

QuCumber: wavefunction reconstruction with neural networks

Matthew J. S. Beach^{1,2}, Isaac De Vlugt², Anna Golubeva^{1,2}, Patrick Huembeli^{1,3}, Bohdan Kulchytskyi^{1,2}, Xiuzhe Luo², Roger G. Melko^{1,2*}, Ejaaz Merali² and Giacomo Torlai^{1,2,4}

¹ Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada

² Department of Physics and Astronomy, University of Waterloo, Ontario N2L 3G1, Canada

³ ICFO-Institut de Ciències Fòniques, Barcelona Institute of Science and Technology, 08860 Castelldefels (Barcelona), Spain

⁴ Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

* rgmelko@uwaterloo.ca

Abstract

As we enter a new era of quantum technology, it is increasingly important to develop methods to aid in the accurate preparation of quantum states for a variety of materials, matter, and devices. Computational techniques can be used to reconstruct a state from data, however the growing number of qubits demands ongoing algorithmic advances in order to keep pace with experiments. In this paper, we present an open-source software package called QuCumber that uses machine learning to reconstruct a quantum state consistent with a set of projective measurements. QuCumber uses a restricted Boltzmann machine to efficiently represent the quantum wavefunction for a large number of qubits. New measurements can be generated from the machine to obtain physical observables not easily accessible from the original data.



Copyright M. J. S. Beach *et al.*

This work is licensed under the Creative Commons Attribution 4.0 International License.

Published by the SciPost Foundation.

Received 02-01-2019

Accepted 05-06-2019

Published 16-06-2019

doi:[10.21468/SciPostPhys.7.1.009](https://doi.org/10.21468/SciPostPhys.7.1.009)



Check for updates

Contents

1	Introduction	1
2	Positive wavefunctions	2
2.1	Setup	3
2.2	Training	4
2.3	Reconstruction of physical observables	5
2.3.1	Diagonal observables	5
2.3.2	Off-diagonal observables	6
2.3.3	Entanglement entropy	8
3	Complex wavefunctions	9
3.1	Setup	11
3.2	Training	12

4 Conclusion	13
A Glossary	14
References	16

1 Introduction

Current advances in fabricating quantum technologies, as well as in reliable control of synthetic quantum matter, are leading to a new era of quantum hardware where highly pure quantum states are routinely prepared in laboratories. With the growing number of controlled quantum degrees of freedom, such as superconducting qubits, trapped ions, and ultracold atoms [1–4], reliable and scalable classical algorithms are required for the analysis and verification of experimentally prepared quantum states. Efficient algorithms can aid in extracting physical observables otherwise inaccessible from experimental measurements, as well as in identifying sources of noise to provide direct feedback for improving experimental hardware. However, traditional approaches for reconstructing unknown quantum states from a set of measurements, such as quantum state tomography, often suffer the exponential overhead that is typical of quantum many-body systems.

Recently, an alternative path to quantum state reconstruction was put forward, based on modern machine learning (ML) techniques [5–10]. The most common approach relies on a powerful generative model called a *restricted Boltzmann machine* (RBM) [11], a stochastic neural network with two layers of binary units. A visible layer \mathbf{v} describes the physical degrees of freedom, while a hidden layer \mathbf{h} is used to capture high-order correlations between the visible units. Given a set of neural network parameters $\boldsymbol{\lambda}$, the RBM defines a probabilistic model described by the parametric distribution $p_{\boldsymbol{\lambda}}(\mathbf{v})$. RBMs have been widely used in the ML community for the pre-training of deep neural networks [12], for compressing high-dimensional data into lower-dimensional representations [13], and more [14]. More recently, RBMs have been adopted by the physics community in the context of representing both classical and quantum many-body states [15, 16]. They are currently being investigated for their representational power [17–19], their relationship with tensor networks and the renormalization group [20–24], and in other contexts in quantum many-body physics [25–27].

In this post, we present QuCumber: a *quantum calculator used for many-body eigenstate reconstruction*. QuCumber is an open-source Python package that implements neural-network quantum state reconstruction of many-body wavefunctions from projective measurement data. Examples of data to which QuCumber could be applied might be magnetic spin projections, orbital occupation number, polarization of photons, or the logical state of qubits. Given a training set of such measurements, QuCumber discovers the most likely compatible quantum state by finding the optimal set of parameters $\boldsymbol{\lambda}$ of an RBM. A properly trained RBM is an approximation of the unknown quantum state underlying the data. It can be used to calculate various physical observables of interest, including measurements that may not be possible in the original experiment.

This post is organized as follows. In Section 2, we introduce the reconstruction technique for the case where all coefficients of the wavefunction are real and positive. We discuss the required format for input data, as well as training of the RBM and the reconstruction of typical observables. In Section 3, we consider the more general case of a complex-valued wavefunction. We illustrate a general strategy to extract the phase structure from data by performing

appropriate unitary rotations on the state before measurements. We then demonstrate a practical reconstruction of an entangled state of two qubits. Note, the detailed theory underlying the reconstruction methods used by QuCumber can be found in the original references [5, 6] and a recent review [28]. A list of useful terms and equations can be found at the end of this post in the Glossary.

2 Positive wavefunctions

We begin by presenting the application of QuCumber to reconstruct many-body quantum states described by wavefunctions $|\Psi\rangle$ with positive coefficients $\Psi(\mathbf{x}) = \langle \mathbf{x} | \Psi \rangle \geq 0$, where $|\mathbf{x}\rangle = |x_1, \dots, x_N\rangle$ is a reference basis for the Hilbert space of N quantum degrees of freedom. The neural-network quantum state reconstruction requires raw data $\mathcal{D} = (\mathbf{x}_1, \mathbf{x}_2, \dots)$ generated through projective measurements of the state $|\Psi\rangle$ in the reference basis. These measurements adhere to the probability distribution given by the Born rule, $P(\mathbf{x}) = |\Psi(\mathbf{x})|^2$. Since the wavefunction is strictly positive, the quantum state is completely characterized by the measurement distribution, i.e. $\Psi(\mathbf{x}) = \sqrt{P(\mathbf{x})}$.

