

Amplitude Ratios and Neural Network Quantum States

Vojtěch Havlíček

IBM Quantum, IBM T.J. Watson Research Center

Neural Network Quantum States (NQS) represent quantum wavefunctions by artificial neural networks. Here we study the wavefunction access provided by NQS defined in [Science, **355**, 6325, pp. 602-606 (2017)] and relate it to results from distribution testing. This leads to improved distribution testing algorithms for such NQS. It also motivates an independent definition of a wavefunction access model: the amplitude ratio access. We compare it to sample and sample and query access models, previously considered in the study of dequantization of quantum algorithms. First, we show that the amplitude ratio access is strictly stronger than sample access. Second, we argue that the amplitude ratio access is strictly weaker than sample and query access, but also show that it retains many of its simulation capabilities. Interestingly, we only show such separation under computational assumptions. Lastly, we use the connection to distribution testing algorithms to produce an NQS with just three nodes that does not encode a valid wavefunction and cannot be sampled from.

Contents

1	Introduction	2
1.1	Results	2
2	Background	3
2.1	Neural Network Quantum States (NQS)	3
2.2	NQS and Restricted Boltzmann Machines (RBMs):	4
2.3	Approximations	4
2.4	The key simplifying assumption	6
3	Amplitude Ratios	6
3.1	Pair-cond and NQS	6
3.2	Amplitude Ratio (AR) Access	8
3.3	SQ vs AR	8
3.4	AR and probabilistic simulation	13
3.5	AR and dequantization?	16
4	Postselection Gadgets	17
4.1	Postselection Gadgets	17
4.2	Not all NQS encode valid quantum states	18
4.3	Other applications?	19

5	Discussion	19
6	Acknowledgements	20
	References	20
A	Additional Postselection Gadgets	22
	A.1 Hamming weight	22
	A.2 Parity	22
B	Pauli Gadgets	23

1 Introduction

Machine learning algorithms have been recently used for simulating quantum many-body systems. A prominent example is a method introduced in Ref. [6], called the *Neural Network Quantum State* (NQS). NQS encodes a quantum wavefunction $\psi_\theta(v) : \{-1, 1\}^n \rightarrow \mathbb{C}$. The encoded wavefunction is defined by complex-valued parameters $\theta = (a, b, W)$, where a are visible biases of the network, b are hidden biases, and W are interaction weights. The wavefunction can be accessed by sampling of the Born distribution $p_\theta(v) := |\psi_\theta(v)|^2$ and computing $\psi_\theta(v)$ up to an unknown and usually inaccessible normalization factor. This allows for efficient computation of ratio of the wavefunctions $\psi_\theta(v)/\psi_\theta(w)$ for two outcomes v, w .

1.1 Results

This work explores the relationship between this wavefunction access and results from conditional distribution testing. This motivates a definition of a wavefunction access model, the amplitude ratio (AR) access. AR can be seen as a variant of the sample and query (SQ) access to wavefunction, which has seen use in analysis of quantum algorithms and dequantization (see for example [25]). Here we argue that AR is a strictly weaker model than the standard SQ. Secondly, the connection to distribution testing allows us to show that there exist small NQS that cannot be sampled from, in an analogy to results about sampling from Restricted Boltzmann Machines [19, 20].

NQS and PCOND: Given a NQS θ and two measurement outcomes $v_0, v_1 \in \{-1, 1\}^n$, it is possible to efficiently compute and sample from $p_\theta(v|v = v_0 \text{ or } v = v_1)$. We observe that this instantiates the *pair-cond* (PCOND) *oracle* used in conditional distribution testing [3, 7]. It is known that PCOND leads to significant improvements in *query complexity* [3, 7, 4]. We show how this improvement enables computational advantages for algorithms for testing NQS Born distributions.

Amplitude ratio (AR) access: We define a version of PCOND for amplitudes, the *amplitude ratio* (AR) *access*. We discuss that it can be understood as a variant the sample and query (SQ) wavefunction access¹ that found many uses in analysis of quantum algorithms and dequantization [25]. We compare AR to SQ with normalized queries and show that it retains many of its simulation capabilities, such as:

¹AR can be seen as an SQ without access to normalization factors. In Ref. [25], Tang uses SQ access to a vector without normalization but their key subroutine requires access to the normalization factor of the vector, resulting in SQ access to normalized data. See Sec.3.3 for discussion.

- a query-efficient fidelity estimator.
- the ability to compute expectation values of sparse observables.

We give evidence that SQ is strictly stronger access model than AR and show that AR is a strictly stronger model than PCOND. We derive a robust version of the fidelity estimator and show how to estimate sparse observables with AR access.

NQS postselection gadgets: We show how to postselect the Born distribution of an NQS by changing its network structure. We call such transformation a NQS postselection gadget and use it to give an NQS with only three nodes (and polynomially-bounded weights and biases) that does not encode a valid wavefunction. As the result implies that the NQS distribution cannot be sampled from, it can be understood as a counterpart to the best known hardness of sampling result for Restricted Boltzmann Machines, which shows that a certain distribution cannot be represented (and hence sampled from) with a polynomially-sized RBM [20]. We briefly discuss other possible application of the gadgets.

2 Background

2.1 Neural Network Quantum States (NQS)

Ref. [6] proposed NQS representation of quantum wavefunctions by a hidden Markov models, largely inspired by Restricted Boltzmann Machines (RBM) [23, 12]. We briefly review it here.

Let $v \in \{-1, +1\}^n$ and define $\psi(v) = f_\theta(v)/Z_\theta$, where:

$$\begin{aligned} f_\theta(v) &= \sum_{h \in \{-1, +1\}^m} \exp(a^\top v + b^\top h + v^\top W h) \\ &= \exp(a^\top v) \prod_{i=1}^m 2 \cosh \left(b_i + \sum_j v_j W_{ji} \right), \end{aligned} \quad (1)$$

and $Z_\theta := \sqrt{\sum_v |f_\theta(v)|^2}$. The parameters:

$$\theta := (a, b, W), \quad (2)$$

are all complex-valued; $a \in \mathbb{C}^n, b \in \mathbb{C}^m, W \in \mathbb{C}^{n \times m}$ and fully specify the model. We denote $\|\theta\|_\infty = \max(\|a\|_\infty, \|b\|_\infty, \|W\|_\infty)$ and assume that $\|\theta\|_\infty \leq \text{poly}(n)$.

Note that $Z_\theta := \sqrt{\sum_v |f_\theta(v)|^2}$ sums over all configurations v and that there is a priori no simple way to evaluate it. In contrast, the numerator $f_\theta(v)$ in Eq. 1 can be evaluated to machine precision² in polynomial time in m and n . This follows by observing that each of the m factors only depends on a $\sum_j v_j W_{ji}$ which has at most n terms. See Fig. 1.

The set of parameters θ is usually found by variational optimization [6]. This relies on sampling from the *Born distribution* $|\psi_\theta(v)|^2$ and gradually updating the network parameters. Sampling from the Born distribution is usually done by “thermalizing” the model with a Markov Chain Monte Carlo; or *Gibbs sampling* from the Born distribution. The same method is then used to sample from the Born distribution of a trained model. There is generally no guarantee that the Markov chains, both during the training and testing phase, converge to the target distribution rapidly for a given θ . Remarkably, we in show in Sec. 4.2 that there exist parameters for which no such Markov chain converges.

²See Sec. 2.3.




Figure 1: NQS is composed of hidden (h) and visible (v) nodes arranged in a bipartite graph. Every node can be in a state $+1$ or -1 . Each edge (i, j) of the graph carries a complex-valued *interaction strength* W_{ij} and every node carries a *local field* a_i (for visible) or b_j (hidden). The model encodes the unnormalized wavefunction into a “marginal amplitude” over hidden nodes (Eq. 1). It can be sampled from using Gibbs sampling.

2.2 NQS and Restricted Boltzmann Machines (RBMs):

NQS were inspired by Restricted Boltzmann Machines (RBMs) [6] and while the two models share many similarities, there are important differences. An RBM represents a *distribution* $p_\theta(v)$ for $v \in \{-1, +1\}^n$ by a marginal over the Gibbs-Boltzmann distribution of an Ising model on a bipartite graph. The model is, in the simplest setting, defined by a set of real-valued weights and biases (local fields). The output distribution is represented as the thermal distribution *marginalized* over the hidden nodes:

$$p_\theta^{\text{RBM}}(v) \propto \sum_{h \in \{-1, +1\}^m} \exp(a^\top v + b^\top h + v^\top W h). \quad (3)$$

Even though Eq. 1 and Eq. 3 look very similar, NQS is *not* straightforwardly represented by a marginal distribution. The model instead implicitly defines a Born distribution:

$$p_\theta^{\text{NQS}}(v) \propto \left| \sum_{h \in \{-1, +1\}^m} \exp(a^\top v + b^\top h + v^\top W h) \right|^2. \quad (4)$$

This allows for *interference* between the summands. Because the network parameters can be complex-valued, such interference can be destructive. This means that Eq. 4 can evaluate to zero. In comparison, the marginal sum for an RBM in Eq 3 is lower-bounded by $2^m \exp(-\|a\|_\infty n)$, which is large for $m \gg n$ and small $\|a\|_\infty$. In such setting, the RBM algorithm cannot faithfully represent zeros in the output distribution. We show in Sec. 4.2 that this implies that there exist parameters for which the NQS does not define a valid distribution.

