Quantum Tensor Networks,
Stochastic Processes, and Weighted Automata

Siddarth Srinivasan[*]
School of Computer Science and Engineering
University of Washington
sidsrini@cs.washington.edu

Sandesh Adhikary[*]
School of Computer Science and Engineering
University of Washington
adhikary@cs.washington.edu

Jacob Miller
Mila & DIRO
Université de Montréal
jmjacobmiller@gmail.com

Guillaume Rabusseau
CCAI chair - Mila & DIRO
Université de Montréal
grabus@iro.umontreal.ca

Byron Boots
School of Computer Science and Engineering
University of Washington
bboots@cs.washington.edu

Abstract

Modeling joint probability distributions over sequences has been studied from many
perspectives. The physics community developed *matrix product states*, a tensor-train
decomposition for probabilistic modeling, motivated by the need to tractably model
many-body systems. But similar models have also been studied in the stochastic
processes and weighted automata literature, with little work on how these bodies of
work relate to each other. We address this gap by showing how stationary or uniform
versions of popular quantum tensor network models have equivalent representations
in the stochastic processes and weighted automata literature, in the limit of infinitely
long sequences. We demonstrate several equivalence results between models used
in these three communities: (i) uniform variants of matrix product states, Born
machines and locally purified states from the quantum tensor networks literature, (ii)
predictive state representations, hidden Markov models, norm-observable operator
models and hidden quantum Markov models from the stochastic process literature,
and (iii) stochastic weighted automata, probabilistic automata and quadratic
weighted automata from the formal languages literature. Such connections may
open the door for results and methods developed in one area to be applied in another.

1 Introduction

Matrix product states (MPS) were first developed by the physics community as compact representations
of often intractable wave functions of complex quantum systems (Perez-Garcia et al., 2006; Klumper
et al., 1993; Fannes et al., 1992), in parallel with the equivalent tensor-train decomposition (Oseledets,
2011) developed in applied mathematics for high-order tensors. These tensor network models have
been gaining popularity in machine learning especially as means of compressing highly-parameterized
models (Novikov et al., 2015; Garipov et al., 2016; Yu et al., 2017; Novikov et al., 2014). There has also
been recent interest in directly connecting ideas and methods from quantum tensor networks to machine
learning (Stoudenmire and Schwab, 2016; Han et al., 2018; Guo et al., 2018; Huggins et al., 2019).

---

[*]denotes equal contribution

In particular, tensor networks have been used for probabilistic modeling as parameterizations of joint probability tensors (Glasser et al., 2019; Miller et al., 2020; Stokes and Terilla, 2019). But the same problem has also been studied from various other perspectives. Notably, observable operator models (Jaeger, 2000) or predictive state representations (PSRs) (Singh et al., 2004) from the machine learning literature and stochastic weighted automata (SWA) (Balle et al., 2014b) from the formal languages literature are approaches to tackle essentially the same problem. While Thon and Jaeger (2015) provide an overview discussing connections between PSRs and SWA, their connection to MPS has not been extensively explored. At the same time, there exist many variants of tensor network models related to MPS that can be used for probabilistic modeling. Glasser et al. (2019) recently provided a thorough investigation of the relative expressiveness of various tensor networks for the *non-uniform* case (where cores in the tensor decomposition need not be identical). However, to the best of our knowledge, similar relationships have not yet been established for the *uniform* case. The inherent weight sharing in uniform tensor networks leads to particularly compact models, especially when learning from highly structured data. The reduced number of parameters in uniform networks can also correspond to a smaller set of quantum gates when implementing these models on resource constrained near-term quantum computers. Despite potential useful properties of uniformity, the tensor-network literature tends to focus more on non-uniform models. We aim to fill this gap by developing expressiveness relationships for uniform variants and how they relate to work in different fields, and derive a collection of results analyzing the relationships in expressiveness between uniform MPS and their various subclasses.

We expect that the connections established in this paper will also open the door for results and methods in one area to be used in another. For instance, one of the proof strategies we adopt is to develop expressiveness relationships between subclasses of PSRs, and show how they also carry over to equivalent uniform tensor networks. Such cross fertilization also takes place at the level of algorithms. For instance, the learning algorithm for locally purified states (LPS) employed in Glasser et al. (2019) does not preserve uniformity of the model across time-steps, or enforce normalization constraints on learned operators. With the equivalence between uniform LPS and hidden quantum Markov models (HQMMs) established in this paper, the HQMM learning algorithm from Adhikary et al. (2020), based on constrained optimization on the Stiefel manifold, can be adapted to learn uniform LPS *while enforcing all appropriate constraints*. Similarly, spectral algorithms that have been developed for stochastic process models such as hidden Markov models (HMMs) and PSRs (Hsu et al., 2012; Siddiqi et al., 2010; Bailly et al., 2009) could potentially be adapted to learn uniform LPS and uniform MPS models. Finally, one of our key motivations is to simply provide a means of translating between similar models developed in different fields. While prior works (Glasser et al., 2019; Kliesch et al., 2014; Critch et al., 2014) have noted similarities between tensor networks, stochastic processes and weighted automata, many formal and explicit connections were still lacking, especially in the context of model expressiveness. We hope that the connections developed in this paper will help bring together complementary advances occurring in these various fields.

Summary of Contributions   In Section 2, we demonstrate that uniform Matrix states (uMPS) are equivalent to predictive state representations and stochastic weighted automata, when taken in the *non-terminating limit* (where we evaluate probabilities sufficiently away from the end of a sequence). Section 3 presents the known equivalence between uMPS with non-negative parameters, HMMs, and probabilistic automata, to show in Section 4 that another subclass of uMPS called Born machines (BM) (Han et al., 2018) is equivalent to norm-observable operator models (NOOM) (Zhao and Jaeger, 2010) and quadratic weighted automata (QWA) (Bailly, 2011). We also demonstrate that uBMs and NOOMs are relatively restrictive model classes in that there are HMMs with no equivalent finite-dimensional uBM or NOOM (HMMs $\nsubseteq$ NOOMs/uBMs). Finally, in Section 5, we analyze a broadly expressive subclass of uMPS known as locally purified states (LPS), demonstrate its equivalence to hidden quantum Markov models, and discuss the open question of how the expressiveness of uLPS relates to that of uMPS. We thus develop a unifying perspective on a wide range of models coming from disparate communities, providing a rigorous characterization of how they are related to one another, as illustrated in Figure 1. An overview of notation and proofs for all theorems are provided in the Appendix.

2   Uniform Matrix Product States

Given a sequence of N observations, where each outcome can take d_i values, the joint probability of any particular sequence $y_1,...,y_N$ can be written using the following tensor-train decomposition, which gives an MPS:

$$P(y_1,...,y_N) \propto \mathrm{MPS}_{y_1,...,y_N} = \mathbf{A}^{[N],y_N} \mathbf{A}^{[N-1],y_{N-1}} ... \mathbf{A}^{[2],y_2} \mathbf{A}^{[1],y_1} \tag{1}$$

where each $\mathbf{A}^{[i]}$ is a three-mode tensor *core* of the MPS containing d_i slices, with the matrix slice associated with outcome Y_i denoted by $\mathbf{A}^{[i],y_i}$. Each slice is a $D_{i+1} \times D_i$ matrix, and the conventional

choice (which we use in this paper) of *open boundary conditions* is to set $D_0 = D_N = 1$ (i.e. $\mathbf{A}^{[1],y_1}$ and $\mathbf{A}^{[N],y_N}$ are column and row vectors respectively). The maximal value of $D = \max_k D_k$ is also called the bond-dimension or the TT-rank (Glasser et al., 2019) of the MPS. MPS with open boundaries are equivalent to tensor train (TT) decompositions, and we will define them over the complex field, a choice common in quantum physics and tensor network settings.

In this paper, we will focus on the "uniform" case of identical cores, i.e., a uniform MPS or uMPS. uMPS models were first developed in the quantum physics community (Perez-Garcia et al., 2006; Vanderstraeten et al., 2019), although employing a different probabilistic correspondence (Born machines as discussed later) than described below. As we will discuss, this corresponds naturally to Markovian dynamical systems; the notion of *past being independent of future given the present* is encoded by the tensor train structure where each core only has two neighbours. While an MPS is inherently defined with respect to a fixed sequence length, a uMPS can be applied to sequences of arbitrary fixed or infinite length (Cirac and Sierra, 2010). As there should be no distinction between the cores at different time steps in a uMPS, a natural representation is to fix two boundary vectors $\vec{\sigma}^\dagger$ († denotes the complex conjugate transpose) and $\vec{\rho}_0$, leading to the following decomposition of the joint probability:

$$P(y_1,...,y_N) = \text{uMPS}_{y_1,...,y_N} = \vec{\sigma}^\dagger \mathbf{A}^{y_N} \mathbf{A}^{y_{N-1}} ... \mathbf{A}^{y_2} \mathbf{A}^{y_1} \vec{\rho}_0 \tag{2}$$

To explore connections with arbitrary-length PSRs and WFAs, we will particularly focus on the *non-terminating limit*. Consider that if we wished to compute the probability of some subsequence from $t = 1, ..., T$ of the N-length uMPS ($T < N$), we could compute $P(y_1,...,y_T) = \sum_{y_N} \cdots \sum_{y_{T+1}} P(y_1,...,y_T,y_{T+1},...,y_N)$. The non-terminating limit is essentially when we consider the uMPS to be infinitely long, i.e., we compute the probability of the subsequence in the limit as $N \to \infty$.

Definition 1 (Non-terminating uMPS). A non-terminating uMPS is an infinitely long uMPS where we can compute the probability of any sequence $P(y_1,...,y_T)$ of length T by marginalizing over infinitely many future observations, i.e. $P(y_1,...,y_T) = \lim_{N\to\infty} \sum_{y_N} \cdots \sum_{y_{T+1}} P(y_1,...,y_T,y_{T+1},...,y_N)$.

This is a natural approach to modeling arbitrary length sequences with Markovian dynamics; intuitively, if given an identical set of tensor cores at each time step, the probability of a sequence should not depend on how far it is from the 'end' of the sequence. Similar notions are routinely used in machine learning and physics. In machine learning, it is common to discard the first few entries of sequences as "burn-in" to allow systems to reach their stationary distribution. In our case, the burn is being applied to the end of the sequence. The non-terminating limit is also similar to the "thermodynamic limit" employed in many-body physics, which marginalizes over an infinite number of future *and* past observations (Vanderstraeten et al., 2019). Such limits reflect the behavior seen in the interior of large systems, and avoid more complicated phenomena which arise near the beginning or end of sequences.

2.1 The Many Names of Matrix Product States

While connections between MPS and hidden Markov models (HMM) have been widely noted, we point out that non-terminating uMPS models have been studied from various perspectives, and are referred to by different names in the literature, such as stochastic weighted finite automata (stochastic WFA) (Balle et al., 2014a), quasi-realizations (Vidyasagar, 2011), observable operator models (OOM) (Jaeger, 2000), and (uncontrolled) predictive state representations (PSR) (Singh et al., 2004). The latter three models are exactly identical (we just refer to them as uncontrolled PSRs in this paper) and come from the stochastic processes perspective, while stochastic WFA are slightly different in their formulation, in that they are more similar to (terminating) uMPS (see below). Thon and Jaeger (2015) detail a general framework of *sequential systems* to study how PSRs and WFA relate to each other.

