

PAPER • OPEN ACCESS

Memory-efficient tracking of complex temporal and symbolic dynamics with quantum simulators

To cite this article: Thomas J Elliott *et al* 2019 *New J. Phys.* **21** 013021

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



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the **collection** - **download the first chapter of every title for free.**



OPEN ACCESS

RECEIVED

4 September 2018

ACCEPTED FOR PUBLICATION

12 December 2018

PUBLISHED

28 January 2019

Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#).

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



PAPER

Memory-efficient tracking of complex temporal and symbolic dynamics with quantum simulators

Thomas J Elliott^{1,2,5} , Andrew J P Garner^{3,4} and Mile Gu^{1,2,4}¹ Complexity Institute, Nanyang Technological University, 637335, Singapore² School of Physical and Mathematical Sciences, Nanyang Technological University, 637371, Singapore³ Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, Boltzmanngasse 3, A-1090 Vienna, Austria⁴ Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543, Singapore⁵ Author to whom any correspondence should be addressed.E-mail: physics@tjelliott.net and mgu@quantumcomplexity.org**Keywords:** quantum simulators, complex processes, computational mechanics, stochastic processes, continuous-time stochastic processes, quantum memories, quantum advantage

Abstract

Tracking the behaviour of stochastic systems is a crucial task in the statistical sciences. It has recently been shown that quantum models can faithfully simulate such processes whilst retaining less information about the past behaviour of the system than the optimal classical models. We extend these results to general temporal and symbolic dynamics. Our systematic protocol for quantum model construction relies only on an elementary description of the dynamics of the process. This circumvents restrictions on corresponding classical construction protocols, and allows for a broader range of processes to be modelled efficiently. We illustrate our method with an example exhibiting an apparent unbounded memory advantage of the quantum model compared to its optimal classical counterpart.

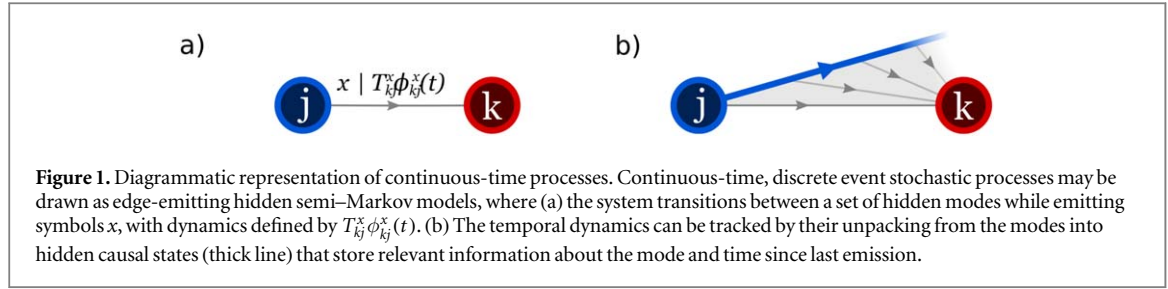
1. Introduction

Continuous-time stochastic processes are omnipresent across the sciences. They are used to model a rich and diverse range of systems [1], such as speech recognition [2, 3], financial time-series [4], neuronal spike trains [5, 6], gene recognition [7], Internet traffic [8], and geophysical processes [9]. Given this broad applicability, our ability to study, simulate, and make predictions using such models is of great import. However, simulations of such models can become highly resource-intensive, in part due to their continuous nature. The information that must be tracked about the past of the process typically diverges with increased precision [6, 10, 11].

Computational mechanics [12, 13] provides a toolset that may be employed in optimising the use of certain resources. Stemming from notions of structural complexity in stochastic processes [14–16], it prescribes a framework for obtaining minimal memory predictive models of a process. This has been applied to a panoply of discrete-time processes [6, 17–24], but only recently have similar studies been made for continuous-time processes, in restricted settings [11, 25].

In parallel, the field of *quantum computational mechanics* has emerged [26–38]. A central result arising from these works is that quantum models of stochastic processes can operate whilst tracking less information about the past than even the optimal classical models [26]. This quantum advantage can be immense, and the gap between quantum and optimal classical memory requirements can grow unbounded [32–34]. As with classical computational mechanics, the focus has largely been on discrete-time processes, and hitherto the quantum computational mechanics of continuous-time processes has been restricted to tracking only a limited set of temporal dynamics (renewal processes), where the times between consecutive emissions are all drawn from the same distribution [34]. Nevertheless, it was found that quantum models of such processes can exhibit unbounded advantages, requiring only finite memory to predict processes that classically need infinite past information to model.

Here, we develop these concepts into a systematic prescription for producing quantum models that simultaneously track temporal and symbolic dynamics in complex stochastic processes, for which the



information about the past that must be stored for future prediction is compressed beyond classical limits. After introducing the computational mechanics framework, we provide a protocol that constructs quantum models of multi-symbol, multi-state stochastic processes in both the discrete- and continuous-time regimes, for which causal equivalence relations are automatically satisfied. This represents the most general construction for quantum models of complex stochastic processes thus far, going far beyond existing quantum constructions and opening the possibility to study a much broader range of complex systems. Moreover, as our construction protocol can be applied to general temporal and symbolic dynamics, it also covers a broader range of systems than existing construction protocols for optimal classical models of stochastic processes. Further, our models exhibit an entropic quantum advantage in the memory required over optimal classical models, even in such cases where systematic construction protocols for these optimal classical models are unknown. We illustrate these features with an example showcasing these advantages, and conclude by highlighting avenues for future development.

2. Framework


We consider continuous-time, discrete event stochastic processes. Such processes are characterised by a sequence (x_n, t_n) detailing what is observed, and when. The emitted symbols $x_n \in \mathcal{A}$ denote the events, while t_n records the time elapsed between events $n-1$ and n . Sequences occur with probabilities drawn from $P(X_n, T_n)$ [39] (upper-case denotes random variables, and lower-case their corresponding realisations). We use shorthand notation $x_n = (x_n, t_n)$, and denote contiguous strings of observations by the concatenation $x_{l:m} = x_l x_{l+1} \dots x_{m-1}$. We restrict our attention to stationary processes, wherein $P(X_{0:L}) = P(X_{s:s+L}) \forall s, L \in \mathbb{Z}$. This framework accommodates emissions that take place either as instantaneous events separated by times t_n , or as continuous emissions with dwell time t_n . We focus primarily on the former, and later discuss how our results may be modified for the latter. Further, though our primary focus will be on continuous-time processes, many of the results can be applied to temporal discretisations of such processes, and we will provide an explicit construction for these coarse-grained analogues.

