Elementary Quantum Cellular Automata

Cooper Jacobus[1,2], Charlie Tolley[1,2]
[1]*Department of Astronomy, UC Berkeley*
[2]*Department of Physics, Lawrence Berkeley National Lab*

Abstract

Cellular automata (CA) are computational models in which simple rules applied to discrete states in a grid produce complex emergent behavior. One well-know 2-dimensional example is Conway's Game of Life, showing how a few basic rules can give rise to intricate patterns and emergent structures. However, 1-dimensional cellular automata, a simpler variant where a string of bits updated iteratively according to local rules, also exhibit fascinating dynamics ranging from stable configurations to chaotic, irreducible behaviors [7]. This project explores quantum cellular automata (QCA), a quantum extension of 1D Cellular Automata, where cells are represented by qubits, allowing for quantum superposition and entanglement [2]. By introducing some quantum nature to these 1D systems, we aim to investigate how quantum mechanics affects emergent behaviors and information flow. Using information theory, particularly mutual information and entanglement entropy, we measure correlations to understand the differences between classical and quantum systems. This approach allows us to assess how quantum effects influence the evolution of CA, potentially leading to unique behaviors and patterns not seen in classical systems. Through this examination, we seek to explore the role of quantum mechanics in complex systems and its implications for quantum information processing and computation.

1 Introduction

Before we proceed with our quantum modification, we must first review the classical model which we are modifying.

1.1 Elementary Cellular Automata

Elementary Cellular Automata (ECA) are a type of 1-dimensional cellular automaton defined by simple, deterministic rules that govern the evolution of a linear array of bits. In ECA, each bit's new state is determined by its current state and the states of its two immediate neighbors [7]. This leads to eight possible neighborhood configurations (000 to 111), each mapped to a resulting bit in the next step. Since each of these configurations can independently produce either a 1 or a 0, there are 256 unique rules for evolving the system. These rules vary widely in their behavior, ranging from simple repetition and stability to chaotic and complex patterns. Rule 30, for example, is renowned for its unpredictable, irreducible complexity, creating intricate, seemingly random sequences from simple initial conditions. This diversity in behavior makes ECA a fascinating subject for studying emergent phenomena, complexity, and computation within a simple rule-based system.

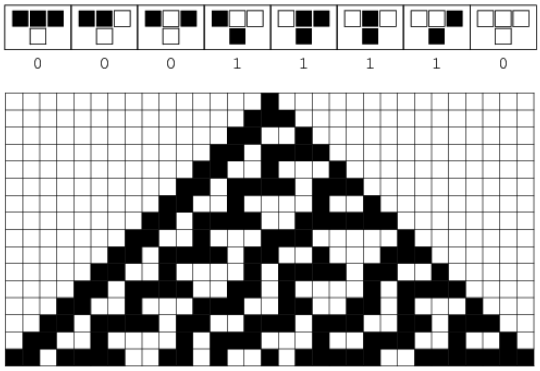

Figure 1: The Elementary Cellular Automata "Rule 30". The 8 individual update rules are shown above, and the evolution of the bit string starting from one central on bit is shown with time flowing down. Although the classical version is entirely deterministic, the behavior is thought to be irreducible and the central column is even used for pseudo-random number generation in computer language Mathematica [7].

1.2 Quantum Modification

We adjust the formulation of Elementary Cellular Automata to allow for the initial conditions of each cell to be in superposition, not just 1 or 0 but with a state-vector anywhere on the Bloch sphere. We translate the rules of the classical cellular automata into the language of quantum gates, so that we can evolve the system in time according to our rules. We check that if we use our quantum rules on a initial state with no superposition, we reproduce the same result as the classical model, but when we introduce superposition we find novel and exciting behavior. Some configurations produce entirely new patterns due to constructive or destructive in-

terference. These systems are chaotic an unpredictable by construction, and we cannot reasonably uncover a generalized understanding of their behavior, but we can examine certain well-defined properties like their degree of initial superposition or the correlation between certain qubits in the output.