The positivity of the wavefunction allows a simple and natural connection between quantum states and classical probabilistic models. QuCumber employs the probability distribution $p_\lambda(\mathbf{x})$ of an RBM (see Eq. 24 of the Glossary) to approximate the distribution $P(\mathbf{x})$ underlying the measurement data. Using contrastive divergence (CD) [12], QuCumber trains the RBM to discover an optimal set of parameters λ that minimize the Kullback-Leibler (KL) divergence between the two distributions (see Eq. 23). Upon successful training ($p_\lambda(\mathbf{x}) \sim P(\mathbf{x})$), we obtain an approximate representation of the target quantum state,

$$\psi_\lambda(\mathbf{x}) \equiv \sqrt{p_\lambda(\mathbf{x})} \simeq \Psi(\mathbf{x}). \quad (1)$$

Note, the precise mathematical form of the marginal distribution $p_\lambda(\mathbf{x})$ defined in terms of an effective energy over the parameters of the RBM is defined in the Glossary.

In the following, we demonstrate the application of QuCumber for the reconstruction of the ground-state wavefunction of the one-dimensional transverse-field Ising model (TFIM). The Hamiltonian is

$$\hat{H} = -J \sum_i \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z - h \sum_i \hat{\sigma}_i^x, \quad (2)$$

where $\hat{\sigma}_i^{x/z}$ are spin-1/2 Pauli operators acting on site i , and we assume open boundary conditions. For this example, we consider a chain with $N = 10$ spins at the quantum critical point $J = h = 1$.

2.1 Setup

Given the small size of the system, the ground state $|\Psi\rangle$ can be found with exact diagonalization. The training dataset \mathcal{D} is generated by sampling the distribution $P(\boldsymbol{\sigma}^z) = |\Psi(\boldsymbol{\sigma}^z)|^2$, obtaining a sequence of $N_S = 10^5$ independent spin projections in the reference basis $\mathbf{x} = \boldsymbol{\sigma}^z$.¹ Each data point in \mathcal{D} consists of an array $\boldsymbol{\sigma}_j^z = (\sigma_1^z, \dots, \sigma_N^z)$ with shape $(N,)$ and should be passed to QuCumber as a numpy array or torch tensor. For example, $\boldsymbol{\sigma}_j^z = \text{np.array}([1, 0, 1, 1, 0, 1, 0, 0, 0, 1])$, where we use $\sigma_j^z = 0, 1$ to represent a spin-down and spin-up state respectively. Therefore, the entire input data set is contained in an array with shape (N_S, N) .

¹ The training dataset can be downloaded from https://github.com/PIQuIL/QuCumber/blob/master/examples/Tutorial1_TrainPosRealWaveFunction/tfim1d_data.txt

Aside from the training data, QuCumber also allows us to import an exact wavefunction. This can be useful for monitoring the quality of the reconstruction during training. In our example, we evaluate the fidelity between the reconstructed state $\psi_\lambda(\mathbf{x})$ and the exact wavefunction $\Psi(\mathbf{x})$. The training dataset, `train_data`, and the exact ground state, `true_psi`, are loaded with the data loading utility as follows:

```
import qucumber.utils.data as data

train_path = "tfim1d_data.txt"
psi_path = "tfim1d_psi.txt"
train_data, true_psi = data.load_data(train_path, psi_path)
```

If `psi_path` is not provided, QuCumber will load only the training data.

Next, we initialize an RBM quantum state $\psi_\lambda(\mathbf{x})$ with random weights and zero biases using the constructor `PositiveWaveFunction`:

```
from qucumber.nn_states import PositiveWaveFunction

state = PositiveWaveFunction(num_visible=10, num_hidden=10)
```

The number of visible units (`num_visible`) must be equal to the number of physical spins N , while the number of hidden units (`num_hidden`) can be adjusted to systematically increase the representational power of the RBM.

The quality of the reconstruction will depend on the structure underlying the specific quantum state and the ratio of visible to hidden units, $\alpha = \text{num_hidden}/\text{num_visible}$. In practice, we find that $\alpha = 1$ often leads to good approximations of positive wavefunctions [6]. However, in the general case, the value of α required for a given wavefunction should be explored and adjusted by the user.

2.2 Training

Once an appropriate representation of the quantum state has been defined, QuCumber trains the RBM through the function `PositiveWaveFunction.fit`. Several input parameters need to be provided aside from the training dataset (`train_data`). These include the number of training iterations (`epochs`), the number of samples used for the positive/negative phase of CD (`pos_batch_size/neg_batch_size`), the learning rate (`lr`) and the number of sampling steps in the negative phase of CD (`k`). The last argument (`callbacks`) allows the user to pass a set of additional functions to be evaluated during training.

As an example of a callback, we show the `MetricEvaluator`, which evaluates a function `log_every` epochs during training. Given the small system size and the knowledge of the true ground state, we can evaluate the fidelity between the RBM state and the true ground-state wavefunction (`true_psi`). Similarly, we can calculate the KL divergence between the RBM distribution $p_\lambda(\mathbf{x})$, and the data distribution $P(\mathbf{x})$, which should approach zero for a properly trained RBM. For the current example, we monitor the fidelity and KL divergence (defined in `qucumber.utils.training_statistics`):

```
from qucumber.callbacks import MetricEvaluator
import qucumber.utils.training_statistics as ts

log_every = 10
space = state.generate_hilbert_space(10)
callbacks = [
    MetricEvaluator(
        log_every,
        {"Fidelity": ts.fidelity, "KL": ts.KL},
```

```

        target_psi=true_psi ,
        space=space ,
        verbose=True ,
    )
]

```

With `verbose=True`, the program will print the epoch number and all callbacks every `log_every` epochs. For the current example, we monitor the fidelity and KL divergence (see Glossary). Note that the KL divergence is only tractable for small systems. The `MetricEvaluator` will compute the KL exactly when provided with a list of all states in the Hilbert space. For convenience these can be generated with `space = state.generate_hilbert_space(10)`.