The observation sequence can be partitioned into past and future. We take 0 as the current emission step, such that x_0 is the next symbol to be emitted, and define $t_{\overleftarrow{0}}(t_{\overrightarrow{0}})$ as the time since the last (until the next) emission, such that $t_0 = t_{\overleftarrow{0}} + t_{\overrightarrow{0}}$. We delineate the past as $\overleftarrow{x} = x_{-\infty:0}(\emptyset, t_{\overleftarrow{0}})$ (\emptyset signifies the event x_0 is currently unknown), and the future as $\overrightarrow{x} = (x_0, t_{\overrightarrow{0}})x_{1:\infty}$ [25].

Such processes can be represented by edge-emitting hidden Semi-Markov models (eeHSMM) [25]. These are defined by a set of hidden modes $\{g\}$, an emission alphabet \mathcal{A} , and a transition dynamic $T_{kj}^x \phi_{kj}^x(t)$. The transition dynamic describes the probability density that the system, upon transitioning to mode j , will subsequently reside in this mode for a time t , at which point it will transition to mode k while emitting symbol x . The $\phi_{kj}^x(t)$ are normalised, such that T_{kj}^x describes the total probability that the system transitions from j to k while emitting x without reference to the time. We can represent such models diagrammatically (figure 1(a)). Semi-Markov [1] refers to the property that the transition dynamic depends only on the current mode and dwell time, such that the causal pair $(g, t_{\overleftarrow{0}})$ gives the fullest possible description for predicting the future of the process that may be obtained from past observations.

We desire models of a process that can faithfully reproduce future statistics given a particular past, and that are *causal*, i.e. contain no information about the future that is not obtainable from the past [37]. Computational mechanics [12, 13] provides a pathway for determining the optimal (storing minimal information) classical models. The basis of computational mechanics are causal states \mathcal{S} , equivalence classes agglomerating pasts with identical future predictions [16]. Two pasts \overleftarrow{x} and \overleftarrow{x}' belong to the same causal state (are causally equivalent \sim_c) iff they have identical conditional future probabilities:

$$\overleftarrow{x} \sim_c \overleftarrow{x}' \Leftrightarrow P(\overrightarrow{x} | \overleftarrow{x}) = P(\overrightarrow{x} | \overleftarrow{x}'). \quad (1)$$



It has been proven that using these causal states as the hidden states of a model provides the optimal classical predictive representation [12]. For discrete-time processes, these form edge-emitting hidden Markov models, while for continuous-time processes one has a continuum of hidden states the system traverses along, jumping into a hidden ‘start’ state (mode of an eeHSM) upon emission (figure 1(b)). The optimal classical predictive models are called ε -machines, and are unifilar: given knowledge of a prior causal state and the observation sequence since, the present causal state is known with certainty [12]. The information required by the ε -machine to track the process is known as the statistical complexity C_μ . This is given by the Shannon entropy (in bits) of the steady-state distribution $P(S)$ over causal states [12, 16]:

$$C_\mu = -\sum_{s \in S} P(s) \log_2(P(s)). \quad (2)$$

This quantity is generically larger than the information shared between the past and future of the process [12], indicating that even these optimal models must store redundant information about the past. Indeed, ε -machines are wasteful whenever there is stochasticity in the transition dynamic of the causal states.

Though much of the focus of computational mechanics has been on discrete-time symbolic processes, recently analogous results have emerged for continuous-time processes tracking purely temporal [11], and both symbolic and temporal dynamics together [25]. While the optimality proofs hold for such processes, systematic construction protocols for finding causal states are known only for a limited set of processes. The first such class are renewal processes. These describe purely temporal dynamics where all emission symbols are identical, and the times between each consecutive pair of emissions are independent and identically distributed according to a common ‘waiting time’ function [11]. Recently [25], the causal architectures of more general processes with complex symbolic and temporal dynamics have been uncovered, albeit only when they satisfy certain (quite stringent) restrictions. Assuming the modes $\{g\}$ are already expressed in the minimal unifilar representation, the process must satisfy:

- (i) unifilarity (i.e. synchronisability) of the modes with regards to the observed symbol sequence alone. That is, the mode the system transitions into on the next emission must depend only on the current mode and emitted symbol, and not the time the emission takes place;
- (ii) emission-time distributions depend only on the current mode [$\phi_{kj}^x(t) = \phi_j(t) \forall k, x$]. That is, the time at which emissions occur are independent of the symbol emitted, given the mode;
- (iii) transition dynamics $T_{kj}^x \phi_j(t)$ must be such that pairs (g_0, t_0^-) are not only sufficient for future prediction, but are also minimal. That is, processes where different pairs (j, t_0^-) and $(k, t_0'^-)$ can become causally equivalent are forbidden. This prohibits processes where two modes, conditioned on times since last emission, lead to identical future statistics. This also rules out processes where the conditional emission-time distribution of a mode becomes periodic after a given time since last emission.

These restrictions are illustrated in figure 2. They also apply to construction protocols for models tracking both temporal and symbolic dynamics together in discrete-time (to our knowledge, no prior work has explicitly covered this latter scenario).

Our quantum models will not be subject to these restrictions. Moreover, unlike these classical works, we shall not employ differential entropies for continuous-time processes, instead preserving the operational meaning of statistical complexity as the information the model must store about the past. A consequence of this is that whenever the temporal dynamic is not wholly memoryless (i.e. the $\phi_{kj}^x(t)$ are not all Poisson distributions) C_μ will diverge, as the ε -machine is ultimately storing a continuous parameter [10, 11, 25, 32, 34].

3. Quantum models for tracking complex temporal and symbolic dynamics

While causal states eliminate redundancy in storing information that distinguishes pasts with identical future statistics, they provide no savings when two pasts have similar, yet non-identical futures. This is because states must be either identical or fully distinguishable in classical information theory. In contrast, quantum information [40] can be encoded into states that are only partly distinguishable, and this may be used to reduce the past information that must be retained [26]. Specifically, causal states can be encoded into quantum states whose overlap increases with the overlap between their corresponding futures. Labelling the information tracked by such quantum models ('q-machines') as M_q , we have that generically $M_q \leq C_\mu$, with equality only for processes with no stochasticity in the hidden states of the model [26] (that is, whenever the ε -machine stores redundant information, a q-machine can mitigate some of this redundancy). This reduced entropy bears operational advantages when considering storage or communication of the states of ensembles of simulators of a process (in some cases the advantage may also manifest in the single-shot case [37]; we do not explore this regime here). The use of M_q rather than C_q denotes that the q-machine may not be optimal, and hence M_q is not necessarily the quantum statistical complexity [31] (but provides an upper bound for it).