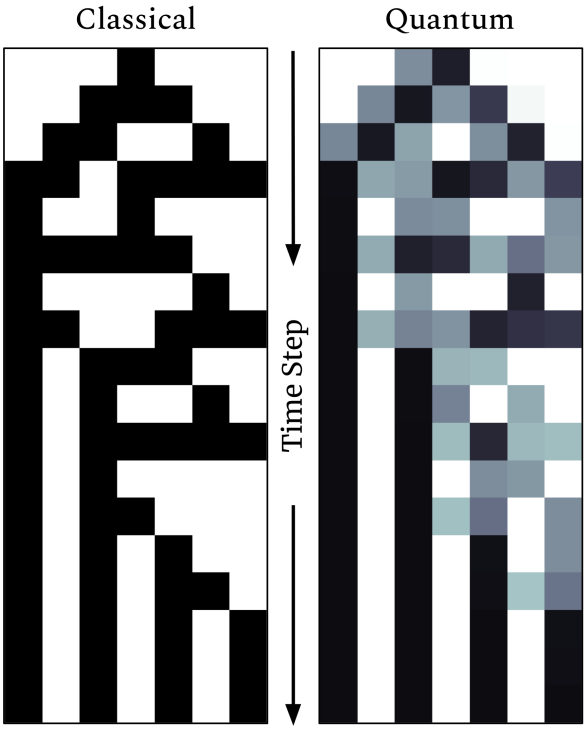

Figure 2: Side by side comparison of the same cellular automata rule (rule 30), both in the classical limit (left) where all cells have been initialized as either 1 or 0, and in the quantum regime (right) where some cells are initialized in superposition. Notice the interference that occurs in the quantum case. A series of vertical stripes is a stable pattern for this rule which emerges in both cases due to the boundary conditions of the small domain.

2 Techniques

2.1 Developing Quantum Rules

Translating the rules of Elementary Cellular Automata (ECA) into the quantum domain requires a careful approach to preserve the system's underlying structure while incorporating quantum mechanics. In classical cellular automata, the state of a cell is influenced by the three-cell neighborhood in the previous step. The quantum equivalent must capture this same principle, yet with quantum gates and qubits. Due to the no coloning theorem, we cannot simply duplicate and modify a register of data, as is done in the classical case. Instead, we initialize the quantum To implement the update rule in a quantum cellular automaton, the triple-controlled NOT (CCCNOT) gate is used. This gate acts on a target qubit, flipping its state if and only if all three control qubits are in the $|1\rangle$ state. We apply the CCCNOT gate using qubits representing the previous state as the control, and with the target qubit representing the next state. In this way we can create quantum cellular automata that respect the same neighborhood-based logic as their classical counterparts. This allows for the exploration of quantum behaviors while maintaining the fundamental rule structure that characterizes Elementary Cellular Automata.

Each of the sub-rules is independent so each of the 8 possible neighborhood states is accounted for and is implemented separately. Although a CCCNOT gate need only be applied for those sub-rules which produce a $|1\rangle$ in the next state as all qubits in the later states are initialized as $|0\rangle$ so sub-rules which yield zero need not be calculated, which accelerates computation.

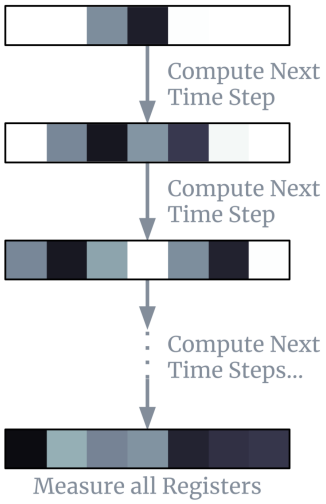

Figure 3: Visual Representation of our time-stepping methodology. All cells remain in superposition and entangled until the very end.

2.2 Circuit Implementation

To track the evolution of a quantum cellular automaton (QCA), it is essential to maintain the complete state information throughout the computation. The implementation described here utilizes a unique qubit for each individual cell at each individual timestep, enabling us to record the entire history of the QCA while avoiding state collapse due to measurement. For instance if we have n cells and m timesteps then we use $n \times m$ qubits to simulate our entire system. This approach imposes a constraint on the total number of timesteps that can be implemented within a given qubit budget,