Now that the metrics to monitor during training have been chosen, we can invoke the optimization with the fit function of `PositiveWaveFunction`.

```

state.fit(
    train_data ,
    epochs=500 ,
    pos_batch_size=100 ,
    neg_batch_size=100 ,
    lr=0.01 ,
    k=5 ,
    callbacks=callbacks ,
)

```

Figure 1 shows the convergence of the fidelity and KL divergence during training. The convergence time will, in general, depend on the choice of hyperparameters. Finally, the network parameters λ , together with the `MetricEvaluator`'s data, can be saved (or loaded) to a file:

```

state.save(
    "filename.pt",
    metadata={
        "fidelity": callbacks[0].Fidelity,
        "KL": callbacks[0].KL,
    },
)
state.load("filename.pt")

```

With this we have demonstrated the most basic aspects of QuCumber regarding training a model and verifying its accuracy. We note that in this example the evaluation utilized the knowledge of the exact ground state and the calculation of the KL divergence, which we reemphasize is tractable only for small system sizes. However, we point out that QuCumber is capable of carrying out the reconstruction of much larger systems. In such cases, users must rely on other estimators to evaluate the training, such as expectation values of physical observables (magnetization, energy, etc). In the following, we show how to compute diagonal and off-diagonal observables in QuCumber.

2.3 Reconstruction of physical observables

In this section, we discuss how to calculate the average value of a generic physical observable \hat{O} from a trained RBM. We start with the case of observables that are diagonal in the reference basis where the RBM was trained. We then discuss the more general cases of off-diagonal observables and entanglement entropies.




Figure 1: The fidelity (left) and the KL divergence (right) during training for the reconstruction of the ground state of the one-dimensional TFIM.

2.3.1 Diagonal observables

We begin by considering an observable with only diagonal matrix elements, $\langle \sigma | \hat{O} | \sigma' \rangle = \mathcal{O}_\sigma \delta_{\sigma\sigma'}$, where for convenience we denote the reference basis $\mathbf{x} = \sigma^z$ as σ unless otherwise stated. The expectation value of \hat{O} is given by

$$\langle \hat{O} \rangle = \frac{1}{\sum_\sigma |\psi_\lambda(\sigma)|^2} \sum_\sigma \mathcal{O}_\sigma |\psi_\lambda(\sigma)|^2. \quad (3)$$

The expectation value can be approximated by a Monte Carlo estimator,

$$\langle \hat{O} \rangle \approx \frac{1}{N_{\text{MC}}} \sum_{k=1}^{N_{\text{MC}}} \mathcal{O}_{\sigma_k}, \quad (4)$$

where the spin configurations σ_k are sampled from the RBM distribution $p_\lambda(\sigma)$. This process is particularly efficient given the bipartite structure of the network which allows the use of block Gibbs sampling.

A simple example for the TFIM is the average longitudinal magnetization per spin, $\langle \hat{\sigma}^z \rangle = \sum_j \langle \hat{\sigma}_j^z \rangle / N$, which can be calculated directly on the spin configuration sampled by the RBM (i.e., the state of the visible layer). The visible samples are obtained with the sample function of the RBM state object:

```
samples = state.sample(num_samples=1000, k=10)
```

which takes the total number of samples (`num_samples`) and the number of block Gibbs steps (`k`) as input. Once these samples are obtained, the magnetization can be calculated simply as

```
magnetization = samples.mul(2.0).sub(1.0).mean()
```

where we converted the binary samples of the RBM back into ± 1 spins before taking the mean.

2.3.2 Off-diagonal observables

We turn now to the case of off-diagonal observables, where the expectation value assumes the following form

$$\langle \hat{O} \rangle = \frac{1}{\sum_\sigma |\psi_\lambda(\sigma)|^2} \sum_{\sigma\sigma'} \psi_\lambda^*(\sigma) \psi_\lambda(\sigma') \mathcal{O}_{\sigma\sigma'}. \quad (5)$$




Figure 2: Reconstruction of the magnetic observables for the TFIM chain with $N = 10$ spins. We show the average longitudinal (left) and transverse (right) magnetization per site obtained by sampling from a trained RBM. The dashed line denotes the results from exact diagonalization.

This expression can once again be approximated with a Monte Carlo estimator,

$$\langle \hat{O} \rangle \approx \frac{1}{N_{\text{MC}}} \sum_{k=1}^{N_{\text{MC}}} \mathcal{O}_{\sigma_k}^{[L]}, \quad (6)$$

of the so-called *local estimator* of the observable:

$$\mathcal{O}_{\sigma_k}^{[L]} = \sum_{\sigma'} \frac{\psi_{\lambda}(\sigma')}{\psi_{\lambda}(\sigma_k)} \mathcal{O}_{\sigma_k \sigma'}. \quad (7)$$

As long as the matrix representation $\mathcal{O}_{\sigma\sigma'}$ is sufficiently sparse in the reference basis, the summation can be evaluated efficiently.

As an example, we consider the specific case of the transverse magnetization for the j -th spin, $\langle \hat{\sigma}_j^x \rangle$, with matrix elements

$$\langle \sigma | \hat{\sigma}_j^x | \sigma' \rangle = \delta_{\sigma'_j, 1-\sigma_j} \prod_{i \neq j} \delta_{\sigma'_i, \sigma_j}. \quad (8)$$

Therefore, the expectation values reduces to the Monte Carlo average of the local observable

$$(\sigma_j^x)^{[L]} = \frac{\psi_{\lambda}(\sigma_1, \dots, 1-\sigma_j, \dots, \sigma_N)}{\psi_{\lambda}(\sigma_1, \dots, \sigma_j, \dots, \sigma_N)}. \quad (9)$$

evaluated on spin configurations σ_k sampled from the RBM distribution $p_{\lambda}(\sigma)$.

