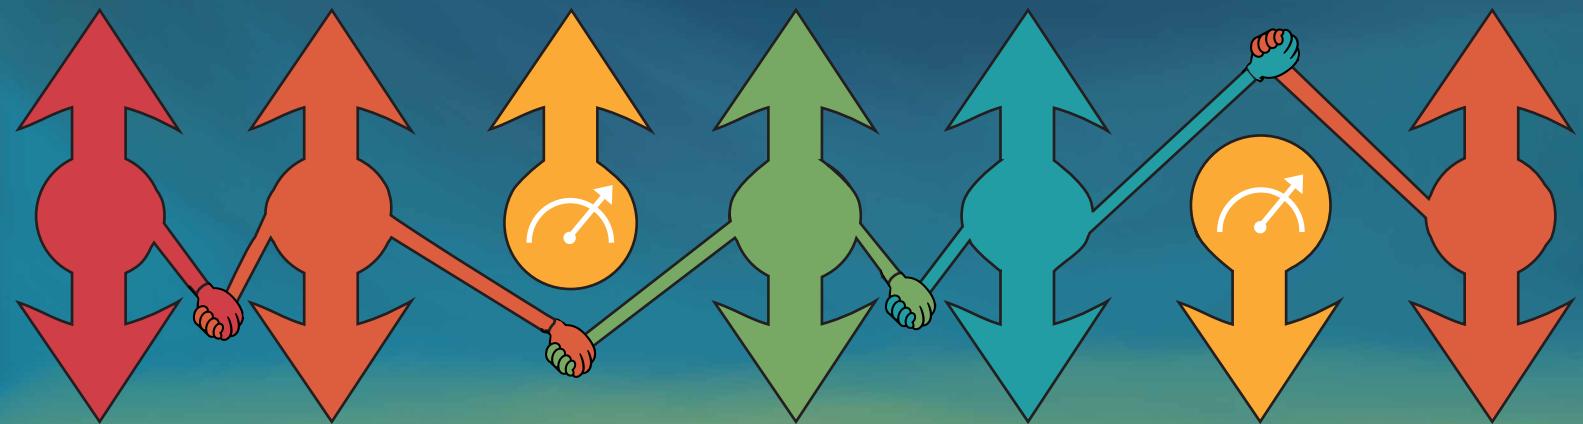


MEASUREMENT- BASED QUANTUM STATE PREPARATION & GROUND STATE SEARCH



DANIEL
ALCALDE PUENTE

DISSERTATION

Measurement-Based Quantum State Preparation and Ground State Search

Inaugural-Dissertation

zur Erlangung des Doktorgrades
der Mathematisch-Naturwissenschaftlichen Fakultät
der Universität zu Köln

vorgelegt von

Daniel Alcalde Puente

aus Madrid



Köln 2025

Berichterstatter: Prof. Dr. Matteo Rizzi
(Gutachter) Prof. Dr. David Gross

Diese Dissertation wurde von der Mathematisch-Naturwissenschaftlichen Fakultät der
Universität zu Köln angenommen.

Abstract

This thesis introduces and systematically analyzes novel measurement-based quantum state preparation protocols in one dimension, leveraging non-unitary dynamics, quantum measurements, and adaptive feedback mechanisms. Additionally, a stochastic sampling-based approach for ground state search using Projected Entangled Pair States (PEPS) in two dimensions is investigated, further expanding the computational toolkit for simulating complex quantum states.

A central contribution is the development of a protocol that employs periodic measurement and resetting of ancilla systems, explicitly applied to the preparation of the one-dimensional Affleck-Kennedy-Lieb-Tasaki (AKLT) state. We identify an optimal reset interval, balancing entanglement generation and convergence speed, and demonstrate robustness against realistic experimental noise such as dephasing.

Furthermore, this thesis introduces a self-learning adaptive measurement-feedback framework within variational quantum circuits (VQCs). The proposed framework autonomously discovers efficient, constant-depth strategies for deterministic preparation of specific AKLT edge states, outperforming existing analytical protocols. This analysis uncovers and mitigates previously unknown optimization challenges unique to measurement-feedback-based variational quantum circuits.

The thesis also advances two-dimensional ground state search by investigating a sampling-based optimization approach for finite PEPS. It demonstrates that PEPS representing physically realistic states are computationally tractable within this sampling framework, in contrast to random states, whose complexity makes them prohibitively difficult to contract. To quantitatively distinguish between computationally easy and hard-to-contract states, a novel diagnostic measure is introduced. The developed methodology is successfully applied to physically significant states, including chiral spin liquids and Rydberg atom arrays with long-range interactions, highlighting its effectiveness in simulating complex states relevant to state-of-the-art cold atom experiments.

Overall, the thesis provides methodological and conceptual insights into measurement-based quantum state preparation, optimization landscapes, and tensor network simulations.

Table of Content

1. Introduction	1
2. Tensor Networks	7
2.1. Matrix Product States	9
2.1.1. Orthogonal and Canonical Forms	10
2.1.2. Measurement and Sampling	12
2.1.3. Injectivity and Frustration-Free Hamiltonians	13
2.1.4. SPT Phases and G-Injectivity	15
2.1.5. The Affleck-Kennedy-Lieb-Tasaki State	16
2.2. Projected Entangled Pair States	18
2.2.1. Contractability and Additional Unpublished Algorithmic Improvements	21
3. Time-Dependent Variational Principle	28
3.1. For an MPS	30
3.2. For General Ansätze using Sampling	32
3.2.1. Expectation Values	32
3.2.2. Evolution	33
3.3. Real-Time Evolution of PEPS: Challenges and Possible Solutions (Unpublished)	34
3.3.1. Gradient Estimation Noise and its Scaling with System Size	36
3.3.2. Possible solutions - Hybrid-sampling TDVP approach	41
3.3.3. Optimizing the Fidelity	43
4. State Preparation	45
4.1. Unitary-only	48
4.1.1. Simple Unitary Preparation of an MPS State	49

4.2. Non-unitary with Only Measurements	50
4.2.1. Designing Lindbladians for Passive State Preparation	51
4.3. Non-unitary with Measurement and Feedback	54
4.3.1. Preparing SPT States Using Feedback and Measurements	55
4.4. Learning Non-unitary with Feedback	57
5. Variational Quantum Circuits	60
5.1. Practical Limitations	62
5.1.1. Barren Plateaus	62
5.1.2. Local Minima	63
5.2. Computing Gradients in VQCs	63
5.2.1. Finite Difference and Stochastic Approximation	63
5.2.2. Parameter-Shift Rule	64
5.2.3. Gradients via MPS	66
6. Neural Architectures	69
6.1. Feedforward Neural Networks	69
6.2. Avoiding Vanishing Gradients	70
6.3. Recurrent Neural Networks and Gated Units	71
6.4. Stacking RNNs	72
6.5. Transformers and Mamba Architectures	74
Publications	75
I. Quantum State Preparation via Engineered Ancilla Resetting	76
II. Learning Feedback Mechanisms for Measurement-Based Variational Quantum State Preparation	99
III. Efficient Optimization and Conceptual Barriers in Variational Finite Projected Entangled-Pair States	126
7. Conclusion	150
Bibliography	153

Publications

- [I] Daniel Alcalde Puente, Felix Motzoi, Tommaso Calarco, Giovanna Morigi, and Matteo Rizzi. “Quantum state preparation via engineered ancilla resetting”. In: *Quantum* 8 (2024), p. 1299.
- [II] Daniel Alcalde Puente and Matteo Rizzi. “Learning Feedback Mechanisms for Measurement-Based Variational Quantum State Preparation”. In: *Quantum* 9 (2025), p. 1792.
- [III] Daniel Alcalde Puente, Erik Lennart Weerda, Konrad Schröder, and Matteo Rizzi. “Efficient optimization and conceptual barriers in variational finite projected entangled pair states”. In: *Phys. Rev. B* 111 (19 2025), p. 195120.

CHAPTER 1

Introduction

Quantum technologies have the potential to profoundly transform information processing, computation, communication, and sensing. They might achieve this by harnessing uniquely quantum mechanical phenomena, such as superposition, entanglement, and coherence [1, 2]. Central to this transformative potential is the ability to efficiently and reliably prepare complex quantum states. Quantum state preparation as a resource is essential across numerous quantum computing algorithms, from resource states for Measurement-Based Quantum Computing [3] to the preparation of topological states for error correction [4]. Beyond its practical significance, the theoretical exploration of quantum state preparation provides fundamental insights into quantum phases of matter, quantum resource theory, and the complexity inherent to various quantum states [5].

However, contemporary quantum devices, known as noisy intermediate-scale quantum (NISQ) systems, face significant limitations. These include environmental decoherence and operational imperfections, which severely constrain their capacity for reliable quantum state preparation [1, 2]. Traditional state preparation protocols predominantly rely on deterministic, unitary quantum circuits composed of local gates. Such methods typically require deep circuit structures, imposing stringent requirements on qubit coherence times and gate fidelity [6]. Given current technological limitations, these requirements remain challenging to fulfill.

To address these issues, alternative variational quantum state preparation methods have emerged. These methods use parameterized quantum circuits optimized through classical feedback loops. Although variational protocols can reduce circuit complexity, they introduce significant practical difficulties. Notably, optimization landscapes often exhibit barren plateaus, characterized by exponentially diminishing gradients and non-convex features that hinder efficient optimization [7, 8], while still suffering from circuit depths that scale with system size. Other approaches, including adiabatic quantum state

preparation and optimized sequences of unitary operations, similarly demand prolonged coherence times and precise control, restricting their experimental practicality [9–11].

Recent advances suggest that measurement-based and dissipative protocols, incorporating non-unitary dynamics, may offer more robust avenues for quantum state preparation. These protocols exploit irreversible quantum operations such as measurement, ancilla reset, or dissipation to steer quantum systems efficiently towards desired target states. Prominent measurement-based techniques include quantum reservoir engineering, quantum state steering via ancillary systems [12–21], and adaptive measurement-feedback protocols [22–33]. Each method involves distinct trade-offs concerning convergence speed, operational complexity, robustness to noise, and feasibility for near-term quantum hardware. A detailed comparative analysis of these methods is provided in Chap. 4.

In our Publication [I], a protocol based on quantum state steering via periodic measurement and resetting of ancillary degrees of freedom is developed and analyzed. Ancilla systems are coherently coupled to the primary quantum system using an engineered Hamiltonian. Periodic measurement and reset operations on the ancillas lead to non-unitary dynamics that guide the system towards the target state. This protocol is explicitly applied to the preparation of the spin-1 Affleck-Kennedy-Lieb-Tasaki (AKLT) state, a symmetry-protected topological (SPT) state characterized by a frustration-free parent Hamiltonian. Numerical results demonstrate an optimal ancilla reset interval, balancing the rate of entanglement generation and convergence speed. Furthermore, the protocol exhibits resilience to deviations from this ideal reset timing and the presence of dephasing noise.

In contrast, adaptive measurement-feedback methods dynamically adjust future operations based on intermediate results, which can significantly accelerate state preparation under ideal¹ conditions.

In situations where no measurement protocol to prepare a certain quantum state is known, self-learning methods in the literature employ greedy strategies [34, 35], where feedback steps are optimized sequentially, based solely on immediate measurement outcomes. Such greedy protocols, while computationally simpler, generally fail to discover globally optimal strategies necessary for efficiently preparing quantum states. For example, they would not be able to learn protocols like the one in Refs. [22–25, 27], where the projective measurements are constructed in such a way as to be correctable with local unitary gates.

Machine learning techniques, such as reinforcement learning (RL), have been proposed to enhance the flexibility and effectiveness of adaptive measurement-feedback protocols

¹Note that in current experimental implementations, real-time feedback is still out of reach.

and could, in theory, learn protocols like the one proposed in Refs. [22–25, 27]. Despite their promise, these methods face considerable computational and practical challenges. These arise primarily from the complexity associated with large Hilbert spaces, limiting their scalability beyond small quantum systems [36–39]. More details about learning protocols can be found in Sec. 4.4.

In Publication [II], we present a novel self-learning approach integrating measurement and conditional feedback within variational quantum circuits. We discovered a new class of local minima related to measurement-based variational circuits and devise strategies to overcome these challenges. This framework is applied to successfully prepare states within the spin-1 AKLT manifold. While analytical protocols exist to statistically prepare states within this manifold [27], the developed learning approach surpasses these methods by achieving deterministic preparation of a specific AKLT state in the manifold, for which no analytical protocol is currently known.

While measurement-feedback techniques have been extensively explored in one-dimensional quantum systems, primarily exploiting the representation power of matrix product states (MPS), their extension to two-dimensional quantum systems presents substantial conceptual and computational challenges. Two-dimensional quantum systems host rich, complex states, including quantum spin liquids and various topological phases, with significant potential for quantum computing and quantum memory applications [40, 41].

To effectively study quantum state preparation in two dimensions, tensor network methods, particularly Projected Entangled Pair States (PEPS), provide powerful numerical frameworks. PEPS naturally capture the complex entanglement structures prevalent in two-dimensional quantum systems. Recent advances in PEPS optimization via the time-dependent variational principle (TDVP), combined with efficient stochastic sampling methods, significantly enhance their computational feasibility and accuracy [42–44].

In Publication [III], computational complexity aspects of simulating quantum measurements within the PEPS framework are thoroughly analyzed, offering a framework to assess, in future work, the classical simulability of measurement-based state preparation protocols in two-dimensional systems². Specifically, the complexity associated with contracting single-layer tensor networks arising from local quantum measurements is examined. Additionally, it is demonstrated that PEPS optimization, combined with stochastic sampling techniques, yields a variationally rigorous and scalable numerical strategy for finite-size quantum systems. The approach is demonstrated on physically non-trivial states, including chiral spin liquids and Rydberg atom arrays with long-range interactions, which are typically difficult to simulate using other methods.

²This aspect is discussed more thoroughly in the Conclusion chapter.

Outline

This thesis is organized as follows: Chapters 2 through 6 introduce essential physical concepts and methodological tools necessary for comprehending the presented publications. These concepts and methods are introduced with a particular emphasis on their relevance to the Publications [I,II,III]. Subsequently, these publications are individually discussed, providing summaries and highlighting their contextual relationship to the thesis as a whole.

In Chap. 2, the key numerical tools employed in this thesis for simulating quantum systems are introduced, emphasizing the role of Tensor Networks. Matrix Product States (MPS) are presented in Sec. 2.1 as a powerful tool for efficiently approximating quantum states in 1D, particularly those satisfying the area law of entanglement. In Sec. 2.1.3, the properties of ground states of frustration-free Hamiltonians are discussed as those are the ones that can be prepared with the scheme presented in Publication [I]. Symmetry-Protected Topological (SPT) phases are explored through the concept of G-injectivity in Sec. 2.1.4, with special emphasis on their unique topological features relevant to Publication [II]. Lastly, Projected Entangled Pair States (PEPS) are introduced in Sec. 2.2 as an extension of MPS to two-dimensional systems, emphasizing their computational challenges and the role of the Vidal gauge in efficient contractions, as discussed in Publication [III]. Note that in Sec. 2.2.1, I further investigate the role of spectral entropy in contraction efficiency and leverage this understanding to introduce an unpublished method that accelerates single-layer PEPS contractions by approximately 50%, which will translate into faster ground state search code.

In Chap. 3, the Time-Dependent Variational Principle (TDVP) is introduced as a central method for simulating the time evolution of parametrized quantum states. The general formalism of TDVP is presented, including its derivation from the Schrödinger equation and the resulting equations of motion for the variational parameters that can in most cases not be solved exactly. Two major cases where solutions can be found are discussed in detail. In Sec. 3.1, the efficient application of TDVP to MPS is presented, where the TDVP equation can be solved approximately through a suitable parametrization of the ansatz. This method is employed for the simulation of the state preparation protocol in Publication [I]. In Sec. 3.2, a sampling-based approach applicable to general ansätze such as finite PEPS and Neural Network States is introduced. This method is central to Publication [III], where it is used for ground state search. Additionally, the chapter discusses the advantages of the recently introduced minimum-step stochastic reconfiguration (minSR) technique [45]. This method enables efficient optimization even in large parameter spaces, as utilized in Publication [III]. In Sec. 3.3, challenges in applying TDVP to finite PEPS and tensor networks are examined in the context of real-time evolution. It is shown analytically that the variance of the gradient estimator

scales linearly with system size for product states, leading to a proportional increase in the required number of Monte Carlo samples. For a PEPS, the situation is found to be significantly worse; the number of samples needed scales with the number of parameters in the ansatz as well as with the entanglement of the PEPS. To address this, a hybrid method is proposed that combines TEBD with simple update as a preconditioning step, followed by a fidelity-optimized sampling refinement procedure that was not done in this thesis due to a lack of time.

In Chap. 4, three strategies for quantum state preparation are presented: purely unitary circuits (Sec. 4.1), measurement-based protocols without feedback (Sec. 4.2), and feedback-assisted schemes (Sec. 4.3). These approaches underpin the results of Publications [I,II], which explore the role of dissipation, measurements, and adaptive control in preparing quantum states. The chapter first discusses unitary methods, including exact circuits for MPS. It then turns to non-unitary protocols based on measurements and engineered dissipation, including the AKLT state preparation protocol from Publication [I], which employs a steering Hamiltonian derived from a tailored Lindbladian. The final section focuses on feedback-based methods, highlighting adaptive strategies developed in Publication [II] through a learning-based approach. Each method is illustrated with a concrete example: an MPS unitary circuit, a dissipative AKLT protocol, and a feedback scheme using Bell measurements and G -injectivity.

In Chap. 5, the structure and training of Variational Quantum Circuits (VQCs) is examined, focusing on their application to quantum state preparation. The chapter introduces hardware-efficient ansätze and analyzes their limitations. These include barren plateaus (Sec. 5.1.1) and local minima (Sec. 5.1.2), which hinder scalability and optimization. Several methods for gradient computation are presented. These include the parameter-shift rule, finite-difference methods, and Simultaneous Perturbation Stochastic Approximation (SPSA). The chapter also discusses the self-learning feedback and measurement strategy employed in Publication [II]. Gradient evaluations in that work were performed using custom MPS-based simulations, implemented in the package `mVQE.jl` [46]. Section 5.2.3 explains how structural features of the hardware-efficient ansatz can be exploited to reduce computational costs.

In Chap. 6, key neural network architectures used in Publication [II] to represent the parameterized feedback function are summarized. The chapter introduces Feedforward Neural Networks (FFNs) and discusses techniques to mitigate vanishing gradients, such as residual connections, specific activation functions (ReLU, SwiGLU), and RMS Normalization. It then details Recurrent Neural Networks (RNNs), particularly Gated Recurrent Units (GRUs), and their stacked application as employed in Publication [II], concluding with a brief overview of Transformer and Mamba architectures as potential future directions.

In the Conclusion in Chap. 7, the thesis's exploration of quantum measurements for state preparation and ground state search is summarized. It recaps key contributions from Publications [I,II,III] and then discusses current limitations and interesting avenues for further research.

CHAPTER 2

Tensor Networks

Tensor networks offer a powerful framework for representing and manipulating quantum many-body states by decomposing large, high-dimensional tensors into structured networks of smaller, interconnected tensors. This decomposition enables efficient computation in otherwise intractable Hilbert spaces by exploiting the locality and entanglement structure inherent in physical systems. Tensor networks reduce computational complexity and enable intuitive visual representations of quantum states and operations. As such, they serve as essential tools in both numerical simulation and theoretical analysis.

Graphical notation is a central feature of tensor network methods. In this representation, tensors are depicted as nodes, and their indices are shown as connecting lines. Contracting indices corresponds to joining lines, providing a compact and intuitive formalism for tensor operations. For example, a vector is represented as:

$$v_i = \overset{i}{\circlearrowleft} v \circlearrowright, \quad (2.1)$$

and the application of a matrix A to this vector is illustrated by:

$$\sum_i A_{j,i} v_i = \overset{j}{\circlearrowleft} A \circlearrowright v \circlearrowright. \quad (2.2)$$

Isometries constitute an important class of tensors used in tensor network constructions. A matrix U is called an isometry if it satisfies $UU^\dagger = \mathbb{1}$, which implies that the rows of U are orthonormal. In graphical notation, isometries are typically represented as triangles,

emphasizing their role as norm-preserving mappings:

$$UU^\dagger = \left(\begin{array}{c} \text{triangle} \\ \text{circle} \end{array} \right) = \mathbb{1} \left(\begin{array}{c} \text{triangle} \\ \text{circle} \end{array} \right) \quad . \quad (2.3)$$

If U is a square matrix, it is unitary, and both $UU^\dagger = \mathbb{1}$ and $U^\dagger U = \mathbb{1}$ hold. If U is rectangular, then only $UU^\dagger = \mathbb{1}$ holds. In this case, U maps a higher-dimensional space to a lower-dimensional one while preserving inner products in the lower-dimensional space. These isometric mappings are fundamental to tensor network algorithms, where they are used to truncate Hilbert spaces while retaining fidelity in physically relevant subspaces.

Operations such as the singular value decomposition (SVD) can also be expressed in graphical notation:

$$A = \text{circle} \quad (2.4)$$

$$= \text{triangle} \text{square} \text{triangle} \quad (2.5)$$

Throughout this thesis, diagonal matrices of the form $S_{i,j} = \delta_{i,j} S_i$ are represented as squares, while general tensors with no special structure are depicted as circles.

This graphical formalism provides a foundation for understanding more advanced tensor network structures, including Matrix Product States (MPS) and Projected Entangled Pair States (PEPS), which are discussed in the following sections.

2.1 Matrix Product States

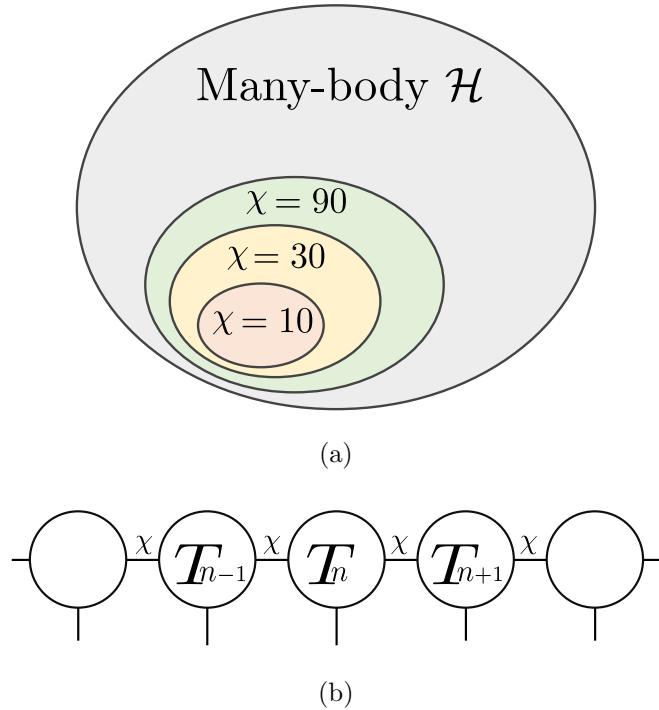


Figure 2.1.: (a) Matrix Product States (MPS) approximate quantum states in the many-body Hilbert space \mathcal{H} . Increasing bond dimension χ systematically captures more states, efficiently representing those satisfying the area law of entanglement. (b) MPS represented as a tensor network, where each tensor T_n is connected by bonds with dimension χ . Larger χ allows representation of states with higher entanglement.

Matrix Product States (MPS) provide a highly efficient representation for quantum states of many-body systems, particularly one-dimensional systems, by decomposing a general quantum state into a chain of local tensors. This representation leverages the area law of entanglement, meaning entanglement, for many relevant states¹, scales only with the boundary between regions, allowing MPS to efficiently approximate states with limited entanglement by spanning a small part of the Hilbert space (see Fig. 2.1a). As a result, MPS significantly mitigate the exponential complexity inherent to quantum many-body problems, making them a powerful computational tool within tensor network methods.

In the context of the research presented in this thesis, MPS played a central role across different methodologies and applications:

¹Notably ground states of gapped Hamiltonians fulfill the area law [47], this is explored in more detail in Sec. 2.1.3.

In Publication [I], MPS were utilized within the Time-Dependent Variational Principle (TDVP) framework (see Sec. 3.1) to simulate the dynamics of quantum systems under the periodic quantum resetting protocols. The flexibility of MPS enabled the simulation of state evolution and the characterization of entanglement growth. These features were critical for identifying optimal resetting intervals and ideal mapping operators.

In Publication [II], MPS served to represent quantum states generated by variational quantum algorithms incorporating measurement and feedback. Moreover, MPS techniques facilitated efficient computation of gradients required for parameter optimization within Variational Quantum Circuits (VQCs), as detailed in Sec. 5.2.3. This enabled scalable numerical simulations of quantum state preparation strategies beyond traditional unitary-only methods.

In Publication [III], MPS were employed via the boundary-MPS method to efficiently contract single-layer 2D tensor networks arising from finite-size Projected Entangled Pair States (PEPS). This approach enabled practical computation of overlaps between PEPS and product states in the computational basis, $\psi(\mathbf{s}) = \langle \mathbf{s} | \psi \rangle$, and is an important part of the algorithm employed for the projective measurement of 2D systems.

Collectively, these applications underscore the versatility and computational advantages of MPS, positioning them as a foundational tool throughout this research.

MPS represent quantum many-body states by decomposing the full state into a product of local tensors:

$$|\psi\rangle = \sum_{s_1, s_2, \dots, s_N} \underbrace{T_1^{s_1} T_2^{s_2} \dots T_N^{s_N}}_{\psi(\mathbf{s})} |s_1 s_2 \dots s_N\rangle, \quad (2.6)$$

where each $T_n^{s_n}$ is a tensor with three indices, except for the boundary tensors T_1 and T_N , which have two. The contraction over internal indices, known as bond indices, is illustrated in Fig. 2.1b. The dimension of these indices, referred to as the bond dimension χ , controls the expressive power of the MPS. For example, product states correspond to $\chi = 1$. More generally, the entanglement entropy across any bipartition is bounded by $S \leq \log(\chi)$, indicating that χ limits the amount of entanglement the MPS can efficiently capture.

2.1.1 Orthogonal and Canonical Forms

The orthogonal form of a Matrix Product State exploits the gauge freedom inherent in its tensor decomposition. Specifically, inserting an invertible matrix M and its inverse on a bond index, such that $T_n^{s_n}, T_{n+1}^{s_{n+1}} \rightarrow T_n^{s_n} M^{-1}, M T_{n+1}^{s_{n+1}}$, leaves the overall quantum state invariant.

Transforming an MPS into its central orthogonal form is achieved through a sequence

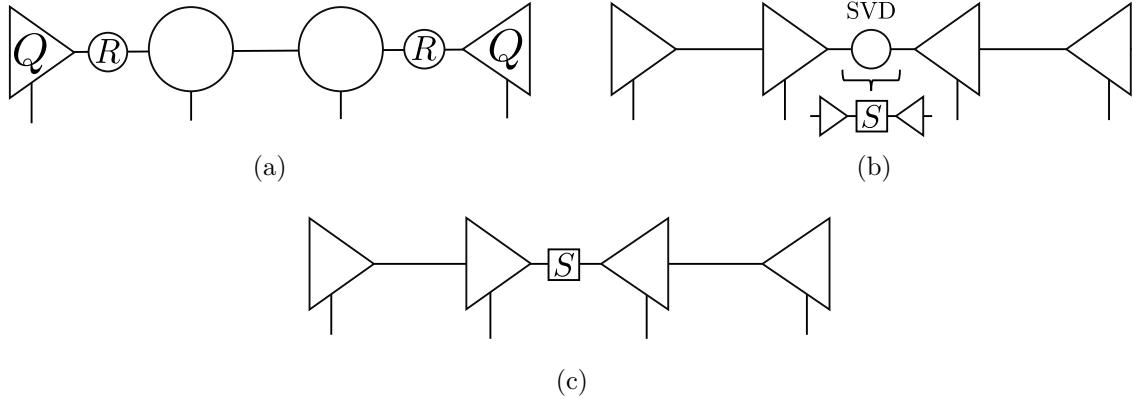


Figure 2.2.: Different orthonormal forms of Matrix Product States.

of QR decompositions, followed by a final singular value decomposition (SVD), as illustrated in Fig. 2.2. In this form, tensors act as isometries directed toward the central tensor. Consequently, computations such as the state norm become straightforward:

$$\langle \psi | \psi \rangle = \sum_{i=1}^{\chi} S_i^2. \quad (2.7)$$

This simplification arises because all isometric tensors to the left and right of the central singular values S contract to the identity.

With this result in hand, reducing the bond dimension χ can be accomplished by truncating the smallest singular values S_i . The fidelity after truncation is:

$$\frac{\sum_{i=1}^{\chi_{\text{new}}} S_i^2}{\sum_{i=1}^{\chi} S_i^2}. \quad (2.8)$$

This procedure can be employed, for example, to reduce the bond dimension of an MPS following the application of a two-body unitary gate as further discussed in Sec. 5.2.3. By specifying a target fidelity, the singular values S_i can be truncated dynamically to achieve the desired accuracy. This method is used extensively in all my publications.

A particularly useful gauge is the canonical form. In this form, tensors are explicitly chosen such that singular values from Schmidt decompositions across each bond are isolated (see Fig. 2.3a). One of the strengths of the canonical form is that it allows for efficient computation of entanglement properties. Given a bipartition of the system into subsystems A and B , the singular values S_i obtained at the cut provide the Schmidt

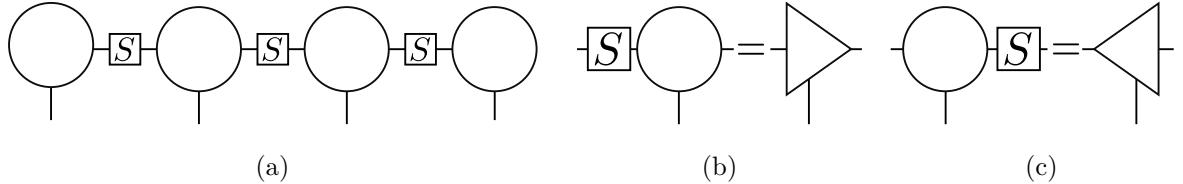


Figure 2.3.: Canonical form. After contracting all S matrices from the left and right towards a central tensor, the MPS is brought into canonical form, with a well-defined orthogonality center at the central tensor.

coefficients. The eigenvalues of the reduced density matrix ρ_A are then:

$$\lambda_i = \frac{S_i^2}{\sum_i S_i^2}, \quad (2.9)$$

and the entanglement entropy is:

$$S = - \sum_i \lambda_i \log(\lambda_i). \quad (2.10)$$

In two-dimensional systems, a canonical gauge analogous to the MPS orthogonal form is not directly available. Instead, an alternative known as the Vidal gauge is employed for finite PEPS, which exhibits a structurally similar form. In this gauge, the singular values no longer directly quantify entanglement or correlations between subsystems. Due to the presence of loops in the PEPS network, their interpretation is more subtle. In Publication [III], a correlation was identified between these singular values and the contractibility of single-layer tensor networks in finite PEPS calculations. This relationship is further analyzed in Sec. 2.2.1.

2.1.2 Measurement and Sampling

Projectively measuring quantum states represented by tensor networks played a major role in all of the Publications this thesis is based on. In Publications [I,II] it was an important tool to perform the ancilla measurements necessary for the implementation of the state preparation algorithms. While in Publication [III] sampling from a PEPS is an integral part of the ground state search algorithm.

Sampling from a wave function represented by an MPS involves expressing the probability distribution in terms of conditional probabilities:

$$P(\mathbf{s}) = P(s_1) \cdot P(s_2|s_1) \cdot P(s_3|s_1, s_2) \dots P(s_N|s_1, s_2, \dots, s_{N-1}). \quad (2.11)$$

Each conditional probability can be explicitly computed from the MPS representation

as follows:

$$P(\mathbf{s}_1) = \sum_{s_2, \dots, s_N} |\psi(\mathbf{s})|^2 \quad (2.12)$$

$$= \begin{array}{c} \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \end{array} \quad (2.13)$$

$$P(\mathbf{s}_1, \mathbf{s}_2) = \begin{array}{c} \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \end{array} \quad (2.14)$$

$$P(s_2|s_1) = \frac{P(s_1, s_2)}{P(s_1)}. \quad (2.15)$$

In practice, the sampling procedure sequentially selects quantum numbers s_i . Initially, the state s_1 is sampled according to $P(s_1)$. Subsequently, state s_2 is sampled from the conditional probability distribution $P(s_2|s_1)$, given the previously determined state s_1 . This iterative process continues until all quantum numbers s_N have been sampled. It is important to note that, when the MPS is in a left-orthogonal form, all right environments from the previous equations simplify to identities, significantly reducing computational overhead.

2.1.3 Injectivity and Frustration-Free Hamiltonians

An Hamiltonian $H = \sum_j h_j$ is called frustration-free if its global ground state $|\psi_g\rangle$ also minimizes each local term individually:

$$h_j |\psi_g\rangle = \epsilon_j^{\min} |\psi_g\rangle, \quad \forall j, \quad (2.16)$$

where ϵ_j^{\min} is the smallest eigenvalue of h_j . In this case, the ground state energy is the sum of the local minimal energies $E = \sum_j \epsilon_j$, and $|\psi_g\rangle$ satisfies all local constraints simultaneously. Frustration-free Hamiltonians play a central role in tensor network theory because their ground states admit exact, compact representations as MPS.

Moreover, many tensor network states admit the so-called parent Hamiltonian construction: given an *injective* MPS, one can explicitly construct a local, frustration-free Hamiltonian for which the MPS is the unique ground state.

An MPS is injective [48, 49] if each of its tensors defines an injective mapping from the virtual (link) indices to the physical indices. More explicitly, consider an MPS tensor

$T_{a,b}^s$ with physical index s and link indices a, b . Injectivity requires that the mapping

$$\Gamma : \mathcal{H}_{\text{link}} \rightarrow \mathcal{H}_{\text{phys}}, \quad \Gamma : |a, b\rangle \mapsto \sum_s T_{a,b}^s |s\rangle \quad (2.17)$$

is injective.

The concept of injectivity extends naturally to MPS with finite injective length. In such cases, injectivity is not evident from an individual tensor but emerges clearly when multiple adjacent tensors are grouped together. The minimal number of contiguous tensors required to form an injective mapping defines the injective length of the MPS. For intuition, an MPS with bond dimension χ and physical dimension d has a minimal injective length of at least $\lceil \log_d \chi \rceil$ so that the mapping can be injective.

Injectivity and finite injective length have important consequences for the structure and physical properties of MPS. Injective MPS generally describe states with finite correlation length and often appear as ground states of gapped local Hamiltonians. Moreover, each injective MPS, as explained before, admits the explicit construction of a corresponding frustration-free parent Hamiltonian [49–51]. The ground state space of this Hamiltonian is exactly spanned by the MPS, and the injective length determines the required interaction range.

This property is particularly relevant for the class of measurement-based preparation protocols introduced in [19–21] and further developed in Publication [I]. States that admit a frustration-free parent Hamiltonian allow the definition of an associated steering Hamiltonian, which drives the system toward its ground state. As a result, a steering Hamiltonian can be explicitly constructed for any injective state.

In practice, the resource requirements of such a measurement-based protocol depend on the characteristics of the MPS. Specifically, the number of ancilla qudits, the structure of the interaction terms, and the interaction range of the steering Hamiltonian scale with both the injective length and the bond dimension.

In contrast, non-injective MPS tensors feature mappings with nontrivial kernels. Familiar examples include the W-state and the GHZ state, whose MPS representations fail to become injective regardless of how many tensors are grouped. These states exhibit long-range correlations and belong to a fundamentally different phase, emphasizing the structural richness beyond injective MPS.

Hence, the concepts of injectivity and injective length offer essential insight into both the classification and the preparability of quantum states represented by MPS.

2.1.4 SPT Phases and G-Injectivity

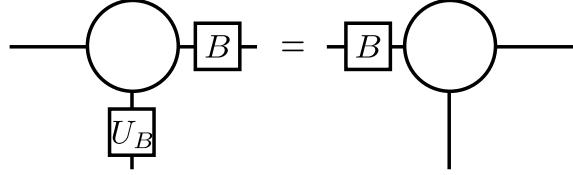


Figure 2.4.: In a G-Injective MPS, through an application of a Unitary U_B a matrix B can be moved between link indices.

Symmetry-protected topological (SPT) phases in one dimension play a central role in the classification of quantum phases of matter. A prominent example is the Affleck-Kennedy-Lieb-Tasaki (AKLT) state, whose identification as an SPT phase offers valuable insights into the state preparation methods discussed in Publications [I,II] and whose properties are detailed in Sec. 2.1.5.

One-dimensional symmetry-protected topological (SPT) phases are gapped quantum states that can be distinguished from trivial product states only in the presence of specific global symmetries [52, 53]. These phases cannot be continuously deformed into trivial states without either closing the energy gap or explicitly breaking the protecting symmetries. It is important to note that if the relevant symmetries are broken, an adiabatic path connecting the SPT phase to a trivial phase becomes possible.

A characteristic feature of SPT phases is the presence of symmetry-protected edge states, which vanish under periodic boundary conditions. These edge states are sensitive to symmetry-breaking operations, which can destroy them even if the bulk gap remains open.

Despite exhibiting exponential decay of correlations, typically of the form $\langle O_i O_j \rangle \sim e^{-|i-j|/\xi}$, SPT phases display a form of hidden order that is invisible to local observables. This topological order is instead revealed through nonlocal string order parameters, which serve as diagnostic tools for detecting the underlying symmetry-protected structure [53]. An example of such a string order parameter is:

$$\lim_{|j-i| \rightarrow \infty} \langle O_i^L e^{i\pi \sum_{k=i+1}^{j-1} O_k} O_j^R \rangle \neq 0, \quad (2.18)$$

indicating nonlocal correlations.

MPS naturally encode the properties of SPT phases through the concept of G-injectivity [47, 49]. An MPS tensor $T^{[i]}$ is called *G*-injective if it fulfills the symmetry relation:

$$U_B(g)_{ij} T^{[j]} = B(g) T^{[i]} B(g)^{-1}, \quad (2.19)$$

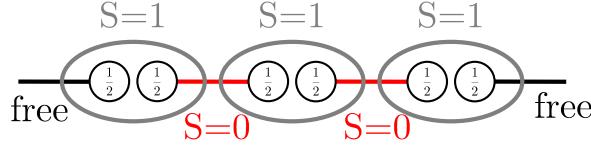


Figure 2.5.: Representation of the AKLT state. Each spin-1 particle is decomposed into two spin- $\frac{1}{2}$ degrees of freedom (gray circles), constrained to form a spin-1. Neighboring spin- $\frac{1}{2}$ particles form singlets (red lines), ensuring a total spin of zero across bonds.

for all group elements g of the symmetry group G . Here, $U_B(g)$ acts on the physical indices, while $B(g)$ acts on the virtual indices of the tensor, as depicted in Fig. 2.4. This structure allows the MPS to capture the SPT properties of the state.

For open boundary conditions, the global ground state of a G -injective MPS is not unique; instead, it exhibits a ground state degeneracy associated with the projective representations of G acting on the virtual edge degrees of freedom. More precisely, the symmetry action on the virtual level can be written as $B(g)_L \otimes B(g)_R$ on the left and right edges, respectively. Since the Hamiltonian is symmetric under the on-site action $U_B(g)$, and the bulk tensors satisfy the G -injectivity condition, the symmetry is effectively pushed to the boundaries and implemented via $B(g)_L$ and $B(g)_R$.

As a result, the edges of an SPT phase host robust zero-energy modes that transform projectively under G . These edge states are protected by the symmetry: as long as the global symmetry is preserved, the projective structure at the edges cannot be removed by any local, symmetric perturbation. However, if the symmetry is explicitly broken, this protection is lifted and the edge modes may be split or gapped out. This characteristic edge behavior is a defining feature of one-dimensional SPT phases.

This edge projective symmetry representation is a defining feature of one-dimensional SPT phases and encapsulates their robust boundary physics.

Moreover, G -injectivity facilitates efficient preparation of SPT states via measurement-based protocols that have a constant circuit depth [27]. In this protocol, after fusion measurements are performed, the target state is prepared up to some errors B in the links of the MPS. These can all be moved to the edge using the G -injectivity (see Fig. 2.4). This is further explored in Sec. 4.3.1.

2.1.5 The Affleck-Kennedy-Lieb-Tasaki State

The Affleck-Kennedy-Lieb-Tasaki (AKLT) model serves as a canonical example of a one-dimensional SPT phase, specifically illustrating the Haldane phase [54]. The AKLT

Hamiltonian is defined as:

$$H_{\text{AKLT}} = \sum_i \left[\frac{1}{2} \mathbf{S}_i \cdot \mathbf{S}_{i+1} + \frac{1}{6} (\mathbf{S}_i \cdot \mathbf{S}_{i+1})^2 + \frac{1}{3} \right], \quad (2.20)$$

which can be expressed as a sum of projectors onto the spin-2 subspace:

$$H_{\text{AKLT}} = \sum_i P_{i,i+1}^{(2)}, \quad (2.21)$$

where $P_{i,i+1}^{(2)}$ projects the combined state of neighboring spin-1 particles onto the total spin-2 subspace. The ground state of this Hamiltonian is the AKLT state, which has zero energy and is characterized by the absence of spin-2 components between neighboring sites. The AKLT MPS tensors are:

$$T^{+1} = +\sqrt{\frac{2}{3}}\sigma^+, \quad T^0 = -\sqrt{\frac{1}{3}}\sigma^z, \quad T^{-1} = -\sqrt{\frac{2}{3}}\sigma^-. \quad (2.22)$$

In the AKLT state, each spin-1 particle is represented as a symmetrized pair of virtual spin- $\frac{1}{2}$ particles. These virtual spin- $\frac{1}{2}$ degrees of freedom form singlet bonds with their counterparts on neighboring sites, as illustrated in Fig. 2.5. This valence bond solid structure ensures that the physical spin-1 particles are entangled in a fixed pattern that reflects the underlying symmetry-protected topological (SPT) order.

When expressed in the computational basis, the AKLT state consists of configurations with alternating spin quantum numbers ± 1 separated by zeros, $|\dots, +1, 0, \dots, 0, -1, \dots\rangle$, a structure that is directly evident from the matrix product state (MPS) representation in Eq. (2.22). This structure can be revealed with the help of a string operator:

$$\lim_{|i-j| \rightarrow \infty} \langle S_i^z e^{i\pi \sum_{k=i+1}^{j-1} S_k^z} S_j^z \rangle. \quad (2.23)$$

Under open boundary conditions, unpaired spin- $\frac{1}{2}$ edge states emerge at both ends of the chain. These edge modes are characteristic of the state's symmetry-protected topological (SPT) nature and are absent under periodic boundary conditions, consistent with the defining features of SPT phases.

The AKLT exhibits several symmetries that underpin its topological properties. Firstly, the MPS tensor possesses an on-site $\text{SO}(3)$ symmetry. Within this group, the dihedral subgroup D_2 , corresponding to π -rotations about the spin axes, enforces a doubly degenerate entanglement spectrum [55].

Secondly, the AKLT state respects bond-centered inversion symmetry, denoted by P .

This symmetry is implemented on the virtual spin- $\frac{1}{2}$ degrees of freedom by the operator $i\sigma^y$. The presence of this symmetry implies that the AKLT chain can not be adiabatically connected to a trivial product state without either breaking inversion symmetry or closing the energy gap.

Thirdly, the AKLT state is invariant under time-reversal symmetry, represented by the operator $U_T = e^{i\pi S^y} K$, where K denotes complex conjugation. On the virtual level, time-reversal acts as $i\sigma^y K$. Similar to the previous symmetries, time-reversal symmetry enforces a two-fold degeneracy in the entanglement spectrum.