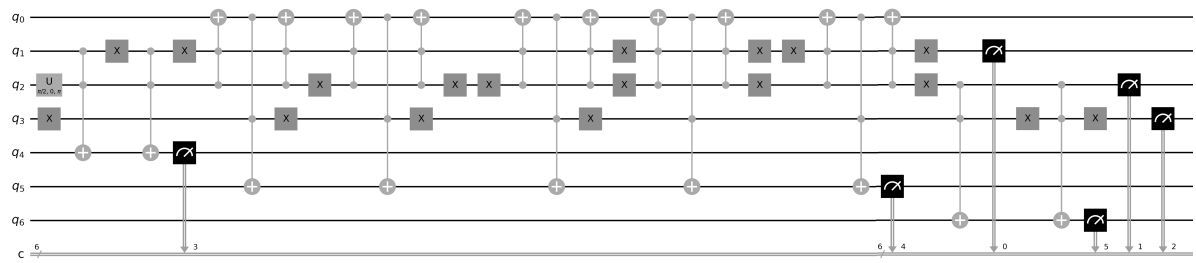

Figure 4: Visualization of a much abridged quantum cellular automata, this shows Rule 30 carried out on a length 3 qubit string over one timestep. Qubits 1,2,3 represent the initial timestep and qubits 4,5,6 represent the next step. Qubit 0 is an ancillary used for the CCCNOT gates. Note that measurements occur only after the last operation on any given qubit.

but it provides a clear representation of the system's evolution from start to finish.

The measurement strategy is key to preserving the quantum information throughout the computation. In this implementation, measurements are conducted only at the end of the final timestep, which prevents state collapse during intermediate steps. This strategy enables the system to evolve according to quantum principles, allowing for superposition and interference without disruption. It also facilitates a comprehensive representation of the entire evolution, as measurements at the final step capture the entire state history.

The initial state of the QCA plays a crucial role in determining its evolution. By initializing all qubits in the computational basis, either in the $|0\rangle$ or $|1\rangle$ state, the QCA reproduces the behavior of classical elementary cellular automata. However, by initializing certain qubits in superposition states, quantum effects like interference become apparent. All qubits are initialized as $|0\rangle$ by default, but we can introduce some superposition to them if we wish.

To study the behavior of these systems in detail, we let the some cells begin in arbitrary superposition state between $|0\rangle$ and $|1\rangle$. Those we wish to introduce superposition to we pass through a U gate to transform them to the state:

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + \sin\left(\frac{\theta}{2}\right)|1\rangle \tag{1}$$

For some desired angle θ between $|0\rangle$ and $|1\rangle$. For the most simple cases, we transform only the center qubit in this way, but to explore more complicated interference patterns we will initialize multiple qubits in superposition, each with their own angle θ from $|0\rangle$.

2.3 Quantifying Complex Behavior

Understanding and measuring the behavior of these systems in a way which is generalized to work on all states is difficult because these systems are complex and sometimes irreducible by construction. To do this in a way which is flexible and interpretable for arbitrary rules and initial conditions, we adopt a standardized traditional metric from information theory called mutual information.

Mutual information describes how much information on random variable contains about another. It essentially quantifies how much measuring one random variable reduces the uncertainty in another, or how well two random variable correlate. It can be calculated from a distribution of random samples of the two random variables and in quantum settings it can be used to account for the correlation of random states [5, 1].

Mutual Information is defined in terms of another well-known quantity known as Shannon Entropy. Shannon Entropy describes the uncertainty (or "surprise") of a random variable, For many possible values x sampled from some probability distribution X, the Shannon Entropy of X is defined:

$$H(X) \equiv -\sum_{x \in X} p(x) \cdot \log_2(p(x)) \tag{2}$$

Here, $p(x)$ can be estimated from a sample distribution. Because $p(x)$ appears in the definition both as a linear term and in a logarithm, the Shannon entropy of a variable is 0 if the variable is not random (ie. if it gives the exact same value every time it is sampled). Likewise the Shannon entropy is maximized when the variable is entirely uncertain, with equal probability of being any of multiple possible values.

We define the Mutual information of two random variables in terms of their individual and joint Shannon Entropies:

$$I(X,Y) = H(X) + H(Y) - H(X,Y) \tag{3}$$

Where $H(X,Y)$ is simply the Shannon Entropy of X and Y combined. The Mutual information is maximized if X and Y are statistically independent, in other words if measuring X give you no additional information about the state of Y.

Substituting in the definition of Shannon Entropy we can see that:

$$I(X,Y) = \sum_{x \in X} \sum_{y \in Y} p(x,y) \cdot \log_2 \left(\frac{p(x,y)}{p(x) \cdot p(y)} \right) \quad (4)$$

Where $p(x,y)$ denotes the probability of getting x and y simultaneously when sampling from distributions X and Y. If x and y are independent then $p(x,y)$ is just $p(x) \cdot p(y)$ so the term in the logarithm is 1 and the contribution to the Mutual information will be 0. Likewise mutual information is maximized when x and y are always somehow correlated. Therefore we can use this metric to quantify the correlation of two of our qubits given a large distribution of sampled from our circuit.

