2 \cup X_3$ . If the tomography boundedness properties hold then we only need to consider fragment tomography for fragments up to a certain size and composition tomography for composites up to a certain size. To do tomography on these fragments and composite fragments are completed into circuits. Hence it follows that we only need to consider circuits smaller than a certain size. This provides our set,  $\mathcal{S}_{\text{tom}} \subset \mathcal{S}_{\text{circuits}}$  of circuits. We will prove a theorem (Theorem 1) saying that if we obtain probabilities  $p_{\mathcal{C}}^n = p_{\mathcal{C}}$  for  $\mathcal{C} \in \mathcal{S}_{\text{tom}}$  then it follows that  $p_{\mathcal{C}}^n = p_{\mathcal{C}}$  for all circuits.

For a circuit,  $\mathcal{C} = \mathcal{F} \cup \bar{\mathcal{F}}$  we can always write the probability as

$$p_{\mathcal{F} \cup \bar{\mathcal{F}}}^n = \mathbf{r}_X^n[\mathcal{F}] \cdot \mathbf{p}_X^n[\bar{\mathcal{F}}] \quad (\text{A1})$$

where we define the ordered set

$$\mathbf{p}^n[\bar{\mathcal{F}}] = (p_{\mathcal{F}^k \cup \bar{\mathcal{F}}}^n : \text{for } k \in \Omega_X) \quad (\text{A2})$$

where  $k \in \Omega_X$  labels the elements,  $\mathcal{F}^k$ , of some tomographic set,  $\mathcal{T}_X^{\text{tom}} \subseteq \mathcal{T}_X$  of minimal possible rank. Here  $\mathcal{T}_X$  is the set of all possible fragments in  $X$ . In general, the choice of tomographic set,  $\mathcal{T}_X^{\text{tom}}$ , is not unique. We can always write (A1) because, in the worst case, we can choose  $\mathcal{T}_X^{\text{tom}} = \mathcal{T}_X$  and then the vector  $\mathbf{r}_X^n[\mathcal{F}]$  is just a list of 0's except with a 1 at position  $k$ . In general, however, there will be some linear relationships between these probabilities so that we can use a proper subset,  $\mathcal{T}_X^{\text{tom}} \subset \mathcal{T}_X$ . We can think of  $\mathbf{p}_X^n[\bar{\mathcal{F}}]$  as the *generalized state* prepared by  $\bar{\mathcal{F}}$  for region  $X$ . And we can think of  $\mathbf{r}_X^n[\mathcal{F}]$  as the *generalized effect* associated with fragment  $\mathcal{F}$  performed in region  $X$ . The choice of tomographic set,  $\mathcal{T}_X^{\text{tom}}$ , must be good for calculating the probability for any circuit in any region  $X \cup \bar{X}$  (that is for any  $\bar{X}$  associated with a fragment  $\bar{\mathcal{F}}$ ).

Using simple linear algebra [35] we can obtain the set  $\{\mathbf{r}_X^n[\mathcal{F}] : \text{for all } \mathcal{F} \in \mathcal{T}_X\}$  if we are given enough empirical information in the form of  $p_{\mathcal{C}}^n$  for  $\mathcal{C} \in \mathcal{S}_X^{\text{tom}}$ . The set,  $\mathcal{S}_X^{\text{tom}}$ , has to be big enough to make this possible—namely it has to generate  $\Omega_X$  linearly independent  $\mathbf{p}_X^n[\bar{\mathcal{F}}]$  vectors. In this case we will say  $\mathcal{S}_X^{\text{tom}}$  is *tomographically complete* for  $X$ . We can be sure this is true simply by choosing  $\mathcal{S}_X^{\text{tom}}$  to be the set of all circuits in  $\mathcal{S}_{\text{all circuits}}$  having elements at positions in  $X$ . This is not very useful however as this set grows very rapidly with  $L$  and  $T$ .

