Dimension reduction in quantum sampling of stochastic processes

Chengran Yang[1 *] Marta Florido-Llinàs[2 3] Mile Gu[1 4] Thomas J. Elliott[5 6 3]

[1] *Nanyang Quantum Hub, Nanyang Technological University, Singapore 637371*
[2] *Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Straße 1, D-85748 Garching, Germany*
[3] *Department of Mathematics, Imperial College London, London SW7 2AZ, United Kingdom*
[4] *Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543*
[5] *Department of Physics & Astronomy, University of Manchester, Manchester M13 9PL, United Kingdom*
[6] *Department of Physics, University of Oxford, Oxford OX1 3PU, United Kingdom*

Abstract. Quantum technologies offer a promising route to the efficient sampling and analysis of stochastic processes, with potential applications across the sciences. Such quantum advantages rely on the preparation of a quantum sample state of the stochastic process, which requires a memory system to propagate correlations between the past and future of the process. Here, we introduce a method of lossy quantum dimension reduction that allows this memory to be compressed, not just beyond classical limits, but also beyond current state-of-the-art quantum stochastic sampling approaches. We investigate the trade-off between the saving in memory resources from this compression, and the distortion it introduces. We show that our approach can be highly effective in low distortion compression of both Markovian and strongly non-Markovian processes alike. We further discuss the application of our results to quantum stochastic modelling more broadly. For more details, please refer to arXiv:2404.10338.

Keywords: Stochastic processes, Quantum sampling, Dimension reduction, Matrix product states

Introduction – Complex stochastic processes abound across the sciences, from evolutionary biology and chemistry [1, 2, 3, 4, 5], through geophysics and astrophysics [6, 7], to financial markets [8, 9, 10, 11], traffic modelling [12], and natural language processing [13, 14]. Given their pivotal role in these fields, it is vital that we can effectively simulate, analyse, and understand stochastic processes. However, the number of possible trajectories that such processes can explore generically grows exponentially over time, limiting the horizon over which we can feasibly study their behaviour. This makes tools and techniques that mitigate this growth in complexity of critical value.

Monte Carlo methods [15, 16, 17] present such a technique. They use generative models to reduce computational resources by sampling from the process one trajectory at a time, and then average over many such sampled trajectories to estimate the expected values of properties of the process. The considerable successes of such techniques notwithstanding, they still suffer from certain drawbacks, such as the need for a (typically large) memory system to carry the information propagated in correlations over time in the process.

Quantum information processing provides a further route to efficient sampling, modelling, and analysis of stochastic processes. Known quantum advantages include quadratic speed-ups in analysing properties such as characteristic functions [18], pricing options [19, 20, 21], enhanced expressivity [22], and significant reductions in the memory required by models [23, 24, 25]. These advantages involve the preparation of quantum sample states ('*q-samples*') that comprise of all possible strings of outputs in (weighted) superposition [26]. Unlike other quantum approaches to superposing classical data states – such as qRAM [27] and variational state preparation [28] – the computational complexity of assembling q-samples of stochastic process trajectories need not grow exponentially; they can be constructed through a local, recurrent circuit structure prescribed by the aforementioned quantum stochastic models [29, 30].

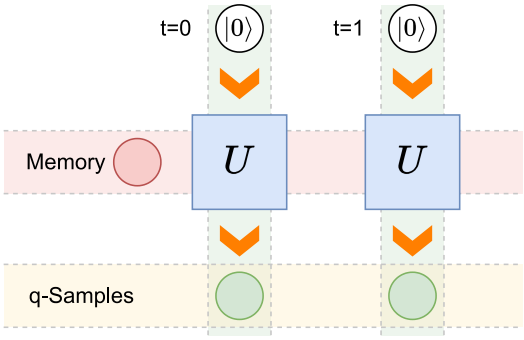

Figure 1: **Recurrent quantum circuit for locally generating quantum sample states (q-samples).** For each timestep a memory system is coupled with a blank ancilla; following an interaction U the ancilla is entangled with the previous timesteps of the q-sample to yield the appropriate marginal q-sample state, whilst the memory is passed forward to interact with the blank ancilla for the subsequent timestep. This sequentially generates the full q-sample state for any desired number of timesteps.