We now provide a systematic protocol that constructs q-machines for simultaneously tracking complex temporal and symbolic dynamics in continuous-time processes. In particular, we show that in general the resultant models:

1. produce accurate future predictions given the past;
2. can be operated in a continuous manner;
3. can be synchronised from the past (are causal);
4. automatically satisfy causal equivalence relations;
5. store less information than any classical model.

As with the classical case, we assume the minimal unifilar modes $\{g\}$ of the process' eeHSM have been determined. This can be achieved using techniques adapted from discrete-time computational mechanics [16, 41], taking the dual x_n as effective emitted symbols. Given these minimal modes, a sufficient, causal set of parameters is the pair (g_0, t_0^-) . A model based on these parameters then need only to identify (with correct probabilities) whether emission occurs in the next infinitesimal interval dt , what the emission and subsequent mode are, and to update to be in the corresponding state. The causal states correspond to groupings of such pairs with identical future predictions. Specifically,

$$(g_0, t_0^-) \sim_e (g'_0, t_0'^-) \Leftrightarrow P(G_1, X_0, T_0^- | g_0, t_0^-) = P(G_1, X_0, T_0^- | g'_0, t_0'^-). \quad (3)$$

We introduce several quantities to characterise the processes. First, we define as shorthand for the transition dynamic

$$\psi_{kj}^x(t) = \sqrt{T_{kj}^x \phi_{kj}^x(t)}. \quad (4)$$

Next, the modal steady-state distribution π_j is defined as the (unique [42]) eigenvector of $\sum_x T_{kj}^x$ with unit eigenvalue, normalised such that $\sum_j \pi_j = 1$. These π_j are the steady-state probabilities that the system is in mode j immediately after an emission. We further define the mode survival probability (the probability that the dwell time in mode j is at least t):

$$\Phi_j(t) = \sum_{xk} \int_t^\infty \psi_{kj}^x(t')^2 dt'. \quad (5)$$

We also define the mode lifetime (the average dwell time for mode j):

$$\tau_j = \sum_{xk} \int_0^\infty t \psi_{kj}^x(t)^2 dt, \quad (6)$$

and average emission lifetime $\tau = \sum_j \pi_j \tau_j$. Finally, the modal and mean firing rates are given by reciprocals of the respective lifetimes: $\mu_j = 1/\tau_j$ and $\mu = 1/\tau$.

With these definitions, we can express

$$P(G_1 = k, X_0 = x, T_0^- = t' | G_0 = j, T_0^- = t) = \frac{\psi_{kj}^x(t + t')^2}{\Phi_j(t)}. \quad (7)$$

From these, we define the associated quantum memory states (QMS) for each pair $(g_0, t_0^-) = (j, t)$:

$$|\varsigma_j(t)\rangle = \sum_{xk} \int_0^\infty \frac{\psi_{kj}^x(t+t')}{\sqrt{\Phi_j(t)}} dt' |t'\rangle |x\rangle |k\rangle, \quad (8)$$

where the function inside the integral is the square root of the conditional probability for the future (equation (7)). The QMS belong to a tripartite composite Hilbert space. The first of these is a continuous space encoding the statistics of the remaining dwell time, while the others are discrete, and correspond to tracking the emitted symbol statistics and subsequent mode respectively. Measurement of the first two spaces in the $|t'\rangle |x\rangle$ basis yields outcomes t' and x with probability density $P(T_0^- = t', X_0 = x | G_0 = j, T_0^- = t)$, and leaves the final subspace in $|k\rangle$, flagging the mode k to which the system transitions after such an emission. Mapping state $|k\rangle$ (with appropriate blank ancillae) to $|\varsigma_k(0)\rangle$ sets the model in the appropriate post-emission QMS. A measurement sweep of the time subspace over the range $[0, \delta t]$ yields a non-emission result with the correct probability, and for such a non-emission, with a relabelling $t' \rightarrow t' - \delta t$, will produce the QMS $|\varsigma_j(t + \delta t)\rangle$. Thus, such measurement sweeps can be used to emulate the passage of time in such models. By performing these measurement sweeps and mappings of flag states for the subsequent post-emission modes, the QMS can model the corresponding stochastic process, fulfilling (C1) and (C2).

As QMS are clearly well-defined for each pair (g_0, t_0^-) , the q-machine is causal, satisfying (C3). Further, as the QMS depend only on the conditional probabilities that define causal equivalence, they satisfy (C4):

$$(g_0, t_0^-) \sim_e (g_0', t_0'^-) \Leftrightarrow |\varsigma_{g_0}(t_0^-) \rangle = |\varsigma_{g_0'}(t_0'^-) \rangle. \quad (9)$$

Thus, QMS corresponding to pasts in the same causal state are identical: *they automatically adopt the causal architecture of the process without the need to explicitly apply the causal equivalence relation*. Further, QMS corresponding to pasts in different causal states generally have non-zero overlap, given by the fidelity of the corresponding conditional probability distributions (equation (7)). Whenever these overlaps are not all either zero or unity (wherein the ε -machine exhibits no stochasticity in its hidden states, and stores no redundant information [26]) the QMS steady-state distribution has lower entropy than that of the classical causal states, and *the q-machine stores less information than the corresponding ε -machine*, satisfying (C5).

The information stored in the q-machine may be calculated by determining the Shannon entropy of the spectrum of the steady-state density matrix

$$\rho_q = \sum_j \int P(j, t) |\varsigma_j(t)\rangle \langle \varsigma_j(t)| dt, \quad (10)$$

where $P(j, t) = \mu \pi_j \Phi_j(t)$ is the steady-state distribution of the QMS (see appendix A). Using a Gram matrix approach [42], we can construct a characteristic equation for the spectrum expressed in terms of the steady-state distribution and overlaps of QMS (see appendix B):

$$\mu \sum_{mkx} \sqrt{\pi_j \pi_k} \iint \psi_{mj}^x(t+a) \psi_{mk}^x(t+b) f_{nk}(b) dt db = \lambda_n f_{nj}(a). \quad (11)$$

The memory required by the q-machine is then given by

$$M_q = \sum_n \lambda_n \log_2(\lambda_n). \quad (12)$$

These λ_n can be determined by solving the integral equation in (11).