For simple cases of our cellular automata we will sometimes wish to compare the mutual information between two qubits (which describes their degree of correlation) to some degree of superposition of some qubit in the initial row of our cellular automata. To quantify this degree of superposition we use the Von Neumann Entropy [6], defined as:

$$S(\rho) = -\operatorname{Tr}(\rho \log \rho) \quad (5)$$

Where ρ is the density matrix of a given state vector $|\psi\rangle$, calculated by:

$$\rho = |\psi\rangle\langle\psi| \quad (6)$$

If a state vector represents a pure state then this metric should be zero as the density matrix will be populated by only 1s and 0s. If we initialize a single qubit in our first row in super position with some angle θ away from pure state $|0\rangle$ according to:

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + \sin\left(\frac{\theta}{2}\right)|1\rangle \quad (7)$$

And we initialize all other qubits in the first row to a pure state, either $|0\rangle$ or $|1\rangle$, then the Von Neumann Entropy of the first row is:

$$S = -\sin^2\theta \cdot \log\left(\sin^2\theta\right) - \cos^2\theta \cdot \log\left(\cos^2\theta\right) \quad (8)$$

Where the degree of superposition is minimized at $\theta = 0$ or pi where $|\psi\rangle = |0\rangle$ or $|1\rangle$ and maximized at $\theta = \pi/2$ where $|\psi\rangle = |+\rangle$, as one would expect. For cases like this where one one qubit is in superposition, we expect the mutual information to be related to the Von Neumann Entropy because the only randomness available for any cells comes from the same place so the cells should be maximally correlated (for example see fig 5). As we allow multiple cells to be in superposition we expect the randomness to interfere and for cells to be more decoherent.

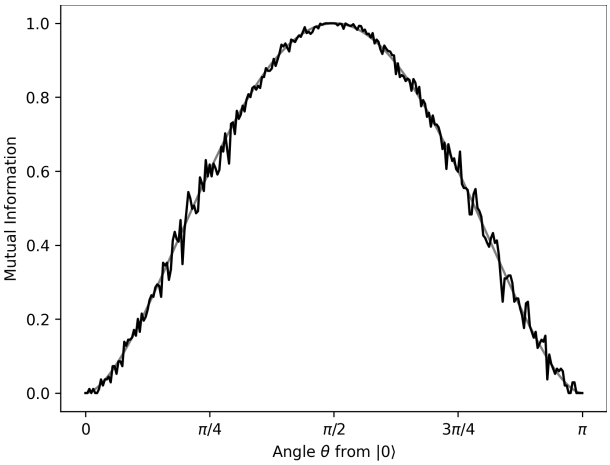

Figure 5: The Mutual Information between the first and last qubits as a function of the superposition angle of the center qubit, plotted alongside the Von Neumann Entropy of the initial state vector. This data is from a family of runs of rule 30 where all but the center qubit are initialized as $|0\rangle$

3 Results

In exploring Wolfram's rule space in the framework of Quantum Cellular Automata, we focused on a few rules well-known for their complex behavior under classical rules. Complexity, as we explored in Section 2.3, characterizes the emergent behavior of the quantum system, which we can explore and quantify, at least in part, with mutual information. Due to a very narrow boundary on the space axis of our 1-dimensional automata, we find emergent behavior that is not a product of the superposition that we intentionally set. This behavior is shown in Figure 2, where in both the classical and quantum cellular automata, we see that stable column structures evolved across the time dimension, which we find stem from the boundary conditions we have imposed. We chose such a slim grid because we were limited by the number of qubits we needed to initialize - our choice to record the full history of the cellular automata at once required that we initialized every single qubit across the time versus space grid. This limit functionally meant that we could only use 7 qubits across and 18 time steps to watch the cellular automata evolve, since the simulations became prohibitively expensive. However, in some cases where the 7 qubit limit meant the mutual information (correlation) was hidden by the dominant boundary condition interference patterns, we were able to extend to a 9-qubit starting condition that provided us with insightful visuals on the complexity of the particular rule.