2.3 Approximations

Here we introduce some notions of approximation that will be used throughout the work. We also state and justify the key assumption (Assumption 1), which ensures that the notion of amplitude ratios is well-defined.

Additive approximations: Let $\epsilon \geq 0$. An ϵ -additive approximation \tilde{R} to a real number R is a real number that satisfies:

$$|R - \tilde{R}| \leq \epsilon. \quad (5)$$

Relative approximations: Let $\epsilon > 0$. An ϵ -relative approximation \tilde{R} to a non-negative real number $R \in \mathbb{R}_0^+$ is a real number that satisfies:

$$(1 + \epsilon)^{-1} R \leq \tilde{R} \leq (1 + \epsilon) R. \quad (6)$$

Lemma 1. Fix $\epsilon > 0$. Let $\tilde{Q} \in \mathbb{R}$ and $\tilde{R} \in \mathbb{R}$ be ϵ -relative approximations to $Q, R \in \mathbb{C}$ respectively. Then \tilde{Q}/\tilde{R} is a 3ϵ -relative approximation to Q/R , $\tilde{Q}\tilde{R}$ is a 3ϵ -relative approximation to QR .

Proof. First note that:

$$\frac{1}{(1+\epsilon)^2} \frac{Q}{R} \leq \frac{\tilde{Q}}{\tilde{R}} \leq (1+\epsilon)^2 \frac{Q}{R}. \quad (7)$$

Set $1+\epsilon' = (1+\epsilon)^2$ as the degree of relative approximation to the ratio, from which $\epsilon' = \epsilon^2 + 2\epsilon \leq 3\epsilon$. Given an ϵ -relative approximation \tilde{Q} to Q , $1/\tilde{Q}$ is an ϵ -relative approximation to $1/Q$. This means that $\tilde{Q}\tilde{R}$ is a 3ϵ -relative approximation to QR . \square

Complex numbers: We will require a similar notion of approximation for complex numbers. To simplify our analysis, we use the following simplification: we assume that the complex numbers are stored in polar form as $e^{i\alpha}R := (\alpha, R)$ for $\alpha \in [0, 2\pi)$ and $R \in \mathbb{R}_0^+$ and that the phase α is stored exactly³. By an ϵ -relative approximation to a complex number $C = e^{i\alpha}R$, we mean a number \tilde{C} that satisfies:

$$(1+\epsilon)^{-1}R \leq \tilde{R} \leq (1+\epsilon)R, \quad (8)$$

and $\alpha = \tilde{\alpha}$. Lemma 1 works for this notion of approximation if we assume that the result of adding two phases is always mapped back to $[0, 2\pi)$.

Machine precision: Machine precision is an upper bound on the relative error ϵ due to rounding in floating point arithmetic: numbers are represented up to a finite number of bits, which leads to *usually* insignificant additive errors. These errors *usually* imply good relative approximation. We define this more precisely now:

Lemma 2. Let $R \in \mathbb{R}^+$ and $\epsilon \in [0, 1)$. Assuming $\epsilon \leq R$, an $\frac{\epsilon^2}{2}$ -additive approximation \tilde{R} to R is also an ϵ -relative approximation to R .

Proof. This follows from:

$$(1+\epsilon)^{-1}R \leq \left(1 - \frac{\epsilon}{2}\right)R \leq \left(R - \frac{\epsilon^2}{2}\right) \leq \tilde{R}, \quad (9)$$

$$\tilde{R} \leq \left(R + \frac{\epsilon^2}{2}\right) \leq \left(1 + \frac{\epsilon}{2}\right)R \leq (1+\epsilon)R. \quad (10)$$

\square

Throughout this work, we will make a simplifying assumption that by *evaluation to machine precision*, we mean evaluation to ϵ -relative error in $\text{poly}(\log(1/\epsilon))$ time and justify it shortly. We will often require evaluation of some quantities to machine precision in $\text{poly}(n, \log(1/\epsilon))$ time, where n is the input size to the problem. From this, we bind the error parameter ϵ to the input size and simply assume that standard functions, such as exp or cosh, can be efficiently evaluated to $2^{-\text{poly}(n)}$ error for arguments with magnitude bounded by some polynomial $\text{poly}(n)$ (similarly to Eq. 1).

The following is an important consequence of Lemma 2:

Corollary 1. If $Q, R \geq 2^{-\text{poly}(n)}$ can be efficiently evaluated to machine precision, then so can $1/Q$, QR and Q/R .

³In practice, it will be only stored to a sufficiently high *additive* precision, using up to $\text{poly}(n)$ bits. We will ignore the issues that arise by assumption of storing α exactly, such as when α is an irrational number.

2.4 The key simplifying assumption

To make use of Corollary 1, we will often need the following assumption regarding wavefunctions:

Assumption 1. *We assume that $|\psi(v)| \geq 2^{-\text{poly}(n)}$ or $\psi(v) = 0$ for some sufficiently large polynomial $\text{poly}(n)$.*

This always holds for any wavefunction represented by NQS, because for $\|\theta\|_\infty \leq \text{poly}_1(n)$, we have $f_\theta(v) \geq 2^{-\text{poly}_2(n)}$ for a suitable polynomial (perhaps distinct from $\text{poly}_1(n)$) and $Z \leq 2^{-n}$. The condition is not automatically guaranteed for other families of wavefunctions and this may become problematic for the access models that we study in the next section. The assumption also guarantees that a machine precision approximation to $\psi(v)$ is always representable by $\text{poly}(n)$ bits.

3 Amplitude Ratios

We first compare the NQS wavefunction access model to the *pair-cond* (PCOND) query access to distributions from Refs. [7, 3]. PCOND queries are strictly more powerful than sampling and we show that they can be simulated efficiently for any NQS distribution. This gives improved algorithms for distribution testing. It also leads to a modification of PCOND for quantum wavefunctions. We define this as amplitude ratio (AR) access and study it independently of NQS. We compare AR to sample and query (SQ) access used in dequantization [25] and probabilistic simulation [27]. We argue that AR is weaker access model than SQ with *normalized* queries, but retains many classical simulation techniques of SQ.

3.1 Pair-cond and NQS

NQS wavefunctions can be efficiently evaluated up to a normalization factor for any $v \in \{-1, +1\}^n$ (Sec. 2.1). This enables computation of amplitude ratios:

Lemma 3. *Let $\psi(v)$ be an NQS over $\{-1, +1\}^n$. The ratio $\psi(i)/\psi(j)$ can be efficiently evaluated to machine precision for any $i, j \in \{-1, +1\}^n$.*

Proof. Observe that $\psi(i)/\psi(j) = f_\theta(i)/f_\theta(j)$. The claim follows from:

$$f_\theta(v) = \exp(a^\top v) \prod_{i=1}^m 2 \cosh \left(b_i + \sum_j^n v_j W_{ji} \right), \quad (11)$$

(see Eq. 1) and noting that each of the m factors only depends on $\sum_j^n v_j W_{ji}$, which has at most n terms. Since we assumed that $\|\theta\|_\infty \leq \text{poly}(n)$ below Eq. 2, $f_\theta(v)$ can be efficiently evaluated to machine precision (Sec. 2.3). The result follows from Corollary 1. \square

There is an analogy between amplitude comparison and *pair-cond* (PCOND) oracle access used in conditional distribution testing [3].

Definition 1 (PCOND [3]). *Let p be a probability distribution over Ω . The PCOND oracle accepts an input set S that is either $S = \Omega$ or $S = \{i, j\}$ for some $i, j \in \Omega$. It returns an element $i \in S$ with probability $p(i)/p(S)$, where $p(S) = \sum_{i \in S} p(i)$.*

Refs. [3, 7] show that PCOND queries lead to significant complexity improvement for some distribution testing tasks. We observe that:

Problem	Random samples	PCOND queries
Is $p_\theta(v)$ ϵ -close to uniform distribution?	$\Theta(2^{n/2}/\epsilon^2)$	$\tilde{O}(1/\epsilon^2)$ [3]
Is $p_\theta(v)$ ϵ -close to a known distribution $q(v)$?	$\Theta(2^{n/2}/\epsilon^2)$	$\tilde{O}(n^4/\epsilon^4)$ [3]
Are $p_{\theta_1}(v), p_{\theta_2}(v)$ ϵ -close?	$\Omega(2^{n/2}/\epsilon^2)$	$\tilde{O}(n^6/\epsilon^{21})$ [3]

Table 1: PCOND results of Ref. [3]. Here ϵ -close means ϵ -close in the total variational distance: $d(p, q)_{\text{TV}} := \frac{1}{2} \|p - q\|_1 = \frac{1}{2} \sum_x |p(v) - q(v)|$. Since Sec. 3.1 shows that any efficiently samplable NQS has an efficient PCOND oracle, the sample complexity improvements translate to improvements in the runtime of distribution testing algorithms for NQS.

Theorem 1. PCOND (Def. 1) can be instantiated efficiently for any NQS that can be efficiently sampled from.

