

Quantum dimension reduction of hidden Markov models

Rishi Sundar^{1,2} and Thomas J. Elliott^{1,3,2}

¹*Department of Physics & Astronomy, [University of Manchester](#), Manchester M13 9PL, United Kingdom*

²*Centre for Quantum Science and Engineering, [University of Manchester](#), Manchester M13 9PL, United Kingdom*

³*Department of Mathematics, [University of Manchester](#), Manchester M13 9PL, United Kingdom*



(Received 23 January 2026; revised 30 March 2026; accepted 12 May 2026; published 25 June 2026)

Hidden Markov models (HMMs) are ubiquitous in time-series modeling, with applications ranging from chemical reaction modeling to speech recognition. These HMMs are often large, with high-dimensional memories. A recently proposed application of quantum technologies is to execute quantum analogs of HMMs. Such quantum HMMs (QHMMs) are strictly more expressive than their classical counterparts, enabling the construction of more parsimonious models of stochastic processes. However, state-of-the-art techniques for QHMM compression, based on tensor networks, are only applicable for a restricted subset of HMMs, where the transitions are deterministic. In this work we introduce a pipeline by which *any* finite, ergodic HMM can be compressed in this manner, providing a route for effective quantum dimension reduction of general HMMs. We demonstrate the method on both a simple toy model, and on a speech-derived HMM trained from data, obtaining favorable memory-accuracy tradeoffs in the examples studied, relative to a simple classical state-merging baseline.

DOI: [10.1103/plt9-cm1x](#)

I. INTRODUCTION

Hidden Markov models (HMMs) provide a compact description of discrete-time stochastic processes in terms of a finite internal memory that evolves stochastically while emitting observable symbols [1,2]. They are used widely across scientific modeling and data analysis, including speech recognition and other sequential data [3–13]. As learned HMMs grow in size and connectivity, a central challenge is to reduce their effective memory cost while retaining the observable statistics that make them useful.

Classical reduction techniques, such as state merging, aim to simplify a model while preserving its predictions [14–18]. In practice, however, these methods can involve large intermediate representations or trade state reduction against loss of longer-range predictive structure when applied to realistic learned models with many states and rich transition patterns. This raises a basic question: Given a complex HMM that fits data well, how much memory is *really* required to reproduce its observable statistics, and can that memory cost be reduced in a controlled way?

Quantum models of stochastic processes provide another route to memory reduction. By encoding predictive information into nonorthogonal quantum memory states, quantum simulators can reproduce the statistics of a classical process while reducing the required memory [19–29]. The q -sample construction makes this connection concrete by

mapping the stationary statistics of a process to a translationally invariant quantum state on an infinite chain, with entanglement across a cut quantifying the memory resources needed for sampling [30–32]. For processes admitting a finite-dimensional q -sample infinite matrix product state (iMPS), standard uniform-MPS algorithms allow controlled bond-dimension truncation, trading a small loss in fidelity rate for a large reduction in bond dimension [32,33], corresponding to quantum models of stochastic processes with significantly reduced memory dimension compared to classical HMMs.

Born-machine approaches provide a useful neighboring point of comparison since both frameworks use tensor-network probability representations based on squared amplitudes. Our focus here is different: We start from a given finite-state HMM and study its compression into a lower-dimensional quantum sequential generator, rather than direct generative learning from data. This places this work more naturally in the QHMM and q -sample setting of recurrent process models, while keeping a clear conceptual link to Born-machine methods. Related work on QHMMs has also considered the direct learning of sequential quantum models from data [34], which is again distinct from the present post-training compression setting.

A key limitation is that existing iMPS compression results apply most cleanly when the underlying process admits a *deterministic* (or *unifilar*) presentation, where the next internal state is fixed by the current state and emitted symbol. In this case the q sample admits a *normal* iMPS representation with a primitive transfer operator and a well-conditioned canonical form [31,32,35]. HMMs learned directly from data, however, are typically nondeterministic, and naive tensor-network constructions can yield non-normal iMPS (or mixed-state tensor-network descriptions) for which stable infinite-system

truncation is not available [32]. This creates a practical gap between realistic learned HMMs and the class of truncated quantum models accessible to current iMPS-based compression pipelines.

Here we bridge that gap. We introduce a dilation that augments the output alphabet of any finite-state, stationary, ergodic HMM so that the resulting model is deterministic while preserving all finite-length word statistics on the original alphabet, without increasing the model's memory dimension. The q sample of the dilated process is guaranteed to be representable as a normal iMPS, enabling stable variational truncation to a target bond dimension $\tilde{d}$. From the compressed tensors we can reconstruct an effective compressed quantum model of the original process.

We demonstrate the method on (i) a tunable nondeterministic source family of toy models and (ii) a speech-derived HMM trained from data. In both cases we obtain a controlled tradeoff between bond dimension and distortion. The compressed iMPS provides a direct specification of a quantum sampler with memory dimension $\tilde{d}$. Our aim is thus to provide a post-training compression pipeline for general finite ergodic HMMs, rather than a direct data-driven training method for compressed quantum generators.

II. BACKGROUND

A. Classical models of stochastic processes

We consider a discrete-time stochastic process $\mathcal{P}$ that generates a bi-infinite sequence of random variables $\{X_t\}_{t \in \mathbb{Z}}$, each taking values in a finite alphabet $\mathcal{X}$ [36]. The process is fully specified by the joint distribution $\Pr(\dots, X_{-1}, X_0, X_1, \dots)$ over all times. In practice, modeling and simulating such a process requires a finite description of the statistical dependencies between past and future [37].

Hidden Markov models (HMMs) provide one such description [1,38]. An HMM is specified by a tuple $(\mathcal{S}, \mathcal{X}, \{T^x\})$, where $\mathcal{S}$ is a finite set of hidden states, $\mathcal{X}$ is the set of observable symbols, and the transition tensor T^x has elements

$$T_{s's}^x = \Pr(S_{t+1} = s', X_t = x \mid S_t = s) \quad (1)$$

for all $s, s' \in \mathcal{S}$ and $x \in \mathcal{X}$. For each fixed state s the matrices satisfy $\sum_{s',x} T_{s's}^x = 1$, so that they define a proper conditional distribution over next states and outputs.

A process has a *deterministic* HMM representation if the next internal state is uniquely determined by the current state and the emitted symbol [2,39]. In terms of the transition structure this means that for every state s and symbol x there is at most one successor state s' with nonzero transition probability:

$$|\{s' : T_{s's}^x > 0\}| \leq 1 \quad \forall s \in \mathcal{S}, x \in \mathcal{X}. \quad (2)$$

Equivalently, given the initial state and the observed symbol sequence, the present internal state is uniquely determined. Deterministic models play a central role in the branch of complexity science known as computational mechanics: the ε machine corresponds to the minimal deterministic HMM that exactly reproduces a process, and its memory cost is the Shannon entropy of the stationary distribution over causal states, the statistical complexity C_μ [2,39], often interpreted as a measure of structure [37,40,41]. Nondeterministic HMMs,

in contrast, can represent the same process with fewer states, but at the cost of ambiguous internal state trajectories conditioned on observations [42,43].

Throughout this work we restrict attention to finite-state, finite-alphabet, stationary, and ergodic processes, so that a unique stationary distribution over internal states exists and time-translation-invariant descriptions are well defined.

B. Quantum models of stochastic processes

Quantum models extend this framework by encoding predictive information into quantum memory states $\{|\sigma_j\rangle\}$ associated with internal configurations of the model [19,23,24]. The simulation proceeds via a joint evolution of the memory system and an output register, implemented by an isometry (equivalently, a unitary on a larger space). Starting from an initial memory state, one applies this evolution and measures the output register to obtain the next symbol x , leaving the memory updated to the postmeasurement state conditioned on that outcome. Repeated application of this procedure generates sequences with the same statistics as the target process. Such a procedure has been experimentally realized in multiple platforms [44–47].

More generally, a quantum hidden Markov model (QHMM) [25,26,28,29,48–50] on an alphabet $\mathcal{X}$ can be specified by the tuple $(\mathcal{X}, \mathcal{H}, \rho_0, \{\mathcal{E}_x\}_{x \in \mathcal{X}})$, where ρ_0 is a density operator on the memory Hilbert space $\mathcal{H}$ and $\{\mathcal{E}_x\}$ is a quantum instrument: each $\mathcal{E}_x$ is completely positive and trace nonincreasing, and $\sum_x \mathcal{E}_x$ is trace preserving. The probability of a word $w = x_1 \dots x_L$ is

$$\Pr(w) = \text{Tr}[\mathcal{E}_{x_L} \circ \dots \circ \mathcal{E}_{x_1}(\rho_0)]. \quad (3)$$

The quantum statistical memory C_q is defined as the von Neumann entropy of the stationary memory state,

$$C_q = S(\rho) = -\text{Tr}(\rho \log_2 \rho), \quad \rho = \sum_j \pi_j |\sigma_j\rangle \langle \sigma_j|, \quad (4)$$

where π_j is the stationary probability of internal configuration j . Many processes admit models for which $C_q \leq C_\mu$, demonstrating a quantum memory advantage, and for some families the separation between C_q and C_μ can be unbounded [51–57]. Finite-dimensional quantum models can be constructed directly from the classical HMM [25,29].

C. Matrix product states for deterministic models

A stochastic process $\mathcal{P}$ over a finite alphabet $\mathcal{X}$ can be encoded into a translationally invariant quantum state via its q sample [31]. Informally, the q sample is the infinite-chain limit of the pure state whose computational-basis amplitudes are square roots of classical word probabilities. For a length- L block $x_{t:t+L} := x_t x_{t+1} \dots x_{t+L-1}$, one considers

$$|P(x_{t:t+L})\rangle := \sum_{x_{t:t+L} \in \mathcal{X}^L} \sqrt{P(x_{t:t+L})} |x_{t:t+L}\rangle, \quad (5)$$

and takes $L \rightarrow \infty$ in a manner compatible with stationarity.