Moreover, the AKLT state satisfies the conditions of G -injectivity under $\text{SO}(3)$ symmetry. The MPS tensors T^i , constructed using spin- $\frac{1}{2}$ Pauli matrices σ^α , fulfill the symmetry relation:

$$U(g)_{ij}^{(\text{spin-1})} T^j = V(g)^{(\text{spin-1/2})} T^i V(g)^\dagger {}^{(\text{spin-1/2})}, \quad (2.24)$$

where $U(g)$ denotes a representation of the group $\text{SO}(3)$ acting on the physical spin-1 indices, and $V(g)$ is a representation of $\text{SU}(2)$ acting on the virtual spin- $\frac{1}{2}$ indices. This projective symmetry at the virtual level underpins the topological protection and classification of the AKLT state [53].

In this thesis, the AKLT state was selected as a target for the state preparation protocols presented in Publications [I,II]. It serves as a minimal nontrivial SPT state that requires a unitary quantum circuit of depth scaling linearly with the system size, making its preparation challenging within the constraints of noisy intermediate-scale quantum (NISQ) devices. Due to its SPT character, the AKLT state holds potential as a resource state in measurement-based quantum computation (MBQC) [3].

Specifically, the AKLT state has been shown to support universal measurement-based quantum computation (MBQC) when defined on certain two-dimensional lattices, such as the honeycomb lattice [56]. In these configurations, the entanglement structure of the AKLT state allows, through measurement, for a mapping to a cluster state. This mapping enables the realization of arbitrary quantum computations through sequences of adaptive measurements. These properties establish the AKLT state as a useful resource for MBQC.

2.2 Projected Entangled Pair States

Projected Entangled Pair States (PEPS) are a natural extension of MPS to two-dimensional quantum systems [42]. Like MPS, PEPS provide an efficient ansatz for quantum many-body states that satisfy an area law for entanglement [57, 58]. Each lattice site is associated with a tensor that carries one physical index and several virtual indices con-

necting to neighboring sites. The wave function amplitude $\psi(\mathbf{s}) = \psi(s_{1,1}, \dots, s_{L,L})$ is obtained by contracting all virtual indices in the tensor network:

$$\psi(\mathbf{s}) = \begin{array}{c} \text{Diagram of a 4x4 PEPS tensor network} \\ \text{A 4x4 grid of white circles (sites) connected by horizontal and vertical lines. Orange dots are placed at the intersections of every second row and every second column, representing boundary conditions or specific indices. The network is fully connected between adjacent sites.} \end{array} . \quad (2.25)$$

A fundamental difference between PEPS and MPS lies in the complexity of their contraction. Exact contraction of PEPS is computationally intractable, with a cost that scales exponentially with system size $\mathcal{O}(D^{2L})$. Consequently, practical computations rely on approximate contraction schemes. These can be broadly classified into two categories: methods based on the double-layer environment approximation [59] and methods based on sampling [43].

For infinite systems, translationally invariant iPEPS have become a widely used tool [60–62]. The Corner Transfer Matrix Renormalization Group (CTMRG) method [41, 63, 64] is commonly employed to contract the double-layer environment surrounding a unit cell of tensors:

$$(\rho_i^R)_{\mathbf{1},\mathbf{r}} = \text{Tr}_{j \neq i}(|\psi\rangle\langle\psi|), \quad (2.26)$$

$$= \begin{array}{c} \text{Diagram of a 4x4 PEPS tensor network with a central unit cell highlighted} \\ \text{The same 4x4 grid as in equation (2.25), but with a central 2x2 unit cell highlighted by red lines. Dotted lines connect the central unit cell to the surrounding environment.} \end{array} . \quad (2.27)$$

CTMRG approximates the infinite environment using corner and edge tensors, which are iteratively computed. This approximation enables the evaluation of reduced density matrices ρ_i^R , which can then be used to compute expectation values. This approach offers a favorable balance between accuracy and computational cost, particularly for

ground states of gapped Hamiltonians on infinite lattices.

For finite systems, finite PEPS (fPEPS) are used. The contraction of fPEPS typically relies on double-layer tensor networks [65, 66]. In this setting, both bra and ket PEPS are contracted to evaluate observables and gradients. Each bond in the double-layer network has dimension D^2 , where D denotes the virtual bond dimension of the original PEPS. It can be shown that the contraction cost scales as $\mathcal{O}(D^{10})$ using boundary-MPS or CTMRG methods adapted for finite systems [67]. Although iPEPS calculations suffer from similar bad scaling, the CTMRG algorithm targets a fixed point of the reduced density matrix in the thermodynamic limit. Since the environments are fixed points, errors in intermediate steps can be corrected and do not accumulate as severely as in the finite case. This makes CTMRG particularly well-suited for infinite systems, where convergence to a translationally invariant environment allows for accurate computation of observables even at moderate bond dimensions. In contrast, for finite systems, the boundary effects and the absence of a true fixed point make the contraction less stable and often more sensitive to truncation errors, especially for larger system sizes or critical states.

To mitigate this high computational cost, sampling-based methods have been developed [43, 44, 59]. These methods estimate observables stochastically by sampling configurations $\mathbf{s} = (s_{1,1}, \dots, s_{L,L})$ from the probability distribution $P(\mathbf{s}) = |\psi(\mathbf{s})|^2 / \langle \psi | \psi \rangle$. Each sample requires only a single-layer contraction, which is significantly more efficient than contracting the full double-layer network². For a more in-depth review about PEPS the curious reader can find it in Ref. [40].

Recent studies have demonstrated the effectiveness of sampling-based optimization for PEPS. In particular, Ref. [44] introduces a variational algorithm based on imaginary-time evolution within the PEPS ansatz. To implement sampling-based Time-Dependent Variational Principle (TDVP) methods [68–71], three components are essential. First, it is necessary to sample configurations $\{\mathbf{s}\}$ from the probability distribution $P(\mathbf{s})$. Second, one must compute the wave function amplitudes $\psi(\mathbf{s}) = \langle \mathbf{s} | \psi \rangle$ and their gradients $\frac{\partial \psi(\mathbf{s}; \theta)}{\partial \theta}$. Third, it must be possible to efficiently evaluate $\psi(\mathbf{s}')$, where \mathbf{s}' denotes a configuration obtained by modifying a small subset of entries in \mathbf{s} .

These components allow for updating the tensor parameters such that the wave function approximately follows the imaginary-time Schrödinger equation. A key advantage of PEPS over other variational ansätze, such as Neural Quantum States [70], is the efficient computation of $\psi(\mathbf{s}')$ using environment tensors. In contrast, for Neural Quantum States, $\psi(\mathbf{s}')$ must be evaluated separately for each term in the Hamiltonian. This results in

²Although double-layer contractions are required to sample from a PEPS, the environment bond dimension can be kept small. Errors arising from sampling an approximate distribution can be systematically corrected using importance sampling.

a significant computational overhead in NQS due to the large number of wave function evaluations required.

Sec. 3.2 presents a detailed discussion of the time evolution procedure, with particular emphasis on the limitations of real-time evolution using PEPS in Sec. 3.3. It is shown that sampling-based approaches for real-time evolution require a number of samples that increases linearly with the number of parameters in the ansatz.

The following section addresses challenges associated with the core computational task of contracting single-layer finite PEPS.

2.2.1 Contractability and Additional Unpublished Algorithmic Improvements

Contracting a single-layer PEPS $\langle \mathbf{s}|\psi \rangle$ for a specific sample has a computational cost that scales as $\mathcal{O}(D^7)$ when using boundary-MPS techniques. This estimate assumes that the boundary-MPS bond dimension satisfies $D_c = \mathcal{O}(D)$ [44, 72], as discussed later in the section. Since the number of samples required for accurate estimation of observables does not grow with D^3 , this approach significantly reduces the total computational cost.

However, this favorable scaling is not universal. As first demonstrated in Refs. [73, 74], the bond dimension D_c for the boundary-MPS can increase exponentially with the linear system size L in the case of random PEPS. Prior work by Gonzalez *et al.* [73] shows that introducing a constant shift to the random tensors restores contractibility. They identify a phase transition in the entanglement of the boundary-MPS required to contract the single-layer PEPS.

In Publication [III], it is demonstrated that this transition corresponds to an entanglement transition in the quantum state represented by the PEPS. On one side of the transition, the PEPS is as highly entangled as permitted by its bond dimension. On the other side, the quantum state is a product state. This understanding is used in Publication [III] to propose an improved random initialization technique for the PEPS.

To address this issue of contractibility, a diagnostic quantity referred to as the spectral entropy H_c was introduced in Publication [III]. This entropy quantifies the contraction difficulty of single-layer PEPS and can be used to assess the feasibility of applying sampling-based optimization methods. Furthermore, H_c could offer insight into the complexity of simulating measurement-based state preparation protocols in two dimensions and could be a starting point to transfer the knowledge gained in Publications [I,II] from 1D into 2D. In this section, we provide insights into the interpretation of H_c and

³The error of a Monte Carlo estimate can be computed with $\sqrt{\text{Var}(O)/N_s}$, where $\text{Var}(O)$ is the variance of the observable and N_s is the number of independent samples. Notably, this expression is independent of how the wave function is represented.

use these insights to propose a modification to the standard contraction algorithm that should achieve a 50% speedup in single-layer contractions.

The boundary-MPS method enables the evaluation of the wave function for a fixed sample:

$$\psi(\mathbf{s}) = \begin{array}{c} \text{Diagram of a 4x4 grid of circles with orange dots at the corners and intersections} \end{array}, \quad (2.28)$$

$$= E_1^t \begin{array}{c} \text{Diagram of a 4x4 grid of circles with orange dots at the corners and intersections} \end{array} \\ M_2 \begin{array}{c} \text{Diagram of a 4x4 grid of circles with orange dots at the corners and intersections} \end{array} \\ = M_3 \begin{array}{c} \text{Diagram of a 4x4 grid of circles with orange dots at the corners and intersections} \end{array}, \quad (2.29)$$

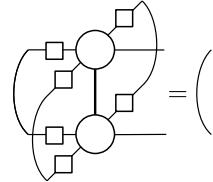
$$M_4 \begin{array}{c} \text{Diagram of a 4x4 grid of circles with orange dots at the corners and intersections} \end{array} \\ E_5^b \begin{array}{c} \text{Diagram of a 4x4 grid of circles with orange dots at the corners and intersections} \end{array} \\ = E_1^t M_2 M_3 M_4 E_5^b. \quad (2.30)$$

Here, E_1^t and E_5^b are boundary MPS, and M_i are matrix product operators (MPOs). The single-layer contraction is performed by sequentially applying MPOs from top to bottom $E_i^t = E_{i-1}^t M_i$. Without truncation, the bond dimension of the boundary MPS grows exponentially $D_c(E_i^t) = D^i$.

In practice, this growth is mitigated by truncating the bond dimension D_c , introducing an approximation error determined by the discarded singular values. Empirically, it has been observed in Publication [III] that relatively small values of D_c suffice. This observation is better understood using the Vidal gauge [75, 76].

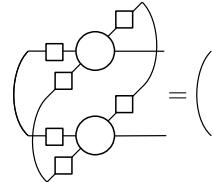
In the Vidal gauge, the tensors are transformed such that the diagonal matrices containing the singular values, denoted S , are explicitly isolated. When these matrices are contracted with the tensor along all but one of its links, the resulting tensor becomes an

isometry with respect to the remaining free link:



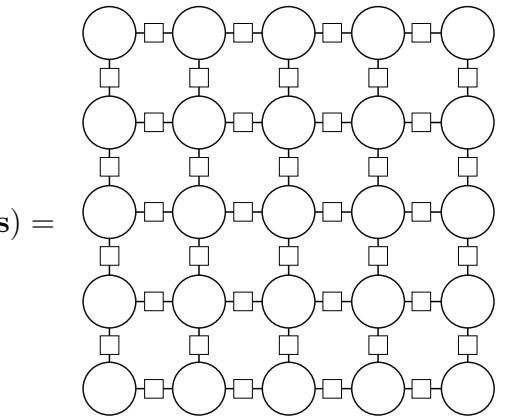
$$= \begin{pmatrix} & \\ & \end{pmatrix} \quad (2.31)$$

For single-layer PEPS, this structure can be adapted as follows:



$$= \begin{pmatrix} & \\ & \end{pmatrix} \quad (2.32)$$

Using the algorithm described in [75, 77], the single-layer PEPS can be transformed into this form for each sample:



$$\psi(\mathbf{s}) = \begin{array}{|c|c|c|c|c|} \hline & \circ & \circ & \circ & \circ \\ \hline \circ & | & | & | & | \\ \hline & \circ & \circ & \circ & \circ \\ \hline \circ & | & | & | & | \\ \hline & \circ & \circ & \circ & \circ \\ \hline \circ & | & | & | & | \\ \hline & \circ & \circ & \circ & \circ \\ \hline \circ & | & | & | & | \\ \hline & \circ & \circ & \circ & \circ \\ \hline \end{array} . \quad (2.33)$$

In Publication [III] we demonstrate that the spectral entropy of the singular values, defined as $H_c = -\sum_i S_i^2 \log(S_i^2)$, correlates with contraction difficulty. This can be observed in Fig. 3,4 of the Publication.

To better understand this, one can contract the vertical singular values into the MPO tensors. This can be done, for example, by computing their square roots and multiplying each square root pair to the corresponding sides (marked red in Eq. 2.35). This operation does not alter the contraction difficulty, as MPO-MPS multiplications are invariant under

such gauge transformations. The resulting form is:

$$\psi(\mathbf{s}) = \begin{array}{c} E_1^{t'} \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ S_1 \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ M'_2 \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ S_2 \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ M'_3 \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ S_3 \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ M'_4 \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ S_4 \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \\ \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \quad \downarrow \\ E_5^{b'} \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \square \circlearrowleft \end{array} , \quad (2.34)$$

$$= E_1^{t'} S_1 M'_2 S_2 M'_3 S_3 M'_4 E_5^{b'} . \quad (2.35)$$

In this structure, the diagonal matrices S_i reduce entanglement⁴ in the boundary MPS E_i^t , while the operators M'_i increase it. This behavior is illustrated in Fig. 2.6a, which shows the entanglement entropy of the boundary MPS for a ground state of the J_1 - J_2 model. The entropy increases upon application of M'_i and decreases after the insertion of S_i . Furthermore, the increase in entropy caused by M'_i can be estimated using the operator bipartite entanglement entropy. This estimate is compared with the actual entropy increase in Fig. 2.6b, where a strong correlation is observed despite the simplicity of the approximation. This highlights the role of the singular value spectrum in determining contractibility. A steeper spectrum implies that the diagonal S matrices remove entanglement from the boundary MPS, while the MPOs M'_i increase the entanglement. This interplay is effectively captured by the spectral entropy H_c .

In the most extreme case, if the spectrum S is given by $S = [1, 0, 0, \dots]$, its application effectively converts the boundary MPS into a product state. As the spectrum becomes flatter, this effect diminishes. A completely flat spectrum, in particular, leaves the entropy of the boundary MPS unchanged. In contrast, applying the MPO M'_i generally increases the entanglement of the boundary MPS. In the worst case, this increase can reach $\log(D)$. As shown in Fig. 2.6b, we find that, in practice, the growth in entanglement is well approximated by the operator entanglement of the MPO.

It is important to note that this observation is heuristic rather than a strict mathematical statement. While the correlation generally holds in practical settings, it might be

⁴Note that the entanglement in the boundary MPS E_i^t is not directly related to the physical entanglement in the quantum state.

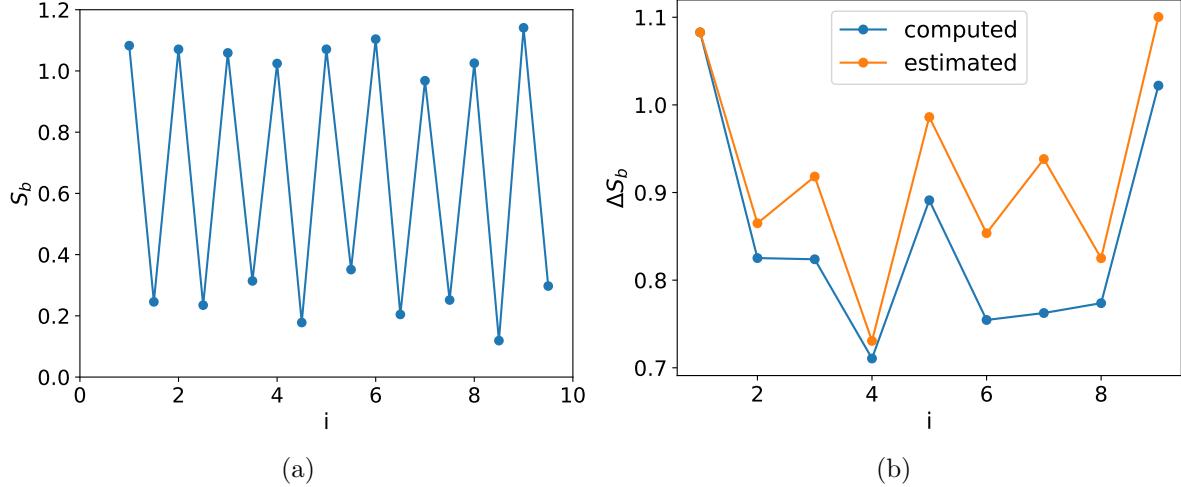


Figure 2.6.: (a) Entanglement entropy of the boundary MPS during the contraction of a single-layer PEPS. The entropy increases after the application of MPO tensors M'_i and decreases following the insertion of diagonal singular value tensors S_i . (b) Comparison between the observed increase in entanglement entropy and the estimated operator bipartite entanglement entropy for each MPO tensor M'_i .

possible to devise a single-layer PEPS for which this relationship does not hold.

To consistently maintain a low-entanglement boundary MPS throughout the contraction process, it is therefore sensible to perform the contraction in the following order:

$$\psi(\mathbf{s}) = (E_1^{l'} S_1)(M'_2 S_2)(M'_3 S_3)(M'_4 S_4) E_5^b. \quad (2.36)$$

This ordering preserves a boundary MPS with low entanglement, which can be efficiently approximated using a small bond dimension.

In Publication [III], this gauge structure was not explicitly enforced. Instead, the PEPS were initialized in a symmetric gauge:

$$\psi(s) = (E_1^{l'} \sqrt{S_1})(\sqrt{S_1} M'_2 \sqrt{S_2})(\sqrt{S_2} M'_3 \sqrt{S_3})(\sqrt{S_3} M'_4 \sqrt{S_4})(\sqrt{S_4} E_5^b), \quad (2.37)$$

and the gauge was left unchanged during optimization.

Fig. 2.7 presents the contraction error:

$$\Delta\psi(D_c, D_c^{\max}) = \frac{\psi(D_c^{\max}) - \psi(D_c)}{\psi(D_c^{\max})}, \quad (2.38)$$

as a function of the contraction bond dimension D_c for several gauge choices. The blue curve corresponds to the original, unaltered gauge. The yellow points represent the symmetric gauge used in Publication [III], which performs identically to the unchanged

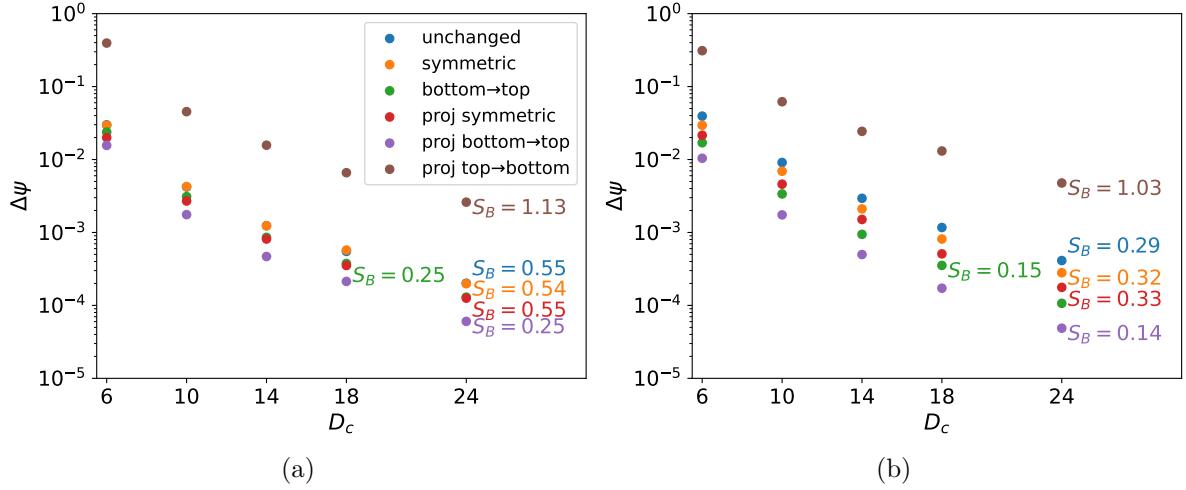


Figure 2.7.: Contraction error as a function of the bond dimension D_c for different gauges. (a) Ground state of the $J_1 - J_2$ model at $J_2 = 0.58$, $L = 10$ and $D = 6$. (b) Random PEPS with an algebraically decaying spectrum $S_i = i^{-1.8}$. The original gauge (blue) and the symmetric gauge used in [III] (yellow) yield nearly identical performance. A directional gauge where singular values are multiplied from bottom to top (green) leads to a modest improvement. (red) Further gains are achieved by first projecting the PEPS to the sample and then applying the symmetric gauge. (violet) The best performance is obtained by contracting each pair as $(M'_i S_i)$, reducing the required bond dimension by 20% for both random PEPS and the ground state of the $J_1 - J_2$, thereby lowering contraction time by approximately 50%. In contrast, (brown) applying the singular values in the reversed order $(S_{i-1} M'_i)$ leads to the poorest performance.

version. A modest improvement is observed when the singular value spectrum is consistently multiplied in a fixed direction, specifically from bottom to top, as indicated by the green data point.

A more substantial improvement is obtained by first projecting the single-layer PEPS for a given sample and then applying the symmetric gauge. This approach outperforms the symmetric gauge applied directly to the unprojected PEPS. The most effective configuration is achieved using the directional scheme described earlier, where singular values and MPOs are contracted as $(M'_i S_i)$ at each step. This method reduces the required bond dimension by approximately 20% to reach the same contraction error as the original gauge. Given that the contraction cost scales as $\mathcal{O}(D^4 D_c^3)$, this should translate into a 50% reduction in computational time. This improvement constitutes a significant algorithmic enhancement with immediate practical implications for large-scale PEPS simulations. In contrast, the worst performance is observed when singular values are applied in the reversed order, $(S_{i-1} M'_i)$.

It is worth noting that this analysis was also conducted for randomly initialized PEPS across various system sizes and bond dimensions. Consistently, a comparable reduction of approximately 20% in the effective bond dimension was observed.

Although computing the gauge is not particularly expensive, it requires iterative calculations at each edge of the PEPS lattice until convergence. Specifically, this involves singular value and eigenvalue decompositions of $D \times D$ matrices, resulting in an overall cost of $\mathcal{O}(N_{\text{iter}} L^2 D^3)$ where N_{iter} are the number of iterations needed. In the examples analyzed, it was sufficient to set $N_{\text{iter}} = 1$ to achieve the 20% reduction in computational cost.

CHAPTER 3

Time-Dependent Variational Principle

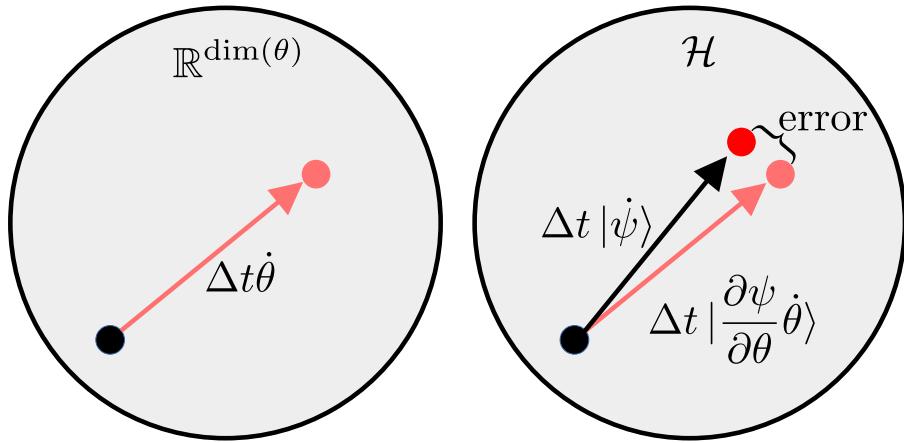


Figure 3.1.: Illustration of a single step of the Time-Dependent Variational Principle (TDVP) as described in Eq. (3.4). The left panel depicts the update $\Delta t \dot{\theta}$ in parameter space $\mathbb{R}^{\dim(\theta)}$. The right panel shows the corresponding evolution in Hilbert space \mathcal{H} : the projected tangent vector $\Delta t |\partial_\theta \psi \dot{\theta}\rangle$ versus the true time evolution $\Delta t |\dot{\psi}\rangle$. The difference between these two vectors quantifies the TDVP projection error.

In Publications [I,III], the time evolution of a parameterized wave function was a central technique. To achieve this, the wave function must follow the Schrödinger equation. The Time-Dependent Variational Principle (TDVP) [68–71] provides a systematic approach for updating the parameters θ of a variational wave function $|\psi(\theta)\rangle$ such that its evolution remains as close as possible to that dictated by the Schrödinger equation. This chapter examines the application of TDVP to MPS (Sec. 3.1) and sampling-based methods (Sec. 3.2) that are used for PEPS [44] and Neural Network States [70]. Furthermore, in Sec. 3.3, I present unpublished results demonstrating that the number of

samples required to accurately perform real-time evolution scales linearly with the number of parameters in the PEPS ansatz and strongly with the amount of entanglement in the PEPS. In the same section, I propose a potential strategy to mitigate this scaling challenge.

Note that although the first two sections of this chapter focus on imaginary-time evolution, the extension to real-time evolution is straightforward and can be implemented by inserting a factor of the imaginary unit into Eq. (3.1).

The equations governing the evolution of the parameters θ are derived from the imaginary-time Schrödinger equation:

$$|\dot{\psi}(\theta)\rangle = -H |\psi(\theta)\rangle , \quad (3.1)$$

under the constraint that the wave function remains normalized,

$$|\psi(\theta)\rangle \rightarrow \frac{|\psi(\theta)\rangle}{\sqrt{\langle\psi(\theta)|\psi(\theta)\rangle}} . \quad (3.2)$$

Applying this constraint and computing partial derivatives yield:

$$\frac{1}{\sqrt{Z}} \left[\left| \frac{\partial\psi}{\partial\theta} \right\rangle - |\psi\rangle \frac{1}{Z} \langle\psi| \frac{\partial\psi}{\partial\theta} \right] \dot{\theta} = -H \frac{|\psi\rangle}{\sqrt{Z}} . \quad (3.3)$$

Since this equation has no exact solution, the goal is to determine $\dot{\theta}$ by minimizing the squared error:

$$\min_{\dot{\theta}} \frac{1}{Z} \left\| \left[\left| \frac{\partial\psi}{\partial\theta} \right\rangle - |\psi\rangle \frac{1}{Z} \langle\psi| \frac{\partial\psi}{\partial\theta} \right] \dot{\theta} + H |\psi\rangle \right\|^2 . \quad (3.4)$$

The optimal parameter update $\dot{\theta}$ depends on the specific variational ansatz. In general, it is obtained by solving the following system of equations:

$$G_{ii'} = \langle \frac{\partial\psi}{\partial\theta_i} | \frac{\partial\psi}{\partial\theta_{i'}} \rangle - \langle \frac{\partial\psi}{\partial\theta_i} | \psi \rangle \langle \psi | \frac{\partial\psi}{\partial\theta_{i'}} \rangle , \quad (3.5)$$

$$F_i = \langle \frac{\partial\psi}{\partial\theta_i} | H | \psi \rangle - \langle \frac{\partial\psi}{\partial\theta_i} | \psi \rangle \langle H | \psi \rangle , \quad (3.6)$$

$$\dot{\theta} = -G^{-1}F , \quad (3.7)$$

where G is the geometric tensor, and F denotes the gradient of the energy.

Depending on the ansatz, direct inversion or full storage of the geometric tensor G may be impractical. The tensor has dimensions $N_p \times N_p$, where N_p is the number of variational parameters. Since N_p can be large, explicit handling of G is often infeasible. However,

for an MPS as demonstrated in Sec. 3.1, G simplifies to a diagonal form when using the correct parametrization. Additionally, Sec. 3.2 presents a sampling-based approach that allows solving Eq. (3.4) approximately without explicitly constructing G .

3.1 For an MPS

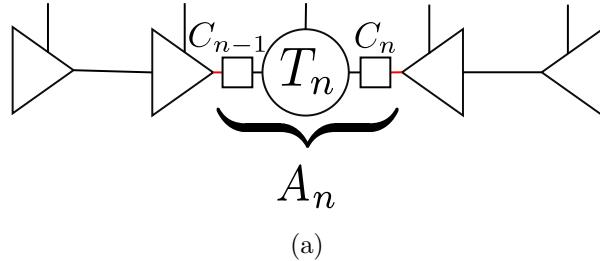


Figure 3.2.: MPS in its orthogonal form.

As indicated previously, explicitly storing and inverting the geometric tensor G for an MPS is computationally impractical due to its size. Specifically, for an MPS parameterized directly by tensors $\theta = \{T_1, \dots, T_L\}$, G contains L^2 elements of the form $\langle \frac{\partial \psi}{\partial T_n} | \frac{\partial \psi}{\partial T_{n'}} \rangle$. Fortunately, by employing a suitable reparameterization of the MPS into its orthogonal form (see Fig. 3.2), most of these elements can be systematically set to zero [69, 71, 78]. This reparameterization reduces the computational complexity significantly. However, the tensors A_n at the orthogonality center are not independent, implying that the resulting set of differential equations must be solved sequentially.

Initially, evaluating $\langle \frac{\partial \psi}{\partial A_n} | \frac{\partial \psi}{\partial A_{n'}} \rangle$ appears equally challenging. The key simplification emerges by imposing the constraint:

$$0 = \boxed{\begin{array}{c} \text{triangular symbol} \\ \text{---} \\ A_n \end{array}} \quad . \quad (3.8)$$

This condition ensures that the norm of the MPS remains unchanged when evolved by \dot{A} :

$$\dot{A}_n \frac{\partial \langle \psi | \psi \rangle}{\partial A_n} = \langle \psi | \frac{\partial \psi}{\partial A_n} \dot{A}_n \rangle = 0. \quad (3.9)$$

Furthermore, this constraint also enforces most of the terms in G to zero:

$$\langle \frac{\partial \psi}{\partial A_n} \dot{A}_n | \frac{\partial \psi}{\partial A_{n'}} \dot{A}_{n'} \rangle = \delta_{n,n'} \langle \frac{\partial \psi}{\partial A_n} \dot{A}_n | \frac{\partial \psi}{\partial A_n} \dot{A}_n \rangle. \quad (3.10)$$

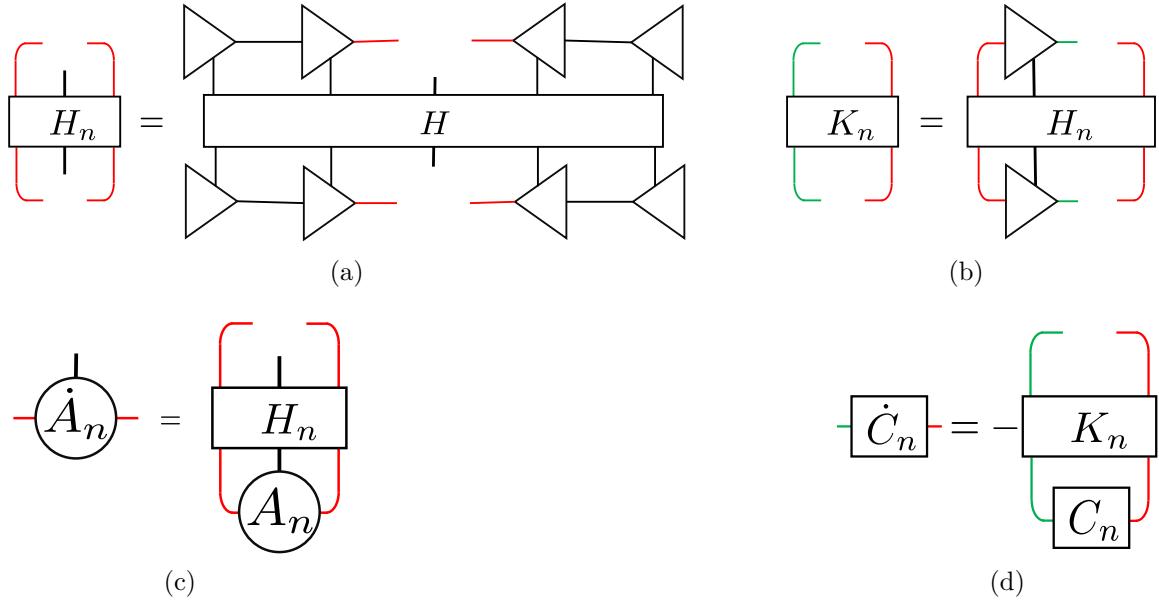


Figure 3.3.: TDVP equations for an MPS. (a,c) H_n is the energy environment to evolve A_n , and (b,c) K_n is the energy environment to evolve C_n .

Consequently, explicit calculation and inversion of the geometric tensor G becomes unnecessary.

Under these conditions, one can derive the ordinary partial differential equations. These coupled equations are typically solved using a sweeping approach, sequentially evolving tensors A_n and C_n as follows:

$$A_n(\Delta t) = e^{-\Delta t H_n} A_n(0), \quad (3.11)$$

$$C_n(\Delta t) = e^{\Delta t K_n} C_n(0), \quad (3.12)$$

where H_n and K_n are defined in Fig. 3.3. This procedure can be interpreted as evolving the A_n tensors forward in time and the C_n tensors backward in time. This correction is necessary because evolving A_n forward in time also implicitly advances the C_n tensors. Each C_n tensor is shared between neighboring A_n tensors and is therefore affected twice when only the A_n tensors are updated. To compensate for this, the C_n tensors must be explicitly evolved backward in time. This correction ensures consistency in the time evolution of the full wave function.

Moreover, the Density Matrix Renormalization Group (DMRG) algorithm for ground state search emerges as the imaginary-time evolution limit of these equations, corresponding to $\Delta t \rightarrow \infty$.

3.2 For General Ansätze using Sampling

The previously discussed procedure for time evolution of MPS is unfortunately not applicable to most wave function ansätze. Therefore, an alternative method known as sampling time-dependent variational principle (sampling TDVP), also referred to as Stochastic Reconfiguration or Quantum Natural Gradient, was developed [44, 68–71]. Alternatively, stochastic gradient descent could be employed for the determination of the ground state; however, this method typically encounters difficulties with convergence to local minima. Therefore, sampling TDVP has been widely adopted for optimizing Neural Quantum States and is particularly relevant for optimizing finite projected entangled pair states (PEPS), as demonstrated in Publication [III].

In order to be able to use this method, the wave function ansatz must satisfy three essential conditions. First, it must allow the evaluation of the wave function amplitude $\psi(\mathbf{s}; \theta) = \langle \mathbf{s} | \psi(\theta) \rangle$. Second, it must enable computation of the gradients with respect to parameters, $\frac{\partial \psi(\mathbf{s}; \theta)}{\partial \theta}$. Third, it must facilitate sampling from the associated probability distribution $P_\psi(\mathbf{s}) = |\psi(\mathbf{s})|^2$.

3.2.1 Expectation Values

Before evolving $|\psi\rangle$, let us first revisit how to compute expectation values with the help of sampling:

$$\langle O \rangle = \langle \psi | \hat{O} | \psi \rangle, \quad (3.13)$$

$$= \sum_{\mathbf{s}, \mathbf{s}'} \frac{1}{Z} \psi^*(\mathbf{s}) \langle \mathbf{s} | \hat{O} | \mathbf{s}' \rangle \psi(\mathbf{s}'), \quad (3.14)$$

$$= \sum_{\mathbf{s}} P_\psi(\mathbf{s}) \underbrace{\sum_{\mathbf{s}'} \langle \mathbf{s} | \hat{O} | \mathbf{s}' \rangle \frac{\psi(\mathbf{s}')}{\psi(\mathbf{s})}}_{O_{\mathbf{s}}^{\text{loc}}}, \quad (3.15)$$

$$= \langle O_{\mathbf{s}}^{\text{loc}} \rangle_{\mathbf{s} \in P_\psi}. \quad (3.16)$$

The quantity $O_{\mathbf{s}}^{\text{loc}}$ denotes the local estimator of the observable \hat{O} corresponding to the configuration \mathbf{s} . Its computation is practical only if the observable \hat{O} has limited support. In this case, the summation over configurations \mathbf{s}' simplifies substantially. Specifically, only configurations \mathbf{s}' differing from \mathbf{s} by a few elements within the support of \hat{O} need

to be considered. For example if $\hat{O} = \sigma_x$ then:

$$O_{\mathbf{s}}^{\text{loc}} = \frac{\langle \mathbf{s} | \sigma_x | \psi \rangle}{\langle \mathbf{s} | \psi \rangle}, \quad (3.17)$$

$$= \frac{\langle \mathbf{s}' | \psi \rangle}{\langle \mathbf{s} | \psi \rangle}. \quad (3.18)$$

Hamiltonians typically have a number of terms that scale proportionally with system size. Therefore, evaluating the expectation value of energy for each sampled configuration \mathbf{s} necessitates computing $\psi(\mathbf{s}')$ multiple times, scaling with the system size. This requirement makes computations for large systems particularly demanding, posing significant challenges if no computational shortcuts are employed.

Fortunately, the finite PEPS ansatz employed in Publication [III] allows many intermediate calculations used to obtain $\psi(\mathbf{s})$ to be reused when evaluating $\psi(\mathbf{s}')$. This strategy considerably reduces computational overhead and thus facilitates dealing with larger system sizes that are very difficult to deal with when, for example, using Neural Quantum states.

3.2.2 Evolution

Computing the geometric tensor G , defined in Eq. (3.5), exactly is typically infeasible for most ansätze due to the requirement of summing over the entire Hilbert space. Therefore, an approximation using Monte Carlo sampling is employed. The minimization problem in Eq. (3.4) is reformulated to consider only the relevant sampled subspace, reducing the dimension to the number of samples N_s :

$$\min_{\dot{\theta}} \sum_s \frac{1}{Z} \left\| \sum_i \underbrace{\left[\langle \mathbf{s} | \frac{\partial \psi}{\partial \theta_i} \rangle - \langle \mathbf{s} | \psi \rangle \frac{1}{Z} \langle \frac{\partial \psi}{\partial \theta_i} | \psi \rangle \right]}_{=: \tilde{O}_{\mathbf{s},i}} \dot{\theta}_i + \underbrace{\langle \mathbf{s} | H | \psi \rangle}_{=: \tilde{E}_s^{\text{loc}}} \right\|^2, \quad (3.19)$$

$$\min_{\dot{\theta}} \sum_s P_\psi(\mathbf{s}) \left\| \sum_i \underbrace{\frac{\tilde{O}_{\mathbf{s},i}}{\psi(\mathbf{s})}}_{=: O_{\mathbf{s},i}} \dot{\theta}_i + \underbrace{\frac{\tilde{E}_s^{\text{loc}}}{\psi(\mathbf{s})}}_{=: E_{\mathbf{s}}^{\text{loc}}} \right\|^2, \quad (3.20)$$

$$\min_{\dot{\theta}} \frac{1}{N_s} \sum_k^{N_s} \left\| \sum_i O_{\mathbf{s}^{(k)},i} \dot{\theta}_i + E_{\mathbf{s}^{(k)}}^{\text{loc}} \right\|^2 \quad \mathbf{s}^{(k)} \in P_\psi(\mathbf{s}). \quad (3.21)$$

This method effectively solves the optimization problem within a sampled region of the Hilbert space. The accuracy of the method thus relies on the generalization of the obtained solution $\dot{\theta}$ from the sampled subset $\{\mathbf{s}^{(k)}\}$ to the entire Hilbert space.

There are two primary strategies for solving this equation. The first, commonly employed until recently, approximates the geometric tensor G using N_s samples:

$$\dot{\theta} = -(\underbrace{O^\dagger O}_{=:G})^{-1} O^\dagger E^{\text{loc}}, \quad (3.22)$$

$$G_{i',i} = \sum_s P_\psi(\mathbf{s}) O_{\mathbf{s},i'}^* O_{\mathbf{s},i}. \quad (3.23)$$

This approach, however, is limited by the dimensions of matrix G , which are $N_p \times N_p$, where N_p is the number of parameters in the ansatz. Storing G becomes impractical for parameter-intensive ansätze, necessitating slow iterative solvers. It is important to note that the sampled G matrix rank is at most N_s .

Recently, Chen and Heyl [45] proposed an alternative method known as minimum-step stochastic reconfiguration (minSR) in the context of neural quantum states. Their method addresses computational difficulties arising for parameter-rich ansätze by leveraging the low rank of G when parameters outnumber samples:

$$\dot{\theta} = O^\dagger (\underbrace{Q O^\dagger}_{=:T})^{-1} E^{\text{loc}}, \quad (3.24)$$

$$T_{\mathbf{s},\mathbf{s}'} = \sum_i O_{\mathbf{s},i} O_{\mathbf{s}',i}^*. \quad (3.25)$$

Here, the matrix T acts exclusively on the sampled subspace and thus has dimensions $N_s \times N_s$. Consequently, it is significantly more computationally efficient to store and invert compared to the $N_p \times N_p$ geometric tensor G . This reduction in computational complexity substantially improves the feasibility of simulations for larger system sizes. The effectiveness of this method for finite PEPS has been demonstrated in Publication [III], reporting performance enhancements of up to 40%.

3.3 Real-Time Evolution of PEPS: Challenges and Possible Solutions (Unpublished)

In this section, we investigate the feasibility of performing real-time evolution of quantum systems using the finite PEPS formalism combined with sampling-based methods. We analytically show that, for product states, the number of samples required scales linearly

with system size. Furthermore, numerical results for random PEPS indicate that the sample complexity scales linearly with the number of variational parameters and exhibits a strong dependence on the state's entanglement. To address this challenge, we propose a hybrid approach that integrates the simple or full update scheme with the TDVP-based sampling method.

Readers are encouraged to consult Sec. 2.2 beforehand, as it introduces essential PEPS-specific concepts relevant to the discussion.

We begin with a brief overview of the two primary algorithms for PEPS time evolution based on the Time-Evolving Block Decimation (TEBD) framework. In this setting, time evolution is mapped onto a sequence of quantum gates applied to the PEPS. These gate applications increase the bond dimension, necessitating a truncation step. Two established schemes for truncation are the simple update and the full update algorithms.

The simple update algorithm for finite PEPS applies each local two-site imaginary- or real-time evolution gate to the local tensors individually. It then uses only the diagonal bond-weight matrices (Schmidt coefficients) in the Vidal gauge to approximate the surrounding environment in a mean-field manner. After applying the gate, the two-site tensor is contracted, followed by a local SVD. The bond is truncated to the target bond dimension by retaining the largest singular values. This method is computationally efficient. However, it neglects loop correlations and introduces systematic errors in long-range entanglement. Consequently, it is not well-suited for accurate real-time evolution.