3.1 9-qubit Simulations

We generated mutual information curves as a function of initial conditions for many different rules and most showed absolutely zero correlation when using 7 qubits, regardless of the superposed initial conditions. In these cases, we chose to find one set of initial conditions that could produce something in the 9-qubit case by taking a very low-resolution and thus very fast simulation of the behavior, and then increasing resolution from there. Each of these 9-qubit simulations evolved for 18 time-steps for a total 162 qubits we initialized at $|0\rangle$ with the top 9 having an initial set of superposition described below.

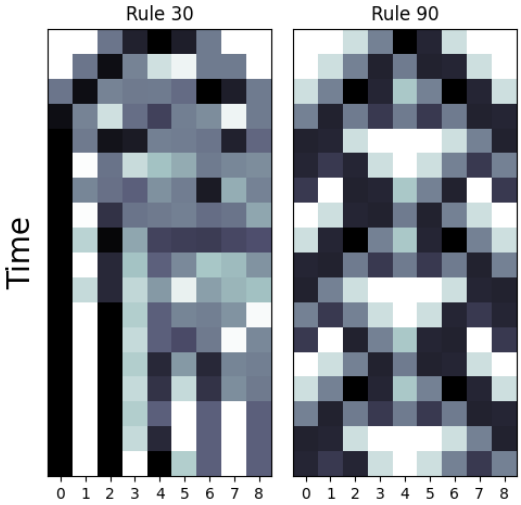

Figure 6: 9 qubit histories for two separate rules. For Rule 30 (also shown alongside it's classical counterpart in Figure 2), we initialized its qubits as follows - qubit 3 at $\frac{\pi}{2}$, qubit 4 at $\frac{-3\pi}{4}$, qubit 6 at $\frac{-3\pi}{4}$, and qubit 7 at $\frac{\pi}{2}$. For Rule 90, we initialized its qubits as - qubit 3 at $\frac{\pi}{4}$, qubit 4 at $\frac{\pi}{2}$, qubit 6 at $\frac{3\pi}{4}$, and qubit 7 at $\frac{\pi}{2}$.

To quantify the complexity of each of these systems, we varied the center qubit from $|0\rangle$ *to* $|1\rangle$, and held constant the initial superposition of the surrounding qubits, finding the MI of each unique system. Rule 30 seems to exhibit an exponential decay of MI, where rule 90 almost follows the Von Neumann Entropy shape of the Figure 5, albeit with a different scale.

3.2 Rule 122

Rule 122, known for the checkerboard pattern it displays as an ECA, proved the easiest to work with

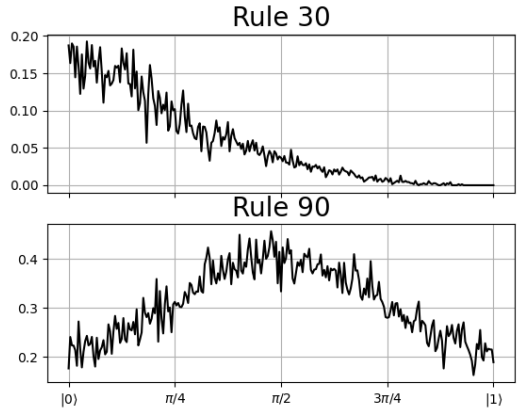

Figure 7: The complexity of rule 30 and 90 with the central qubit varying from $|0\rangle$ *to* $|1\rangle$ and the other qubits held in the superposition as shown in Figure 6.

as it produced the most varied complexity behavior with a 7-qubit initial state. This allowed me to use the `Python threading` module to calculate the 7 systems shown in Figure 8 at once. Many of these

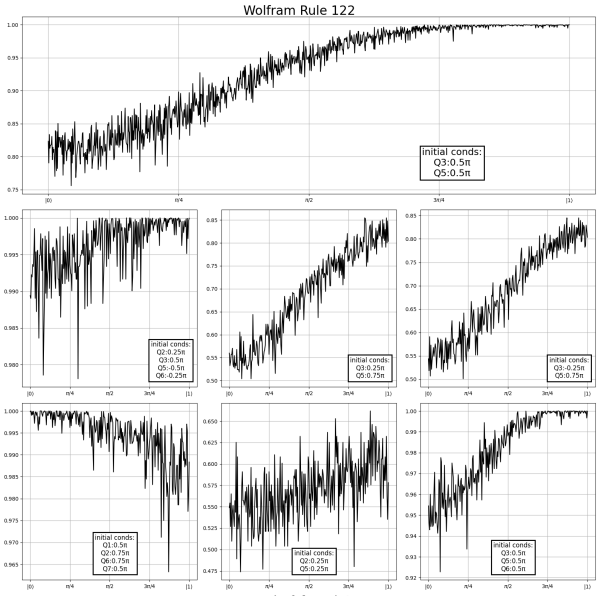

Figure 8: Mutual Information of the first and last qubit of the 7-qubit system after 18 time steps. The initial conditions surrounding the central qubit are listed on each component graph.