When the q sample admits an infinite matrix product state (iMPS) representation with site tensors $\{A^x\}_{x \in \mathcal{X}}$ and transfer

matrix

$$E = \sum_{x \in \mathcal{X}} A^x \otimes (A^x)^*, \quad (6)$$

all finite-block measurement statistics are obtained from a *double-layer* contraction. Concretely, if the leading eigenvalue of E is nondegenerate, then boundaries are irrelevant in the thermodynamic limit and one may express the block distribution in terms of the leading left and right eigenmatrices V_l and V_r of E as [31]

$$P_{\text{MPS}}(x_{t:t+L}) = \text{Tr}(A^{x_{t+L-1}} \dots A^{x_t} V_l A^{x_t} \dots A^{x_{t+L-1}} V_r). \quad (7)$$

Equation (7) is the appropriate probability-extraction formula for q -sample iMPS, and it replaces any single-layer expression of the form $\langle l | A^{x_t} \dots A^{x_{t+L-1}} | r \rangle$.

For unifilar HMMs (in particular, ε machines), Yang *et al.* show that the choice

$$A_{s's}^x = \sqrt{T_{s's}^x} \quad (8)$$

fully represents the process, in the sense that $P_{\text{MPS}}(x_{t:t+L}) = P(x_{t:t+L})$ for all L , and moreover that the process is ergodic if and only if the transfer matrix E has a nondegenerate leading eigenvalue [31]. In addition, the canonical form of an iMPS can be constructed systematically from V_l and V_r via factorizations $V_r = W_r W_r^\dagger$ and $V_l = W_l^\dagger W_l$ followed by an singular value decomposition (SVD) of $W_l W_r$ to obtain Schmidt coefficients and canonical tensors [31].

In this work, we will use Eq. (7) as the operative link between iMPS tensors and classical block statistics, and we explicitly gauge fix our tensors into canonical form before interpreting them as Kraus operators. A compatible canonical-form framework and completeness relations are reviewed in Ref. [32] (Appendix A), where left- and right-canonical tensors satisfy $\sum_x (A_l^x)^\dagger A_l^x = \mathbb{I}$ and $\sum_x A_r^x (A_r^x)^\dagger = \mathbb{I}$, enabling a channel (instrument) interpretation of the site tensors [32].

D. Difficulties with nondeterministic HMMs

Direct application of MPS methods to nondeterministic HMMs encounters significant challenges. A naive construction might attempt to use the same elementwise square-root mapping $T^x \mapsto A^x$ as above. For nondeterministic models this typically yields a transfer operator E that is not primitive: it may have multiple dominant eigenvalues or leading eigenvectors that are not strictly positive. The resulting iMPS is non-normal, so its canonical form is either not unique or poorly conditioned. Without normality, standard truncation methods and tangent-space-based infinite MPS algorithms become inapplicable or numerically unstable, and straightforward singular-value truncations along the virtual bonds often produce suboptimal approximations [32].

From the perspective of quantum implementations, one can represent general nondeterministic processes using matrix product density operators (MPDOs) or mixed-state tensor networks [25,34,58,59]. These constructions are more flexible but less suited to controlled dimension reduction, and they do not directly provide the normal iMPS structure required by existing compression schemes. In the remainder of this work we address these difficulties by constructing, for any finite nondeterministic HMM, a dilated deterministic process whose

q -sample iMPS is normal by construction. This provides the starting point for stable variational compression and for the compressed representations that we analyze in the following sections.

III. THEORETICAL RESULTS

Our aim is to start from a general finite-state, discrete-alphabet, stationary HMM and obtain a normal infinite matrix product state (iMPS) that can be compressed by standard tensor-network methods. The overall pipeline has four steps: (i) dilate the original HMM to a deterministic process on an augmented output alphabet; (ii) construct its q -sample iMPS; (iii) variationally truncate this iMPS to a reduced bond dimension; and (iv) reconstruct an effective model for the original outputs and quantify the approximation quality using the coemission divergence rate (CDR).

Throughout we assume that the original HMM is finite state, finite alphabet, stationary, and ergodic.

A. Dilation to a deterministic process

Let $(\mathcal{S}, \mathcal{X}, \{T^x\})$ be an HMM for a stationary, ergodic process $\mathcal{P}$. In general this model need not be deterministic: for a given state s and symbol x , there may be several possible successor states s' with $T_{s's}^x > 0$.

We quantify this local branching by

$$k(s, x) = |\{s' \in \mathcal{S} : T_{s's}^x > 0\}| \quad (9)$$

and define

$$d_y = \max_{s \in \mathcal{S}, x \in \mathcal{X}} k(s, x), \quad (10)$$

so that $1 \leq d_y \leq |\mathcal{S}|$. We introduce an auxiliary output alphabet $\mathcal{Y} = \{y_1, \dots, y_{d_y}\}$ and choose a labeling function

$$f : \mathcal{S} \times \mathcal{S} \times \mathcal{X} \rightarrow \mathcal{Y}, \quad (11)$$

which assigns a label $y = f(s, s', x)$ to each allowed transition ($s \rightarrow s'$) emitting symbol x . The only constraint on f is that it be injective for each fixed (s, x) :

$$s' \neq s'' \implies f(s, s', x) \neq f(s, s'', x) \quad \forall s \in \mathcal{S}, x \in \mathcal{X}. \quad (12)$$

Such a labeling always exists because d_y is at least as large as the maximal branching $k(s, x)$.

The dilated model has the same hidden-state space $\mathcal{S}$ and a composite output alphabet $\mathcal{X} \times \mathcal{Y}$. Its transition tensor is

$$T_{s's}^{(x,y)} = \begin{cases} T_{s's}^x, & \text{if } T_{s's}^x > 0 \text{ and } y = f(s, s', x) \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

This dilation has three key structural properties:

(1) **Determinism:** for each state s and composite symbol (x, y) there is at most one successor state s' with $T_{s's}^{(x,y)} > 0$.

(2) **Preservation of observable statistics:** marginalizing over the auxiliary outputs y reproduces the same finite-length block probabilities on $\mathcal{X}$ as the original process $\mathcal{P}$.

(3) **Ergodicity:** if the original HMM is ergodic, then the dilated HMM is also ergodic since the underlying Markov chain on $\mathcal{S}$, obtained by summing over outputs, is unchanged.

Proofs are collected in Appendix B. The dilation therefore yields a finite-state, stationary, ergodic, and deterministic representation of $\mathcal{P}$ on the augmented alphabet $\mathcal{X} \times \mathcal{Y}$.

B. Constructing the q -sample iMPS

Given the deterministic dilated HMM $(\mathcal{S}, \mathcal{X} \times \mathcal{Y}, \{T^{(x,y)}\})$, we construct its q -sample iMPS [31] by taking the elementwise principal square root of the transition tensor,

$$A_{s's}^{(x,y)} = \sqrt{T_{s's}^{(x,y)}}, \quad (14)$$

for all $s, s' \in \mathcal{S}$ and $(x, y) \in \mathcal{X} \times \mathcal{Y}$. Treating the composite symbol (x, y) as a single physical index, these tensors define a translationally invariant iMPS on an infinite chain with physical dimension $d_{\text{phys}} = |\mathcal{X}| |\mathcal{Y}|$.

The probability of a finite block $(x_1, y_1), \dots, (x_L, y_L)$ is obtained by contracting the iMPS with stationary boundary data, for example, in the standard transfer-matrix form. By construction and by the preservation property above, the marginal distribution over $(x_1, \dots, x_L)$ coincides with that of the original process $\mathcal{P}$.

The corresponding transfer operator for the iMPS is

$$E = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} A^{(x,y)} \otimes (A^{(x,y)})^*, \quad (15)$$

where $*$ denotes complex conjugation. Since the dilated HMM is finite state, stationary, ergodic, and deterministic, the associated q -sample iMPS has a *primitive* transfer operator with a unique leading eigenvalue $\eta_0 = 1$ and strictly positive left and right eigenvectors [31]. The iMPS is therefore *normal*, and blocking a finite number of sites yields an injective MPS with a unique canonical form and minimal bond dimension [35]. In practical computations we bring $\{A^{(x,y)}\}$ into mixed canonical form using standard gauge-fixing procedures.

The overall construction, from the original transition tensor through dilation to the site tensors of the iMPS and the associated transfer operators, is illustrated schematically in Fig. 1.

C. Variational truncation at reduced bond dimension

The iMPS defined by $\{A^{(x,y)}\}$ is an exact representation of the dilated process, with bond dimension equal to the number of internal states $d_s = |\mathcal{S}|$ of the original HMM. To obtain a more compact model we seek an approximate iMPS with a reduced bond dimension $\tilde{d} < d_s$ that is as close as possible, in the many-body sense, to the original.

We adopt a uniform tangent-space variational approach for iMPS with fixed bond dimension [32,33]. Given the normal iMPS $\{A^{(x,y)}\}$ with bond dimension d_s , we project it onto the manifold of translationally invariant iMPS with bond dimension $\tilde{d}$ by solving the associated optimization problem for a new set of tensors $\{\tilde{A}^{(x,y)}\}$. In practice, this is implemented as an effective eigenvalue problem in the tangent space of the target manifold, solved iteratively using standard MPS routines.