The full update, in contrast, approximates the environment surrounding each bond prior to truncation. This is typically achieved using methods such as boundary matrix product states or CTMRG. The split two-site tensors are chosen such that they maximize the fidelity with the target evolved state. Although this method yields higher accuracy than the simple update, it incurs significantly greater computational cost. The cost scales unfavorably with both the bond dimension and system size. In many cases, even the full update does not achieve sufficient accuracy, and further variational optimization is required to maximize the fidelity between the truncated and the untruncated PEPS.

While real-time evolution with infinite Projected Entangled Pair States (iPEPS) has advanced significantly in recent years, progress for finite PEPS has been comparatively limited. Variational optimization strategies, such as full updates using environment tensors obtained via CTMRG or boundary Matrix Product Operator (MPO) methods, have been extensively developed for iPEPS [79–81]. These techniques have enabled simulations of time dynamics in infinite systems, particularly when combined with symmetry exploitation, such as $U(1)$ or $SU(2)$, to reduce computational costs and stabilize evolution [82, 83]. Time-evolving block decimation (TEBD) with simple updates remains a

widely used method for short-time dynamics due to its computational efficiency, despite neglecting environmental entanglement [84, 85]. Fidelity-maximizing local variational updates provide improved accuracy and have been employed in the reliable simulation of quantum dynamics across a range of systems. These include studies of correlation spreading in Bose-Hubbard systems [85], critical dynamics in valence bond spin liquids [82], thermal state preparation [80], global quenches in the Heisenberg model [86], and many-body localization [87].

In contrast, real-time evolution for finite PEPS remains largely unexplored in the literature. Foundational techniques, such as Trotter decompositions for TEBD and variational optimization, exist for finite systems and have been applied to ground state computations [79]. However, no published work to date demonstrates an efficient evolution protocol tailored specifically for finite PEPS.

The only known attempt is by Murg *et al.* [88], who time-evolved an 11×11 lattice of hardcore bosons. Due to the use of the full update scheme in combination with variational optimization using sweeping, they are limited to a bond dimension of $D = 4$. The poor scaling of the method is evident, as one time step required 55 hours on 2006-era hardware. Although current hardware would reduce the wall time considerably, the unfavorable scaling with bond dimension and system size remains a fundamental obstacle.

This lack of development highlights a gap in the tensor network literature. Finite PEPS are the natural ansatz for finite-sized, experimentally relevant systems. Yet, robust and scalable algorithms for their real-time evolution are not currently available. Closing this gap is essential for enabling accurate simulations of dynamics in realistic finite geometries.

A natural approach is to adapt the sampling TDVP method employed in Publication [III]. However, as will be demonstrated through a simple example, this method fails without translational invariance. Specifically, the number of samples required scales linearly with system size. Therefore, in this chapter we propose an alternative that combines simple update with the sampling approach.

3.3.1 Gradient Estimation Noise and its Scaling with System Size

Consider a simple example of an N -qubit product state $|\psi(\theta)\rangle = \bigotimes_i^N |\psi_i(\theta^i)\rangle$, evolved under a Hamiltonian consisting of single-site operators, such as $H = \sum_i \sigma_i^z$. Although the time evolution of this product state can be computed analytically, we demonstrate that, under the sampling-based TDVP algorithm, the number of samples required for a single time evolution step scales linearly with system size. This is not immediately obvious and may seem counterintuitive.

Two quantities are required to perform the time evolution: the energy gradient

$$\frac{\partial \langle H \rangle}{\partial \theta} = \frac{\partial \langle \psi(\theta) | [H - \langle H \rangle] | \psi(\theta) \rangle}{\partial \theta}, \quad (3.26)$$

and the geometric tensor

$$G = \left\langle \frac{\partial \psi(\theta)}{\partial \theta} \left| \frac{\partial \psi(\theta')}{\partial \theta'} \right. \right\rangle. \quad (3.27)$$

For clarity, we consider the case where each θ^i is a vector with components θ_j^i , where $i \in \{1, \dots, N\}$ is the spatial index and $j \in \{0, 1\}$ distinguishes components of the local parameter. The wave function is expressed such that $\langle s_i | \psi_i \rangle = \theta_{s_i}^i$. Within the sampling-based TDVP framework, the energy gradient is computed as follows:

$$E_{\mathbf{s}}^{\text{loc}} = \frac{\langle \mathbf{s} | H | \psi \rangle}{\langle \mathbf{s} | \psi \rangle}, \quad O_{(i,j),\mathbf{s}} = \frac{\langle \mathbf{s} | \frac{\partial \psi}{\partial \theta_j^i} \rangle}{\langle \mathbf{s} | \psi \rangle}, \quad (3.28)$$

$$\frac{\partial \langle H \rangle}{\partial \theta_j^i} = \langle O_{(i,j),\mathbf{s}}^* (E_{\mathbf{s}}^{\text{loc}} - \langle E_{\mathbf{s}}^{\text{loc}} \rangle) \rangle_{\mathbf{s} \sim P(\mathbf{s})}. \quad (3.29)$$

For product states, both O and E^{loc} can be computed explicitly:

$$\langle \mathbf{s} | \psi \rangle = \prod_i \theta_{s_i}^i, \quad E_{\mathbf{s}}^{\text{loc}} = \sum_i^N (1 - 2s_i), \quad O_{(i,j),\mathbf{s}} = \delta_{j,s_i} \frac{1}{\theta_{s_i}^i}. \quad (3.30)$$

The average of the local energy reads:

$$\langle E_{\mathbf{s}}^{\text{loc}} \rangle = \sum_i^N (1 - 2p_i), \quad p_i := \sum_{\mathbf{s}} P(\mathbf{s}) s_i = |\theta_1^i|^2, \quad 1 - p_i = |\theta_0^i|^2, \quad (3.31)$$

To analyze the scaling behavior of the gradient average, consider the expectation value of the gradient estimator:

$$\langle O_{(i,j),\mathbf{s}}^* (E_{\mathbf{s}}^{\text{loc}} - \langle E_{\mathbf{s}}^{\text{loc}} \rangle) \rangle = \sum_s P(\mathbf{s}) \delta_{j,s_i} \frac{1}{\theta_j^{i*}} \left(\sum_k^N (1 - 2s_k) - \sum_k^N (1 - 2p_k) \right), \quad (3.32)$$

$$= \frac{|\theta_j^i|^2}{\theta_j^{i*}} (2(p_i - j)), \quad (3.33)$$

$$= 2 \theta_j^i (p_i - j). \quad (3.34)$$

In Eq. (3.32), most terms cancel out when applying the definition in Eq. 3.31. This already reveals that the number of terms contributing to the gradient estimator scales

linearly with N . While cancellations occur at the level of the mean, this is not the case for the variance.

To quantify the error, we define the estimator $X_{i,j}(\mathbf{s}) = O_{(i,j),\mathbf{s}}(E_{\mathbf{s}}^{\text{loc}} - \langle E_{\mathbf{s}}^{\text{loc}} \rangle)$ and compute its variance:

$$\text{Var}[X_{i,j}] = \langle X_{i,j}^2 \rangle - \langle X_{i,j} \rangle^2 , \quad (3.35)$$

$$X_{i,j}^2 = \delta_{j,s_i} \frac{1}{|\theta_j^i|^2} \left(\sum_k (1 - 2s_k) - \sum_k (1 - 2p_k) \right)^2 , \quad (3.36)$$

$$\langle X_{i,j}^2 \rangle = \mathbb{E}_{s_{\neq i}} [(A + B)^2] , \quad (3.37)$$

where we define

$$A_{i,j} = 2(p_i - j) , \quad (3.38)$$

$$B_i = \sum_{k \neq i} [(1 - 2s_k) - (1 - 2p_k)] . \quad (3.39)$$

Since $\langle B_j \rangle = 0$ and $\text{Var}(1 - 2s_k) = 4p_k(1 - p_k)$, we obtain

$$\langle X_{i,j}^2 \rangle = 4(p_i - j)^2 + \sum_{k \neq i} 4p_k(1 - p_k) , \quad (3.40)$$

$$\text{Var}[X_{i,j}] = 4(p_i - j)^2(1 - |\theta_j^i|^2) + \sum_{k \neq i} 4p_k(1 - p_k) . \quad (3.41)$$

This result shows that although the estimator is localized at site i , its variance involves a sum over all other sites $k \neq i$. For generic parameters p_k , this sum contributes a term that scales linearly with the system size N :

$$\text{Var}[X_{i,j}] = \mathcal{O}(N) . \quad (3.42)$$

Consequently, the number of samples required to achieve a fixed error in the gradient estimate scales linearly with N . This result holds despite the simplicity of the underlying system. It highlights that even in trivial cases, the Monte Carlo estimation of TDVP gradients may lead to significant sampling overhead. Note that the same analysis can be done for the geometric tensor, where no such scaling is encountered.

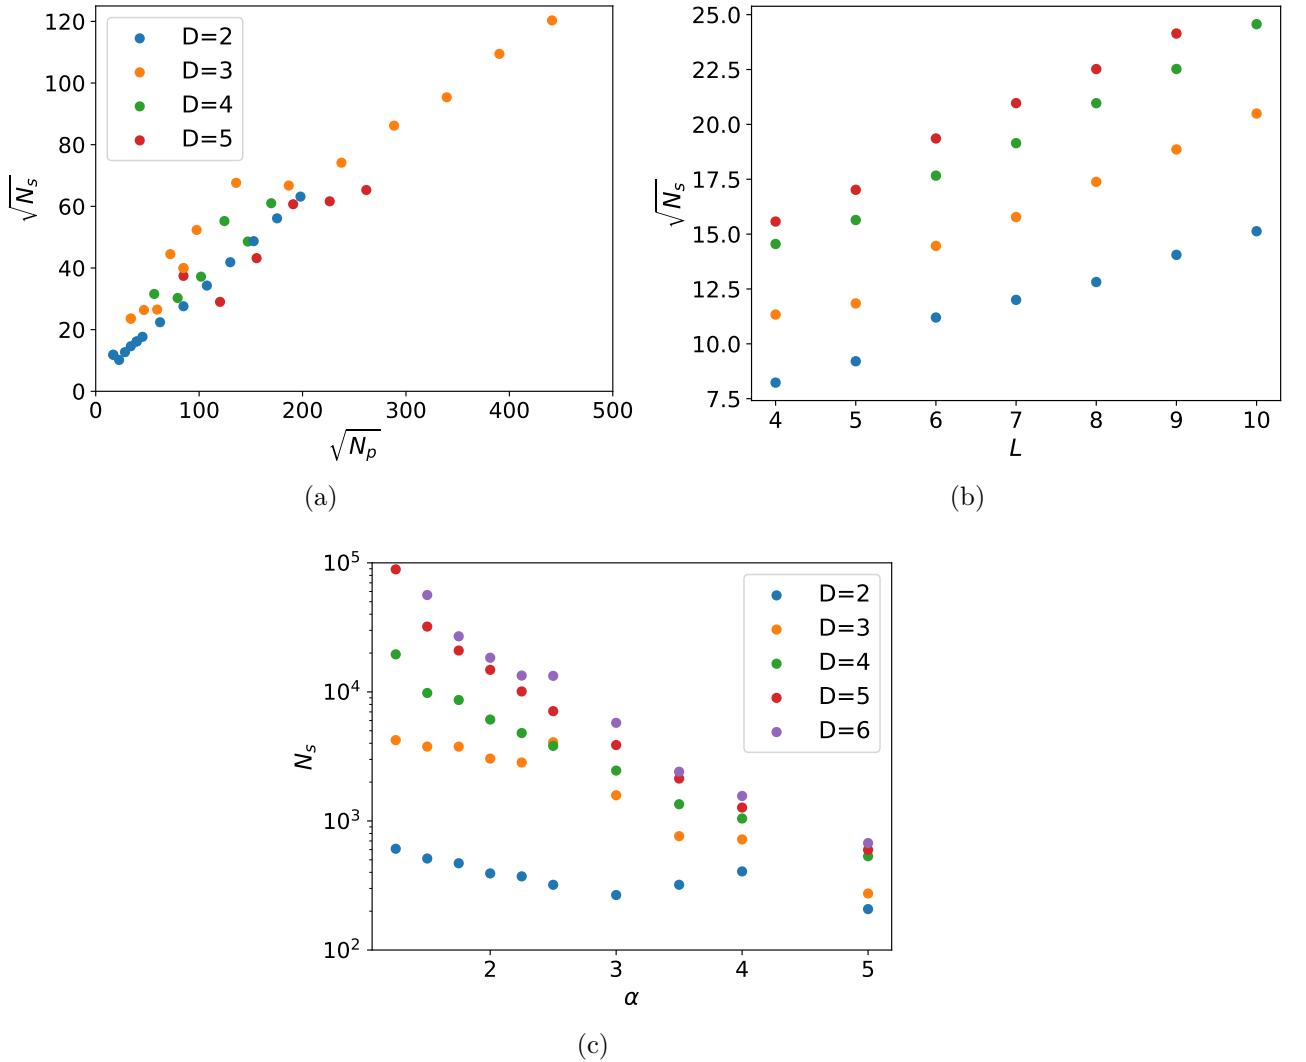


Figure 3.4.: Number of samples N_s needed to obtain a Median noise-to-signal ratio of the gradients of one for a PEPS with bond dimension D for different system sizes. The tensors were drawn from a random distribution, and their open indices were multiplied by an algebraically decaying spectrum $S_i = i^{-\alpha}$ with (a) $\alpha = 3$ and (b) $\alpha = 10$. For small α (subfigure (a)), the number of samples N_s scales linearly with the number of parameters N_p . For large α (subfigure (b)), the PEPS is close to a product state. In this case, N_s scales linearly with system size L^2 and exhibits an offset dependent on the bond dimension. (c) N_s is plotted against α for $L = 8$ and $D = 4$.

Random PEPS

This analysis can also be performed numerically for random PEPS. To this end, the noise-to-signal ratios (NSR) of the gradient components are computed as:

$$\text{NSR}_i = \frac{\text{std}(O_{i,s}^* E_s^{\text{loc}})}{\sqrt{N_s} \langle O_{i,s}^* E_s^{\text{loc}} \rangle}, \quad (3.43)$$

where N_s denotes the number of samples. The calculation of the NSR is stable only when its value is significantly smaller than one. Since NSR_i contains many elements, some components may exceed one, or even diverge, and thus distort the average. To mitigate this, the median(NSR) is used to assess the overall noisiness of the gradients.

In principle, a representative estimate of the NSR requires averaging over many independent random PEPS. This process is computationally demanding because a large sample size ($N_s = 10^6$) is necessary to achieve sufficient accuracy¹. In this preliminary investigation, only a single random PEPS was analyzed per system size and bond dimension. To isolate the dependence of the NSR on system size and bond dimension, the following procedure was employed. First, a PEPS tensor with bond dimension $D_{\text{max}} = 6$ was drawn from a Gaussian distribution and used consistently across all experiments. Second, the open indices of the tensor were rescaled by an algebraically decaying spectrum $S_i = i^{-\alpha}$ to tune the entanglement. Third, the bond dimension was truncated to the target value D . Finally, the same tensor was repeated on an $L \times L$ lattice, with random boundary vectors applied to the open edges. This convoluted initialization of the random PEPS ensures that states at different system sizes exhibit similar properties. Consequently, any observed variations in the NSR can only be weakly attributed to differences in the initial states.

Note that in Fig. 2b of Publication [III] the relationship between α and the entanglement of the quantum state is shown explicitly.

Fig. 3.4 shows the number of samples N_s required to achieve a median NSR of one² for a simple Hamiltonian consisting of a sum of σ_z operators, as a function of various quantities. When the entanglement is large (small $\alpha = 3$ in subfigure (a)), N_s scales linearly with the number of parameters N_p . For weakly entangled states (large $\alpha = 10$ in subfigure (b)), N_s scales linearly with the system size L^2 , and is largely independent of bond dimension. This observation suggests that for entangled random PEPS, the number of samples required for accurate gradient estimation scales as $N_s \sim L^2 D^4$.

This trend is further illustrated in Fig. 3.4c, where the number of samples is plotted as a function of α for fixed system size $L = 8$ and various bond dimensions. At large

¹With $N_s = 10^6$ the worst median NSR in our calculations was 0.1.

²In practice, a lower NSR is desirable, but a threshold of one represents the minimum acceptable value.

α , where entanglement is low, the required number of samples is nearly independent of bond dimension. As α decreases and entanglement increases, the number of samples needed grows substantially. This indicates a clear relationship between entanglement and the difficulty of computing accurate gradients for finite PEPS. In the case of $L = 8$ and $D = 6$, the number of samples required to achieve a fixed NSR increases by a factor of approximately 250 when transitioning from weakly entangled to strongly entangled states.

This is significantly more demanding than the linear scaling with system size observed for product states. As such, we do not expect PEPS to perform well for real-time evolution, where accurate gradients are required. To draw definitive conclusions, however, these experiments should ideally be repeated for physically relevant states that are the actual targets of real-time evolution.

It remains unclear whether this unfavorable scaling with entanglement is specific to sampling-based PEPS or also applies to Neural Quantum States (NQS), where time evolution is performed similarly and has demonstrated significant success in simulating out-of-equilibrium dynamics in two dimensions [89]. It would be important to carry out a comparable analysis for NQS to determine whether they exhibit the same sensitivity to entanglement. This would clarify whether the observed scaling is an inherent feature of the underlying poly-linear structure of tensor networks or a more general limitation of variational sampling approaches.

3.3.2 Possible solutions - Hybrid-sampling TDVP approach

As discussed previously, TEBD combined with the simple update (SU), full update (FU), or variational compression techniques such as those introduced in Ref. [88] remains the only viable non-sampling-based approach currently in use. However, the SU and FU methods serve as suboptimal compression schemes and often yield imprecise results. Specifically, these methods are restricted to local truncation of the PEPS. For non-local truncation, variational compression like in Ref. [88] remains the only available option. While more accurate, this approach is computationally prohibitive for large bond dimensions. In particular, Ref. [88] proposes first performing a full update, followed by a variational optimization step that maximizes the overlap between the truncated and untruncated states using a sweeping algorithm. This procedure involves contracting double-layer PEPS networks with an effective bond dimension of DD_{TEBD} , where D_{TEBD} is the bond dimension after TEBD evolution and before truncation.

To address these limitations, we propose using the TEBD approximation with the SU/FU truncation scheme as an initialization step, following the approach of Ref. [88]. We then variationally optimize the PEPS to approximate the time-evolved state using a sampling-based TDVP method. This approach is expected to be less computationally

demanding than a full variational sweeping algorithm employed in Ref. [88].

The objective is to find the PEPS that best represents the time-evolved state $|\psi_{\delta t}\rangle = e^{-iH\delta t}|\psi\rangle$. The proposed method proceeds in two stages:

First, TEBD with either SU or FU is used to produce an approximate state

$$|\psi_{\delta t}^{\text{SU/FU}}\rangle \underset{\text{SU/FU}}{\approx} \text{Trotter}(e^{-iH\delta t})|\psi\rangle.$$

This state serves as a coarse approximation to $|\psi_{\delta t}\rangle$. In the second stage, this approximation is refined using a sampling-based optimization.

We consider two strategies for refinement:

1. Use sampling to optimize the PEPS parameters of $|\psi_{\delta t}^{\text{SU/FU}}\rangle$ such that the distance between $|\psi_{\delta t}^{\text{SU/FU}}\rangle$ and $\text{Trotter}(e^{-iH\delta t})|\psi\rangle$ is minimized. This mimics the sweeping procedure from Murg *et al.* [88] but replaces costly double-layer contractions with sampling, which would reduce computational overhead significantly.
2. Minimize the distance between $|\psi_{\delta t}^{\text{SU/FU}}\rangle$ and the first-order approximation $|\psi\rangle - i\delta t H|\psi\rangle$.

The second approach is expected to be both more accurate and more efficient. It avoids the evaluation of terms such as $\langle \mathbf{s} | \text{Trotter}(e^{-iH\delta t}) |\psi\rangle$, which require contraction of tensor networks with large bond dimensions. In total, this method is significantly less costly than variational sweeping as employed in Ref. [88].

It is important to assess whether the SU or FU method provides a better approximation of the time-evolved state compared to leaving the initial state unchanged. This condition is satisfied only if

$$|\langle \psi_{\delta t} | \psi_{\delta t}^{\text{SU/FU}} \rangle|^2 > |\langle \psi_{\delta t} | \psi \rangle|^2, \quad (3.44)$$

which must be verified for each problem instance. This criterion can serve as a basis for selecting the more computationally intensive FU algorithm over the simpler SU scheme during initialization.

The remaining challenge lies in determining how to optimize the fidelity between two PEPS using sampling-based techniques.

Algorithm 1 Hybrid-Sampling TDVP of a parameterized PEPS $\psi(\theta)$

```

1: function HYBRIDTDVP( $|\psi(\theta)\rangle$ ,  $H$ ,  $\delta t$ )
2:    $|\psi(\theta_{\delta t})\rangle \leftarrow \text{TEBD}_{\text{SU/FU}}(|\psi(\theta)\rangle, H, \delta t)$ 
3:    $\text{argmax}_{\theta_{\delta t}} |\langle\psi(\theta_{\delta t})| [\psi\rangle - i\delta t H |\psi\rangle]|^2$ 
4:    $\triangleright$  How the second step is performed using sampling is explained in the next
   section.
5:   return  $|\psi(\theta_{\delta t})\rangle$ 
6: end function

```

3.3.3 Optimizing the Fidelity

To optimize a parameterized state $|\psi(\theta)\rangle$ so that its fidelity with a target state $|\psi_{\oplus}\rangle$ is maximized, several approaches can be considered. This section introduces a differential equation for evolving $|\psi\rangle$ that enables such optimization using sampling, analogous to time evolution under the Schrödinger equation.

We begin by minimizing the fidelity between normalized wave functions:

$$\frac{\partial}{\partial \langle\psi|} [\langle\psi|\psi_{\oplus}\rangle \langle\psi_{\oplus}|\psi\rangle] = |\psi_{\oplus}\rangle \langle\psi_{\oplus}|\psi\rangle . \quad (3.45)$$

Gradient descent on this functional yields the differential equation:

$$|\dot{\psi}\rangle = |\psi_{\oplus}\rangle \langle\psi_{\oplus}|\psi\rangle . \quad (3.46)$$

However, this equation does not conserve the norm of $|\psi\rangle$, since $\langle\psi|\dot{\psi}\rangle \neq 0$. This issue can be addressed by adding a correction term:

$$|\dot{\psi}\rangle = |\psi_{\oplus}\rangle \langle\psi_{\oplus}|\psi\rangle - |\psi\rangle \langle\psi|\psi_{\oplus}\rangle \langle\psi_{\oplus}|\psi\rangle . \quad (3.47)$$

If the overlap between $|\psi\rangle$ and $|\psi_{\oplus}\rangle$ is small, convergence becomes slow. This can be mitigated by rescaling the right-hand side, for instance, by a factor $|\langle\psi_{\oplus}|\psi\rangle|$:

$$|\dot{\psi}\rangle = |\psi_{\oplus}\rangle \frac{\langle\psi_{\oplus}|\psi\rangle}{|\langle\psi_{\oplus}|\psi\rangle|} - |\psi\rangle |\langle\psi_{\oplus}|\psi\rangle| . \quad (3.48)$$

Alternatively, by fixing the global phase of ψ , a more natural formulation is obtained:

$$|\dot{\psi}\rangle = |\psi_{\oplus}\rangle - |\psi\rangle \langle\psi|\psi_{\oplus}\rangle . \quad (3.49)$$

This equation becomes particularly convenient when $|\psi_{\oplus}\rangle = |\psi\rangle - i\delta t H |\psi\rangle$, yielding a

rescaled version of the Schrödinger equation:

$$|\dot{\psi}\rangle = -i\delta t(H - \langle H \rangle) |\psi\rangle . \quad (3.50)$$

To convert this into a sampling-based update scheme, we reparameterize $|\psi\rangle$ as $|\psi\rangle \rightarrow \frac{|\psi(\theta)\rangle}{\sqrt{\langle\psi(\theta)|\psi(\theta)\rangle}}$ and define the sampling distribution $P(\mathbf{s}) = \frac{|\langle\mathbf{s}|\psi(\theta)\rangle|^2}{\langle\psi(\theta)|\psi(\theta)\rangle}$. The relevant quantities are then:

$$F_{\mathbf{s}}^{\text{loc}} := \frac{\langle \mathbf{s} | \psi_{\oplus} \rangle}{\langle \mathbf{s} | \psi(\theta) \rangle} \sqrt{\frac{\langle \psi(\theta) | \psi(\theta) \rangle}{\langle \psi_{\oplus} | \psi_{\oplus} \rangle}} , \quad (3.51)$$

$$\langle F_{\mathbf{s}}^{\text{loc}} \rangle_{\mathbf{s} \in P(\mathbf{s})} = \frac{\langle \psi(\theta) | \psi_{\oplus} \rangle}{\sqrt{\langle \psi(\theta) | \psi(\theta) \rangle \langle \psi_{\oplus} | \psi_{\oplus} \rangle}} , \quad (3.52)$$

$$\dot{\theta} = G^{-1} O^{\dagger} (F_{\mathbf{s}}^{\text{loc}} - \langle F_{\mathbf{s}}^{\text{loc}} \rangle) . \quad (3.53)$$

Note that $\langle F_{\mathbf{s}}^{\text{loc}} \rangle_{\mathbf{s} \in P(\mathbf{s})}^2$ gives the fidelity between the two states.

Finally, the ratio of the norms can be evaluated using the same set of samples:

$$\frac{\langle \psi_{\oplus} | \psi_{\oplus} \rangle}{\langle \psi | \psi \rangle} = \left\langle \frac{|\langle \mathbf{s} | \psi_{\oplus} \rangle|^2}{|\langle \mathbf{s} | \psi \rangle|^2} \right\rangle_{\mathbf{s} \in P(\mathbf{s})} . \quad (3.54)$$

It is important to note that sampling from $P(\mathbf{s})$ provides low-variance estimators only when $|\psi\rangle$ and $|\psi_{\oplus}\rangle$ are close. If the states differ significantly, it is advisable to either employ an alternative probability distribution or adopt a different evolution equation.

CHAPTER 4

State Preparation

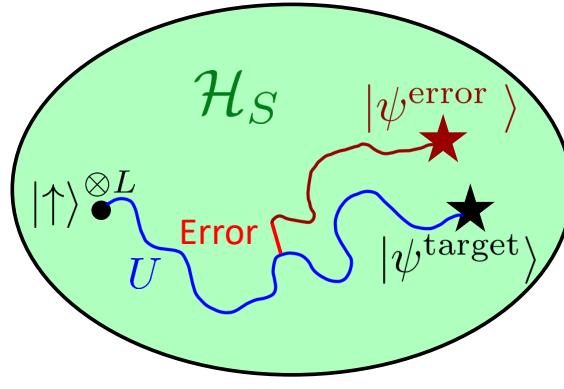
Quantum state preparation is a central task in quantum information science. It is essential for quantum computing, quantum simulation, and the initialization of many-body quantum systems [5].

State preparation plays a vital role in measurement-based quantum computing (MBQC), where computation proceeds through local measurements on an entangled resource state [3]. It is also fundamental for quantum error correction, which requires precise preparation of topological or stabilizer code states [4].

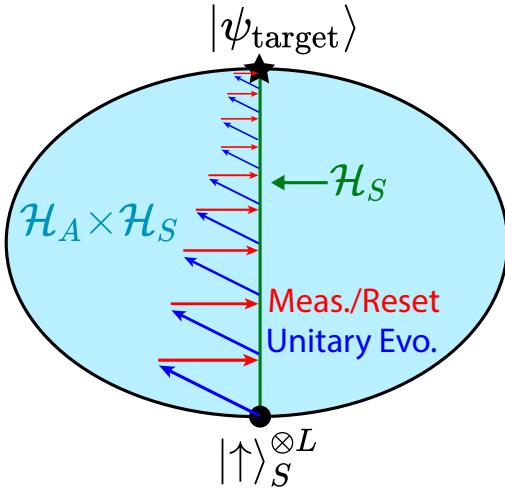
This chapter analyzes three main classes of quantum state preparation protocols. These are unitary circuits, measurement-based protocols without feedback, and feedback-assisted protocols, as illustrated in Fig. 4.1. While the analysis in this chapter focuses on one-dimensional systems as studied in Publications [I,II], many relevant states reside in two spatial dimensions. Extending these methods to higher-dimensional systems remains a central challenge. In this context, the techniques developed in Publication [III] offer a framework to assess, in future work, the classical simulability of measurement-based state preparation protocols in two-dimensional systems.

A fundamental question in quantum many-body physics is the following: given a quantum platform, which quantum states of matter can it prepare? From a resource theoretic perspective, answering this question yields insight into both operational capabilities and the intrinsic structure of the target state [5].

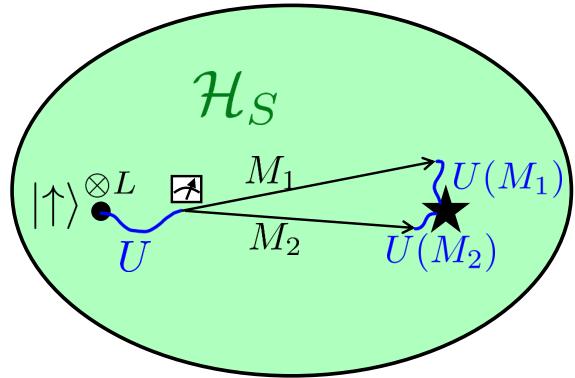
This chapter classifies state preparation methods according to the role of measurements and feedback. The first class comprises unitary circuits. These prepare the target state through deterministic sequences of gates. Their depth generally scales with system size and, as such, is sensitive to noise [7, 8]. The second class involves measurement and reset operations without feedback. These protocols exploit dissipation or ancilla measurements to drive the system into the desired state [19–21]. The third class introduces



(a)



(b)



(c)

Figure 4.1.: Illustration of the three approaches to quantum state preparation explored in this chapter.

(a) In unitary-only protocols, circuit depth typically scales poorly with system size, leading to a high likelihood of errors during preparation. (b) Measurement-assisted state preparation, as developed in Publication [I], guides the system toward the target state regardless of its location in Hilbert space. The initial state $|\uparrow\rangle_S^{\otimes L}$ lies entirely within the system Hilbert space \mathcal{H}_S (green line). Unitary evolution under the steering Hamiltonian for a duration δt (blue arrows) entangles the system with ancillas, resulting in a state within the combined space $\mathcal{H}_S \otimes \mathcal{H}_A$ (light blue surface). Measurement or reset of the ancillas (red arrows) projects the system back into \mathcal{H}_S , moving it closer to the target state. Repeated application leads to convergence toward $|\psi^{\text{target}}\rangle$. (c) In feedback-based protocols, measurement outcomes are not discarded but instead used to inform and adapt subsequent unitaries. The unitary operation is conditioned on the result M , as autonomously learned in Publication [II].

feedback. Here, measurement outcomes determine subsequent operations. This adaptivity can significantly accelerate convergence and reduce resource requirements [22–25, 27, 32, 33]. Note that these three classes are depicted in Fig. 4.1.

These approaches are governed by the constraints of quantum mechanics and must conform to quantum channels described by completely positive trace-preserving (CPTP) maps. The evolution of a density matrix $\rho_1 \rightarrow \rho_2$ under such a channel is expressed as:

$$\rho_2 = \sum_i K^i \rho_1 K^{i\dagger}, \quad \sum_i K^{i\dagger} K^i = \mathbb{1}, \quad (4.1)$$

where $\{K^i\}$ are Kraus operators that capture both unitary and non-unitary transformations. They are trace-preserving:

$$\text{Tr}(\rho_2) = \sum_i \text{Tr}(K^i \rho_1 K^{i\dagger}) = \text{Tr}\left(\rho_1 \sum_i K^{i\dagger} K^i\right) = \text{Tr}(\rho_1), \quad (4.2)$$

and preserve positive definiteness:

$$\langle \psi | \rho_2 | \psi \rangle = \sum_i \langle \psi | K^i \rho_1 K^{i\dagger} | \psi \rangle = \sum_i \underbrace{\langle \psi_i | \rho | \psi_i \rangle}_{\geq 0} \geq 0. \quad (4.3)$$

This chapter introduces three broad classes of state preparation protocols. The first class comprises unitary-only circuits, which implement state preparation using deterministic sequences of gates. These protocols realize special cases of CPTP maps where the channel consists of a single unitary Kraus operator $K^1 = U$.

The second class is explored in Publication [I] and includes measurement-based protocols without feedback. These methods leverage the intrinsic non-unitary nature of quantum measurements to drive the system toward a target state. In such protocols, the state is entangled with an ancilla and then projected through measurement, effectively realizing non-unitary operations of the form:

$$K^M = \langle M |_A U_{SA} | 0 \rangle_A, \quad (4.4)$$

where U_{SA} is a joint system-ancilla¹ unitary, and $|M\rangle_A$ is the measurement outcome on the ancilla. As the outcome is not used to condition future operations, the system evolves under an ensemble of K^M drawn randomly by quantum mechanics with a probability of:

$$p_i = \text{Tr}(K^i \rho K^{i\dagger}), \quad (4.5)$$

¹The subscript A/S indicates that the operator acts on the ancilla's/system's Hilbert space.

forming a positive operator-valued measure (POVM). An important example of this approach is presented in Publication [I], where a steering Hamiltonian is constructed to implement a measurement-induced protocol for preparing the AKLT state.

The third class extends this framework by incorporating feedback. These protocols use classical information from measurement outcomes to condition subsequent operations, effectively implementing adaptive control. Feedback enables the design of CPTP maps where each Kraus operator depends on a sequence of operations, such as

$$K^M = U_S^{(2)}(M) \langle M|_A U_{SA}^{(1)} |0\rangle_A \quad (4.6)$$

where $U_S^{(2)}(M)$ denotes a system unitary conditioned on the measurement result M . Feedback-based protocols can substantially accelerate convergence, increase fidelity, and reduce the required circuit depth. In Publication [II], such adaptive strategies are learned automatically within variational quantum circuits, enabling preparation protocols beyond the reach of unitary-only or measurement-only approaches.

The remainder of this chapter provides a detailed analysis of each approach. The discussion emphasizes both theoretical considerations and practical implications for near-term quantum devices.

4.1 Unitary-only

Purely unitary approaches for state preparation are governed by well-established theoretical bounds and practical constraints. Sequential circuit methods can exactly prepare all MPS using strictly local gates, though this typically requires linear circuit depth scaling with the system size ($T \propto N$) [90]; this protocol is explained in detail at the end of this section. Alternatively, constant-depth circuits that use quasilocal gates, with support growing (poly-)logarithmically in the system size, can approximate injective MPS. However, converting these quasilocal gates into strictly local gates generally leads to an exponential increase in circuit depth [91, 92].

Another viable strategy is adiabatic preparation, where a quantum system is slowly evolved from a product state to the desired target state, leveraging guaranteed spectral gaps typically present in topologically trivial phases (see Sec. 2.1.4). This approach allows for preparation in a circuit depth scaling polynomially with $\log(N/\epsilon)$ [93–95], where ϵ denotes the maximum allowable error in the fidelity between the prepared state and the target state. Specifically, for translationally invariant injective MPS, optimal strictly local gate algorithms have achieved circuit depths of $O(\log(N/\epsilon))$ [32].

Note that variational methods can also be used to find quantum circuits for the preparation of target states of interest. This is further discussed in Chap. 5.

4.1.1 Simple Unitary Preparation of an MPS State

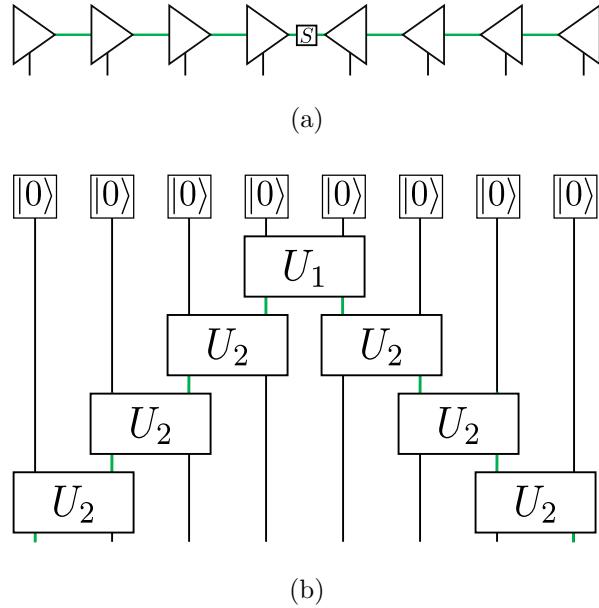


Figure 4.2.: Preparation of an MPS with bond dimension equal to the physical dimension employing only unitary operations. Green lines represent the MPS bond dimension, and black lines represent the physical dimension.

This section presents the preparation of an orthogonalizable MPS using a circuit with depth scaling linearly with the system size. The procedure begins by transforming the MPS into its central orthogonal form, as illustrated in Fig. 4.2a. When the physical and bond dimensions are identical, the MPS can be decomposed into gates satisfying the following relations:

$$\begin{array}{c} 0 \quad 0 \\ \hline U_1 \end{array} = \begin{array}{c} \text{---} \quad \text{---} \\ \boxed{S} \end{array} , \quad (4.7)$$

$$\begin{array}{c} 0 \\ \hline U_2 \end{array} = \begin{array}{c} \text{---} \quad \text{---} \\ \text{triangle} \end{array} . \quad (4.8)$$

In these diagrams, green lines represent the bond indices of the tensor, while black lines correspond to the physical indices. The labels 0 indicate that only the zeroth component of the respective physical indices is selected, that is, the component $U_{0,0,i,j}$ of the tensor is used. The unitary U_1 is chosen such that it initializes the entanglement between the left and right parts of the system through the Schmidt spectrum S . There exist many unitaries fulfilling Eq. (4.7); the simplest to implement should be selected. For U_2 , since

the right-hand side of Eq. (4.8) defines an isometry, a corresponding unitary U_2 satisfying the equation can always be constructed.

If the bond dimension χ exceeds the physical dimension d , the unitaries U_1 and U_2 must act on a larger number of qudits. In this case, the number of required qudits scales as $\log_d(\chi)$.

4.2 Non-unitary with Only Measurements

Quantum state preparation based on unitary operations combined with measurements, without adaptive feedback, has emerged as a versatile framework. This approach was employed for state preparation in Publication [I]. The present section reviews foundational theoretical proposals and recent experimental realizations of passive state preparation. The focus lies on methods that incorporate measurements, sequential ancilla-based protocols, and engineered dissipation mechanisms.

One foundational experimental technique is heralded entanglement generation, where measurements on ancillary systems project primary qubits into entangled states. Originally proposed for matter qubits coupled via photonic links, this approach enables probabilistic yet high-fidelity creation of Bell pairs and GHZ states [96–102]. A successful measurement outcome heralds entanglement without requiring corrective operations.

Complementary to measurement-based protocols, engineered dissipation provides a route to quantum state preparation by constructing Lindblad operators that have the target state as a steady state. Seminal theoretical frameworks for dissipative state engineering were developed in Refs. [12–14]. The dissipative dynamics autonomously remove energy and correct deviations from target states without active feedback.

The dissipative preparation of SPT phases, particularly the AKLT state, has recently attracted significant attention. Symmetry-respecting Lindblad processes have been employed to stabilize the AKLT ground state while preserving the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry [19–21]. The frustration-free nature of the AKLT Hamiltonian, where each local term is minimized independently, facilitates efficient local stabilization [13]. This approach is used in Publication [I] and is further discussed in the next section.

Another avenue is simulated cooling, where the system is coupled to an actively cooled reservoir [15–18]. These protocols offer the advantage of being applicable even when the ground state is unknown. However, they are more challenging to implement in practice, as engineering suitable system-reservoir couplings that do not disturb the system’s energy spectrum is difficult.

Fully passive dissipative encoding can prepare topologically ordered states, such as the toric code. By engineering local mapping operators that detect and annihilate anyonic

excitations, systems can be driven into topologically ordered steady states from arbitrary initial conditions [103, 104]. Such encoders are promising for robust preparation of quantum memories, as they naturally correct local errors during the encoding procedure.

In summary, a wide range of quantum states, from simple entangled pairs to topologically ordered codes and SPT phases, can be prepared using combinations of unitary operations, fixed-basis measurements, and engineered dissipation without relying on adaptive feedback. These passive strategies are particularly suited to near-term quantum devices, where coherence times and control fidelity remain limited. In the next section, one of these approaches is showcased.

4.2.1 Designing Lindbladians for Passive State Preparation

This section provides an intuition on how to design a Lindbladian that steers a system into a target state, and how it can be rewritten as a Hamiltonian evolution with periodic measurements. In Publication [I], this protocol is investigated and improved for the AKLT case.

One example illustrating the use of measurements in state preparation is derived from imaginary time evolution. If a random initial state $|\psi_0\rangle$ is evolved under the imaginary time evolution e^{-tH} , the ground state will be reached, for sufficiently large times. Unfortunately, this operation is not trace-preserving, as can be seen by expressing the evolution in terms of the density matrix:

$$\dot{\rho} = -\{H, \rho\} \quad (4.9)$$

$$\text{Tr}(\dot{\rho}) = -2 \text{Tr}(H\rho). \quad (4.10)$$

To preserve normalization while maintaining the structure of the evolution, an additional term can be introduced that subtracts the energy expectation value:

$$\dot{\rho} = -\{H, \rho\} - 2 \text{Tr}(H\rho)\rho. \quad (4.11)$$

This modification results in a nonlinear differential equation, as the evolution now depends explicitly on the state ρ . Implementing this process on a quantum device requires knowledge of the evolving quantum state. This idea forms the basis of quantum imaginary time evolution (QITE) [105], where measurements are used to characterize² ρ to find the best unitary operations to implement the equation above.

To maintain linearity of the differential equation, it is necessary to abandon the require-

²The number of measurements required to characterize the quantum state depends significantly on its entanglement. States with higher entanglement typically require exponentially more measurements to achieve accurate characterization.

ment of exactly following imaginary time dynamics and instead ensure that the target state is a steady state of the evolution. This can be achieved by introducing a dissipative term:

$$\dot{\rho} = -\{H, \rho\} + \sum_i M_i \rho M_i^\dagger \quad (4.12)$$

$$\text{Tr}(\dot{\rho}) = 0 \quad \forall \rho \quad (4.13)$$

$$\Rightarrow H = \frac{1}{2} \sum_i M_i^\dagger M_i \quad (4.14)$$

This Lindblad form is particularly useful because Ref. [21] provides a method to implement this evolution using unitaries and measurements.

To ensure convergence to the ground state, two conditions must be satisfied. The energy must decrease during the evolution, such that $\text{Tr}(H\dot{\rho}) < 0$ for all states other than the ground state $\rho_\oplus = |\psi_\oplus\rangle\langle\psi_\oplus|$. Second, the ground state must be a dark state of the evolution, satisfying $\dot{\rho}_\oplus = 0$.

The second condition can be met by shifting the Hamiltonian spectrum such that $H|\psi_\oplus\rangle = 0$. In this case, the evolution of the target state simplifies to:

$$\dot{\rho}_\oplus = \sum_i M_i \rho_\oplus M_i^\dagger. \quad (4.15)$$

For frustration-free Hamiltonians composed of local terms, $H = \sum_l O_l$, which can be minimized independently, it is straightforward to construct local operators M_l such that $M_l |\psi_\oplus\rangle = 0$. This approach was first introduced by Verstraete *et al.* [14], who also observed that any injective MPS possesses a frustration-free parent Hamiltonian, as discussed in Sec. 2.1.3. Consequently, injective MPS are amenable to preparation via this method. The protocol was subsequently analyzed in detail by Zhou *et al.* [19] in the context of the AKLT state, showing that it can indeed be prepared using this dissipative approach.