To obtain a more useful notion we define a *bounded set* of circuits or fragments as one for which each element fits inside a box with bounded spatial and temporal dimensions which do not scale with  $L$  and  $T$ . Consider the following property.

**Fragment tomography boundedness.** We say we have fragment tomography boundedness if, for any bounded set of fragments, there exists a bounded and tomographically complete set of circuits.

We can motivate the assumption that this property holds by finiteness and locality. First, note that the operationally accessed part of the Hilbert space associated with the inputs and outputs for any  $\mathcal{F} \in \mathcal{T}$  should be finite so only require a finite number of circuits for tomography. Furthermore, by locality we should be able to do tomography on this Hilbert space by means of circuits that are not too much bigger than the fragments.

Consider a composite region,  $X_1 \cup X_2$  (where  $X_1$  and  $X_2$  are disjoint). Then, for the circuit  $\mathcal{C} = \mathcal{F}_1 \cup \mathcal{F}_2 \cup \bar{\mathcal{F}}$ , we can write the probability as

$$p_{\mathcal{F}_1 \cup \mathcal{F}_2 \cup \bar{\mathcal{F}}}^n = \mathbf{r}_{X_1 \cup X_2}^n[\mathcal{F}_1 \cup \mathcal{F}_2] \cdot \mathbf{p}_{X_1 \cup X_2}^n[\bar{\mathcal{F}}] \quad (\text{A3})$$



where

$$\mathbf{p}_{X_1 \cup X_2}^n[\bar{\mathcal{F}}] = (p_{\mathcal{F}^{k_1} \cup \mathcal{F}^{k_2} \cup \bar{\mathcal{F}}}^n : \text{for } k \in \Omega_{X_1 \cup X_2}). \quad (\text{A4})$$

It can be shown [23] that for the composite region  $X_1 \cup X_2$  we can choose a tomographic set of fragments of the form  $\mathcal{F}^{k_1} \cup \mathcal{F}^{k_2}$  where

$$\Omega_{X_1 \cup X_2} \subseteq \Omega_{X_1} \times \Omega_{X_2} \quad (\text{A5})$$

and, correspondingly, that

$$\mathbf{r}_{X_1 \cup X_2}^n[\mathcal{F}_1 \cup \mathcal{F}_2] = \Lambda_{X_1 \cup X_2}^{X_1, X_2} (\mathbf{r}_{X_1}^n[\mathcal{F}_1] \otimes \mathbf{r}_{X_2}^n[\mathcal{F}_2]) \quad (\text{A6})$$

where  $\Lambda_{X_1 \cup X_2}^{X_1, X_2}$  (which we called a composition tomogram in section 5) linearly projects vectors,  $\mathbf{r}_{X_1}^n[\mathcal{F}] \otimes \mathbf{r}_{X_2}^n[\mathcal{F}]$  living in a space of dimension  $|\Omega_{X_1} \times \Omega_{X_2}|$  down to the vectors  $\mathbf{r}_{X_1 \cup X_2}^n[\mathcal{F}_1 \cup \mathcal{F}_2]$  living in a space of dimension  $|\Omega_{X_1 \cup X_2}|$ . If we do fragment tomography to obtain  $\mathbf{r}_{X_1}^n[\mathcal{F}_1]$ ,  $\mathbf{r}_{X_2}^n[\mathcal{F}_2]$ , and  $\mathbf{r}_{X_1 \cup X_2}^n[\mathcal{F}_1 \cup \mathcal{F}_2]$  then we can find an appropriate  $\Omega_{X_1 \cup X_2}$  set and solve for  $\Lambda_{X_1 \cup X_2}^{X_1, X_2}$ . This last step completes the composition tomography for  $X_1 \cup X_2$ .