Each iteration involves contractions of the transfer operator E with trial tensors of bond dimension $\tilde{d}$. We need never form E explicitly as a dense $d_s^2 \times d_s^2$ matrix; instead, we apply it as a linear map with cost that scales polynomially in d_s , $\tilde{d}$, and the physical dimension $d_{\text{phys}} = |\mathcal{X}| |\mathcal{Y}|$. The auxiliary




FIG. 1. Tensor-network representations of stochastic models and their compression. (a) A transition matrix T_{ij}^x of a HMM can be represented as a rank-3 tensor; analogously, (b) the transition matrix of the dilated process $T_{ij}^{(x,y)}$ can be represented by a rank-4 tensor, as can (c) its elementwise square root $A_{ij}^{(x,y)}$. An array of these form an injective iMPS when the two output legs are grouped, enabling the application of iMPS compression methods to deduce a dimension-reduced iMPS with site tensors $\tilde{A}_{ij}^{(x,y)}$. The (d) transfer matrix of the reduced iMPS defines a quantum channel $\tilde{E}$ when a Kronecker δ is applied to the two visible output legs, implementing a compressed model of the original process. This can then (e) be paired with the original process to efficiently calculate their CDR.

alphabet size satisfies $|\mathcal{Y}| = d_y \leq |\mathcal{S}|$, so in the worst case the local physical dimension of the dilated iMPS can grow substantially compared to the original alphabet size $|\mathcal{X}|$. This enlarged local dimension is the computational price paid for obtaining a deterministic presentation whose q sample admits a normal iMPS, and it can become the dominant bottleneck in the variational truncation step for highly branching models. A conservative dense upper bound for a single local transfer-operator application is therefore of order $O(|\mathcal{X}| |\mathcal{Y}| d_s^2 \tilde{d}^2)$, with the precise prefactor depending on the chosen gauge and eigensolver implementation. In practice, however, the effective cost is governed by the number of allowed transitions inherited from the original HMM, rather than by the full product $|\mathcal{X}| |\mathcal{Y}|$; exploiting this sparsity replaces the dense worst-case factor $|\mathcal{X}| |\mathcal{Y}|$ by a cost scaling linearly in the nonzero tensor contributions. For sparse learned HMMs this can be substantially smaller than the worst-case scaling. The outcome is a compressed normal iMPS, specified by tensors $\{\tilde{A}^{(x,y)}\}$, that approximates the original dilated process while retaining translational invariance. For algorithmic details of how the variational truncation approach for iMPS is performed, we refer the reader to Ref. [33], with its application

to quantum compression of deterministic HMMs discussed in Ref. [32]

D. Reconstructing an effective model

The truncated tensors $\{\tilde{A}^{(x,y)}\}$ define a translationally invariant iMPS approximation to the dilated (x, y) process. From this we can construct a compressed quantum model of the dilated process, from which we can in turn construct a dimension-reduced QHMM of the original process by coarse graining over the dilation.

A crucial technical point is that Kraus (instrument) normalization is a gauge property of the iMPS representation. Following standard iMPS canonicalization, we first gauge fix the truncated tensors into a canonical form obtained from the leading left and right eigenmatrices of the truncated transfer operator [31,32]. In left-canonical form the site tensors satisfy the completeness relation

$$\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} (\tilde{A}^{(x,y)})^\dagger \tilde{A}^{(x,y)} = \mathbb{I}, \quad (16)$$

and hence define a completely positive trace-preserving (CPTP) map on the bond space [32]. This is the sense in which the truncated iMPS specifies a physically valid sequential quantum generator.

We now group the Kraus operators by the observed symbol x by summing over the auxiliary label y . For each $x \in \mathcal{X}$ define the completely positive (CP) map

$$\tilde{\mathcal{E}}_x(\rho) := \sum_{y \in \mathcal{Y}} \tilde{A}^{(x,y)} \rho \tilde{A}^{(x,y)\dagger}. \quad (17)$$

Equation (16) implies that the unconditional channel $\tilde{\mathcal{E}} := \sum_x \tilde{\mathcal{E}}_x$ is trace preserving. Let $\tilde{\rho}_\star$ denote its stationary state, i.e., a fixed point satisfying

$$\tilde{\rho}_\star = \sum_{x \in \mathcal{X}} \tilde{\mathcal{E}}_x(\tilde{\rho}_\star). \quad (18)$$

Then $\{\tilde{\mathcal{E}}_x\}$ forms a quantum instrument that generates a stationary process $\tilde{\mathcal{P}}$ on $\mathcal{X}$ via the usual update rule: conditioned on observing x_t , the bond state updates as

$$\tilde{\rho}_{t+1} = \frac{\tilde{\mathcal{E}}_{x_t}(\tilde{\rho}_t)}{\text{Tr } \tilde{\mathcal{E}}_{x_t}(\tilde{\rho}_t)}, \quad \Pr_{\tilde{\mathcal{P}}}(x_t) = \text{Tr } \tilde{\mathcal{E}}_{x_t}(\tilde{\rho}_t). \quad (19)$$

Consequently, the probability of a word $w = x_1 \dots x_L$ is

$$\Pr_{\tilde{\mathcal{P}}}(w) = \text{Tr}(\tilde{\mathcal{E}}_{x_L} \circ \dots \circ \tilde{\mathcal{E}}_{x_1}(\tilde{\rho}_\star)). \quad (20)$$

The dimension-reduced QHMM is then specified by the tuple $(\mathcal{X}, \tilde{\mathcal{H}}, \tilde{\rho}_0, \{\tilde{\mathcal{E}}_x\}_{x \in \mathcal{X}})$, where $\tilde{\mathcal{H}}$ is the Hilbert space of the memory of the compressed QHMM, and $\tilde{\rho}_0$ is the initial state (generally, taken to be the steady state $\tilde{\rho}_\star$).

E. Quantifying distance from exact process

For notational convenience we pass to the Liouville representation. Assign column-stacking vectorization vec defined on matrix units by

$$\text{vec}(|i\rangle\langle j|) := |i\rangle \otimes |j\rangle. \quad (21)$$

With this convention one has the standard identity

$$\text{vec}(A\rho B) = (A \otimes B^\top) \text{vec}(\rho), \quad (22)$$

where $^\top$ denotes transpose in the computational basis. Write $|\rho\rangle\rangle := \text{vec}(\rho) \in \mathbb{C}^{\tilde{d}^2}$. Then the CP map $\tilde{\mathcal{E}}_x$ in Eq. (17) is represented by a linear operator $G^{(x)}$ via

$$|\rho'\rangle\rangle = G^{(x)}|\rho\rangle\rangle, \quad G^{(x)} = \sum_{y \in \mathcal{Y}} \tilde{A}^{(x,y)} \otimes (\tilde{A}^{(x,y)})^*, \quad (23)$$

where * denotes complex conjugation. The trace functional is $\langle\langle \omega | := |\mathbb{I}\rangle\rangle^\dagger$, so that $\langle\langle \omega | |\rho\rangle\rangle = \text{Tr}(\rho)$. Hence, the triple $(\mathcal{X}, |\rho_\star\rangle\rangle, \{G^{(x)}\})$ defines a finite-state linear generator, a generalized HMM (GHMM) [29,38], where for any word $w = x_1 \dots x_L$,

$$\Pr_{\tilde{\mathcal{P}}}(w) = \langle\langle \omega | G^{(x_L)} \dots G^{(x_1)} |\rho_\star\rangle\rangle. \quad (24)$$

The operators $G^{(x)}$ need not be elementwise non-negative or stochastic; the requirement is simply that Eq. (24) yields valid word probabilities.

To quantify how well $\tilde{\mathcal{P}}$ approximates the original process $\mathcal{P}$, we use the coemission divergence rate (CDR) [60]. Let P and Q be stationary, ergodic processes described by finite-state linear generators (either HMMs or GHMMs) $\{L^x\}$ and $\{\hat{L}^x\}$. We define the associated transfer operators on the product state spaces by

$$\begin{aligned} E_P &:= \sum_{x \in \mathcal{X}} L^x \otimes L^x, \\ E_Q &:= \sum_{x \in \mathcal{X}} \hat{L}^x \otimes \hat{L}^x, \\ E_{PQ} &:= \sum_{x \in \mathcal{X}} L^x \otimes \hat{L}^x \end{aligned} \quad (25)$$

and let μ_P, μ_Q , and μ_{PQ} denote their leading eigenvalues. The CDR is then given by

$$R_C(P, Q) = -\frac{1}{2} \log_2 \frac{\mu_{PQ}}{\sqrt{\mu_P \mu_Q}}. \quad (26)$$

While calculated from GHMM representations, this expression depends only on the induced word distributions, valid for any stationary and ergodic processes [60].

In our setting we take P to be the original HMM process $\mathcal{P}$ on $\mathcal{X}$, with generators $L^x = T^x$, and Q to be the compressed process $\tilde{\mathcal{P}}$ generated by $\hat{L}^x = G^{(x)}$ as defined above. Their CDR, $R_C(\mathcal{P}, \tilde{\mathcal{P}})$, is the main observable figure of merit we report in our numerics.

F. Analytic bounds via fidelity divergence rates

The pipeline above produces, for each target bond dimension $\tilde{d}$, a compressed process $\tilde{\mathcal{P}}$ together with its CDR $R_C(\mathcal{P}, \tilde{\mathcal{P}})$, which we evaluate numerically from finite-state generators. Analytically, the dimension-reduction results of Ref. [32] control *fidelity-type* errors of the underlying q -sample iMPS. We therefore derive rigorous upper bounds on a $\mathcal{X}$ -level *classical fidelity divergence rate* (CFDR), obtained from the dilated q samples by data processing. This CFDR is a different, but closely analogous, distinguishability measure from the CDR R_C used in our plots.