systems appear to asymptote to a maximal MI of 1, where only one starts at a maximal MI and decreases as the center qubit goes to $|1\rangle$, highlighting the variety in complex emergent behavior that even a single rule can produce.

4 Discussion

We are seriously limited in the number of qubits we can implement, this often results in stable patterns like those seen in figure 2 which result from interaction with the boundary. All boundaries are treated as if they are zero, but themselves cannot be updated so they produce a kind of artificially stable pattern which dominates the behavior especially at late time. To get around this, future work could employ periodic boundary conditions or explore other implementations which do not track the entire history and can therefore use more qubits for the automata's width. The overwhelming number of qubits needed for our original implementation also imposes a barrier toward implementing more sophisticated metrics, such as quantum information [3] and quantum discord [4]. We attempted to implement such metrics but quickly discovered that they required petabytes of memory to be deployed on our 128-qubit circuit. This provides further motivation to explore more efficient implementations which do not require as many qubits.

Implementations of this type are complicated as a result of the no-cloning theorem. A single row of qubits representing the 1D automata cannot be used alone as the updates could not occur simultaneously without disrupting the state. But perhaps two rows of qubits could be used and alternated between. Or, in special cases, the entire update rule could be encapsulated in a single (very large) unitary operator which could be applied repeatedly to simulate the timestepping. Such an implementation would only contain information about the final row when measured but could be run for an arbitrary number of timesteps. We suggest that a QCA circuit of this nature could be used to greater effect as it could be much wider and therefore not suffer from edge effects as immediately and because the smaller memory footprint would allow for the efficient employment of more informative metrics. We also suggest the possibility of measuring the phase of the cellular automata or even developing an expanding family of rules using Hadamard gates in addition to X's to introduce even more diverse behavior.

5 Conclusion

We have explored a novel approach to implementing famous classical cellular automata rules in a quantum framework. We developed an algorithm which can translate the classical rules into a quantum circuit is a generalized and flexible way. We demonstrate that our method reproduces expected classical behavior when fed classical inputs but can produce new and intriguing phenomena when initialized in superposition.

We explore the space of possible rules and initial conditions and find that the diversity of classical cellular automata is expanded into an entirely new direction of superposition and quantum interference. We explore the correlation of separate qubits as a function of initial conditions and find hints of tractable behavior but as of yet have no understanding of their nature other than that they are very diverse and unpredictable!

6 Acknowledgements

The authors would like to thank John Paul (JP) Marceaux for his guidance and insight and for his eagerness to collaborate on this work and push it to be the best it possibly could.

7 Contributions

C.J. Developed the algorithm which constructs and executes the Quantum Cellular Automata and the computational methods which quantify superposition and correlation. C.T. Explored the rule space and tabulated interesting results and made plots of correlation vs initial condition. Both authors contributed to the writing of this report.

References

[1] Silviu Guiasu. *Information Theory with Applications*. McGraw-Hill, New York, 1977.

[2] Logan E Hillberry, Matthew T Jones, David L Vargas, Patrick Rall, Nicole Yunger Halpern, Ning Bao, Simone Notarnicola, Simone Montangero, and Lincoln D Carr. Entangled quantum cellular automata, physical complexity, and goldilocks rules. *Quantum Science and Technology*, 6(4):045017, sep 2021.

[3] Asutosh Kumar. Multiparty quantum mutual information: An alternative definition. *Phys. Rev. A*, 96:012332, Jul 2017.

[4] Harold Ollivier and Wojciech H. Zurek. Quantum discord: A measure of the quantumness of correlations. *Phys. Rev. Lett.*, 88:017901, Dec 2001.

[5] Claude Elwood Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, July 1948.

[6] John von Neumann. *Mathematical Foundations of Quantum Mechanics*. Princeton University Press, new edition edition, 2018.

[7] Stefan Wolfram. *A New Kind of Science*. Wolfram Media, 2002.