Composition tomography extends in the obvious way to more than two regions. Thus, we can obtain  $\Lambda_{X_1 \cup X_2 \cup X_3}^{X_1, X_2, X_3}$ ,  $\Lambda_{X_1 \cup X_2 \cup X_3 \cup X_4}^{X_1, X_2, X_3, X_4}$ , and so on. It turns out that these  $\Lambda$ 's are related by mathematical identities if the  $\Omega$  sets associated with the  $X$ 's have certain relationships with one another (see section 23 of [23]). Consider the following property.

**Composition tomography boundedness.** We will say we have composition tomography boundedness for circuits formed on a given lattice using a given gate set,  $G$ , if we can obtain  $\Lambda_X^{X_1, X_2, \dots}$  for any  $X$  from mathematical identities concerning only  $\Lambda$ 's pertaining to boxes bounded by some constant size  $(\Delta L_{\text{bounded}}, \Delta T_{\text{bounded}})$  (which do not increase with  $L$  and  $T$ ).

Here we say that  $\Lambda_{X_1 \cup X_2 \cup \dots \cup X_M}^{X_1, X_2, \dots, X_M}$  pertains to a box bounded by  $(\Delta L_{\text{bounded}}, \Delta T_{\text{bounded}})$  if  $X_1 \cup X_2 \cup \dots \cup X_M$  fits inside such a box. From the results in section 29 of [23] it is clear this conjecture is true for a specific gate set,  $G$ , consisting of gates for which the associated operators span the full space of operators acting on the Hilbert space associated with two qubits. In this case,  $|G| = 256$ . The gate set illustrated in the above table has only  $|G| = 14$ . We conjecture that circuits formed on the above lattice with this gate set also have the property of composition tomography boundedness. More generally, we conjecture that circuits formed with respect to any lattice with any universal gate set have this property.

We can calculate the probability for any circuit,  $\mathcal{C}$ , using only  $\mathbf{r}$  vectors in the following way. First we consider

$$p_{\mathcal{F} \cup \bar{\mathcal{F}}}^n = \mathbf{r}_X^n[\mathcal{F}] \cdot \mathbf{p}_X^n[\bar{\mathcal{F}}] \quad (\text{A7})$$

where  $X$  is the region for fragment  $\mathcal{F}$ , and  $\bar{\mathcal{F}}$  is another fragment. Also consider

$$\text{prob}(\bar{\mathcal{F}}|\mathcal{T}_{\text{choice}}) = \sum_l p_{\mathcal{F}[l] \cup \bar{\mathcal{F}}}^n = \sum_l \mathbf{r}_X^n[\mathcal{F}[l]] \cdot \mathbf{p}_X^n[\bar{\mathcal{F}}] \quad (\text{A8})$$

where  $\mathcal{F} = \mathcal{F}[1]$  and  $\mathcal{T}_{\text{choice}} = \{\mathcal{F}[l] : \text{forall } l\}$  are a set of fragments for  $X$  with mutually exclusive and exhaustive set of outcomes so that their probabilities must add to one. We denote this set with the subscript 'choice' because it represents the choice we make in  $X$  (a different choice would correspond to a different mutually exclusive and exhaustive set). The vector  $\mathbf{p}_X[\bar{\mathcal{F}}]$  depends on  $\bar{\mathcal{F}}$  and is independent of  $\mathcal{F}[l]$  in region  $X$ . By Bayes rule

$$\text{prob}^n(\mathcal{F}|\bar{\mathcal{F}}, \mathcal{T}_{\text{choice}}) = \frac{\mathbf{r}_X^n[\mathcal{F}] \cdot \mathbf{p}_X^n[\bar{\mathcal{F}}]}{\sum_l \mathbf{r}_X^n[\mathcal{F}[l]] \cdot \mathbf{p}_X^n[\bar{\mathcal{F}}]} \quad (\text{A9})$$