Predictive State Representations We write the stochastic process defined by an n–dimensional predictive state representation over a set of discrete observations $\mathcal{O}$ as a tuple $(\mathbb{C}^n, \vec{\sigma}, \{\boldsymbol{\tau}_y\}_{y\in\mathcal{O}}, \vec{x}_0)$. The initial state $\vec{x}_0 \in \mathbb{C}^n$ is normalized, as enforced by the linear evaluation functional $\vec{\sigma}$, i.e., $\vec{\sigma}^\dagger \vec{x}_0 = 1$, and the observable operators are constrained to have normalized marginals over observations $\vec{\sigma}^\dagger \sum_y \boldsymbol{\tau}_y = \vec{\sigma}^\dagger$, i.e., $\vec{\sigma}^\dagger$ is a fixed point of the *transfer operator* $\sum_y \boldsymbol{\tau}_y$. The probability of arbitrary length sequences $y_1, ..., y_T \in \mathcal{O}^T$ is computed as $\vec{\sigma}^\dagger \boldsymbol{\tau}_{y_T} ... \boldsymbol{\tau}_{y_1} \vec{x}_0$, which should be non-negative for any sequence. Note that we simply require this to hold for a valid PSR; we do not explicitly enforce constraints to ensure this. This joint probability computation is identical to Equation 2, where evaluation functional $\vec{\sigma}$ and the initial state $\vec{x}_0$ are analogous to the left and right boundary vectors of the uMPS, and the *observable operators* $\boldsymbol{\tau}_y$ correspond to the matrix slices $\mathbf{A}^y$. In this sense, both uMPS and PSRs define tensor-train decompositions of joint distributions for a given fixed number of observations T. The only difference is that a uMPS does not require its evaluation functional to be the fixed point of its transfer operator. However, as we now discuss, any arbitrary uMPS evaluation
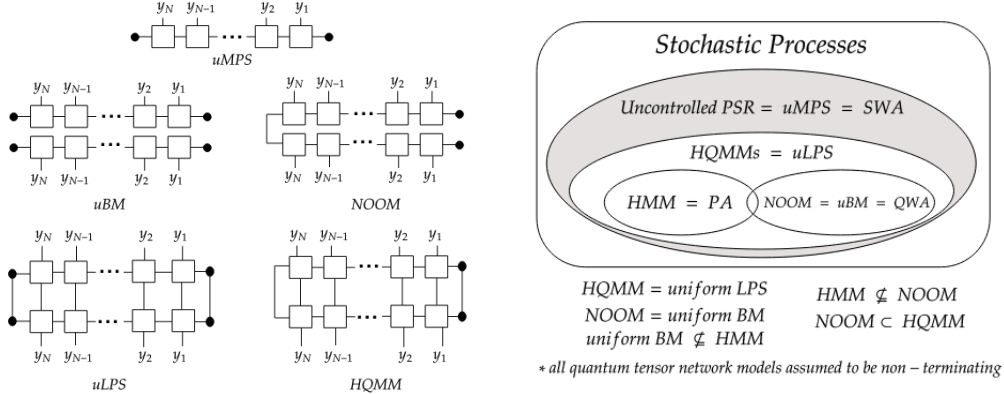
Figure 1: (Left) Tensor network diagrams (Right) Subset relationships between stochastic process models, non-terminating uniform quantum tensor networks, and weighted automata, along with a summary of new relationships established in this paper. The grey region is potentially empty.

functional will eventually converge to the fixed point of its transfer operator in the non-terminating limit. The fixed point then becomes the *effective* evaluation functional of the uMPS in this limit.

Since PSRs were formulated with dynamical systems in mind, we typically consider sequences of *arbitrary* length, whose probabilities are determined via a hidden state which evolves under a time-invariant update rule: the state update conditioned on an observation y_t is computed as $\vec{x}_t = \frac{\boldsymbol{\tau}_{y_t} \vec{x}_{t-1}}{\vec{\sigma}^T \boldsymbol{\tau}_{y_t} \vec{x}_{t-1}}$ and the probability of an observation y_t is $P(y_t|\vec{x}_t) = \vec{\sigma}^T \boldsymbol{\tau}_{y_t} \vec{x}_t$. This allows us to deal more flexibly with arbitrary length sequences as compared to a generic uMPS. This flexibility for arbitrary-length sequences is precisely why we consider non-terminating uMPS: we can compute the conditional probability of a sequence $P(y_t|y_{1:t-1})$ by marginalizing over all possible future observations with a relatively simple equation:

$$P(y_t|y_{1:t-1}) = \frac{\sum_{y_N,...,y_{t+1}} P(y_N,...,y_{t+1},y_t,y_{t-1},...y_1)}{\sum_{y_N,...,y_t} P(y_N,...,y_t,y_{t-1},...,y_1)} = \frac{\vec{\sigma}^\dagger \left(\sum_y \boldsymbol{\tau}_y\right)^{N-(t+1)} \boldsymbol{\tau}_{y_t}...\boldsymbol{\tau}_1 \vec{\rho}_0}{\vec{\sigma}^\dagger \left(\sum_y \boldsymbol{\tau}_y\right)^{N-t} \boldsymbol{\tau}_{y_{t-1}}...\boldsymbol{\tau}_1 \vec{\rho}_0} \quad (3)$$

Thus the *effective* evaluation functional $\vec{\sigma}^\dagger \left(\sum_y \boldsymbol{\tau}_y\right)^{N-t}$ is a function of time and so different at every time step in general. However, if the transfer operator $\boldsymbol{\tau} = \sum_y \boldsymbol{\tau}_y$ has some fixed point $\vec{\sigma}_*^\dagger$, i.e., $\vec{\sigma}_*^\dagger \boldsymbol{\tau} = \vec{\sigma}_*^\dagger$, then the effective evaluation functional at timestep t (which is $N-t$ steps away from the left boundary of the uMPS) will eventually converge to $\vec{\sigma}_*$ given a long enough sequence.[2] So, as long as we remain sufficiently far from the end of a sequence, the particular choice of the the left boundary vector does not matter.

Given that the non-terminating limit effectively replaces the uMPS evaluation functional $\vec{\sigma}$ by the fixed point $\vec{\sigma}_*$, consider what happens if we require $\vec{\sigma} = \vec{\sigma}_*$ to begin with, as is the case for PSRs. In this case, our effective evaluation functional remains independent of t, permitting a simple recursive state update rule that does not require fixing a prior sequence length or marginalizing over future observations. In this sense, a non-terminating uMPS is strictly equivalent to a PSR, a relationship which we will see holds between several other model families.

Theorem 2.1. *Non-terminating uniform matrix product states are equivalent to uncontrolled predictive state representations.*

If we do not consider the non-terminating limit of a uMPS, the subsequence length and boundary choice will affect the probability computed. Then, we technically only have an equivalence with PSRs for a fixed sequence length (when the evaluation functional is a fixed point of the transfer operator) and no notion of recursive state update.

---

[2]In fact, the effective evaluation functional will converge at an exponential rate towards the fixed point, so that $\|\vec{\sigma}_t - \vec{\sigma}_*\|^2 = \mathcal{O}(\exp \frac{N-t}{\xi})$, with a "correlation length" $\xi \simeq (1-|\lambda_2|/|\lambda_1|)^{-1}$ set by the ratio of the largest and second largest eigenvalues of the transfer operator (Orús, 2014). Transfer operators with non-degenerate spectra can always be rescaled to have a unique fixed point, while matrices with degenerate spectra form a measure zero subset.

Stochastic Weighted Automata A weighted automaton (WA) is a tuple $(\mathbb{C}^n, \vec{\sigma}, \{\boldsymbol{\tau}_y\}_{y \in \mathcal{O}}, \vec{x}_0)$ which computes a function $f(y_1,...,y_T) = \vec{\sigma}^\dagger \boldsymbol{\tau}_{y_T}...\boldsymbol{\tau}_{y_1} \vec{x}_0$. In contrast with PSRs, no constraints are enforced on the weights of a weighted automaton in general. A weighted automaton is *stochastic* if it computes a probability distribution. These models constitute another class of models equivalent to PSRs and uMPS, and represent probability distributions over sequences of symbols from an alphabet $\mathcal{O}$. It is worth mentioning that the semantics of the probabilities computed by PSRs and stochastic WAs can differ: while PSR typically maintain a recursive state and are used to compute the probability of a given sequence conditioned on some past sequence, stochastic WA are often used to compute the joint distributions over the set of all possible finite length sequences (just as in uMPS). We refer the reader to Thon and Jaeger (2015) for a unifying perspective.

3 Non-Negative uMPS, Hidden Markov Models, and Probabilistic Automata

We first point out a well-known connection between hidden Markov models (HMM) and matrix product states (Kliesch et al., 2014; Critch and Morton, 2014). We refer to uMPS where all tensor cores and boundary vectors are non-negative as *non-negative uMPS*. HMMs have been extensively studied in machine learning and are a common approach to modeling discrete-observation sequences where an unobserved hidden state undergoes Markovian evolution (where the future is independent of past given present) and emits observations at each time-step (Rabiner and Juang, 1986). HMMs can be thought of as a special case of PSRs where all parameters are constrained to be non-negative. Such models are usually characterized by an initial *belief state* $\vec{x}_0$ and column-stochastic transition and emission matrices $\mathbf{A}$ and $\mathbf{C}$. Formally, we give the following definition:

Definition 2 (Hidden Markov Model)**.** An $n-$dimensional hidden Markov model for a set of discrete observations $\mathcal{O}$ is a stochastic process described by the tuple $(\mathbb{R}^n, \mathbf{A}, \mathbf{C}, \vec{x}_0)$. The transition matrix $\mathbf{A} \in \mathbb{R}^{n \times n}_{\geq 0}$ and the emission matrix $\mathbf{C} \in \mathbb{R}^{|\mathcal{O}| \times n}_{\geq 0}$ are non-negative and column stochastic i.e. $\vec{\mathbf{1}}^T \mathbf{A} = \vec{\mathbf{1}}^T \mathbf{C} = \vec{\mathbf{1}}^T$. The initial state $\vec{x}_0 \in \mathbb{R}^n_{\geq 0}$ is also non-negative and is normalized $||\vec{x}_0||_1 = \vec{\mathbf{1}}^T \vec{x} = 1$.

The state transitions through the simple linear update $\vec{x}'_t = \mathbf{A} \vec{x}_{t-1}$. To condition on observation y, we construct the diagonal matrix $\mathrm{diag}(\mathbf{C}_{(y,:)})$ from the y^{th} row of $\mathbf{C}$, and perform a normalized update $\vec{x}_t | y_t = \frac{\mathrm{diag}(\mathbf{C}_{(y,:)}) \vec{x}'_t}{\vec{\mathbf{1}}^T \mathrm{diag}(\mathbf{C}_{(y,:)}) \vec{x}'_t}$. This multi-step filtering process can be simplified using an alternative representation with *observable operators* as $\mathbf{T}_y = \mathrm{diag}(\mathbf{C}_{(y,:)}) \mathbf{A}$, where we rewrite the normalization constraints on operators as $\vec{\mathbf{1}}^T \sum_y \mathbf{T}_y = \vec{\mathbf{1}}^T$. Then, we recover a recursive state update $\vec{x}_t = \frac{\mathbf{T}_{y_{t-1}} \vec{x}_{t-1}}{\vec{\mathbf{1}}^T \mathbf{T}_{y_{t-1}} \vec{x}_{t-1}}$. Clearly, HMMs are a special case of PSRs where the parameters are restricted to be non-negative and the evaluation functional $\vec{\sigma}$ is set to be $\vec{\mathbf{1}}$. Alternatively, we can also view HMMs as a special case of uMPS when the left boundary $\vec{\sigma}^\dagger = \vec{\mathbf{1}}^T$, the right boundary $\vec{\rho}_0 = \vec{x}_0$, and the tensor core slice $\mathbf{A}^y = \mathbf{T}_y$.

