1 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 Miller et al. (2021) for the case of uBMs). The probability is calculated as follows

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

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.[1]

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 1 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 1, we see that the numerator and denomenator are rescaled 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 $\boldsymbol{\tau}^\dagger$. □

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[2] that the two eigenvalues of $\boldsymbol{\tau}$ with greatest magnitude, λ_* and λ_2, satisfy $|\lambda_*| > |\lambda_2|$.

[1] uMPS which 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 (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. This procedure is also called 'blocking' in the condensed-matter literature (Cirac et al., 2017).

[2] 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.

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 rescaling 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 $\boldsymbol{\tau}^\dagger$ has a unique fixed-point operator $\vec{\sigma}_*$ which is the vectorization of a full-rank (and consequently, invertible) positive definite matrix $\boldsymbol{\sigma}_*$.

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

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

In the equations above, we have used the facts that (a) $\boldsymbol{\sigma}_*$, $\boldsymbol{\sigma}_*^{1/2}$, and $\boldsymbol{\sigma}_*^{-1/2}$ are Hermitian, so that complex conjugation acts as $\overline{\boldsymbol{\sigma}_*^{1/2}} = (\boldsymbol{\sigma}_*^{1/2})^T$, (b) $\left(\mathbf{Z}^T \otimes \mathbf{X}\right)\vec{Y} = \vec{W}$ with $\mathbf{W} = \mathbf{XYZ}$, for any matrices $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, 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 $\vec{\sigma}'^T = \vec{\sigma}^T\mathbf{S}$.*

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 $Y = y_1, \ldots, y_k$ (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) $\mathcal{M}' = (\mathbb{R}^n, \vec{\mathbb{I}}, \{\boldsymbol{\phi}_y\}, \vec{\psi}_0)$. 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., $\vec{\psi}_k = \text{vec}(\vec{\psi}_k \vec{\psi}_k^T)$ with $\|\vec{\psi}_k\|_2 = 1$.

But we can also write NOOM state in terms of its equivalent PSR-HMM as

$$\begin{aligned}\vec{\psi}_k &= \Phi\Pi^+\mathbf{S}^{-1}\vec{x}_k \\ &= \Phi\Pi^+\mathbf{S}^{-1}(\alpha\vec{x}_i + \beta\vec{x}_j) \\ &= \Phi\Pi^+\mathbf{S}^{-1}\left(\alpha\left(\mathbf{S}\Pi\Phi^+\vec{\psi}_i\right) + \beta\left(\mathbf{S}\Pi\Phi^+\vec{\psi}_j\right)\right) \\ &= \alpha\vec{\psi}_i + \beta\vec{\psi}_j\end{aligned}$$

If $\vec{\psi}_i$ and $\vec{\psi}_j$ are valid NOOM states, we have that $\vec{\psi}_k = \alpha\vec{\psi}_i + \beta\vec{\psi}_j = \alpha\text{vec}(\vec{\psi}_i\vec{\psi}_i^T) + \beta\text{vec}(\vec{\psi}_j\vec{\psi}_j^T) = \text{vec}(\alpha\vec{\psi}_i\vec{\psi}_i^T + \beta\vec{\psi}_j\vec{\psi}_j^T)$. 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 $\vec{\psi}_i$ and $\vec{\psi}_j$ 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

$$\vec{x}_0 = \begin{bmatrix}1\\0\end{bmatrix} \quad \boldsymbol{\tau}_1 = \begin{bmatrix}0.25 & 0.5\\0.75 & 0\end{bmatrix} \quad \boldsymbol{\tau_2} = \begin{bmatrix}0 & 0\\0 & 0.5\end{bmatrix} \tag{2}$$

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

$$\vec{x}_0 = \begin{bmatrix}1\\0\end{bmatrix} \quad \vec{x}_1 = \begin{bmatrix}0.25\\0.75\end{bmatrix} \quad \vec{x}_2 = \begin{bmatrix}0.7\\0.3\end{bmatrix} = 0.6\vec{x}_0 + 0.4\vec{x}_1 \tag{3}$$

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.

2 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 Section 3.1). 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.[3] 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.

3 Controlled Stochastic Processes

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

[3] 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.

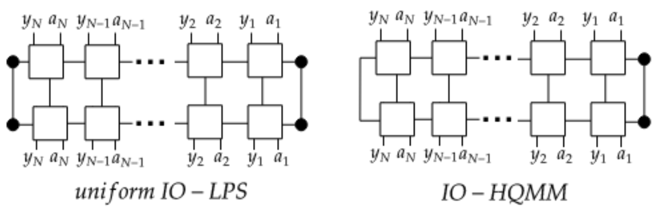

Figure 1: Tensor network diagrams for controlled processes.

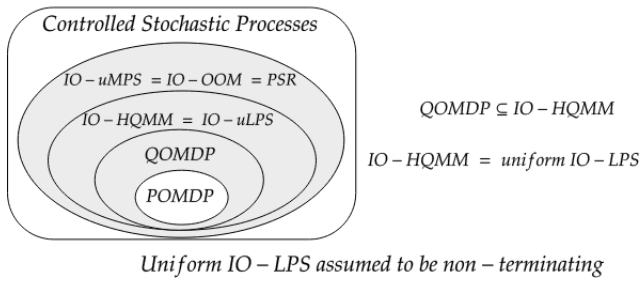

Figure 2: Subset relationships between models, and new relationships established in this paper. The grey sections represent potentially non-strict subsets.

3.1 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[4]:

Definition 1 (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 $|\vec{x}_0|_1 = 1$.

POMDPs can also be defined 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 $P(a_1y_1, \ldots, a_N, yN) = \vec{\mathbf{1}}^T\mathbf{T}^{a_k}_{y_k}\ldots\mathbf{T}^{a_1}_{y_1}\vec{x}$; exactly the same as HMMs, but with controls.