This recurrent circuit is illustrated in Fig. 1. It works by propagating a memory system – which carries information about the correlations between the past and future of the process – to interact with blank ancillae and sequentially assemble the q-sample state timestep by timestep. The circuit depth – and number of ancillae qubits – thus scales only linearly with the number of timesteps. However, as complex stochastic processes

*yangchengran92@gmail.com

often exhibit strong temporal correlations, the requisite dimension of the memory system can grow quite large.

Specifically, we develop a systematic approach to *lossy quantum dimension reduction*, whereby given a target q-sample circuit we determine a new circuit of fixed memory dimension that assembles an approximation to the original q-sample. Our approach is inspired by matrix product state (MPS) truncation techniques [31, 32], and exploits a correspondence between MPSs and q-samples [33]. We show that this approach yields high-fidelity approximations to q-samples corresponding to both Markovian and highly non-Markovian (i.e., strong temporal correlations) stochastic processes. We further discuss the applications of this approach in stochastic modelling [34], demonstrating that we can achieve significant compression beyond both current state-of-the-art lossless quantum dimension reduction for quantum models, and lossy compression for classical models, whilst retaining highly-accurate reconstruction of the output statistics.

Processes and Models – For discrete-time, discrete-event stochastic processes, at each timestep t the corresponding event $x_t \in \mathcal{X}$ is described by a random variable X_t, drawn from a joint distribution $P(\ldots, X_t, X_t + 1, \ldots)$ [35]. Sequences of events are typically correlated, and we use the shorthand $x_{t:t'}$ to denote the contiguous sequence $x_t, x_{t+1}, \ldots, x_{t'-1}$. Here we consider stationary (time-invariant) stochastic processes, such that $t \in \mathbb{Z}$, and all marginal distributions of any length are shift invariant (i.e., $P(X_{0:L}) = P(X_{m:m+L}) \forall m \in \mathbb{Z}, L \in \mathbb{N}$.

Such stochastic processes can be generated sequentially by a *model*, consisting of an encoding function $f : \overleftarrow{\mathcal{X}} \to \mathcal{M}$ that maps from sequences of past observations to a set of memory states $\mathcal{M}$, and an update function $\Lambda : \mathcal{M} \to \mathcal{X} \times \mathcal{M}$ that acts on the current memory state to (stochastically) produce the next output and update the memory state accordingly. Amongst all classical unifilar models that reproduce the exact statistics of the process, the provably-memory-minimal (both in terms of information stored and memory dimension) is the *ε-machine*, prescribed by an encoding function f_ε that satisfies $f_\varepsilon(\overleftarrow{x}) = f_\varepsilon(\overleftarrow{x}') \Leftrightarrow P(\overrightarrow{X}|\overleftarrow{x}) = P(\overrightarrow{X}|\overleftarrow{x}')$ [36]. The corresponding memory states $s \in \mathcal{S}$ are referred to as the *causal states* of the process [37].

Curiously, it is possible to push the memory cost below classical limits when using quantum encoding and update functions, even though we are considering classical stochastic processes [38, 39, 29, 30, 40]. Such quantum models use an encoding function that maps to quantum states, and the update function is a quantum channel. These memory states are defined implicitly by the update function, which can be expressed in terms of a unitary operator U:

$$U \left|\sigma_j\right\rangle \left|0\right\rangle = \sum_x \sqrt{P(x|j)} e^{i\varphi_{xj}} \left|\sigma_{\lambda(x,j)}\right\rangle \left|x\right\rangle, \qquad (1)$$

where $P(x|j)$ is the probability that the next symbol is x given we are in causal state s_j, $\lambda(x, j)$ is an update rule that outputs the updated memory state label, and $\{\varphi_{xj}\}$ are a set of real numbers.