Hence, iff

$$\mathbf{r}_X^n[\mathcal{F}] \text{ is proportional to } \sum_l \mathbf{r}_X^n[\mathcal{F}[l]] \quad (\text{A10})$$

then  $\text{prob}^n(\mathcal{F}|\bar{\mathcal{F}}, \mathcal{T}_{\text{choice}})$  is independent of  $\bar{\mathcal{F}}$ . We denote this probability by  $p_{\mathcal{F}|\mathcal{T}_{\text{choice}}}^n$  and it is given by

$$\mathbf{r}_X^n[\mathcal{F}] = p_{\mathcal{F}|\mathcal{T}_{\text{choice}}}^n \sum_l \mathbf{r}_X^n[\mathcal{F}[l]]. \quad (\text{A11})$$

We can apply this approach to the case where  $\mathcal{F}$  and  $\bar{\mathcal{F}}$  are circuits. If we have ideal circuits then it is true in quantum theory (and circuit theories in general [36]) that

$$p_{\mathcal{C} \cup \bar{\mathcal{C}}} = p_{\mathcal{C}} p_{\bar{\mathcal{C}}}. \quad (\text{A12})$$

From this property it can be proven that the proportionality condition in (A10) must hold and hence the probability for a circuit is given by

$$\mathbf{r}_X[\mathcal{C}] = p_{\mathcal{C}} \sum_I \mathbf{r}_X[\mathcal{C}[I]]. \quad (\text{A13})$$

Here we have dropped the possible dependence of  $p_{\mathcal{C}}$  on the choice of mutually exclusive circuit set (it is possible to prove that there is no such dependence). For non-ideal circuits, however, the proportionality condition of (A10) may fail even when  $\mathcal{F}$  is a circuit because we have not fully trained it to behave like a circuit. Nevertheless, if the proportionality condition holds approximately, then we can bound the probability [26].

We now prove the following theorem.

**Theorem 1.** If we have fragment tomography boundedness and composition tomography boundedness, then there exists a bounded set of circuits  $\mathcal{S}_{\text{tom}}$  such that if, at some iteration  $n = \tilde{n}$ , we have

$$p_{\mathcal{C}}^{\tilde{n}} = p_{\mathcal{C}} \quad \text{for all } \mathcal{C} \in \mathcal{S}_{\text{tom}}$$

then  $p_{\mathcal{C}}^{\tilde{n}} = p_{\mathcal{C}}$  for any  $\mathcal{C} \in \mathcal{S}_{\text{circuits}}$ .