Let $\{\lambda_k\}_{k=1}^{d_s}$ denote the *bond spectrum* of the dilated q -sample iMPS across a single cut, i.e., the eigenvalues of the stationary bond state ρ_* in canonical form (equivalently, squared Schmidt singular values), ordered so that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{d_s}$ and $\sum_k \lambda_k = 1$. For a truncated bond dimension $\tilde{d}$ define the discarded tail weight

$$\varepsilon_{\tilde{d}} := \sum_{k > \tilde{d}} \lambda_k. \quad (27)$$

We write $R_F(|P_{xy}\rangle, |\tilde{P}_{xy}\rangle)$ for the quantum fidelity divergence rate (QFDR) between the exact and truncated (x, y) -dilated q -sample states, and $R_F(\mathcal{P}, \tilde{\mathcal{P}})$ for the CFDR on $\mathcal{X}$, defined by

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) := - \lim_{L \rightarrow \infty} \frac{1}{2L} \log_2 \sum_{\vec{x} \in \mathcal{X}^L} \sqrt{P^{(L)}(\vec{x}) \tilde{P}^{(L)}(\vec{x})}. \quad (28)$$

By applying dephasing in the (x, y) basis followed by tracing out the auxiliary register Y , fidelity monotonicity implies the data-processing bound

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq R_F(|P_{xy}\rangle, |\tilde{P}_{xy}\rangle), \quad (29)$$

proved in Appendix B. Reference [32] then implies that, for sufficiently small $\varepsilon_{\tilde{d}}$, there exists a constant $c > 0$ (depending on the local physical dimension and the spectral gap of the transfer operator) such that

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq c \varepsilon_{\tilde{d}}. \quad (30)$$

The tail weight can in turn be bounded in terms of the bond entropy

$$H(\lambda) := - \sum_k \lambda_k \log_2 \lambda_k, \quad (31)$$

giving (Appendix B)

$$\varepsilon_{\tilde{d}} \leq \frac{H(\lambda)}{\log_2 \tilde{d}} \quad (\tilde{d} \geq 2), \quad (32)$$

and hence an entropy-based guarantee

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq c \frac{H(\lambda)}{\log_2 \tilde{d}}. \quad (33)$$

Finally, we connect $H(\lambda)$ to a simple algebraic quantity depending on the dilation. Let K be the “slice” matrix formed by horizontally concatenating all nonzero site tensors $A^{(x,y)}$, viewed as linear maps on the bond space,

$$K = [A^{(x_1, y_1)} \ A^{(x_2, y_2)} \ \dots]. \quad (34)$$

The column space of K contains the support of the stationary bond state, so $\text{rank}(K)$ upper bounds the number of nonzero eigenvalues of ρ_* , and therefore

$$H(\lambda) \leq \log_2 \text{rank } K. \quad (35)$$

Substituting into Eq. (33) yields the structural, label-aware bound

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq c \frac{\log_2 \text{rank } K}{\log_2 \tilde{d}}. \quad (36)$$

In the numerical results below we report the CDR R_C because it is efficiently computable from finite-state generators: evaluating Eq. (26) requires only the leading eigenvalues




FIG. 2. HMM representation of the generalized N -state tunable nondeterministic source (TNS). Transition labels show “probability | emitted symbol,” with the auxiliary label symbol in red.

of the product-space transfer operators, which is a standard sparse-eigenvalue problem. The bounds above instead provide rigorous certificates for the classical fidelity divergence rate R_F , a Bhattacharyya-type overlap measure, that is directly controlled by iMPS truncation theorems. Both R_C and R_F vanish if and only if the two processes are identical, but they are distinct measures of process discrepancy and are not in general proportional. Without additional assumptions one should not interpret these as direct upper bounds on R_C . Thus, the analytic results provide rigorous worst-case control for the CFDR, while the CDR serves as the main empirical diagnostic in our numerics.

IV. EXAMPLES AND ANALYSIS

We now demonstrate and analyze the dilation-compression methodology on two examples. First, we consider a *tunable nondeterministic source* (TNS), a generalization of a nondeterministic generator [61] that has previously been used to exhibit quantum reductions in memory [25]. Here we use it as a controlled test bed for how compression performance depends on internal parameters and on the entanglement structure induced by the dilation. Second, we apply the method to a hidden Markov model trained on a real speech dataset and compare the compressed representations to a standard classical reduction baseline.

A. Illustrative model: N -state simple nondeterministic source

As a controlled test bed we use a generalized N -state TNS with internal states $\mathcal{S} = \{0, 1, \dots, N-1\}$ and a single parameter $p \in (0, 1)$ tuning the transition structure. For generic p and $N \geq 2$ the model is nondeterministic: conditioning on the emitted symbol does not uniquely fix the successor state. The transition structure is shown schematically in Fig. 2.

Starting from this TNS HMM we construct the dilated deterministic model according to Sec. III A, build its q -sample iMPS as in Sec. III B, and perform variational truncation at a sequence of reduced bond dimensions $\tilde{d}$ following Sec. III C. For each choice of (N, p) and target $\tilde{d}$ we reconstruct the effective model for the original alphabet and compute the coemission divergence rate $R_C(\mathcal{P}, \tilde{\mathcal{P}})$ between the original and compressed processes, as described in Sec. III E.

Figure 3 shows the resulting CDR as a function of the truncated bond dimension $\tilde{d}$ for several representative parameter choices. Each curve corresponds to a fixed pair (N, p) . In all cases the CDR decreases systematically as $\tilde{d}$ increases, indicating that the variational procedure produces families of compressed models that converge toward the original process as more virtual resources are allowed. The rate of




FIG. 3. TNS model compression results: coemission divergence rate R_C between the original process and the reconstructed effective model as a function of truncated bond dimension $\tilde{d}$. Each curve corresponds to a different choice of the number of states N and internal parameter p .

decrease depends strongly on both N and p : some parameter regimes reach a given CDR at much smaller $\tilde{d}$ than others, reflecting differences in how strongly the process resists low-dimensional compression. This behavior is also in line with the analytic picture of Sec. III F, where the loss induced by truncation is controlled by the tail of the bond spectrum of the (dilated) iMPS.

To understand this dependence more structurally, we examine the Schmidt spectra of the dilated iMPS before truncation. For each (N, p) we bring the iMPS into canonical form and compute the singular values across a single bond, which quantify the bipartite entanglement across that cut. Figure 4 shows the Schmidt coefficients as a function of their index for a fixed N and varied p . The decay patterns vary markedly: some spectra are rapidly decaying (concentrating most weight into a few leading modes), while others are much flatter. This also motivates the use of the variational method, as standard SVD truncation methods which work by discarding singular values work poorly on flat distributions.




FIG. 4. Schmidt spectra of the dilated iMPS of the TNS for fixed $N = 15$ and varied p , ordered by decreasing magnitude. The shape of the spectra can vary, but parameter regimes with faster-decaying spectra are expected to be more amenable to truncation to small bond dimension.

Comparing Figs. 3 and 4, one finds that parameter choices with more rapidly decaying Schmidt spectra tend to achieve low CDR at smaller $\tilde{d}$, while flatter spectra require larger $\tilde{d}$ to reach comparable accuracy. This is consistent with the intuition that the entanglement encoded in the dilated iMPS is a practical proxy for the structural complexity that survives the dilation and resists compression.

B. Application to a speech-derived HMM

To assess the practical utility of the method we apply it to a hidden Markov model trained on a subset of the Google speech commands dataset [13]. Each audio clip is mapped to a sequence of acoustic feature vectors using Mel-frequency cepstral coefficients (MFCCs), including time-derivative features [62]. We then discretize these features by vector quantization using MiniBatch k means, yielding a finite-output alphabet [63]. Finally, we fit a categorical (edge-emitting) HMM by maximum likelihood via the expectation-maximization (or *Baum-Welch*) procedure [3,64]. The resulting learned model is strongly nondeterministic, making it a realistic test case for the dilation step.

We take this trained HMM as the target process $\mathcal{P}$ and apply our dilation-compression pipeline to obtain a compressed quantum model of the process. To benchmark against a classical reduction pipeline, we also perform greedy state merging directly on the original learned HMM. In each merge step we choose the pair of states that minimizes a stationary-weighted KL objective defined on the one-step predictive distributions $p_s(x, s') = T_{s's}^x$, and we update the merged state by a simple stationary-mixture lumping rule. We then evaluate the CDR between the original HMM and each merged HMM at the corresponding state counts.

We emphasize that this merging routine is intended as a transparent and computationally inexpensive *baseline* rather than a best-in-class classical reduction method. It is greedy, uses a local one-step objective, and does not optimize sequence-level divergences such as the CDR directly; more sophisticated classical reductions (e.g., global objectives, multistep lookahead, or spectral methods) may achieve stronger performance in some regimes. The purpose of the baseline is to provide a concrete reference curve against which the effect of the dilation-compression pipeline can be assessed on the same learned model.

Figure 5 compares the quantum reduction curve (memory dimension $\tilde{d}$) with the classical merging curve (retained HMM state count). In the strongly compressed regime (memory dimension $\lesssim 20$), the quantum models achieve substantially smaller CDR, often by orders of magnitude, than this greedy classical baseline at the same nominal memory dimension. Notably, the quantum curve exhibits a sharp improvement around $\tilde{d} \approx 12$, consistent with the variational truncation needing only few states to capture the dominant modes of the dilated iMPS. By contrast, the classical baseline only approaches the numerical floor once the retained state count is sufficiently large (here around $\gtrsim 22$ states), suggesting that a moderate number of learned states can be merged with negligible sequence-level distortion, whereas more aggressive classical reductions rapidly degrade performance. We note that this comparison is against one specific classical baseline




FIG. 5. Speech-derived HMM compression results: comparison of quantum dimension reduction and a classical greedy state-merging baseline. Horizontal axis reports memory dimension as classical state count or quantum bond dimension. Vertical axis reports CDR R_C (bits per symbol, logarithm scale). Values below 10^{-9} are displayed at 10^{-9} for visibility.

on the same learned model; more sophisticated classical reduction methods may narrow the gap in some regimes.