Proof. Assume that the NQS encodes a wavefunction $\psi_\theta(v)$ over $\{-1, 1\}^n$. From Eq. 1, we have that $\psi_\theta(v) = f_\theta(v)/Z_\theta$, where $f_\theta(v)$ can be computed efficiently to machine precision.

To show that NQS allows PCOND queries, fix $\Omega = \{-1, 1\}^n$ in the definition of PCOND.

1. On input $S = \Omega$, output a sample from the Born distribution encoded by the NQS.
2. On input $S = \{i, j\}$ for $i, j \in \Omega$, compute:

$$r = \frac{|f_\theta(i)|^2}{|f_\theta(i)|^2 + |f_\theta(j)|^2} = \left(1 + \left|\frac{f_\theta(j)}{f_\theta(i)}\right|^2\right)^{-1}. \quad (12)$$

Return i or j according to a r -biased coin flip (if both $|f_\theta(i)|^2$ and $|f_\theta(j)|^2$ are zero, return one of i, j u.a.r).

The ratio in Eq. 12 is not computed exactly, but we assume that the deviation from its *exact* value will be only observed for exponentially many coin flips in the input size. Since we are interested in algorithms that make at most polynomially many PCOND queries, we neglect this. \square

Theorem 1 enables efficient implementation of algorithms for NQS distributional testing; see Tab. 1. The algorithms test the total variation distance of two NQS Born distributions and have exponential advantage in their runtime compared to sampling algorithms.

Procedure compare The key algorithmic tool used in Ref. [3] is a procedure *compare* that uses PCOND queries for estimating probability ratios. Given that the ratio $f_\theta(i)/f_\theta(j)$ is not too large or too small, the procedure *compare* outputs its $1/\text{poly}(n)$ additive approximation with high probability of success using $\text{poly}(n)$ many PCOND queries [3]. For NQS, this ratio can be efficiently computed to machine precision, which offers further simplifications of the algorithms. This motivates the definition of *amplitude ratio* wavefunction access model introduced in the next section.

PCOND and RBM: PCOND can be also instantiated for RBMs. The PCOND algorithms of Ref. [3] apply with little modification to RBMs and imply significant speedups for some distribution testing tasks. Results of Ref. [3] were presented as a part of a theoretical analysis in conditional property testing and do not have runtimes that would make

them immediately practical. No natural and efficient instantiation of the PCOND oracle (such as the one in the context of RBMs) has been known to the authors in mid 2021 [2]. It is therefore possible that the algorithms could be optimized and perhaps used in practice.

Previous quantum information work on PCOND: PCOND access was studied from the quantum computing viewpoint by Sardharwalla, Strelchuk and Jozsa in Ref. [22], but the AR definition (introduced below) is new. The difference between their definition and AR is that their quantum extension of PCOND oracle (PQCOND) provides access to conditional probabilities associated with the underlying distribution – their oracles are defined at a level of *quantum states*, while the AR access is always defined with respect to some fixed basis/*wavefunction*. The authors give a PQCOND version of the *compare* procedure from Ref. [3] and show that PQCOND queries have polynomial improvements upon many of the PCOND distribution testing results. They also derive results on boolean distribution testing and quantum spectrum testing. There does not seem to be an obvious connection between AR access and PQCOND, but it would be interesting to understand this better.

3.2 Amplitude Ratio (AR) Access

Because NQS allows for computation of *amplitude ratios* to high precision, they should offer stronger access to the wavefunction than PCOND. We study this as a wavefunction access model, independently of NQS.

Definition 2 (Exact AR). *Let $\psi : \Omega \rightarrow \mathbb{C}$ be a wavefunction over Ω . The AR oracle accepts as an input either an ordered pair $S = (i, j)$ or $S = \Omega$. If $S = \Omega$, it returns a random sample from the Born distribution $|\psi(v)|^2$ over $v \in \{-1, 1\}^n$. If $S = (i, j)$, it returns $\psi(i)/\psi(j)$. If the ratio diverges, it returns a special symbol ‘DIV’ and it conventionally returns 1 for $\psi(i) = \psi(j) = 0$.*

Definition 2 assumes that the queries are returned exactly. This becomes problematic if $\psi(i)/\psi(j)$ becomes irrational number as the query result is an infinitely-long output. This issue can be dealt with as follows:

Definition 3 (AR). *Let $\psi : \Omega \rightarrow \mathbb{C}$ be a wavefunction over Ω subject to Assumption 1. For $\epsilon \in [0, 1)$, the AR(ϵ) oracle accepts as an input either an ordered pair $S = (i, j)$ or $S = \Omega$. If $S = \Omega$, it returns a random sample from the Born distribution $|\psi(v)|^2$ over $v \in \{-1, 1\}^n$. If $S = (i, j)$, it returns an ϵ -relative approximation to $\psi(i)/\psi(j)$.⁴ If the ratio diverges, it returns ‘DIV’ and it conventionally returns 1 for $\psi(i) = \psi(j) = 0$. We write $AR := AR(\epsilon)$ if ϵ scales as $2^{-poly(n)}$ for some polynomial $poly(n)$.⁵*

3.3 SQ vs AR

We now compare AR with another type of access to quantum wavefunctions, the *sample and query* (SQ) access. SQ was defined in Ref. [25], but was previously used as *computational tractability* with additional computational requirements in Ref. [27]. There is a difference between the two definitions: Ref. [25] defines SQ access *without* normalizing

⁴To make the queries deterministic, one can force this to be the approximation with the shortest binary expansion.

⁵The AR oracle can be defined rigorously in the oracle Turing machine model. Let the author know if you ever need this.

the queries, but subsequently assumes knowledge of the normalization factor (see for example Prop 4.2. [25]), while Ref. [27] assumes normalization (as well as efficiency) in the definition of computationally tractable states (Ref. [27], Def. 1). Tang’s definition of SQ also does not treat the underlying object as a wavefunction, but more generally as a real valued vector with ℓ_2 -norm sampling access. This presentation puts less emphasis on the need for the normalization factor. Here we use the SQ access model assuming that *the queries yield normalized amplitudes*, but impose no efficiency constraints.

Definition 4 (Exact SQ). *Let $\psi : \Omega \rightarrow \mathbb{C}$ be a wavefunction over Ω . The wavefunction has SQ access if $\psi(i)$ can be computed for any $i \in \{-1, +1\}^n$ and its Born distribution $|\psi(i)|^2$ can be sampled from.*

The above definition has the same problem as *exact* AR (Def. 2): if the result of the query is an irrational number, the result of the query will not have bounded size. We update it as follows:

Definition 5 (SQ). *Let $\psi : \Omega \rightarrow \mathbb{C}$ be a wavefunction over Ω subject to Assumption 1. For $\epsilon \in [0, 1)$, the SQ(ϵ) oracle accepts as an input either a singleton set $S = \{i\}$ or $S = \Omega$. If $S = \Omega$, it returns a random sample from the Born distribution $|\psi(v)|^2$ over $v \in \{-1, 1\}^n$. If $S = \{i\}$, it returns an ϵ -relative approximation to $\psi(i)$ with the shortest binary expansion. We write $SQ := SQ(\epsilon)$ if ϵ scales as $2^{-\text{poly}(n)}$ for some polynomial $\text{poly}(n)$.*

AR and unnormalized SQ Tang used SQ with unnormalized queries in [25] but subsequently assumed knowledge of the normalizing factor to the wavefunction throughout the work. Without knowledge of the normalization factor, the only information about the vector with such access comes from amplitude ratios and sampling. This way, the definition is essentially the same as AR – so why bother with AR?

The key reason for defining AR is to emphasize that the normalization factor of the wavefunction is simply unavailable, aside from its empirical estimate by sampling. It seems problematic to guarantee this (at least somewhat) rigorously: does evaluation of the state up to an “arbitrary” normalization factor always force you to not use it? *That subtlety aside, one can also see AR as unnormalized SQ access and the rest of this section as comparison between normalized and unnormalized variants of SQ.*

AR is not stronger than SQ AR can be simulated by SQ:

Lemma 4. *Under Assumption 1, an AR(ϵ) query can be simulated by two SQ($\epsilon/3$) queries.*

Proof. Let $\tilde{\psi}(i)$ and $\tilde{\psi}(j)$ be results of querying SQ($\epsilon/3$) on i and j respectively. To simulate AR(ϵ), output $\tilde{\psi}(i)/\tilde{\psi}(j)$. This is an ϵ -approximation to $\psi(i)/\psi(j)$ by Lemma 1. \square

Evidence that AR is weaker than SQ: To show that AR access is *in some sense* weaker than SQ, we show conditional separation between variants of the two models under *efficiency* constraints. SQ is related to the concept of computationally tractable (CT) states [27]:

Definition 6 (Computationally Tractable (CT) states). *An n -qubit wavefunction $\psi(v)$ over $v \in \{-1, +1\}^n$, subject to Assumption 1, is computationally tractable (CT) if both queries in Def. 5 can be implemented with a $\text{poly}(n)$ -time randomized algorithm.*⁶

This can be seen as SQ with efficient classical sampler and efficient classical algorithm for the amplitude queries. The key capability of CT states is an efficient algorithm that, given two CT wavefunctions $\psi(v)$ and $\phi(v)$, approximates $\langle \psi | \phi \rangle$ in polynomial time to inverse-polynomial precision (see Theorem 3. and Lemma 3. in Ref. [27] or equivalently Prop 4.8. in Ref. [25]). This enables estimation of constant-local bounded observables on CT states or simulation of sparse quantum circuits. Techniques related to CT framework were used in dequantization algorithms in Ref. [25] or quantum algorithm analysis in Ref. [11].