Probabilistic Automata Lastly, non-negative uMPS are equivalent to probabilistic automata from formal language theory (Denis and Esposito, 2008, Section 4.2), which are in essence weighted automata where transition matrices need to satisfy stochasticity constraints. The strict equivalence between probabilistic automata and HMMs is proved in Dupont et al. (2005, Proposition 8) (see also Section 2.2 in Balle et al. (2014b)). In addition, it is known that non-negative uMPS are strictly less expressive than general uMPS for representing probability distributions; a proof of this result in the context of formal languages can be found in Denis and Esposito (2008). We give a brief discussion of this next.

The Negative Probability Problem and the Positive Realization Problem As noted by several authors (Jaeger, 2000; Adhikary et al., 2020), PSRs lack a constructive definition; the definition of a PSR simply demands that the probabilities produced by the model be non-negative without describing constraints that can achieve this. Indeed, this is the cost of relaxing the non-negativity constraint on HMMs; it is undecidable whether a given set of PSR parameters will assign a negative probability to some arbitrary-length sequence (Denis and Esposito, 2004; Wiewiora, 2008), an issue known as the *negative probability problem* (NPP). While this undecidability is an inconvenient feature of PSRs, it turns out that constraining PSRs to have only non-negative entries comes with a reduction in expressive power; there are finite (state/bond) dimensional PSRs/uMPS which have no equivalent finite-dimensional HMM representations for arbitrary length sequences (for example, the probability clock in Jaeger (2000)). The general question of which uncontrolled PSRs/uMPS have equivalent finite-dimensional HMMs (though not always discussed in those terms) is referred to by some as the *positive realization problem* (Benvenuti and Farina, 2004; Vidyasagar, 2011).

4 Uniform Born Machines, Norm-Observable Operator Models, and Quadratic Weighted Automata

Born machines (BMs) (Han et al., 2018) are a popular class of quantum tensor networks that model probability densities as the absolute-square of the outputs of a tensor-train decomposition, and hence always output valid probabilities. As with uMPS, we will work with uniform Born machines (uBMs) (Miller et al., 2020), for which the joint probability of N discrete random variables $\{Y_i\}_{i=1}^N$ is computed as follows (with boundary vectors $\vec{\alpha}$ and $\vec{\omega}_0$ sandwiching a sequence of identical cores $\mathbf{A}$):

$$P(y_1,...,y_N) = \text{uBM}_{y_1,...,y_N} = \left| \vec{\alpha}^{\,\dagger} \mathbf{A}^{y_N}...\mathbf{A}^{y_1} \vec{\omega}_0 \right|^2 \tag{4}$$

We can re-write this decomposition showing uBMs to be special kinds of uMPS/PSR:

$$\text{uBM}_{y_1,...,y_N} = \left| \vec{\alpha}^{\,\dagger} \mathbf{A}^{y_N}...\mathbf{A}^{y_1} \vec{\omega}_0 \right|^2 = \vec{\alpha}^\dagger \mathbf{A}^{y_N}...\mathbf{A}^{y_1} \vec{\omega}_0 \, \vec{\omega}_0^\dagger (\mathbf{A}^{y_1})^\dagger...(\mathbf{A}^{y_N})^\dagger \vec{\alpha} = \vec{\sigma}^\dagger \boldsymbol{\tau}_{y_N}...\boldsymbol{\tau}_{y_1} \vec{\rho}_0 \tag{5}$$

where $\boldsymbol{\tau}_y = \overline{\mathbf{A}^y} \otimes \mathbf{A}^y$, $\vec{\rho}_0 = \overline{\vec{\omega}_0} \otimes \vec{\omega}_0$, and $\vec{\sigma} = \overline{\vec{\alpha}} \otimes \vec{\alpha}$. This makes it clear that uBMs are a special class of MPS/PSRs, where the observable operators $\boldsymbol{\tau}_y$ and boundary conditions must satisfy the additional requirement of having unit Kraus-rank[3] (i.e., a symmetric unit Schmidt rank decomposition) (Wolf, 2012).

Norm Observable Operator Models Motivated by the NPP for PSRs, Zhao and Jaeger (2010) introduce norm-observable operator models or NOOMs. Coming from the PSR literature, they were designed to model joint distributions of observations as well as recursive state-updates to obtain conditional probabilities (analogous to PSRs in Section 2.1). They bear a striking resemblance to uniform Born machines (uBMs) and the connection has not been previously explored. Both NOOMs and uBMs associate probabilities with quadratic functions of the state vector, with NOOMs directly using the squared 2-norm of the state to determine observation probabilities. While NOOMs were originally defined on the reals, we use a more general definition over complex numbers.

Definition 3 (Norm-observable operator model). An n-dimensional norm-observable operator model for a set of discrete observations $\mathcal{O}$ is a stochastic process described by the tuple $(\mathbb{C}^n, \{\boldsymbol{\phi}_y\}_{y\in\mathcal{O}}, \vec{\psi}_0)$. The initial state $\vec{\psi}_0 \in \mathbb{C}^n$ is normalized by having unit 2-norm i.e. $\|\vec{\psi}_0\|_2^2 = 1$. The operators $\boldsymbol{\phi}_y \in \mathbb{C}^{n\times n}$ satisfy $\sum_y \boldsymbol{\phi}_y^\dagger \boldsymbol{\phi}_y = \mathbb{I}$.

These models avoid the NPP by using the 2-norm of the state to recover probability which, unlike for HMMs, is insensitive to the use of negative parameters in the matrices $\boldsymbol{\phi}_y$. We write the joint probability of a sequence as computed by a NOOM and manipulate it using using the relationship between 2-norm and trace to show:

$$P(y_1,...,y_N) = \left\| \boldsymbol{\phi}_{y_N}...\boldsymbol{\phi}_{y_1} \vec{\psi}_0 \right\|_2^2 = \text{tr}(\boldsymbol{\phi}_{y_N}...\boldsymbol{\phi}_{y_1} \vec{\psi}_0 \vec{\psi}_0^\dagger (\boldsymbol{\phi}_{y_1})^\dagger...(\boldsymbol{\phi}_{y_N})^\dagger) = \vec{\mathbb{I}}^T \boldsymbol{\tau}_{y_N} ... \boldsymbol{\tau}_{y_1} \vec{\rho}_0 \tag{6}$$

where $\boldsymbol{\tau}_y = \overline{\boldsymbol{\phi}_y} \otimes \boldsymbol{\phi}_y \in \mathbb{C}^{n^2 \times n^2}$ and $\vec{\rho}_0 = \overline{\vec{\psi}_0} \otimes \vec{\psi}_0$. Equation 6 shows that NOOMs are a special subset of PSRs/MPS, as every finite-dimensional NOOM has an equivalent finite-dimensional PSR $(\mathbb{C}^{n^2}, \vec{\mathbb{I}}, \{\overline{\boldsymbol{\phi}_y} \otimes \boldsymbol{\phi}_y\}_{y\in\mathcal{O}}, \overline{\vec{\psi}_0} \otimes \vec{\psi}_0)$ (Zhao and Jaeger, 2010). From a quantum mechanical perspective, the unit rank constraint on NOOM initial states can be framed as requiring the initial state to be a pure density matrix. We can also recursively update the state conditioned on observation y_t as $\vec{\psi}_t = \frac{\boldsymbol{\phi}_{y_t} \vec{\psi}_{t-1}}{\|\boldsymbol{\phi}_{y_t} \vec{\psi}_{t-1}\|_2}$, where y_t is observed with probability $P(y_t|\vec{\psi}_t) = \|\boldsymbol{\phi}_{y_t} \vec{\psi}_{t-1}\|_2^2$.

Non-terminating uniform BMs are NOOMs Note that the NOOM joint distribution in Equation 6 is almost identical to that of uBMs in Equation 5, with $\boldsymbol{\tau}_y$ and $\vec{\rho}_0$ having unit Kraus rank; but the left boundary / evaluation functional $\vec{\sigma} = \overline{\vec{\alpha}} \otimes \vec{\alpha}$ is replaced by $\vec{\sigma} = \vec{\mathbb{I}}$ and necessarily is full Kraus rank. So how can we reconcile these nearly identical models? Similar to our approach in Section 2.1, we can consider the uBM for an infinitely long sequence where the exact specification of the left boundary / evaluation functional ceases to matter in the non-terminating limit; the effective evaluation functional converges to the fixed point of the transfer operator and we have a notion of recursive state update. Assuming that the uBM transfer operator is a similarity transform away from a trace-preserving quantum channel (i.e., which is normalized by satisfying $\sum_y \boldsymbol{\phi}_y^\dagger \boldsymbol{\phi}_y = \mathbb{I}$; see appendix for more details), we have that an arbitrary evaluation functional (with unit Kraus-rank) of such a uBM will eventually converge to $\vec{\mathbb{I}}$, the NOOM's evaluation functional:

---

[3] An operator $\mathbf{A}$ has unit Kraus-rank if it has a decomposition $\mathbf{A} = \overline{\mathbf{X}} \otimes \mathbf{X}$.

Theorem 4.1. *Non-terminating uniform Born machines are equivalent to norm observable operator models.*

With the above equivalence, we now turn to the question of how the expressiveness of uBM/NOOMs compares to non-negative uMPS/HMMs, As we have seen, they are all special classes of uMPS, but with different constructions. Glasser et al. (2019) studied the expressiveness of *non-uniform* BMs, showing that there are finite-dimensional non-uniform BMs that cannot be modeled by finite dimensional non-uniform HMMs, and conjecture that the reverse direction is also true. Zhao and Jaeger (2010) showed by example the existence of a NOOM (and so a non-terminating uBM) that cannot be modeled by any finite-dimensional HMM. However, they left open the question of whether HMMs were a subclass of NOOMs. We answer this question in the following theorem, which also implies the latter corollary through its equivalence with non-terminating uBM.

Theorem 4.2 (HMM $\not\subseteq$ NOOM). *There exist finite-dimensional hidden Markov models that have no equivalent finite-dimensional norm-observable operator model.*

Corollary 4.1 (uBM $\not\subseteq$ HMM and HMM $\not\subseteq$ uBM). *There exist finite-dimensional non-terminating uniform Born machines that have no equivalent finite-dimensional hidden Markov models, and vice-versa.*

Quadratic Weighted Automata   Finally, we note that Quadratic Weighted Automata (QWA) (Bailly, 2011), developed in the stochastic weighted automata literature, are equivalent to uBM. Bailly (2011) suggest that QWA $\not\subseteq$ HMM and that HMM $\not\subseteq$ QWA, but do not provide a proof. To the best of our knowledge, the proof we provide is the first to formally show the non-equivalence of QWA and HMM.