In practical implementations, the resource requirements for this measurement-based preparation scheme are determined by the properties of the underlying MPS. Specifically, the number of ancilla qudits and the interaction range of the steering Hamiltonian both scale with the injective length and bond dimension.

To construct the operators M_i , one proceeds by computing the eigenvalues and eigenvectors of each local O_l , and grouping them into excited $|\Theta_{l,k'}^{\text{undesired}}\rangle$ and non-excited states $|\Theta_{l,k}^{\text{desired}}\rangle$ depending on their eigenvalues. Notably, $\langle\Theta_{l,k}^{\text{undesired}}|\psi_\oplus\rangle = 0$ if the Hamiltonian is indeed frustration-free. The operators M_l are then designed to locally map the entire

set of undesired states into the desired subspace:

$$M_{l,\alpha} = \sum_{k,k'} C_{k,k'}^{\alpha} |\Theta_{l,k}^{\text{desired}}\rangle\langle\Theta_{l,k'}^{\text{undesired}}| . \quad (4.16)$$

Several degrees of freedom are available in this construction, which can be exploited to accelerate the steering. Specifically, in Sec. 5 of Publication [I] we find that if the mapping operators are chosen such that they commute as much as possible, this greatly speeds up convergence.

The other condition requires that the Lindbladian evolution decreases the energy expectation value for any ρ different from ρ_{\oplus} , that is, $\text{Tr}(H\rho) < 0$:

$$\sum_i \text{Tr}\left(M_i^{\dagger} H M_i \rho\right) < \text{Tr}(H^2 \rho) \quad \forall \rho, \quad (4.17)$$

$$\sum_{i,j} \text{Tr}\left(M_i^{\dagger} M_j^{\dagger} M_j M_i \rho\right) < \sum_{i,j} \text{Tr}\left(M_i^{\dagger} M_i M_j^{\dagger} M_j \rho\right) \quad \forall \rho, \quad (4.18)$$

$$\sum_{i,j} \text{Tr}\left(\left[M_i^{\dagger}, M_j^{\dagger}\right] M_j M_i \rho\right) + \text{Tr}\left(M_j^{\dagger} \left[M_i^{\dagger}, M_j\right] M_i \rho\right) < 0. \quad (4.19)$$

If the operators M_i are local and act on distinct sites, all commutator terms vanish. When two or more operators act on the same site, the inequality must be verified explicitly for each site. This verification is tractable, as it reduces to analyzing the corresponding local Hilbert space. Given the flexibility in choosing M_i , one can optimize the inequality to enhance convergence. This condition, however, only ensures rapid convergence in the dissipative limit. In contrast, when the dynamics are implemented via unitary evolution with periodic measurements, the degree of commutativity among the M_i becomes an additional relevant factor.

Ref. [21] further shows that by incorporating ancillas, the Lindbladian in Eq. (4.12) can be rewritten as a real-time evolution with periodic measurements on the ancillas. The joint system-ancilla evolution is governed by a Hamiltonian of the form:

$$H_{SA} = \sum_{l,\alpha} M_{l,\alpha}^S \otimes D_{l,\alpha}^{A\dagger} + \text{h.c.} \quad (4.20)$$

Here, the $D_{l,\alpha}^{A\dagger}$ are ancilla operators that map the ancilla into an orthogonal subspace relative to its initial state $|\Phi^A\rangle$.

After a time interval δt , the ancillas are reset to their initial state $|\Phi^A\rangle$, effectively realizing a quantum reset operation. The period δt is crucial in determining the convergence speed of the protocol. In the weak-coupling, or Markovian, limit where $\delta t \ll 1$, the

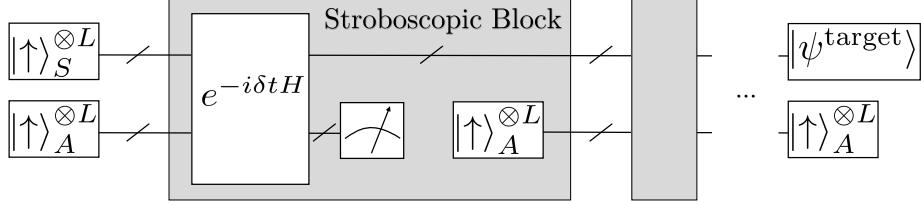


Figure 4.3.: Illustration of the steering protocol used in Publication [I]. The system is initialized in the product state $|\uparrow\rangle_S^{\otimes L}$. Unitary evolution under the steering Hamiltonian given in Eq. (4.20) for a duration δt entangles the system with the ancillas. Subsequent measurement or reset of the ancillas projects the system closer to the target state. Repeated application of this procedure leads to convergence toward the desired state $|\psi_t\rangle$.

system dynamics are described by an effective Lindblad master equation:

$$\dot{\rho} = \delta t \sum_{l,\alpha} \left(M_{l,\alpha}^S \rho M_{l,\alpha}^{S\dagger} - \frac{1}{2} \{ M_{l,\alpha}^{S\dagger} M_{l,\alpha}^S, \rho \} \right). \quad (4.21)$$

This limit simplifies the dynamics but does not lead to the fastest convergence. In the limit $\delta t \rightarrow 0$, the dynamics are effectively frozen, similar to the quantum Zeno effect. Increasing δt is therefore beneficial, up to an optimal point, as further discussed in [I]. Stronger entanglement between the system and the ancillas enhances convergence, although beyond a certain point, the entanglement becomes detrimental.

Furthermore, Publication [I] discusses strategies for designing the operators $M_{l,\alpha}^S$ to optimize the convergence rate.

4.3 Non-unitary with Measurement and Feedback

Adaptive measurement-feedback methods leverage past measurement results to guide future actions, significantly improving convergence speed. A prominent example is quantum error correction, where adaptive feedback enables real-time correction based on measurement outcomes, stabilizing quantum states against decoherence. These techniques form the foundation of fault-tolerant quantum computing, and are closely related to the measurement-based protocols discussed here.

Beyond error correction, adaptive feedback methods have been effectively applied to prepare diverse quantum states, including fracton states, Schrödinger's cat states, and certain non-Abelian topological orders [22–26]. Some of these states can even be created with a single measurement round when specific algebraic criteria are satisfied [26].

The effectiveness of adaptive measurement strategies is further demonstrated by deterministic constant-depth schemes that use measurements to fuse smaller resource states.

Such schemes have achieved high fidelities experimentally, as exemplified by the preparation of the AKLT state [27]. However, theoretical analyses impose rigorous constraints on the applicability of fusion-based methods, closely linked to the entanglement structure of the target state. Specifically, preparing MPS via fusion measurements typically requires them to be symmetry-protected topological (SPT) phases with a flat entanglement spectrum [28–30].

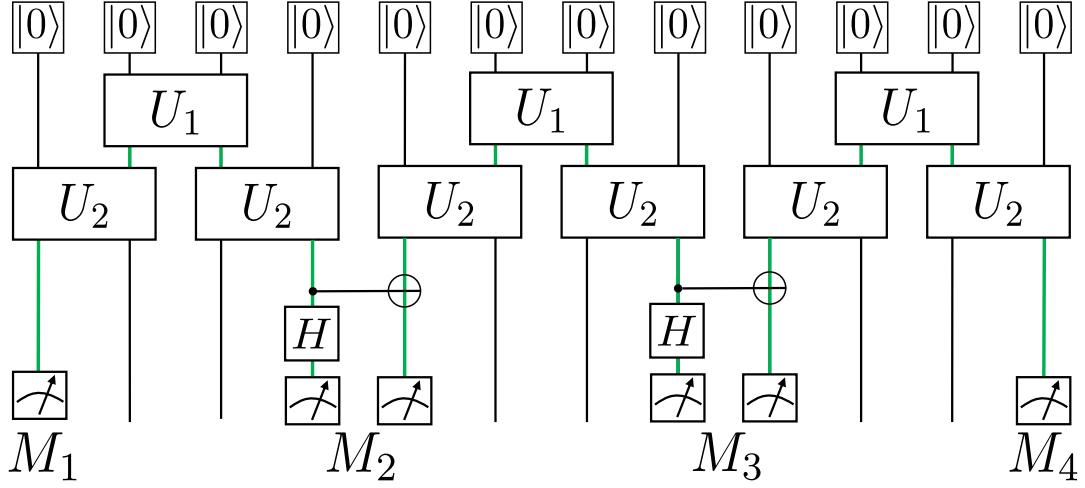
More broadly, adaptive quantum circuits enable constant-depth preparation of various MPS classes, including symmetry-protected, symmetry-broken, and states with finite Abelian, non-Abelian, or continuous symmetries [31]. For translationally invariant injective MPS, adaptive measurement-based protocols can even reduce circuit depth scaling from logarithmic to double logarithmic, with complexity $O(\log \log(N/\epsilon))$, if exact preparation is relaxed [32]. Similarly, probabilistic or approximate protocols can achieve constant depth independent of system size, exemplified by Dicke and W-states [33].

A practical consideration for implementing feedback-based protocols is the overhead introduced by measurement and classical feedback. While the discussion above assumes idealized conditions with instantaneous feedback, current quantum hardware still faces significant latency challenges. Active research is addressing these limitations, particularly in the context of real-time quantum error correction [106–108]. For example, in a recent experiment on real-time error correction, Google reported a QEC cycle time of $1.1 \mu\text{s}$ and an average feedback latency of $63 \mu\text{s}$, which corresponds to a backlog of approximately 57 cycles [106]. These constraints currently limit the applicability of feedback-based preparation schemes to small system sizes or require buffering and delayed correction strategies. However, the rapid pace of experimental progress suggests that these latencies will decrease in the near future, enabling the implementation of real-time feedback protocols.

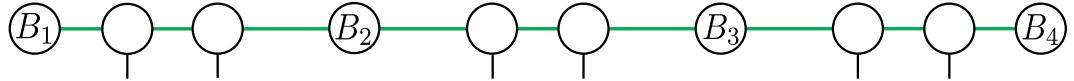
4.3.1 Preparing SPT States Using Feedback and Measurements

The first example to consider in state preparation with measurements and feedback is the GHZ state. In this protocol, a Hadamard gate is first applied to each qubit, placing the system into a uniform superposition $|+\rangle^{\otimes N}$. Subsequently, the $Z_i Z_{i+1}$ operators are measured between neighboring qubits. These measurements collapse the system into a superposition of the form $\frac{1}{\sqrt{2}}[|s\rangle + |\bar{s}\rangle]$, where s is a random bit string and \bar{s} denotes its bitwise complement. Based on the measurement outcomes, corrective σ_x operations are applied to specific qubits such that $|0 \dots 0\rangle = \bigotimes_{s_i \neq 0} U_i^x |s_1, \dots, s_L\rangle$.

The second example is the protocol by Smith *et al.* [27] for preparing G-injective MPS. In Publication [II], this scheme is extended using a learning-based approach that incorporates measurements and feedback within variational quantum circuits. As a first application, the learning protocol is applied to the preparation of the AKLT state, for



(a)



(b)

Figure 4.4.: Preparation of a G-injective MPS with G-injective length of two using fusion measurements like first described by Smith *et al.* [27]. (a) Small versions of the state are prepared using U_1 and U_2 . Then, the ancillas are measured in the Bell basis, yielding one of four possible outcomes. One outcome properly merges the two MPS, whereas the others create an error quantified by $B_i = B(M_i)$. If the state is G-injective, and the defects B are elements of G then they can be moved to the edges of the lattice using Eq. (4.22).

which an analytically derived scheme by Smith *et al.* exists.

The key property enabling the Smith *et al.* protocol is that, for G-injective MPS, an operator in the group G acting on the bond dimension can be transferred from one side of the tensor to the other:

$$\begin{array}{c} \text{---} \\ | \quad | \quad | \\ \text{---} \end{array} \quad B_g \quad = \quad \begin{array}{c} \text{---} \\ | \quad | \quad | \\ \text{---} \end{array} \quad . \quad (4.22)$$

SPT states are examples of G-injective states, as discussed in Sec. 2.1.4, and are therefore amenable to preparation using this scheme.

The algorithm proceeds by first preparing small segments of the target state, leaving the MPS bond dimension (green) open (see Fig. 4.4a). This initial preparation is carried out using the method described in Sec. 4.1.1. A Bell measurement is then performed

on the ancillas associated with the bond dimension. This measurement merges the two small MPS into a larger state, as illustrated in Fig. 4.4b. The resulting state reproduces the target MPS structure up to local defects, denoted by B_i . These defects can be systematically propagated to one end of the chain by applying local operators derived from the G-injectivity property in Eq. (4.22). This procedure can be performed in parallel for all sites.

Using this approach, it is possible to prepare G-injective states, and consequently SPT states, such as the AKLT state. Although the prepared state exhibits the correct bulk properties, it may contain random edge states. This arises because, after all correction operations are applied, residual errors accumulate at the boundary of the MPS.

The learning algorithm developed in Publication [II] is capable of autonomously discovering this preparation protocol. Furthermore, it can improve upon it by deterministically preparing a specific edge mode of the AKLT state, a task that cannot be achieved with the original analytic method.

4.4 Learning Non-unitary with Feedback

Autonomous learning of measurement and feedback protocols for quantum state preparation has been explored in the literature through two primary approaches.

In the first approach, reinforcement learning (RL) is used to dynamically determine both the optimal timing of measurements and the corresponding unitary operations conditioned on measurement outcomes. This method has demonstrated effectiveness in simple settings, such as single-particle systems [36, 37] and two-particle systems [38, 39]. Furthermore, deep RL techniques have enabled the development of feedback strategies in cavity QED platforms [109], and systems with nonlinear potentials under continuous weak measurements [37]. Despite these advances, applying RL-based protocols to multi-qubit systems remains challenging due to the increasing complexity and poor scalability of such methods. In response to these limitations, a recent work [110] has investigated transformer-based neural network architectures, which have shown promise in learning robust feedback strategies from sequences of measurements, even under noisy conditions.

In the second approach, greedy optimization techniques have been proposed for the preparation of multi-qubit states [34, 35]. These protocols optimize unitary operations incrementally, immediately after each measurement, to maximize fidelity at each step. Ancilla systems are coupled to the target system via unitary interactions and then measured projectively. The corresponding unitary gates are optimized independently at

each step:

$$|\psi'_i(\theta_i, \dots)\rangle = U(\theta_i) |\psi_{i-1}(\theta_{i-1}, \dots)\rangle, \quad (4.23)$$

$$\theta_i = \min_{\theta_i} \text{loss}(|\psi'_i(\theta_i, \dots)\rangle),$$

$$P(M) = |\langle M|_A |\psi'_i(\theta_i, \dots)\rangle|^2, \quad (4.24)$$

$$|\psi_i(\theta_i, \dots)\rangle = \frac{1}{\sqrt{P(M)}} \langle M|_A |\psi'_i(\theta_i, \dots)\rangle, \quad M \in P(M).$$

This strategy exhibits two main limitations. First, experimental realization requires real-time simulation of the system to optimize unitaries based on measurement outcomes, which imposes considerable computational overhead. Second, the greedy optimization criterion only accounts for immediate outcomes and does not take into account future possible operations. As a result, these methods do not scale efficiently beyond a few qubits and fail to implement more sophisticated feedback protocols, such as the deterministic strategy proposed by Smith *et al.* [27], which would involve the joint optimization of unitaries for all possible measurement outcomes, as demonstrated in Publication [II].

Recent work [111, 112], including the approach described in Publication [II], addresses these challenges by employing variational quantum circuits (VQCs) augmented with projective measurements and conditional feedback. In this framework, the entire feedback strategy is optimized simultaneously, enabling the identification of optimal completely positive trace-preserving (CPTP) maps and corresponding feedback operations:

$$\rho_1(\theta_1) = U_1(\theta_1) \rho_0 U_1^\dagger(\theta_1), \quad (4.25)$$

$$\rho_1^M(\theta_1) = |0\rangle_A \langle M| \rho_1(\theta_1) |M\rangle_A \langle 0|, \quad (4.26)$$

$$\rho_2(\theta_1, W) = \sum_M U_2(\theta_2) \rho_1^M(\theta_1) U_2^\dagger(\theta_2), \quad (4.27)$$

$$\text{with } \theta_2 = f(M; W).$$

A key distinction between the greedy and our non-greedy approach lies in the optimization strategy. In the greedy protocol, the fidelity of $\rho_1(\theta_1)$ with respect to the target state is optimized first. Only after a measurement is performed is the subsequent fidelity of $\rho_2(\theta_2)$ optimized, conditioned on the measurement outcome. In contrast, our method jointly optimizes both θ_1 and the feedback parameters W by directly maximizing the fidelity of the final state $\rho_2(\theta_1, W)$ ³. This simultaneous optimization enables the protocol to select measurements that ensure all possible outcomes yield states that can be efficiently mapped to the target. As a result, the learned strategy is significantly more

³As a consequence of the joint optimization, the intermediate state $\rho_1(\theta_1)$ may have negligible overlap with the target state. This contrasts with greedy approaches, which explicitly maximize the fidelity of intermediate states at each step.

flexible and capable of implementing complex state preparation protocols beyond the scope of greedy methods.

CHAPTER 5

Variational Quantum Circuits

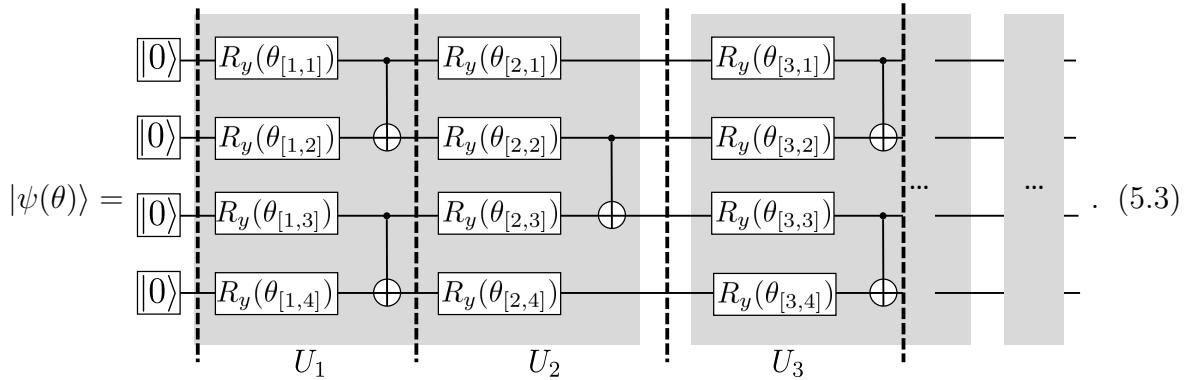
Variational Quantum Circuits (VQCs) are hybrid quantum-classical algorithms where a parameterized quantum circuit $U(\theta)$ acts on an initial state $|0\rangle^{\otimes N}$ to produce a variational quantum state:

$$|\psi(\theta)\rangle = U(\theta)|0\rangle^{\otimes N}. \quad (5.1)$$

The parameters θ are optimized to minimize a classical loss function $C(\theta)$, often defined as the expectation value of an observable O :

$$C(\theta) = \langle\psi(\theta)|O|\psi(\theta)\rangle. \quad (5.2)$$

A common choice for $U(\theta)$ is the hardware-efficient ansatz, which is composed of alternating layers of parameterized single-qubit rotations and entangling gates. Its structure can be written as:



VQCs serve as the foundation for various near-term quantum algorithms, including quantum state preparation, Variational Quantum Eigensolvers (VQE), Quantum Approximation, and more.

mate Optimization Algorithms (QAOA), and quantum machine learning tasks such as classification and generative modeling.

Originally, the hybrid structure of VQCs was considered ideal for Noisy Intermediate-Scale Quantum (NISQ) devices. These circuits are typically shallow, minimizing the effects of decoherence and hardware noise. The classical optimizer performs most of the computational work, reducing quantum resource demands. However, this apparent advantage does not translate to scalability, which remains limited by fundamental optimization challenges.

Two central obstacles are the presence of barren plateaus, which are regions of exponentially vanishing gradients, and the existence of complex optimization landscapes with numerous local minima.

In Publication [II], the expressivity of the variational ansatz is extended by incorporating measurement and feedback operations aiming to learn a protocol of the type discussed in Sec. 4.3. The process proceeds as follows:

$$\rho_1(\theta_1) = U_1(\theta_1)\rho_0U_1^\dagger(\theta_1), \quad (5.4)$$

$$\rho_1^M(\theta_1) = |0\rangle_A\langle M|\rho_1(\theta_1)|M\rangle_A\langle 0|, \quad (5.5)$$

$$\rho_2(\theta_1, W) = \sum_M U_2(\theta_2)\rho_1^M(\theta_1)U_2^\dagger(\theta_2), \quad (5.6)$$

$$\text{with } \theta_2 = f(M; W).$$

Initially, a parameterized unitary $U_1(\theta_1)$ is applied to the input state ρ_0 . A projective measurement is then performed, yielding an outcome M . The resulting post-measurement state $\rho_1^M(\theta_1)$ depends on the observed measurement outcome. A second unitary $U_2(\theta_2)$ is subsequently applied, where the parameters θ_2 are determined by a feedback function $f(M; W)$. This function maps the measurement outcome M to a parameter vector θ_2 based on a set of learnable parameters W .

The utility of this measurement and feedback ansatz is discussed in more detail in Sec. 4.3. Notably, as demonstrated in Publication [II], this circuit structure introduces a new type of local minima that is specific to feedback-based variational circuits. These minima can be mitigated by appropriately tuning the learning rates and by regularizing the ancillary systems used in the protocol.

5.1 Practical Limitations

5.1.1 Barren Plateaus

A barren plateau refers to a region in the parameter space where the gradients of the loss function decay exponentially with the number of qubits [113]. As a result, estimating gradients from a finite number of measurements becomes infeasible. Formally, the variance of gradient estimators scales as [113, 114]:

$$|\partial C| \sim e^{-\alpha N}, \quad (5.7)$$

where N denotes the number of qubits. To maintain a constant signal-to-noise ratio, the required number of measurements grows exponentially with N . Intuitively, barren plateaus emerge due to the concentration of measure in high-dimensional Hilbert spaces. Random circuits tend to explore the space uniformly, leading the output state to become nearly orthogonal to directions along which parameters can be improved.

The occurrence of barren plateaus significantly restricts scalability, even if the employed ansatz has sufficient expressivity. The locality of the loss function strongly influences gradient behavior. Global loss functions, such as fidelity with respect to an N -qubit target state, commonly lead to barren plateaus [113]. Conversely, local loss functions, defined via subsystem overlaps, can mitigate exponential gradient decay [114] for shallow circuits.

For sufficiently expressive ansätze approximating unitary 2-designs, gradients concentrate near zero due to concentration-of-measure effects [113, 115]. Even hardware-efficient ansätze, designed explicitly for NISQ hardware, can manifest barren plateaus when the circuit depth increases significantly¹.

Several mitigation strategies exist, including the use of local or incremental loss functions [114], layer-wise or problem-informed ansatz training [116], and symmetry-preserving initialization methods [117]. These methods help but do not solve the inherent problems of optimizing in the exponentially large Hilbert space.

Vanishing gradients are less problematic when computed using classical methods. In such cases, gradients can be rescaled by a suitable factor without additional computational cost. This is possible because classical gradient computation does not rely on sampling. As a result, classical methods avoid the exponential cost associated with barren plateaus in quantum settings, where gradients are estimated through repeated measurements.

¹In hardware-efficient ansätze with local cost functions, the average gradient magnitude decays exponentially with circuit depth until it reaches a minimum value. This minimum gradient magnitude scales exponentially with the system size [113].

The primary challenge in classical approaches lies in the representability of the wave function. As the circuit depth increases, so does the entanglement, making the wave function increasingly difficult to represent and optimize. The computational cost of classical simulations scales exponentially with circuit depth due to the growing complexity of the Hilbert space.

In summary, both classical and quantum gradient computations face exponential overheads stemming from the size of the Hilbert space. However, the source of the overhead differs: classical methods are limited by representational complexity, whereas quantum methods are constrained by sampling noise and the presence of barren plateaus.

5.1.2 Local Minima

In addition to barren plateaus, optimization may also fail due to the presence of numerous local minima in the loss landscape. Shallow circuits are particularly susceptible to this issue, as they are often underparameterized and lack sufficient variational freedom to represent the target states accurately. This limitation results in a rugged optimization landscape populated with suboptimal local minima [8, 118].

This issue was especially relevant in Publication [II], where the primary objective was to learn shallow circuits capable of preparing quantum states efficiently. The challenge lies in balancing the expressiveness of the circuit with the need to maintain a low depth suitable for noise-free state preparation.

In contrast, deeper circuits may overcome local minima through increased expressivity, but face heightened risk of barren plateaus if the ansatz is highly unstructured or strongly entangling [114, 117]. Thus, the complexity of the optimization landscape significantly depends on the ansatz expressivity. Underparameterized ansätze feature numerous poor-quality minima, while overparameterized ansätze exhibit flatter landscapes with most minima approaching global optimality [119].

To mitigate local minima, strategies include problem-inspired ansatz designs that embed physical constraints and symmetries [118]. Carefully controlled overparameterization can also smooth optimization landscapes, but must be balanced to avoid barren plateaus [119].

5.2 Computing Gradients in VQCs

5.2.1 Finite Difference and Stochastic Approximation

Classical finite-difference methods approximate gradients numerically by evaluating the loss function at slightly perturbed parameter values. The gradient approximation for a

parameter θ_i is given by:

$$\frac{\partial C(\theta)}{\partial \theta_i} \approx \frac{C(\dots, \theta_i + \epsilon, \dots) - C(\dots, \theta_i - \epsilon, \dots)}{2\epsilon}, \quad (5.8)$$

where ϵ is a small shift. While simple, finite differences introduce truncation errors and are sensitive to measurement noise. Specifically, the number of shots necessary scales quadratically in ϵ . Choosing an optimal shift size ϵ is crucial and challenging [120].

In contrast, the Simultaneous Perturbation Stochastic Approximation (SPSA) method estimates the gradient using random perturbations across all parameters simultaneously. SPSA significantly reduces the computational overhead, requiring only two evaluations of C per gradient estimation, independent of the parameter count. Specifically, each iteration of SPSA involves randomly selecting a perturbation vector Δ_i independently sampled from a symmetric Bernoulli distribution, typically $\Delta_i \in \{\pm 1\}$. The gradient estimator is computed as:

$$\frac{\partial C(\theta)}{\partial \theta_i} \approx \left\langle \frac{C(\theta + \epsilon\Delta) - C(\theta - \epsilon\Delta)}{2\epsilon\Delta_i} \right\rangle, \quad (5.9)$$

where ϵ controls the perturbation magnitude. SPSA is robust against measurement noise due to the stochastic averaging effect but inherently yields noisy gradient estimates. Thus, practical implementations typically average several SPSA iterations to obtain reliable optimization steps [121, 122].

5.2.2 Parameter-Shift Rule

The parameter-shift rule enables the exact evaluation of gradients of quantum observables on quantum hardware, up to statistical measurement noise, without relying on finite difference approximations. It applies to parameterized gates generated by Hermitian operators with eigenvalues ± 1 , such as the Pauli operators P . For a gate of the form $U(\theta_i) = e^{-i\theta_i P/2}$, the gradient of a cost function $C(\theta)$ with respect to θ_i is given by [120]:

$$\frac{\partial C(\theta)}{\partial \theta_i} = \frac{1}{2} \left[C(\dots, \theta_i + \frac{\pi}{2}, \dots) - C(\dots, \theta_i - \frac{\pi}{2}, \dots) \right]. \quad (5.10)$$

This expression requires two evaluations of C per parameter. Consequently, the number of function evaluations scales linearly with the number of parameters N_p , that is, as $O(N_p)$. In typical hardware-efficient ansätze, the number of parameters scales as $N_p \sim dL$, where L is the system size and d is the circuit depth. The cost of a single evaluation of C also scales linearly with d , leading to an overall scaling of $O(Ld^2)$ for the gradient computation.

While the method is exact in the absence of noise, it is sensitive to statistical fluctuations in the measurement outcomes. This sensitivity becomes especially problematic when the gradient magnitudes are small. In the presence of barren plateaus, where gradients vanish exponentially with system size, the effective cost of obtaining meaningful gradients increases exponentially [120].

To derive the parameter shift rule, consider a quantum circuit in which a unitary gate $U(\theta_i) = e^{-i\theta_i P/2}$ acts on an initial state $|\psi\rangle$, followed by a measurement of an observable H . The expectation value of the circuit is

$$C(\theta) = \langle \psi | U^\dagger(\theta_i) H U(\theta_i) | \psi \rangle. \quad (5.11)$$

Differentiating $C(\theta)$ with respect to θ_i yields

$$\frac{\partial C(\theta)}{\partial \theta_i} = \langle \psi | \left(\frac{d}{d\theta_i} U^\dagger(\theta_i) \right) H U(\theta_i) | \psi \rangle + \langle \psi | U^\dagger(\theta_i) H \left(\frac{d}{d\theta_i} U(\theta_i) \right) | \psi \rangle. \quad (5.12)$$

The derivatives of the unitary and its adjoint are:

$$\frac{d}{d\theta_i} U(\theta_i) = -\frac{i}{2} P U(\theta_i), \quad (5.13)$$

$$\frac{d}{d\theta_i} U^\dagger(\theta_i) = \frac{i}{2} U^\dagger(\theta_i) P. \quad (5.14)$$

Substituting these expressions gives:

$$\frac{\partial C(\theta)}{\partial \theta_i} = \frac{i}{2} \langle \psi | U^\dagger(\theta_i) P H U(\theta_i) | \psi \rangle - \frac{i}{2} \langle \psi | U^\dagger(\theta_i) H P U(\theta_i) | \psi \rangle, \quad (5.15)$$

$$= \frac{i}{2} \langle \psi | U^\dagger(\theta_i) [P, H] U(\theta_i) | \psi \rangle. \quad (5.16)$$

This expression involves the commutator of P and H . Since P has eigenvalues ± 1 , the commutator can be rewritten [123] as:

$$[P, H] = -i \left[U^\dagger \left(\frac{\pi}{2} \right) H U \left(\frac{\pi}{2} \right) - U^\dagger \left(-\frac{\pi}{2} \right) H U \left(-\frac{\pi}{2} \right) \right]. \quad (5.17)$$

Substituting this into the previous expression gives the parameter-shift rule:

$$\frac{\partial C(\theta)}{\partial \theta_i} = \frac{1}{2} \left[C \left(\theta_i + \frac{\pi}{2} \right) - C \left(\theta_i - \frac{\pi}{2} \right) \right]. \quad (5.18)$$

This confirms the validity of the parameter-shift rule for gates generated by Hermitian operators with two distinct eigenvalues, such as the Pauli matrices.

5.2.3 Gradients via MPS

The gradients required for optimizing the feedback protocol described in Publication [III] were computed using MPS. Custom code was developed for this purpose because the general-purpose gradient algorithms in `ITensors.jl` [124] are not well-suited for performance in specialized applications. The algorithms described below were implemented in the package `mVQE.jl` [46] and used in the simulations for Publications [II].

Before addressing the gradient computation, it is necessary to describe how the quantum circuit defined by the hardware-efficient ansatz in Eq. (5.3) acts on an MPS. This evolution can be performed either exactly or approximately.

Exact

Exact evolution is possible because the CNOT gate admits a matrix product operator (MPO) representation with bond dimension 2:

$$\text{CNOT} = \begin{bmatrix} \text{C}_1 & 0 \\ 0 & \text{C}_2 \end{bmatrix}^2. \quad (5.19)$$

Using this decomposition, the output state of the full circuit in Eq. (5.3) can be exactly represented as an MPS with bond dimension $\chi = 2^{d/2}$, where d is the circuit depth.

This exact method is particularly effective for small d , where it may outperform approximate approaches. The evolution consists solely of matrix multiplications, which are computationally efficient and well-optimized on modern GPU architectures. All CNOT gates can be applied in parallel through a single matrix multiplication. Similarly, the R_y gates are applied via batched matrix multiplications. In contrast, approximate methods involve repeated SVDs and QR decompositions, which are less efficient on current GPUs. Note that gradients can be computed using any automatic differentiation engine.

Approximate

As discussed in Sec. 2.1, the bond dimension of an MPS can be reduced through truncation, which maintains an orthogonal form but decreases fidelity.

Assume the state at step i is given by $|\psi_i\rangle = U_i |\psi_{i-1}\rangle$, and the MPS is initially in left-orthogonal form. The application of a small number of unitaries approximately preserves this structure. The bond dimension can then be reduced by performing a left-to-right sweep with alternating SVD and QR decompositions.

For the hardware-efficient ansatz, the evolution at each layer consists of R_y and CNOT gates. When applying the R_y gates one can exploit that they are local and unitary and

as such do not increase the bond dimension or shift the orthogonality center:

$$U_i(\theta_i) = \text{CNOT}^{\otimes L/2} R_y^{\otimes L}(\theta_i), \quad (5.20)$$

$$|\psi'_i\rangle = R_y^{\otimes L}(\theta_i) |\psi_{i-1}\rangle, \quad (5.21)$$

and can therefore be applied directly. However, this structure is not exploited in the `ITensors.jl` [124] implementation, where the MPS is reorthogonalized after each R_y gate application, introducing unnecessary overhead.

The second step is the application of the CNOT gates, which requires additional tensor operations:

$$|\psi_i\rangle = \text{CNOT}^{\otimes L/2} |\psi'_{i-1}\rangle, \quad (5.22)$$

$$= \begin{array}{c} \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \end{array} \begin{array}{c} \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \end{array} \begin{array}{c} \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \end{array} \begin{array}{c} \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \\ \text{CNOT} \end{array} \quad , \quad (5.23)$$

$$= \begin{array}{c} \text{CNOT} \\ \text{CNOT} \end{array} \begin{array}{c} \text{CNOT} \\ \text{CNOT} \end{array} \begin{array}{c} \text{CNOT} \\ \text{CNOT} \end{array} \quad . \quad (5.24)$$

Following this step, small singular values in S can be discarded using a fidelity cutoff. The orthogonality center is then shifted leftward by a QR decomposition.

To evaluate fidelity with a target state, all unitary layers are applied after each other:

$$\langle \psi^{\text{target}} | \psi(\theta) \rangle = \underbrace{\langle \psi^{\text{target}} | U_{2d}(\theta_{2d}) \dots U_3(\theta_3) U_2(\theta_2) U_1(\theta_1) | 0 \rangle}_{\langle \psi_3^t |} . \quad (5.25)$$

The gradient can then be computed efficiently using the resulting structure:

$$\frac{\partial \langle \psi^{\text{target}} | \psi(\theta) \rangle}{\partial \theta_{2,3}} = \langle \psi_3^t | \frac{\partial U(\theta_2)}{\partial \theta_{2,3}} | \psi_1 \rangle , \quad (5.26)$$

$$= \underbrace{\langle \psi_3^t | \text{CNOT}^{\otimes L/2} \left[\prod_{i \neq 3} R_y(\theta_{2,i}) \right]}_{\langle \psi_3^{t'} |} \frac{\partial R_y(\theta_{2,3})}{\partial \theta_{2,3}} | \psi_1 \rangle , \quad (5.27)$$

$$= \begin{array}{c} \text{Circuit Diagram} \\ \text{with } R_y(\theta_{2,1}), R_y(\theta_{2,2}), \frac{\partial R_y(\theta_{2,3})}{\partial \theta_{2,3}}, R_y(\theta_{2,4}) \text{ gates} \\ \text{and } \langle \psi_3^{t'} | \text{ and } | \psi_1 \rangle \end{array} . \quad (5.28)$$

The target state $|\psi^{\text{target}}\rangle$ is evolved backward in time, while the initial product state $|0\rangle$ is evolved forward. Once both representations are available, gradients with respect to all R_y gates in a specific layer can be computed by constructing the corresponding environments.

The `ITensors.jl` [124] implementation does not exploit the commutativity of gates within each layer. Moreover, it computes derivatives with respect to the CNOT gates, which are not parametrized. These inefficiencies are addressed in the `mVQE.jl` [46] implementation.

CHAPTER 6

Neural Architectures

In this section, key neural network architectures used in Publication [II] are summarized. These architectures are employed to represent the parametrized feedback function $\theta_2 = f(M; W)$, which maps the measurement outcomes M to the parameters θ_2 for the next unitary operation (see Chap. 5).

6.1 Feedforward Neural Networks

Feedforward neural networks (FFNs) form the basis of most neural network models. They consist of sequential layers that alternate between linear transformations and nonlinear activation functions:

$$h^l(h^{l-1}; W^l, b^l) = \phi(W^l h^{l-1} + b^l), \quad (6.1)$$

where ϕ is a nonlinear activation function, such as the sigmoid or ReLU, and W^l, b^l are trainable weights and biases [125]. The universal approximation theorem states that a single hidden layer with sufficient width can approximate any continuous function on a compact domain [126].

In practice, several such layers are composed to form a deep neural network $y = f(x; W, b)$, which maps an input vector x to an output vector y . The network is trained using a task-specific loss function. In Publication [II], the loss is defined as the fidelity with respect to the target quantum state.

Training is performed using gradient descent. Gradients are calculated via backpropagation, which recursively applies the chain rule through the computational graph.

FFNs can model complex functions and extract hierarchical features. However, deep FFNs often suffer from vanishing gradients similar to those in VQE, which limit effective

learning [127]. Several architectural improvements have been proposed to address this issue.

6.2 Avoiding Vanishing Gradients

Neural networks face similar challenges as Variational quantum circuits, vanishing gradients. Deep VQCs often exhibit barren plateaus, where gradients vanish with increasing system size [113]. Several solutions exist for NNs. However, no analogous strategy is currently known for variational quantum circuits (see Sec. 5.1.1).

Residual Connections

Residual connections improve the training of deep networks by modifying layer outputs as

$$y = x + f(x), \quad (6.2)$$

where f is a learnable transformation such as a convolution or RNN block [128]. The identity mapping helps preserve gradient magnitude during backpropagation, improving stability in deep networks. Each residual connection refines the input representation by learning a correction term $f(x)$. This allows successive layers to focus on incremental improvements, which facilitates convergence and gradient flow in deep architectures.

Activation Functions

The selection of activation and normalization methods strongly affects the training dynamics and expressivity of neural networks.

The ReLU activation function is defined as $\phi(x) = \max(0, x)$. For positive inputs, it yields a constant gradient. For negative inputs, the gradient is zero. This property helps to mitigate the vanishing gradient problem in deep neural networks [129]. Unlike activation functions such as the sigmoid, ReLU does not cause repeated multiplication of small constants during backpropagation. As a result, the gradient remains stable over many layers, facilitating effective training of deep models.

The SwiGLU activation function is a gated variant of the GLU (Gated Linear Unit), designed to enhance expressivity through multiplicative interactions. It is often used in the MLP blocks of transformer architectures. Given an input vector x , the SwiGLU function is defined as:

$$\text{SwiGLU}(x; W, V, b, c, \beta) = \text{Swish}_\beta(xW + b) \odot (xV + c), \quad (6.3)$$

where W, V are learned weight matrices, b, c are learned bias vectors, and β is the

Swish scaling parameter (typically set to 1). The function $\text{Swish}_\beta(z) = z \cdot \sigma(\beta z)$ applies a smooth, non-monotonic gating to the input. The symbol \odot denotes element-wise multiplication.

This formulation splits the linear transformation of the input into two parallel branches: one gated via the Swish activation, and one linear. The gating mechanism allows the model to dynamically control feature flow, effectively deciding which components of the representation to retain or suppress. This improves model capacity, supports better gradient flow in deep networks, and leads to faster convergence in large-scale learning tasks [130].

Root Mean Square Normalization

Root Mean Square Normalization (RMSNorm) normalizes inputs based on their root-mean-square magnitude:

$$\text{RMS}(x; \gamma) = \gamma \frac{x_{i,t}}{\sqrt{\frac{1}{n} \sum_{i=1}^n x_{i,t}^2}}. \quad (6.4)$$

Here, γ is a learnable scaling parameter. The index i runs over the feature dimension, while t denotes the temporal or spatial position [131]. RMSNorm stabilizes training by normalizing across features without subtracting the mean, offering lower computational complexity than LayerNorm. By maintaining consistent input scales across layers, RMSNorm improves gradient flow and reduces the risk of exploding or vanishing gradients, particularly in deep or recurrent architectures.

6.3 Recurrent Neural Networks and Gated Units

Recurrent neural networks (RNNs) extend FFNs to sequential data. They maintain a hidden state that is updated at each time step:

$$h_t = \phi(Wx_t + Uh_{t-1} + b), \quad (6.5)$$

where ϕ is a nonlinear activation, and W , U , and b are trainable parameters [132]. RNNs are well-suited to modeling data with temporal dependencies.

An RNN can be viewed as a deterministic finite-state machine. The evolution of its hidden state is governed by the current input and internal gating. This viewpoint reveals a structural connection to tensor networks, which are in some sense linear versions of this.

Wu *et al.* [133] showed that any matrix product state (MPS) can be exactly represented

by an RNN with linear hidden-state updates. The hidden dimension of the RNN corresponds to the MPS bond dimension. The sequential updates of the RNN mirror the contraction structure of the MPS.

This structural similarity highlights a connection between RNN-based architectures and one-dimensional quantum states described by tensor networks. The main difference lies in the gating mechanism. While MPS updates are linear, RNNs introduce nonlinear control.

The quantum measurement and feedback policy described by Smith *et al.* [27] can be described by a finite-state machine. Therefore, RNNs are a natural choice for learning such feedback functions. In Publication [II], RNNs are trained to map sequences of measurement outcomes M_t to parameters for the next unitary gate sequence $U(\theta_{t,d})$.

Despite their expressivity, vanilla RNNs often suffer from vanishing or exploding gradients during training [134]. Gated architectures such as the Gated Recurrent Unit (GRU) [135] address this issue.

A GRU uses an update gate z_t and a reset gate r_t defined as

$$z_t = \sigma(W_z x_t + U_z h_{t-1}), \quad (6.6)$$

$$r_t = \sigma(W_r x_t + U_r h_{t-1}), \quad (6.7)$$

which control how much of the past state is retained or reset. The candidate state is computed as

$$\tilde{h}_t = \tanh(W_h x_t + U_h(r_t \odot h_{t-1})), \quad (6.8)$$

and the final hidden state is updated according to

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t, \quad (6.9)$$

where \odot denotes element-wise multiplication. This architecture enables selective memory updates and improves gradient stability across long sequences. In Publication [II], the GRU demonstrated superior performance compared to both standard RNNs and LSTMs.

6.4 Stacking RNNs

The RNN architecture employed in Publication [II] is illustrated in Fig. 6.1. This design is inspired by the LLaMA 3 model [136], with the self-attention mechanism replaced by a recurrent layer. The model was adapted to the quantum feedback setting to map sequences of measurement outcomes to adaptive quantum gate parameters.