The QMS can be adapted to model coarse-grained versions of a process, wherein continuous time is replaced by small, discrete timesteps. Analogous to classical coarse-graining, in which all states within a given timestep δt are merged into a single state prior to applying the causal equivalence relation, we likewise merge the QMS in a particular timestep to form coarse-grained QMS, with probabilities defined by integrals of probability densities over timesteps. The corresponding coarse-grained QMS are given by

$$|\tilde{\varsigma}_j(n)\rangle = \sum_{xk} \sum_{n'} \frac{\tilde{\psi}_{kj}^x(n+n')}{\sqrt{\Phi(n\delta t)}} |n'\rangle |x\rangle |k\rangle, \quad (13)$$

where $\tilde{\psi}_{kj}^x(n)^2 = \int_{n\delta t}^{(n+1)\delta t} \psi_{kj}^x(t)^2 dt$. Outcome n from measurement of the first subspace corresponds to $t_0^- = n\delta t$. As the coarse-grained states will generally not be mutually orthogonal, a quantum advantage remains. The continuous-time case is recovered in the limit $\delta t \rightarrow 0$.

For continuously-emitting processes, the emitted symbols are determined immediately after transition events, rather than immediately before. The current emission is then known throughout the current step. An appropriate modification to the protocol is to make a measurement of the symbol subspace in the new QMS immediately after the previous transition event occurs, rather than on the old QMS. The QMS are now also labelled by the current symbol, with

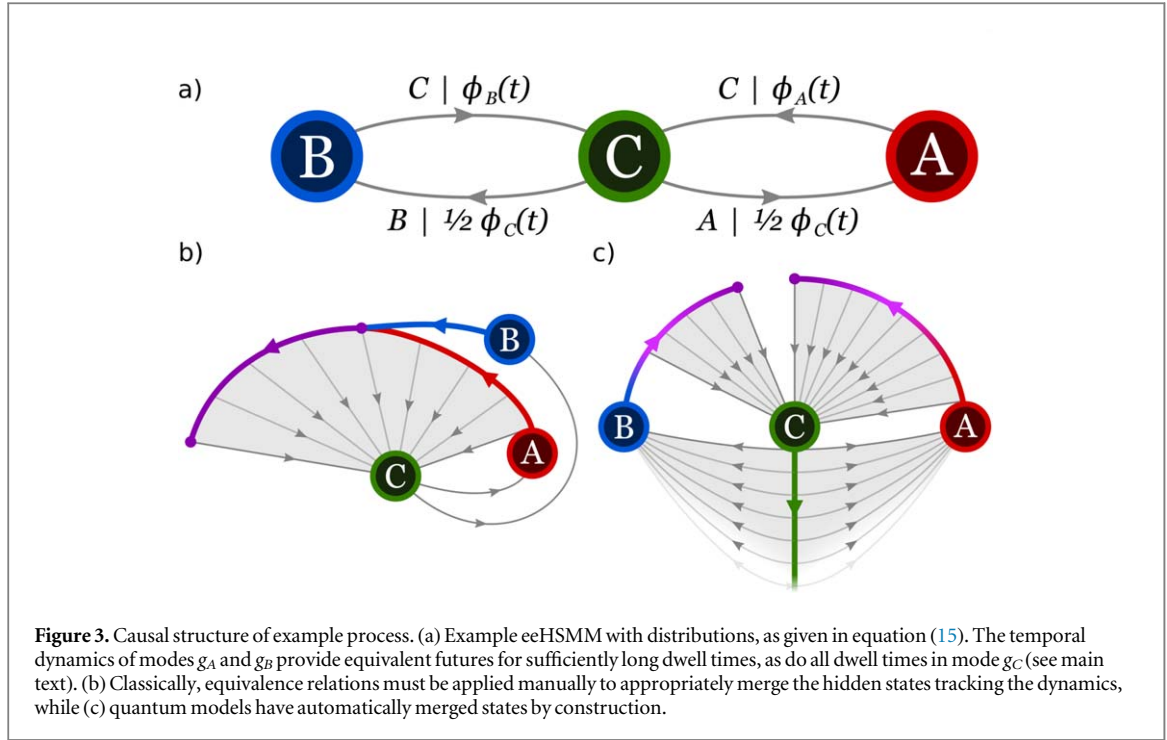


Figure 3. Causal structure of example process. (a) Example eeHSM with distributions, as given in equation (15). The temporal dynamics of modes g_A and g_B provide equivalent futures for sufficiently long dwell times, as do all dwell times in mode g_C (see main text). (b) Classically, equivalence relations must be applied manually to appropriately merge the hidden states tracking the dynamics, while (c) quantum models have automatically merged states by construction.

$$|\varsigma_{jx}(t)\rangle = \sum_k \int \delta_{kj}^x \frac{\sqrt{\phi_{kj}^x(t+t')}}{\sqrt{\Phi_j^x(t)}} dt' |t'\rangle |x\rangle |k\rangle, \quad (14)$$

where $\delta_{kj}^x = 1$ if T_{kj}^x is non-zero (and zero otherwise), and $\Phi_j^x(t) = \sum_k \int_t^\infty \delta_{kj}^x \phi_{kj}^x(t') dt'$. As with the discrete event case, one can calculate the steady-state density matrix of these states, and from its spectrum determine the amount of information tracked by the q-machine.

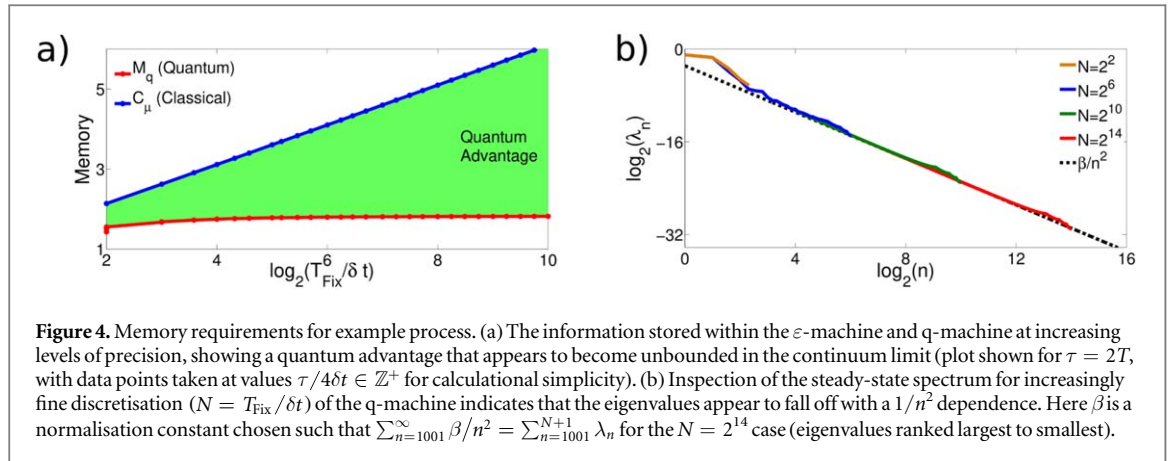
4. Example process with complex temporal and symbolic dynamics exhibiting extreme quantum advantage