Analogously to CT states, we define *amplitude ratio (AR) states* and show that their *fidelities* and expectation values on constant-local observables can be also efficiently computed. They subsume CT states, which suggests that the CT requirement can be relaxed to AR in many applications.

Definition 7 (Amplitude Ratio (AR) states). *An n -qubit wavefunction $\psi(v)$ over $v \in \{-1, +1\}^n$, subject to Assumption 1, is an amplitude ratio (AR) state if queries to AR in Def. 3 can be implemented with a $\text{poly}(n)$ -time randomized algorithm.*

We show that, subject to Assumption 1, all CT states are AR states. Our *evidence* that SQ access is somewhat stronger than the AR access follows from that for *machine precision*, not all AR states are CT, unless $\#\text{P} = \text{FBPP}$.⁷ This is shown using the separation between (exact) counting and uniform sampling by Jerrum, Valiant and Vazirani (Sec. 4 [15]).

Theorem 2. *All CT states are AR. There exist AR states that are not CT unless $\#\text{P} = \text{FBPP}$.*

Proof. (All CT states are AR): Both models perform AR and SQ queries in $\text{poly}(n)$ time. The rest follows from Lemma 4.

(There is an exact AR state that is not CT): The proof uses the observation that uniform sampling of solutions to a given boolean formula over n variables in disjunctive normal form (DNF) is easy, while their exact enumeration is $\#\text{P}$ -complete (Sec. 4 of Ref. [14]). We consider a quantum state that is a uniform superposition over satisfying assignments to the boolean formula and show that its Born distribution can be easily sampled from. The normalization factor of such state counts the number of satisfying solutions to the formula, which is $\#\text{P}$ -complete to compute exactly. We show that a good relative approximation to the amplitude determines this quantity.

Given a boolean formula in DNF over n variables, interpret $v \in \{-1, +1\}^n$ as an assignment to its variables: if $v_i = +1$, then the i -th variable is true and if $v_i = -1$, the i -th variable is false. Define the DNF formula-“state” as follows:

$$\psi_{\text{DNF}}(v) = \begin{cases} \frac{1}{\sqrt{Z}} & \text{if } v \text{ satisfies the formula.} \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

⁶Van den Nest uses exact queries, but also writes that his definition generalizes straightforwardly to computation with exponential additive precision. This definition is a modification of Def. 1, Ref. [27], which also does not need the Assumption 1

⁷ $\#\text{P}$ is the class of function problems that enumerate accepting paths to NP problems [26]. FBPP is the class of functions computable in polynomial time with a probabilistic Turing machine.

Notice that Z is the number of variable assignments that satisfy the formula.

For any input v , the predicate $|\psi_{\text{DNF}}(v)| > 0$ can be tested by plugging in the variable assignments into the formula. Any non-zero amplitude evaluates to $1/\sqrt{Z}$, from which it is possible to compute the amplitude ratio as required by exact AR.

The Born distribution $|\psi_{\text{DNF}}(v)|^2$ is a uniform distribution over the satisfying assignments to the boolean formula. This distribution can be sampled from exactly by a polynomial-time randomized algorithm described in Sec. 4 of Ref. [15]. It works as follows: For a DNF boolean formula $F = F_1 \vee F_2 \dots \vee F_m$ in n variables, where F_i is a conjunction of literals for all $i \in [1, m]$, let $S_j \subseteq \{-1, +1\}^n$ be the set of satisfying assignments to F_j . Note that $|S_j|$ is easy to compute, because F_j is only satisfied if it fixes a subset of the variables involved in it and the remaining variables can take arbitrary values. Let $S = \bigcup_j S_j$ the set of all satisfying assignments to F . The aim is to sample uniformly over S , which is achieved by Algorithm 1.

Algorithm 1 Uniform sampling of DNF assignments [15]

for $i \in 1 \dots m$ **do**:

$j \leftarrow$ random integer from $[1, m]$ chosen proportionally to $|S_j| / \sum_k |S_k|$.

$a \leftarrow$ random element of S_j .

$N \leftarrow |\{k \in [1, m] : a \in S_k\}|$.

With $1/N$ probability, output a and halt.

end for

The algorithm does not halt with small probability. If it does not halt, rerun the for-loop. The output of the algorithm is a uniformly random satisfying variable assignment. This implies that $\psi_{\text{DNF}}(v)$ is an exact-AR state.⁸

The family of states ψ_{DNF} is not exact-CT unless $\#\text{P} = \text{FBPP}$, because $\psi_{\text{DNF}}(v) = 1/\sqrt{Z}$ for any satisfying variable assignment v . Evaluating this exactly is a $\#\text{P}$ -complete problem, because Z is the number of satisfying solutions to the boolean formula. So unless $\text{FBPP} = \#\text{P}$, $\psi_{\text{DNF}}(v)$ is an exact-AR state that is not exact-CT.

It remains to show that $\psi_{\text{DNF}}(v)$ is not CT. To do this, we show that CT can compute Z to sufficient accuracy to determine it exactly. First note that $Z \leq 2^n$. We want to choose ϵ so that the ϵ -relative approximations to Z and $Z + 1$ can be distinguished for any $0 \leq Z \leq 2^n$. This happens if $(1 + \epsilon)Z < (1 + \epsilon)^{-1}(Z + 1)$. Since $\epsilon < 1$, we have that $(1 + \epsilon)^2 < 1 + 3\epsilon$ and $(1 + 3\epsilon)Z < Z + 1 \implies (1 + \epsilon)^2 Z < Z + 1$. We can therefore choose $\epsilon \leq 2^{n-2}$ to recover Z exactly. It follows that if $\psi_{\text{DNF}}(v)$ was CT, Z could be computed exactly by a polynomial time randomized algorithm (FBPP). This problem is however $\#\text{P}$ -complete, so this gives a contradiction unless $\#\text{P} = \text{FBPP}$. \square

We remark that the above argument won't work for $\epsilon = 1/\text{poly}(n)$. The reason is that there is a polynomial time randomized algorithm for approximating the number of DNF satisfying clauses to relative error ϵ that outputs the solution with high probability in $\text{poly}(n, 1/\epsilon)$ time [17].⁹ This is worse approximation than what is assumed in the definition of AR states, which shows that one has to use different argument to separate $\text{CT}(\epsilon)$ from $\text{AR}(\epsilon)$ for ϵ scaling as $1/\text{poly}(n)$. Similar argument to the above leads to an argument that *could* separate SQ and AR in terms of their query complexity:

⁸As discussed in Ref. [15] Sec. 4, Algorithm 1 cannot be implemented on a probabilistic Turing Machine. A modification of it, also in [15] Sec. 4, deals with this subtlety and leads to an FBPP algorithm.

⁹Such algorithm is called *fully-polynomial randomized approximation scheme* (FPRAS)

Lemma 5. *Given SQ access to ψ_{DNF} from the previous theorem, a single query suffices to approximate Z to machine precision.*

Proof. (sketch) Given ψ_{DNF} and an SQ access to it, the SQ algorithm draws a sample v from $|\psi_{\text{DNF}}(v)|^2$ and evaluates $\psi_{\text{DNF}}(v) = 1/\sqrt{Z}$, which gives a machine precision approximation to Z . \square

Conjecture 1. *Z for ψ_{DNF} can be query-efficiently estimated by AR to no better than $1/\text{poly}(n)$ relative error.*

The reasoning behind the conjecture is the following: Given AR access to ψ_{DNF} , any amplitude ratio query on a pair of nonzero amplitudes evaluates to 1. Any amplitude ratio on pairs of zero and nonzero amplitudes gives either 0 or ‘DIV’. By finding a v with zero amplitude and a single w with a nonzero amplitude, AR can detect if $\psi_{\text{DNF}}(z) > 0$ for any $z \in \{-1, +1\}^n$. Possibly the best way to approximate Z with this access *seems to be* variants of importance/nested sampling (see for example [17, 24, 8, 13]), which at best lead to $\text{poly}(n)$ -sample algorithm that estimates Z to $1/\text{poly}(n)$ relative error.¹⁰ It would be extremely surprising if these methods were not asymptotically optimal. I don’t have a proof though.

Theorem 3. *Assuming Conjecture 1, there is a task that requires just one SQ query, but at least $\text{poly}(n)$ -many AR queries.*

AR is stronger than sample access This is shown by noting that AR access implies PCOND access and then referencing the known results from conditional distribution testing that separate PCOND and sample access.

Lemma 6. *AR simulates PCOND.*

Proof. See discussion above Eq. 12. \square

AR access is stronger than PCOND access and PCOND is strictly stronger than sample access (Tab 1). It is worth noting that the algorithms of [2] work even if the AR ratios can be approximated only to $1/\text{poly}(n)$ additive error. It may be therefore interesting to study weaker variants of AR.