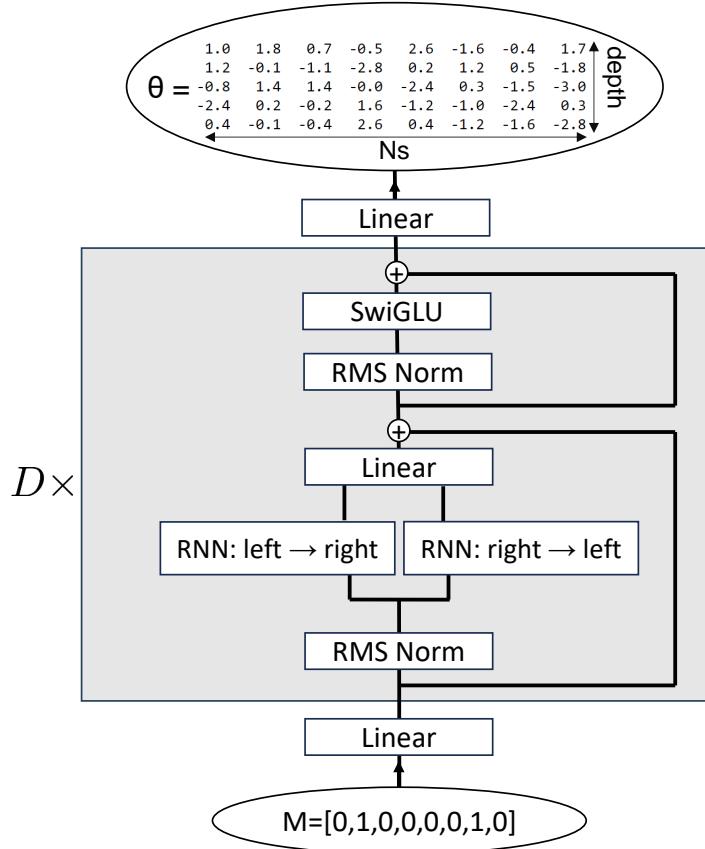


Figure 6.1.: The architecture used in Publication [II] which is shown in Fig. 8 of that work. The input measurements pass through RMS normalization, a bidirectional RNN, and SwiGLU layers, generating the output angles θ for feedback unitaries. The grey box indicates repeated application of these core layers up to depth D .

The architecture consists of multiple stacked layers of bidirectional or unidirectional GRUs. Each RNN layer receives input processed by RMSNorm and outputs hidden representations that are passed through a SwiGLU activation function. These representations are then fed into the next RNN block in the stack. The overall depth of the model is denoted by D , which refers to the number of such RNN-SwiGLU blocks applied sequentially.

For bidirectional RNNs, the forward and backward passes are concatenated at each layer to form the input for the subsequent layer. This enhances the model's ability to extract correlations across both left and right measurement outcomes. In contrast, unidirectional RNNs were augmented with a one-dimensional convolution layer (kernel size 5) as the first input layer. This convolutional preprocessing ensures that the unidirectional RNN receives sufficient contextual information from neighboring measurements, particularly those to the right of a given position.

To address the vanishing gradient problem associated with training deep recurrent models, GRUs were used in place of vanilla RNNs. As mentioned in Sec. 6.3, GRUs allow for selective memory retention and controlled forgetting, enabling the model to propagate gradients more effectively across long sequences. Empirical comparisons with LSTM units showed that GRUs performed better in this setting.

While the stacked RNN model demonstrated strong results for moderate system sizes, its effectiveness diminished for longer sequences. This limitation stems from the inherent difficulty RNNs face in modeling long-range dependencies. Consequently, alternative architectures such as transformers or Mamba-based models are recommended for future work to address these challenges more effectively.

6.5 Transformers and Mamba Architectures

Transformers use self-attention to model pairwise dependencies between sequence elements. Each token attends to all others, enabling global context modeling in a single computational step [137]. The use of residual and normalization layers improves training stability and gradient flow.

Self-attention has a quadratic cost in sequence length, which limits its scalability. This has led to alternative designs such as Mamba, which builds on structured state-space models.

Mamba replaces attention with input-conditioned state transitions. Its recurrence operates in linear time and memory, making it well suited for long sequences [138].

The recurrence in Mamba is linear and sequential. This structure resembles the contraction pattern of MPS. This analogy suggests applications in quantum-inspired models. Mamba could offer a promising approach to modeling long-range correlations, such as those encountered in Publication [II], where measurement outcomes exhibit non-local dependencies. Its linear complexity and structural resemblance to MPS make it a natural extension of the ideas explored in earlier sections, bridging classical sequence models with tensor-network representations.

Publications

PUBLICATION I

Quantum State Preparation via Engineered Ancilla Resetting

Authors:	Daniel Alcalde Puente , Felix Motzoi, Tommaso Calarco, Giovanna Morigi, Matteo Rizzi
Year:	2024
Journal:	Quantum
Volume/Issue:	8, 1299
DOI:	10.22331/q-2024-03-27-1299

Summary

Quantum state preparation is a central task in quantum simulation and computation. On noisy intermediate-scale quantum (NISQ) devices, it remains a major bottleneck due to environmental decoherence and limited coherence times [1, 2]. Traditional unitary approaches for state preparation typically require deep circuits (see Sec. 4.1) and high-fidelity gates, which are challenging to implement experimentally [6, 9].

To address these limitations, this work analyzes a measurement-based protocol that employs periodically reset ancillary qubits to steer the system toward a target quantum state. The protocol induces effective non-unitary dynamics through coherent system-ancilla coupling combined with reset operations (see Sec. 4.2). It extends earlier studies [19, 21] by investigating finite reset intervals, moving beyond the idealized instantaneous reset limit.

The protocol is applied to the preparation of the Affleck-Kennedy-Lieb-Tasaki (AKLT) state, a symmetry-protected topological (SPT) state relevant for quantum memory and computation (see Sec. 2.1.4). Exploiting the frustration-free property of the AKLT

Hamiltonian, local mapping operators are constructed to guide the system into the desired state through repeated ancilla measurements and resets.

Numerical simulations using MPS and quantum trajectory methods characterize the protocol across weak-coupling (Markovian) and strong-coupling (non-Markovian) regimes. An optimal reset interval is identified that balances rapid convergence with low entanglement entropy production. Furthermore, criteria for constructing effective mapping operators based on commutation relations are introduced to improve convergence.

The protocol is shown to be robust against realistic noise sources, including dephasing. In addition, adaptive stopping conditions based on ancilla measurement outcomes are proposed to enhance practical implementation.

This work complements the adaptive measurement-feedback approach presented in Publication [II] by offering an analytical alternative. It eliminates the need for conditional operations, which are currently difficult to realize on quantum hardware. Moreover, it lays the groundwork for extending measurement-based preparation protocols to two-dimensional systems.

In particular, it would be of interest to simulate the protocol on a honeycomb lattice, leveraging insights from Publication [III]. The AKLT state on a honeycomb lattice is known to serve as a universal resource for measurement-based quantum computation [56].

Future work should also explore the integration of variational optimization techniques, similar to those employed in Publication [II], to enhance convergence and robustness when analytical direction is unavailable.

Author's Contribution

Daniel Alcalde Puente designed all algorithms in collaboration with Matteo Rizzi. He implemented the algorithms and performed all numerical simulations and analyses. He wrote the entire initial draft of the manuscript, with the exception of Section 2.1. Throughout the project, he received guidance and regularly discussed key ideas and results with his supervisors.

Data Availability

The simulation data presented in this publication is stored on PGI-8's backed-up data servers at Forschungszentrum Jülich (FZJ) and can be made available upon request from either the author of this thesis or the Institute for Quantum Control (PGI-8).

Quantum state preparation via engineered ancilla resetting

Daniel Alcalde Puente ^{1,2}, Felix Motzoi ¹, Tommaso Calarco ^{1,2,3}, Giovanna Morigi ⁴, and Matteo Rizzi ^{1,2}

¹Forschungszentrum Jülich, Institute of Quantum Control, Peter Grünberg Institut (PGI-8), 52425 Jülich, Germany

²Institute for Theoretical Physics, University of Cologne, 50937 Köln, Germany

³Dipartimento di Fisica e Astronomia, Università di Bologna, 40127 Bologna, Italy

⁴Theoretical Physics, Department of Physics, Saarland University, 66123 Saarbrücken, Germany

In this theoretical investigation, we examine the effectiveness of a protocol incorporating periodic quantum resetting for preparing ground states of frustration-free parent Hamiltonians. This protocol uses a steering Hamiltonian that enables local coupling between the system and ancillary degrees of freedom. At periodic intervals, the ancillary system is reset to its initial state. For infinitesimally short reset times, the dynamics can be approximated by a Lindbladian whose steady state is the target state. For finite reset times, however, the spin chain and the ancilla become entangled between reset operations. To evaluate the protocol, we employ Matrix Product State simulations and quantum trajectory techniques, focusing on the preparation of the spin-1 Affleck-Kennedy-Lieb-Tasaki state. Our analysis considers convergence time, fidelity, and energy evolution under different reset intervals. Our numerical results show that ancilla system entanglement is essential for faster convergence. In particular, there exists an optimal reset time at which the protocol performs best. Using a simple approximation, we provide insights into how to optimally choose the mapping operators applied to the system during the reset procedure. Furthermore, the protocol shows remarkable resilience to small deviations in reset time and dephasing noise. Our study suggests that stroboscopic maps using quantum resetting may offer advantages over alternative methods, such as quantum reservoir engineering and quantum state steering protocols, which rely on Markovian dynamics.

1 Introduction

Quantum technologies hold immense potential for addressing significant challenges in quantum simulation, communication, and information processing. Consequently, the preparation of fiducial quantum states and the development of noise-resistant state preparation routines are of paramount importance, especially for noisy intermediate-scale quantum (NISQ) devices [1].

Quantum state preparation by means of unitary quantum circuits often faces challenges due to the required circuit complexity [2], which may demand large circuit depths and thus require low error rates, unattainable in the NISQ era. Various approaches have been proposed, starting from adiabatic techniques [3, 4, 5] and sequences of unitary transformations [6, 7].

An alternative, and possibly more robust, strategy implements protocols based on non-unitary dynamics. One strategy goes under the name of reservoir engineering [8, 9, 10, 11, 12] and aims at designing effective master equations whose steady state is the target state of the quantum state preparation. Quantum reservoir engineering has been proposed for quantum computing [11] and for robust preparation of many-body quantum states [13, 14]. Designing Markovian master equations, however, limits to performing operations in the weak coupling limit, leading to slow convergence speeds, which could be counterbalanced by including time-dependent controls [15].

Further strategies pursue quantum state preparation using adaptive measurements, as for example in Ref. [16, 17]. These protocols are efficient, yet require the use of feedback based on non-local measurements, which is not yet widely accessible on current quantum platforms.

In particular, local feedback control has only been achieved experimentally in a few cases, and it underperforms coherent entanglement gates by about two orders of magnitude [18, 19, 20]. Preparing non-local states between spatially separated elements [21, 22, 23] or applying feedback to larger lattices [24] is expected to create an even larger gap.

A procedure for preparing quantum states from a fiducial one, which is not adaptive and is based solely on local operations, can be achieved through quantum state steering via measurements and reset of ancillary degrees of freedom [25, 26, 27, 28]. It is known that for unitary dynamics the use of ancillas can dramatically improve the speed of unitary quantum circuits [7, 29, 30, 31]. Likewise, coherent feedback is known to have improved best-case scaling compared to measurement-based approaches for state preparation [32]. Here, we make use of ancilla coherence to accelerate dissipative dynamics, where the ancillas act as a non-Markovian memory within the engineered reservoir [33]. Thus, in this study, we explore the transition from quantum reservoir engineering to quantum state steering via the central tool of a stroboscopic map, which we turn from representing just an unraveling of the Lindblad master equation to becoming the key to go beyond.

Here, we test this idea on a protocol introduced by Roy *et al.* [28] for quantum state steering into states with frustration-free parent Hamiltonians. Specifically, we will be investigating the AKLT state, a spin-1 symmetry-protected topological state (SPT) [34, 35] with spin- $\frac{1}{2}$ edge states, is used in this study as a test case for the proposed protocol. This state has potential applications in the study of phases of matter and as a resource for Measurement-Based Quantum Computing [36]. The protocol is applicable to any state with a frustration-free parent Hamiltonian, like for example the cluster state. In this work, we select the AKLT state for both its simplicity and its potential usefulness. In a nutshell, we utilize measurements and resets to eliminate undesired portions of the Hilbert space from the available dynamical trajectories of the quantum state. Thus, we effectively drive the system into the target state of interest.

In order to do so, and as illustrated in Fig. 1, the system is coupled to ancillas with a time-

independent Hamiltonian for a time period of δt . The protocol exploits the frustration-free nature of the parent Hamiltonian, enabling the writing of local operators that map from locally excited states to locally unexcited states. Coupling these local operators to ancilla qubits results in a bias towards the desired non-unitary dynamics if the ancillas are reset periodically. On a different note to previous works, we determine the convergence time as a function of the resetting time δt , and investigate the stroboscopic map in regimes beyond the Markovian limit. Note that in the weak-coupling/Markovian limit, our protocol exhibits similarities to the Zeno effect, as rapid successive measurements effectively freeze the state's evolution. However, a key difference arises in stronger coupling regimes, where some evolution is permitted, and the ancillas are reset to their initial states, distinguishing the protocol from the Zeno effect setup.

This study is organized as follows: In Sec. 2, the potential of a stroboscopic map, defined as $\rho(t + \delta t) = \Lambda_{\delta t} [\rho(t)]$, for preparing desired quantum states is investigated. In Sec. 3, it is demonstrated how the protocol can be used to steer into the AKLT state, and a criterion for selecting optimal mapping operators that minimize both entanglement generation and convergence times is proposed, since the choice of mapping operators is non-unique.

In Sec. 4, full simulations of the system and ancillas are conducted using Matrix Product States (MPS) and quantum trajectories to examine the stroboscopic map in regimes beyond the purely Markovian limit. Furthermore, optimal measurement intervals are determined, and two distinct limits of the system's evolution are identified: the *weak-coupling* and the *strong-coupling* limits. In the weak-coupling limit, the system exhibits Markovian behavior akin to Lindblad dynamics, with an increased convergence speed for larger values of δt . This behavior persists up to a certain threshold for δt .

Conversely, in the strong-coupling limit, the convergence rate decreases as the time between successive measurements and resets grows. Remarkably, a broad optimum between these two regimes is uncovered, the position of which can be estimated using a simplified model of commuting mapping operators. In Sec. 5, the simplified model of commuting mapping operators

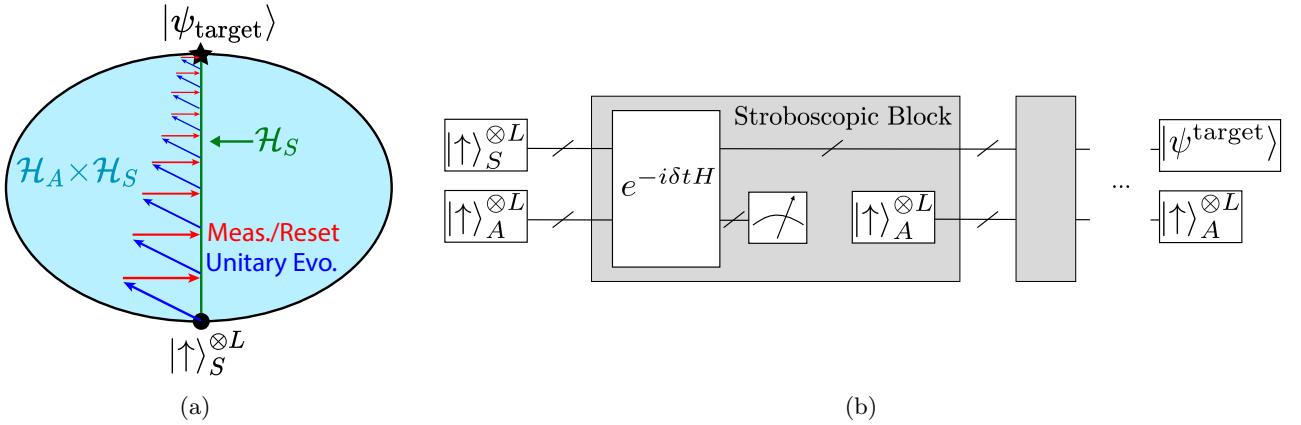


Figure 1: This figure illustrates the stroboscopic map. Subfigure (a) depicts the quantum systems evolution, where the state is initially $|\uparrow\rangle_S^{\otimes L}$, which resides entirely within the system's Hilbert space (green line). The unitary evolution with the steering Hamiltonian for a duration of δt (blue arrows) entangles the system and the ancillas, leading the state to occupy the combined Hilbert space $\mathcal{H}_S \times \mathcal{H}_A$ (light blue surface). Measurement or resetting of the ancillas (red arrows) project back the state into \mathcal{H}_S , bringing the system closer to the target state. Upon repeated application of this procedure, the system converges to the target state $|\psi_{\text{target}}\rangle$. Subfigure (b) represents the same process as a quantum circuit.

is employed to pinpoint optimal mapping operators and to shed further light on the relationship between minimizing the entropy generated by the protocol and the convergence time.

Lastly, the stability of the protocol to dephasing noise and the introduction of a stopping time selection scheme based on information obtained from the ancilla measurements is discussed in Sec. 6.

In conclusion, our study provides valuable insight into the different regimes of state preparation beyond traditional Markovian dynamics and the potential applications of the stroboscopic map for quantum state preparation.

2 The protocol

In this Section, we summarize the general features of our protocol for ancilla-assisted quantum state preparation, as illustrated in Fig. 1. The idea has its roots in quantum reservoir engineering [8, 9, 10, 11, 12]. The dynamics are designed such as to pull a quantum system into a target state ρ_0 as the result of the interplay between coherent and incoherent dynamics. The target state is thus the steady state. Assuming that the system's dynamics can be described by a linear map Λ , the target state is a fixed point of the map:

$$\rho_0 = \Lambda[\rho_0]. \quad (1)$$

In our work, the system is a many-body system, and the target state is a symmetry-protected topological state. In order to benchmark our analysis we analyze the quantum state preparation of a spin-1 chain in the AKLT state. Below we review the general idea for generating the map Λ and then describe the specific implementation for the case here considered.

2.1 Time-periodic master equation

In order to generate Λ , the system S is coherently coupled to an ancilla A . The ancilla A subsequently experiences non-unitary dynamics, characterized by a periodic resetting of its state to a reference state. The composite Hilbert space of the system and ancilla is denoted by $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_A$. The time-dependent master equation governing the density operator $\chi(t)$ in \mathcal{H} is represented by $\mathcal{L}(t)$, given as:

$$\mathcal{L}(t)\chi(t) = \frac{1}{i\hbar}[H, \chi(t)] + \sum_n \delta(t - n\delta t)\mathcal{K}\chi(t) \quad (2)$$

In this equation, the Hamiltonian H describes the interaction between the system and the ancilla. The superoperator \mathcal{K} , which acts periodically on the ancilla with a period of δt , is responsible for resetting its state, and assumed to act on a much faster timescale than the dynamics. Master equations such as Eq. 2 have been explored in

the literature to describe the dynamics of masers pumped by beams of atoms [37, 38, 39, 40, 41] and, more generally, in order to extend Floquet theory to open quantum systems [42]. In some cases, the ancilla is reset to different states as a function of n [40].

In the context of microwave cavity quantum electrodynamics, the operator \mathcal{K} captures the effect of an atom interacting with the maser at short intervals of time. When these interaction times are much shorter than the other time scales of the dynamics, they are approximated by effective kicks occurring at the instants $t_n = n\delta t$ [39, 41]. In our case, instead, the interaction time between the system and the ancilla stretches over a finite interval δt , in which they become entangled. These dynamics are encompassed by the Hamiltonian

$$H = H_S + H_A + H_{SA} \quad (3)$$

where H_S and H_A denote the Hamiltonian of subsystems S and A, respectively, and H_{SA} their interaction.

The dynamics are generated by Eq. 2, with the additional constraint that superoperator \mathcal{K} leads to an instantaneous resetting of the state of the ancilla to a reference state. The density matrix immediately after the instant of time t_n is given by

$$\chi(t_n^+) = \rho(t_n) \otimes \varrho^A, \quad (4)$$

where $t_n^\pm = \lim_{\epsilon \rightarrow 0^+} t_n \pm \epsilon$. In this equation, ϱ^A is the ancilla's reference state, and $\rho(t_n)$ is the system's state obtained by tracing out the ancilla degrees of freedom from the density operator immediately before the resetting:

$$\rho(t_n) = \text{Tr}_A\{\chi(t_n^-)\}. \quad (5)$$

The protocol consists of identifying the operator H and the superoperator \mathcal{K} that leads to the desired map Λ .

Some general considerations are in order. We first notice that, for finite periods δt , the asymptotic state of the system ρ_{SS} is periodic, namely, $\rho_{SS}(t + \delta t) = \rho_{SS}(t)$. The state ρ_0 is the time-average of $\rho_{SS}(t)$ over a period δt and the map Λ is the one corresponding to the time-averaged dynamics, $\rho_0 = \int_0^{\delta t} d\tau \rho_{SS}(\tau)/\delta t$. The limit $\delta t \rightarrow 0$ leads to a time-independent master equation [39]. This master equation can be reduced to a time-independent Born-Markov master equation for

the density matrix of the system ρ when the ancilla-system coupling is sufficiently weak: Assuming that J is the characteristic frequency characterizing the system-ancilla coupling, then the weak-coupling limit corresponds to $J\delta t \ll 1$ [41]. The strong-coupling limit, on the other hand, can be characterized by limit cycles. The latter have been discussed in microwave cavity quantum electrodynamics [37, 38, 41] and were instrumental in order to prepare Fock states of a microwave resonator [43, 44]. Limit cycles in the time-periodic master equation have been also proposed for the quantum state preparation of trapped ions [45] and ultracold atoms [46]. In the following, we consider a protocol for preparing a spin-1 chain into the AKLT state and systematically analyze its efficiency as a function of the periodicity δt , ranging from the weak-coupling to the strong-coupling limit.

2.2 Designing the system-ancilla coupling

Now, consider the density matrix $\chi(t)$. Its stroboscopic evolution is expressed as:

$$\chi(t_n^+) = \text{Tr}_A\{e^{-iH\delta t}\chi(t_{n-1}^+)e^{iH\delta t}\} \otimes \varrho^A. \quad (6)$$

A significant challenge in quantum reservoir engineering is identifying suitable Hamiltonian operators H to design the target dynamics. For a map Λ , corresponding to a Lindblad time-independent master equation, one must identify the Kraus operators and construct the associated Hamiltonian. This procedure is similar to unraveling the master equation [47, 48]. A convenient approach is outlined in [28], which applies the von Neumann theory of measurement to develop protocols that steer quantum systems towards pure target states. According to this approach, $H_S = H_A = 0$ and the ancilla-system Hamiltonian adopts the form:

$$H_{SA} = J \sum_{l,\alpha} M_{l,\alpha}^S \otimes D_{l,\alpha}^{A\dagger} + \text{h.c} \quad (7)$$

Here the parameter J is the coupling strength, $D_{l,\alpha}^{A\dagger}$ are local operators that map the ancillas into a state orthogonal to the initial state, and $M_{l,\alpha}^S$ maps the system into the target subspace. More specifically:

$$D_{l,\alpha}^{A\dagger} = |\overline{\Phi_{l,\alpha}^A}\rangle \langle \Phi_{l,\alpha}^A| \quad (8)$$

$$M_{l,\alpha}^S = |\psi_{l,\alpha}^{\text{target}}\rangle \langle \psi_{l,\alpha}^{\text{undesired}}|. \quad (9)$$

In these equations, $|\Phi^A\rangle$ represents the initial ancilla state, and $|\overline{\Phi}_{l,\alpha}^A\rangle$ are ancilla states chosen to be orthogonal to the initial state $\varrho^A = |\Phi^A\rangle\langle\Phi^A|$ the ancilla is periodically reset to. The states $|\psi_{l,\alpha}^{\text{undesired}}\rangle$ span the undesired subspace, while the states $|\psi_{l,\alpha}^{\text{target}}\rangle$ only need to span part of the target subspace. The index l represents the lattice site where the operator is applied, and α denotes the operator type. The target state is a pure state, $|\psi_0\rangle$, and resides outside the local subspaces of all the $|\psi_{l,\alpha}^{\text{undesired}}\rangle$, rendering it a fixed point of the dynamics ($[H, \rho_0 |\Phi^A\rangle\langle\Phi^A|] = 0$) and ensuring convergence. Note that the $|\psi_{l,\alpha}^{\text{target}}\rangle$ do not need to span the entire target subspace. For the protocol to converge to the target state, two conditions must be met. First, the mapping operators must steer out of the undesired subspace, so that any state that is in the undesired subspace will be mapped to a state closer to the target state. Second, the mapping operators must be chosen so that their effects at neighboring lattice sites do not disrupt each other, as this can potentially lead to the dynamics being trapped in metastable states. We formalize "not disrupting each other" as mapping operators at nearest neighbors having a sufficiently small commutation relation (see Sec. 5). This will be demonstrated in the next sections for the AKLT state, where the $|\psi_{l,\alpha}^{\text{target}}\rangle$ actually span only a part of the target subspace.

Let us first consider the weak-coupling limit or Markovian limit, $J\delta t \ll 1$. Then, the system's dynamics $\rho(t) = \text{Tr}_A(\chi(t))$ can be described by an effective Lindblad equation [39, 41, 28] :

$$\partial_t \rho(t) = J\delta t \sum_{l,\alpha} \left(M_{l,\alpha}^S \rho(t) M_{l,\alpha}^{S\dagger} - \frac{1}{2} \{ M_{l,\alpha}^{S\dagger} M_{l,\alpha}^S, \rho(t) \} \right). \quad (10)$$

where time is now rescaled by $1/J$ to be unitless $Jt \rightarrow t$. In this limit, the mapping operators become effective jump operators that map into the target state. This equation suggests, however, that the rate of convergence towards the steady state increases with δt , and is thus expected to be found in the strong-coupling regime, where Eq. 10 does not apply.

Generally, determining the optimal period δt for the fastest convergence rate is a complex task and often requires numerical approaches [41]. However, in a specific limit, an explicit derivation

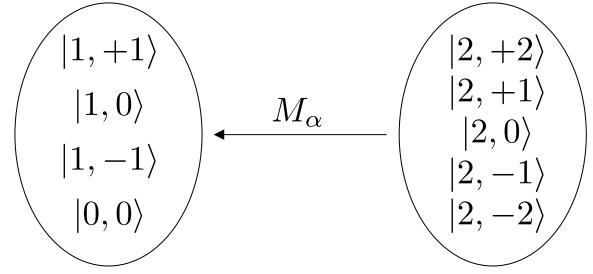


Figure 2: The mapping operators M_α are defined such that all spin states resulting from the combination of two spin-1 sites map from the spin-2 subspace to the spin-(0,1) subspace.

of an expression for the fidelity evolution for generic δt is possible. This requires constraining the mapping operators to satisfy the commutation relations:

$$[M_{l,\alpha}^S, M_{l',\alpha'}^S] = 0 \quad \wedge \quad [M_{l,\alpha}^S, M_{l',\alpha'}^{S\dagger}] = 0 \quad \forall l \neq l'.$$

In this case, the fidelity $\mathcal{F}(t) = \langle \psi_0 | \rho(t) | \psi_0 \rangle$ evolves according to:

$$\begin{aligned} \mathcal{F}(t + \delta t) = & \mathcal{F}(t) \\ & + \sin(J\delta t)^2 \sum_{l,\alpha} \langle \psi_0 | M_{l,\alpha}^S \rho(t) M_{l,\alpha}^{S\dagger} | \psi_0 \rangle. \end{aligned} \quad (11)$$

This equation demonstrates that selecting $J\delta t = \frac{\pi}{2}$ leads to the fastest convergence. Although it is only possible to steer into simple quantum states when using operators that satisfy the commutation relation above, this result offers valuable insights into mapping operators that do not entirely fulfill the commutation relations. Specifically, in Sec. 5 we derive mapping operators that optimize the protocol efficiency taking this into consideration.

Before concluding this section, it is important to highlight a key observation. While in the weak-coupling limit, the ancilla and system essentially exist in a separable state due to the Markovian nature of the limit [49], they become entangled in the strong-coupling limit, revealing the protocol's non-Markovian nature. In this regard, the quantum nature of the ancilla is crucial to the protocol. Furthermore, the reset operation of the ancilla state is effectively a quantum resetting [50, 51, 52].

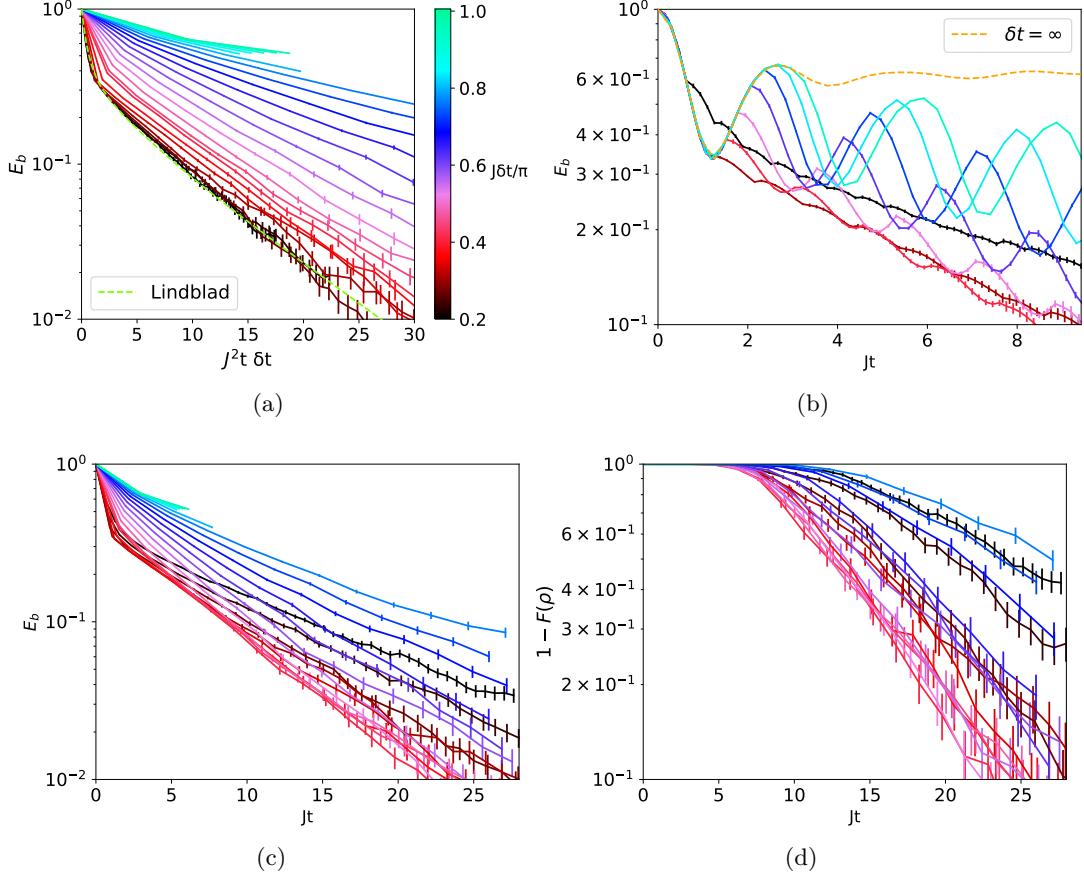


Figure 3: This Figure shows the convergence of the protocol for different measurement intervals δt in a system of size $L_s = 15$. Figure a) compares the dynamics obtained from projectively measuring the ancillas with the effective Lindblad dynamics ($J\delta t \rightarrow 0$). It is expected that the measurement-induced behavior matches the Lindblad dynamics up to a factor of $J\delta t$ in time (Eq. 10), and this seems to be approximately the case for $\delta t < \frac{\pi}{4}$. Note the rescaled x-axis by a factor of $J\delta t$ to make the comparison to the Lindblad dynamics feasible. Figure b) shows the periodic behavior of the energy per bond $E_b(t) = \frac{1}{L_s-1} \text{Tr} [H_{\text{AKLT}} \rho(t)]$. Figure c) shows the smoothed version of the energy, where no data is displayed when the dynamics cannot be accurately computed. Lastly, Figure d) shows the average infidelity $1 - F(\rho) = 1 - \sum_i^4 \langle \text{AKLT} |_i \rho | \text{AKLT} \rangle_i$ with respect to the four AKLT states.

3 Mapping Operators For The AKLT State

The 1D spin-1 Affleck-Kennedy-Lieb-Tasaki (AKLT) state is a relevant example for evaluating the effectiveness of the stroboscopic map in preparing quantum states. The AKLT state is known for its symmetry-protected topological (SPT) [34, 35] properties and spin- $\frac{1}{2}$ edge states. It has also been proposed as a resource in measurement-based quantum computing [36]. The AKLT state can be prepared with the stroboscopic map due to the frustration-free nature of its parent Hamiltonian, which allows for the use of local operators to map excited states to unexcited ones.

The AKLT state is a ground state of the fol-

lowing Hamiltonian:

$$\begin{aligned}
 H_{\text{AKLT}} &= \sum_l P_l^2 (\vec{S}_l + \vec{S}_{l+1}) \\
 &= \sum_l \left[\frac{1}{2} \vec{S}_l \cdot \vec{S}_{l+1} + \frac{1}{6} (\vec{S}_l \cdot \vec{S}_{l+1})^2 + \frac{1}{3} \right]
 \end{aligned} \tag{12}$$

Here, $P_{l,l+1}^2 = P^2(\vec{S}_l + \vec{S}_{l+1})$ are operators that project two neighboring spin-1 \vec{S}_l and \vec{S}_{l+1} into the spin-2 subspace. Note that the Hamiltonian was chosen to be dimensionless. Four distinct AKLT states exist in the subspace spanned by the (zero, one) basis, characterized by the property $P_{l,l+1}^2 | \text{AKLT} \rangle_i = 0 \quad \forall l$.

To prepare the AKLT state using a stroboscopic map, local operators are required to map

from the spin-2 subspace to the spin-(0,1) subspace. To accomplish this, a set of operators M_α must be defined such that all states in the spin-2 subspace are mapped to the spin-(0,1) subspace (see Fig. 2). There are many different choices for the mapping operators M_α , but different choices will lead to drastically different convergence speeds. In the main text, we used the mapping operators:

$$\begin{aligned} M_{l,1}^S &= \left[|1,1\rangle\langle 2,2| + \frac{1}{\sqrt{2}} |1,0\rangle\langle 2,0| + \right. \\ &\quad \left. |1,-1\rangle\langle 2,-2| \right]_{l,l+1} \\ M_{l,2}^S &= \left[|1,1\rangle\langle 2,1| + \frac{1}{\sqrt{2}} |1,0\rangle\langle 2,0| + \right. \\ &\quad \left. |1,-1\rangle\langle 2,-1| \right]_{l,l+1}. \end{aligned} \quad (13)$$

These operators were chosen based on an ad hoc condition, to try to conserve quantum numbers as much as possible and thus improve experimental realizability. This condition excludes mapping operators that significantly change the quantum numbers, such as mapping from angular momentum 2 to -1, i.e. $|1,-1\rangle\langle 2,2|$. Note that these mapping operators are not steering into the $|0,0\rangle$ state, this is not an issue as long as they are steering away from all the states in the spin two subspace.

In the rest of this paper, we will analyze the protocol using the mapping operators in Eq. 13. Before concluding this section we note that the choice of these operators is not optimal. The mapping operators, in fact, can be further optimized on the basis of two important considerations.

The first observation to make is that all the states in the spin-2 subspace should be mapped out with equal strength, leading to the imposition of a condition that the mapping operators must sum to the Hamiltonian:

$$H_{\text{AKLT}} = \sum_l \left(M_{l,1}^{S\dagger} M_{l,1}^S + M_{l,2}^{S\dagger} M_{l,2}^S \right). \quad (14)$$

In the limit of $J\delta t \rightarrow 0$, combining this with Eq. 10 yields:

$$\begin{aligned} \partial_t \rho(t) &= -\frac{J\delta t}{2} \{ H_{\text{AKLT}}, \rho(t) \} \\ &\quad + J\delta t \sum_{l,\alpha} M_{l,\alpha}^S \rho(t) M_{l,\alpha}^{S\dagger}. \end{aligned} \quad (15)$$

The first term in Eq. 15 corresponds to imaginary time evolution, while the second term maps the state from the undesired to the desired subspace. When measuring the ancillas, effectively the imaginary time evolution term is applied when the ancillas are measured to be in the initial state $|\Phi^A\rangle$, and the second term is effectively applied when they are found to be in an excited state. Note that the mapping operators in Eq. 13 satisfy this first condition.

The second consideration involves the commutation relations of the mapping operators at different lattice sites. If the mapping operators satisfy $[M_{l,\alpha}^S, M_{l',\alpha'}^S] = 0$ and $[M_{l,\alpha}^S, M_{l',\alpha'}^{S\dagger}] = 0$ for all $l \neq l'$, the dynamics will follow a sine function, allowing the target state to be reached after only one measurement with $J\delta t = \frac{\pi}{2}$. However, it is not possible to choose mapping operators for the AKLT state that fulfill these commutation relations.

Instead, the convergence time can be further optimized by minimizing the commutator between the mapping operators. The relationship between the commutation relation and convergence time is further explored in Sec. 5.

To steer a state into an AKLT state with L_s qutrits, a total of $L = 2L_s - 1$ qutrits are necessary. The ancillas, which are also spin ones, are initialized/reset to the spin-up state and are evolved according to:

$$D_{l,1}^{A\dagger} = |0\rangle\langle \uparrow|_l \quad (16)$$

$$D_{l,2}^{A\dagger} = |\downarrow\rangle\langle \uparrow|_l \quad (17)$$

$$|\Phi^A\rangle = |\uparrow\uparrow\uparrow\dots\rangle. \quad (18)$$

where $|\downarrow\rangle$, $|\uparrow\rangle$, and $|0\rangle$ are the eigenstates of the spin-1 S_z operator.

4 The Stroboscopic Protocol: results

In this section, we compare the dynamics of the stroboscopic protocol with those found in the weak-coupling limit governed by Markovian dynamics, which can be described by a Lindblad master equation. Notably, the Markovian limit exhibits characteristics of quantum reservoir engineering, providing a basis for comparing the dynamics of the stroboscopic protocol to those observed in reservoir engineering. Simulation results reveal two distinct limits for the δt parameter in the stroboscopic protocol. The first limit,

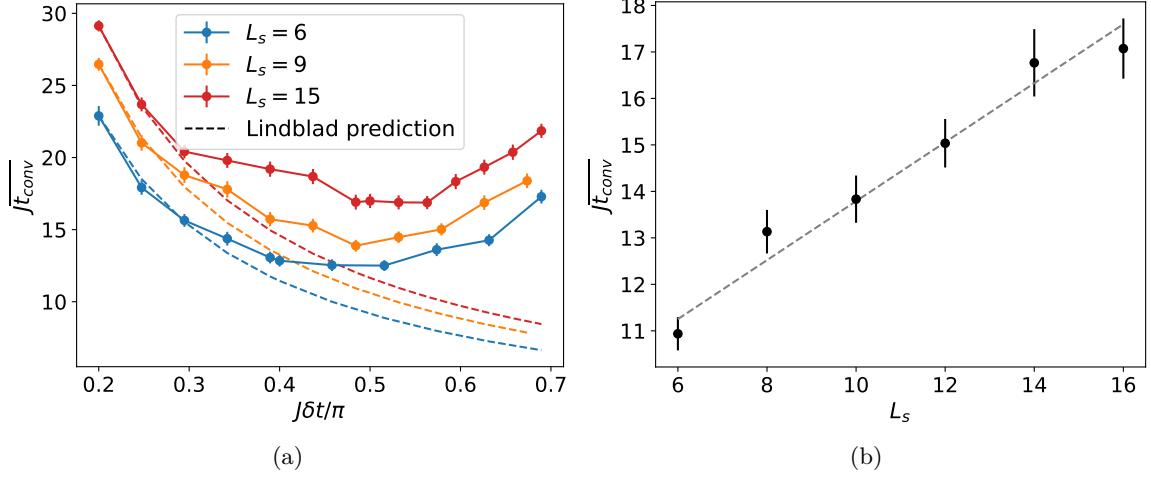


Figure 4: The figure illustrates the convergence time, denoted as t_{conv} . This is defined as the time it takes for the ancillas to stop transitioning to states other than their initial spin-up state, effectively stopping their flipping into any excited states. This marks the beginning of the linear decay of the temperature of the system. t_{conv} is plotted as a function of (a) the measurement interval δt and (b) the system size L_s . In panel (a), we see that t_{conv} decreases as $J\delta t$ increases until a minimum is reached around $\frac{\pi}{2}$, while the Lindblad approximation (dashed lines) would predict a decrease of $t_{\text{conv}} \sim \frac{1}{\delta t}$ for small δt . Panel (b) shows that t_{conv} increases linearly with L_s at a fixed time step of $J\delta t = \pi/2$, as indicated by the gray regression line.

referred to as the weak-coupling limit, adheres to a Lindbladian dynamic, while the second regime, the strong-coupling limit, demonstrates convergence at a slower rate. Furthermore, we identify an intermediate regime that exhibits the fastest convergence and resembles the exceptional case where mapping operators commute. Note that all simulations were performed with the help of the AbelianSymTensor library in Fortran.

The quantum state is initialized as a chain of alternating system and ancilla qutrits in a product state, $|\psi\rangle = |\uparrow\uparrow\dots\rangle$. To time-evolve the state, the Time-Dependent Variational Principle (TDVP) algorithm [53] is employed, with the state represented by a Matrix Product State (MPS). After each time interval δt , the ancilla qutrits are measured and reset to the spin-up state. This measurement and reset process projects the quantum state closer to the target state. Through repeated iterations and a sufficiently small δt , the state converges to the target state. This process is simulated using quantum trajectories [54].

4.1 Weak-Coupling Limit

In the weak-coupling limit, where $J\delta t < \frac{\pi}{4}$, the dynamics resemble those of Markovian dynamics. Consequently, it is expected that the conver-

gence speed should be proportional to the measurement interval, and so the convergence time should be inversely proportional $t_{\text{conv}} \propto \frac{1}{\delta t}$ (see Eq. 10). The Lindblad limit is compared to simulations at finite values of δt in Fig. 3(a). For both paradigms, the energy of the AKLT Hamiltonian per bond, $E_b(t) = \frac{1}{L_s-1} \text{Tr}[H_{\text{AKLT}}\rho(t)]$, converges to the AKLT state exponentially fast. As anticipated, for small measurement intervals $J\delta t < \frac{\pi}{4}$, the evolution at finite δt matches the solution of the effective Lindbladian for $J\delta t \rightarrow 0$. Outside of the Markovian limit, the energy per bond decays more slowly than predicted by the Lindbladian but still converges to the AKLT state.

To gain a deeper understanding of the energy evolution for finite values of δt , one can refer to Fig. 3(b). The yellow dashed line illustrates the energy evolution in the absence of measurements, offering insight into the convergence speed for various δt values. This line exhibits damped energy oscillations that would continue if measurements were to cease at any point during the protocol. The speed of convergence to the target state is determined by the energy value at which the damped oscillation is interrupted by measurement and reset, initiating a new cycle of damped energy oscillations that will be interrupted once another measurement is performed.

Two limits corresponding to different values of δt can be observed in this damped oscillation. In the weak-coupling limit, the energy initially decreases when the Hamiltonian is turned on. In the absence of measurements, the intermediate regime is entered, where the energy follows a sinusoidal pattern until it reaches saturation at a value below the initial energy in the strong-coupling limit. When the ancillas are reset, the sinusoidal energy decrease starts again from the beginning, which is different from the case without measurements. The black line in Fig. 3(b), representing the weak-coupling limit with a resetting interval of $J\delta t = 0.2\pi$, behaves similarly to the Lindbladian as discussed earlier. Since the measurement interval for the black line is fast, it is not affected by the sinusoidal nature of the measurement-free evolution. In contrast, larger values of δt in the intermediate regime $0.4\pi < J\delta t < 0.6\pi$ can exploit the sinusoidal structure to minimize the energy in each step of the process. However, in the strong-coupling limit $0.6\pi < J\delta t$, larger values of δt , measured after the sinusoidal structure has already started to increase, resulting in slower convergence times.