QuCumber provides an interface for sampling off-diagonal observables in the `ObservableBase` class. Thorough examples are available in the tutorial section in the [documentation](#).² As an example, σ^x can be written as an observable class with

```
import torch
from qcumber.utils import cplx
from qcumber.observables import ObservableBase

class SigmaX(ObservableBase):
    def apply(self, nn_state, samples):
```

²The observables tutorial is available at https://qcumber.readthedocs.io/en/stable/_examples/Tutorial3_DataGeneration_CalculateObservables/tutorial_sampling_observables.html

```

psi = nn_state.psi(samples)
psi_ratio_sum = torch.zeros_like(psi)

for i in range(samples.shape[-1]): # sum over sites
    flip_spin(i, samples) # flip the spin at site i
    # add ratio psi_(-i) / psi to the running sum
    psi_flip = nn_state.psi(samples)
    psi_ratio = cplx.elementwise_division(psi_flip, psi)
    psi_ratio_sum.add_(psi_ratio)
    flip_spin(i, samples) # flip it back

# take real part and divide by number of spins
return psi_ratio_sum[0].div_(samples.shape[-1])

```

The value of the observable can be estimated from a set of samples with:

```
SigmaX().statistics_from_samples(state, samples)
```

which produces a dictionary containing the mean, variance, and standard error of the observable. Similarly, the user can define other observables like the energy.

The reconstruction of two magnetic observables for the TFIM is shown in Fig. 2, where a different RBM was trained for each value of the transverse field h . In the left plot, we show the average longitudinal magnetization per site, which can be calculated directly from the configurations sampled by the RBM. In the right plot, we show the off-diagonal observable of transverse magnetization. In both cases, QuCumber successfully discovers an optimal set of parameters λ that accurately approximate the ground-state wavefunction underlying the data.

2.3.3 Entanglement entropy

A quantity of significant interest in quantum many-body systems is the degree of entanglement between a subregion A and its complement \bar{A} . Numerically, measurement of bipartite entanglement entropy is commonly accessed through the computation of the second Rényi entropy $S_2 = -\ln \text{Tr}(\rho_A^2)$. When one has access to a pure state wavefunction $\psi_\lambda(\mathbf{x})$, Rényi entropies can be calculated as an expectation value of the “Swap” operator [29],

$$S_2 = -\ln \langle \widehat{\text{Swap}}_A \rangle. \quad (10)$$

It is essentially an off-diagonal observable that acts on an extended product space consisting of two independent copies of the wavefunction, $\psi_\lambda(\mathbf{x}) \otimes \psi_\lambda(\mathbf{x})$, referred to as “replicas”. As the name suggests, the action of the Swap operator is to swap the spin configurations in region A between the replicas,

$$\widehat{\text{Swap}}_A |\sigma_A, \sigma_{\bar{A}}\rangle_1 \otimes |\sigma'_A, \sigma'_{\bar{A}}\rangle_2 = |\sigma'_A, \sigma'_{\bar{A}}\rangle_1 \otimes |\sigma_A, \sigma_{\bar{A}}\rangle_2. \quad (11)$$

Here the subscript of the ket indicates the replica index, while the two labels inside a ket, such as $\sigma_A, \sigma_{\bar{A}}$, describe the spins configurations within the subregion and its complement.

In QuCumber, the Swap operator is implemented as a routine within the entanglement observable unit,

```

def swap(s1, s2, A):
    _s = s1[:, A].clone()
    s1[:, A] = s2[:, A]
    s2[:, A] = _s
    return s1, s2

```


where s_1 and s_2 are batches of samples produced from each replica, and A is a list containing the indices of the sites in subregion A . While ideally those samples should be entirely independent, in order to save computational costs, QuCumber just splits a given batch into two equal parts and treats them as if they were independent samples. This is implemented within the SWAP observable,

```
class SWAP(ObservableBase):
    def __init__(self, A):
        self.A = A

    def apply(self, nn_state, samples):
        _ns = samples.shape[0] // 2
        samples1 = samples[:_ns, :]
        samples2 = samples[_ns : _ns * 2, :]

        psi_ket1 = nn_state.psi(samples1)
        psi_ket2 = nn_state.psi(samples2)

        psi_ket = cplx.elementwise_mult(psi_ket1, psi_ket2)
        psi_ket_star = cplx.conjugate(psi_ket)

        samples1_, samples2_ = swap(
            samples1, samples2, self.A
        )
        psi_bra1 = nn_state.psi(samples1_)
        psi_bra2 = nn_state.psi(samples2_)

        psi_bra = cplx.elementwise_mult(psi_bra1, psi_bra2)
        psi_bra_star = cplx.conjugate(psi_bra)

        return cplx.real(
            cplx.elementwise_division(
                psi_bra_star, psi_ket_star
            )
        )
```

Note the similarity in the implementation to that for the transverse magnetization observable from the last section, once the amplitude of a sample is substituted with the product of amplitudes drawn from each replica.

Using this observable, we can estimate the Rényi entropy of the region containing the first 5 sites in the chain using Eq. 10,

```
A = [0, 1, 2, 3, 4]
swap_ = SWAP(A)
swap_stats = swap_.statistics_from_samples(state, samples)
S_2 = -np.log(swap_stats["mean"])
```

We apply this measurement procedure to a TFIM chain with results shown in Fig. 3. As was the case with the magnetization observables, the trained RBM gives a good approximation to the second Rényi entropy for different subregion A sizes. Being a basis-independent observable, this constitutes a useful test on the ability of QuCumber to capture the full wavefunction from the information contained in a single-basis dataset for TFIM.




Figure 3: The second Rényi entropy for the TFIM chain with $N = 10$ spins. The number of sites in the entangled bipartition A is indicated by the horizontal axis. The markers indicate values obtained through the “Swap” operator applied to the samples from a trained RBM. The dashed line denotes the result from exact diagonalization.

3 Complex wavefunctions

For positive wavefunctions, the probability distribution underlying the outcomes of projective measurements in the reference basis contains all possible information about the unknown quantum state. However, in the more general case of a wavefunction with a non-trivial sign or phase structure, this is not the case. In this section, we consider a target quantum state where the wavefunction coefficients in the reference basis can be complex-valued, $\Psi(\boldsymbol{\sigma}) = \Phi(\boldsymbol{\sigma})e^{i\theta(\boldsymbol{\sigma})}$. We continue to choose the reference basis as $\boldsymbol{\sigma} = \boldsymbol{\sigma}^z$. We first need to generalize the RBM representation of the quantum state to capture a generic complex wavefunction. To this end, we introduce an additional RBM with marginalized distribution $p_\mu(\boldsymbol{\sigma})$ parameterized by a new set of network weights and biases $\boldsymbol{\mu}$. We use this to define the quantum state as:

$$\psi_{\lambda\mu}(\boldsymbol{\sigma}) = \sqrt{p_\lambda(\boldsymbol{\sigma})}e^{i\phi_\mu(\boldsymbol{\sigma})/2}, \tag{12}$$

where $\phi_\mu(\boldsymbol{\sigma}) = \log(p_\mu(\boldsymbol{\sigma}))$ [6]. In this case, the reconstruction requires a different type of measurement setting. It is easy to see that projective measurements in the reference basis do not convey any information on the phases $\theta(\boldsymbol{\sigma})$, since $P(\boldsymbol{\sigma}) = |\Psi(\boldsymbol{\sigma})|^2 = \Phi^2(\boldsymbol{\sigma})$.