To illustrate our results, we employ our q-machine construction protocol to study an example process that violates the restrictions on current systematic ε -machine construction protocols. Consider the following conceptual scenario allegorising a process: Charlie owns a device that exhibits a constant breaking probability within any fixed interval of time. Upon breakage, he takes the device to either Alice or Bob for repair, after which the device is returned to Charlie. The time taken by Alice and Bob to fix the device is a random variable; Alice's fixing time distribution is different to Bob's, but has the same average. Emissions herald when the device changes hands, and indicate the new holder.

Mathematically, this process can be represented by three modes, forming the eeHSM depicted in figure 3(a). Modes g_A and g_B each emit symbol C at times drawn from $\phi_A(t)$ and $\phi_B(t)$ respectively, upon which the system transitions to mode g_C . Mode g_C emits symbols A or B with equal probability at times drawn from Poisson distribution $\phi_C(t)$ and transitions into the corresponding mode g_A or g_B . We define the distributions:

$$\begin{aligned} \phi_A(t) &= \begin{cases} \frac{1}{T_{\text{Fix}}} & 0 \leq t \leq T_{\text{Fix}}, \\ 0 & \text{elsewhere} \end{cases} \\ \phi_B(t) &= \begin{cases} \frac{2}{T_{\text{Fix}}} & \frac{T_{\text{Fix}}}{4} \leq t \leq \frac{3T_{\text{Fix}}}{4}, \\ 0 & \text{elsewhere} \end{cases} \\ \phi_C(t) &= \frac{e^{-t/T_{\text{Brk}}}}{T_{\text{Brk}}}. \end{aligned} \quad (15)$$

We see that this process exhibits causal equivalence between certain of its causal pairs, and as such falls outside the class for which systematic ε -machine construction methods are currently known. Specifically, because $\phi_C(t)$ is a Poisson distribution, all dwell times within this mode have identical conditional futures, and so only a single causal state is needed to describe occupation of this mode [10, 11, 34]:



$$(g_C, t) \sim_e (g_C, t') \forall t, t' \geq 0. \quad (16)$$

Further, we see that the conditional futures for modes g_A and g_B can become identical for certain combinations of dwell times; specifically, we have

$$\left(g_A, t + \frac{T_{\text{Fix}}}{2}\right) \sim_e \left(g_B, t + \frac{T_{\text{Fix}}}{4}\right) \forall t \geq 0. \quad (17)$$

These both exemplify violations of restriction (iii). The conceptual scenario can be easily extended to describe a process in violation of restriction (ii): suppose Charlie's device has two possible faults, with different failure rates, and the choice between Alice and Bob is decided by which of the faults occurred. Restriction (i) would be violated if Charlie were to merely announce that the device has broken (instead of who will be fixing it) and the choice between Alice and Bob is determined by how long the device took to fail.

The relative simplicity of this example allows us to perform the appropriate application of the equivalence relations to merge causal pairs into causal states on an ad hoc basis, as is displayed in figure 3(b). Notably however, in contrast to this we can exploit the self-merging nature of the QMS construction, and blindly assign QMS for each causal pair. These QMS will automatically satisfy the causal equivalence relations, and will not incur any penalty in the information the q-machine must track. This is depicted in figure 3(c).

As described in appendix C, we can calculate the relevant modal and mean lifetimes and firing rates of the process, the coarse-grained causal states and QMS, and their respective steady-state probabilities. Solving the coarse-grained analogue of equation (11) we can determine the steady-state density matrix spectrum, and hence calculate M_q for the particular level of coarse-graining. We can further calculate C_μ at the same coarse-graining for comparison. In figure 4(a) we show that with increasingly fine coarse-graining C_μ diverges, while M_q appears to converge towards a finite value. This suggests an unbounded advantage of the q-machine over the ε -machine, similar to that seen for renewal processes [34]. Inspection of the steady-state density matrix spectrum indicates that it falls off approximately as $1/n^2$ (figure 4(b)). When the spectrum has this dependence, the associated entropy is finite [34] providing further support that M_q is bounded in the continuum limit.

5. Discussion

We have proposed a systematic construction protocol for quantum models of general complex continuous-time stochastic processes that automatically adopt the process' causal architecture, and exhibit an entropic memory advantage over their optimal classical counterparts. Moreover, our model construction protocol can be applied more generally than corresponding protocols for optimal classical models, without any restrictions on the combined symbolic and temporal dynamics. This allows our models to be utilised to study a much broader range of processes, such as the aforementioned examples in our introduction [1–9].

While the information cost of a quantum model constructed by our protocol is lower than any classical model, there is no claim of optimality over all quantum models. By accounting for longer-range temporal correlations [27, 28, 35], it has been found that discrete-time q-machines can reduce their information cost, and analogous constructs may be possible in the continuous-time case, providing further memory savings. Other reductions may be possible by exploiting the possibility of using complex amplitudes in the QMS. Even so, our current models provide an upper bound on the information cost for the optimal quantum model, and a lower bound is given by the mutual information between the past and future of the process [12, 26].

It is interesting to consider how to implement the q-machine as a simulator. Our earlier work on q-machines for continuous-time renewal processes [34] discussed how they might be realised by using the position of a

particle as the continuous variable tracking time, and the motion of the particle towards a detector as the measurement sweep. One can envisage the more general multi-mode, multi-symbol processes here might be implemented similarly, using internal states of the particle for the mode and symbol subspaces of the QMS. We leave specific details of particular experimental implementations as an open question for future work.