5   Locally Purified States and Hidden Quantum Markov Models

While uBMs/NOOMs are constructive models guaranteed to return valid probabilities, they still aren't expressive enough to capture all HMMs, a fairly general class. Hence, it may be desirable to identify a construction that is more expressive than these models but still gives valid probabilities. Locally Purified States (LPS) were proposed as a tensor-network model of discrete multivariate probability distributions inspired from techniques used in the simulation of quantum systems. Glasser et al. (2019) show that these models are not only strictly more expressive than non-uniform HMMs, but also correspond directly to local quantum circuits with ancillary qubits – serving as a guide to design quantum circuits for probabilistic modeling. We arrive at the LPS model from the MPS model essentially by marginalizing over an additional mode – called the "purification dimension" – in each of the MPS tensors. The rank of an LPS, also called its puri-rank, is defined the same way as the bond dimension (or TT-rank) for the MPS. The corresponding uniform LPS defines the unnormalized probability mass function over N discrete random variables $\{Y_i\}_{i=1}^N$ as follows:

$$P(y_1,...,y_N) = \left(\sum_{\beta_L}\overline{\mathbf{K}}_{\beta_L,L}^T \otimes \mathbf{K}_{\beta_L,L}^T\right)\left(\sum_{\beta}\overline{\mathbf{K}}_{\beta,y_N} \otimes \mathbf{K}_{\beta,y_N}\right)\cdots\cdots\left(\sum_{\beta}\overline{\mathbf{K}}_{\beta,y_1} \otimes \mathbf{K}_{\beta,y_1}\right)\left(\sum_{\beta_R}\overline{\mathbf{K}}_{\beta_R,R} \otimes \mathbf{K}_{\beta_R,R}\right)$$
(7)

Hidden Quantum Markov Models   Hidden quantum Markov models (HQMMs) were developed by Monras et al. (2010) as a quantum generalization of hidden Markov models that can model joint probabilities of sequences and also allow for recursive state updates we have described previously. Srinivasan et al. (2018b,a) specifically develop HQMMs by constructing quantum analogues of classical operations on graphical models, and show that HQMMs are a more general model class compared to HMMs. Adhikary et al. (2020) on the other hand develop HQMMs by relaxing the unit Kraus-rank constraint on NOOM operators and initial state. We give a formal definition of these models here (noting that the Choi matrix $\mathbf{C}_y$ is a particular reshuffling of the sum of superoperators $\mathbf{L}_y$ defined below, see Adhikary et al. (2020)):

Definition 4 (Hidden Quantum Markov Models). An n^2-dimensional hidden quantum Markov model for a set of discrete observations $\mathcal{O}$ is a stochastic process described by the tuple $(\mathbb{C}^{n^2}, \vec{\mathbb{I}}, \{\mathbf{L}_y\}_{y\in\mathcal{O}}, \vec{\rho})$. The initial state $\vec{\rho}_0 \in \mathbb{C}^{n^2}$ is a vectorized unit-trace Hermitian PSD matrix of arbitrary rank, so $\vec{\mathbb{I}}^T\vec{\rho}_0 = 1$. The Liouville operators $\mathbf{L}_y \in \mathbb{C}^{n^2 \times n^2}$ (with corresponding Choi matrices $\mathbf{C}_y$) are trace-preserving (TP) i.e. $\vec{\mathbb{I}}^T\left(\sum_y \mathbf{L}_y\right) = \vec{\mathbb{I}}^T$, and completely positive (CP) i.e. $\mathbf{C}_y \succeq 0$.

The CP-TP condition on the operator $\mathbf{L}_y$ implies that we can equivalently write it via the Kraus decomposition as $\mathbf{L}_y = \left(\sum_\beta \mathbf{K}_{\beta,y} \otimes \overline{\mathbf{K}}_{\beta,y}\right)$, using Kraus operators $\mathbf{K}_{\beta,y}$ (Kraus, 1971; Adhikary

et al., 2020). Intuitively, what makes HQMMs more general than NOOMs is that its state can be a vectorized density matrix of arbitrary rank and the superoperators can have arbitrary Kraus-rank, while NOOMs require both these ranks to be 1. With this in mind, we can write and manipulate the joint probability of a sequence of N observations as:

$$P(y_1,...,y_N) = \vec{\mathbb{I}}^T \mathbf{L}_{y_1} \cdots \mathbf{L}_{y_N} \vec{\rho}_0 = \vec{\mathbb{I}}^T \left(\sum_\beta \mathbf{K}_{\beta,y_N} \otimes \overline{\mathbf{K}}_{\beta,y_N} \right) ... \left(\sum_\beta \overline{\mathbf{K}}_{\beta,Y_1} \otimes \mathbf{K}_{\beta,y_1} \right) \vec{\rho}_0 \quad (8)$$

The joint probability computation makes it clear that HQMMs are a class of PSRs, and the manipulation shows how they are equivalent to a uLPS where the left boundary vector is $\vec{\mathbb{I}}^T$. We compute the recursive state update conditioned on observation y as $\vec{\rho}_{t+1} = \frac{\mathbf{L}_y \vec{\rho}_t}{\vec{\mathbb{I}}^T \mathbf{L}_y \vec{\rho}_t}$ and the probability of an observation y is $P(y|\rho_t) = \vec{\mathbb{I}}^T \mathbf{L}_y \vec{\rho}$.

Non-terminating Uniform LPS are HQMMs Equation 8 shows that every HQMM is a uLPS, but we also consider in what sense every uLPS is an HQMM: the transfer operator of arbitrary CP maps with unit spectral radius[4] is a similarity transform away from that of a CP-TP map (Perez-Garcia et al., 2006), so $\vec{\mathbb{I}}$ is related to such a fixed point by such a similarity transform. Thus, every non-terminating uLPS has an equivalent HQMM and allows for an HQMM-style recursive state update. This is the same reasoning behind the equivalence between non-terminating uBMs (with CP maps) and NOOMs (with CP-TP maps).

Theorem 5.1. *Non-terminating uniform locally purified states are equivalent to hidden quantum Markov models.*

While it is already known that HMMs are a strict subset of HQMMs (since HQMMs also contain NOOMs which cannot always be modeled by a HMM), Adhikary et al. (2020) left open the possibility that every HQMM could have an equivalent NOOM with a possibly higher dimensional state space. In light of Theorem 4.2, we can say this is not possible as NOOMs do not capture HMMs, while HQMMs do. We defer a longer discussion of how the expressiveness of HQMMs/uLPS compares to PSR/uMPS to the appendix (see Appendix C), with the simple remark that we are not currently aware of a uMPS without an equivalent uLPS [5]. We are not aware of any proposals from the weighted automata literature that are analogous to these uLPS/HQMMs.

Corollary 5.1 (NOOM $\subset$ HQMM). *Finite dimensional norm-observable operator models are a strict subset of finite dimensional hidden quantum Markov models.*

6   Conclusion

We presented uniform matrix product states and their various subclasses, and showed how they relate to previous work in the stochastic processes and weighted automata literature. In discussing the relative expressiveness of various models, we discuss if we can find an *equivalent finite-dimensional parameterization* in another model class, but we do not discuss the relative compactness of various parameterizations. Glasser et al. (2019) do discuss this for the non-uniform case, and this could be an interesting direction to explore for the uniform case. We also speculate that the connections laid out here may allow spectral learning algorithms commonly used for PSRs and weighted automata (Hsu et al., 2012; Balle et al., 2014a; Hefny et al., 2015) suitable for learning uMPS, and an algorithm for constrained optimization on the Stiefel manifold (Adhikary et al., 2020) suitable for learning uLPS *with appropriate constraints*. Future work will involve adapting these algorithms so they can be transferred between the two fields.

We can extend our analyses to *controlled* stochastic processes. Controlled generalizations of uMPS may be developed through matrix product operators (Murg et al., 2008; Chan et al., 2016) that append an additional open index at each core of a uMPS, which we can associate with actions. We can also develop input-output versions of uniform tensor networks and uncontrolled stochastic process models, similar to input-output OOMs from Jaeger (1998). We briefly describe such extensions for HQMMs and uLPSs in Appendix D, showing that they generalize recently proposed quantum versions of partially observable Markov decisions processes (Barry et al., 2014; Ying and Ying, 2014; Cidre, 2016). With this connection, we also find that the undecidability of perfect planning (determining if there exists a policy that can deterministically reach a goal state from an arbitrary initial state in finite steps) established for quantum POMDPs by Barry et al. (2014) also extends to these generalizations. We leave a longer discussion for future work.

---

[4]This condition is necessary for probability distributions such as Equations 5 and 7 to be properly normalized.

[5]The results in Glasser et al. (2019) and De las Cuevas et al. (2013) show that MPS are more expressive than LPS in the *non-uniform* case, but their technique cannot be easily adapted to the uniform case.

References

Adhikary, S., Srinivasan, S., Gordon, G., and Boots., B. (2020). Expressiveness and learning of hidden quantum markov models. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*.

Bailly, R. (2011). Quadratic weighted automata: Spectral algorithm and likelihood maximization. In *Asian Conference on Machine Learning*, pages 147–163.

Bailly, R., Denis, F., and Ralaivola, L. (2009). Grammatical inference as a principal component analysis problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 33–40.

Balle, B., Carreras, X., Luque, F. M., and Quattoni, A. (2014a). Spectral learning of weighted automata. *Machine learning*, 96(1-2):33–63.

Balle, B., Hamilton, W., and Pineau, J. (2014b). Methods of moments for learning stochastic languages: Unified presentation and empirical comparison. In *International Conference on Machine Learning*, pages 1386–1394.

Barry, J., Barry, D. T., and Aaronson, S. (2014). Quantum partially observable markov decision processes. *Phys. Rev. A*, 90:032311.

Benvenuti, L. and Farina, L. (2004). A tutorial on the positive realization problem. *IEEE Transactions on automatic control*, 49(5):651–664.

Chan, G., Keselman, A., Nakatani, N., Li, Z., and White, S. (2016). Matrix product operators, matrix product states, and ab initio density matrix renormalization group algorithms. *The Journal of chemical physics*, 145 1:014102.

Cidre, G. A. (2016). Planning in a quantum system. *Masters Thesis, Carnegie Mellon University*.

Cirac, J. and Sierra, G. (2010). Infinite matrix product states, conformal field theory, and the haldane-shastry model. *Physical Review B*, 81:104431.

Critch, A. and Morton, J. (2014). Algebraic geometry of matrix product states. *Symmetry Integrability and Geometry-methods and Applications*, 10:095.

Critch, A., Morton, J., et al. (2014). Algebraic geometry of matrix product states. *SIGMA. Symmetry, Integrability and Geometry: Methods and Applications*, 10:095.

Cuevas, G. D. L., Cirac, J., Schuch, N., and Pérez-García, D. (2017). Irreducible forms of matrix product states: Theory and applications. *Journal of Mathematical Physics*, 58:121901.

De las Cuevas, G., Schuch, N., Pérez-García, D., and Cirac, J. I. (2013). Purifications of multipartite states: limitations and constructive methods. *New Journal of Physics*, 15(12):123021.

Denis, F. and Esposito, Y. (2004). Learning classes of probabilistic automata. In *International Conference on Computational Learning Theory*, pages 124–139. Springer.

Denis, F. and Esposito, Y. (2008). On rational stochastic languages. *Fundamenta Informaticae*, 86(1, 2):41–77.

Dupont, P., Denis, F., and Esposito, Y. (2005). Links between probabilistic automata and hidden markov models: probability distributions, learning models and induction algorithms. *Pattern recognition*, 38(9):1349–1371.

Eisert, J., Mueller, M., and Gogolin, C. (2012). Quantum measurement occurrence is undecidable. *Physical review letters*, 108 26:260501.

Evans, D. E. and Høegh-Krohn, R. (1977). Spectral properties of positive maps on c*-algebras. *Preprint series: Pure mathematics http://urn. nb. no/URN: NBN: no-8076*.