However, there is still a cost that grows with the complexity of the process – the size of the memory system. In practical terms, this can be quantified by the number of qubits required for the memory, i.e., $D_q := \log_2(d)$, where d is the dimension of the memory. This dimension is upper-bounded by the number of memory states, though can be lower when they exhibit linear dependencies [30, 23]. Also of relevance is the amount of information stored within the memory, i.e., the von Neumann entropy of the memory states. That is, $C_q := S(\rho)$, where $\rho := \sum_j P(j) \left|\sigma_j\right\rangle \left\langle\sigma_j\right|$ is the steady-state of the memory system. Clearly, we must have that $C_q \leq D_q$; for many stochastic processes we even find that $C_q \ll D_q$ [41, 42, 43, 44]. This indicates that much of the memory capacity is under-utilised given the amount of information that must be stored. In turn, this suggests that it may be possible to drastically reduce the memory dimension whilst discarding very little information, and thus assemble high-fidelity approximations to q-samples with drastically reduced memory size.

Distortion Measure – To quantify the accuracy of our approximate q-samples, we will use the quantum fidelity divergence rate (QFDR) [32]. This extends the notion of quantum fidelity (i.e., state overlap) from finite-dimensional states to infinite length chains. The QFDR R_F therefore captures the rate at which this distortion reduces the fidelity on a per timestep basis. Specifically,

$$R_F\left(\left|P\right\rangle, \left|P'\right\rangle\right) := -\lim_{L\to\infty} \frac{1}{2L} \log_2 F\left(\left|P_L\right\rangle, \left|P'_L\right\rangle\right), \qquad (2)$$

where $F(\left|\psi\right\rangle, \left|\phi\right\rangle) = \left|\left\langle\psi|\phi\right\rangle\right|$ is the quantum fidelity between states, $\left|P_L\right\rangle$ is shorthand for the L-length q-sample $\left|P(X_{0:L})\right\rangle$, and $\left|P\right\rangle$ the $L \to \infty$ limit of this.

Theoretical Bounds – The direct form of truncation for MPSs is to reduce them to their canonical form, keep the $\tilde{d}$ largest Schmidt coefficients at each bond (where $\tilde{d}$ is the desired final bond size), discard the rest, and then uniformly rescale the remaining Schmidt coefficients such that the resulting MPS is appropriately normalised [45, 46, 47, 48]. For the case of q-samples, described by iMPS, we can upper-bound the QFDR R_F in terms of the magnitude of the discarded Schmidt coefficients.

Theorem 1 *Consider a q-sample $\left|P\right\rangle$ for which the iMPS representation has d Schmidt coefficients $\{\lambda_j\}$ labelled in decreasing order. For any truncated dimension $\tilde{d}$, there always exists a q-sample $\left|\tilde{P}\right\rangle$ that can be constructed sequentially with a memory of at most dimension $\tilde{d}$ that satisfies*

$$R_F(\left|P\right\rangle, \left|\tilde{P}\right\rangle) \leq \frac{\epsilon_{\tilde{d}}}{2(1 - L\epsilon_{\tilde{d}}) \ln 2} + O\left(\frac{1}{L}\right) \qquad (3)$$

for any $L \in \mathbb{N}$, where $\epsilon_{\tilde{d}} := \sum_{k=\tilde{d}+1}^{d} \lambda_k$.

Moreover, by setting $L \sim 1/2\epsilon_{\tilde{d}}$ we obtain that $R_F(\left|P\right\rangle, \left|\tilde{P}\right\rangle) \leq O(\epsilon_{\tilde{d}})$. Thus, we see that the QFDR grows only linearly with the sum of discarded Schmidt

coefficients. By choosing $\tilde{d}$ such that $\epsilon_{\tilde{d}}$ is sufficiently small, we can therefore control the distortion between the target q-sample and its reduced memory approximation.