Finally, let us remark on the apparent unbounded advantage in our example. In earlier work we have postulated that this is likely a typical feature of continuous-time q-machines [34]. We recapitulate the argument [32, 34] here. Consider a coarse-graining with timesteps δt . With further refinement, the new QMS will typically have large overlap with the existing temporally-adjacent QMS. At very fine coarse-graining, the new states will be almost identical to those existing, and so the increase in information cost will be ever-decreasing, resulting in the observed convergence of M_q . In contrast, the mutual orthogonality of the classical states leads to a logarithmic divergence in information cost. An enticing problem for future work is to develop methods of coarse-graining that exploit such quantum features. We expect that such a quantum coarse-graining could be used for near-exact simulation with extreme memory advantages even in the single-shot regime. It would also be interesting to consider the extension to input-output processes [43, 44] operating in continuous-time.

Acknowledgments

This work was funded by the Lee Kuan Yew Endowment Fund (Postdoctoral Fellowship), Singapore Ministry of Education Tier 1 grant RG190/17, FQXi Large Grants: The role of quantum effects in simplifying adaptive agents and Observer-dependent complexity: The quantum-classical divergence over ‘what is complex?’, John Templeton Foundation grant 53914, and Singapore National Research Foundation Fellowship NRF-NRFF2016-02. T.J.E. thanks the Centre for Quantum Technologies for their hospitality.

Appendix A. Derivation of steady-state distribution for causal pairs

We here derive the steady-state probability distribution of the causal pair (j, t) , where j is the current mode, and t the time since last emission. As stated in the main text, this distribution is given by

$$P(G_0 = j, T_{\overleftarrow{0}} = t) = \mu \pi_j \Phi_j(t). \quad (\text{A1})$$

Let us first decompose

$$P(G_0 = j, T_{\overleftarrow{0}} = t) = P(T_{\overleftarrow{0}} = t | G_0 = j) P(G_0 = j), \quad (\text{A2})$$

and find each of the constituent probabilities.

Beginning with the former, recall that the conditional probability density of emission from mode j at time t' from the present given a prior wait time of t is given by

$$P(T_{\overleftarrow{0}} = t' | T_{\overleftarrow{0}} = t, G_0 = j) = \sum_{kx} \frac{T_{kj}^x \phi_{kj}^x(t + t')}{\Phi_j(t)}. \quad (\text{A3})$$

The conditional survival probability in mode j to time $t + t'$ is thus given by $\Phi_j(t + t') / \Phi_j(t)$. Clearly then, the probability of progressing without an emission attenuates by the survival probability, and we have

$$P(T_{\overleftarrow{0}} = t + t' | G_0 = j) = \frac{\Phi_j(t + t')}{\Phi_j(t)} P(T_{\overleftarrow{0}} = t | G_0 = j). \quad (\text{A4})$$

This is satisfied for all $t, t' \geq 0$, and hence

$$P(T_{\overleftarrow{0}} = t | G_0 = j) = \Phi_j(t) P(T_{\overleftarrow{0}} = 0 | G_0 = j) \quad \forall t \geq 0. \quad (\text{A5})$$

We can further find $P(T_{\overleftarrow{0}} = 0 | G_0 = j)$ by normalising this distribution:

$$P(T_{\overleftarrow{0}} = 0 | G_0 = j)^{-1} = \int_0^\infty \Phi_j(t) dt = \mu_j^{-1}. \quad (\text{A6})$$

Thus,

$$P(T_{\overleftarrow{0}} = t | G_0 = j) = \mu_j \Phi_j(t). \quad (\text{A7})$$

Next, we must determine the probability of a particular mode being occupied in the steady-state. Recall that π_j is the steady-state probability that after an emission the system transitions into mode j . Since the expected time of remaining in mode j is given by the modal lifetime τ_j , we have that the steady-state probability of being in mode j is proportional to $\pi_j \tau_j$. By noticing that the sum of this over all modes is the average emission lifetime τ , we hence obtain

$$P(G_0 = j) = \frac{\pi_j \tau_j}{\tau} = \frac{\mu \pi_j}{\mu_j}. \quad (\text{A8})$$

Thus, putting equations (A7) and (A8) together, we have

$$P(G_0 = j, T_{\tau_0} = t) = \mu \pi_j \Phi_j(t). \quad (\text{A9})$$

Appendix B. Gram matrix and derivation of the characteristic equation

A useful tool for calculating the spectrum of a density matrix is offered by the Gram matrix [42]. The Gram matrix provides an alternative operator that possesses the same spectrum as the original operator of interest, but that may be more amenable for extracting said spectrum. To obtain the Gram matrix, we consider a purification of the QMS steady-state ensemble:

$$|\Psi\rangle_{qG} = \sum_j \int \sqrt{P(j, t)} |\varsigma_j(t)\rangle |t\rangle |j\rangle dt. \quad (\text{B1})$$

By performing a partial trace over the second and third subsystems we obtain the original steady-state density matrix ρ_q . On the other hand, tracing out the first subsystem leaves us with

$$\rho_G = \sum_{jj'} \iint \sqrt{P(j, t) P(j', t')} |\varsigma_j(t)\rangle |\varsigma_{j'}(t')\rangle |t\rangle |j\rangle \langle t'| \langle j'| dt dt'. \quad (\text{B2})$$

The resulting operator ρ_G is called the Gram matrix, and as noted above, has the same spectrum as ρ_q . It has a corresponding characteristic equation

$$\rho_G f_n = \lambda_n f_n, \quad (\text{B3})$$

where f_n are the eigenfunctions, and λ_n the spectrum. Expanding this in terms of the QMS and steady-state distributions, we obtain the characteristic equation of (11), which may be solved to obtain the spectrum λ_n , and hence calculate the information M_q stored by the q-machine.

Appendix C. Further calculational details for the example process

Here we provide the detailed derivations of the properties of the example in the main text. Recall the definition of the process, as depicted in figure 3(a), whereby two modes g_A and g_B each emit symbol C while transitioning to mode g_C at a time drawn from $\phi_A(t)$ or $\phi_B(t)$ respectively, while mode g_C emits A or B and transitions to the corresponding mode with equal probability, at a time drawn from $\phi_C(t)$. The respective distributions are defined by

$$\begin{aligned} \phi_A(t) &= \begin{cases} \frac{1}{T_{\text{Fix}}} & 0 \leq t \leq T_{\text{Fix}}; \\ 0 & \text{elsewhere} \end{cases} \\ \phi_B(t) &= \begin{cases} \frac{2}{T_{\text{Fix}}} & \frac{T_{\text{Fix}}}{4} \leq t \leq \frac{3T_{\text{Fix}}}{4}; \\ 0 & \text{elsewhere} \end{cases} \\ \phi_C(t) &= \frac{e^{-t/T_{\text{Brk}}}}{T_{\text{Brk}}}, \end{aligned} \quad (\text{C1})$$

as illustrated in figure 5.