AR is stronger than PCOND

Lemma 7. *AR is stronger than PCOND.*

Proof. (sketch) AR can decide if $|\psi(i)/\psi(j)|^2 \leq 2^{-n}$ in a single query, but the same task requires exponentially many PCOND queries. This is because the PCOND model can estimate the ratio only by sampling, which is limited by the usual concentration bounds. These imply a lower bound of exponentially many queries. \square

One can object that the above comparison is rather unfair, because AR computes the amplitude ratio to a high degree of precision with a single query, while PCOND can only do so query-efficiently to $1/\text{poly}(n)$ additive error. We strengthen the above lemma to separation from a version of high-precision PCOND that can query for ratios of Born distribution as follows:

¹⁰FPRAS runs in $\text{poly}(n, 1/\epsilon)$, but we assume $\epsilon = 1/\text{poly}(n)$ binding. An AR access algorithm that achieves this precision for this problem is a variant of the Karp-Luby-Madras Markov Chain [17].

Lemma 8. *AR is stronger than a version of PCOND that can query ratios of Born distributions to machine precision.*

Proof. First, this version of PCOND can be simulated by AR by squaring all amplitude ratio queries. Let $\psi = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ and $\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and note that:

$$|\langle \psi | \phi \rangle|^2 = \sum_{y,x} \psi^*(x) \psi(y) \phi^*(y) \phi(x) = \sum_{x,y} |\psi(x)|^2 |\phi(y)|^2 \frac{\psi(y) \phi(x)}{\psi(x) \phi(y)} \quad (14)$$

$$= \frac{1}{4} \left[\frac{\psi(0) \phi(0)}{\phi(0) \psi(0)} + \frac{\psi(1) \phi(0)}{\phi(1) \psi(0)} + \frac{\psi(0) \phi(1)}{\phi(0) \psi(1)} + \frac{\psi(1) \phi(1)}{\phi(1) \psi(1)} \right] = 0, \quad (15)$$

which can be computed with 2 AR queries. The states have the same, uniform, Born distributions, so all PCOND ratio queries evaluate to 1. The PCOND therefore cannot evaluate the overlap of ψ and ϕ . The PCOND access model therefore cannot evaluate the overlap of ψ and ϕ . \square

Summary Results of Sec. 3.3 can be summarized as :

$$\text{SQ} \stackrel{*}{\supset} \text{AR} \supset \text{PCOND} \supset \text{SAMP}, \quad (16)$$

since PCOND access is separated from sampling access to a Born distribution of a wavefunction by Tab 1, AR is separated from PCOND access by Lemma 8, and SQ is *almost surely* (that's why the star) separated from AR by Theorems 2 and 3. By separation, we mean that there exists at least one problem that can be query efficiently solved by one of the classes, but not by the other.

3.4 AR and probabilistic simulation

We now show that AR states retain many simulation capabilities of CT states. Most of the results follow from standard algorithms used with NQS that implicitly used the AR access. We improve some of them, for example by robustification of the AR fidelity estimator, to give the closest possible analogues to the previous results for CT states [27].

AR fidelity estimator: Given two amplitude ratio (AR) states ψ and ϕ , there is a randomized algorithm that approximates their fidelity $|\langle \psi | \phi \rangle|^2$ in polynomial time to inverse-polynomial precision. Medvidovic and Carleo use the following estimator for NQS in Ref. [21]:

$$|\langle \phi | \psi \rangle|^2 = \sum_{x,y} \phi^*(x) \psi(x) \psi^*(y) \phi(y) \quad (17)$$

$$= \frac{Z_\phi}{Z_\psi} \frac{Z_\psi}{Z_\phi} \sum_x |\phi(x)|^2 \frac{f_\psi(x)}{f_\phi(x)} \sum_y |\psi(y)|^2 \frac{f_\phi(y)}{f_\psi(y)} \quad (18)$$

$$= \sum_x |\phi(x)|^2 \frac{f_\psi(x)}{f_\phi(x)} \sum_y |\psi(y)|^2 \frac{f_\phi(y)}{f_\psi(y)}. \quad (19)$$

Every term in the summation uses two AR ratio queries, which can be seen from:

$$|\langle \phi | \psi \rangle|^2 = \sum_{x,y} |\phi(x)|^2 |\psi(y)|^2 \left[\frac{f_\psi(x)}{f_\phi(x)} \right] \left[\frac{f_\phi(y)}{f_\psi(y)} \right]. \quad (20)$$

This can be operationally understood as follows: sample $(x, y) \sim |\phi(x)|^2|\psi(y)|^2$ (which is a valid product distribution) and compute the product of AR ratios on ψ for (x, y) and ϕ for (y, x) .

We now show that the estimator has finite variance and give a robust version of it with fast concentration around the mean. Set:

$$G(x, y) = \left[\frac{\psi(x)\phi(y)}{\psi(y)\phi(x)} \right] = \left[\frac{f_\psi(x)f_\phi(y)}{f_\psi(y)f_\phi(x)} \right], \quad (21)$$

and notice that:

$$\mathbb{E}[|G|^2] = \sum_{x,y} |\phi(x)\psi(y)|^2 |G(x, y)|^2 = 1. \quad (22)$$

From this, the variance becomes:

$$\sigma^2[G] = \mathbb{E}[|G|^2] - \mathbb{E}[|G|]^2 \leq 1. \quad (23)$$

While the main utility of $G(x, y)$ is a simplification of the variance analysis, it could also allow for additional cancellation between $f_\psi(x)$ and $f_\phi(y)$ that could not be exploited in the product of means estimator of Ref. [21].

Robust AR Fidelity Estimator: Despite the fact that the estimator has a finite variance, the random variable $G(x, y)$ is unbounded because it contains a ratio of wavefunctions evaluated at two distinct points and may explode if either $f_\psi(y)$ or $f_\phi(x)$ are close to zero. If we are unlucky enough to obtain such outlier in our empirical estimation, it can significantly skew the statistics. It is therefore desirable to use an estimator that is less affected by outliers – such estimators are often called *robust*. We show how to make the above estimator robust using the median of means amplification.

Theorem 4 (Median of means estimator). *Let k, ℓ be two integers. Define the empirical mean \bar{G} as: $\bar{G} = \sum_i G_i/k$. Compute ℓ such empirical means and use the median as the estimator. Then:*

$$\Pr(|\text{med}[\{\bar{G}\}_{i=1}^\ell] - \mathbb{E}[G]| > \epsilon) \leq \exp\left(-2\ell \left[\frac{1}{2} - \frac{\sigma^2}{k\epsilon^2}\right]^2\right) \quad (24)$$

Proof. We have $\sigma^2(\bar{G}) = \sigma^2[G]/k$. By Chebyshev inequality:

$$p_* =: \Pr\left(|\bar{G} - \mathbb{E}[G]| \geq \epsilon\right) \leq \frac{\sigma^2}{k\epsilon^2} := p. \quad (25)$$

The median of means condition in Eq. 24 is violated if the majority of empirical means violate the Chebyshev condition in Eq. 25. The probability of this happening is at most:

$$\Pr\left(|\text{med}[\{\bar{G}\}_{i=1}^\ell] - \mathbb{E}[G]| > \epsilon\right) = \sum_{k=\lceil \frac{\ell}{2} \rceil}^{\ell} \binom{\ell}{k} p_*^k (1-p_*)^{\ell-k} \quad (26)$$

$$\leq \sum_{k=\lceil \frac{\ell}{2} \rceil}^{\ell} \binom{\ell}{k} p^k (1-p)^{\ell-k}, \quad (27)$$

where the inequality follows by monotonicity. We can bound this by:

$$\sum_{k=\lceil \frac{\ell}{2} \rceil}^{\ell} \binom{\ell}{k} p^k (1-p)^{\ell-k} \leq \exp\left(-2\ell \left(p - \frac{1}{2}\right)^2\right), \quad (28)$$

where we used a tail-bound on the binomial distribution. \square

This technique is commonly used in computer science (see Eg. [15] or a recent review [18]) and was previously used with SQ access/CT states by Tang [25], where it was presented as a standard technique. An alternative way to make the SQ estimators robust was used by Van den Nest in Ref. [27], but it does not work for AR. A corollary of Theorem 4 is a polynomial-time robust estimator of the overlap:

Corollary 2 (Robust AR fidelity estimator). *There is an algorithm that estimates $|\langle \psi | \phi \rangle|^2$ to ϵ -additive error with probability $1 - e^{-n}$ using $64n/\epsilon^2$ AR queries.*

Proof. Use the median of means estimator in Theorem 4. One evaluation of the random variable $G(x, y)$ in Eq. 21 costs two AR queries. Use the median of ℓ empirical means of $G(x, y)$ (each over k evaluations of the random variable) as the estimator of the overlap. Setting $k = 4/\epsilon^2$ and $\ell = 8n$ in Theorem 4 gives:

$$\exp\left(-16n\left(\frac{1}{2} - \frac{\sigma^2}{4}\right)^2\right) < \exp(-n), \quad (29)$$

which follows from:

$$\left(\frac{1}{2} - \frac{\sigma^2}{4}\right)^2 \in \left[\frac{1}{16}, \frac{1}{4}\right]. \quad (30)$$

The overall number of AR queries that achieves this is at most $2\ell k = 64n/\epsilon^2$. \square

We briefly compare this estimator to the CT estimator of Ref. [27]. The algorithms achieve the same goal and their asymptotic query complexity in ϵ and n are the same. The above fidelity estimator however does not require computation of the amplitude, but only amplitude ratios, which is a computationally easier problem as argued in Theorem 2.