The general strategy to learn a phase structure is to apply a unitary transformation \mathcal{U} to the state $|\Psi\rangle$ before the measurements, such that the resulting measurement distribution $P'(\boldsymbol{\sigma}) = |\Psi'(\boldsymbol{\sigma})|^2$ of the rotated state $\Psi'(\boldsymbol{\sigma}) = \langle \boldsymbol{\sigma} | \mathcal{U} | \Psi \rangle$ contains fingerprints of the phases $\theta(\boldsymbol{\sigma})$ (Fig. 4). In general, different rotations must be independently applied to gain full information on the phase structure. We make the assumption of a tensor product structure of the rotations, $\mathcal{U} = \bigotimes_{j=1}^N \hat{\mathcal{U}}_j$. This is equivalent to a local change of basis from $|\boldsymbol{\sigma}\rangle$ to $\{|\boldsymbol{\sigma}^b\rangle = |\sigma_1^{b_1}, \dots, \sigma_N^{b_N}\rangle\}$, where the vector \mathbf{b} identifies the local basis b_j for each site j . The target wavefunction in the new basis is given by

$$\Psi(\boldsymbol{\sigma}^b) = \langle \boldsymbol{\sigma}^b | \Psi \rangle = \sum_{\boldsymbol{\sigma}} \langle \boldsymbol{\sigma}^b | \boldsymbol{\sigma} \rangle \langle \boldsymbol{\sigma} | \Psi \rangle = \sum_{\boldsymbol{\sigma}} \mathcal{U}(\boldsymbol{\sigma}^b, \boldsymbol{\sigma}) \Psi(\boldsymbol{\sigma}), \tag{13}$$

and the resulting measurement distribution is

$$P_b(\boldsymbol{\sigma}^b) = \left| \sum_{\boldsymbol{\sigma}} \mathcal{U}(\boldsymbol{\sigma}^b, \boldsymbol{\sigma}) \Psi(\boldsymbol{\sigma}) \right|^2. \tag{14}$$




Figure 4: Unitary rotations for two qubits. (left) Measurements on the reference basis. (right) Measurement in the rotated basis. The unitary rotation (the Hadamard gate on qubit σ_0) is applied after state preparation and before the projective measurement.

To clarify the procedure, let us consider the simple example of a quantum state of two qubits:

$$|\Psi\rangle = \sum_{\sigma_0, \sigma_1} \Phi_{\sigma_0 \sigma_1} e^{i\theta_{\sigma_0 \sigma_1}} |\sigma_0 \sigma_1\rangle, \quad (15)$$

and rotation $\mathcal{U} = \hat{H}_0 \otimes \hat{I}_1$, where \hat{I}_1 is the identity operator and

$$\hat{H} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad (16)$$

is called the *Hadamard gate*. This transformation is equivalent to rotating the qubit σ_0 from the reference σ_0^z basis to the σ_0^x basis. A straightforward calculation leads to the following probability distribution of the projective measurement in the new basis $|\sigma_0^x, \sigma_1\rangle$:

$$P_b(\sigma_0^x, \sigma_1) = \frac{\Phi_{0\sigma_1}^2 + \Phi_{1\sigma_1}^2}{4} + \frac{1 - 2\sigma_0^x}{2} \Phi_{0\sigma_1} \Phi_{1\sigma_1} \cos(\Delta\theta), \quad (17)$$

where $\Delta\theta = \theta_{0\sigma_1} - \theta_{1\sigma_1}$. Therefore, the statistics collected by measuring in this basis implicitly contains partial information on the phases. To obtain the full phases structure, additional transformations are required, one example being the rotation from the reference basis to the σ_j^y local basis, realized by the elementary gate

$$\hat{K} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -i \\ 1 & i \end{bmatrix}. \quad (18)$$

3.1 Setup

We now proceed to use QuCumber to reconstruct a complex-valued wavefunction. For simplicity, we restrict ourselves to two qubits and consider the general case of a quantum state with random amplitudes $\Phi_{\sigma_0 \sigma_1}$ and random phases $\theta_{\sigma_0 \sigma_1}$. This example is available in the online tutorial.³ We begin by importing the required packages:

```
from qucumber.nn_states import ComplexWaveFunction
import qucumber.utils.unitaries as unitaries
import qucumber.utils.cplx as cplx
```

³The tutorial for complex wavefunctions can be found at https://qucumber.readthedocs.io/en/stable/_examples/Tutorial2_TrainComplexWaveFunction/tutorial_qubits.html

Since we are dealing with a complex wavefunction, we load the corresponding module `ComplexWaveFunction` to build the RBM quantum state $\psi_{\lambda\mu}(\sigma)$. Furthermore, the following additional utility modules are required: the `utils.cplx` backend for complex algebra, and the `utils.unitaries` module which contains a set of elementary local rotations. By default, the set of unitaries include local rotations to the σ^x and σ^y basis, implemented by the \hat{H} and \hat{K} gates respectively.