Fig. 3(c),(d) further demonstrates the convergence behavior of the energy/infidelity for different values of δt . Both quantities show the fastest convergence when in the intermediate regime with $J\delta t_{\text{opt}} \approx \frac{\pi}{2}$. This is consistent with the observations made in the previous paragraph.

It should be noted that for values of δt in the strong-coupling limit, accurate simulations of the system are not possible due to technical limitations. As measurements become less frequent, the quantum state accumulates more entanglement, which cannot be adequately represented by an MPS (see Sec. 4.3). As a result, the lines in Fig. 3(c),(d) are discontinued once these technical limitations arise.

4.2 Convergence times

To further examine the convergence behavior, we define the convergence time as the moment when the ancillas cease to flip into an excited state (i.e., any state other than the initial state $|\uparrow\rangle$) and instead stay in their initial state. This event occurs when the system's excitations are insufficient to induce transitions excitations in the ancillas, allowing only the imaginary time evolution term in Eq. 15 to act on the state. As a result, the

cessation of ancilla flipping corresponds to the cooling of the state, leading to a reduction in the inverse temperature of the state, denoted by $\Delta\beta(t) = t \cdot J\delta t$. Once the state is considered converged, an exponential decrease in energy, characterized by $E(t) \propto e^{-Jt \cdot \delta t \Delta E}$, where ΔE is the energy gap of the AKLT Hamiltonian, is observed (see App. A).

Fig. 4(a) displays the convergence time for the system as a function of δt for different system sizes. The optimal value of δt is found to be around $J\delta t_{\text{opt}} \approx \frac{\pi}{2}$ for all sizes, corresponding to the intermediate regime where the sinusoidal pattern of the energy evolution enables faster convergence. In contrast, smaller values of δt in the weak-coupling limit and larger values in the strong-coupling limit result in slower convergence times due to not exploiting the sinusoidal structure. It is evident that in the weak-coupling limit, the procedure will always converge with a convergence time proportional to $\frac{1}{\delta t}$ (see the dotted lines in Fig. 4(a)), as the Lindbladian guarantees it. However, the possibility of convergence in the strong-coupling limit remains uncertain and will be further explored in the following section.

The energy evolution behavior in the intermediate regime can be further understood by considering the case where the mapping operators commute, $[M_{l,\alpha}^S, M_{l+1,\alpha'}^S] = 0$ (see Sec. 2). In this case, the optimal value for δt is $J\delta t_{\text{opt}} = \frac{\pi}{2}$, and the energy fluctuates periodically with a sinusoidal pattern. This observation aligns with those made in the intermediate regime. However, if the mapping operators do not commute, the sinusoidal behavior may be altered, leading to slightly different convergence times in the intermediate regime and a complete departure from the commuting case when entering the strong-coupling limit. This commutation relation plays a significant role in finding the optimal mapping operators (see Sec. 5).

The system's convergence time, analyzed in Fig. 4(b) as a function of the system size, exhibits a linear increase with the size of the chain. This observation is attributed to the delocalized nature of the AKLT state, featuring two spin-half edge states. Preparing the AKLT state requires entangling the two edges of the chain, which, for a local protocol, scales linearly with the system size at best. Although protocols employing non-local feedback from measurement results, as demon-

strated in [16], can achieve non-scaling convergence speeds, these methods have limitations, such as the inclusion of non-local measurement-based feedback, which is not yet widely accessible on current quantum platforms. In light of this, we focus on a minimalist solution for stabilizing these topological states in the NISQ era.

4.3 Strong-Coupling limit

In the previous section, the strong-coupling limit was characterized by values of δt at which the system's energy ceases to display a sinusoidal pattern and instead approaches a constant energy value during each iteration of the protocol. This section focuses on examining the system's behavior under these values of δt , with particular emphasis on the role of entanglement entropy. Entanglement entropy is computed by partitioning a quantum system into two parts, A and B, and determining the von Neumann entropy of the reduced density matrix for part A. This measure quantifies the extent of quantum correlations between subsystems and is essential for investigating properties of quantum many-body systems, such as phase transitions and entanglement scaling. In this study, attention is directed toward the average entanglement entropy across the pure states in different trajectories characterized by the probabilistic ancilla measurement outcomes.

It is a priori unclear whether the protocol will continue to converge to the AKLT state for larger time intervals, δt , between measurements. Fig. 3(c) demonstrates that convergence decelerates as δt increases, but it remains uncertain whether the protocol will cease converging at a finite δt . In the absence of measurements, i.e., $\delta t \rightarrow \infty$, the system is expected not to converge to the AKLT state. Moreover, in the no-measurement scenario, the entanglement of the quantum system is anticipated to grow ballistically until it reaches the system's maximal entropy capacity, $\frac{L}{2} \log(3)$. This section investigates if there exists a finite value of the measurement time interval δt , at which measuring and resetting the ancillas fail to reduce the system's entropy, resulting in exponentially scaling convergence times with system size, as observed in the study of the Measurement Induced Entanglement Phase Transition (MIEPT).

Models featuring MIEPT typically have two different elements: a measurement element that

reduces entropy and a Hamiltonian evolution or random circuit element that increases entropy. In our protocol, the entropy-reducing element is the reset of the ancillas, and the entropy-increasing element is the Hamiltonian evolution. In systems exhibiting MIEPT, measurements can counter-balance entanglement growth induced by Hamiltonian evolution up to a certain ratio between the two elements. In such systems, if an insufficient number of measurements are performed, an entanglement phase transition may arise, leading to entanglement growth.

Should a phase transition manifest in the stroboscopic protocol, it would likely materialize as a deviation in the entanglement growth, potentially impeding convergence to the target state. Nonetheless, our analysis employing the entanglement-extrapolation technique outlined in App. B suggests that this model does not exhibit a phase transition. Instead, it displays characteristics of a volume law phase even for small values of δt .

The system's entanglement initially increases from its initial value, reaching a maximum at a time t_{\max} and an entanglement entropy S_{\max} (see Fig. 5(a)). Subsequently, the entropy decreases until it reaches the entanglement entropy of the AKLT state. The AKLT state's entropy lies between $\log(2)$ and $\log(4)$, depending on the entanglement degree of the two edge states. This pattern is observed for all values of δt , with the sole difference being that larger values of δt yield higher values of S_{\max} . Our aim is to ascertain whether there exists a value of δt for which the system's entanglement begins to increase ballistically. This effect would cause S_{\max} to increase to $\frac{L}{2} \log(3)$ and would exponentially slow down convergence to the target state. The increase in entropy would also hinder our ability to accurately simulate the system, as the maximum entanglement entropy that an MPS can represent is constrained by the bond dimension to $S \leq \log(\text{bond})$. Nevertheless, this limitation is partially alleviated by our entanglement-extrapolation technique, detailed in App. B.

Fig 5(b) displays the maximal entanglement entropy reached during the evolution, S_{\max} , for various system sizes, L_s . The results reveal that S_{\max} approaches the entropy of the AKLT state as the time step δt nears zero and increases linearly for large δt . As the system size expands,

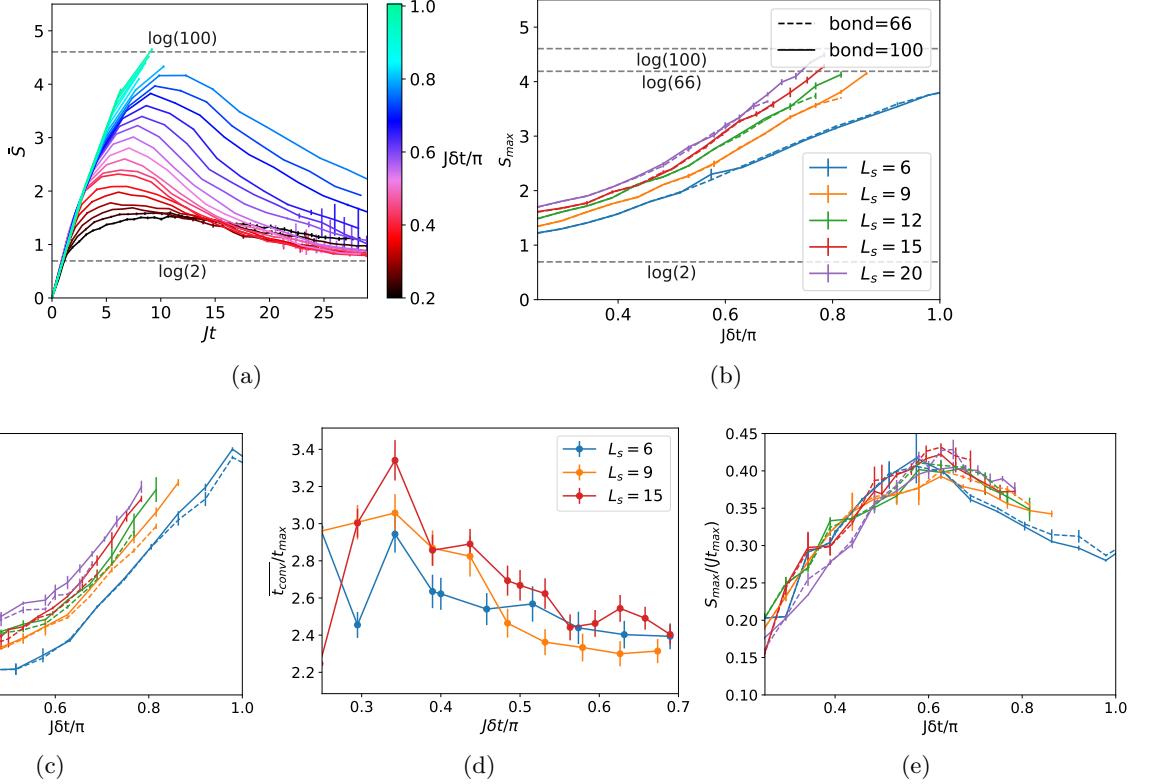


Figure 5: (a) Shows the average extrapolated entanglement \bar{S} for different measurement intervals δt at a system size of $L_s = 15$. Simulations at large δt produce more entanglement than can be accurately represented with a bond dimension of 100 and are thus discontinued when they no longer match simulations with a lower bond dimension. (b) Presents the maximum entanglement reached for different δt , which are effectively the maxima of Fig. (a) for simulations with bond dimensions of 66 (dotted line) and 100 (continuous line) (c) Shows the time required to reach the maximal entanglement. (d) Compares the convergence time to the time taken to reach the maximal entropy. (e) Displays the entanglement buildup rate to the maximum S_{\max}/t_{\max} .

S_{\max} grows approximately linearly, indicating a volume law phase. Additionally, the time required to reach S_{\max} , as depicted in Fig. 5(c), appears to increase linearly with large δt . Importantly, no lattice size-dependent discontinuity is observed in either t_{\max} , S_{\max} , or their derivatives, suggesting the absence of a phase transition. Intriguingly, even without biasing the dynamics by resetting the ancillas, the entanglement entropy attains a stationary value similar to S_{\max} observed in the reset-and-measure protocol. The only distinction to the data shown here is that the entropy does not decrease to the target state because the dynamics are no longer biased without the reset. Nonetheless, even in this modified protocol, there is no evidence of a phase transition.

Remarkably, the time required to reach maximal entanglement exhibits a behavior similar to the convergence time, with a minimum around

$J\delta t = \frac{\pi}{2}$. For large δt , t_{\max} can be accurately calculated, as shorter simulations are necessary. The same linear behavior is observed for t_{\max} as for the convergence time t_{conv} when δt is large. Specifically, t_{conv} appears to be two to three times larger than t_{\max} (see Fig. 5(d)). Since t_{\max} does not diverge at any point, it is plausible that t_{conv} will not diverge either. If this trend persists for larger δt , it would imply that the protocol converges after a fixed number of measurements for larger δt .

The analysis of the quantity $\frac{S_{\max}}{Jt_{\max}}$, illustrated in Fig. 5(e), provides valuable insights. This quantity serves multiple purposes: it reflects the average rate at which the system attains maximal entanglement entropy, and it remains invariant with respect to lattice size, suggesting the absence of a phase transition. Initially, this quantity increases due to the minimum at $J\delta t = \frac{\pi}{2}$ in t_{\max} , but it later converges to the ratio of the

corresponding slopes of the linear functions S_{\max} and t_{\max} in the strong-coupling limit. Examining this quantity further corroborates the conclusion that no phase transition exists in the system, thus emphasizing the protocol's reliability in converging to the target state.

Simulations were performed using bond dimensions of 66 and 100 to validate the findings. Owing to the necessity of executing 128 distinct trajectories for each parameter set, only small bond dimensions were employed. In particular, 13,000 simulations were required to obtain the data presented in Fig. 5(b).

5 Optimal Mapping Operators

In this section, we analyze the influence of the choice of mapping operators M_α on the dynamics and convergence time required to prepare a target state. To ensure that the mapping operators map into the AKLT state, they must satisfy two conditions: first, they must map away from the five states in the spin-2 subspace and into the spin-(0,1) subspace (see Fig. 2); second, they must contribute to the Hamiltonian as outlined in Eq. 14.

If the mapping operators satisfy the commutation relations $[M_{l,\alpha}^S, M_{l',\alpha'}^S] = 0$ and $[M_{l,\alpha}^S, M_{l',\alpha'}^{S\dagger}] = 0$ for all $l \neq l'$, the dynamics exhibit a sinusoidal pattern, and the target state can be reached after a single measurement at $J\delta t = \frac{\pi}{2}$ time, as discussed in Sec. 2. In practice, however, it is not feasible to satisfy these conditions for complex target states such as the AKLT state.

If the commutation relations are not satisfied, the operators interfere with each other, leading to the formation of undesired entanglement in parts of the system orthogonal to the target state. The relationship between entanglement growth and convergence times is explored in the previous section.

In order to minimize entanglement generation and thus reduce convergence times, the goal is to approximate the commutation relation as closely as possible. This goal is achieved by minimizing the commutation measure:

$$\text{comm}(M) = \sum_{l \neq l', \alpha, \alpha'} \| [M_{l,\alpha}, M_{l',\alpha'}] \|_2 + \| [M_{l,\alpha}, M_{l',\alpha'}^\dagger] \|_2. \quad (19)$$

For two nearest-neighbor mapping operators that are identical for all lattice sites, the measure simplifies to:

$$\begin{aligned} \text{comm}(M) = & \| [M_{l,1}, M_{l+1,1}] \|_2 + \| [M_{l,1}, M_{l+1,2}] \|_2 \\ & + \| [M_{l,2}, M_{l+1,2}] \|_2 + \| [M_{l,1}, M_{l+1,1}^\dagger] \|_2 \\ & + \| [M_{l,1}, M_{l+1,2}^\dagger] \|_2 + \| [M_{l,2}, M_{l+1,2}^\dagger] \|_2. \end{aligned} \quad (20)$$

Note that all mapping operators are also constrained by Eq. 14.

In the previous sections, the mapping operators:

$$\begin{aligned} {}_1M_{l,1}^S &= \left[|1,1\rangle\langle 2,2| + \frac{1}{\sqrt{2}} |1,0\rangle\langle 2,0| + |1,-1\rangle\langle 2,-2| \right]_{l,l+1} \\ {}_1M_{l,2}^S &= \left[|1,1\rangle\langle 2,1| + \frac{1}{\sqrt{2}} |1,0\rangle\langle 2,0| + |1,-1\rangle\langle 2,-1| \right]_{l,l+1} \end{aligned} \quad (21)$$

have been used. These mapping operators have a commutation measure of $\text{comm}({}_1M) = 18.2$ which is higher than the other mapping operators we will introduce in this section. The ${}_1M$ operators were chosen for symmetry reasons before the relationship between the commutation measure and the convergence time was found.

The effectiveness of the commutation measure as a criterion for selecting optimal mapping operators is demonstrated by analyzing a parameterized set of operators:

$$\begin{aligned} {}_2M_{l,1}^S(\alpha) &= \left[|1,1\rangle\langle 2,2| + \alpha |0,0\rangle\langle 2,-1| \right. \\ &\quad \left. + |1,-1\rangle\langle 2,-2| \right]_{l,l+1} \\ {}_2M_{l,2}^S(\alpha) &= \left[|1,1\rangle\langle 2,1| + \sqrt{1-\alpha^2} |1,-1\rangle\langle 2,-1| \right. \\ &\quad \left. + |1,0\rangle\langle 2,0| \right]_{l,l+1}. \end{aligned} \quad (22)$$

Fig. 6(c) illustrates the relationship between the convergence time and commutation measure for different values of α . Both measures exhibit a U-shaped pattern, with a minimum at $\alpha = 0.404$. This observation suggests that the commutation measure can serve as a reliable indicator for selecting optimal mapping operators.

To optimize the mapping operators, we employed gradient descent techniques to minimize the commutation measure. Each mapping operator was parameterized using a 4×5 matrix,

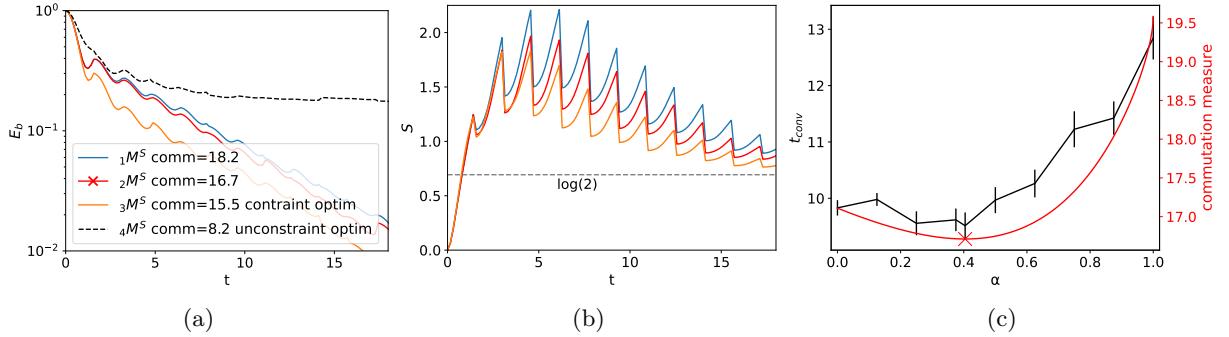


Figure 6: This figure shows the convergence behavior of different mapping operators at a system size of $L = 10$. (a) shows the energy per bond for four different mapping operators: the (1) blue curve represents the mapping operators used in the rest of the paper, the (2) red curve represents an improved version of the mapping operators, the (3) yellow curve represents mapping operators optimized to minimize the commutation measure, and the (4) gray dotted curve represents mapping operators optimized to minimize the commutation measure while not enforcing the equation (Eq. 14). (b) shows the entanglement entropy for the different mapping operators. (c) shows the convergence time and the commutation measure for mapping operators ${}_2M(\alpha)$ for different values of α . The minimum value for both measures is observed at $\alpha = 0.404$, marked by a red cross in the legend of sub-figure (a) and in sub-figure (c).

where each matrix element maps a state from the spin-2 subspace to the spin-(0,1) subspace.

The resulting optimized operators were minimized to a commutation measure of $\text{comm}({}_3M) = 15.5$ with the form:

$$\begin{aligned} \alpha &= 0.8482 \\ {}_3M_{l,1}^S &= \left[|1,1\rangle\langle 2,2| + |1,0\rangle\langle 2,1| \alpha \right. \\ &\quad \left. + |0,0\rangle\langle 2,1| \sqrt{1-\alpha^2} \right. \\ &\quad \left. + |1,-1\rangle\langle 2,0| \frac{1}{\sqrt{2}} \right]_{l,l+1} \\ {}_3M_{l,2}^S &= \left[|1,-1\rangle\langle 2,-2| + |1,0\rangle\langle 2,-1| \alpha \right. \\ &\quad \left. + |0,0\rangle\langle 2,-1| \sqrt{1-\alpha^2} \right. \\ &\quad \left. + |1,1\rangle\langle 2,0| \frac{1}{\sqrt{2}} \right]_{l,l+1}. \end{aligned} \quad (23)$$

The table below presents the results for convergence time and commutation measures associated with the three aforementioned mapping operators.

	comm	t_{conv}
${}_1M$	18.2	10.94
${}_2M$	16.7	9.51
${}_3M$	15.5	7.66

A significant observation is the correlation between the decrease in commutation measure and the decrease in convergence time. The energy

evolution for these mapping operators can be examined in Fig. 6(a). Furthermore, Fig. 6(b) shows that mapping operators with lower commutation measures produce less undesired entanglement, as hypothesized.

It is worth mentioning that the causal relationship between the commutation relation and convergence time was established only after most of the computationally expensive simulations were completed. In the previous section, the mapping operators in Eq. 21 were used. They perform approximately 20% worse than the fully optimized ones Eq. 21 but, no qualitative change between the time evolution was found. Consequently, the original choice of mapping operators was retained to conserve computational resources.

An important observation is that not enforcing the Hamiltonian constraint, as shown in Eq. 14, allows for smaller commutation measures. However, the resulting dynamics exhibit slow convergence, as shown by the black dotted line in Fig. 6(a). This slow convergence is due to the weak mapping of certain states in the spin-2 subspace to the spin-(0,1) subspace. Such cases occur when the mapping operators sum to an operator with eigenvalues unequal to one for eigenvectors in the spin-2 subspace. It should be noted that the choice of optimal mapping operators is not unique. Several alternatives with comparably small commutation measures perform equally well; however, the simplest-looking option is presented here. It is possible to apply the same

procedure with three mapping operators instead of two. Nonetheless, using three mapping operators requires additional ancilla qutrits and results in higher commutation measures due to the increased number of terms in the commutation relation. An alternative mapping approach to map away from the spin-2 subspace is to use different operators for even and odd sites. This method results in a reduced commutation measure of $\text{comm}(M) = 11.5$. However, the convergence rate of the operators obtained by this approach is slow, suggesting that further conditions must be satisfied for fast convergence rates. These conditions are likely to be satisfied automatically if the operators used for odd and even sites are identical.

6 Dephasing Noise

We test the protocol's resilience to dephasing errors, which erases quantum entanglement during the dynamics. We consider here Markovian processes described by the Lindbladian

$$\dot{\chi} = -i[H, \chi] + \epsilon \sum_{l, \alpha \in \{\uparrow, 0, \downarrow\}} \left[P_l^\alpha \chi P_l^{\alpha\dagger} - \frac{1}{3} \chi \right], \quad (24)$$

which affects both the system and ancillas through a Lindbladian. Here, the P^α operators project the system and ancillas into spin-up, spin-zero, or spin-down states.

Dephasing is implemented by performing random projection measurements without resetting, with probability $P(\epsilon) = 1 - e^{-\epsilon dt}$ applied to both the system and the ancillas.

A metric for evaluating the performance of the protocol is the stationary mean energy $\overline{E}_b = \frac{1}{k(L-1)} \sum_{i=N-k}^N \text{Tr}(H\chi(t_i))$, which indicates the degree of convergence to the target state upon reaching equilibrium (see Fig. 7(a)). For error rates of $\epsilon = 0.015$ and 0.030 , the optimal δt remains approximately $\frac{\pi}{2}$. This is consistent with the optimal δt obtained in the absence of errors, leading to the conclusion that the optimal δt is independent of the error rate.

It is crucial to evaluate the robustness of the protocol against different noise levels. For this purpose, both the infidelity and the mean energy per bond at equilibrium are plotted against the error rate in Fig. 7(b). Stopping the protocol

at a predetermined time when equilibrium is assumed to be reached is inefficient. Instead, the information obtained from measuring the ancillas can be used to determine when to stop the evolution of the state, this is done in a similar fashion to Matthies *et al.* [27]. The state's evolution is stopped when no ancillas have been measured in an excited state for a duration of $t_{\text{wait}} = 4\delta t$, at which point fidelity and energy per bond are measured. Importantly, extending the waiting time would not substantially improve fidelity to the target state, as the primary factor affecting performance at large t_{wait} is the likelihood of an error occurring during the final Hamiltonian evolution. This simple selection mechanism enhances fidelity by approximately a factor of three. This allows the self-correcting nature of the protocol to achieve even better fidelities, exceeding a fidelity of $F = 0.99$ at $\epsilon = 10^{-3}$, with the help of the stopping time selection scheme.

Fig. 7(c) shows the average time for the system to exhibit no ancilla flips for $t_{\text{wait}} = 4\delta t$. For small error rates, these values converge to the error-free convergence time discussed in Sec. 4.2. For larger error rates, an exponential growth with the error rate is observed. This exponential growth is expected, as higher error rates make it exponentially less likely to encounter a period where no excited ancillas are measured.

7 Conclusion

In conclusion, this study investigated the potential applications of a stroboscopic map for quantum state preparation, based on the periodic resetting of ancillary degrees of freedom. We focused on the protocol introduced by Gefen *et al.* [28] for unravelling quantum state steering into states with frustration-free parent Hamiltonians, and went beyond the regimes discussed there. The AKLT state was chosen as a test case due to its simplicity, topological symmetry-protected properties, and potential usefulness for quantum computation. Although the protocol is applicable to other states with similar features, such as the cluster state, the investigation primarily centered on the AKLT state. This study provided insights into different regimes dictated by the ancilla resetting time.

Two distinct limits of the system's evolution were identified: the weak-coupling limit, re-

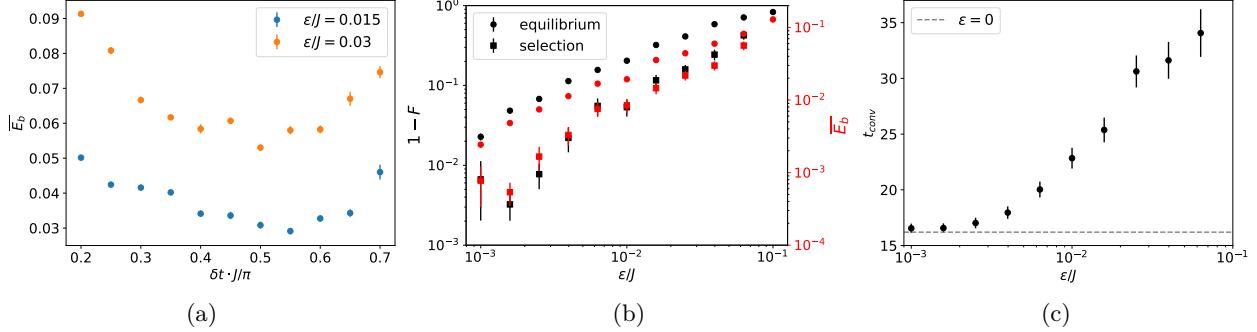


Figure 7: The figure demonstrates the dependencies of mean bond energy \bar{E}_b , infidelity $1 - F$, and convergence time t_{conv} on error rates ϵ and measurement time intervals δt for a system size of $L = 10$. In subplot (a), ϵ is held constant at 0.015 and 0.03, while δt varies. Subplots (b) and (c) employ a fixed δt of $\frac{\pi}{2}$, with ϵ serving as the variable parameter. In (b), equilibrium values for energy per bond are calculated using $\bar{E}_b = \frac{1}{k(L-1)} \sum_{i=N-k}^N \text{Tr}(H\chi(t_i))$. The selection scheme's energies (red) and infidelities (black) are assessed by measuring observables after four consecutive evolution and measurement cycles without ancilla flips. In (c), the duration required to measure four successive cycles without any ancilla flipping is depicted for different error rates. At low error rates t_{conv} converges to noise-free convergence time.

sembling Markovian dynamics and described by a Lindblad equation, and the strong-coupling limit, characterized by increasing entanglement entropy production. A broad optimum between these two limits was found around $J\delta t \simeq \frac{\pi}{2}$, a value identified in a simplified model of commuting mapping operators. In turn, a measure of non-commutativity was used to determine optimal mapping operators: the higher such a measure is, the more entanglement entropy is produced, and the slower the convergence.

In the strong-coupling limit, despite the evident peak of generated entanglement, we identified no phase transition in the model. To this end, we introduced an entanglement-extrapolation technique and highlighted that no lattice size dependent discontinuity emerged in either the time to reach the peak or its value. We, therefore, concluded that the protocol converges also at large δt , albeit slower and slower.

Furthermore, we demonstrated that the protocol is resistant to dephasing noise, achieving a fidelity $F > 0.99$ for an error rate of $\epsilon = 10^{-3}$. The optimal resetting time remained approximately $J\delta t = \frac{\pi}{2}$, and the protocol exhibited a self-correcting nature. A stopping time selection scheme based on the ancilla measurement results was introduced, which improved fidelity by approximately a factor of three.

It is worth noting that the recent work by Smith *et al.* [16] presents a method that outperforms the approach in preparing specifically

the AKLT state. However, their method relies on non-local measurement-based feedback, which is not yet widely accessible on current quantum platforms. This work contributes to the understanding of alternative or complementary state preparation methods, where the roles of measurement and non-locality may be removed.

Future research directions include studying the behavior of the protocol for different Hamiltonians and target states, as well as exploring alternative strategies for choosing the mapping operators, such as incorporating feedback and using machine learning methods. Overall, the results can contribute to the development of noise-resistant state preparation routines, especially in the context of noisy intermediate-scale quantum (NISQ) devices.

Acknowledgements

We thank M. Buchhold, S. Diehl, Y. Gefen, I. Gorny, C. Koch, and K. Snizhko for inspiring discussions on the topic. We acknowledge funding by the German Federal Ministry of Education and Research (BMBF) for support under the thematic programme “Quantum technologies – from the basics to the market”, project number 13N16202 “Noise in Quantum Algorithms (NiQ)”. This work was also partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), Project-ID 429529648 – TRR 306 QuCoLiMa (“Quan-

tum Cooperativity of Light and Matter") and Project-ID 277101999, within the CRC network TRR 183. F.M. and M.R. are further supported by the DFG under Germany's Excellence Strategy – Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC 2004/1 – 390534769. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at the Jülich Supercomputing Centre (JSC) (Grant NeTeNeSyQuMa) and the FZ Jülich for computing time on JURECA (institute project PGI-8) [55].

References

- [1] John Preskill. “Quantum Computing in the NISQ era and beyond”. *Quantum* **2**, 79 (2018).
- [2] Jens Eisert. “Entangling power and quantum circuit complexity”. *Physical Review Letters* **127**, 020501 (2021). url: <https://doi.org/10.1103/physrevlett.127.020501>.
- [3] Tameem Albash and Daniel A. Lidar. “Adiabatic quantum computation”. *Rev. Mod. Phys.* **90**, 015002 (2018).
- [4] Pimonpan Sompet, Sarah Hirthe, Dominik Bourgund, Thomas Chalopin, Julian Bibo, Joannis Koepsell, Petar Bojović, Ruben Verresen, Frank Pollmann, Guillaume Salomon, et al. “Realizing the symmetry-protected haldane phase in fermi–hubbard ladders”. *NaturePages* 1–5 (2022). url: <https://doi.org/10.1038/s41586-022-04688-z>.
- [5] Zhi-Yuan Wei, Daniel Malz, and J. Ignacio Cirac. “Efficient adiabatic preparation of tensor network states”. *Physical Review Research* **5** (2023).
- [6] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf. “Sequential generation of entangled multiqubit states”. *Phys. Rev. Lett.* **95**, 110503 (2005).
- [7] Felix Motzoi, Michael P Kaicher, and Frank K Wilhelm. “Linear and logarithmic time compositions of quantum many-body operators”. *Physical review letters* **119**, 160503 (2017). url: <https://doi.org/10.1103/physrevlett.119.160503>.
- [8] J. F. Poyatos, J. I. Cirac, and P. Zoller. “Quantum reservoir engineering with laser cooled trapped ions”. *Phys. Rev. Lett.* **77**, 4728–4731 (1996).
- [9] Susanne Pielawa, Giovanna Morigi, David Vitali, and Luiz Davidovich. “Generation of einstein-podolsky-rosen-entangled radiation through an atomic reservoir”. *Phys. Rev. Lett.* **98**, 240401 (2007).
- [10] S. Diehl, A. Micheli, A. Kantian, B. Kraus, H. P. Büchler, and P. Zoller. “Quantum states and phases in driven open quantum systems with cold atoms”. *Nature Physics* **4**, 878–883 (2008).
- [11] Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. “Quantum computation and quantum-state engineering driven by dissipation”. *Nature Physics* **5**, 633–636 (2009).
- [12] SG Schirmer and Xiaoting Wang. “Stabilizing open quantum systems by markovian reservoir engineering”. *Physical Review A* **81**, 062306 (2010).
- [13] Giovanna Morigi, Jürgen Eschner, Cecilia Cormick, Yiheng Lin, Dietrich Leibfried, and David J. Wineland. “Dissipative quantum control of a spin chain”. *Phys. Rev. Lett.* **115**, 200502 (2015).
- [14] Leo Zhou, Soonwon Choi, and Mikhail D Lukin. “Symmetry-protected dissipative preparation of matrix product states”. *Physical Review A* **104**, 032418 (2021). url: <https://doi.org/10.1103/physreva.104.032418>.
- [15] Felix Motzoi, Eli Halperin, Xiaoting Wang, K Birgitta Whaley, and Sophie Schirmer. “Backaction-driven, robust, steady-state long-distance qubit entanglement over lossy channels”. *Physical Review A* **94**, 032313 (2016). url: <https://doi.org/10.1103/physreva.94.032313>.
- [16] Kevin C. Smith, Eleanor Crane, Nathan Wiebe, and S.M. Girvin. “Deterministic constant-depth preparation of the aklt state on a quantum processor using fusion measurements”. *PRX Quantum* **4** (2023).
- [17] Nathanan Tantivasadakarn, Ryan Thorngren, Ashvin Vishwanath, and Ruben Verresen. “Long-range entanglement from

measuring symmetry-protected topological phases” (2021). url: <https://arxiv.org/abs/2112.01519>.

[18] Clément Sayrin, Igor Dotsenko, Xingxing Zhou, Bruno Peaudecerf, Théo Rybarczyk, Sébastien Gleyzes, Pierre Rouchon, Mazyar Mirrahimi, Hadis Amini, Michel Brune, et al. “Real-time quantum feedback prepares and stabilizes photon number states”. *Nature* **477**, 73–77 (2011). url: <https://doi.org/10.1038/nature10376>.

[19] R Vijay, Chris Macklin, DH Slichter, SJ Weber, KW Murch, Ravi Naik, Alexander N Korotkov, and Irfan Siddiqi. “Stabilizing rabi oscillations in a superconducting qubit using quantum feedback”. *Nature* **490**, 77–80 (2012). url: <https://doi.org/10.1038/nature11505>.

[20] D Riste, M Dukalski, CA Watson, G De Lange, MJ Tiggelman, Ya M Blanter, Konrad W Lehnert, RN Schouten, and L DiCarlo. “Deterministic entanglement of superconducting qubits by parity measurement and feedback”. *Nature* **502**, 350–354 (2013). url: <https://doi.org/10.1038/nature12513>.

[21] Hideo Mabuchi. “Continuous quantum error correction as classical hybrid control”. *New Journal of Physics* **11**, 105044 (2009). url: <https://doi.org/10.1088/1367-2630/11/10/105044>.

[22] Joseph Kerckhoff, Hendra I Nurdin, Dmitri S Pavlichin, and Hideo Mabuchi. “Designing quantum memories with embedded control: photonic circuits for autonomous quantum error correction”. *Physical Review Letters* **105**, 040502 (2010). url: <https://doi.org/10.1103/physrevlett.105.040502>.

[23] Leigh Martin, Felix Motzoi, Hanhan Li, Mohan Sarovar, and K Birgitta Whaley. “Deterministic generation of remote entanglement with active quantum feedback”. *Physical Review A* **92**, 062321 (2015). url: <https://doi.org/10.1103/physreva.92.062321>.

[24] Google Quantum AI. “Suppressing quantum errors by scaling a surface code logical qubit”. *Nature* **614**, 676–681 (2023). url: <https://www.nature.com/articles/s41586-022-05434-1>.

[25] Daniel Burgarth and Vittorio Giovannetti. “Mediated homogenization”. *Phys. Rev. A* **76**, 062307 (2007).

[26] Daniel Burgarth and Vittorio Giovannetti. “Full control by locally induced relaxation”. *Phys. Rev. Lett.* **99**, 100501 (2007).

[27] Anne Mattheis, Mark Rudner, Achim Rosch, and Erez Berg. “Programmable adiabatic demagnetization for systems with trivial and topological excitations” (2022). url: <https://arxiv.org/abs/2210.17256>.

[28] Sthitadhi Roy, JT Chalker, IV Gornyi, and Yuval Gefen. “Measurement-induced steering of quantum systems”. *Physical Review Research* **2**, 033347 (2020). url: <https://doi.org/10.1103/physrevresearch.2.033347>.

[29] Christopher Moore and Martin Nilsson. “Parallel quantum computation and quantum codes”. *SIAM journal on computing* **31**, 799–815 (2001). url: <https://doi.org/10.1137/s0097539799355053>.

[30] Rodney Van Meter and Kohei M Itoh. “Fast quantum modular exponentiation”. *Physical Review A* **71**, 052320 (2005). url: <https://doi.org/10.1103/physreva.71.052320>.

[31] Bhaskar Gaur, Edgard Muñoz-Coreas, and Himanshu Thapliyal. “A logarithmic depth quantum carry-lookahead modulo (2n - 1) adder”. In *Proceedings of the Great Lakes Symposium on VLSI 2023*. Pages 125–130. (2023).

[32] Kurt Jacobs, Xiaoting Wang, and Howard M Wiseman. “Coherent feedback that beats all measurement-based feedback protocols”. *New Journal of Physics* **16**, 073036 (2014).

[33] Ángel Rivas, Susana F Huelga, and Martin B Plenio. “Entanglement and non-markovianity of quantum evolutions”. *Physical review letters* **105**, 050403 (2010). url: <https://doi.org/10.1103/physrevlett.105.050403>.

[34] Ruben Verresen, Roderich Moessner, and Frank Pollmann. “One-dimensional symmetry protected topological phases and their transitions”. *Physical Review B* **96**, 165124 (2017). url: <https://doi.org/10.1103/physrevb.96.165124>.

[35] Frank Pollmann and Ari M Turner. “Detection of symmetry-protected topological phases in one dimension”. *Physical review b*

86, 125441 (2012). url: <https://doi.org/10.1103/physrevb.86.125441>.

[36] Gavin K Brennen and Akimasa Miyake. “Measurement-based quantum computer in the gapped ground state of a two-body hamiltonian”. *Physical review letters* **101**, 010502 (2008). url: <https://doi.org/10.1103/physrevlett.101.010502>.

[37] P. Filipowicz, J. Javanainen, and P. Meystre. “Theory of a microscopic maser”. *Phys. Rev. A* **34**, 3077–3087 (1986).

[38] John J. Slosser and Pierre Meystre. “Tangent and cotangent states of the electromagnetic field”. *Phys. Rev. A* **41**, 3867–3874 (1990).

[39] Hans-Jürgen Briegel and Berthold-Georg Englert. “Macroscopic dynamics of a maser with non-poissonian injection statistics”. *Phys. Rev. A* **52**, 2361–2375 (1995).

[40] Thomas Wellens, Andreas Buchleitner, Burkhard Küpper, and Hans Maassen. “Quantum state preparation via asymptotic completeness”. *Phys. Rev. Lett.* **85**, 3361–3364 (2000).

[41] Susanne Pielawa, Luiz Davidovich, David Vitali, and Giovanna Morigi. “Engineering atomic quantum reservoirs for photons”. *Phys. Rev. A* **81**, 043802 (2010).

[42] M Hartmann, D Poletti, M Ivanchenko, S Denisov, and P Hänggi. “Asymptotic floquet states of open quantum systems: the role of interaction”. *New Journal of Physics* **19**, 083011 (2017).

[43] M. Weidinger, B. T. H. Varcoe, R. Heerlein, and H. Walther. “Trapping states in the micromaser”. *Phys. Rev. Lett.* **82**, 3795–3798 (1999).

[44] B. T. H. Varcoe, S. Brattke, M. Weidinger, and H. Walther. “Preparing pure photon number states of the radiation field”. *Nature* **403**, 743–746 (2000).

[45] G. Morigi, J. I. Cirac, M. Lewenstein, and P. Zoller. “Ground-state laser cooling beyond the lamb-dicke limit”. *Europhysics Letters* **39**, 13 (1997).

[46] G. Morigi, J. I. Cirac, K. Ellinger, and P. Zoller. “Laser cooling of trapped atoms to the ground state: A dark state in position space”. *Phys. Rev. A* **57**, 2909–2914 (1998).

[47] Jean Dalibard, Yvan Castin, and Klaus Mølmer. “Wave-function approach to dissipative processes in quantum optics”. *Phys. Rev. Lett.* **68**, 580–583 (1992).

[48] R. Dum, P. Zoller, and H. Ritsch. “Monte carlo simulation of the atomic master equation for spontaneous emission”. *Phys. Rev. A* **45**, 4879–4887 (1992).

[49] T. S. Cubitt, F. Verstraete, W. Dür, and J. I. Cirac. “Separable states can be used to distribute entanglement”. *Phys. Rev. Lett.* **91**, 037902 (2003).

[50] Édgar Roldán and Shamik Gupta. “Path-integral formalism for stochastic resetting: Exactly solved examples and shortcuts to confinement”. *Phys. Rev. E* **96**, 022130 (2017).

[51] B. Mukherjee, K. Sengupta, and Satya N. Majumdar. “Quantum dynamics with stochastic reset”. *Phys. Rev. B* **98**, 104309 (2018).

[52] R. Yin and E. Barkai. “Restart expedites quantum walk hitting times”. *Phys. Rev. Lett.* **130**, 050802 (2023).

[53] Jutho Haegeman, J Ignacio Cirac, Tobias J Osborne, Iztok Pižorn, Henri Verschelde, and Frank Verstraete. “Time-dependent variational principle for quantum lattices”. *Physical review letters* **107**, 070601 (2011). url: https://doi.org/10.1007/3-540-10579-4_20.

[54] Andrew J. Daley. “Quantum trajectories and open many-body quantum systems”. *Advances in Physics* **63**, 77–149 (2014).

[55] Jülich Supercomputing Centre. “Jureca: Data centric and booster modules implementing the modular supercomputing architecture at jülich supercomputing centre”. *Journal of large-scale research facilities* **7**, A182 (2021).

[56] Artur Garcia-Saez, Valentin Murg, and Tzu-Chieh Wei. “Spectral gaps of affleck-kennedy-lieb-tasaki hamiltonians using tensor network methods”. *Physical Review B* **88**, 245118 (2013). url: <https://doi.org/10.1103/physrevb.88.245118>.

A Imaginary time evolution

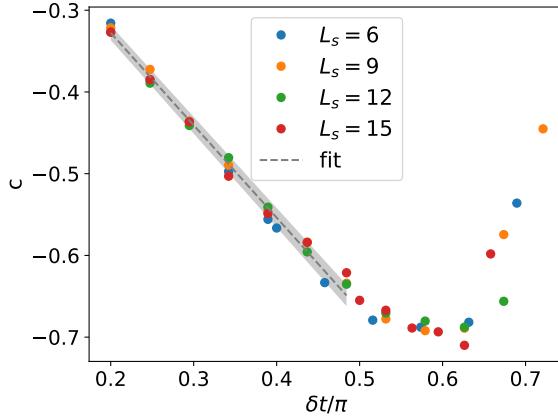


Figure 8: The figure shows the exponential coefficient c of the energy evolution $E(t) = e^{ct+a}$ for different values of δt after the system has reached the convergence time t_{conv} . A linear fit of $c(\delta t)$ for $L_s = 15$ gives a slope of $\widetilde{\Delta E} = -0.351 \pm 0.009$, which is consistent with the theoretical value for the energy gap.