Fannes, M., Nachtergaele, B., and Werner, R. F. (1992). Finitely correlated states on quantum spin chains. *Communications in mathematical physics*, 144(3):443–490.

Garipov, T., Podoprikhin, D., Novikov, A., and Vetrov, D. (2016). Ultimate tensorization: compressing convolutional and fc layers alike. *arXiv preprint arXiv:1611.03214*.

Glasser, I., Sweke, R., Pancotti, N., Eisert, J., and Cirac, I. (2019). Expressive power of tensor-network factorizations for probabilistic modeling. In *Advances in Neural Information Processing Systems*, pages 1496–1508.

Guo, C., Jie, Z., Lu, W., and Poletti, D. (2018). Matrix product operators for sequence-to-sequence learning. *Physical Review E*, 98(4):042114.

Han, Z.-Y., Wang, J., Fan, H., Wang, L., and Zhang, P. (2018). Unsupervised generative modeling using matrix product states. *Physical Review X*, 8(3):031012.

Hefny, A., Downey, C., and Gordon, G. J. (2015). Supervised learning for dynamical system learning. In *Advances in neural information processing systems*, pages 1963–1971.

Heller, A. (1965). On stochastic processes derived from markov chains. *The Annals of Mathematical Statistics*, 36(4):1286–1291.

Helton, J. W. and Nie, J. (2009). Sufficient and necessary conditions for semidefinite representability of convex hulls and sets. *SIAM Journal on Optimization*, 20(2):759–791.

Hsu, D., Kakade, S. M., and Zhang, T. (2012). A spectral algorithm for learning hidden markov models. *Journal of Computer and System Sciences*, 78(5):1460–1480.

Huggins, W., Patil, P., Mitchell, B., Whaley, K. B., and Stoudenmire, E. M. (2019). Towards quantum machine learning with tensor networks. *Quantum Science and technology*, 4(2):024001.

Ito, H., Amari, S.-I., and Kobayashi, K. (1992). Identifiability of hidden markov information sources and their minimum degrees of freedom. *IEEE transactions on information theory*, 38(2):324–333.

Jaeger, H. (1998). Discrete-time, discrete-valued observable operator models: A tutorial. Technical report, German National Research Center for Information Technology.

Jaeger, H. (2000). Observable operator models for discrete stochastic time series. *Neural Computation*, 12(6):1371–1398.

Kliesch, M., Gross, D., and Eisert, J. (2014). Matrix-product operators and states: Np-hardness and undecidability. *Physical review letters*, 113(16):160503.

Klumper, A., Schadschneider, A., and Zittartz, J. (1993). Matrix product ground states for one-dimensional spin-1 quantum antiferromagnets. *EPL*, 24:293–297.

Kraus, K. (1971). General state changes in quantum theory. *Annals of Physics*, 64(2):311–335.

Miller, J., Rabusseau, G., and Terilla, J. (2020). Tensor networks for probabilistic sequence modeling. *arXiv preprint arXiv:2003.01039*.

Monras, A., Beige, A., and Wiesner, K. (2010). Hidden quantum markov models and non-adaptive read-out of many-body states. *arXiv preprint arXiv:1002.2337*.

Monras, A. and Winter, A. (2016). Quantum learning of classical stochastic processes: The completely positive realization problem. *Journal of Mathematical Physics*, 57(1):015219.

Murg, V., Cirac, J., Pirvu, B., and Verstraete, F. (2008). Matrix product operator representations. *New Journal of Physics*, 12:025012.

Nielsen, M. A. and Chuang, I. L. (2011). *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, USA, 10th edition.

Novikov, A., Podoprikhin, D., Osokin, A., and Vetrov, D. P. (2015). Tensorizing neural networks. In *Advances in neural information processing systems*, pages 442–450.

Novikov, A., Rodomanov, A., Osokin, A., and Vetrov, D. (2014). Putting mrfs on a tensor train. In *International Conference on Machine Learning*, pages 811–819.

Orús, R. (2014). A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349:117–158.

Oseledets, I. (2011). Tensor-train decomposition. *SIAM J. Sci. Comput.*, 33:2295–2317.

Perez-Garcia, D., Verstraete, F., Wolf, M. M., and Cirac, J. I. (2006). Matrix product state representations. *arXiv preprint quant-ph/0608197*.

Pineau, J., Gordon, G., and Thrun, S. (2003). Point-based value iteration: An anytime algorithm for pomdps. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, IJCAI'03, page 1025–1030, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Rabiner, L. and Juang, B. (1986). An introduction to hidden markov models. *IEEE ASSP Magazine*, 3:4–16.

Scheiderer, C. (2018). Spectrahedral shadows. *SIAM Journal on Applied Algebra and Geometry*, 2(1):26–44.

Siddiqi, S., Boots, B., and Gordon, G. (2010). Reduced-rank hidden markov models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 741–748.

Singh, S., James, M. R., and Rudary, M. R. (2004). Predictive state representations: A new theory for modeling dynamical systems. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, UAI '04, pages 512–519, Arlington, Virginia, United States. AUAI Press.

Srinivasan, S., Downey, C., and Boots, B. (2018a). Learning and inference in hilbert space with quantum graphical models. In *Advances in Neural Information Processing Systems*, pages 10338–10347.

Srinivasan, S., Gordon, G., and Boots, B. (2018b). Learning hidden quantum markov models. In *International Conference on Artificial Intelligence and Statistics*, pages 1979–1987.

Stinespring, W. (1955). Positive functions on *-algebras. *Proceedings of the American Mathematical Society*, 6.

Stokes, J. and Terilla, J. (2019). Probabilistic modeling with matrix product states. *Entropy*, 21(12):1236.

Stoudenmire, E. and Schwab, D. J. (2016). Supervised learning with tensor networks. In *Advances in Neural Information Processing Systems*, pages 4799–4807.

Thon, M. and Jaeger, H. (2015). Links between multiplicity automata, observable operator models and predictive state representations | a unified learning framework. *Journal of Machine Learning Research*, 16:103–147.

Vanderstraeten, L., Haegeman, J., and Verstraete, F. (2019). Tangent-space methods for uniform matrix product states. *SciPost Physics Lecture Notes*.

Vidyasagar, M. (2011). The complete realization problem for hidden markov models: a survey and some new results. *Mathematics of Control, Signals, and Systems*, 23(1-3):1–65.

Wiewiora, E. W. (2008). *Modeling probability distributions with predictive state representations*. PhD thesis, UC San Diego.

Wolf, M. M. (2012). Quantum channels & operations: Guided tour. *Lecture notes available at http://www-m5. ma. tum. de/foswiki/pub M*, 5.

Ying, S. and Ying, M. (2014). Reachability analysis of quantum markov decision processes. *Inf. Comput.*, 263:31–51.

Yu, R., Zheng, S., Anandkumar, A., and Yue, Y. (2017). Long-term forecasting using higher order tensor rnns. *arXiv: Learning*.

Zhao, M.-J. and Jaeger, H. (2010). Norm Observable Operator Models. *Neural Computation*, 22(7):1927–1959.

Notation   We use bold-face for matrix and tensor operators (e.g. $\mathbf{A}$), arrows over symbols to denote vectors (e.g. $\vec{x}$), and plain non-bold symbols for scalars. The vector-arrows are also used to indicate vectorization (column-first convention) of matrices. We frequently make use of the ones matrix $\mathbf{1}$ (filled with 1s) and the identity matrix $\mathbb{I}$, as well as their vectorizations $\vec{\mathbf{1}}$ and $\vec{\mathbb{I}}$. We use overhead bars to denote complex conjugates (e.g. $\bar{\mathbf{A}}$) and $\dagger$ for the conjugate transpose ($\bar{\mathbf{A}}^T = \mathbf{A}^\dagger$). Finally, $\mathrm{tr}(\cdot)$ denotes the trace operation applied to matrices, and $\otimes$ denotes the Kronecker product.

A   The Negative Probability Problem and the Expressiveness of Finite PSRs

Here, we continue the discussion from the end of Section 3.

Negative Probability Problem   PSRs lack a constructive definition, and in general it is undecidable if a given candidate set of parameters is a valid PSR that will not assign a negative probability to any sequence; this is the negative probability problem. A similar issue arises in the many-body physics setting, where the analogous question of whether a general matrix product operator describes a non-negative quantum density operator is also undecidable (Kliesch et al., 2014). In the special case where all PSR parameters are non-negative, we have a sufficient condition for generating valid probabilities, namely that the PSR is a hidden Markov model. Otherwise, the best available approach to characterize valid states for PSRs is whether they define a pointed convex cone (that includes the initial state) which is closed under its operators, and all points in it generate valid probabilities (Heller, 1965; Jaeger, 2000; Adhikary et al., 2020).

Positive Realization Problem   Despite the undecidability of PSRs, they are more expressive than HMMs. The general question of which uncontrolled PSRs have an equivalent finite-dimensional HMM is known as the positive realization problem (Benvenuti and Farina, 2004). A common approach is to use the result that a PSR has an equivalent finite-dimensional HMM if and only if the aforementioned convex cone of valid initial states $\{\vec{x}_0\}$ for a set of given operators $\vec{\sigma}^\dagger, \{\boldsymbol{\tau}_y\}$ is k-polyhedral for some finite k (Jaeger, 2000).

There has been some work trying to investigate whether it is possible to maintain the superior expressiveness of uMPS/PSRs while avoiding the undecidability issue. Zhao and Jaeger (2010); Bailly (2011); Adhikary et al. (2020) explore this question in the machine learning context, while Glasser et al. (2019) consider this problem from the quantum tensor network perspective. We will explore these proposals shortly. When discussing the relative expressiveness of a model compared to a uMPS/PSR, if its bond dimension (i.e. state dimension) grows with sequence length, we say there is no equivalent parameterization of the uMPS/PSR distribution in this model class.