We continue by loading the data⁴ into QuCumber, which is done using the `load_data` function of the data utility:

```
train_path = "qubits_train.txt"
train_bases_path = "qubits_train_bases.txt"
psi_path = "qubits_psi.txt"
bases_path = "qubits_bases.txt"

train_samples, true_psi, train_bases, bases = data.load_data(
    train_path, psi_path, train_bases_path, bases_path
)
```

As before, we may load the true target wavefunction from `qubits_psi.txt`, which can be used to calculate the fidelity and KL divergence. In contrast with the positive case, we now have measurements performed in different bases. Therefore, the training data consists of an array of qubits projections $(\sigma_0^{b_0}, \sigma_1^{b_1})$ in `qubits_train_samples.txt`, together with the corresponding bases (b_0, b_1) where the measurement was taken, in `qubits_train_bases.txt`. Finally, QuCumber loads the set of all the bases appearing in the training dataset, stored in `qubits_bases.txt`. This is required to properly configure the various elementary unitary rotations that need to be applied to the RBM state during the training. For this example, we generated measurements in the following bases:

$$(b_0, b_1) = (z, z), (x, z), (z, x), (y, z), (z, y). \quad (19)$$

Finally, before the training, we initialize the set of unitary rotations and create the RBM state object. In the case of the provided dataset, the unitaries are the \hat{H} and \hat{K} gates. The required dictionary can be created with `unitaries.create_dict()`. By default, when `unitaries.create_dict()` is called, it will contain the identity, the \hat{H} gate, and the \hat{K} gate, with the keys Z, X, and Y, respectively. It is possible to add additional gates by specifying them as

```
U = torch.tensor([[<re_part>], [<im_part>]], dtype=torch.double)
unitary_dict = unitaries.create_dict(<unitary_name>=U)
```

where `re_part`, `im_part`, and `unitary_name` are to be specified by the user.

We then initialize the complex RBM object with

```
state = ComplexWaveFunction(
    num_visible=2, num_hidden=2, unitary_dict=unitary_dict
)
```

The key difference between positive and complex wavefunction reconstruction is the requirement of additional measurements in different basis. Despite this, loading the data, initializing models, and training the RBMs are all very similar to the positive case, as we now discuss.

⁴ The training dataset can be downloaded from https://github.com/PIQuIL/QuCumber/blob/master/examples/Tutorial2_TrainComplexWaveFunction/

3.2 Training

Like in the case of a positive wavefunction, for the complex case QuCumber optimizes the network parameters to minimize the KL divergence between the data and the RBM distribution. When measuring in multiple bases, the optimization now runs over the set of parameters (λ, μ) and minimizes the sum of KL divergences between the data distribution $P(\sigma^b)$ and the RBM distribution $|\psi_{\lambda\mu}(\sigma^b)|^2$ for each bases b appearing in the training dataset [6]. For example, if a given training sample is measured in the basis (x, z) , QuCumber applies the appropriate unitary rotation $\mathcal{U} = \hat{H}_0 \otimes \hat{I}_1$ to the RBM state before collecting the gradient signal.

Similar to the case of positive wavefunction, we generate the Hilbert space (to compute fidelity and KL divergence) and initialize the callbacks

```
state.space = nn_state.generate_hilbert_space(2)
callbacks = [
    MetricEvaluator(
        log_every,
        {"Fidelity": ts.fidelity, "KL": ts.KL},
        target_psi=true_psi,
        bases=bases,
        verbose=True,
        space=state.space,
    )
]
```

The training is carried out by calling the fit function of ComplexWaveFunction, given the set of hyperparameters

```
state.fit(
    train_samples,
    epochs=100,
    pos_batch_size=10,
    neg_batch_size=10,
    lr=0.05,
    k=5,
    input_bases=train_bases,
    callbacks=callbacks,
)
```

In Fig. 5 we show the total KL divergence and the fidelity with the true two-qubit state during training. After successfully training a QuCumber model, we can once again compute expectation values of physical observables, as discussed in Section 2.3.

4 Conclusion

We have introduced the open source software package QuCumber, a quantum calculator used for many-body eigenstate reconstruction. QuCumber is capable of taking input data representing projective measurements of a quantum wavefunction, and reconstructing the wavefunction using a restricted Boltzmann machine (RBM).

Once properly trained, QuCumber can produce a new set of measurements, sampled stochastically from the RBM. These samples, generated in the reference basis, can be used to verify the training of the RBM against the original data set. More importantly, they can be used to calculate expectation values of many physical observables. In fact, any expectation value typically estimated by conventional Monte Carlo methods can be implemented as an estimator in QuCumber. Such estimators may be inaccessible in the reference basis, for example. Or, they




Figure 5: Training a complex RBM with QuCumber on random two-qubit data. We show the fidelity (left), and KL divergence (right), as a function of the training epochs.

may be difficult or impossible to implement in the setup for which the original data was obtained. This is particularly relevant for experiments, where it is easy to imagine many possible observables that are inaccessible, due to fundamental or technical challenges.

Future versions of QuCumber, as well as the next generation of quantum state reconstruction software, may explore different generative models, such as variational autoencoders, generative adversarial networks, or recurrent neural networks. The techniques described in this paper can also be extended to reconstruct mixed states, via the purification technique described in Reference [7]. In addition, future techniques may include hybridization between machine learning and other well-established methods in computational quantum many-body physics, such as variational Monte Carlo and tensor networks [8].

Acknowledgements

We acknowledge M. Albergo, G. Carleo, J. Carrasquilla, D. Sehayek, and L. Hayward Sierens for stimulating discussions. We thank the Perimeter Institute for Theoretical Physics for the continuing support of PIQuIL.

Author contributions Authors are listed alphabetically. For an updated record of individual contributions, consult the repository at <https://github.com/PIQuIL/QuCumber/graphs/contributors>.

Funding information This research was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Research Chair program, and the Perimeter Institute for Theoretical Physics. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used in this work. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Economic Development, Job Creation and Trade. P. H. acknowledges support from ICFOstepstone, funded by the Marie Skłodowska-Curie Co-funding of regional, national and international programmes (GA665884) of the European Commission, as well as by the Severo Ochoa 2016–2019’ program at ICFO (SEV–2015–0522), funded by the Spanish Ministry of Economy, Industry, and Competitiveness (MINECO).

A Glossary

This section contains an overview of terms discussed in the document which are relevant for RBMs. For more detail we refer the reader to the code documentation on <https://qucumber.readthedocs.io/en/stable/>, and References [12, 30].