However, this does not by itself carry any guarantees that such a $\tilde{d}$ is much smaller than the original memory dimension d. In the following theorem and corollary we establish a relationship between the information cost of the q-sample, the truncated memory dimension, and the size of the sum of discarded Schmidt coefficients.

Theorem 2 *Given an iMPS with $d \geq 3$ non-zero Schmidt coefficients, and its approximation by truncation to the largest $\tilde{d} \geq 3$ Schmidt coefficients, the sum of the truncated Schmidt coefficients $\epsilon_{\tilde{d}}$ is bounded as follows*

$$\epsilon_{\tilde{d}} \leq \frac{H(\lambda)}{\frac{\tilde{d}-2}{d-\tilde{d}}\log_2(d-\tilde{d}) + \log_2(\tilde{d})}, \tag{4}$$

where $H(\lambda)$ is the Shannon entropy of the Schmidt coefficients.

By recasting this theorem in terms of q-samples and by further loosening Eq. (4) we obtain the following corollary:

Corollary 1 *Consider a q-sample $|P\rangle$ for which the iMPS representation has $d \geq 3$ non-zero Schmidt coefficients and can be constructed sequentially by a memory with an information cost C_q. Consider also an approximation to this q-sample $|\tilde{P}\rangle$ with truncated dimension $\tilde{d} \geq 3$ formed by truncating the $d - \tilde{d}$ smallest Schmidt coefficients of $|P\rangle$. The sum of the truncated Schmidt coefficients $\epsilon_{\tilde{d}}$ satisfies*

$$\epsilon_{\tilde{d}} \leq \frac{C_q}{\tilde{D}_q}, \tag{5}$$

where $\tilde{D}_q := \log_2(\tilde{d})$.

Further, by combining this with Theorem 1, we have that for any q-sample $|P\rangle$ of memory dimension $d \geq 3$ and information cost C_q, there exists a q-sample $|\tilde{P}\rangle$ of memory dimension $\tilde{d}$ that approximates $|P\rangle$ satisfying

$$R_F(|P\rangle, |\tilde{P}\rangle) \leq O\left(\frac{C_q}{\tilde{D}_q}\right), \tag{6}$$

where $\tilde{D}_q = \log_2(\tilde{d})$. This indicates that if $C_q \ll D_q$, there is significant scope for truncating the memory dimension without introducing significant error, affirming our earlier intuition.

Computational Approach – To carry out our quantum dimension reduction for q-sample construction, we use the *variational truncation approach* for iMPS, based on tangent-space methods [49, 32, 50]. Given an injective iMPS $\{A^x\}$ of bond dimension d, this approach seeks to find an injective iMPS $\{\tilde{A}^x\}$ of dimension $\tilde{d}$ of maximum fidelity with the target iMPS. This is equivalent to finding the iMPS representation of a q-sample $|\tilde{P}\rangle$ with memory cost $\tilde{D}_q = \log_2(\tilde{d})$ that minimises the QFDR

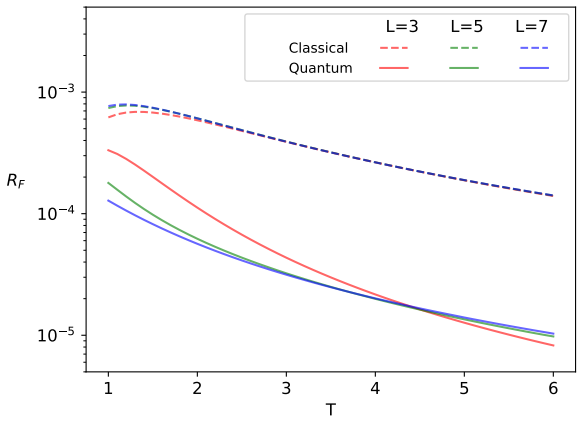

Figure 2: **Divergence rates for 1 (qu)bit q-samples and models of the Dyson-Ising chain process.**

with respect to a target q-sample $|P\rangle$ with memory cost $D_q = \log_2(d)$. We then use the iMPS representation of the approximating q-sample to prescribe a circuit that can be used to construct it. We apply our algorithm to an example known as the Dyson-Ising spin chain.