B   Proofs of Theorems

Theorem 2.1. *Non-terminating uniform matrix product states are equivalent to uncontrolled predictive state representations.*

Proof. Consider a uMPS $\{\vec{\sigma}, \{\boldsymbol{\tau}_y\}_y, \vec{\rho}_0\}$ defined over sequences of length N and define the transfer operator $\boldsymbol{\tau} := \sum_y \boldsymbol{\tau}_y$. Say we want to compute the probability of observing y_t at time t, conditioned on past observations $y_{1:t-1}$. This requires us to not only condition on the past observations, but also marginalize over all possible future sequences $y_{t:N}$ (for example, see Theorem 1 in Miller et al. (2020) for the case of uBMs). The probability is calculated as follows

$$P(y_t \mid y_{1:t}) = \frac{\vec{\sigma}^\dagger \boldsymbol{\tau}^{N-t} \boldsymbol{\tau}_{y_t} \left(\boldsymbol{\tau}_{y_{t-1}} \cdots \boldsymbol{\tau}_{y_1} \vec{\rho}_0 \right)}{\vec{\sigma}^\dagger \boldsymbol{\tau}^{N-t} \boldsymbol{\tau} \left(\boldsymbol{\tau}_{y_{t-1}} \cdots \boldsymbol{\tau}_{y_1} \vec{\rho}_0 \right)} = \frac{\vec{\sigma}_t^\dagger \boldsymbol{\tau}_{y_t} \left(\boldsymbol{\tau}_{y_{t-1}} \cdots \boldsymbol{\tau}_{y_1} \vec{\rho}_0 \right)}{\vec{\sigma}_t^\dagger \boldsymbol{\tau} \left(\boldsymbol{\tau}_{y_{t-1}} \cdots \boldsymbol{\tau}_{y_1} \vec{\rho}_0 \right)} \tag{9}$$

Here, we have defined the *effective* evaluation functional $\vec{\sigma}_t = \left(\boldsymbol{\tau}^{N-t} \right)^\dagger \vec{\sigma}$ at time t. Intuitively, this represents the evaluation functional after having marginalized over all possibilities for the remaining $N-t$ time steps. We perform this marginalization via the transfer operator $\boldsymbol{\tau}^\dagger$. As $N \to \infty$ for fixed t, the trajectory of the effective evaluation functional will be strongly determined by the spectral properties of $\boldsymbol{\tau}$. Here, we restrict ourselves to transfer operators where the magnitudes of the two largest eigenvalues are distinct. Note that this is not a strong requirement, given that matrices with degenerate spectra form a measure zero subset of general square matrices.[6]

---

[6]Uniform MPS that violate this non-degeneracy condition are associated with "(anti-) ferromagnetic order" in quantum-many body physics, and can produce *periodic* behavior in the limit of non-terminating sequences

Now if the top eigenvalue of $\boldsymbol{\tau}^\dagger$ is $\lambda_* = 1$, $\vec{\sigma}_t$ will eventually converge to $\vec{\sigma}_*$, the corresponding fixed point of $\boldsymbol{\tau}^\dagger$, owing to the second largest eigenvalue of $\boldsymbol{\tau}^\dagger$ satisfying $|\lambda_2| < 1$. The probability computation in Equation 9 then reduces to that of a PSR with evaluation functional $\vec{\sigma}_*$. Note that in a PSRs, the $\vec{\sigma}$ is chosen precisely to be the fixed point of the adjoint transfer operator via the normalization requirement $\vec{\sigma}^\dagger \boldsymbol{\tau} = \vec{\sigma}^\dagger$, forcing $\lambda_* = 1$.

If $\lambda_* \neq 1$ for a uMPS model, we can simply rescale our matrices $\boldsymbol{\tau}$ to obtain a properly normalized transfer operator, by replacing $\boldsymbol{\tau}_y$ with $\boldsymbol{\tau}_y/\lambda_*$. Making this substitution in Equation 9, we see that the numerator and denominator are re-scaled by the same constant λ_*^{t-N}, leaving the overall probability distribution unchanged. As before, this new model produces exactly the same probabilities as a PSR with the evaluation functional $\vec{\sigma}_*$, for $\vec{\sigma}_*$ the fixed point of $\boldsymbol{\tau}^\dagger$. $\qquad\square$

Theorem 4.1. *Non-terminating uniform Born machines are equivalent to norm observable operator models.*

Proof. Note that uniform Born machines and norm observable operator models differ only in their evaluation functionals, and that NOOM transfer operators are required to be trace-preserving. While uBMs can have an arbitrary Kraus-rank 1 evaluation functional, NOOMs are restricted to the higher-rank identity evaluation functional $\vec{\mathbb{I}}$. Both models have operators of the form $\boldsymbol{\tau}_y = \boldsymbol{\phi}_y \otimes \boldsymbol{\phi}_y$, which are completely positive with Kraus-rank 1. Therefore, we need to show that an arbitrary uBM is equivalent to one where the transfer operator $\boldsymbol{\tau} = \sum_y \boldsymbol{\tau}_y$ is trace-preserving, which is the same as its adjoint $\boldsymbol{\tau}^\dagger$ being unital, i.e. having the identity $\vec{\mathbb{I}}$ as a fixed point (Nielsen and Chuang, 2011). With this fact demonstrated, the same argument used in the proof of Theorem 2.1 to prove convergence of the effective functional of a uMPS to the fixed point $\vec{\sigma}_*^\dagger$ can be applied here. This has the effect of replacing the original Kraus-rank 1 functional of the uBM by the identity $\vec{\mathbb{I}}$ in the non-terminating limit, completing the conversion from uBM to NOOM.

If the uBM transfer operator $\boldsymbol{\tau}$ is already trace-preserving then we are done, so assume it is not. We make the generic assumption[7] that the two eigenvalues of $\boldsymbol{\tau}$ with greatest magnitude, λ_* and λ_2, satisfy $|\lambda_*| > |\lambda_2|$. By replacing all operators $\boldsymbol{\phi}_y$ by $\boldsymbol{\phi}_y/\sqrt{\lambda_*}$, we convert the uBM into one whose transfer operator $\boldsymbol{\tau}$ has leading eigenvalue of magnitude 1, a re-scaling which leaves the joint probability distributions unchanged. In this case, the quantum Perron-Frobenius theorem (Evans and Høegh-Krohn, 1977) then ensures that $\lambda_* = 1$, and that $\tau^\dagger$ has a unique fixed-point operator $\vec{\sigma}_*$ which is the vectorization of a full-rank (and consequently, invertible) positive definite matrix σ_*.

A similarity transformation of $S = \sigma_*^{1/2}$ can then be applied to the uBM matrices, replacing $\boldsymbol{\phi}_y$ with $\boldsymbol{\phi}_y' = S\boldsymbol{\phi}_y S^{-1}$. This similarity transformation ensures the new transfer operator $\boldsymbol{\tau}'$ is trace-preserving, as demonstrated by the following:

$$
\begin{aligned}
\boldsymbol{\tau}'^\dagger \vec{\mathbb{I}} &= \left(\sum_y \overline{\boldsymbol{\phi}_y'}^\dagger \otimes \boldsymbol{\phi}_y'^\dagger \right) \vec{\mathbb{I}} = \left(\sum_y (\sigma_*^{-1/2})^T \overline{\boldsymbol{\phi}_y}^\dagger (\sigma_*^{1/2})^T \otimes \sigma_*^{-1/2} \boldsymbol{\phi}_y^\dagger \sigma_*^{1/2} \right) \vec{\mathbb{I}} \\
&= \left((\sigma_*^{-1/2})^T \otimes \sigma_*^{-1/2} \right) \left(\sum_y \overline{\boldsymbol{\phi}_y}^\dagger \otimes \boldsymbol{\phi}_y^\dagger \right) \left((\sigma_*^{1/2})^T \otimes \sigma_*^{1/2} \right) \vec{\mathbb{I}} \\
&= \left((\sigma_*^{-1/2})^T \otimes \sigma_*^{-1/2} \right) \boldsymbol{\tau}^\dagger \vec{\sigma}_* \\
&= \left((\sigma_*^{-1/2})^T \otimes \sigma_*^{-1/2} \right) \vec{\sigma}_* \\
&= \vec{\mathbb{I}}
\end{aligned}
\tag{10}
$$

(Cuevas et al., 2017). Although we don't give the full details here, such degenerate uMPS can still be converted to equivalent PSR by redefining the observation space to consist of k-tuples of adjacent observations, for k the periodicity of the uMPS.

[7]This is similar to the assumption made in proving Theorem 2.1, and is valid everywhere outside of a measure-zero subset of uBMs. In the case of degenerate eigenvalues, the conversion from uBMs to NOOMs can still be achieved given a slight redefinition of the observation space to combine together k adjacent observations.

In the equations above, we have used the facts that (a) σ_*, $\sigma_*^{1/2}$, and $\sigma_*^{-1/2}$ are Hermitian, so that complex conjugation acts as $\overline{\sigma_*^{1/2}} = (\sigma_*^{1/2})^T$, (b) $(Z^T \otimes X)\vec{Y} = \vec{W}$ with $W = XYZ$, for any matrices X, Y, Z, and (c) $\boldsymbol{\tau}^\dagger \vec{\sigma_*} = \vec{\sigma_*}$

We have demonstrated that the similarity-transformed uBM now possesses a trace-preserving transfer operator, which by the arguments above ensures it is a valid NOOM in the non-terminating sequence limit. □

Theorem 4.2. *[HMM $\not\subseteq$ NOOM] There exist finite-dimensional hidden Markov models that have no equivalent finite-dimensional norm-observable operator model.*

We first introduce a lemma (Ito et al., 1992; Vidyasagar, 2011; Thon and Jaeger, 2015) that will help us in our proof. It tells us that two equivalent PSRs of the same dimension are simply a similarity transform away from one another.

Lemma 1 (Thon and Jaeger (2015))**.** *Suppose $(\mathbb{C}^n, \vec{\sigma}, \{\boldsymbol{\tau}_y\}_{y \in \mathcal{O}}, \vec{x}_0)$ and $(\mathbb{C}^n, \vec{\sigma}', \{\boldsymbol{\tau}'_y\}_{y \in \mathcal{O}}, \vec{x}'_0)$ are two equivalent PSR representations, i.e., they generate the same sequence of probabilities. Then, there exists some non-singular $\mathbf{S} \in \mathbb{C}^{n \times n}$ such that $\vec{x}'_0 = \mathbf{S}^{-1}\vec{x}_0$, $\boldsymbol{\tau}'_y = \mathbf{S}^{-1}\boldsymbol{\tau}_y \mathbf{S}$, and $\vec{\sigma}'^\dagger = \vec{\sigma}^\dagger \mathbf{S}$.*

Proof. We construct a class of HMMs for which there are no equivalent NOOMs and give a proof by contradiction. Let $\mathcal{A} = (\mathbb{R}^p, \vec{\sigma}, \{\boldsymbol{\tau}_y\}_{y \in \mathcal{O}}, \vec{x}_0)$ be a minimal PSR equivalent to some hidden Markov model $\mathcal{M} = (\mathbb{R}^m, \mathbf{A}, \mathbf{C}, \vec{p}_0)$ for which some future state reachable from the initial state can be written as a convex combination of some previously reached states, i.e., $\vec{x}_k = \alpha \vec{x}_i + \beta \vec{x}_j$ for some $\alpha, \beta > 0$, $\alpha + \beta = 1$ and some $i, j, k \in \mathbb{N}$ with $i < j < k$ and $\vec{x}_k = \frac{\boldsymbol{\tau}_{y_k} \cdots \boldsymbol{\tau}_{y_1} \vec{x}_0}{\vec{\sigma}^T \boldsymbol{\tau}_{y_k} \cdots \boldsymbol{\tau}_{y_1} \vec{x}_0}$ for some sequence of observations $Y = y_1, ..., y_k$ (and similarly for $\vec{x}_j$ and $\vec{x}_i$ which truncate the observations at y_j and y_i respectively).

Suppose there exists an equivalent NOOM (represented in its vectorized form) $\mathcal{M}' = (\mathbb{R}^n, \vec{\mathbb{I}}, \{\boldsymbol{\phi}_y\}_{y \in \mathcal{O}}, \vec{\psi}_0)$. Let $\mathcal{A}' = (\mathbb{R}^p, \vec{\sigma}', \{\boldsymbol{\tau}'_y\}_{y \in \mathcal{O}}, \vec{x}'_0)$ be a minimal PSR computing the same distribution as $\mathcal{M}'$.

Then, by Lemma 1, we have some similarity transform $\mathbf{S}$ such that $\vec{\sigma}'^T = \vec{\sigma}^T \mathbf{S}$, $\boldsymbol{\tau}'_y = \mathbf{S}^{-1}\boldsymbol{\tau}_y \mathbf{S}$, and $\vec{x}'_0 = \mathbf{S}^{-1}\vec{x}_0$. Thon and Jaeger (2015) show that there are matrices Φ, Π that relate the NOOM to its minimal representation as $\mathcal{A}' = \Pi\Phi^+ \mathcal{M}'\Phi\Pi^+$ (where $+$ represents the Moore-Penrose pseudoinverse). This allows us to relate the NOOM with the minimal PSR representation of its equivalent HMM as $\vec{\mathbb{I}}^T \Phi\Pi^+ \mathbf{S}^{-1} = \vec{\sigma}^T$, $\boldsymbol{\tau}_y = \mathbf{S}\Pi\Phi^+ \boldsymbol{\phi}_y \Phi\Pi^+ \mathbf{S}^{-1}$, and $\vec{x}_0 = \mathbf{S}\Pi\Phi^+ \vec{\psi}_0$.