From these distributions, we can straightforwardly calculate the modal lifetimes and steady-state distributions using the definitions in the main text. These are as follows:

$$\begin{aligned} \tau_A = \tau_B &= \frac{T_{\text{Fix}}}{2}; & \tau_C &= T_{\text{Brk}}; \\ \pi_A = \pi_B &= \frac{1}{4}; & \pi_C &= \frac{1}{2}. \end{aligned} \quad (\text{C2})$$

The average lifetime is $\tau = (T_{\text{Fix}} + 2T_{\text{Brk}})/4$. We can also calculate the mode survival probabilities:

$$\begin{aligned} \Phi_A(t) &= \begin{cases} 1 - \frac{t}{T_{\text{Fix}}} & 0 \leq t \leq T_{\text{Fix}}; \\ 0 & \text{elsewhere} \end{cases} \\ \Phi_B(t) &= \begin{cases} 1 & 0 \leq t < \frac{T_{\text{Fix}}}{4} \\ \frac{3}{2} - \frac{2t}{T_{\text{Fix}}} & \frac{T_{\text{Fix}}}{4} \leq t \leq \frac{3T_{\text{Fix}}}{4}; \\ 0 & \text{elsewhere} \end{cases} \\ \Phi_C(t) &= e^{-t/T_{\text{Brk}}}. \end{aligned} \quad (\text{C3})$$




Figure 5. Emission distributions for example. Dwell times of equations (C1) describing the temporal dynamics of the example.

The corresponding steady-state probabilities are:

$$\begin{aligned}
 P(g_A, t) &= \begin{cases} \frac{1 - \frac{t}{T_{\text{Fix}}}}{T_{\text{Fix}} + 2T_{\text{Brk}}} & 0 \leq t \leq T_{\text{Fix}}, \\ 0 & \text{elsewhere} \end{cases} \\
 P(g_B, t) &= \begin{cases} \frac{1}{T_{\text{Fix}} + 2T_{\text{Brk}}} & 0 \leq t < \frac{T_{\text{Fix}}}{4} \\ \frac{\frac{3}{2} - \frac{2t}{T_{\text{Fix}}}}{T_{\text{Fix}} + 2T_{\text{Brk}}} & \frac{T_{\text{Fix}}}{4} \leq t \leq \frac{3T_{\text{Fix}}}{4} \\ 0 & \text{elsewhere} \end{cases} \\
 P(g_C, t) &= \frac{2e^{-t/T_{\text{Brk}}}}{T_{\text{Fix}} + 2T_{\text{Brk}}}. \quad (\text{C4})
 \end{aligned}$$

While no general prescription is presently known for implementing the causal equivalence relations for continuous-time processes in classical systems [25], their relative simplicity in this example allows for them to be applied on an ad hoc basis, as displayed in figure 3(b). We can merge all times within mode g_C into a single causal state, as well as all pairs of states $(g_A, t + T_{\text{Fix}}/2)$ and $(g_B, t + T_{\text{Fix}}/4)$ for $t \geq 0$. For clarity, in the following we label states by a specific mode only for the unmerged states, and use g_M to represent states in the merge between the g_A and g_B pairs. The corresponding causal states have steady-state probabilities given by

$$\begin{aligned}
 P(A, t) &= \frac{1 - \frac{t}{T_{\text{Fix}}}}{T_{\text{Fix}} + 2T_{\text{Brk}}} & 0 \leq t < \frac{T_{\text{Fix}}}{2} \\
 P(B, t) &= \frac{1}{T_{\text{Fix}} + 2T_{\text{Brk}}} & 0 \leq t < \frac{T_{\text{Fix}}}{4} \\
 P(M, t) &= \frac{3\left(\frac{1}{2} - \frac{t}{T_{\text{Fix}}}\right)}{T_{\text{Fix}} + 2T_{\text{Brk}}} & 0 \leq t < \frac{T_{\text{Fix}}}{2} \\
 P(C) &= \frac{2T_{\text{Brk}}}{T_{\text{Fix}} + 2T_{\text{Brk}}}. \quad (\text{C5})
 \end{aligned}$$

As discussed in the main text, the continuous nature of the distributions leads to the classical statistical complexity C_μ being infinite. However, we can still investigate how the complexity grows with increasingly fine discretisation, and compare this to the quantum case. For calculational simplicity, we use timesteps δt such that $T_{\text{Fix}}/4\delta t$ is an integer, which ensures that each timestep does not span across any borders where the behaviour of the probability distributions change (i.e. there are no ‘partially merged’ timesteps). When discretising, we note that there is a small $O(\delta t)$ correction to the effective lifetimes and firing rates [10], which can be appropriately accounted for by inclusion of a normalisation factor \mathcal{N} such that $\mathcal{N}\sum_{n=0}^{\infty} \Phi(n\delta t) = \mu\pi_j\tau_j$. Calculating the Shannon entropy of the discretised distributions, we obtain C_μ at each level of precision, as plotted in figure 4(a) for the case $T_{\text{Fix}} = 2T_{\text{Brk}}$.

For the q-machine, we must first construct the discretised QMS, as prescribed in the main text:

$$\begin{aligned}
 |\zeta_{g_A}(n)\rangle &= \sum_{n'=0}^{\frac{T_{\text{Fix}}}{\delta t} - (n+1)} \frac{\sqrt{\delta t}}{\sqrt{\frac{T_{\text{Fix}}}{\delta t} - n}} |n'\rangle |C\rangle |g_C\rangle & n = 0, 1, \dots, \frac{T_{\text{Fix}}}{\delta t} - 1 \\
 |\zeta_{g_B}(n)\rangle &= \sum_{n'=\max(0, \frac{T_{\text{Fix}}}{4\delta t} - n)}^{\frac{3T_{\text{Fix}}}{4\delta t} - (n+1)} \frac{\sqrt{\delta t}}{\sqrt{\frac{3T_{\text{Fix}}}{4\delta t} - \max(\frac{T_{\text{Fix}}}{4\delta t}, n)}} |n'\rangle |C\rangle |g_C\rangle & n = 0, 1, \dots, \frac{3T_{\text{Fix}}}{4\delta t} \\
 |\zeta_{g_C}(n)\rangle &= \sum_{n'=0}^{\infty} \frac{e^{-n'\delta t/2T_{\text{Brk}}} \sqrt{1 - e^{-\delta t/T_{\text{Brk}}}}}{\sqrt{2}} |n'\rangle (|A\rangle |g_A\rangle + |B\rangle |g_B\rangle) & n \in \mathbb{N}. \quad (\text{C6})
 \end{aligned}$$