3.2 Input-Output Extensions of Stochastic Processes

In the main body of the paper, we considered *uncontrolled* PSRs, which are equivalent to the observable operator models (OOMs) defined in Jaeger (2000). 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 essentially consist of a separate OOM for every action. With the same motivation, we can also define analogous input-output HQMMs as follows

Definition 2 (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 $(\mathbb{C}^{n^2}, \vec{\mathbb{I}}, \{\mathbf{L}^a_y\}_{y\in\mathcal{O},a\in\mathcal{A}}, \vec{\rho}_0)$. The set of operators $\{\mathbf{L}^a_y\}_{y\in\mathcal{O}}$ for every action a form a set of CP-TP Liouville operators. 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: $P(a_1y_1\ldots a_ky_k) = \vec{\mathbb{I}}^T\mathbf{L}^{a_k}_{y_k}\ldots\mathbf{L}^{a_1}_{y_1}\vec{\rho}_0$ and $\vec{\rho}\,' = \frac{\mathbf{L}^{a_k}_{y_k}\ldots\mathbf{L}^{a_1}_{y_1}\vec{\rho}_0}{\vec{\mathbb{I}}^T\mathbf{L}^{a_k}_{y_k}\ldots\mathbf{L}^{a_1}_{y_1}\vec{\rho}_0}$ Similarly, we can define input-output uLPSs and input-output uMPSs

[4]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.

in exactly the same way, with non-terminating variants being equivalent to IO-HQMMs and IO-OOMs/PSRs respectively. When viewed as a tensor network in Figure 1, 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 $\boldsymbol{\rho}$. Ignoring reward functions, QOMDPs can be defined as follows

Definition 3. An n-dimensional quantum observable Markov decision process for sets of discrete actions $\mathcal{A}$ and observations $\mathcal{O}$ is a tuple $(\mathbb{C}^n, \{\mathbf{K}_y^a\}_{a\in\mathcal{A},y\in\mathcal{O}}, \boldsymbol{\rho}_0)$, 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 $\boldsymbol{\rho}_0$ is a unit-trace Hermitian PSD density matrix of arbitrary rank.

The probability of an action-observation sequence $a_1y_1\cdots a_Ty_T$ for a QOMDP is then taken to be $\mathrm{Tr}\left(\mathbf{K}_{y_T}^{a_T}\ldots\mathbf{K}_{y_1}^{a_1}\;\rho_0\;\mathbf{K}_{y_1}^{a_1\,\dagger}\ldots\mathbf{K}_{y_T}^{a_T\,\dagger}\right)$. To draw connections with HQMMs, we represent the vectorized version of the QOMDP update equations as follows:

$$\vec{\rho}\,' = \mathrm{vec}\left(\frac{\mathbf{K}_{o_T}^{a_T}\ldots\mathbf{K}_{o_1}^{a_1}\;\boldsymbol{\rho}\;\mathbf{K}_{o_1}^{a_1\,\dagger}\ldots\mathbf{K}_{o_T}^{a_T\,\dagger}}{\mathrm{Tr}\left(\mathbf{K}_{o_T}^{a_T}\ldots\mathbf{K}_{o_1}^{a_1}\;\boldsymbol{\rho}\;\mathbf{K}_{o_1}^{a_1\,\dagger}\ldots\mathbf{K}_{o_T}^{a_T\,\dagger}\right)}\right) = \frac{\left(\overline{\mathbf{K}}_{y_T}^{a_T}\otimes\mathbf{K}_{y_T}^{a_T}\right)\ldots\left(\overline{\mathbf{K}}_{y_1}^{a_1}\otimes\mathbf{K}_{y_1}^{a_1}\right)\vec{\rho}}{\vec{\mathbb{I}}\left(\overline{\mathbf{K}}_{y_T}^{a_T}\otimes\mathbf{K}_{y_T}^{a_T}\right)\ldots\left(\overline{\mathbf{K}}_{y_1}^{a_1}\otimes\mathbf{K}_{y_1}^{a_1}\right)\vec{\rho}} = \frac{\mathbf{L}_{y_T}^{a_T}\ldots\mathbf{L}_{y_1}^{a_1}\vec{\rho}_0}{\vec{\mathbb{I}}^T\mathbf{L}_{y_T}^{a_T}\ldots\mathbf{L}_{y_1}^{a_1}\vec{\rho}_0}$$

Notice that this expression is identical to the update equation for IO-HQMMs. The only difference is 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 like NOOMs; however, unlike NOOMs, QOMDPs allow latent states of arbitrary Kraus-ranks – i.e. both mixed and pure states. We thus arrive at the following theorem characterizing the expressiveness of QOMDPs

Theorem 3.1. *QOMDPs ⊆ IO-HQMMS = non-terminating uniform IO-LPS ⊆ IO-OOMs = PSRs*

Quantum Markov decision processes were also studied by Ying and Ying (2014), who arrived at essentially the same model as 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 Stinesprig 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 is stated as follows: given an arbitrary initial state and a goal state, is there a sequence of operators that will leave the system in a goal 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. 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 3.2 (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-HQMM, (non-terminating uniform) IO-LPS, IO-OOM or PSR in the goal state 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. □

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.*

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

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. I., Perez-Garcia, D., Schuch, N., and Verstraete, F. (2017). Matrix product density operators: Renormalization fixed points and boundary theories. *Annals of Physics*, 378:100–149.

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.

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.*

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.

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.

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.

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

Miller, J., Rabusseau, G., and Terilla, J. (2021). Tensor networks for probabilistic sequence modeling. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics.*

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.

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.

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

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.

Stinespring, W. (1955). Positive functions on *-algebras.

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.

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.

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

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