Now, by the NOOM evolution rules, the NOOM state at timestep k for the sequence Y is $\vec{\psi}_k = \frac{\boldsymbol{\phi}_{y_k} \cdots \boldsymbol{\phi}_{y_1} \vec{\psi}_0}{\vec{\mathbb{I}}^T \boldsymbol{\phi}_{y_k} \cdots \boldsymbol{\phi}_{y_1} \vec{\psi}_0}$ and the probability of any given observation at that time-step is $P(y | \vec{\psi}_k) = \vec{\mathbb{I}}^T \boldsymbol{\phi}_y \vec{\psi}_k$. Further, note that the NOOM state must be a vectorized rank-1 matrix whose eigenvector has unit ℓ_2 norm, i.e., $\vec{\psi}_k = \text{vec}(\psi_k \psi_k^T)$ with $\|\psi_k\|_2 = 1$.

But we can also write NOOM state in terms of its equivalent PSR-HMM as
$$\vec{\psi}_k = \Phi\Pi^+ \mathbf{S}^{-1} \vec{x}_k = \Phi\Pi^+ \mathbf{S}^{-1}(\alpha \vec{x}_i + \beta \vec{x}_j) = \Phi\Pi^+ \mathbf{S}^{-1}\left(\alpha\left(\mathbf{S}\Pi\Phi^+ \vec{\psi}_i\right) + \beta\left(\mathbf{S}\Pi\Phi^+ \vec{\psi}_j\right)\right) = \alpha\vec{\psi}_i + \beta\vec{\psi}_j$$

If $\vec{\psi}_i$ and $\vec{\psi}_j$ are valid NOOM states, we have that $\vec{\psi}_k = \alpha\vec{\psi}_i + \beta\vec{\psi}_j = \alpha\text{vec}(\psi_i \psi_i^T) + \beta\text{vec}(\psi_j \psi_j^T) = \text{vec}(\alpha\psi_i \psi_i^T + \beta\psi_j \psi_j^T)$. However, $\alpha\vec{\psi}_i + \beta\vec{\psi}_j$ is not the vectorization of a rank-1 matrix in general. In particular, $\vec{\psi}_k$ has unit Kraus-rank only if ψ_i and ψ_j are linearly dependent, which is true only if $\vec{x}_i$ and $\vec{x}_j$ are linearly dependent. Thus, whenever $\vec{x}_i$ and $\vec{x}_j$ are linearly independent, $\vec{\psi}_k$ cannot have unit Kraus-rank. But as normalized HMM states, $\vec{x}_i$ and $\vec{x}_j$ are always linearly independent, unless they are exactly the same. Hence, a contradiction. Thus, for such an HMM, there is no equivalent NOOM.

□

NOOMs are a restrictive model class The proof above essentially argues that if an HMM had an equivalent NOOM, then anytime a reachable HMM state can be written as a convex combination of some other reachable states, the equivalent NOOM state should also admit a representation as a convex combination of the NOOM-equivalent reachable states, but such a representation violates the condition that NOOM states have unit Kraus-rank, and hence there cannot be an equivalent NOOM. Here, we

provide an example of an HMM that can have a state be a linear (here, convex) combination of two prior linearly independent states

$$\vec{x}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \qquad \boldsymbol{\tau}_1 = \begin{bmatrix} 0.25 & 0.5 \\ 0.75 & 0 \end{bmatrix} \qquad \boldsymbol{\tau_2} = \begin{bmatrix} 0 & 0 \\ 0 & 0.5 \end{bmatrix} \qquad (11)$$

Then, for the sequence $Y = (1,1)$:

$$\vec{x}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \qquad \vec{x}_1 = \begin{bmatrix} 0.25 \\ 0.75 \end{bmatrix} \qquad \vec{x}_2 = \begin{bmatrix} 0.7 \\ 0.3 \end{bmatrix} = 0.6\vec{x}_0 + 0.4\vec{x}_1 \qquad (12)$$

Since $\vec{x}_0$ and $\vec{x}_1$ are linearly independent, we know that such an HMM cannot have an equivalent NOOM because of NOOM's state rank constraints. In this case, we see that if an HMM reaches a state that lies strictly inside the convex hull of other reachable states, it rules out a NOOM representation. This constitutes a fairly expansive class of HMMs, suggesting that NOOMs cannot model a wide variety of HMMs, making NOOMs a restrictive model class. This core insight holds for any quantum model that updates and maintains only pure quantum states. Since HQMMs allow states to be vectorizations of arbitrary rank matrices (and so convex combinations of other reachable states), we do not run into the same issues.

Corollary 4.1 (uBM $\not\subseteq$ HMM and HMM $\not\subseteq$ uBM). *There exist finite-dimensional uniform Born machines that have no equivalent finite-dimensional hidden Markov models, and vice-versa*

Proof. For the first relationship uBM $\not\subseteq$ HMM, note that we already know that NOOM $\not\subseteq$ HMM. Zhao and Jaeger (2010) demonstrate this by designing the NOOM probability clock model, where the latent state (and therefore conditional probabilities) exhibit oscillatory behavior. The negative entries in NOOM operators allow their top eigenvalues to be complex valued, which allow for such oscillatory behavior. On the other hand, HMMs with non-negative operators with real eigenvalues cannot produce such oscillations. Geometrically, valid finite-dimensional HMM states are restricted to form polyhedral cones, while the oscillations in the probability clock require state dynamics on non-polyhedral (or infinitely generated polyhedral) cones. From Theorem 4.1, we know that a non-terminating uBM with the same operators as the probability clock NOOM will generate identical probabilities, and thus produce the same dynamics as the probability clock NOOM. These are then instances of uBMs that cannot be modeled by a finite dimensional HMM.

For the second relationship HMM $\not\subseteq$ uBM, note that Theorem 4.2 tells us that there exist HMMs that cannot be modeled by finite dimensional NOOMs. The same HMMs cannot be modeled by non-terminating finite dimensional uBMs, as these are equivalent to NOOMs. □

Theorem 5.1. *Non-terminating uniform locally purified states are equivalent to hidden quantum Markov models.*

We follow the same approach as in Theorem 4.1. Uniform LPS models and HQMMs differ only in their evaluation functionals, and that HQMMs operators are trace-preserving. While uLPSs can have an arbitrary evaluation functional, HQMMs are restricted to the identity evaluation functional $\vec{\mathbb{I}}$. Both models have operators of the form $\boldsymbol{\tau}_y = \sum_y \overline{\mathbf{K}_y} \otimes \mathbf{K}_y$, which are completely positive operators.

As discussed in Theorem 4.1, the transfer operator $\boldsymbol{\tau}$ can be rescaled and similarity transformed into one that is trace-preserving. In the limit of non-terminating sequences, the evaluation functional of this transformed model will then converge to $\vec{\mathbb{I}}$, the fixed point of $\boldsymbol{\tau}^\dagger$. Therefore, this similarity transform allows us to map a non-terminating uLPS to an HQMM, proving that non-terminating uLPSs are equivalent to HQMMs.

Corollary 5.1 (NOOM $\subset$ HQMM). *Finite dimensional norm-observable operator models are a strict subset of finite dimensional hidden quantum Markov models.*

We know from Theorem 4.2 that there exist HMMs that cannot be modeled by finite-dimensional NOOMs. However, since all HMMs can be modeled by finite dimensional HQMMs, the same HMMs serve as instances of HQMMs that cannot be modeled by NOOMs. This, combined with the fact that NOOM $\subseteq$ HQMM (Adhikary et al., 2020), give us that NOOM $\subset$ HQMM.

C Expresiveness of HQMMs (uLPS)

With Corollary 5.1, we see that HQMMs are a particularly expressive *constructive* class of PSRs that avoid the NPP by design. What makes HQMMs more expressive than HMMs is that they still admit infinitely generated (i.e., non-polyhedral) cones of valid states (in reference to the discussion in Appendix

A). Geometrically, we can think of the valid unit-Kraus rank (or equivalently pure density matrix) initial NOOM states as the extremal points of a *spectraplex*, which is the intersection of the affine space of unit trace matrices with the convex cone of PSD matrices. The arbitrary Schmidt rank (or mixed density matrices) admitted as initial states for HQMMs fills the entire spectraplex (Adhikary et al., 2020).

The natural question to ask is whether HQMMs can model *any* finite dimensional PSRs, or if we lose any expressiveness in using HQMMs or LPSs? Glasser et al. (2019) provide results for the non-uniform case: there exist finite-dimensional non-uniform MPS that have no equivalent finite-dimensional LPS. A similar answer for the *uniform* case is as yet unknown. Investigating this from the HQMM perspective, Monras and Winter (2016) term the question of which PSRs have equivalent HQMMs as the *completely positive realization problem*. They argue that a necessary and sufficient condition for a PSR to have an equivalent HQMM is if the operators characterizing the PSR come from a semi-definite representable (SDR) cone that also defines the convex cone of valid states.[8] They suggest that if we could show that every valid cone for a PSR satisfied the SDR condition, we could show that HQMMs are equivalent to PSRs via the Helton-Nie conjecture; while it is the case that every SDR set is convex and semi-algebraic, the converse, known as the Helton-Nie conjecture was only recently shown to be false (Scheiderer, 2018).

Since convex semi-algebraic sets are SDR sets in a wide variety of cases, it may still be the case that HQMMs are equivalent to PSRs. We thus pose two open questions that must be answered to obtain a full characterization of HQMMs relative to PSRs: first, whether the convex cone characterizing the PSR is semi-algebraic, and second, if it satisfies any sufficient conditions for being SDR (Helton and Nie, 2009), and if not, how common such conditions are. Nevertheless, convex and semi-algebraic cones that are SDR are a broad class and the results thus far show that HQMMs are the most expressive known subset of PSRs. Indeeed, we do not know any examples of PSRs that do not have an equivalent finite-dimensional HQMM.

D Controlled Stochastic Processes

In this paper, we have discussed various models of stochastic processes, including counter parts in quantum tensor networks and weighted automata. These models are limited in the sense that have no notion of control – an agent can make observations of a system but cannot affect it. We now extend our analysis to *controlled* stochastic processes; connections between these models are illustrated in Figure 2.

D.1 Partially Observable Markov Decision Processes

Partially observable Markov decision processes (POMDPs) can be viewed as natural extensions of hidden Markov models to controlled stochastic processes. Formally, we give the following definition[9]:

Definition 5 (Partially Observable Markov Decision Processes)**.** An n-dimensional partially observable Markov decision process for sets of discrete actions $\mathcal{A}$ and observations $\mathcal{O}$ is a tuple $(\mathbb{R}^n, \{\mathbf{A}^a\}_{a \in \mathcal{A}}, \{\mathbf{C}^a\}_{a \in \mathcal{A}}, \vec{x}_0)$, where the matrices $\mathbf{A}^a \in \mathbb{R}^{n \times n}_{\geq 0}$ and $\mathbf{C}^a \in \mathbb{R}^{|\mathcal{O}| \times n}_{\geq 0}$ are non-negative column-stochastic transition and emission matrices corresponding to the action a such that $\vec{\mathbf{1}}^T \mathbf{A}^a = \vec{\mathbf{1}}^T \mathbf{C}^a = \vec{\mathbf{1}}^T$. The initial state $\vec{x}_0 \in \mathbb{R}^n_{\geq 0}$ is also non-negative and satisfies $|\vec{x}_0|_1 = \vec{\mathbf{1}}^T \vec{x}_0 = 1$.