C. Impact of the labeling function

The dilation procedure of Sec. III A contains a genuine degree of freedom: the choice of the labeling function $f : \mathcal{S} \times \mathcal{S} \times \mathcal{X} \rightarrow \mathcal{Y}$ used to assign auxiliary outputs to transitions. Different admissible labelings (all satisfying the injectivity constraint) can induce different entanglement structures in the dilated iMPS. Since truncation performance is sensitive to the bond spectrum, it is natural to ask how strongly the choice of f affects compressibility in practice.

To probe this, we fix a TNS instance and compare several simple labeling strategies that differ only in how they assign the auxiliary symbols y to the allowed transitions, while keeping the auxiliary alphabet size $|\mathcal{Y}|$ fixed. For each choice of f we construct the corresponding dilated iMPS, perform variational truncation across a range of bond dimensions $\tilde{d}$, reconstruct the effective models, and compute the CDR.

Figure 6(a) shows $R_C(\mathcal{P}, \tilde{\mathcal{P}})$ as a function of $\tilde{d}$ for the different labeling strategies. The curves differ markedly: some labelings yield low CDR already at small $\tilde{d}$, while others require substantially larger bond dimensions to reach comparable accuracy. In this sense the labeling is not merely cosmetic; it acts as an optimization parameter that can materially change the achievable CDR- $\tilde{d}$ tradeoff. This observation also matches the structural picture in Sec. III F, where label-dependent quantities (such as the slice structure) can upper bound entropic tails and hence influence error guarantees.

The corresponding Schmidt spectra of the dilated iMPS for the same labelings are shown in Fig. 6(b). Different assignments of y produce distinct decay patterns. Labelings with more rapidly decaying spectra tend to yield the lowest CDR at small $\tilde{d}$, while flatter spectra yield poorer compression performance, reinforcing the role of the induced entanglement structure as a practical predictor of compressibility. Structurally, different admissible labelings all preserve the original




FIG. 6. (a) Impact of labeling function on CDR for TNS model with fixed parameters (N, p) . Each curve corresponds to a different strategy for assigning auxiliary labels y to transitions. (b) Schmidt spectra of the dilated iMPS for each of the above labeling strategies; different assignments of auxiliary labels y produce distinct entanglement spectra that correlate with compression performance.

HMM exactly before truncation, but they redistribute the nonzero transition probabilities across the (x, y) slices of the dilated tensor. This reorganization changes the column structure of the slice matrix K and hence the Schmidt spectrum of the dilated iMPS, which in turn controls the bond-dimension-versus-error tradeoff after truncation. The specific labeling heuristics used in our numerical study are documented in Appendix A.

A full optimization of f over all admissible assignments is a nontrivial combinatorial problem and is beyond the scope of this work. The examples here show, however, that even simple heuristic labelings can significantly change the practical tradeoff between bond dimension and CDR, suggesting that learning or optimizing good labeling strategies is an important direction for future research.

V. DISCUSSION AND CONCLUSION

We introduced a general framework for mapping non-deterministic HMMs to matrix product state representations, enabling techniques for their implementation with compressed quantum models to be applied. We achieve with a dilation procedure that takes any finite, stationary, ergodic

HMM to a deterministic model on an enlarged output alphabet that preserves all observable statistics on the original symbols. The q sample of this dilated process is representable by a normal iMPS with bond dimension $d_s = |\mathcal{S}|$, which can be compressed using standard uniform-MPS tangent-space methods. From the compressed tensors we reconstruct quantum model of the original process, with reduced memory dimension.

We applied this pipeline to a TNS toy model and to a speech-derived HMM learned from data. In both cases we obtain families of compressed models that interpolate smoothly between the original process and low-bond approximations. Empirically the CDR decreases monotonically with $\tilde{d}$, and there is a broad regime in which $\tilde{d} \ll d_s$ while $R_C(\mathcal{P}, \tilde{\mathcal{P}})$ remains small. This supports the practical value of the dilation-compression pipeline as a dimension-reduction tool for practical quantum compression of HMMs.

The examples also clarify how compressibility depends on the entanglement structure of the dilated q sample. For the TNS family, different choices of the internal parameter p produce markedly different Schmidt spectra: rapidly decaying spectra admit accurate truncations at small $\tilde{d}$, while flatter spectra require larger bond dimensions to achieve comparable CDR. This supports the view that entanglement in the dilated iMPS is a useful proxy for structural complexity that resists compression.

We emphasize that this work establishes a general compression route for arbitrary finite ergodic HMMs and demonstrates it on representative examples, rather than providing a comprehensive large-scale benchmarking study. The computational cost of the pipeline is dominated by the variational truncation step, whose per-iteration cost scales with the bond dimensions d_s and $\tilde{d}$ and the effective physical dimension of the dilated iMPS. A conservative dense estimate for the local contractions entering one transfer-operator application is $O(|\mathcal{X}| |\mathcal{Y}| d_s^2 \tilde{d}^2)$, so because the auxiliary dilation can enlarge the physical dimension up to $|\mathcal{X}| |\mathcal{Y}|$ with $|\mathcal{Y}| \leq |\mathcal{S}|$, this cost can grow substantially for highly branching models with many states. In practice, the sparsity of the original transition tensor provides a significant mitigation since only nonzero tensor contributions enter the contraction cost, effectively replacing the worst-case factor $|\mathcal{X}| |\mathcal{Y}|$ by the number of allowed transitions. Nevertheless, scaling the pipeline to HMMs with hundreds or thousands of states will require careful exploitation of sparsity and potentially approximate contraction schemes, and we regard this as an important direction for future engineering of the method.

The most immediate extensions of this work are practical. First, the labeling function f is a genuine design degree of freedom: our results show that it can reshape the entanglement structure of the dilated q sample and substantially change the achievable CDR at fixed $\tilde{d}$. This suggests an algorithmic direction in which the dilation is not treated as a fixed pre-processing step, but as something to be chosen (or learned) to produce dilations that are intrinsically easier to compress.

Second, the dilation-compression pipeline provides a systematic route from large learned HMMs to small quantum samplers with an explicit and interpretable resource parameter $\tilde{d}$. This opens the door to using tensor-network diagnostics (bond spectra, entropies, canonical form data) as practical

predictors of which learned models are likely to admit aggressive compression without significantly distorting sequence statistics.

It is also useful to situate this work relative to tensor-network Born-machine approaches [65–67], which also use squared amplitudes of matrix product states to define probability distributions. Born-machine methods are typically framed as direct generative learning from data, without requiring an intermediate classical model. By contrast, our pipeline takes a given HMM as input and compresses it into a lower-dimensional quantum sequential generator with explicit structural control over the approximation error. These are parallel approaches with different modeling emphases. Training compressed quantum sequential models directly from raw data, bypassing the intermediate HMM, is an interesting open direction that is more naturally aligned with Born-machine methodology than with the present post-training compression setting. It would be interesting in future work to explore whether the compressed generators produced by our pipeline could also serve as informed initializations for subsequent data-driven optimization.

Finally, while we reported CDR because it is operationally meaningful and efficiently computable from finite-state generators, the analytic truncation guarantees established here apply to a fidelity-type divergence rate for the dilated q sample, which induces a corresponding classical fidelity-type quantity after dephasing and marginalization. There is, to our knowledge, no known general theorem directly linking these two quantities: they are based on different notions of process discrepancy, and the analytic results should therefore not be interpreted as direct bounds on CDR. At the same time, in the examples studied here they exhibit similar qualitative behavior, which is why CDR remains a useful empirical diagnostic while preserving a clear distinction from the analytically controlled quantity.

ACKNOWLEDGMENT

This work was funded by the University of Manchester Dame Kathleen Ollerenshaw Fellowship.

DATA AVAILABILITY

Data is available upon request to the corresponding author.

APPENDIX A: PROPERTIES OF THE DILATION

In this Appendix we collect the technical arguments underlying the structural properties of the dilation. We work throughout with the dilated q -sample iMPS for the process over (x, y) and its compressed version at bond dimension $\tilde{d}$, as constructed in Secs. III B and III C.

We start by recording the basic structural properties of the dilation defined in Sec. III A. Recall that the original HMM $(\mathcal{S}, \mathcal{X}, \{T^x\})$ is mapped to a dilated HMM $(\mathcal{S}, \mathcal{X} \times \mathcal{Y}, \{T^{(x,y)}\})$ with transition tensor

$$T_{s's'}^{(x,y)} = \begin{cases} T_{s's'}^x & \text{if } T_{s's'}^x > 0 \quad \text{and } y = f(s, s', x), \\ 0, & \text{otherwise} \end{cases} \quad (\text{A1})$$

where the labeling function $f : \mathcal{S} \times \mathcal{S} \times \mathcal{X} \rightarrow \mathcal{Y}$ is assumed to be injective in its second argument for each fixed (s, x) .

Lemma A1 (Deterministic of the dilation). For each state $s \in \mathcal{S}$ and composite symbol $(x, y) \in \mathcal{X} \times \mathcal{Y}$ there is at most one successor $s' \in \mathcal{S}$ with $T_{s's}^{(x,y)} > 0$.

Proof. Fix s and x . For each s' with $T_{s's}^x > 0$ the construction assigns a single label $y = f(s, s', x)$. The injectivity of f for fixed (s, x) implies that $s' \neq s''$ gives $f(s, s', x) \neq f(s, s'', x)$. Hence, for any pair (x, y) there is at most one s' such that $T_{s's}^{(x,y)} > 0$. ■

Lemma A2 (Preservation of observable statistics). Let $\mathcal{P}$ be the process on $\mathcal{X}$ generated by the original HMM and let $\mathcal{P}_{\text{dil}}$ be the process on $\mathcal{X}$ obtained from the dilated HMM by marginalizing over $\mathcal{Y}$. Then for all block lengths L and strings $(x_1, \dots, x_L)$,

$$\Pr_{\mathcal{P}}(x_1, \dots, x_L) = \Pr_{\mathcal{P}_{\text{dil}}}(x_1, \dots, x_L). \quad (\text{A2})$$

Proof. For any s, s' and x we have

$$\sum_{y \in \mathcal{Y}} T_{s's}^{(x,y)} = \sum_{y: y=f(s,s',x)} T_{s's}^x = T_{s's}^x, \quad (\text{A3})$$

since at most one y contributes. Thus, the conditional distribution of (S_{t+1}, X_t) given S_t is unchanged when we marginalize over Y_t . Iterating this equality along the chain yields equality of all finite-length block probabilities on $\mathcal{X}$. ■

Lemma A3 (Ergodicity of the dilation). If the original HMM is stationary and ergodic, then the dilated HMM is also stationary and ergodic.