- *Batch*: A subset of data upon which the gradient is computed and the network parameters are adjusted accordingly. A smaller batch size often results in a more stochastic trajectory, while a large batch size more closely approximates the exact gradient and has less variance.
- *Biases*: Adjustable parameters in an RBM, denoted by b_j and c_i in Eq. (20).
- *Contrastive divergence*: An approximate maximum-likelihood learning algorithm for RBMs [12]. CD estimates the gradient of the effective energy (20) with respect to model parameters by using Gibbs sampling to compare the generated and target distributions.
- *Energy*: The energy of the joint configuration (\mathbf{v}, \mathbf{h}) of a RBM is defined as follows:

$$E_{\lambda}(\mathbf{v}, \mathbf{h}) = -\sum_{j=1}^{n_v} b_j v_j - \sum_{i=1}^{n_h} c_i h_i - \sum_{ij} h_i W_{ij} v_j. \quad (20)$$

- *Effective energy*: Obtained from the energy by tracing out the hidden units \mathbf{h} ; often called the “free energy” in machine learning literature.

$$\mathcal{E}_{\lambda}(\mathbf{v}) = -\sum_{j=1}^{n_v} b_j v_j - \sum_{i=1}^{n_h} \log \left[1 + \exp \left(\sum_j W_{ij} v_j + c_i \right) \right]. \quad (21)$$

- *Epoch*: A single pass through an entire training set. For example, with a training set of 1,000 samples and a batch size of 100, one epoch consists of 10 updates of the parameters.
- *Gibbs sampling*: A Monte Carlo algorithm that samples from the conditional distribution of one variable, given the state of other variables. In an RBM, the restricted weight connectivity allows Gibbs sampling between the visible “block”, conditioned on the hidden “block”, and vice versa.
- *Hidden units*: There are n_h units in the hidden layer of the RBM, denoted by the vector $\mathbf{h} = (h_1, \dots, h_{n_h})$. The number of hidden units can be adjusted to increase the representational capacity of the RBM.
- *Hyperparameters*: A set of parameters that are not adjusted by a neural network during training. Examples include the learning rate, number of hidden units, batch size, and number of training epochs.
- *Joint distribution*: The RBM assigns a probability to each joint configuration (\mathbf{v}, \mathbf{h}) according to the Boltzmann distribution:

$$p_{\lambda}(\mathbf{v}, \mathbf{h}) = \frac{1}{Z_{\lambda}} e^{-E_{\lambda}(\mathbf{v}, \mathbf{h})}. \quad (22)$$

- *KL divergence*: The Kullback-Leibler divergence, or relative entropy, is a measure of the “distance” between two probability distributions P and Q , defined as:

$$\text{KL}(P \parallel Q) = \sum_{\mathbf{v}} P(\mathbf{v}) \log \frac{P(\mathbf{v})}{Q(\mathbf{v})}. \quad (23)$$

The KL divergence between two identical distributions is zero. Note that it is not symmetric between P and Q .

- *Learning rate*: The step size used in the gradient descent algorithm for the optimization of the network parameters. A small learning rate may result in better optimization but will take more time to converge. If the learning rate is too high, training might not converge or will find a poor optimum.
- *Marginal distribution*: Obtained by marginalizing out the hidden layer from the joint distribution via

$$p_{\lambda}(\mathbf{v}) = \frac{1}{Z_{\lambda}} \sum_{\mathbf{h}} e^{-E_{\lambda}(\mathbf{v}, \mathbf{h})} = \frac{1}{Z_{\lambda}} e^{-\mathcal{E}_{\lambda}(\mathbf{v})}. \quad (24)$$

- *QuCumber*: A quantum calculator used for many-body eigenstate reconstruction.
- *Parameters*: The set of weights and biases $\lambda = \{\mathbf{W}, \mathbf{b}, \mathbf{c}\}$ characterizing the RBM energy function. These are adjusted during training.
- *Partition function*: The normalizing constant of the Boltzmann distribution. It is obtained by tracing over all possible pairs of visible and hidden vectors:

$$Z_{\lambda} = \sum_{\mathbf{v}, \mathbf{h}} e^{-E_{\lambda}(\mathbf{v}, \mathbf{h})}. \quad (25)$$

- *Restricted Boltzmann Machine*: A two-layer network with bidirectionally connected stochastic processing units. “Restricted” refers to the connections (or weights) between the visible and hidden units: each visible unit is connected with each hidden unit, but there are no intra-layer connections.
- *Visible units*: There are n_v units in the visible layer of the RBM, denoted by the vector $\mathbf{v} = (v_1, \dots, v_{n_v})$. These units correspond to the physical degrees of freedom. In the cases considered in this paper, the number of visible units is equal to the number of spins N .
- *Weights*: W_{ij} is the symmetric connection or interaction between the visible unit v_j and the hidden unit h_i .

References

- [1] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow and J. M. Gambetta, *Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets*, Nature **549**, 242 (2017), doi:[10.1038/nature23879](https://doi.org/10.1038/nature23879).
- [2] N. Moll et al., *Quantum optimization using variational algorithms on near-term quantum devices*, Quantum Sci. Technol. **3**, 030503 (2018), doi:[10.1088/2058-9565/aab822](https://doi.org/10.1088/2058-9565/aab822).
- [3] H. Bernien et al., *Probing many-body dynamics on a 51-atom quantum simulator*, Nature **551**, 579 (2017), doi:[10.1038/nature24622](https://doi.org/10.1038/nature24622).