This section presents an analysis of the dynamics' evolution beyond t_{conv} , which represents the time when the ancillas no longer flip into an excited state. The lack of ancilla flipping effectively causes only the imaginary time evolution term in Eq. 15 to act on the state. If this assumption holds true, the energy gap of the AKLT Hamiltonian can be readily obtained. By comparing this energy gap to values reported in the literature, we can confirm the accuracy of our assumption regarding the protocol reducing to imaginary time evolution.

At t_{conv} , the state closely approximates the AKLT state, allowing us to assume a low-temperature thermal state:

$$\rho = \frac{|E_0\rangle\langle E_0|}{1 + e^{-\Delta E \beta_0}} + |E_1\rangle\langle E_1| \frac{e^{-\Delta E \beta_0}}{1 + e^{-\Delta E \beta_0}} \quad (25)$$

where $|E_1\rangle$ denotes the first excited state, ΔE represents the energy gap, and β_0 is the inverse temperature. Evolving the state with the imaginary time evolution term in Eq. 15 results in:

$$e^{-\frac{t}{2}H_{\text{AKLT}}\delta t} \rho e^{-\frac{t}{2}H_{\text{AKLT}}\delta t} = \frac{|E_0\rangle\langle E_0|}{1 + e^{-\Delta E \beta_0}} + \frac{|E_1\rangle\langle E_1|}{1 + e^{\Delta E(\beta_0 + \delta t \cdot t)}} \quad (26)$$

$$E(t) = \frac{\Delta E}{1 + e^{\Delta E(\beta_0 + \delta t \cdot t)}} \quad (27)$$

Consequently, if β_0 is sufficiently large, the energy should scale as $E(t) \propto e^{-\Delta E \delta t \cdot t}$. To compare this relationship with the simulation data, the exponential factor $c(\delta t)$ can be calculated by fitting the exponential $E(t) = e^{ct+a}$ to the energy of the trajectories after t_{conv} . This fitting produces the data shown in Fig. 8, which displays a linear decrease until it starts to flatten near $\delta t = \frac{\pi}{2}$. The linear part was fit with $c(\delta t) \approx -\delta t \widetilde{\Delta E} + b$, yielding $\widetilde{\Delta E} = 0.351 \pm 0.009$. This value is consistent with the energy gap of the AKLT Hamiltonian, $\Delta E = 0.3501$ [56], thus corroborating our comprehension of the dynamics.

B Entanglement extrapolation

In this section, we introduce a novel technique to extrapolate the entanglement of a Matrix Product State (MPS). This method is crucial for determining the validity duration of an MPS simulation and for addressing the question of whether an entanglement phase transition can be observed.

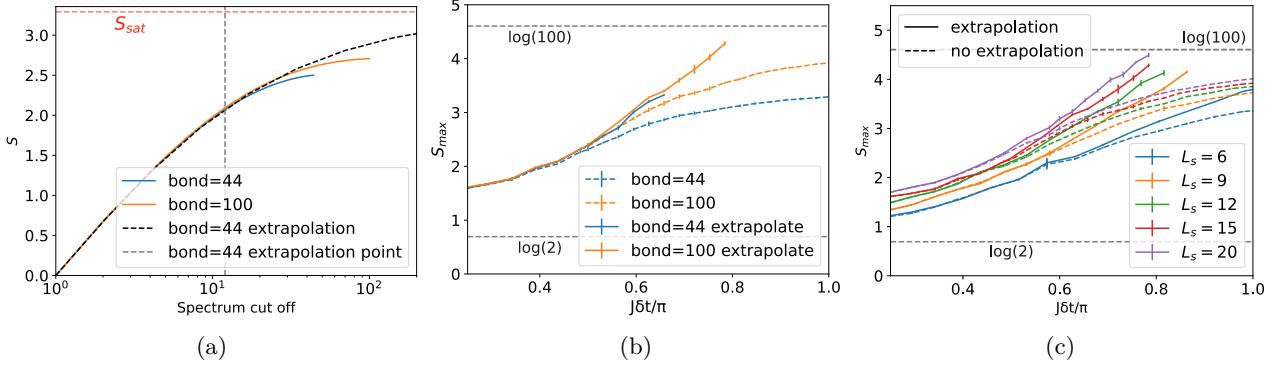


Figure 9: (a) The entanglement is extrapolated from a bond dimension of 44 using Eq. 29. (b) The maximal entanglement achieved in a simulation for a system size of $L_s = 15$ is shown. By comparing extrapolated and non-extrapolated results, we can see that the extrapolated values at a bond dimension of 44 match the more expensive simulation at a bond dimension of 100. (c) The same plot as in (b) is shown, but this time comparing different system sizes L_s .

The entanglement of a quantum system can be computed by bisecting the system and performing a Singular Value Decomposition (SVD) $|\psi\rangle = \sum_i^{2^{L/2}} \lambda_i |\psi_L^i\rangle \otimes |\psi_R^i\rangle$. Once the singular values are obtained, the entanglement entropy can be calculated using the expression $S = \sum_i^{2^{L/2}} \lambda_i^2 \log(\lambda_i^2)$. The tensor network technique involves retaining only the largest λ_i , so that the new entanglement is

$$p_i = \frac{\lambda_i^2}{\sum_i^{\text{cut}} \lambda_i^2}$$

$$S(\text{cut}) = \sum_i^{\text{cut}} p_i \log(p_i). \quad (28)$$

If the discarded λ_i are excessively large, the simulation is no longer valid.

Our approach approximates the magnitude of the discarded eigenvalues through a four-step process. First, simulations are conducted at both small and large bond dimensions. Second, the entanglement is computed for various cutting points at both bond dimensions. Third, the saturation entanglement entropy is extrapolated by fitting $S(\text{cut})$ with $\text{cut} < 0.9 \text{bond}_{\text{small}}$:

$$S(\text{cut}) = S_{\text{sat}} \tanh \left(\frac{\log(\text{cut}) \sigma_1}{S_{\text{sat}}} + \frac{\log(\text{cut})^2 \sigma_2}{S_{\text{sat}}} + \frac{\sqrt{\log(\text{cut})} \sigma_3}{S_{\text{sat}}} \right). \quad (29)$$

Lastly, the resulting saturation points S_{sat} are compared for both bond dimensions. If they match, the extrapolation is deemed valid, and a new extrapolation is performed with $S(\text{cut})$ with $\text{cut} < 0.9 \text{bond}_{\text{large}}$.

It is important to note that the magnitude of the σ_1 term indicates the number of eigenvalues needed to reach the saturation point, while the terms σ_2 and σ_3 slightly adjust the curvature of the saturation curve and are nearly zero for most extrapolations. Adding more terms to the expansion does not affect the saturation point S_{sat} . An example of this procedure is illustrated in Fig. 9(a).

In Fig. 9(b), the maximal entanglement achieved in the simulation S_{max} is presented for bond dimensions of 44 and 100, with extrapolation both enabled and disabled. It is evident that although the non-extrapolated values diverge, the extrapolated ones yield consistent results. Given that conducting a simulation at bond dimension 100 is eight times more resource-intensive than at bond dimension 44, these findings are promising.

Lastly, Fig. 9(c) displays the maximal entropy observed in the simulation for various reset intervals δt across different lattice sizes. The values of S_{max} without employing the extrapolation technique

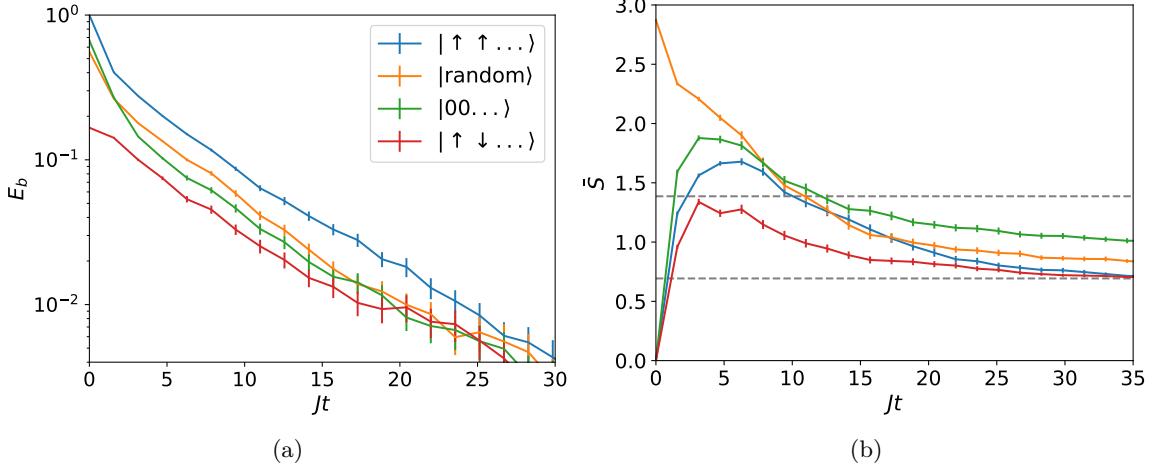


Figure 10: The figure shows the measurement-induced steering of three different product states (blue: all spins pointing upwards, green: all spins in the zero component, and red: alternating spin up and down), as well as a random initial MPS state with a bond dimension of 44 (yellow). Panel (a) shows that states with lower initial energy converge faster, as indicated by the smoothed energy per bond E_b . Panel (b) reveals that the average bipartite entanglement entropy increases and then decreases for the product states but steadily decreases for the random states, resulting in entropies between $\log(2)$ and $\log(4)$. This indicates a different mixture of the four AKLT states based on the initial state.

begin to flatten as they approach the maximal entropy that the MPS can represent $\log(100)$. In contrast, the extrapolated values increase linearly until their validity can no longer be verified. This demonstrates the importance of applying the extrapolation technique to ensure the reliability of the results.

C Different initial states

In this appendix, we explore the robustness of our results by considering four different initial states. These states consist of: all spins pointing up (blue); all spins in the zero component (green); alternating spin up and down (red); and a random initial MPS state with a bond dimension of 44 (yellow). As seen in Fig. 8(a), states with lower initial energy converge faster, as expected since they are closer to the ground state. Fig. 8(b) shows that the average bipartite entanglement entropy increases and then decreases for the product states, but steadily decreases for the random states. This results in entropies between $\log(2)$ and $\log(4)$, indicating a different mixture of the four AKLT states depending on the initial state. We attribute this to the fact that the product state with all spin components in the zero state is closer to the AKLT state in which the two edge modes are entangled.

PUBLICATION II

Learning Feedback Mechanisms for Measurement-Based Variational Quantum State Preparation

Authors: Daniel Alcalde Puente, Matteo Rizzi
Year: 2024
Journal: Quantum
Volume/Issue: 9, 1792
DOI: [10.22331/q-2025-07-11-1792](https://doi.org/10.22331/q-2025-07-11-1792)

Summary

The efficient preparation of quantum states is a cornerstone for advancing quantum technologies. Standard variational quantum circuits (VQCs), discussed in Chap. 5, offer a learnable approach but frequently require circuit depths that scale unfavorably with system size. As an alternative, measurement-based protocols incorporating adaptive feedback (Sec. 4.3) have emerged as promising approaches, potentially achieving constant-depth preparation. However, the design of such adaptive protocols remains a nontrivial task, and existing autonomous learning methods, including reinforcement learning and greedy optimization, face significant scalability and optimality limitations (see Sec. 4.4).

In Publication [II], a method that autonomously discovers efficient state preparation strategies that exploit non-unitary quantum dynamics, surpassing purely unitary approaches (Sec. 4.1), is presented.

The key contributions and findings of this work are as follows:

A novel measurement-based VQC architecture is developed, incorporating intermediate projective measurements and classical feedback. This approach enables the simultaneous learning of both the optimal unitary operations and the feedback strategies required to steer the quantum state towards a target state.

A new class of local minima specific to measurement-based variational circuits is identified. Mitigation strategies are proposed, including the use of different update frequencies for pre-measurement and feedback unitaries, as well as ancilla-based regularization techniques. These strategies partially alleviate the newly identified optimization challenges and build upon the broader discussion of local minima presented in Sec. 5.1.2.

The self-learning protocol is showcased by preparing the spin-1 Affleck-Kennedy-Lieb-Tasaki (AKLT) state, a canonical example of a symmetry-protected topological (SPT) phase (see Sec. 2.1.4). Comparison with known analytical measurement-based protocols, such as that proposed by Smith et al. [27], reveals that the learned approach can identify comparable solutions, often utilizing slightly shallower pre-measurement circuits.

Scalability is further addressed by extending the framework to larger system sizes through translationally invariant VQC ansätze and the use of recurrent neural networks (RNNs) to parameterize the adaptive feedback mechanism. This demonstrates a viable pathway for addressing increasingly large quantum systems.

Notably, the method is capable of deterministically preparing a specific AKLT edge state (with both edge spins up), for which no analytical deterministic low-depth protocol was previously known. This result highlights the framework’s potential to autonomously discover new quantum algorithms beyond existing analytical methods.

While the protocol is capable of representing efficient quantum state preparation strategies, the abundance of local minima presents a significant challenge to training. This difficulty is particularly pronounced in shallow circuits, which are desirable for implementation on near-term devices as they are less prone to noise, but are especially susceptible to underparameterization and poor optimization landscapes (see Sec. 5.1.2). Consequently, while shallow circuits are an essential objective, their vulnerability to local minima suggests they must be approached with caution.

Numerical simulations of quantum dynamics and gradient computations for the VQCs are performed using matrix product state (MPS) techniques (Sec. 2.1), with gradient evaluation strategies detailed in Sec. 5.2.3. These simulations are implemented using the custom-developed software package `mVQE.jl` [46].

This publication complements the dissipative state preparation approach explored in Publication [I] (Sec. 4.2), emphasizing adaptive measurement and feedback over purely dissipative dynamics. It provides a robust framework for the design and optimization

of non-unitary quantum state preparation protocols, advancing the feasibility of implementing such methods on near-term quantum hardware. Future research may extend this learning-based framework to two-dimensional quantum systems, potentially employing finite PEPS representations (Sec. 2.2) as developed in Publication [III]. This would be particularly relevant for practical applications, as many quantum states of interest arise in two-dimensional systems.

In particular, the extension to topological states such as the toric code presents a compelling direction. These states are essential in the context of fault-tolerant quantum computation due to their nonlocal encoding of quantum information. Adapting the feedback-based learning framework to prepare such states could allow not only for their accurate preparation but also for the deterministic selection of the encoded logical qubit state. This would parallel the results achieved for the AKLT edge mode and could overcome limitations of passive preparation protocols, which lack control over logical state selection. However, due to the inherent long-range entanglement of topological states, a speedup over unitary-only methods might not be feasible.

Another difficulty might be that to optimize the circuit efficiently in this setting, it may be necessary to relax the constraint on circuit shallowness, as overly shallow ansätze are especially prone to local minima that hinder convergence.

Author’s Contribution

Daniel Alcalde Puente led the conceptual development and implementation of all algorithms. He carried out all numerical simulations and analyses. He wrote the first draft of the manuscript. Throughout the project, he received guidance and regularly discussed key ideas and results with his supervisor.

Data Availability

All simulations were performed using the MPS formalism to represent the wave function; the code and resulting data are available on Zenodo [139]. The Julia library `mVQE.jl` [46], developed specifically for these simulations, uses the `ITensors.jl` [124] library as its backbone.

Learning Feedback Mechanisms for Measurement-Based Variational Quantum State Preparation

Daniel Alcalde Puente  ^{1,2} and Matteo Rizzi  ^{1,2}

¹Forschungszentrum Jülich, Institute of Quantum Control, Peter Grünberg Institut (PGI-8), 52425 Jülich, Germany

²Institute for Theoretical Physics, University of Cologne, 50937 Köln, Germany

This work introduces a self-learning protocol that incorporates measurement and feedback into variational quantum circuits for efficient quantum state preparation. By combining projective measurements with conditional feedback, the protocol learns state preparation strategies that extend beyond unitary-only methods, leveraging measurement-based shortcuts to reduce circuit depth. Using the spin-1 Affleck-Kennedy-Lieb-Tasaki state as a benchmark, the protocol learns high-fidelity state preparation by overcoming a family of measurement-induced local minima through adjustments of parameter update frequencies and ancilla regularization. Despite these efforts, optimization remains challenging due to the strongly non-convex optimization landscapes inherent to variational circuits. The approach is extended to larger systems using translationally invariant ansätze and recurrent neural networks for feedback, demonstrating scalability. Additionally, the successful preparation of a specific AKLT state with desired edge modes highlights the potential to discover new state preparation protocols where none currently exist. These results indicate that integrating measurement and feedback into variational quantum algorithms provides a promising framework for quantum state preparation.

1 Introduction

Quantum technologies have significant potential to address key challenges in quantum simulation, communication, and information processing. As such, the efficient preparation of high-

fidelity quantum states and the creation of robust state preparation protocols are critical, particularly for noisy intermediate-scale quantum (NISQ) devices [1]. Additionally, quantum measurements are becoming an integral part of many new circuit paradigms [2, 3, 4, 5], emphasizing the importance of incorporating measurement and feedback mechanisms into quantum state preparation. In this work, we explore the autonomous learning of variational quantum circuits with measurement and feedback for quantum state preparation.

Variational quantum circuits (VQCs) [6] hold promise for a wide range of quantum computing applications. A parameterized quantum circuit is constructed from gates parameterized by their angles. Its parameters are optimized to minimize a cost function, often infidelity, in the context of state preparation. Preparing one-dimensional long-range entangled quantum states using local two-qubit variational quantum circuits (VQCs) typically requires a circuit depth that scales linearly with the system size.

However, measurement-based quantum circuits can be analytically derived for constant-depth state preparation by exploiting the non-unitary character of quantum measurements and the use of adaptive feedback. These methods have been applied to a broad class of target states, including toric code states, symmetry-protected topological (SPT) phases, fracton phases, and Schrödinger's cat states [7, 8, 9, 10], as well as certain non-Abelian topological orders [11]. Adaptive measurement strategies further extend these capabilities by enabling the fusion of small resource states into complex many-body targets. This approach has led to deterministic, constant-depth protocols that have been demonstrated experimentally, for instance, in the preparation of the AKLT

state [12]. The feasibility of such fusion-based schemes depends on the entanglement structure of the target state. Theoretical results identify features such as flat entanglement spectra in SPT phases as crucial for exact preparation [13, 14, 15], and also extend to symmetry-broken and topologically ordered phases [16]. Allowing for approximate rather than exact preparation further reduces circuit complexity. For example, translationally invariant matrix product states can be prepared approximately with depth scaling as $\mathcal{O}(\log\log(N/\epsilon))$ [17], where ϵ is the allowed error. In addition, probabilistic schemes permit constant-depth preparation of states such as Dicke and W-states by accepting non-deterministic outcomes [18].

Incorporating these ideas, this work extends VQCs to include non-unitary measurements and feedback, enabling the protocol to learn fundamentally different state preparation techniques than standard VQCs and potentially reducing the required circuit depth. These measurement-based protocols exploit the non-unitary nature of quantum measurements to construct non-unitary gates. However, one cannot apply non-unitary gates directly; rather, one must design a set of non-unitary gates that, when combined, form a completely positive trace-preserving (CPTP) map from which one gate is randomly applied. The challenge is to design a CPTP map that, through the application of conditional feedback, ensures that the desired quantum state is achieved regardless of the specific non-unitary gate that is randomly selected by quantum mechanics. In this work, this task is learned automatically using VQCs.

Despite their potential, VQCs face significant challenges that hinder their practical implementation. One prominent problem is the occurrence of barren plateaus in the optimization landscape of deep circuits, where the gradients of the parameters approach zero, rendering the optimization process infeasible. While measurement-based methods at shorter circuit depths may avoid these plateaus, they introduce new optimization challenges. Shallow quantum circuits, while less prone to barren plateaus, can be difficult to optimize due to the presence of local minima [19, 20]. Local minima can be partially mitigated in specific cases, such as the AKLT model, by incorporating symmetry constraints into the

variational ansatz [21]. However, there is no general strategy for addressing this issue across different models. In this work, additional local minima that are unique to measurement-based VQCs are discussed and mitigation strategies are discussed.

Autonomous learning of measurement and feedback protocols for state preparation have been studied in the literature with two different methods.

In the first method, reinforcement learning is used to develop a complete policy that determines dynamically both when measurements should be performed and which unitary operations should be applied based on those measurements. This approach shows success in the single-particle case [22, 23] and for two-particle systems [24, 25]. However, we note that given the evidence presented in the papers we expect these methods to have significant challenges when extended to multi-qubit systems.

In the second method, greedy optimization techniques have been proposed [26, 27] to prepare multi-qubit states. These methods periodically optimize unitaries to maximize fidelity after each weak measurement. More precisely some ancillas A are introduced and coupled to the system with unitaries and the measured projectively. Each of these unitaries is learned independently like: $|\psi'_i(\theta_i, \dots)\rangle = U(\theta_i)|\psi_{i-1}(\theta_{i-1}, \dots)\rangle$ with $\theta_i = \min_{\theta_i} \text{loss}(|\psi'_i(\theta_i, \dots)\rangle)$, and $|\psi_i(\theta_i, \dots)\rangle = \langle M|_A |\psi'_i(\theta_i, \dots)\rangle$ sampled from $M \in \|\langle M|_A |\psi'_i(\theta_i, \dots)\rangle\|^2$. This makes them suffer from two major shortcomings. First, implementing these methods experimentally requires running a simulated version of the experiment in parallel to optimize the unitary operations on the fly depending on the measurement results obtained in the experiment. Second, the cost function used in these approaches is inherently greedy, reacting to measurement results rather than proactively incorporating them into the optimization. Consequently, these limitations hinder the preparation of states with more than a few qubits and prevent the learning of feedback mechanisms such as those proposed by Smith *et al.* [12], where the pre-measurement unitary circuit is designed in such a way that after the measurement all possible resulting states can be mapped to the target state using two-body gates. For this, all unitaries must

be learned at the same time for all possible measurement outcomes, which is done in our work.

Our learning technique can handle larger systems than the aforementioned protocols using reinforcement learning because it has a fixed structure (VQCs), simplifying the learning problem. However, since our protocol is non-greedy and learns the complete feedback step by learning all unitaries at the same time, it can effectively learn which CPTP map and corresponding feedback to apply automatically, and as such can learn more intricate protocols like the one by Smith *et al.* [12].

Mixing variational circuits and measurement has also been discussed in the realm of Measurement-based quantum computing (MBQC), where VQE can be used to decide which basis to measure with during the circuit [28].

Similar concepts have been explored in the study of error-correcting codes, where reinforcement learning algorithms attempt to automate the discovery and implementation of error-correcting protocols. The complexity of this task forces its decomposition into subtasks that are handled by different reinforcement learning agents, such as the decoding of known stabilizer codes [29, 30, 31], or the spatial deformation of codes for better logical error rates [32]. Although learning the complete error-correcting task simultaneously would be ideal, it remains highly challenging and is only accomplished with strategies that do not scale with system size in [33] for a small system of four qubits. This highlights the difficulty of learning full measurement and feedback-based protocols, which our work addresses in the context of state preparation.

During the final preparation of this manuscript, various manuscripts were published using a similar framework. In [34], a similar approach is employed but without measurement feedback, relying solely on projective measurements as a non-unitary operation. In [35, 36], the protocols largely mirror those presented here, with minor variations, but applied to the GHZ state. In Yan *et al.* [35], the authors propose a method to experimentally estimate the gradients of the protocol's parameters on a Quantum Computer. In Alam *et al.* [36], a Density Matrix Renormalization Group

(DMRG)-inspired sweeping optimization is introduced to avoid measurement-induced local minima in the optimization landscape. For completeness, we have added a short analysis of the GHZ state using our methodology in A.

However, we expect this to worsen the above-mentioned vanilla-VQC local minima [19, 20], which should still be present in addition to measurement-induced ones. Our internal experiments on vanilla-VQCs (inspired by Pollmann *et al.* [37]) indicate indeed that sweeping does not scale well with system size and tends to get trapped in local minima for larger systems. Therefore, the application of a pure sweeping approach to larger measurement-based VQCs remains unclear to us and constitutes an interesting direction for future research. Consequently, we did not pursue this approach in our work, although a hybrid approach could be beneficial. Instead, we found an explanation for why the measurement-induced local minima occur and addressed them directly for a system size of 16 qubits.

This paper is organized as follows: In Sec. 2, we present our self-learning protocol that integrates measurement and feedback into VQCs for efficient quantum state preparation. The protocol employs a sequence of parameterized unitaries and projective measurements, using feedback from measurement outcomes to inform subsequent operations. This approach allows the protocol to learn non-unitary state preparation techniques that can reduce the circuit depth. In Sec. 3 and 4, we apply our protocol to the preparation of the spin-1 AKLT state, using it as a benchmark to evaluate the learning capabilities of our approach, which if optimized naively is plagued by local minima. We propose two strategies to mitigate these: adjusting the parameter update frequencies between the initial unitary and feedback operations (detailed in Sec. 4.1), and introducing ancilla regularization to promote a more uniform distribution of measurement outcomes (discussed in Sec. 4.2). In Sec. 5, we compare the learned protocol to the analytically derived protocol proposed by Smith *et al.* [12]. We analyze the similarities and differences using quantum mutual information, highlighting how the learned protocol can achieve similar or better performance with potentially shallower circuits by requiring less mutual information prior

to measurement. In Sec. 6, we extend our approach to larger systems by employing a translationally invariant ansatz and utilizing a Recurrent Neural Network (RNN) for the feedback function. We demonstrate our protocol’s scalability and discuss the RNN’s performance, but note challenges in performing optimal corrections for large systems. In Sec. 7, we explore the preparation of a specific AKLT state with both edge modes in the spin-up configuration—a task for which no known deterministic, low-depth protocol exists. We show that our learning protocol can discover such a state preparation strategy, highlighting its potential to find new protocols where none currently exist. In conclusion, this study shows the promise of integrating measurement and feedback into variational quantum algorithms to discover quantum state preparation algorithms.

2 Protocol Description

Our protocol consists of a sequence of parameterized unitary transformations and projective quantum measurements to prepare a desired target state $|\psi_t\rangle$. This process is iterative, using feedback from measurement outcomes to inform subsequent quantum operations. Its simplest form, with only one feedback round in the ansatz, is reminiscent of error correction and can represent state preparation protocols like [7, 11, 8, 12, 18, 16, 13, 15]. Conversely, if the feedback is removed completely (set it to a constant), then the operation $U(\theta)$ with the addition of measurement and reset, will be able to represent passive steering strategies like [38, 39, 40]. This emphasizes the representative capacity of this approach. The potential of this generalized, learnable protocol is twofold: first, to discover methods for preparing quantum states where no effective protocol currently is known; and second, to identify faster and more optimal approaches for state preparation where established protocols are already known.

2.1 Framework

Our protocol, shown in Fig. 1(b), is described by the following sequence of operations:

$$\rho_1(\theta_1) = U_1(\theta_1)\rho_0 U_1^\dagger(\theta_1), \quad (1)$$

$$\rho_1^M(\theta_1) = |0\rangle_A\langle M|\rho_1(\theta_1)|M\rangle_A\langle 0|, \quad (2)$$

$$\rho_2(\theta_1, W) = \sum_M U_2(\theta_2)\rho_1^M(\theta_1)U_2^\dagger(\theta_2), \quad (3)$$

$$\text{with } \theta_2 = f(M; W)$$

where ρ_0 is the initial state of the system, which we will choose to be $\rho_0 = |0\rangle_{SA}\langle 0|$, $U(\theta)$ is a unitary operation parameterized by θ , and M are the measurement outcomes. The operator $|0\rangle_A\langle M|$ projects the ancilla qubits to the measured state M which is then reset to a product state composed of zeros. The feedback function $f(M; W)$ with learnable parameters W is the key ingredient of the protocol as it decides on the parameters of subsequent unitary operations based on previous measurement outcomes.

For a more generalized approach that can be applied iteratively for multiple rounds of feedback, we extend the notation as follows:

$$\rho_i^M(\theta_1, W_2, \dots, W_i) = |0\rangle_A\langle M|\rho_i|M\rangle_A\langle 0| \quad (4)$$

$$\rho_i(\theta_1, W_2, \dots, W_i) = \sum_M U_i(\theta_i)\rho_{i-1}^M U_i^\dagger(\theta_i) \quad (5)$$

$$\text{with } \theta_i = f_i(M; W_i)$$

This iterative relation frames any general feedback protocol within our proposed scheme, where different functions $f_i(M; W_i)$ decide on subsequent unitary operations. Depending on the ansatz it is possible to use different parameterized unitaries U_i for different feedback sweeps.

The efficacy and applicability of the proposed protocol will first be evaluated in the context of the AKLT state preparation, demonstrating its ability to learn feedback strategies. Specifically, we aim to automatically reproduce the deterministic low-depth protocol outlined in [12], which will be further discussed in Sec. 3.

2.2 Feedback Mechanism

The success of our protocol is critically dependent on the function $f(M; W)$, which parametrizes the feedback loop. It maps measurement outcomes onto new parameters for unitary transformations. This function is vital, as it chooses subsequent

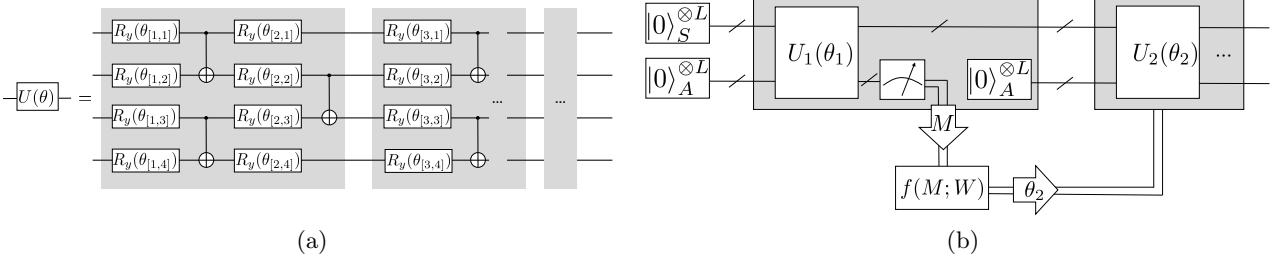


Figure 1: Illustration of the quantum feedback control protocol. (a) Depicts the hardware-efficient ansatz used in the construction of the unitary $U(\theta)$, showcasing a series of parameterized rotation gates $R_y(\theta_{i,j})$ arranged in alternating layers with CNOT gates. (b) Outlines the quantum-classical feedback loop, starting with the initial state preparation of the system S and ancilla A in $|0\rangle$ states. The application of $U(\theta_1)$ is followed by measurement, and the measurement results M are fed into a function that outputs the parameters $\theta_2 = f(M; W)$ for the next unitary operation $U(\theta_2)$. This loop implements the adaptive adjustment of parameters based on measurement results, which is central to the feedback control strategy. Note that the optimal parameters $\{\theta_1, W\}$ are learned using the gradient descent algorithm.

quantum gates depending on the measurement results.

For small ancilla qubit Hilbert spaces, $f(M; W)$ can be effectively represented using a tabular approach. In this representation, $f(M; W) = W_M$, with W_M being a set of learnable vectors, each of them storing the angles for the parameterized unitaries that should be applied in response to a particular measurement outcome M . This tabular method is straightforward and computationally manageable when the ancilla Hilbert space remains small.

However, as the ancilla Hilbert space expands, the tabular method becomes impractical because of the exponential increase in the number of potential measurement outcomes. For larger systems, a neural network offers a more sophisticated and scalable representation for $f(M; W)$. The chosen architecture for our implementation includes both SwiGLU (Swish-Gated Linear Unit) and bidirectional Recurrent Neural Network (RNN) layers, which were chosen as they are adept at modeling spatial correlations. These are used in Sec. 6 to obtain a translationally invariant feedback ansatz and their structure is explained in detail in App. C.

3 AKLT State(s)

The 1D spin-1 Affleck-Kennedy-Lieb-Tasaki (AKLT) states [41] provide a useful benchmark for evaluating the learning capabilities of our protocol. These states are well-known for their symmetry-protected topological (SPT) proper-

ties, characterized by two spin- $\frac{1}{2}$ edge states [42, 43].

The presence of these two free spin- $\frac{1}{2}$ edge states results in a four-fold degeneracy of the AKLT state. Furthermore, the AKLT state has been proposed as a valuable resource for measurement-based quantum computing [44]. If just local unitaries are employed, the preparation time scales linearly with the system size: by using also measurements, Smith *et al.* have shown that a protocol with depth independent from the system size is possible [12]. This consists of preparing small AKLT states and fusing them using measurement and deterministic feedback based on the measurement outcomes. This method is deeply rooted in the exact representation of AKLT states as matrix product states.

The analytically derived protocol in [12] can prepare a random state in the manifold spanned by the four AKLT states within a single round of measurement and feedback. Thus, our first task will be to try to learn an equivalent protocol without any prior knowledge of [12], to better understand the challenges faced by the learning algorithm.

Under the restriction of using a single round of measurements, the fidelity to the AKLT man-

ifold, spanned by $|\psi_i\rangle$, can be computed via:

$$F(\theta_1, W) = \sum_i^4 \langle \psi_i |_S \rho_2(\theta_1, W) |\psi_i\rangle_S \quad (6)$$

$$= \sum_i^4 \sum_M \|\langle \psi_i |_S U_2^{(S)}(\theta_2) \langle M |_A U_1^{(SA)}(\theta_1) | 0\rangle_{SA}\|^2 \quad \text{with } \theta_2 = f(M; W) \quad (7)$$

where $U_1^{(SA)}(\theta)$ represents a hardware-efficient ansatz (see Fig. 1(a)), and $U_2^{(S)}(\theta)$ is a similar ansatz acting only on the system qubits. Ideally, the feedback operation would use an ansatz identical to U_{SA}^1 ; however, implementing this approach with a short circuit depth after measurement introduces insurmountable local minima. This issue commonly arises in shallow VQCs when the initial wave function is not a product state, making optimization highly challenging. Therefore, instead of replicating $U_1^{(SA)}$, a simpler ansatz using only two-qubit gates is selected for the feedback step. Further details are provided in App. D.

For the sake of simplicity, we have first restricted ourselves to just employing tabular feedback $f(M; W) = W_M$ as it offers fewer possible points of failure than the use of a Neural Network, which we later analyze in Sec. 6. The objective is to optimize this protocol by minimizing the infidelity. We employ a gradient descent approach, specifically using the ADAM optimizer [45], for this purpose. However, a straightforward application of this optimization method leads to entrapment in local minima, as do other commonly used optimization methods like Natural Gradient, LBFGS, and gradient-free methods. These local minima are distinct both from the barren plateaus encountered in long and wide circuits and from local minima encountered in shallow circuits [19, 20]. To circumvent this class of local minima specific to feedback protocols, we implement two novel strategies as described in Sec. 4.1 and Sec. 4.2.

Note that in this work the spin-1 is mapped to qubits with the mapping $|+\rangle \rightarrow |10\rangle$, $|0\rangle \rightarrow |00\rangle$ and $|-\rangle \rightarrow |01\rangle$ as done also in [12], as it reduces entanglement between the two spin $\frac{1}{2}$, without affecting the properties of the final state. Other encoding choices are also possible and may impact the circuit's efficiency, as discussed in Ref. [46], where various mappings were analyzed in the context of Trotterization.

4 Local minima and how to avoid them

During the optimization process, it was observed that the probability distribution $P(M; \theta) = \text{Tr}_S [\langle M |_A \rho_i(\theta) | M \rangle_A]$ for measuring the ancillas in $|M\rangle_A$ in particular bit-string M in the z-base tends to favor a specific value of M' . This results in $P(M; \theta) = \delta(M, M')$ being a delta function, so that when the ancillas are measured, they always give the same measurement result M' . This results in the measurement operation having no effect on the quantum state, indicating that the optimization algorithm effectively found a way to bypass the measurement operation and as such the feedback altogether.

This can be quantified by the Shannon entropy H or the harder-to-compute entanglement entropy S , defined as:

$$\rho_1^A = \text{Tr}_S (U_1(\theta_1) | 0 \rangle_{SA} \langle 0 | U_1^\dagger(\theta_1)) \quad (8)$$

$$P(M) = \langle M | \rho_1^A | M \rangle \quad (9)$$

$$H = - \sum_M P(M) \log_2(P(M)) \quad (10)$$

$$S = - \text{Tr}_A [\rho_1^A \log(\rho_1^A)]. \quad (11)$$

Both of them turn out to be small, or even zero, at the local minima encountered along a naive minimization of the loss function in Eq. 6, as can be seen in the orange curve in Fig. 2(b). In the next two sections, two strategies to mitigate local minima are discussed.

4.1 Different update frequencies

We conjecture that the extreme sharpening of the measurement probability distribution $P(M)$ occurs due to an imbalance between the intrinsic learning rates of the initial unitary $U_1(\theta_1)$ and of the feedback $U_2(\theta_2)$, with the latter being too slow. The first attempt to improve the optimization involves therefore an ad-hoc increase of the update frequency of the W parameters that define the feedback.

Our reasoning is based on the following consideration: in order to achieve a pure state at the end of the protocol, either there is a single possible measurement outcome $P(M) = \delta_{M, M'}$ (and thus $H = 0$) or the feedback process has to distill the mixed state $\rho_1^M = \sum_M |\psi(M)\rangle \langle \psi(M)|$ (with $|\psi(M)\rangle_S = \langle M |_A U_1 | 0 \rangle_{SA}$) ($H \neq 0$) into a pure state through conditional feedback. It is much

easier to fulfill the first condition than the second one. As a consequence, if the feedback part of the protocol is updated too slowly, it may appear convenient for the optimizer to simply tune $U_1(\theta_1)$ to prepare the best possible state where the ancillas are in a product state and the action of the measurement is irrelevant. Otherwise stated, either all the information extracted from the system using the measurement needs to be effectively used to reduce the mixedness of the final state, or else the optimization algorithm will decide to reduce how much the system and ancillas are entangled. This would also favor H to settle mostly on integer values of H , as it is easier to fully use an ancilla or not use it at all. To address this issue, the rate of learning of the parameters of the feedback operation can be increased relative to that of the parameters of the initial unitary. This adjustment can be implemented in several ways: here, we choose to update the parameters of the feedback operation more frequently, as described in Alg. 1.

Algorithm 1 Optimization algorithm that updates the feedback parameters W more frequently than the initial unitary parameters θ_1 . This approach aims to help the feedback mechanism to adapt more rapidly, preventing the optimization process from ignoring the measurement and feedback steps that lead to local minima.

```

function UPDATE PARAMETERS( $\theta_1, W$ , loss)
     $\theta_1^{\text{new}} \leftarrow \text{ADAM\_update}(\theta_1, \nabla_{\theta_1} \text{loss}(\theta_1, W))$ 
     $W^{\text{new}} \leftarrow W$ 
    for 1 to update_freq do
         $W^{\text{new}} \leftarrow$ 
         $\text{ADAM\_update}(W^{\text{new}}, \nabla_W \text{loss}(\theta_1, W^{\text{new}}))$ 
    end for
    return  $\theta_1^{\text{new}}, W^{\text{new}}$ 
end function

```

To test the hypothesis, in Fig. 2 the feedback protocol is optimized for the AKLT state with 16 physical qubits and 8 ancilla qubits for different update frequencies of the feedback step. The system size 16 is chosen as it is large enough that every optimization run gets trapped in local minima, making this system size a compelling challenge for analysis.

If the update frequency of the feedback parameters is set to 0 (see blue line in Fig. 2), they are not updated at all. Consequently, the pre-

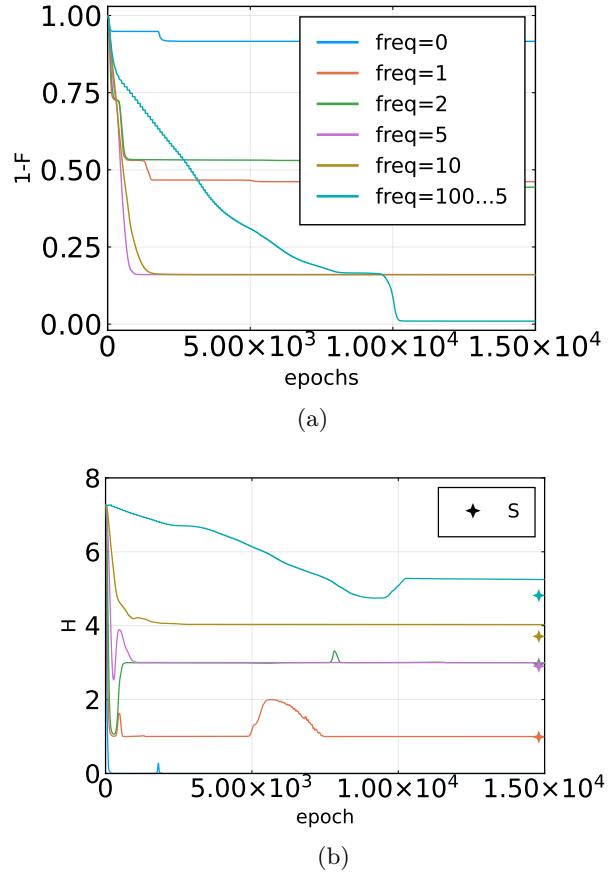


Figure 2: Influence of feedback update frequency on local minima encountered during the optimization of the feedback protocol to prepare any of the four AKLT states with a system size of 16 and 8 ancillas, configured in a repeating ASSSA pattern where A stands for ancilla and S for the system. Every epoch, the gradient is calculated using all possible measurement outcomes. The parameters of the feedback operation are updated more frequently using gradients. The label "freq=100..5" indicates that the update frequency decreases linearly from 100 to 5 over 10⁴ epochs. (a) Infidelity to the AKLT states: higher update frequencies prevent the protocol from getting trapped in local minima. (b) Shannon entropy H of the measurement probability distribution $P(M)$ and entanglement entropy between the ancilla and the system before measurement at the end of the optimization process (indicated by the star marker): higher update frequencies result in higher Shannon and entanglement entropies. Note that the entanglement entropy between the ancilla and the system was computed at the end of the optimization only once as it is expensive to calculate.

measurement unitaries are forced to prepare an ancilla state that, when measured, does not affect the final state, resulting in $H = 0$.

With the standard approach, updated frequency equal to one, the optimization gets

trapped in a local minimum with $H = 1$, corresponding to one bit of information contained in the ancilla distribution. The entanglement entropy S between system and ancilla at the end of the optimization is also one, confirming that one bit of information is extracted from the system with the measurement.

Higher update frequencies for the feedback parameters lead to improved infidelities and both higher Shannon entropy H and entanglement entropy S , which lie close to each other, showing that our initial conjecture seems to be correct. The low entropy is caused by the pre-measurement unitary trying to bypass the measurement step. And interestingly the Shannon entropy mostly settles to integer values, showing that as expected these are attractors of the optimization dynamics.

The most effective protocol with this strategy is achieved by initially setting the update frequency to 100 and then linearly decreasing it to 5 over the first 10^4 epochs achieving the highest value of H and showing that high update frequencies are mostly important at the beginning of the optimization process. The optimal update frequency should be chosen such that the feedback operation parameters are close to $\min_W \text{loss}(\theta_1, W)$.