Estimating Sparse Observables: Given AR access, there is an algorithm for estimating expectation value of a (hermitian) observable O expanded in the same basis as the wavefunction. This algorithm is well-known in computational physics as local observable estimation:

$$\langle \psi | O | \psi \rangle = \sum_k \sum_j \psi(j) \psi^*(j) O_{jk} \frac{\psi(k)}{\psi(j)} = \sum_j |\psi(j)|^2 \left[\sum_k O_{jk} \frac{f_\psi(k)}{f_\psi(j)} \right]. \quad (31)$$

This can be interpreted as sampling from the Born distribution $|\psi(j)|^2$ and querying the $f_\psi(k)/f_\psi(j)$ AR ratios for every k that appears in:

$$X := \sum_k O_{jk} [f_\psi(k)/f_\psi(j)]. \quad (32)$$

This means that if O has at most $poly(n)$ non-zero entries in each row, we can compute this estimator in polynomial time. As previously, we bound the variance of this estimator. We have that:

$$\sigma^2[X] = \mathbb{E}[|X|^2] - \mathbb{E}[X]^2 \leq \mathbb{E}[|X|^2], \quad (33)$$

and

$$\mathbb{E}[|X|^2] = \sum_j |\psi(j)|^2 \sum_{k, \ell} O_{jk} O_{j\ell}^* \frac{\psi(k)}{\psi(j)} \frac{\psi^*(\ell)}{\psi^*(j)} \quad (34)$$

$$\leq \sum_{j, k, \ell} O_{jk} O_{j\ell}^* \psi(k) \psi^*(\ell) \quad (35)$$

$$= \sum_{k, \ell} O_{\ell k}^2 \psi^*(\ell) \psi(k) = \langle \psi | O^2 | \psi \rangle. \quad (36)$$

We have that $\langle \psi | O^2 | \psi \rangle \leq \sum_{\lambda} \lambda^2 |\langle \lambda | \psi \rangle|^2 \leq \lambda_{\max}^2$, where λ_{\max} is the largest eigenvalue of O . For hermitian matrices, the largest eigenvalue coincides with the operator norm $\|O\| = \lambda_{\max}$, from which we have that $\langle \psi | O^2 | \psi \rangle \leq \|O\|^2$. It follows that: $\sigma[X] \leq \|O\|^2$.

Theorem 5 (Estimating Sparse Observables with AR). *There is an algorithm that estimates $\langle \psi | O | \psi \rangle$ for $\|O\| \leq 1$ to ϵ -additive error with probability at least $1 - e^{-n}$ using at most $32sn/\epsilon^2$, where s is the row-sparsity of O .*

Proof. Use the median of means estimator in Theorem 4. One evaluation of the random variable X costs at most s queries, where s is the row sparsity of O . The rest follows from Corollary 2. \square

3.5 AR and dequantization?

SQ access was studied in dequantization and it is natural to ask if AR can lead to some improvements in that framework. We outline mostly negative results for improvements of Ref. [25] using AR.

Ref. [25], Proposition 4.2. gives an algorithm for estimating the inner product $\langle x, y \rangle$ of two real vectors x and y , using the knowledge of the normalization constant of one of the vectors and error that depends on the normalization factor of both. Assuming $\langle x, y \rangle \geq 0$ (i.e. for non-negative vectors), the AR fidelity algorithm gives a good estimator of the inner product without knowledge of the normalizing factors.

Proposition 4.3. of Ref. [25] crucially depends on computing the rejection sampling filter:

$$r_i = \frac{(Vw)^2(i)}{k \sum_j^k w_j^2 V_{ij}^2} = \left[k \sum_{\ell}^k \left(\sum_j^k \left[\frac{V_{ij} \omega_j}{V_{i\ell} \omega_{\ell}} \right] \right)^{-2} \right]^{-1}. \quad (40)$$

This can be computed with $O(k^2)$ AR ratio queries to V , assuming that the entire matrix has been normalized with the same (possibly unknown) normalization factor. This assumption is most likely too strong to be useful in the context of Tang's algorithm.

¹¹The cancellation of $|\psi(j)|^2$ is problematic for all values for which $\psi(j) = 0$, because the random variable becomes unbounded on the values outside of the support of $|\psi(j)|^2$. The expectation value then depends on the values that the observable O_{jk} takes on samples outside of the support of $|\psi(j)|^2$, which is undesirable. Let Ω be the domain of $|\psi(j)|^2$ and $\Sigma := \{j \in \Omega : |\psi(j)|^2 > 0\}$ be its support. To alleviate the above issue, it is more natural to define:

$$\mathbb{E}[|X|^2]' := \sum_{j \in \Sigma} |\psi(j)|^2 \left| \sum_{k \in \Omega} \frac{\psi_k}{\psi(j)} O_{jk} \right|^2, \quad (37)$$

such that the cancellation of $|\psi(j)|^2$ is well-defined. We can write this as:

$$\mathbb{E}[|X|^2]' = \sum_{j \in \Sigma} \left(\underbrace{\sum_{k \in \Omega} \psi_k O_{jk}}_{\phi_j} \right) \left(\underbrace{\sum_{\ell \in \Omega} \psi_{\ell} O_{j\ell}}_{\phi_j^*} \right)^* = \sum_{j \in \Sigma} |\phi_j|^2 \quad (38)$$

$$\leq \sum_{j \in \Omega} |\phi_j|^2 = \langle \psi | O^2 | \psi \rangle. \quad (39)$$

Irrespectively of the definition of $\mathbb{E}[|X|^2]$, the inequality following Eq. 36 holds. I want to thank Giuseppe Carleo for pointing out this subtlety.

Lastly, the modified version of the FKV algorithm [9] (Algorithm 2 in [25]) crucially relies on sampling from the distribution induced by norms of the matrix columns (normalization factors). There does not seem to be a simple way to circumvent this with the AR access. It would be really interesting to see if some of these limitations could be avoided in different dequantization algorithms.

4 Postselection Gadgets

Here we explore a different way of accessing the wavefunction that can be implemented with an NQS: postselection gadgets. Postselection gadgets are maps between NQs that allow for a different set of conditional queries to the Born distribution encoded in the NQS called *subcube conditional queries* in Ref. [4]. In contrast to PCOND, postselection gadgets cannot be instantiated efficiently for an arbitrary NQS. We use this property to show that there is an NQS with just three nodes that does not encode valid wavefunction and cannot be sampled from – a counterpart to similar results for RBMs [19, 20]. It is possible that the gadgets may have applications beyond this, but the analysis seems to be beyond reach of the techniques known to the author.

4.1 Postselection Gadgets

Let $|\psi_\theta(v|r)|^2 := p_\theta(v|r)$ where $v \in \{-1, 1\}^n$ and $r \in \{-1, 1, \star\}^n$ be the distribution $|\psi_\theta(v)|^2 := p_\theta(v)$, conditioned on the event that for all non- \star bits of r , the output bits of v are fixed to bits of r . A *postselection gadget* is a function that transforms a NQS θ to another NQS $\tilde{\theta}$ (possibly with additional nodes), such that $|\psi_\theta(v|r)|^2 = |\psi_{\tilde{\theta}}(v)|^2$.

Theorem 6. *There is a polynomial-time algorithm that, given $\theta = (a, b, W)$ and $r \in \{1, -1, \star\}^n$, outputs $\tilde{\theta}$, such that $p_{\tilde{\theta}}(v) = p_\theta(v | [r_i \neq \star] \implies [r_i = v_i]) = p_\theta(v|r)$.*

Proof. The following proof uses notation used in Sec. 2.1. For every non- \star bit of r , introduce a hidden node g and attach it to the corresponding visible node. Set the bias on this hidden node to $i(\pi/4)$ and couple it to the visible node with strength $-i(\pi/4)r_i$ (see Fig. 2). This gives:

$$f_{\tilde{\theta}}(v) = \left(\prod_{j:r_j \neq \star} 2 \cosh \left[i \frac{\pi}{4} (r_j - v_j) \right] \right) \times f_\theta(v). \quad (41)$$

Let F be the event¹² that $v_1 = r_1, v_2 = r_2, \dots, v_k = r_k$. Because $2 \cosh(i\pi/2) = 0$, we have that:

$$f_{\tilde{\theta}}(v \in F) = 2^k f_\theta(v), \quad f_{\tilde{\theta}}(v \notin F) = 0. \quad (42)$$