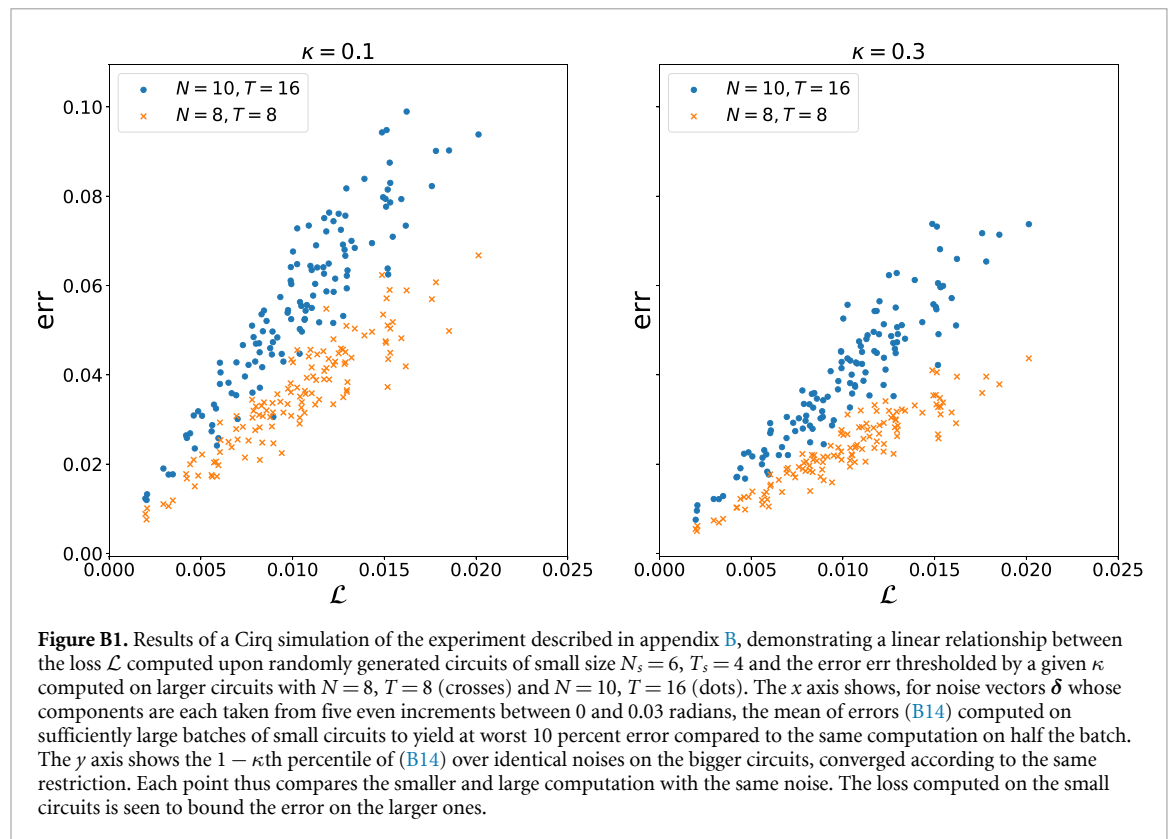
Recall that  $p_{\mathcal{C}}$  are the probabilities for idealised circuits when the gates  $g$  are actually those in our UGS,  $G$ . We can obtain these probabilities by calculation using Quantum Theory. To calculate  $p_{\mathcal{C}}^{\tilde{n}}$  for an arbitrary circuit we need to be able to calculate the  $\mathbf{r}_X^{\tilde{n}}$  vectors for the region associated with  $\mathcal{C}$ . We can do this in the following way. If we have  $\Lambda_X^{X_1, X_2, \dots}$  where  $X_1, X_2, \dots$  live inside bounded boxes and we can also calculate  $\mathbf{r}_{X_1}, \mathbf{r}_{X_2}, \dots$  for these same bounded boxes and then we can calculate  $\mathbf{r}_X$  vectors for region  $X$ , and from these we can calculate  $p_{\mathcal{C}}$ . We can obtain these  $\mathbf{r}_{X_i}$  vectors from probabilities for bounded circuits by fragment tomography boundedness. If we have composition tomography boundedness then we can obtain  $\Lambda_X^{X_1, X_2, \dots}$  from  $\Lambda$ 's that pertain to bounded boxes. Furthermore, we can determine these  $\Lambda$ 's that live in bounded boxes by constructing circuits that live in bigger, but still bounded, boxes by fragment tomographic locality. Thus, all the circuits we need to do tomography live in bounded boxes. If we obtain the probabilities  $p_{\mathcal{C}}^{\tilde{n}} = p_{\mathcal{C}}$  for this set of circuits then, by the above argument, we will obtain  $\mathbf{r}_X^{\tilde{n}}[\mathcal{F}] = \mathbf{r}_X[\mathcal{F}]$  for any  $X$  and  $\mathcal{F}$ . We can calculate the probability for any circuit using these  $\mathbf{r}$  vectors and hence this proves the theorem.