- [4] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong and C. Monroe, *Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator*, Nature **551**, 601 (2017), doi:[10.1038/nature24654](https://doi.org/10.1038/nature24654).
- [5] G. Torlai and R. G. Melko, *Learning thermodynamics with Boltzmann machines*, Phys. Rev. B **94**, 165134 (2016), doi:[10.1103/PhysRevB.94.165134](https://doi.org/10.1103/PhysRevB.94.165134).
- [6] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko and G. Carleo, *Neural-network quantum state tomography*, Nat. Phys. **14**, 447 (2018), doi:[10.1038/s41567-018-0048-5](https://doi.org/10.1038/s41567-018-0048-5).
- [7] G. Torlai and R. G. Melko, *Latent space purification via neural density operators*, Phys. Rev. Lett. **120**, 240503 (2018), doi:[10.1103/PhysRevLett.120.240503](https://doi.org/10.1103/PhysRevLett.120.240503).
- [8] J. Carrasquilla, G. Torlai, R. G. Melko and L. Aolita, *Reconstructing quantum states with generative models*, Nat. Mach. Intell. **1**, 155 (2019), doi:[10.1038/s42256-019-0028-1](https://doi.org/10.1038/s42256-019-0028-1).
- [9] D. T. Lennon, H. Moon, L. C. Camenzind, L. Yu, D. M. Zumbühl, G. A. D. Briggs, M. A. Osborne, E. A. Laird and N. Ares, *Efficiently measuring a quantum device using machine learning* (2018), [arXiv:1810.10042](https://arxiv.org/abs/1810.10042).
- [10] C. Kim, J.-K. Kevin Rhee, W. Lee and J. Ahn, *Mixed quantum state dynamics estimation with artificial neural network*, 2018 International conference on information and communication technology convergence (ICTC), 740 (2018), doi:[10.1109/ICTC.2018.8539575](https://doi.org/10.1109/ICTC.2018.8539575).
- [11] P. Smolensky, *Information processing in dynamical systems: Foundations of harmony theory*, In Parallel distributed processing: explorations in the microstructure of cognition, **Volume 1: Foundations**, 194 (1986), ISBN 0-262-68053-X.
- [12] G. E. Hinton, *Training products of experts by minimizing contrastive divergence*, Neural Comput. **14**, 1771 (2002), doi:[10.1162/089976602760128018](https://doi.org/10.1162/089976602760128018).
- [13] G. E. Hinton and R. R. Salakhutdinov, *Reducing the dimensionality of data with neural networks*, Science **313**, 504 (2006), doi:[10.1126/science.1127647](https://doi.org/10.1126/science.1127647).
- [14] Y. LeCun, Y. Bengio and G. Hinton, *Deep learning*, Nature **521**, 436 (2015), doi:[10.1038/nature14539](https://doi.org/10.1038/nature14539).
- [15] G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, Science **355**, 602 (2017), doi:[10.1126/science.aag2302](https://doi.org/10.1126/science.aag2302).
- [16] G. Carleo, Y. Nomura and M. Imada, *Constructing exact representations of quantum many-body systems with deep neural networks*, Nat. Commun. **9**, 5322 (2018), doi:[10.1038/s41467-018-07520-3](https://doi.org/10.1038/s41467-018-07520-3).
- [17] X. Gao and L.-M. Duan, *Efficient representation of quantum many-body states with deep neural networks*, Nat. Commun. **8**, 662 (2017), doi:[10.1038/s41467-017-00705-2](https://doi.org/10.1038/s41467-017-00705-2).
- [18] K. Choo, G. Carleo, N. Regnault and T. Neupert, *Symmetries and many-body excitations with neural-network quantum states*, Phys. Rev. Lett. **121**, 167204 (2018), doi:[10.1103/PhysRevLett.121.167204](https://doi.org/10.1103/PhysRevLett.121.167204).
- [19] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez and J. Ignacio Cirac, *Neural-network quantum states, string-bond states, and chiral topological states*, Phys. Rev. X **8**, 011006 (2018), doi:[10.1103/PhysRevX.8.011006](https://doi.org/10.1103/PhysRevX.8.011006).
- [20] P. Mehta and D. J. Schwab, *An exact mapping between the variational renormalization group and deep learning* (2014), [arXiv:1410.3831](https://arxiv.org/abs/1410.3831).

- [21] M. Koch-Janusz and Z. Ringel, *Mutual information, neural networks and the renormalization group*, Nat. Phys. **14**, 578 (2018), doi:[10.1038/s41567-018-0081-4](https://doi.org/10.1038/s41567-018-0081-4).
- [22] S. Iso, S. Shiba and S. Yokoo, *Scale-invariant feature extraction of neural network and renormalization group flow*, Phys. Rev. E **97**, 053304 (2018), doi:[10.1103/PhysRevE.97.053304](https://doi.org/10.1103/PhysRevE.97.053304).
- [23] P. M. Lenggenhager, Z. Ringel, S. D. Huber and M. Koch-Janusz, *Optimal renormalization group transformation from information theory* (2018), [arXiv:1809.09632](https://arxiv.org/abs/1809.09632).
- [24] J. Chen, S. Cheng, H. Xie, L. Wang and T. Xiang, *Equivalence of restricted Boltzmann machines and tensor network states*, Phys. Rev. B **97**, 085104 (2018), doi:[10.1103/PhysRevB.97.085104](https://doi.org/10.1103/PhysRevB.97.085104).
- [25] Y. Nomura, A. S. Darmawan, Y. Yamaji and M. Imada, *Restricted Boltzmann machine learning for solving strongly correlated quantum systems*, Phys. Rev. B **96**, 205152 (2017), doi:[10.1103/PhysRevB.96.205152](https://doi.org/10.1103/PhysRevB.96.205152).
- [26] S. Weinstein, *Neural networks as "hidden" variable models for quantum systems* (2018), [arXiv:1807.03910](https://arxiv.org/abs/1807.03910).
- [27] Y. Zheng, H. He, N. Regnault and B. A. Bernevig, *Restricted Boltzmann machines and matrix product states of 1D translational invariant stabilizer codes*, Phys. Rev. B **99**, 155129 (2019), doi:[10.1103/PhysRevB.99.155129](https://doi.org/10.1103/PhysRevB.99.155129).
- [28] G. Torlai and R. G. Melko, *Machine learning quantum states in the NISQ era* (2019), [arXiv:1905.04312](https://arxiv.org/abs/1905.04312).
- [29] M. B. Hastings, I. González, A. B. Kallin and R. G. Melko, *Measuring Renyi entanglement entropy in quantum Monte Carlo simulations*, Phys. Rev. Lett. **104**, 157201 (2010), doi:[10.1103/PhysRevLett.104.157201](https://doi.org/10.1103/PhysRevLett.104.157201).
- [30] G. E. Hinton, *A practical guide to training restricted Boltzmann machines*, In Neural Networks: Tricks of the Trade 599, Springer Berlin Heidelberg (2012), doi:[10.1007/978-3-642-35289-8_32](https://doi.org/10.1007/978-3-642-35289-8_32).