Let $E \subseteq \{-1, +1\}^n$ be an arbitrary event. We have that:

$$\begin{aligned} p_{\tilde{\theta}}(E) &= p_{\tilde{\theta}}(E \cap F) + p_{\tilde{\theta}}(E \cap \neg F) \\ &= p_{\tilde{\theta}}(E|F)p_{\tilde{\theta}}(F) + p_{\tilde{\theta}}(E|\neg F)p_{\tilde{\theta}}(\neg F). \end{aligned} \quad (43)$$

From Eq. 42, it follows that:

$$p_{\tilde{\theta}}(\neg F) = \frac{\sum_{v \in \neg F} |f_{\tilde{\theta}}(v)|^2}{\sum_{v \in \{-1, +1\}^n} |f_{\tilde{\theta}}(v)|^2} = 0. \quad (44)$$

¹²An event $F \subseteq \Omega$ is a subset of domain Ω of the probability distribution.

and

$$p_{\tilde{\theta}}(E|F) = \frac{\sum_{v \in E \cap F} |f_{\tilde{\theta}}(v)|^2}{\sum_{v \in F} |f_{\tilde{\theta}}(v)|^2} = \frac{\sum_{v \in E \cap F} 4^k |f_{\theta}(v)|^2}{\sum_{v \in F} 4^k |f_{\theta}(v)|^2} = p_{\theta}(E|F). \quad (45)$$

Hence, from Eq. 43: $p_{\tilde{\theta}}(E) = p_{\theta}(E|F)$. Thus, the output distribution of the augmented state is *exactly* the conditional of the original distribution. See Fig. 2. \square




Figure 2: To sample an RBM distribution over $\{-1, 1\}^3$ conditioned on $r = 1 \star \star$, we can augment the RBM by an extra hidden node that only couples to the visible node v_1 with a bias $i\pi/4$ and coupling strength $-i\pi/4$.

Note that the function $\theta \mapsto \theta'$ is easy to compute. Postselection gadgets allow for sampling from a subset of conditional distributions of the distribution encoded into the NQS. Additional examples of postselection gadgets can be found in the appendix. The construction almost trivially extends to Deep Boltzmann Machines [10, 5]. While the analysis was inspired by its RBM counterpart by Long and Servedio in Ref. [19], their gadget construction does not straightforwardly extend to RBMs.

4.2 Not all NQS encode valid quantum states

We show that many NQS do not encode valid quantum states. This follows from the fact that the postselection gadget allows postselection on probability zero events. Any NQS can be modified by adding two hidden nodes g_1 and g_2 , as in Sec. 4, that fix the value of the visible node v_1 to 1 and -1 respectively. Eq. 41 then gives $f_{\tilde{\theta}}(v) = 0$, which also means that the encoded “wavefunction” is zero. The resulting NQS then does not encode any wavefunction because it cannot be normalized. The smallest NQS for which this works is one with two hidden nodes and one visible node. Such NQS naturally cannot be sampled from by any algorithm.




Figure 3: A NQS that does not encode a valid quantum state can be constructed by appending two auxiliary nodes, each of which fixes the value of v_1 to $+1$ and -1 respectively.

There does not seem to be any analogous *simple* construction for restricted Boltzmann machines with real valued coefficients. The reason is, as shown earlier, that RBMs cannot encode zeroes in the output probability because each outcome probability is lower-bounded by $2^m \exp(-\|a\|_{\infty} n)$ and the contribution of any “gadget” hidden node to the output RBM probability is a factor of $\cosh(x)$ for some real x , which is lower-bounded by 1. We remark that existence of such “zero-valued” NQS is however not significant for applications in which the network is trained by sampling from output distributions of a sequence of NQSs. Some care must be however taken when encoding quantum states directly.

4.3 Other applications?

The postselection gadget can be in principle used for postselection without retries. However, even when the original NQS θ can be easily sampled from, it may be (and it is very likely in some cases) that the Gibbs sampling algorithm for the modified NQS $\tilde{\theta}$ won't be efficient anymore. Still, one toy example that may suggest that this could be more efficient than resampling (but that also seems useless) is the case in which the encoded distribution is a product distribution - fixing output bits of such distribution using gadgets will not lead to any slowdown.

In the case where the sampling algorithm remains efficient for the appropriate sequence of conditionals, one can also obtain a *very crude* multiplicative approximation of the normalizing factor of the NQS wavefunction, essentially by retracing the RBM algorithm presented in Ref. [19]. The question of characterizing what conditional gadgets allow for efficient sampling however remains widely open.

Lastly, Ref. [16] used a similar gadget construction to simulate universal random circuits with NQs. The postselection gadgets from Thm. 6 can be seen as extension of their result. See Appendix A for additional postselection gadgets and Appendix B for encoding of Pauli gates.

5 Discussion

We studied the access model offered by neural network quantum states (NQS), which, along with connection to previous results from conditional distribution testing, motivated the definition of amplitude ratio (AR) access model. We related AR to sample and query (SQ) access and showed that it retains some of the simulation capabilities of SQ. We gave some evidence that AR may be weaker than SQ and showed that existing results in distribution testing imply that AR is stronger than sample access. We then considered alternative access to the NQS wavefunctions by means of *subcube conditional* queries and showed that even small NQS may not encode valid distributions. Our work left several questions open:

- It would be interesting to further explore the connections between AR and dequantization and understand if it is possible to meaningfully relax the SQ normalization requirement.
- Both definitions of CT and AR states assumed existence of a classical randomized sampler for the Born distribution. One may thus ask if there are any nontrivial states in the *quantum-classical generalization* of CT and AR states, where we can efficiently sample from using quantum algorithm, but still classically estimate the ratios or amplitudes. Such states may be useful for construction of quantum algorithms based on conditional property testing results in Ref. [3, 7]. It might be for example interesting to understand how this plays with some of the known supremacy results in which approximating a target amplitude is known to be $\#P$ -hard, yet there is an efficient quantum algorithm that samples the output [1]. It may be that, while the amplitudes are hard to compute, approximating their ratios remains tractable.
- It would be also interesting to understand, perhaps numerically, if and for what problems could the NQS postselection gadgets provide advantage when compared to postselection by resampling in some of the applications of NQS.

6 Acknowledgements

I want to sincerely thank Ashley Montanaro and Noah Linden for their help. I also want to thank Giuseppe Carleo, Sridhar Arunachalam, Sergii Strelchuk, James Stokes and Juani Bermejo-Vega for their suggestions and discussion.

References

- [1] Scott Aaronson and Alex Arkhipov. “The Computational Complexity of Linear Optics”. In: *STOC ’11*. San Jose, California, USA: Association for Computing Machinery, 2011. DOI: 10.1145/1993636.1993682.
- [2] Clement Cannone. Personal communication. Jan. 10, 2021.
- [3] Clément L. Canonne, Dana Ron, and Rocco A. Servedio. “Testing Probability Distributions using Conditional Samples”. In: *SIAM Journal on Computing* 44.3 (2015), pp. 540–616. DOI: 10.1137/130945508.
- [4] Clement L. Canonne et al. “Random Restrictions of High Dimensional Distributions and Uniformity Testing with Subcube Conditioning”. In: *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA ’21. Virtual Event, Virginia: Society for Industrial and Applied Mathematics, 2021, pp. 321–336.
- [5] Giuseppe Carleo, Yusuke Nomura, and Masatoshi Imada. “Constructing exact representations of quantum many-body systems with deep neural networks”. In: *Nature Communications* 9.1 (Dec. 2018), p. 5322. DOI: 10.1038/s41467-018-07520-3.
- [6] Giuseppe Carleo and Matthias Troyer. “Solving the quantum many-body problem with artificial neural networks”. In: *Science* 355.6325 (Feb. 2017), pp. 602–606. DOI: 10.1126/science.aag2302.
- [7] Sourav Chakraborty et al. “On the Power of Conditional Samples in Distribution Testing”. In: *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*. ITCS ’13. Berkeley, California, USA: Association for Computing Machinery, 2013, pp. 561–580. DOI: 10.1145/2422436.2422497.
- [8] Martin Dyer, Alan Frieze, and Ravi Kannan. “A Random Polynomial-Time Algorithm for Approximating the Volume of Convex Bodies”. In: *J. ACM* 38.1 (Jan. 1991), pp. 1–17. DOI: 10.1145/102782.102783.
- [9] Alan Frieze, Ravi Kannan, and Santosh Vempala. “Fast Monte-Carlo Algorithms for Finding Low-Rank Approximations”. In: *J. ACM* 51.6 (Nov. 2004), pp. 1025–1041. DOI: 10.1145/1039488.1039494.
- [10] Xun Gao and Lu-Ming Duan. “Efficient representation of quantum many-body states with deep neural networks”. In: *Nature Communications* 8.1 (2017), p. 662. DOI: 10.1038/s41467-017-00705-2.
- [11] Vojtech Havlicek and Sergii Strelchuk. “Quantum Schur Sampling Circuits can be Strongly Simulated”. In: *Phys. Rev. Lett.* 121 (6 Aug. 2018), p. 060505. DOI: 10.1103/PhysRevLett.121.060505.
- [12] Geoffrey E. Hinton. “Training Products of Experts by Minimizing Contrastive Divergence”. In: *Neural Computation* 14.8 (Aug. 2002), pp. 1771–1800. DOI: 10.1162/089976602760128018.
- [13] Mark Huber. “Approximation algorithms for the normalizing constant of Gibbs distributions”. In: *The Annals of Applied Probability* 25.2 (Apr. 2015). DOI: 10.1214/14-aap1015.