## Appendix B. Numerical motivation of random sampling assumption

In this appendix we motivate the random circuit sampling assumption (the text surrounding (5)) by way of a simple numerical model. We imagine an experimenter with access to a quantum computer composed of  $N_q$  qubits. The computer is able to implement the Clifford gate set, CNOT,  $H$ ,  $T$ ,  $S$  and  $I$ , subject to noise modelled by a noise vector  $\delta = (\delta_x, \delta_y, \delta_z)$ . The actually-implemented single-site gates differ from the experimenter's choice by a phase rotation about each axis of the Bloch sphere. The magnitude of each rotation is drawn from a Gaussian distribution about 0 whose standard deviation is given by its respective component of  $\delta$ .

The experimenter seeks to construct random circuits  $\mathcal{S}_{\text{rand}}$  such that the loss function (5) – the RMS error between batches of observed and actual circuit measurement probabilities – computed on small circuits places a bound on finite percentiles of those which would be computed for arbitrarily large circuits via (5). The idea is to motivate that knowledge of how to implement random small circuits grants knowledge of how to implement arbitrary circuits, including much larger ones. In fact Theorem 1 proves this, so that the goal of the experiment is to motivate that the results persist despite error in the quantum circuit.

The experimenter constructs their  $\mathcal{S}_{\text{rand}}$  as follows. A 'small' circuit size of  $N_s$  qubits and  $T_s$  timesteps is selected. The qubits are initialized in the  $|0\rangle$  state. At each timestep either a CNOT or an identity gate is applied to each qubit pair. One of  $I$ ,  $H$ ,  $S$ , or  $T$  gates, along with the Gaussian noise, is also randomly selected and applied with equal probability. The squared modulus of the resulting wavefunction then yields the probability vector  $P_{\text{actual}}$  for measurements of this circuit in the computational basis. A second circuit, identical except that  $\delta = 0$ , is generated to yield a corresponding  $P_{\text{expected}}$ .



The error  $e_i$  of the  $i$ th run is then estimated as the L2 norm of the difference between  $P_{\text{actual}}$  and  $P_{\text{expected}}$ :

$$e_i = \mathbf{norm}|P_{\text{actual}} - P_{\text{expected}}|. \quad (\text{B14})$$

We then estimate the loss  $\mathcal{L}$  computed on the small circuit for a given noise vector  $\delta$  as the mean of all  $e_i$  computed on some batch. This batch is expanded by increments of 100 until the mean of the batch and of half the batch disagree at most by 10 percent.

A ‘large’ circuit size  $(N, T)$  is then chosen, and batches of  $e_i$  computed upon these for the same various  $\delta$  as used on the small one. We then choose a few values of the  $\kappa$  appearing in the random circuit sampling assumption, and compute the  $1 - \kappa$ th percentile of the  $e_i$ ’s on the batch. The batch is again expanded until each percentile disagrees with that computed upon half the batch by at most 10 percent. We use this computation as an estimate of the error  $\text{err}$  in (5) for the given  $\kappa$ .

We consider the random circuit sampling assumption to have been falsified if a predictable relationship between  $\mathcal{L}$  and  $\text{err}$  fails to develop having computed the two for each of various  $\delta$ . But in figure B1, which presents the results of a simulation of the experiment detailed here conducted using Google’s Cirq framework, we instead see a clearly linear relationship, indicating that study of the smaller circuit in fact yields information about the larger one despite the error.

## ORCID iD

Adam G M Lewis  <https://orcid.org/0000-0002-0604-195X>

## References

- [1] Chuang I L and Nielsen M A 1997 *J. Mod. Opt.* **44** 2455
- [2] O’Brien J L, Pryde G J, Gilchrist A, James D F V, Langford N K, Ralph T C and White A G 2004 *Phys. Rev. Lett.* **93** 080502
- [3] Poyatos J F, Cirac J I and Zoller P 1997 *Phys. Rev. Lett.* **78** 390
- [4] Artiles L M, Gill R D and Guta M I 2005 *J. Royal Stat. Soc.* **67** 109
- [5] D’Ariano G M, Paris M G A and Sacchi M F 2003 *Adv. Imag. Elect. Phys.* **128** 205
- [6] Nielsen M A and Chuang I L 2011 *Quantum Computation and Quantum Information* 10th edn (New York: Cambridge University Press)
- [7] Knill E *et al* 2008 *Phys. Rev. A* **77** 012307
- [8] Emerson J, Silva M, Moussa O, Ryan C, Laforest M, Baugh J, Cory D G and Laflamme R 2007 *Science* **317** 1893
- [9] Emerson J, Alicki R and Życzkowski K 2005 *J. Opt. B: Quantum Semiclass. Opt.* **7** S347
- [10] Magesan E, Gambetta J M and Emerson J 2011 *Phys. Rev. Lett.* **106** 180504