Proof. The underlying Markov chain on $\mathcal{S}$, obtained by summing over outputs, is the same in both models:

$$\sum_{x,y} T_{s's}^{(x,y)} = \sum_x T_{s's}^x. \quad (\text{A4})$$

Irreducibility and aperiodicity of this chain are therefore inherited from the original HMM, and the stationary distribution over $\mathcal{S}$ carries over unchanged. Hence, the dilated HMM is stationary and ergodic. ■

Lemma A4 (Square-root tensors reproduce word statistics for unifilar generators). Let $(\mathcal{S}, \mathcal{A}, \{T^a\}_{a \in \mathcal{A}})$ be a finite-state, edge-emitting unifilar HMM for a stationary process on alphabet $\mathcal{A}$ (in our application $\mathcal{A} = \mathcal{X} \times \mathcal{Y}$). Write the transition probabilities as $T_{s's}^a = \Pr(S_{t+1} = s', A_t = a \mid S_t = s)$. Assume the underlying hidden-state Markov chain $P := \sum_{a \in \mathcal{A}} T^a$ is ergodic with stationary distribution π .

Define site matrices $A^a \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ by

$$A_{s's}^a := \sqrt{T_{s's}^a}. \quad (\text{A5})$$

Let $\rho_\star := \text{diag}(\pi)$ be the diagonal matrix with entries π_s . Then for every word $w = a_1 a_2 \dots a_L \in \mathcal{A}^L$,

$$\Pr_{\text{HMM}}(w) = \text{Tr}(A^{a_L} \dots A^{a_1} \rho_\star A^{a_1^\dagger} \dots A^{a_L^\dagger}). \quad (\text{A6})$$

Equivalently, the completely positive maps $\mathcal{E}_a(\rho) := A^a \rho A^{a^\dagger}$ define a quantum instrument whose word probabilities from the stationary bond state $\rho_\star$ agree exactly with the original unifilar HMM.

Proof. Fix a word $w = a_1 \dots a_L$ and write $M_w := A^{a_L} \dots A^{a_1}$. Since $\rho_\star$ is diagonal,

$$\begin{aligned} \text{Tr}(M_w \rho_\star M_w^\dagger) &= \sum_{s_0 \in \mathcal{S}} (\rho_\star)_{s_0 s_0} \sum_{s_L \in \mathcal{S}} |(M_w)_{s_L s_0}|^2 \\ &= \sum_{s_0 \in \mathcal{S}} \pi_{s_0} \sum_{s_L \in \mathcal{S}} |(M_w)_{s_L s_0}|^2. \end{aligned} \quad (\text{A7})$$

Now use unifilarity: for each current state s and symbol a there is at most one successor state $\delta(s, a)$ with $T_{\delta(s,a),s}^a > 0$. Hence, for fixed s_0 , there is at most one compatible internal-state path $s_1, s_2, \dots, s_L$ satisfying $s_t = \delta(s_{t-1}, a_t)$ for all t . Therefore, each column s_0 of M_w has at most one nonzero entry, located at row $s_L = \delta(\dots \delta(s_0, a_1) \dots, a_L)$ when the path exists. Moreover, along that path,

$$(M_w)_{s_L s_0} = \prod_{t=1}^L \sqrt{T_{s_t s_{t-1}}^{a_t}} \implies |(M_w)_{s_L s_0}|^2 = \prod_{t=1}^L T_{s_t s_{t-1}}^{a_t}. \quad (\text{A8})$$

Summing over s_L in (A7) therefore removes the final-state index without introducing cross terms:

$$\sum_{s_L} |(M_w)_{s_L s_0}|^2 = \prod_{t=1}^L T_{s_t s_{t-1}}^{a_t}, \quad (\text{A9})$$

with the understanding that the product is zero if the path does not exist. Substituting back into (A7) gives

$$\text{Tr}(M_w \rho_\star M_w^\dagger) = \sum_{s_0 \in \mathcal{S}} \pi_{s_0} \prod_{t=1}^L T_{s_t s_{t-1}}^{a_t}. \quad (\text{A10})$$

The right-hand side is exactly the standard unifilar-HMM expression for the stationary word probability $\Pr_{\text{HMM}}(w)$ (sum over initial state weighted by π , with the internal path fixed by unifilarity). This proves (A6). ■

Relation to the q -sample construction. Lemma A4 shows that the square-root tensors always reproduce the *classical* word statistics for any unifilar presentation when probabilities are extracted via the induced instrument. For predictive presentations (e.g., ε machines) this iMPS further coincides with the q -sample state in the sense of Ref. [31]; for general nonpredictive unifilar presentations one should not assume this stronger identification, even though the measured statistics agree.

Implemented labeling heuristics

The dilation of Sec. III A requires choosing an admissible labeling function f satisfying the injectivity condition (12). All admissible choices produce a deterministic HMM that exactly reproduces the original process before truncation; they differ only in how auxiliary labels $y \in \mathcal{Y}$ are assigned to the nonzero outgoing transitions from each (s, x) pair, and hence in the tensor structure seen by the subsequent variational truncation.

In our numerical study (Sec. IV C) we implement the following four heuristic strategies. In each case, for a fixed source state s and emitted symbol x , let $s'_1, \dots, s'_k$ denote the destination states with $T_{s's}^x > 0$, and let $|\mathcal{Y}| = d_y$ be the minimal auxiliary alphabet size.

(1) *Sequential*. Sort the destination states by state index in ascending order and assign auxiliary labels $y = 0, 1, \dots, k - 1$ in that order. This is a deterministic baseline that depends only on the state-labeling convention.

(2) *Probability descending*. Sort the destination states by their transition probabilities $T_{s's}^x$ in decreasing order and assign $y = 0$ to the highest-probability transition, $y = 1$ to the next, and so on.

(3) *Probability ascending*. As above, but with the destination states sorted in increasing order of transition probability, so that $y = 0$ is assigned to the least probable transition.

(4) *Random permutation*. Assign auxiliary labels according to a pseudorandom permutation of the destination states. When a base seed is supplied, the assignment is reproducible.

All four strategies use the same minimal auxiliary alphabet size d_y and produce dilated tensors satisfying $\sum_y T_{s's}^{(x,y)} = T_{s's}^x$ for all s, s', x . They are therefore all exact prior to truncation. As illustrated in Fig. 6, the choice of strategy can materially affect the Schmidt spectrum of the dilated iMPS and hence the compression performance.

APPENDIX B: ERROR BOUNDS

1. From quantum fidelity rate to a classical Bhattacharyya rate

Let $|P_{xy}\rangle$ and $|\tilde{P}_{xy}\rangle$ denote the infinite-chain q -sample states of the exact and truncated dilated processes over $\mathcal{X} \times \mathcal{Y}$. For concreteness we define the quantum fidelity divergence rate (QFDR) as

$$R_F(|P_{xy}\rangle, |\tilde{P}_{xy}\rangle) = - \lim_{L \rightarrow \infty} \frac{1}{2L} \log_2 F(\rho_L, \tilde{\rho}_L), \quad (\text{B1})$$

where ρ_L and $\tilde{\rho}_L$ are the reduced density operators on L sites and $F(\rho, \sigma)$ is the quantum fidelity.

The following lemma relates QFDR to a classical fidelity divergence rate on $\mathcal{X}$, defined by the blockwise Bhattacharyya coefficient $\sum_{\vec{x}} \sqrt{P^{(L)}(\vec{x}) \tilde{P}^{(L)}(\vec{x})}$.

Lemma B1 (Classical fidelity rate bounded by QFDR).

Let $\mathcal{P}$ and $\tilde{\mathcal{P}}$ be the processes on $\mathcal{X}$ obtained by marginalizing the exact and truncated dilated q samples over $\mathcal{Y}$. Define the classical fidelity divergence rate

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) := - \lim_{L \rightarrow \infty} \frac{1}{2L} \log_2 \sum_{\vec{x} \in \mathcal{X}^L} \sqrt{P^{(L)}(\vec{x}) \tilde{P}^{(L)}(\vec{x})}. \quad (\text{B2})$$

Then

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq R_F(|P_{xy}\rangle, |\tilde{P}_{xy}\rangle). \quad (\text{B3})$$

Proof. Consider the reduced density operators ρ_L and $\tilde{\rho}_L$ on L sites. Apply two completely positive trace-preserving (CPTP) maps to both states: (i) dephasing in the computational basis $\{|x_1, y_1; \dots; x_L, y_L\rangle\}$; and (ii) partial trace over the auxiliary outputs $y_1, \dots, y_L$. Let Φ denote the composition of these maps. Fidelity is monotone under CPTP maps, so

$$F(\rho_L, \tilde{\rho}_L) \leq F(\Phi(\rho_L), \Phi(\tilde{\rho}_L)). \quad (\text{B4})$$

The outputs $\Phi(\rho_L)$ and $\Phi(\tilde{\rho}_L)$ are diagonal density matrices on $\mathcal{X}^L$ encoding the classical block distributions $P^{(L)}$ and $\tilde{P}^{(L)}$. Their quantum fidelity reduces to the classical fidelity coefficient $\sum_{\vec{x}} \sqrt{P^{(L)}(\vec{x}) \tilde{P}^{(L)}(\vec{x})}$. Taking $-\frac{1}{2L} \log_2(\cdot)$ and passing to $L \rightarrow \infty$ yields the claim. ■