As with HMMs, we can also define POMDPs via observable operators $\mathbf{T}^a_y = \text{diag}(\mathbf{C}^a_{(y,:)}) \mathbf{A}^a$ corresponding to taking action a and observing y. The state update is done as $\vec{x}_t = \frac{\mathbf{T}^a_y \vec{x}_{t-1}}{\vec{\mathbf{1}}^T \mathbf{T}^a_y \vec{x}_{t-1}}$, and probabilities are calculated as $P(a_1 y_1, ..., a_k y_k) = \vec{\mathbf{1}}^T \mathbf{T}^{a_k}_{y_k} ... \mathbf{T}^{a_1}_{y_1} \vec{x}$; exactly the same as HMMs, but with controls.

D.2 Input-Output Extensions of Stochastic Processes

In Section 2.1, we considered *uncontrolled* PSRs, which are equivalent to the observable operator models (OOMs) defined in Jaeger (2000) (Thon and Jaeger, 2015). Augmenting these models with a notion of control, we obtain input-output OOMs (IO-OOMs) (Jaeger, 1998) or controlled PSRs (Singh et al., 2004), which are essentially comprised of a separate OOM for every action. Following the same motivation, we can also analogously define *input-output* HQMMs as follows

---

[8]The intersection of the cone of positive semi-definite matrices with an affine subspace is called a spectrahedron, and linear maps of spectrahedra are called *spectrahedral shadows* or semi-definitely representable sets. These are the feasible regions of a semidefinite program.

[9]We have ignored the reward function commonly associated with controlled stochastic processes to draw a clearer comparison with uncontrolled processes. Our analysis can be equivalently applied with reward functions.

Definition 6 (Input-Output HQMM). An n^2-dimensional input-output hidden quantum Markov model for sets of discrete actions $\mathcal{A}$ and observations $\mathcal{O}$ is a controlled stochastic process given by the tuple $(\mathbb{C}^{n^2}, \vec{\mathbb{I}}, \{\mathbf{L}_y^a\}_{y\in\mathcal{O}, a\in\mathcal{A}}, \vec{\rho}_0)$. The set of operators $\{\mathbf{L}_y^a\}_{y\in\mathcal{O}}$ for every action a form a set of CP-TP Liouville operators as in Definition 4. The initial state $\vec{\rho}_0$ is a vectorized unit-trace Hermitian PSD matrix of arbitrary rank.

The state update and probability computations for IO-HQMMs are done exactly the same as with HQMMs, but now with operators indexed by both observations and actions at each time step:

$$P(a_1 y_1 ... a_k y_k) = \vec{\mathbb{I}}^T \mathbf{L}_{y_k}^{a_k} ... \mathbf{L}_{y_1}^{a_1} \vec{\rho}_0 \qquad\qquad \vec{\rho}' = \frac{\mathbf{L}_{y_k}^{a_k} ... \mathbf{L}_{y_1}^{a_1} \vec{\rho}_0}{\vec{\mathbb{I}}^T \mathbf{L}_{y_k}^{a_k} ... \mathbf{L}_{y_1}^{a_1} \vec{\rho}_0}$$

We can also define input-output uLPS and input-output uMPS in exactly the same way, with non-terminating variants being equivalent to IO-HQMMs and IO-OOMs/PSRs respectively. When viewed as tensor network in Figure 2, we immediately see that such IO-HQMMs and uniform IO-LPSs have the same structure as matrix product operators (Chan et al., 2016; Murg et al., 2008) which can be viewed as MPSs with an additional open index at each core.

Quantum Observable Markov Decision Processes are Input-Ouput HQMMs   Barry et al. (2014) developed quantum observable Markov decision processes (QOMDPs) as a strict generalization of classical POMDPs, by swapping the belief state vector $\vec{x}$ with the density matrix of a quantum state ρ. Ignoring reward functions, QOMDPs can be defined as follows

Definition 7. An n^2-dimensional quantum observable Markov decision process for sets of discrete actions $\mathcal{A}$ and observations $\mathcal{O}$ is a tuple $(\mathbb{C}^{n^2}, \vec{\mathbb{I}}, \{\mathbf{K}_y^a\}_{a\in\mathcal{A}, y\in\mathcal{O}}, \rho_0)$, where the set of operators $\{\mathbf{K}_y^a \in \mathbb{C}^{n\times n}\}_{y\in\mathcal{O}}$ for each action a forms a quantum channel with $\sum_{y\in\mathcal{O}} \mathbf{K}_y^{a\dagger} \mathbf{K}_y^a = \mathbb{I}$. The initial state ρ_0 is a unit-trace Hermitian PSD density matrix of arbitrary rank.

The probability of an action-observation sequence $a_1 y_1 \cdots a_T y_T$ for a QOMDP is then taken to be $\mathrm{Tr}\left(\mathbf{K}_{y_T}^{a_T} ... \mathbf{K}_{y_1}^{a_1} \rho_0 \mathbf{K}_{y_1}^{a_1\dagger} ... \mathbf{K}_{y_T}^{a_T\dagger}\right)$. To draw connections with HQMMs, we represent the vectorized version of the QOMDP update equations as follows:

$$\vec{\rho}' = \mathrm{vec}\left(\frac{\mathbf{K}_{y_T}^{a_T} ... \mathbf{K}_{y_1}^{a_1} \rho_0 \mathbf{K}_{y_1}^{a_1\dagger} ... \mathbf{K}_{y_T}^{a_T\dagger}}{\mathrm{Tr}\left(\mathbf{K}_{y_T}^{a_T} ... \mathbf{K}_{y_1}^{a_1} \rho_0 \mathbf{K}_{y_1}^{a_1\dagger} ... \mathbf{K}_{y_T}^{a_T\dagger}\right)}\right) = \frac{(\overline{\mathbf{K}}_{y_T}^{a_T} \otimes \mathbf{K}_{y_T}^{a_T}) ... (\overline{\mathbf{K}}_{y_1}^{a_1} \otimes \mathbf{K}_{y_1}^{a_1}) \vec{\rho}_0}{\vec{\mathbb{I}}(\overline{\mathbf{K}}_{y_T}^{a_T} \otimes \mathbf{K}_{y_T}^{a_T}) ... (\overline{\mathbf{K}}_{y_1}^{a_1} \otimes \mathbf{K}_{y_1}^{a_1}) \vec{\rho}_0} = \frac{\mathbf{L}_{y_T}^{a_T} ... \mathbf{L}_{y_1}^{a_1} \vec{\rho}_0}{\vec{\mathbb{I}}^T \mathbf{L}_{y_T}^{a_T} ... \mathbf{L}_{y_1}^{a_1} \vec{\rho}_0}$$

$$(13)$$

The above update equation is identical to the update equation for IO-HQMMs; the only difference being that the QOMDP operators are restricted to have unit Kraus-rank: $\mathbf{L}_y^a = \overline{\mathbf{K}}_y^a \otimes \mathbf{K}_y^a$. This makes them somewhat alike NOOMs; however, unlike NOOMs, QOMDPs allow latent states of arbitrary Schmidt-ranks – i.e. both mixed and pure states. We thus arrive at the following theorem characterizing the expressiveness of QOMDPs

Theorem D.1. *QOMDPs $\subseteq$ IO-HQMMS $=$ non-terminating uniform IO-LPS $\subseteq$ IO-OOMs $=$ PSRs*

Quantum Markov decision processes were also studied by Ying and Ying (2014), who arrived at a formulation similar to QOMDPs. Cidre (2016) proposed an alternate formulation, called Quantum MDPs (QuaMDPs), along with an associated point-based value iteration algorithm similar to that used to learn classical POMDPs (Pineau et al., 2003). Although connections to QOMDPs were not considered in the original work, QuaMDPs are a special case of QOMDPs. The filtering process in IO-HQMMs and QOMDPs corresponds to positive operator valued measurements, which are generalizations of the more restrictive projection valued measurements used in QuaMDPs. One can reduce positive operator valued measurements to projection based measurements via the Stinespring dilation theorem (Stinespring, 1955), but this requires the system interacting with an ancillary sub-system; this is not the case in QuaMDPs.

Undecidability of Perfect Planning   In moving from POMDPs to QOMDPs, Barry et al. (2014) find an apparent classical-quantum separation in the problem of perfect goal state reachability. They consider particular instances of controlled processes where certain states are labeled as goals, and are set to be absorbing – the transition probabilities from these states to any other state is zero; we will refer to these as *goal oriented* models.

In such models, the perfect goal state reachability problem can be stated as follows: given an arbitrary initial state and a goal state, is there a sequence of valid operations that will leave the system in a goal
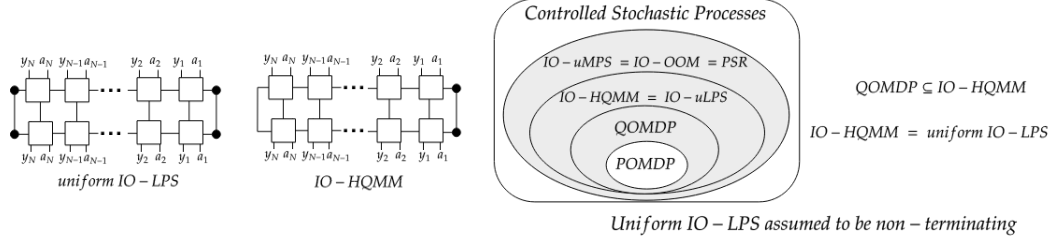
Figure 2: (Left) Tensor network diagrams (Right) Subset relationships between models, and new relationships established in this paper. The grey sections represent potentially non-strict subsets and are potentially empty.

state with probability 1 in a finite number of steps. In other words: is there a policy that will take the agent deterministically from some initial state to a goal state in a finte number of steps? Barry et al. (2014) show that this problem is undecidable for QOMDPs, even though it is decidable for POMDPs. Formally, the undecidability of perfect goal state reachability in QOMDPs is a consequence of the undecidability of the quantum measurement occurence problem (Eisert et al., 2012). Intuitively, Barry et al. (2014) point out that the decidability of the same problem for POMDPs boils down to its non-negativity constraints. However, negative parameters are also present in *classical* generalizations of POMDPs, namely IO-OOMs or PSRs, as well as other quantum models including IO-HQMMs or IO-LPS. Thus, we would expect this problem should be undecidable for these models as well. Indeed, we can show this to be the case using the subset relationships we have established between these models and QOMDPs.

Theorem D.2 (Perfect Goal State Reachability). *Given an initial state and a goal state, it is undecidable whether there exists a policy that will leave a goal-oriented IO-HQMMs, (non-terminating uniform) IO-LPS, IO-OOM or PSR in the goal state with probability* 1 *in a finite number of steps.*

Proof. QOMDPs are contained within the class of IO-HQMMs or IO-uLPS, which are, in turn, contained within the class of IO-OOMs or PSRs. If a problem is undecidable for a model class, it must be undecidable in general for any other class that contains it as a subset. Therefore, we immediately see that the undecidability carries over to these models as well. □