- [11] Merkel S T, Gambetta J M, Smolin J A, Poletto S, Córcoles A D, Johnson B R, Ryan C A and Steffen M 2013 *Phys. Rev. A* **87** 062119
- [12] Blume-Kohout R, Gamble J K, Nielsen E O, Mizrahi J, Sterk J D, and Maunz P 2013
- [13] di Matteo O 2019 PIRSA:19070074 (available at <http://pirsa.org/19070074/>)
- [14] Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N and Lloyd S 2017 *Nature* **549** 195
- [15] Schuld M, Sinayskiy I and Petruccione F 2015 *Contemp. Phys.* **56** 172
- [16] Jaffali H and Oeding L 2019 (arXiv:1908.10247)
- [17] Carrasquilla J, Torlai G, Melko R and Aolita L 2019 *Nat. Mach. Intell.* **1** 155
- [18] Beach M J S, Vlugt I D, Golubeva A, Huembeli P, Kulchytskyy B, Luo X, Melko R G, Merali E and Torlai G 2019 *SciPost Phys.* **7** 9
- [19] Torlai G, Mazzola G, Carrasquilla J, Troyer M, Melko R and Carleo G 2018 *Nat. Phys.* **14** 447
- [20] Quek Y, Fort S, and Ng H K 2018 (arXiv:1812.06693)
- [21] Xu Q and Xu S 2018 (arXiv:1811.06654)
- [22] Hardy L 2012 *Philos. Trans. R. Soc. A* **370** 3385–417
- [23] Hardy L 2005 (arXiv:0509120)
- [24] Hardy L 2007 *J. Phys. A: Math. Theor.* **40** 3081
- [25] Hardy L 2009 *Quantum Reality, Relativistic Causality and Closing the Epistemic Circle* (Berlin: Springer) pp 379–401
- [26] Markes S and Hardy L 2011 *J. Phys.: Conf. Ser.* (Bristol: IOP Publishing) vol 306 p 012043
- [27] Cun Y L, Bengio Y and Hinton G 2015 *Nature* **521** 436
- [28] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (Cambridge, MA: MIT Press) ([www.deeplearningbook.org](http://www.deeplearningbook.org).)
- [29] Rumelhart D, Hinton G and Williams R 1986 *Nature* **323** 533
- [30] Rumelhart D E and McClelland J L 1987 Learning internal representations by error propagation *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations* (Cambridge, MA: MIT Press) pp 318–62 (<https://ieeexplore.ieee.org/document/6302929>)
- [31] Hochreiter S and Schmidhuber J 1997 *Neural Comput.* **9** 1735
- [32] Sherstinsky A 2018 (arXiv:1808.03314) CoRR abs/1808.03314
- [33] Mattoussi H, Mauro J M, Goldman E R, Anderson G P, Sundar V C, Mikulec F V and Bawendi M G 2000 *J. Am. Chem. Soc.* **122** 12142
- [34] Garca J, Medeiros-Ribeiro G, Schmidt K, Ngo T, Feng J, Lorke A, Kotthaus J and Petroff P 1997 *Appl. Phys. Lett.* **71** 2014
- [35] Mana P G 2003 quant-ph/0305117
- [36] Hardy L 2011 (arXiv:1104.2066)