- [14] Mark Jerrum. “Random Generation of Combinatorial Structures from a Uniform Distribution (Extended Abstract)”. In: *Proceedings of the 12th Colloquium on Automata, Languages and Programming*. Berlin, Heidelberg: Springer-Verlag, 1985, pp. 290–299.
- [15] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. “Random generation of combinatorial structures from a uniform distribution”. In: *Theoretical Computer Science* 43 (1986), pp. 169–188. DOI: [https://doi.org/10.1016/0304-3975\(86\)90174-X](https://doi.org/10.1016/0304-3975(86)90174-X).
- [16] Bjarni Jónsson, Bela Bauer, and Giuseppe Carleo. “Neural-network states for the classical simulation of quantum computing”. In: *arXiv e-prints*, arXiv:1808.05232 (Aug. 2018), arXiv:1808.05232. DOI: 10.48550/ARXIV.1808.05232.
- [17] Richard M Karp, Michael Luby, and Neal Madras. “Monte-Carlo approximation algorithms for enumeration problems”. In: *Journal of Algorithms* 10.3 (1989), pp. 429–448. DOI: [https://doi.org/10.1016/0196-6774\(89\)90038-2](https://doi.org/10.1016/0196-6774(89)90038-2).
- [18] Matthieu Lerasle. “Lecture Notes: Selected topics on robust statistical learning theory”. In: *arXiv e-prints*, arXiv:1908.10761 (Aug. 2019), arXiv:1908.10761. DOI: 10.48550/ARXIV.1908.10761.
- [19] Philip M. Long and Rocco A. Servedio. “Restricted Boltzmann Machines Are Hard to Approximately Evaluate or Simulate”. In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. ICML’10. Haifa, Israel: Omnipress, 2010, pp. 703–710.
- [20] James Martens et al. “On the Representational Efficiency of Restricted Boltzmann Machines”. In: *Advances in Neural Information Processing Systems 26*. Ed. by C. J. C. Burges et al. Curran Associates, Inc., 2013, pp. 2877–2885.
- [21] Matija Medvidović and Giuseppe Carleo. “Classical variational simulation of the Quantum Approximate Optimization Algorithm”. In: *npj Quantum Information* 7, 101 (Jan. 2021), p. 101. DOI: 10.1038/s41534-021-00440-z.
- [22] Imdad S. B. Sardharwalla, Sergii Strelchuk, and Richard Jozsa. “Quantum Conditional Query Complexity”. In: *Quantum Info. Comput.* 17.7–8 (June 2017), pp. 541–567.
- [23] P. Smolensky. “Information Processing in Dynamical Systems: Foundations of Harmony Theory”. In: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations*. Cambridge, MA, USA: MIT Press, 1986, pp. 194–281.
- [24] Daniel Štefankovič, Santosh Vempala, and Eric Vigoda. “Adaptive Simulated Annealing: A near-Optimal Connection between Sampling and Counting”. In: *J. ACM* 56.3 (May 2009). DOI: 10.1145/1516512.1516520.
- [25] Ewin Tang. “A Quantum-Inspired Classical Algorithm for Recommendation Systems”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 217–228. DOI: 10.1145/3313276.3316310.
- [26] L.G. Valiant. “The complexity of computing the permanent”. In: *Theoretical Computer Science* 8.2 (1979), pp. 189–201. DOI: [https://doi.org/10.1016/0304-3975\(79\)90044-6](https://doi.org/10.1016/0304-3975(79)90044-6).
- [27] Maarten Van Den Nest. “Simulating Quantum Computers with Probabilistic Methods”. In: *Quantum Info. Comput.* 11.9–10 (Sept. 2011), pp. 784–812.

A Additional Postselection Gadgets

A.1 Hamming weight




Figure 4: To sample an RBM distribution conditioned on $\sum_i v_i = k$ for some $k \in [n]$, we use a two step process. We augment the RBM by an extra node coupled to all visible nodes with weight $W_k = \frac{i\pi}{k}$ and bias $b = -i\frac{\pi}{4}$. We do this all $n - 1$ times, excluding only the sector we are interested in.

Start with n visible nodes. Add a single hidden node g_k coupled to all of the visible nodes with weight $W = i\pi/n$ and $b = i\pi/2 - ik\pi/n$. This encodes:

$$f_k(v) = 2 \cosh \left[\sum_i^n W v_i + b \right] = 2 \sin \left[\frac{\pi}{n} \left(\sum_j^n v_j - k \right) \right]. \quad (46)$$

Notice that $f_k(v) = 0$ if $\sum_j v_j = k$. By adding $(n - 1)$ such nodes, we have that $F_k(v) := \prod_{j=1}^n f_j(v)/f_k(v)$. This function satisfies:

$$F_k(v) := \begin{cases} 0 & \text{if } \sum_i v_i \neq k, \\ 2^{n-1} \sin \left[\frac{\pi}{n} (|v| - 1) \right] \sin \left[\frac{\pi}{n} (|v| - 2) \right] \dots & \end{cases} \quad (47)$$

By modifying the visible biases, this can be normalized such that:

$$F'_k(v) := \begin{cases} 0 & \text{if } \sum_i v_i \neq k. \\ c & \text{otherwise.} \end{cases} \quad (48)$$

for any constant c .

A.2 Parity

Martens et. al showed in Ref. [20] how to approximately represent parity (or any symmetric function) with an RBM of $m = n^2 + 1$ hidden nodes. Here's an explicit construction that represents parity with only two hidden nodes of NQS:




Figure 5: Fixing parity with two hidden nodes.

Choose $W = i\pi/2$ for all edges from hidden to visible nodes and $\text{Im}(b) = 0$ for all hidden biases. This gives:

$$f_\theta(v) := 4 \sin^2 \left(\frac{\pi}{2} \sum_i^n v_i \right) = \begin{cases} 0 & \text{if } \sum_i^n v_i \text{ is even.} \\ 4 & \text{if } \sum_i^n v_i \text{ is odd.} \end{cases} \quad (49)$$

This represents a distribution that is uniform over all strings with odd parity and zero on the even parity ones and encodes parity distribution by only two hidden nodes. This gadget can be used to fix parity of NQS samples. Similar gadget has been used in Ref. [5].

B Pauli Gadgets

We can use the same construction to postselection gadgets to simulate action of Pauli unitaries on the NQS wavefunction. This has application to simulation of quantum circuits in Ref. [16]. The Pauli Z gadget appeared in [16], but the others seem useful as a reference, so I will leave it here. Along with the CZ gate and Hadamard approximation of [16], the gadgets below give encoding of X, Y, Z, H, CZ universal gateset, useful for simulating quantum circuits with NQS.

Pauli Z: The action of a Pauli Z_1 operator applied to the first qubit on the NQS wavefunction can be understood as the following function on the amplitudes (recall that we chose $v \in \{-1, +1\}^n$):

$$f_\theta(v) \mapsto i^{(1+v_1)} f_\theta(v). \quad (50)$$

As an example, consider just two qubits, where we expect Z_1 to act as follows:

$$\begin{aligned} Z_1[f_\theta(++)] &= f_\theta(++), & Z_1[f_\theta(-+)] &= -f_\theta(-+), \\ Z_1[f_\theta(-+)] &= f_\theta(-+), & Z_1[f_\theta(++)] &= -f_\theta(++). \end{aligned} \quad (51)$$

Let $Z_s = \bigotimes_{k=1}^n Z^{s_k}$ for $s \in \{0, 1\}^n$. A gadget that applies this operator is defined as follows:

$$f_{Z_s}(v) = \prod_{g \in [n]} \left[2 \cosh \left(i \frac{\pi}{2} (1 - v_g) \right) \right]^{s_k} \times f(v). \quad (52)$$

It can be constructed similarly to the postselection gadget by adding a hidden node to all nodes that should have Pauli gate applied to them, fixing their bias to $b = (i\frac{\pi}{2})$ and interaction strength to $-i\frac{\pi}{2}$.

Pauli X: The action of a Pauli X_1 operator applied to the first qubit on the NQS wavefunction can be understood as a permutation on the amplitudes. As an example, consider the action of X_1 :

$$\begin{aligned} X_1[f_\theta(++)] &= f_\theta(-+), & X_1[f_\theta(-+)] &= f_\theta(++), \\ X_1[f_\theta(-+)] &= f_\theta(++), & X_1[f_\theta(++)] &= f_\theta(-+). \end{aligned} \quad (53)$$

This transformation can be implemented by a map on the parameters, $\theta \mapsto \theta$, such that $a_1 \mapsto -a_1$ and $W_{1i} \mapsto -W_{1i}$ for all i .

Pauli Y: The gadget that applies this operator to the first qubit is defined as the action on the parameters as $a_1 \mapsto -a_1$ and $W_{1i} \mapsto -W_{1i} \forall i$ and an additional phase factor, similarly to **Pauli Z** gadget.

Pauli tensors: Combination of the above Pauli gadgets can be used to implement arbitrary Pauli tensor strings by adding up to n hidden nodes.