Consider the thermal state of Dyson-Ising Hamiltonian at temperature T, whose dynamics are described by the (classical) Hamiltonian $H_{\mathrm{DI}} = \sum_{j,k} J(j,k)\sigma_j\sigma_k$. The probability of a given spin configuration $\overleftrightarrow{\sigma}$ is given by $P(\overleftrightarrow{\sigma}) = \frac{1}{Z}e^{-\frac{H_{\mathrm{DI}}^{(L)}}{T}}$, where the normalisation Z is the partition function. This can be viewed as a stochastic process, where we sweep sequentially along the chain spin by spin, and consider the distribution for the state of the next spin given the previous spins seen [41]. It can be shown that the corresponding conditional distribution for a given site is a function of only the last L spins – it is a Markov order L process. Thus, as L is increased the non-Markovianity of the process increases.

As can be seen in Fig. 2(a), for high T our reduced memory q-samples achieve QFDRs below 10^{-3}qubits/timestep, even as L increases. Intriguingly, for low T, the QFDR actually *decreases* with L, despite the increasing non-Markovianity. While we lack a rigorous argument, we suggest that this could be because the truncated q-sample circuits tend to display infinite Markov order behaviour in their statistics, and so in certain circumstances the reduced memory q-samples may more naturally fit to processes with infinite Markov order.

Conclusion – Our approach to quantum dimension reduction for quantum sampling and quantum stochastic modelling may provide a route to leveraging quantum approaches for more efficient feature extraction [51, 52], a task of vital importance in a world that is becoming ever-increasingly data-rich and data-intensive. We have substantiated that quantum models can in effect do more with less, requiring smaller memories to capture expressivity not possible with larger classical memories. Moreover, by reducing the quantum resources required for such tasks, we bring them increasingly into reach of quantum processors of the nearer future.

References

[1] David Kulp, David Haussler, M G Reese, and Eeckman F H. A generalized hidden Markov model for the recognition of human genes in DNA. In *Proceedings of the 4th International Conference on Intelligent Systems for Molecular Biology*, pages 134–142, 1996.

[2] Volker Grimm, Eloy Revilla, Uta Berger, Florian Jeltsch, Wolf M Mooij, Steven F Railsback, Hans-Hermann Thulke, Jacob Weiner, Thorsten Wiegand, and Donald L DeAngelis. Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *science*, 310(5750):987–991, 2005.

[3] Daniel T Gillespie. Stochastic simulation of chemical kinetics. *Annual Reviews of Physical Chemistry*, 58:35–55, 2007.

[4] Darren J Wilkinson. Stochastic modelling for quantitative description of heterogeneous biological systems. *Nature Reviews Genetics*, 10(2):122–133, 2009.

[5] Peter E Smouse, Stefano Focardi, Paul R Moorcroft, John G Kie, James D Forester, and Juan M Morales. Stochastic modelling of animal movement. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1550):2201–2211, 2010.

[6] Elsa Garavaglia and Raffaella Pavani. About earthquake forecasting by Markov renewal processes. *Methodology and Computing in Applied Probability*, 13(1):155–169, 2011.

[7] Stuart Bartlett, Jiazheng Li, Lixiang Gu, Lana Sinapayen, Siteng Fan, Vijay Natraj, Jonathan H Jiang, David Crisp, and Yuk L Yung. Assessing planetary complexity and potential agnostic biosignatures using epsilon machines. *Nature Astronomy*, 6(3):387–392, 2022.

[8] W Brian Arthur. Complexity and the economy. *Science*, 284(5411):107–109, 1999.

[9] Jan Bulla and Ingo Bulla. Stylized facts of financial time series and hidden semi-Markov models. *Computational Statistics & Data Analysis*, 51(4):2192–2209, 2006.

[10] Joongwoo Brian Park, Jeong Won Lee, Jae Suk Yang, Hang Hyun Jo, and Hie Tae Moon. Complexity analysis of the stock market. *Physica A: Statistical Mechanics and its Applications*, 379(1):179–187, 2007.