2. Truncation and quantum fidelity rate

Let $\{\lambda_k\}_{k=1}^{d_s}$ be the Schmidt coefficients of the dilated iMPS across a fixed bond, ordered so that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{d_s}$ and $\sum_k \lambda_k = 1$. For a target bond dimension $\tilde{d}$ we define the discarded tail weight

$$\varepsilon_{\tilde{d}} := \sum_{k > \tilde{d}} \lambda_k. \quad (\text{B5})$$

The effect of discarding Schmidt weight at a single bond of a normal iMPS and replacing it with a variationally optimal truncated iMPS is analyzed in Ref. [32]. We state the relevant scaling result:

Lemma B2 (QFDR versus Schmidt tail). For the normal dilated iMPS and its variational truncation at bond dimension $\tilde{d}$, there exists a constant $c > 0$ (depending on the local physical dimension and the spectral gap of the transfer operator) such that, for sufficiently small $\varepsilon_{\tilde{d}}$,

$$R_F(|P_{xy}\rangle, |\tilde{P}_{xy}\rangle) \leq c \varepsilon_{\tilde{d}}. \quad (\text{B6})$$

Sketch. Discarding Schmidt weight $\varepsilon_{\tilde{d}}$ and constructing the optimal variational approximation at bond dimension $\tilde{d}$ perturbs the iMPS by an amount linear in $\varepsilon_{\tilde{d}}$ in appropriate norms. For a normal iMPS with a primitive transfer operator, a finite-block overlap bound can be derived in terms of $\varepsilon_{\tilde{d}}$ and the spectral gap. Passing to the infinite-chain limit yields the stated bound on the fidelity divergence rate. The full argument is given in Ref. [32]. ■

Combining Lemmas B1 and B2 yields

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq c \varepsilon_{\tilde{d}}. \quad (\text{B7})$$

3. Tail weight and Schmidt entropy

We now bound the tail weight $\varepsilon_{\tilde{d}}$ in terms of the Schmidt entropy

$$H(\lambda) := - \sum_k \lambda_k \log_2 \lambda_k. \quad (\text{B8})$$

Lemma B3 (Entropy controls tail). Let $\{\lambda_k\}$ be a probability distribution in nonincreasing order and let $\varepsilon_{\tilde{d}} = \sum_{k > \tilde{d}} \lambda_k$ for some integer $\tilde{d} \geq 2$. Then

$$\varepsilon_{\tilde{d}} \leq \frac{H(\lambda)}{\log_2 \tilde{d}}. \quad (\text{B9})$$

Proof. For $k > \tilde{d}$ we have $\lambda_k \leq \lambda_{\tilde{d}}$. Since $\sum_{k \leq \tilde{d}} \lambda_k \leq 1$, it follows that $\tilde{d} \lambda_{\tilde{d}} \leq 1$ and hence $\lambda_{\tilde{d}} \leq 1/\tilde{d}$. Thus, $\lambda_k \leq 1/\tilde{d}$ for all $k > \tilde{d}$.

Restricting the entropy sum to the tail gives

$$\begin{aligned} H(\lambda) &= - \sum_k \lambda_k \log_2 \lambda_k \geq - \sum_{k > \tilde{d}} \lambda_k \log_2 \lambda_k \\ &= \sum_{k > \tilde{d}} \lambda_k \log_2 \frac{1}{\lambda_k}. \end{aligned} \quad (\text{B10})$$

For $k > \tilde{d}$ we have $\lambda_k \leq 1/\tilde{d}$, so $\log_2(1/\lambda_k) \geq \log_2 \tilde{d}$, and therefore

$$H(\lambda) \geq \sum_{k > \tilde{d}} \lambda_k \log_2 \tilde{d} = \varepsilon_{\tilde{d}} \log_2 \tilde{d}. \quad (\text{B11})$$

Rearranging yields the claimed inequality. ■

Substituting Lemma B3 into Eq. (B7) gives the entropy-based bound quoted in the main text:

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq c \frac{H(\lambda)}{\log_2 \tilde{d}}. \quad (\text{B12})$$

4. Slice matrix and support of the stationary bond state

To relate $H(\lambda)$ to an algebraic quantity depending on the dilation, we consider the stationary bond density operator $\rho_\star$ of the dilated iMPS. In canonical form, $\rho_\star$ is the fixed point of the bond channel

$$\mathcal{E}(\rho) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} A^{(x,y)} \rho A^{(x,y)\dagger}. \quad (\text{B13})$$

The eigenvalues of $\rho_\star$ are the Schmidt coefficients $\{\lambda_k\}$ across the corresponding bond.

Define the slice matrix K by horizontally concatenating all nonzero site tensors $A^{(x,y)}$, viewed as linear maps on the bond space:

$$K = [A^{(x_1,y_1)} \ A^{(x_2,y_2)} \ \dots]. \quad (\text{B14})$$

The column space of K is the span of the ranges of the individual $A^{(x,y)}$.

Lemma B4 (Support contained in the slice span). Let $\rho_\star$ be the stationary bond state of the dilated iMPS. Then the support of $\rho_\star$ is contained in the column space of K , and

hence

$$\text{rank}(\rho_\star) \leq \text{rank}(K). \quad (\text{B15})$$

Proof. Let ρ_0 be any initial positive operator with full support. Iterating the bond channel gives $\rho_{n+1} = \mathcal{E}(\rho_n)$. Each term $A^{(x,y)} \rho_n A^{(x,y)\dagger}$ has columns in the range of $A^{(x,y)}$, and therefore in the column space of K . Hence, ρ_n has support contained in that space for all $n \geq 1$. For a primitive channel $\mathcal{E}$, the sequence ρ_n converges to the unique fixed point $\rho_\star$ as $n \rightarrow \infty$. The column space of K is closed, so $\rho_\star$ also has support contained in it. The rank bound follows. ■

The Schmidt entropy is the von Neumann entropy of $\rho_\star$ in bits. For any density matrix of rank r , the entropy is maximized by the uniform distribution on its support and satisfies $H(\lambda) \leq \log_2 r$. Applying this with $r = \text{rank}(\rho_\star)$ and using Lemma B4 yields

$$H(\lambda) \leq \log_2 \text{rank}(\rho_\star) \leq \log_2 \text{rank}(K). \quad (\text{B16})$$

5. Combined slice-rank bound

Combining Eq. (B12) with (B16) gives the slice-rank bound

$$R_F(\mathcal{P}, \tilde{\mathcal{P}}) \leq c \frac{\log_2 \text{rank } K}{\log_2 \tilde{d}}, \quad (\text{B17})$$

which appears as Eq. (36) in the main text. The rank of K depends explicitly on the chosen labeling function f through the set of nonzero slices $A^{(x,y)}$, making this a genuinely label-aware certificate that links the achievable compression error to the structure induced by the dilation.

-
- [1] L. Rabiner and B. Juang, An introduction to hidden Markov models, *IEEE ASSP Mag.* **3**, 4 (1986).
 - [2] C. R. Shalizi and J. P. Crutchfield, Computational mechanics: Pattern and prediction, structure and simplicity, *J. Stat. Phys.* **104**, 817 (2001).
 - [3] L. R. Rabiner, A tutorial on hidden Markov models and selected applications in speech recognition, *Proc. IEEE* **77**, 257 (1989).
 - [4] P. Baldi, Y. Chauvin, T. Hunkapiller, and M. A. McClure, Hidden Markov models of biological primary sequence information., *Proc. Natl. Acad. Sci. USA* **91**, 1059 (1994).
 - [5] Z. Ghahramani and M. I. Jordan, Factorial hidden Markov models, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 1996), pp. 472–478.
 - [6] S. Fine, Y. Singer, and N. Tishby, The hierarchical hidden Markov model: Analysis and applications, *Mach. Learn.* **32**, 41 (1998).
 - [7] K. Seymore, A. McCallum, and R. Rosenfeld, Learning hidden Markov model structure for information extraction, in *AAAI-99 Workshop on Machine Learning for Information Extraction* (AAAI Press, Washington, DC, 1999), pp. 37–42.
 - [8] A. Krogh, B. Larsson, G. Von Heijne, and E. L. Sonnhammer, Predicting transmembrane protein topology with a hidden Markov model: Application to complete genomes, *J. Mol. Biol.* **305**, 567 (2001).
 - [9] M. Stanke and S. Waack, Gene prediction with a hidden Markov model and a new intron submodel, *Bioinformatics* **19**, ii215 (2003).
 - [10] C. Karlof and D. Wagner, Hidden Markov model cryptanalysis, in *International Workshop on Cryptographic Hardware and Embedded Systems* (Springer, Berlin, 2003), pp. 17–34.
 - [11] R. Bhar and S. Hamori, *Hidden Markov Models: Applications to Financial Economics* (Springer, New York, 2004), Vol. 40.
 - [12] S. Gammelmark, K. Mølmer, W. Alt, T. Kampschulte, and D. Meschede, Hidden Markov model of atomic quantum jump dynamics in an optically probed cavity, *Phys. Rev. A* **89**, 043839 (2014).
 - [13] P. Warden, Speech commands: A dataset for limited-vocabulary speech recognition, [arXiv:1804.03209](https://arxiv.org/abs/1804.03209).
 - [14] A. Stolcke and S. M. Omohundro, Best-first model merging for hidden Markov model induction, [arXiv:cmp-lg/9405017](https://arxiv.org/abs/cmp-lg/9405017).
 - [15] A. Stolcke and S. M. Omohundro, Best-first model merging for hidden Markov model induction, ICSI Technical Report TR-94-003 (International Computer Science Institute, Berkeley, CA, 1994).
 - [16] P. Dupont, F. Denis, and Y. Esposito, Links between probabilistic automata and hidden Markov models: Probability distributions, learning models and induction algorithms, *Pattern Recognit.* **38**, 1349 (2005).
 - [17] P. Dupont, Probabilistic DFA inference using Kullback-Leibler divergence and Alergia, in *Proceedings of the International Conference on Machine Learning (ICML)* (ACM Press, New York, 2000).
 - [18] S. Singh, M. R. James, and M. R. Rudary, Predictive state representations: A new theory for modeling dynamical systems, [arXiv:1207.4167](https://arxiv.org/abs/1207.4167).