Due to the self-merging nature of the QMS, this blind construction protocol is automatically consistent with the causal equivalence relations satisfied by the classical causal states. By constructing a density matrix of these QMS, with weighting given by the appropriate steady-state probabilities equations (C4) (or rather, discretised analogues thereof), we can find its spectrum, and the associated Shannon entropy provides the quantum

memory requirement M_q . These are plotted for increasingly fine levels of discretisation in figure 4(a) alongside the corresponding classical memory requirements, where it is clear that the q-machine requires less information than the optimal classical machine, and appears to converge to a bounded value, much like was previously found for renewal processes [34].

ORCID iDs

Thomas J Elliott  <https://orcid.org/0000-0001-5392-9241>

Andrew J P Garner  <https://orcid.org/0000-0002-3747-9997>

References

- [1] Yu S-Z 2010 *Artif. Intell.* **174** 215
- [2] Levinson S E 1986 *Comput. Speech Lang.* **1** 29
- [3] Rabiner L R 1989 *Proc. IEEE* **77** 257
- [4] Bulla J and Bulla I 2006 *Comput. Stat. Data Anal.* **51** 2192
- [5] Gerstner W and Kistler W M 2002 *Spiking Neuron Models: Single Neurons, Populations, Plasticity* (Cambridge: Cambridge University Press)
- [6] Marzen S E, De Weese M R and Crutchfield J P 2015 *Frontiers Comput. Neurosci.* **9** 105
- [7] Kulp D, Haussler D, Reese M G and Eeckman F H 1996 *Proc. 4th Int. Conf. on Intelligent Systems for Molecular Biology* pp 134–42
- [8] Yu S-Z, Liu Z, Squillante M S, Xia C and Zhang L 2002 *21st IEEE Int. Performance, Computing, and Communications Conf., 2002 (IEEE)* pp 65–72
- [9] Garavaglia E and Pavani R 2011 *Methodol. Comput. Appl. Probab.* **13** 155
- [10] Marzen S E and Crutchfield J P 2015 *Entropy* **17** 4891
- [11] Marzen S and Crutchfield J P 2017 *J. Stat. Phys.* **168** 109
- [12] Shalizi C R and Crutchfield J P 2001 *J. Stat. Phys.* **104** 817
- [13] Crutchfield J P 2012 *Nat. Phys.* **8** 17
- [14] Grassberger P 1986 *Int. J. Theor. Phys.* **25** 907
- [15] Zambella D and Grassberger P 1988 *Complex Syst.* **2** 269
- [16] Crutchfield J P and Young K 1989 *Phys. Rev. Lett.* **63** 105
- [17] Crutchfield J P and Feldman D P 1997 *Phys. Rev. E* **55** R1239
- [18] Tino P and Koteles M 1999 *IEEE Trans. Neural Netw.* **10** 284
- [19] Palmer A J, Fairall C W and Brewer W A 2000 *IEEE Trans. Geosci. Remote Sens.* **38** 2056
- [20] Clarke R W, Freeman M P and Watkins N W 2003 *Phys. Rev. E* **67** 016203
- [21] Park J B, Lee J W, Yang J-S, Jo H-H and Moon H-T 2007 *Physica A* **379** 179
- [22] Li C-B, Yang H and Komatsuzaki T 2008 *Proc. Natl Acad. Sci.* **105** 536
- [23] Haslinger R, Klinkner K L and Shalizi C R 2010 *Neural Comput.* **22** 121
- [24] Kelly D, Dillingham M, Hudson A and Wiesner K 2012 *PloS One* **7** e29703
- [25] Marzen S E and Crutchfield J P 2017 *J. Stat. Phys.* **169** 303
- [26] Gu M, Wiesner K, Rieper E and Vedral V 2012 *Nat. Commun.* **3** 762
- [27] Mahoney J R, Aghamohammadi C and Crutchfield J P 2016 *Sci. Rep.* **6** 20495
- [28] Riechers P M, Mahoney J R, Aghamohammadi C and Crutchfield J P 2016 *Phys. Rev. A* **93** 052317
- [29] Aghamohammadi C, Mahoney J R and Crutchfield J P 2016 *Phys. Lett. A* **381** 1223
- [30] Palsson M S, Gu M, Ho J, Wiseman H M and Pryde G J 2017 *Sci. Adv.* **3** e1601302
- [31] Suen W Y, Thompson J, Garner A J P, Vedral V and Gu M 2017 *Quantum* **1** 25
- [32] Garner A J P *et al* 2017 *New J. Phys.* **19** 103009
- [33] Aghamohammadi C, Mahoney J R and Crutchfield J P 2017 *Sci. Rep.* **7** 6735
- [34] Elliott T J and Gu M 2018 *npj Quantum Inf.* **4** 18
- [35] Binder F C, Thompson J and Gu M 2018 *Phys. Rev. Lett.* **120** 240502
- [36] Ghafari Jouneghani F, Gu M, Ho J, Thompson J, Suen W Y, Wiseman H M and Pryde G J 2017 arXiv:1711.03661
- [37] Thompson J, Garner A J P, Mahoney J R, Crutchfield J P, Vedral V and Gu M 2018 *Phys. Rev. X* **8** 031013
- [38] Aghamohammadi C, Loomis S P, Mahoney J R and Crutchfield J P 2018 *Phys. Rev. X* **8** 011025
- [39] Khintchine A 1934 *Math. Annalen* **109** 604
- [40] Nielsen M A and Chuang I 2000 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press)
- [41] Shalizi C R 2001 Causal architecture, complexity and self-organization in the time series and cellular automata *PhD Thesis* University of Wisconsin-Madison
- [42] Horn R A and Johnson C R 1990 *Matrix Analysis* (Cambridge: Cambridge University Press)
- [43] Barnett N and Crutchfield J P 2015 *J. Stat. Phys.* **161** 404
- [44] Thompson J, Garner A J P, Vedral V and Gu M 2017 *npj Quantum Inf.* **3** 6