[11] Jae Suk Yang, Wooseop Kwak, Taisei Kaizoji, and In Mook Kim. Increasing market efficiency in the stock markets. *European Physical Journal B*, 61(2):241–246, 2008.

[12] Sven Maerivoet and Bart De Moor. Cellular automata models of road traffic. *Physics reports*, 419(1):1–64, 2005.

[13] Stephen E Levinson. Continuously variable duration hidden Markov models for automatic speech recognition. *Computer Speech & Language*, 1(1):29–45, 1986.

[14] Lawrence R. Rabiner. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

[15] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 04 1970.

[16] Walter R Gilks, Sylvia Richardson, and David Spiegelhalter. *Markov chain Monte Carlo in practice*. CRC press, 1995.

[17] Reuven Y Rubinstein and Dirk P Kroese. *Simulation and the Monte Carlo method*, volume 10. John Wiley & Sons, 2016.

[18] Carsten Blank, Daniel K Park, and Francesco Petruccione. Quantum-enhanced analysis of discrete stochastic processes. *npj Quantum Information*, 7:1–9, 2021.

[19] Patrick Rebentrost, Brajesh Gupt, and Thomas R Bromley. Quantum computational finance: Monte carlo pricing of financial derivatives. *Physical Review A*, 98(2):022321, 2018.

[20] Stefan Woerner and Daniel J Egger. Quantum risk analysis. *npj Quantum Information*, 5(1):1–8, 2019.

[21] Nikitas Stamatopoulos, Daniel J Egger, Yue Sun, Christa Zoufal, Raban Iten, Ning Shen, and Stefan Woerner. Option pricing using quantum computers. *Quantum*, 4:291, 2020.

[22] Xun Gao, Eric R. Anschuetz, Sheng-Tao Wang, J. Ignacio Cirac, and Mikhail D. Lukin. Enhancing generative models via quantum correlations. *Phys. Rev. X*, 12:021037, May 2022.

[23] Thomas J Elliott, Chengran Yang, Felix C Binder, Andrew J P Garner, Jayne Thompson, and Mile Gu. Extreme dimensionality reduction with quantum modeling. *Physical Review Letters*, 125:260501, 2020.

[24] T J Elliott. Quantum coarse graining for extreme dimension reduction in modeling stochastic temporal dynamics. *PRX Quantum*, 2:020342, 2021.

[25] Thomas J. Elliott, Mile Gu, Andrew J. P. Garner, and Jayne Thompson. Quantum adaptive agents with efficient long-term memories. *Phys. Rev. X*, 12:011007, Jan 2022.

[26] Maria Schuld, Francesco Petruccione, Maria Schuld, and Francesco Petruccione. Quantum advantages. *Supervised Learning with Quantum Computers*, pages 127–137, 2018.

[27] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum random access memory. *Physical Review Letters*, 100(16):160501, April 2008.

[28] Anirban N. Chowdhury, Guang Hao Low, and Nathan Wiebe. A Variational Quantum Algorithm for Preparing Quantum Gibbs States, January 2020.

[29] Felix C Binder, Jayne Thompson, and Mile Gu. Practical unitary simulator for non-Markovian complex processes. *Physical Review Letters*, 120(24):240502, 2018.

[30] Qing Liu, Thomas J. Elliott, Felix C. Binder, Carlo Di Franco, and Mile Gu. Optimal stochastic modeling with unitary quantum dynamics. *Phys. Rev. A*, 99:062110, Jun 2019.

[31] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of physics*, 349:117–158, 2014.

[32] Bram Vanhecke, Maarten Van Damme, Jutho Haegeman, Laurens Vanderstraeten, and Frank Verstraete. Tangent-space methods for truncating uniform mps. *SciPost Physics Core*, 4(1):004, 2021.