- [19] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Quantum mechanics can reduce the complexity of classical models, *Nat. Commun.* **3**, 762 (2012).
- [20] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, Occam's quantum strop: Synchronizing and compressing classical cryptic processes via a quantum channel, *Sci. Rep.* **6**, 20495 (2016).
- [21] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, Minimized state complexity of quantum-encoded cryptic processes, *Phys. Rev. A* **93**, 052317 (2016).
- [22] C. Aghamohammadi, S. P. Loomis, J. R. Mahoney, and J. P. Crutchfield, Extreme quantum memory advantage for rare-event sampling, *Phys. Rev. X* **8**, 011025 (2018).
- [23] F. C. Binder, J. Thompson, and M. Gu, Practical unitary simulator for non-Markovian complex processes, *Phys. Rev. Lett.* **120**, 240502 (2018).
- [24] Q. Liu, T. J. Elliott, F. C. Binder, C. Di Franco, and M. Gu, Optimal stochastic modeling with unitary quantum dynamics, *Phys. Rev. A* **99**, 062110 (2019).
- [25] T. J. Elliott, Memory compression and thermal efficiency of quantum implementations of nondeterministic hidden Markov models, *Phys. Rev. A* **103**, 052615 (2021).
- [26] S. Adhikary, S. Srinivasan, J. Miller, G. Rabusseau, and B. Boots, Quantum tensor networks, stochastic processes, and weighted automata, in *International Conference on Artificial Intelligence and Statistics* (PMLR, New York, 2021), pp. 2080–2088.
- [27] T. J. Elliott, M. Gu, A. J. P. Garner, and J. Thompson, Quantum adaptive agents with efficient long-term memories, *Phys. Rev. X* **12**, 011007 (2022).
- [28] M. Zonnios, A. Boyd, and F. Binder, Quantum generation of stochastic processes: Spectral invariants and memory bounds, *New J. Phys.* **27**, 064507 (2025).
- [29] P. M. Riechers and T. J. Elliott, Identifiability and minimality bounds of quantum and post-quantum models of classical stochastic processes, *arXiv:2509.03004*.
- [30] M. Schuld and F. Petruccione, Quantum advantages, in *Supervised Learning with Quantum Computers*, Quantum Science and Technology (Springer, Cham, 2018), pp. 127–137.
- [31] C. Yang, F. C. Binder, V. Narasimhachar, and M. Gu, Matrix product states for quantum stochastic modeling, *Phys. Rev. Lett.* **121**, 260602 (2018).
- [32] C. Yang, M. Florido-Llinàs, M. Gu, and T. J. Elliott, Dimension reduction in quantum sampling of stochastic processes, *npj Quantum Inf.* **11**, 34 (2025).
- [33] L. Vanderstraeten, J. Haegeman, and F. Verstraete, Tangent-space methods for variational optimization of uniform matrix product states, *SciPost Phys. Lect. Notes* **7** (2019).
- [34] S. Srinivasan, G. Gordon, and B. Boots, Learning hidden quantum Markov model, *arXiv:1710.09016*.
- [35] D. Pérez-García, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations, *Quantum Inf. Comput.* **7**, 401 (2007).
- [36] A. Khintchine, Korrelationstheorie der stationären stochastischen Prozesse, *Math. Ann.* **109**, 604 (1934).
- [37] J. P. Crutchfield, Between order and chaos, *Nat. Phys.* **8**, 17 (2012).
- [38] D. R. Upper, Theory and algorithms for hidden Markov models and generalized hidden Markov models, Ph.D. thesis, University of California, Berkeley, 1997.
- [39] J. P. Crutchfield and K. Young, Inferring statistical complexity, *Phys. Rev. Lett.* **63**, 105 (1989).
- [40] J. P. Crutchfield, The calculi of emergence: Computation, dynamics and induction, *Phys. D (Amsterdam)* **75**, 11 (1994).
- [41] J. P. Crutchfield and D. P. Feldman, Statistical complexity of simple one-dimensional spin systems, *Phys. Rev. E* **55**, R1239 (1997).
- [42] W. Löhr and N. Ay, Non-sufficient memories that are sufficient for prediction, in *International Conference on Complex Sciences* (Springer, Berlin, 2009), pp. 265–276.
- [43] J. B. Ruebeck, R. G. James, J. R. Mahoney, and J. P. Crutchfield, Prediction and generation of binary Markov processes: Can a finite-state fox catch a Markov mouse? *Chaos* **28**, 013109 (2018).
- [44] M. S. Palsson, M. Gu, J. Ho, H. M. Wiseman, and G. J. Pryde, Experimentally modeling stochastic processes with less memory by the use of a quantum processor, *Sci. Adv.* **3**, e1601302 (2017).
- [45] F. Ghafari, N. Tischler, J. Thompson, M. Gu, L. K. Shalm, V. B. Verma, S. W. Nam, R. B. Patel, H. M. Wiseman, and G. J. Pryde, Dimensional quantum memory advantage in the simulation of stochastic processes, *Phys. Rev. X* **9**, 041013 (2019).
- [46] F. Ghafari, N. Tischler, C. Di Franco, J. Thompson, M. Gu, and G. J. Pryde, Interfering trajectories in experimental quantum-enhanced stochastic simulation, *Nat. Commun.* **10**, 1630 (2019).
- [47] K.-D. Wu, C. Yang, R.-D. He, M. Gu, G.-Y. Xiang, C.-F. Li, G.-C. Guo, and T. J. Elliott, Implementing quantum dimensionality reduction for non-Markovian stochastic simulation, *Nat. Commun.* **14**, 2624 (2023).
- [48] A. Monras, A. Beige, and K. Wiesner, Hidden quantum Markov models and non-adaptive read-out of many-body states, *arXiv:1002.2337*.
- [49] A. Monras and A. Winter, Quantum learning of classical stochastic processes: The completely positive realization problem, *J. Math. Phys.* **57**, 015219 (2016).
- [50] M. Fanizza, J. Lumbreras, and A. Winter, Quantum theory in finite dimension cannot explain every general process with finite memory, *Commun. Math. Phys.* **405**, 50 (2024).
- [51] A. J. P. Garner, Q. Liu, J. Thompson, V. Vedral, and M. Gu, Provably unbounded memory advantage in stochastic simulation using quantum mechanics, *New J. Phys.* **19**, 103009 (2017).
- [52] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield, Extreme quantum advantage when simulating classical systems with long-range interaction, *Sci. Rep.* **7**, 6735 (2017).
- [53] T. J. Elliott and M. Gu, Superior memory efficiency of quantum devices for the simulation of continuous-time stochastic processes, *npj Quantum Inf.* **4**, 18 (2018).
- [54] T. J. Elliott, A. J. P. Garner, and M. Gu, Memory-efficient tracking of complex temporal and symbolic dynamics with quantum simulators, *New J. Phys.* **21**, 013021 (2019).
- [55] T. J. Elliott, C. Yang, F. C. Binder, A. J. Garner, J. Thompson, and M. Gu, Extreme dimensionality reduction with quantum modeling, *Phys. Rev. Lett.* **125**, 260501 (2020).
- [56] T. J. Elliott, Quantum coarse graining for extreme dimensionality reduction in modeling stochastic temporal dynamics, *PRX Quantum* **2**, 020342 (2021).

- [57] T. J. Elliott and M. Gu, Embedding memory-efficient stochastic simulators as quantum trajectories, *Phys. Rev. A* **109**, 022434 (2024).
- [58] L. Banchi, Accuracy vs memory advantage in the quantum simulation of stochastic processes, *Mach. Learn.: Sci. Technol.* **5**, 025036 (2024).
- [59] S. Srinivasan, S. Adhikary, J. Miller, G. Rabusseau, and B. Boots, Quantum tensor networks, stochastic processes, and weighted automata, [arXiv:2010.10653](https://arxiv.org/abs/2010.10653).
- [60] C. Yang, F. C. Binder, M. Gu, and T. J. Elliott, Measures of distinguishability between stochastic processes, *Phys. Rev. E* **101**, 062137 (2020).
- [61] S. E. Marzen and J. P. Crutchfield, Informational and causal architecture of discrete-time renewal processes, *Entropy* **17**, 4891 (2015).
- [62] S. B. Davis and P. Mermelstein, Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences, *IEEE Trans. Acoust. Speech Signal Process.* **28**, 357 (1980).
- [63] D. Sculley, Web-scale K-means clustering, in *Proceedings of the 19th International Conference on World Wide Web (WWW)* (ACM Press, New York, 2010), pp. 1177–1178.
- [64] A. P. Dempster, N. M. Laird, and D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *J. R. Stat. Soc. B* **39**, 1 (1977).
- [65] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Un-supervised generative modeling using matrix product states, *Phys. Rev. X* **8**, 031012 (2018).
- [66] E. M. Stoudenmire, Learning relevant features of data with multi-scale tensor networks, *Quantum Sci. Technol.* **3**, 034003 (2018).
- [67] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and J. I. Cirac, Expressive power of tensor-network factorizations for probabilistic modeling, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Red Hook, NY, 2019), pp. 1496–1508.