[33] Chengran Yang, Felix C. Binder, Varun Narasimhachar, and Mile Gu. Matrix product states for quantum stochastic modeling. *Phys. Rev. Lett.*, 121:260602, Dec 2018.

[34] James P Crutchfield. Between order and chaos. *Nature Physics*, 8(1):17–24, 2012.

[35] Alexander Khintchine. Korrelationstheorie der stationären stochastischen Prozesse. *Mathematische Annalen*, 109(1):604–615, 1934.

[36] Cosma Rohilla Shalizi and James P Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *Journal of Statistical Physics*, 104(3):817–879, 2001.

[37] James P Crutchfield and Karl Young. Inferring statistical complexity. *Physical Review Letters*, 63(2):105, 1989.

[38] Mile Gu, Karoline Wiesner, Elisabeth Rieper, and Vlatko Vedral. Quantum mechanics can reduce the complexity of classical models. *Nature Communications*, 3:762, 2012.

[39] John R. Mahoney, Cina Aghamohammadi, and James P. Crutchfield. Occam's Quantum Strop: Synchronizing and Compressing Classical Cryptic Processes via a Quantum Channel. *Scientific Reports*, 6(1):20495, 2016.

[40] Farzad Ghafari, Nora Tischler, Jayne Thompson, Mile Gu, Lynden K Shalm, Varun B Verma, Sae Woo Nam, Raj B Patel, Howard M Wiseman, and Geoff J Pryde. Dimensional quantum memory advantage in the simulation of stochastic processes. *Physical Review X*, 9(4):041013, 2019.

[41] Cina Aghamohammadi, John R. Mahoney, and James P. Crutchfield. Extreme Quantum Advantage when Simulating Classical Systems with Long-Range Interaction. *Scientific Reports*, 7(1):6735, sep 2017.

[42] Andrew J P Garner, Qing Liu, Jayne Thompson, Vlatko Vedral, and Mile Gu. Unbounded memory advantage in stochastic simulation using quantum mechanics. *New Journal of Physics*, 19(10):103009, 2017.

[43] Thomas J Elliott and Mile Gu. Superior memory efficiency of quantum devices for the simulation of continuous-time stochastic processes. *npj Quantum Information*, 4:18, 2018.

[44] Thomas J Elliott, Andrew J P Garner, and Mile Gu. Memory-efficient tracking of complex temporal and symbolic dynamics with quantum simulators. *New Journal of Physics*, 21(1):013021, jan 2019.

[45] Guifré Vidal. Efficient classical simulation of slightly entangled quantum computations. *Physical review letters*, 91(14):147902, 2003.

[46] Andrew John Daley, Corinna Kollath, Ulrich Schollwöck, and Guifré Vidal. Time-dependent density-matrix renormalization-group using adaptive effective hilbert spaces. *Journal of Statistical Mechanics: Theory and Experiment*, 2004(04):P04005, 2004.

[47] Steven R White and Adrian E Feiguin. Real-time evolution using the density matrix renormalization group. *Physical review letters*, 93(7):076401, 2004.

[48] R. Orús and G. Vidal. Infinite time-evolving block decimation algorithm beyond unitary evolution. *Physical Review B - Condensed Matter and Materials Physics*, 78(15):155117, 2008.

[49] V. Zauner-Stauber, L. Vanderstraeten, M. T. Fishman, F. Verstraete, and J. Haegeman. Variational optimization algorithms for uniform matrix product states. *Phys. Rev. B*, 97:045145, Jan 2018.

[50] Laurens Vanderstraeten, Jutho Haegeman, and Frank Verstraete. Tangent-space methods for uniform matrix product states. *SciPost Physics Lecture Notes*, (7), 2019.

[51] Isabelle Guyon and André Elisseeff. An introduction to feature extraction. In *Feature extraction: foundations and applications*, pages 1–25. Springer, 2006.

[52] Samina Khalid, Tehmina Khalil, and Shamila Nasreen. A survey of feature selection and feature extraction techniques in machine learning. In *2014 science and information conference*, pages 372–378. IEEE, 2014.