4.2 Ancilla Regularization

Another approach to prevent the measurement distribution from becoming trapped in local minima, characterized by $H = 0$, involves encouraging broader exploration of the solution space. This is achieved by regularizing the distribution $P(M)$, by adding a term to the loss function that encourages a more uniform probability distribution. The regularization term $l_R(M; \theta)$ is defined as follows:

$$d(M; \theta) = 1 + \frac{\log_2(P(M; \theta))}{N_a} \quad (12)$$

$$l_R(M; \theta) = \begin{cases} 0 & \text{if } -c < d < c \\ (d(M; \theta) + c)^2 & \text{if } d < -c \\ (d(M; \theta) - c)^2 & \text{if } c < d \end{cases} \quad (13)$$

$$l_R(\theta) = \frac{1}{2^{N_a}} \sum_M l_R(M; \theta) \quad (14)$$

where the distance $d(M)$ tells us how far the measurement sample M is from having the proba-

bility 2^{-N_a} where N_a is the number of ancillas. The window width c was chosen so that if $l_R = 0$ then $\frac{\max_M P(M)}{\min_M P(M)} < r$ with $c = \frac{\log_2(r)}{2N_a}$. In our experiments, the ratio $r = 2$ is chosen, ensuring that when $l_R(\theta) = 0$, the largest and smallest measurement probabilities differ by a factor of at most two. This promotes a near-uniform distribution of measurement results, mitigating the risk of the protocol becoming trapped in low- H regions that cannot effectively utilize feedback.

Fig. 3 shows the effect of adding the regularization term to the loss function during optimization. With the regularization enabled, the Shannon entropy directly takes its maximal value of $H = N_a$. With a feedback update frequency of 1 the infidelity reaches a local minimum with lower infidelity than without regularization.

When the update frequency is doubled to 2 or higher, entrapment in local minima occurs less frequently, showing the effectiveness of the regularization procedure. However, even though less frequently and at lower infidelities local minima can still appear. The optimization with an update frequency of 10 in Fig. 3(a) is a good example. The reason for the improvement can be correlated to the high Shannon and entanglement entropies reached due to the regularization.

Note that choosing a large update frequency does not provide significant additional benefits in avoiding local minima, but it does cause the optimization to take much longer to converge.

5 Comparing learned and reference protocol

In this section, the similarities and differences between the learned and the analytically derived protocol by Smith *et al.* [12] are examined. The primary tool for characterizing the feedback protocol is the analysis of intermediate states using quantum mutual information between two qubits:

$$I(j, j') = S(\rho_i^R) + S(\rho_{j'}^R) - S(\rho_{j,j'}^R) \quad (15)$$

$$\rho_j^R = \text{Tr}_{\bar{j}}(|\psi\rangle\langle\psi|), \quad (16)$$

where $\text{Tr}_{\bar{j}}$ traces out the entire system other than the j -th qubit. This quantity measures the extent to which two qubits are entangled when the rest of the system is traced out. Furthermore, we examine whether the learned feedback on a certain

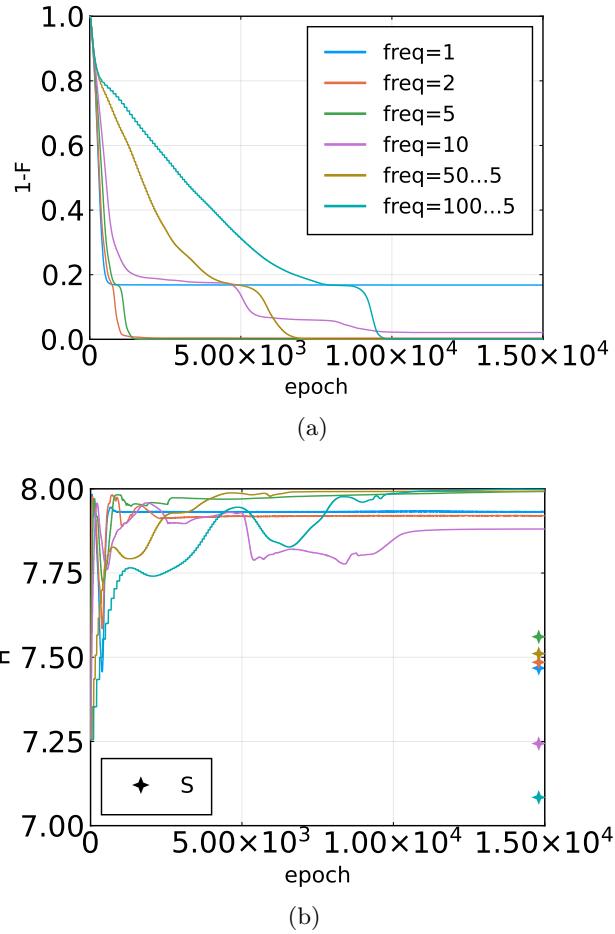


Figure 3: Same optimization like in Fig. 2 but with the ancilla regularization turned on.

qubit pair only depends on the measurement outcomes of the qubits on its left/right correctability.

The analytically derived protocol begins by preparing small AKLT states whose boundaries are entangled with adjacent ancillas. This is shown in Fig. 4(a), where the mutual information of that state reveals 6x6 blocks of strong entanglement. The ancillas are then measured in the Bell basis (see Fig. 4(b)), where the boundary conditions of the small AKLT states are merged. This process introduces defects into the AKLT state that depend on the measurement results. These defects, which can not be seen in the mutual information, can then be corrected during the feedback step (see Fig. 4(c)), allowing the system to randomly reach any of the four AKLT states.

The learned approach, on the other hand, is not limited to the block structure of the analytically derived protocol. While it learns a similar structure, it exhibits greater flexibility by generating entanglement beyond the 6x6 sub-

lattice. This additional flexibility can become particularly valuable when preparing other quantum states, as the learned protocol can exploit this ability to improve its performance.

Notably, the learned protocol needs to generate less mutual information prior to measurement (see Fig. 4(d)) compared to the analytically derived one. The analytically derived protocols feedback gates are restricted to the creation or removal of excitations in the AKLT state in the feedback step. A comparable mutual information pattern emerges when the learned protocol is constrained to the same correction gates. Interestingly, this shows not only that there are many (almost equivalent) variants of the protocol, but also that most of them require less mutual information before measurement than the analytically derived protocol.

In Fig. 4(g), both the learned and analytically derived protocol initially exhibit a maximum mutual information length, $\bar{I}(d)$, of 6 before measurement, reflecting the presence of well-structured, localized patches of mutual information generated by short-range unitaries. After measurement, the mutual information decays exponentially, similar to the AKLT state. This behavior demonstrates how both protocols effectively fuse these localized patches into a unified state, transforming short-range correlations into exponentially decaying ones.

Finally note that the smaller mutual information is consistent across all learned protocols, making them easier to implement using hardware-efficient approaches at a smaller circuit depth of 7. If, instead of learning the protocol from scratch, one attempts to replicate the analytically derived protocol by maximizing $f(\theta) = \langle 0 | U_1^{\text{smith}\dagger} U_1^{\text{learned}}(\theta) | 0 \rangle$, a circuit depth of 8 is required. Note that the minimum theoretical circuit depth to entangle the system S and ancilla A arranged in the pattern ASSSA is 6.

Another aspect of interest is whether the feedback is left- or right-correctable. This property implies that the feedback parameters $\theta_2^{i,j} = f(M_1, \dots, M_{j-1})$ depend only on measurements from one side. For example, this is the case for the analytically derived protocol of Smith *et al.* [12]. To test if the learned protocol is left correctable, we examine how the pre-measurement unitary $U_1(\theta_1)$, which is initially learned without any constraints on left correctability, can be cor-

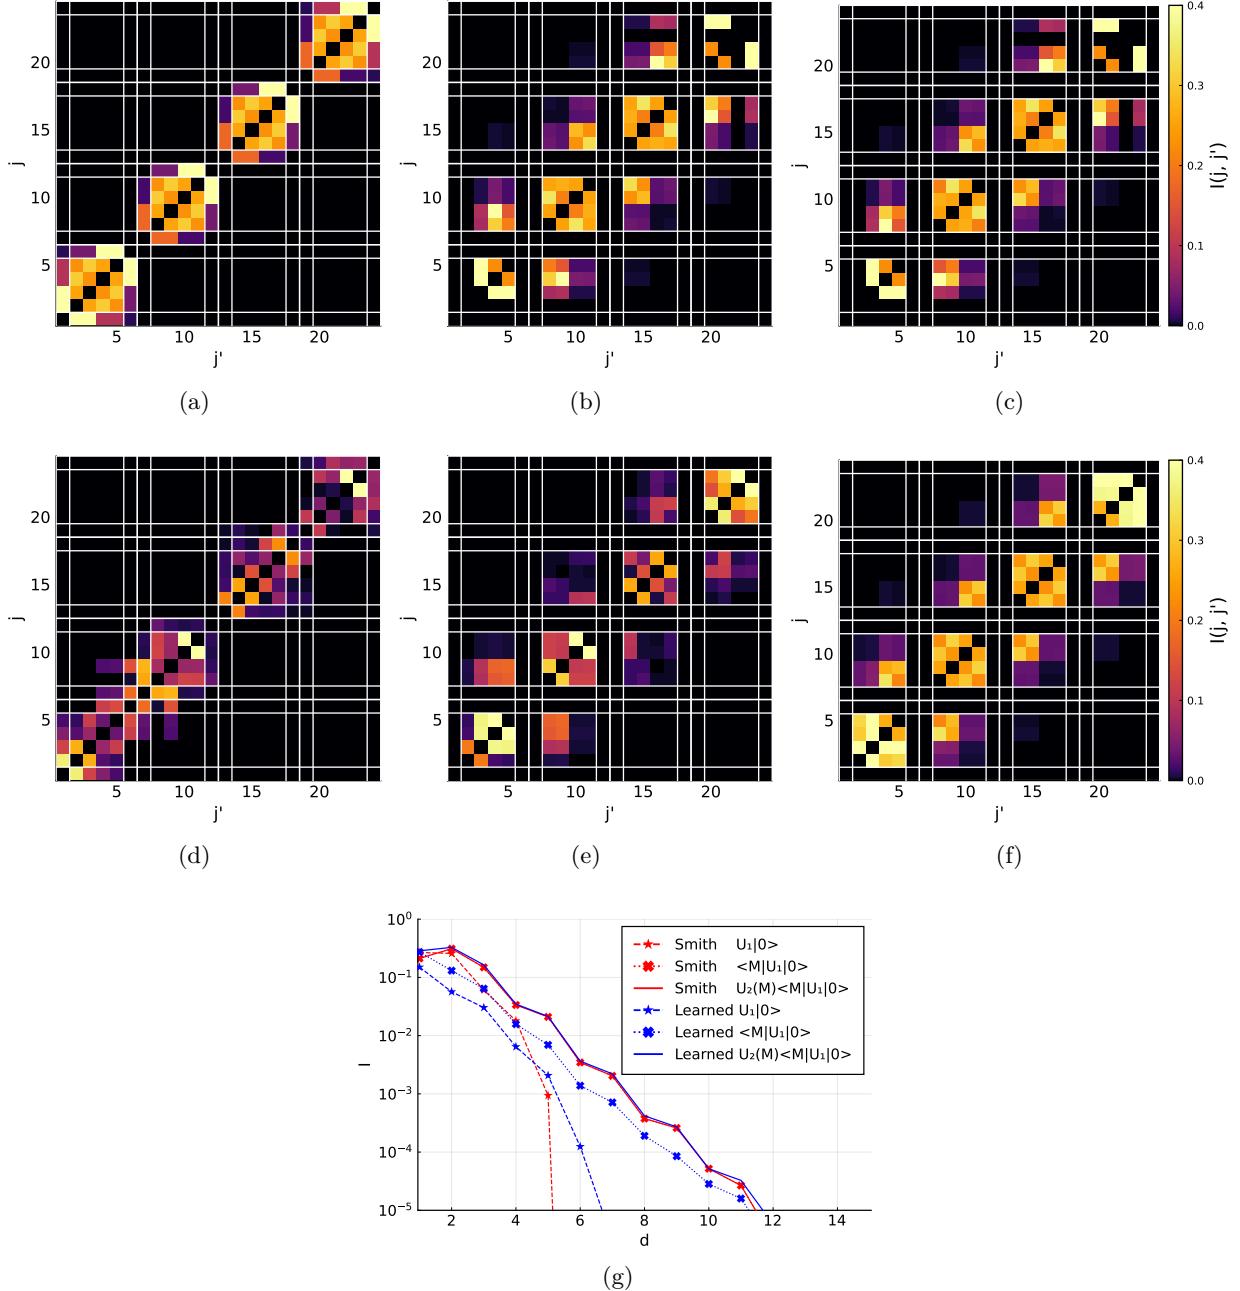


Figure 4: Mutual Information $I(j, j')$ between different sites of the quantum state after the three operations in the protocol for the analytically derived protocol (a,b,c) and the learned protocol (d,e,f), where the white lines signify that the corresponding qubit is an ancilla. First, both ancillas and system are initialized in the $|0\rangle$ state and the first unitary is applied $U|0\rangle_A|0\rangle_S$ (a, d). Then a measurement is performed $\langle M|_A U|0\rangle_A|0\rangle_S$ (b, e). Finally, the resulting state is corrected conditionally on the measurement outcome $U_S(M)\langle M|_A U_{SA}|0\rangle_S|0\rangle_A$. (g) Is the averaged mutual information $\bar{I}(d) = \frac{1}{N(d)} \sum_j I(j, j' \pm d)$ at distance d for both the analytically derived and learned protocol.

rected with left correctable feedback. We do this by constraining a new feedback unitary $U_2(\theta'_2)$ to be learned under the requirement of being left correctable (see App. E). The results indicate that U_1 is indeed left correctable, even though no such constraint is explicitly applied during its training. Furthermore, this is the case for all pro-

ocols found in this work.

6 Periodic Ansatz & RNN

The ansatz presented in this section seeks to autonomously learn a translationally invariant strategy with the help of Recurrent Neural Net-

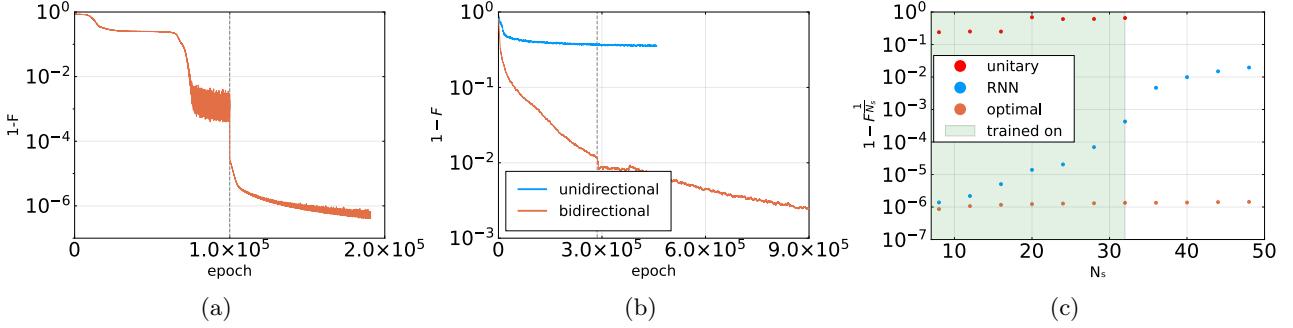


Figure 5: Figure (a) presents the optimization of the protocol for a system size of $N_s = 16$, utilizing a translationally invariant ansatz for the pre-measurement unitary U_1 , with an RNN serving as the feedback function. The ADAM optimizer's learning rate was reduced from 10^{-3} to 10^{-5} after 10^5 epochs. Figure (b) shows the RNN performance following re-initialization, with the angles of the pre-measurement unitary frozen. The RNN was trained concurrently on system sizes $N_s = [8, 12, 16, 20, 24, 28, 32]$, and the curve was smoothed with a running mean as the raw data was noisy. The blue curve shows the training for a unidirectional RNN and the orange curve for a bidirectional RNN. The gradient is calculated every epoch for 144 measurement outcomes. In Figure (c), instead of the infidelity the infidelity per site $1 - F^{1/N_s}$ is plotted to be able to compare system sizes meaningfully. The red points indicate the infidelity reached when only optimizing unitaries, removing the ancillas and measurements, the blue points indicate the performance of the RNN feedback on 1,000 measurement outcomes. In contrast, the orange points show further optimization of the feedback angles by minimizing the infidelity until convergence only for the 1,000 measurement outcomes. This comparison reveals that although the RNN does not learn the optimal feedback strategy, it performs well for small system sizes, and the learned strategy successfully extrapolates to larger sizes, even though better performance was hoped for.

works (RNN) to prepare the four-fold AKLT state manifold. The advantage of this approach lies in its ability to be trained on a set of fixed system sizes and subsequently extrapolated to larger ones. This ansatz builds upon the methods introduced in previous sections but incorporates two key modifications. First, the ansatz is modified so that the angles of the pre-measurement unitary are set to repeat with a periodicity of six, i.e., $\theta_{i,j} = \theta_{[i,j \bmod 6]}$. This is a natural choice as the qubits repeat in the ancilla qubit A and system qubit S pattern ASSSSA. Second, the feedback mechanism is designed to be system-size independent by replacing the tabular feedback approach with an RNN f . The RNN architecture comprises five layers, alternating between Gated Recurrent Units (GRUs) to propagate measurement information spatially and Swish-Gated Linear Units (SwiGLUs) to process information locally. For these experiments, we selected a hidden dimension of $d_h = 60$. Additional architectural details can be found in App. C.

The optimization procedure for $N_s = 16$, combining the translationally invariant pre-measurement unitary U_1 with bidirectional RNN-based feedback, is illustrated in Fig. 5(a). Using the translationally invariant ansatz introduces lo-

cal minima that trap the optimization process. To mitigate this, the optimization was repeated three times until a successful run avoided these local minima. Once the process escapes the local minima, further optimization requires only adjusting the learning rate when necessary. The procedure was halted upon reaching a sufficiently low infidelity of 10^{-6} .

However, the protocol at this stage is limited to the system size $N_s = 16$ since the RNN has not yet generalized to other sizes. To address this, we freeze the pre-measurement unitary angles and reinitialize the RNN. The RNN was then trained simultaneously on a range of system sizes, $N_s = [8, 12, 16, 20, 24, 28, 32]$, using the following loss function:

$$|\psi_{N_s}(M)\rangle = \langle M|_A U_1(\theta_1) |0\rangle_S^{\otimes N_s} |0\rangle_A^{\otimes N_s} \quad (17)$$

$$P_{N_s}(M) = \langle \psi_{N_s}(M) | \psi_{N_s}(M) \rangle \quad (18)$$

$$F(M, N_s; W) = \frac{|\langle \psi_{\text{target}} | U_2(\theta_2 = f(M; W)) | \psi_{N_s}(M) \rangle|^2}{P_{N_s}(M)} \quad (19)$$

$$\text{loss}(\theta) = 1 - \frac{1}{\mathcal{N}(N_s)} \sum_{N_s} \langle F(M, N_s; W)^{1/N_s} \rangle_{M \in P_{N_s}(M)} \quad (20)$$

This loss function averages over sampled measurement results at different system sizes, and uses the fidelity per site in order to avoid system size artifacts.

Initially, a unidirectional RNN was employed, as analysis using the algorithm in App. E indicated that the protocol was left-correctable. However, as shown in Fig. 5(b), the unidirectional RNN yielded poor performance. Surprisingly, no clear explanation has been identified for this underperformance. Only after switching to a bidirectional RNN did the model achieve satisfactory results, with average infidelity across all trained system sizes reducing to $3 \cdot 10^{-3}$ after $9 \cdot 10^5$ epochs. During training, as a test, the learning rate was adjusted to a cosine schedule [47] at $2.9 \cdot 10^5$ epochs, which further improved performance and was kept as a result. The optimization was ultimately halted at epoch $9 \cdot 10^5$ even though the loss was still decreasing as the training time became prohibitively long. At the last epoch, the infidelity at each system size was assessed individually and plotted in Fig. 5(c).

In this analysis, we use the infidelity per site, defined as $1 - F^{1/N_s}$, instead of the total infidelity as the figure of merit. This choice reflects the exponential scaling of fidelity with system size and allows for a more meaningful comparison across different N_s . For example, even when preparing a product state using single-site unitaries a modest constant per-site error will lead to a fidelity that decays exponentially with N_s , $F = (1 - \epsilon)^{N_s}$. Using fidelity when comparing different system sizes would mask the actual per-site performance of the protocol. Note here that for small ϵ the infidelity grows linearly with system size $1 - F = 1 - (1 - \epsilon)^{N_s} \approx \epsilon N_s$. The resulting protocol whose performance can be seen in Fig. 5(c), prepares states with infidelities that grow exponentially with system size and the performance deteriorates faster than expected for larger system sizes.

To better understand what is causing the poor performance the correction angles were further optimized for 1,000 values of M to find the best possible feedback for them.

It is observed that the infidelity per site for the optimal correction hardly scales with system size. This indicates that the RNN could still be improved, but nonetheless, it gives good performance for system sizes that are small enough. In

Fig. 5(b) one can see that the average infidelity is still decreasing, but due to the expense of the gradient calculations and the many steps required, the RNN was not optimized until convergence.

The slow convergence could be due to vanishing gradients caused by the RNN architecture. It is possible that other architectures like Transformers [48] or Mamba [49], which have better gradient flow could learn faster. This is left for future investigations.

7 Other states/ a specific AKLT state

In this section, we shift the focus from the random preparation of AKLT states to the consistent production of a specific AKLT state. Without loss of generality, our target is the AKLT state with both edge modes in the spin-up configuration. This is a challenging task as there is no known deterministic short-circuit protocol for this purpose.

The study of this particular state preparation serves two important purposes. Firstly, it demonstrates the feasibility of deterministic state preparation using feedback mechanisms. Secondly, it highlights the potential for discovering novel protocols through learning-based approaches. Currently, the Smith *et al.* protocol remains the most effective method for this operation, achieving the desired spin-up edge mode state with a 50% probability, through repeated AKLT state preparations. This is the case as initially each of the 4 AKLT edge states can be prepared with 25% probability and the edges can easily flipped simultaneously like $|\uparrow\uparrow\rangle \leftrightarrow |\downarrow\downarrow\rangle$, $|\downarrow\uparrow\rangle \leftrightarrow |\uparrow\downarrow\rangle$.

We optimize the protocol for the preparation of a single AKLT state using the methodology described in Sec. 4.2. Fig. 6(a) shows that the optimization process often gets trapped in local minima, as evidenced by the different results of different random seeds. To mitigate this issue, alternative strategies were investigated; however, these approaches did not yield significant improvements in the final fidelities. Interestingly, one optimization run achieves low infidelity due to a favorable random seed (see the violet line in Fig. 6(a)). This "lucky" run was obtained by adding gradually decreasing random noise to the parameters (θ_1, W) during the optimization process. This demonstrates the potential for high-

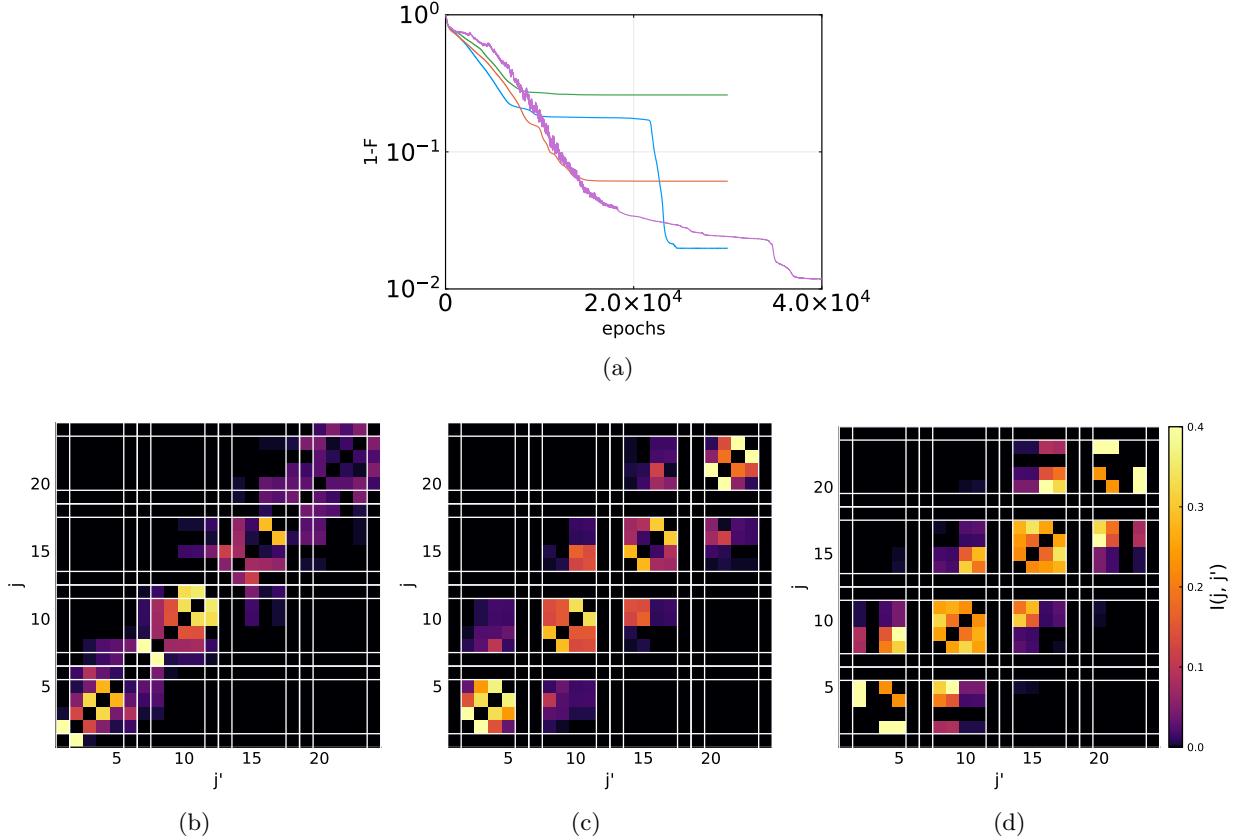


Figure 6: (a) Infidelity evolution during protocol optimization for preparing the AKLT state with both edge modes in the spin-up state. Blue, orange, and green lines represent optimizations with different initial seeds. The violet line shows a particularly successful "lucky" run that achieved significantly lower infidelity, likely due to finding a more optimal path in the parameter space. This demonstrates the potential for high-fidelity state preparation when favorable parameters are found. (b-d) Mutual Information $I(j, j')$ between different sites of the quantum state after each operation in the learned protocol. White lines indicate ancilla qubits. (b) After applying the first unitary: $U_1 |0\rangle_A |0\rangle_S$, where both ancillas and system are initialized in the $|0\rangle$ state. (c) After measurement: $\langle M|_A U_1 |0\rangle_A |0\rangle_S$. (d) After the feedback step: $U_2(\theta_2 = f(M; W)) \langle M|_A U_1 |0\rangle_S |0\rangle_A$, resulting in the prepared AKLT state with spin-up edge modes.

fidelity state preparation when advantageous parameter configurations are found and the importance of further improving optimization techniques. Consequently, this "lucky" protocol warrants further analysis.

Fig. 6(b) shows the mutual information of the intermediate state before the ancilla measurement in the "lucky" protocol. The mutual information shows a distinctive behavior compared to the one observed in previous sections. In particular, the right side of the chain shows weak mutual information before the measurement. Fig. 6(c) shows that entanglement does not fully propagate through the chain until after the measurement. This observation suggests a greater reliance on measurements for entanglement propagation in this protocol than for the simpler optimization objective of the previous

sections. The difference between the pre- and post-measurement mutual information distributions highlights the critical role of quantum measurement in this more complex protocol.

The protocol was found to be left/right correctable which is surprising. The random preparation of any of the 4 AKLT states was expected to be left correctable as the analytically derived protocol is left correctable. But for this protocol, we expected that the feedback gates would need to be dependent on both the left and right sides.

8 Conclusion

A self-learning protocol for quantum state preparation is presented that integrates measurement and feedback in variational quantum circuits. By incorporating projective measurements and con-

ditional feedback, the protocol learns efficient state preparation strategies beyond unitary-only methods, especially where measurement-based shortcuts reduce circuit depth.

Using the 1D spin-1 AKLT states as a benchmark, the protocol successfully learned a feedback mechanism to prepare these states with high fidelity. Notably, the learned approach required less mutual information prior to measurement than the analytically derived protocol of Smith *et al.*, indicating a potential for shallower circuits.

To address local minima during optimization, two strategies were implemented: adjusting parameter update frequencies to balance learning rates between the initial unitary and feedback operations, and introducing ancilla regularisation to promote uniform measurement results. These strategies effectively mitigated local minima and improved the performance of the feedback mechanism.

The protocol was extended to larger systems using a translationally invariant approach and a recurrent neural network (RNN) for feedback. While the RNN did not fully capture optimal corrections for large systems, it performed well for smaller sizes and showed some potential for generalization, even though better extrapolation performance was expected.

The preparation of a specific AKLT state with both edge modes in the spin-up configuration was also explored - a task for which no known deterministic, low-depth protocol exists. The results demonstrated the possibility of learning such a protocol, highlighting the potential of the approach to discover new state preparation methods.

Moreover, a significant potential application of our work lies in addressing open questions in the classification of phases of matter via finite-depth unitaries and feedback. Previous works, such as [50, 18], have introduced the concept of classifying quantum phases using measurement and feedback protocols. However, certain questions remain unresolved—for instance, whether all two-dimensional topological phases become trivial when feedback is added. It is suspected that some, like the Fibonacci anyon phase, remain non-trivial even with feedback. Our algorithm provides a framework that could be used to explore these questions numerically.

A practical consideration for implementing

our protocol regards the overhead introduced by measurement and classical feedback. While our simulations assume ideal conditions with instantaneous feedback, current hardware presents significant latency challenges. Active research is addressing this challenge, particularly in the context of quantum error correction [51, 52, 53]. For instance, in a recent experiment on real-time quantum error correction [51], Google reports a QEC cycle time of $1.1 \mu\text{s}$, with a real-time decoder achieving an average feedback latency of $63 \mu\text{s}$ —corresponding to a backlog of approximately 57 cycles. These constraints currently limit the applicability of feedback-based protocols to small system sizes or require buffering and delayed correction strategies. Nevertheless, the rapid pace of experimental progress suggests that these latencies will be reduced in the near future, enabling real-time feedback as envisioned in our approach.

In conclusion, the incorporation of measurement and feedback into variational quantum algorithms offers a promising framework for quantum state preparation. This approach addresses optimization challenges but does not completely solve them and extends quantum state preparation tools by exploiting the non-unitary effects of measurements. Future work should focus on using these learning techniques to design protocols with multiple rounds of feedback since most analytical protocols rely on only one round of measurements, and we see potential in the multi-round approach, which we imagine will be particularly useful in the presence of errors. Another important extension of this work is the implementation of such a protocol in an experiment, once measurement-based feedback becomes available. We envision further fine-tuning the measurement protocol in the experiment using reinforcement learning techniques so that it can learn to mitigate errors present on real hardware.

Code and simulation data— All simulations were performed using the MPS formalism to represent the wave function; the code and resulting data are available on Zenodo [54]. The Julia library mVQE.jl [55], developed specifically for these simulations, uses the ITensors.jl [56] library as its backbone.

Note added.— During the final stages of preparing this manuscript, we became aware of related work by Alam *et al.* [36] and Yan *et al.*

[35], which employs a variational ansatz similar to the one presented here. Additional details on the distinctions between the two approaches have been included in the [Introduction](#).

Acknowledgements

We thank Giovanna Morigi and Lorenzo Piroli for inspiring discussions on the topic. We acknowledge funding by the German Federal Ministry of Education and Research (BMBF) for support under the thematic programme “Quantum technologies – from the basics to the market”, project number 13N16202 “Noise in Quantum Algorithms (NiQ)”. This work was also partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) via Project-ID 277101999 – CRC network TRR 183 (“Entangled states of matter”) and under Germany’s Excellence Strategy – Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC 2004/1 – 390534769. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at the Jülich Supercomputing Centre (JSC) (Grant NeTeNeSyQuMa) and the FZ Jülich for computing time on JURECA (institute project PGI-8) [57].

References

- [1] John Preskill. “Quantum Computing in the NISQ era and beyond”. *Quantum* **2**, 79 (2018).
- [2] Matthew PA Fisher, Vedika Khemani, Adam Nahum, and Sagar Vijay. “Random quantum circuits”. *Annual Review of Condensed Matter Physics* **14**, 335–379 (2023). url: <https://doi.org/10.1146/annurev-conmatphys-031720-030658>.
- [3] Guo-Yi Zhu, Nathanan Tantivasadakarn, Ashvin Vishwanath, Simon Trebst, and Ruben Verresen. “Nishimori’s cat: stable long-range entanglement from finite-depth unitaries and weak measurements”. *Physical Review Letters* **131**, 200201 (2023). url: <https://doi.org/10.1103/physrevlett.131.200201>.
- [4] Sergey Bravyi, Isaac Kim, Alexander Kliensch, and Robert Koenig. “Adaptive constant-depth circuits for manipulating non-abelian anyons” (2022). url: <https://doi.org/10.48550/arXiv.2205.01933>.
- [5] Hiroki Sukeno and Tzu-Chieh Wei. “Quantum simulation of lattice gauge theories via deterministic duality transformations assisted by measurements”. *Physical Review A* **109**, 042611 (2024). url: <https://doi.org/10.1103/physreva.109.042611>.
- [6] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. “Variational quantum algorithms”. *Nature Reviews Physics* **3**, 625–644 (2021). url: <https://doi.org/10.1038/s42254-021-00348-9>.
- [7] Nathanan Tantivasadakarn, Ryan Thorngren, Ashvin Vishwanath, and Ruben Verresen. “Long-range entanglement from measuring symmetry-protected topological phases”. *Phys. Rev. X* **14**, 021040 (2024).
- [8] Tsung-Cheng Lu, Leonardo A. Lessa, Isaac H. Kim, and Timothy H. Hsieh. “Measurement as a shortcut to long-range entangled quantum matter”. *PRX Quantum* **3**, 040337 (2022).
- [9] Ruben Verresen, Nathanan Tantivasadakarn, and Ashvin Vishwanath. “Efficiently preparing schrödinger’s cat, fractons and non-abelian topological order in quantum devices,(2021)” (2021). url: <https://doi.org/10.48550/arXiv.2112.03061>.
- [10] Lorenzo Piroli, Georgios Styliaris, and J Ignacio Cirac. “Quantum circuits assisted by local operations and classical communication: Transformations and phases of matter”. *Physical Review Letters* **127**, 220503 (2021). url: <https://doi.org/10.1103/physrevlett.127.220503>.
- [11] Nathanan Tantivasadakarn, Ruben Verresen, and Ashvin Vishwanath. “Shortest route to non-abelian topological order on a quantum processor”. *Phys. Rev. Lett.* **131**, 060405 (2023).
- [12] Kevin C. Smith, Eleanor Crane, Nathan Wiebe, and S.M. Girvin. “Deterministic constant-depth preparation of the aklt state on a quantum processor using fu

sion measurements”. *PRX Quantum* **4**, 020315 (2023).

[13] David T. Stephen and Oliver Hart. “Preparing matrix product states via fusion: Constraints and extensions”. *Phys. Rev. B* **Pages** – (2025).

[14] Laurens Lootens, Clement Delcamp, Dominic Williamson, and Frank Verstraete. “Low-depth unitary quantum circuits for dualities in one-dimensional quantum lattice models”. *Phys. Rev. Lett.* **134**, 130403 (2025).

[15] Rahul Sahay and Ruben Verresen. “Classifying one-dimensional quantum states prepared by a single round of measurements”. *PRX Quantum* **6**, 010329 (2025).

[16] Kevin C. Smith, Abid Khan, Bryan K. Clark, S.M. Girvin, and Tzu-Chieh Wei. “Constant-depth preparation of matrix product states with adaptive quantum circuits”. *PRX Quantum* **5**, 030344 (2024).

[17] Daniel Malz, Georgios Styliaris, Zhi-Yuan Wei, and J Ignacio Cirac. “Preparation of matrix product states with log-depth quantum circuits”. *Physical Review Letters* **132**, 040404 (2024). url: <https://doi.org/10.1103/physrevlett.132.040404>.

[18] Lorenzo Piroli, Georgios Styliaris, and J. Ignacio Cirac. “Approximating many-body quantum states with quantum circuits and measurements”. *Phys. Rev. Lett.* **133**, 230401 (2024).

[19] Eric R Anschuetz and Bobak T Kiani. “Quantum variational algorithms are swamped with traps”. *Nature Communications* **13**, 7760 (2022). url: <https://doi.org/10.1038/s41467-022-35364-5>.

[20] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. “Cost function dependent barren plateaus in shallow parametrized quantum circuits”. *Nature communications* **12**, 1791 (2021). url: <https://doi.org/10.1038/s41467-021-21728-w>.

[21] Chufan Lyu, Zuoheng Zou, Xusheng Xu, Man-Hong Yung, and Abolfazl Bayat. “Variational simulation of d -level systems on qubit-based quantum simulators” (2024). url: <https://doi.org/10.48550/arXiv.2405.05051>.

[22] Zhikang T. Wang, Yuto Ashida, and Masahito Ueda. “Deep reinforcement learning control of quantum cartpoles”. *Phys. Rev. Lett.* **125**, 100401 (2020).

[23] Sangkha Borah, Bijita Sarma, Michael Kewming, Gerard J. Milburn, and Jason Twamley. “Measurement-based feedback quantum control with deep reinforcement learning for a double-well nonlinear potential”. *Phys. Rev. Lett.* **127**, 190403 (2021).

[24] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsoutsios, and M. H. Devoret. “Model-free quantum control with reinforcement learning”. *Phys. Rev. X* **12**, 011059 (2022).

[25] Matteo Puviani, Sangkha Borah, Remmy Zen, Jan Olle, and Florian Marquardt. “Non-markovian feedback for optimized quantum error correction”. *Phys. Rev. Lett.* **134**, 020601 (2025).

[26] Yaroslav Herasymenko, Igor Gornyi, and Yuval Gefen. “Measurement-driven navigation in many-body hilbert space: Active-decision steering”. *PRX Quantum* **4**, 020347 (2023).

[27] Samuel Morales, Yuval Gefen, Igor Gornyi, Alex Zazunov, and Reinhold Egger. “Engineering unsteerable quantum states with active feedback”. *Phys. Rev. Res.* **6**, 013244 (2024).

[28] Albie Chan, Zheng Shi, Luca Dellantonio, Wolfgang Dür, and Christine A Muschik. “Measurement-based infused circuits for variational quantum eigensolvers”. *Physical Review Letters* **132**, 240601 (2024). url: <https://doi.org/10.1103/physrevlett.132.240601>.

[29] Ryan Sweke, Markus Kesselring, Evert van Nieuwenburg, and Jens Eisert. “Reinforcement learning decoders for fault-tolerant quantum computation”. *Machine Learning: Science and Technology* **2** (2020).

[30] Philip Andreasson, Joel Johansson, Simon Liljestrand, and Mats Granath. “Quantum error correction for the toric code using deep reinforcement learning”. *Quantum* **3**, 183 (2019).

[31] Laia Domingo Colomer, Michalis Skotiniotis, and Ramon Muñoz-Tapia. “Reinforcement learning for optimal error correction of toric codes”. *Physics Letters A* **384**, 126353 (2020).

[32] Hendrik Poulsen Nautrup, Nicolas Delfosse,

Vedran Dunjko, Hans J. Briegel, and Nicolai Friis. “Optimizing Quantum Error Correction Codes with Reinforcement Learning”. *Quantum* **3**, 215 (2019).

[33] Thomas Fösel, Petru Tighineanu, Talitha Weiss, and Florian Marquardt. “Reinforcement learning with neural networks for quantum feedback”. *Phys. Rev. X* **8**, 031084 (2018).

[34] Yigal Ilin and Itai Arad. “Dissipative variational quantum algorithms for gibbs state preparation”. *IEEE Transactions on Quantum Engineering* (2024). url: <https://doi.org/10.1109/TQE.2024.3511419>.

[35] Yuxuan Yan, Muzhou Ma, You Zhou, and Xiongfeng Ma. “Variational locc-assisted quantum circuits for long-range entangled states”. *Phys. Rev. Lett.* **134**, 170601 (2025).

[36] Faisal Alam and Bryan K Clark. “Learning dynamic quantum circuits for efficient state preparation” (2024). url: <https://doi.org/10.48550/arXiv.2410.09030>.

[37] Frank Pollmann, Vedika Khemani, J Ignacio Cirac, and Shivaji Lal Sondhi. “Efficient variational diagonalization of fully many-body localized hamiltonians”. *Physical Review B* **94**, 041116 (2016). url: <https://doi.org/10.1103/physrevb.94.041116>.

[38] Daniel Alcalde, Felix Motzoi, Tommaso Calarco, Giovanna Morigi, and Matteo Rizzi. “Quantum state preparation via engineered ancilla resetting”. *Quantum* **8**, 1299 (2024). url: <https://doi.org/10.22331/q-2024-03-27-1299>.

[39] Sthitadhi Roy, JT Chalker, IV Gornyi, and Yuval Gefen. “Measurement-induced steering of quantum systems”. *Physical Review Research* **2**, 033347 (2020). url: <https://doi.org/10.1103/physrevresearch.2.033347>.

[40] Anne Matthes, Mark Rudner, Achim Rosch, and Erez Berg. “Programmable adiabatic demagnetization for systems with trivial and topological excitations”. *Quantum* **8**, 1505 (2024).

[41] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. “Rigorous results on valence-bond ground states in antiferromagnets”. *Phys. Rev. Lett.* **59**, 799–802 (1987).

[42] Ruben Verresen, Roderich Moessner, and Frank Pollmann. “One-dimensional symmetry protected topological phases and their transitions”. *Physical Review B* **96**, 165124 (2017). url: <https://doi.org/10.1103/PhysRevB.96.165124>.

[43] Frank Pollmann and Ari M Turner. “Detection of symmetry-protected topological phases in one dimension”. *Physical review b* **86**, 125441 (2012). url: <https://doi.org/10.1103/PhysRevB.86.125441>.

[44] Gavin K Brennen and Akimasa Miyake. “Measurement-based quantum computer in the gapped ground state of a two-body hamiltonian”. *Physical review letters* **101**, 010502 (2008). url: <https://doi.org/10.1103/PhysRevLett.101.010502>.

[45] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization” (2014). url: <https://doi.org/10.48550/arXiv.1412.6980>.

[46] Nicolas PD Sawaya, Tim Menke, Thi Ha Kyaw, Sonika Johri, Alán Aspuru-Guzik, and Gian Giacomo Guerreschi. “Resource-efficient digital quantum simulation of d-level systems for photonic, vibrational, and spin-s hamiltonians”. *npj Quantum Information* **6**, 49 (2020). url: <https://doi.org/10.1038/s41534-020-0278-0>.

[47] Ilya Loshchilov and Frank Hutter. “Sgdr: Stochastic gradient descent with warm restarts” (2016). url: <https://doi.org/10.48550/arXiv.1608.03983>.

[48] A Vaswani. “Attention is all you need”. *Advances in Neural Information Processing Systems* (2017). url: <https://doi.org/10.48550/arXiv.1706.03762>.

[49] Albert Gu and Tri Dao. “Mamba: Linear-time sequence modeling with selective state spaces” (2023). url: <https://doi.org/10.48550/arXiv.2312.00752>.

[50] Nathanan Tantivasadakarn, Ashvin Vishwanath, and Ruben Verresen. “Hierarchy of topological order from finite-depth unitaries, measurement, and feedforward”. *PRX Quantum* **4**, 020339 (2023). url: <https://doi.org/10.1103/prxquantum.4.020339>.

[51] Rajeev Acharya, Dmitry A Abanin, Laleh Aghababaie-Beni, Igor Aleiner, Trond I Andersen, Markus Ansmann, Frank Arute, Kunal Arya, Abraham Asfaw, Nikita Astrakhantsev, et al. “Quantum error cor-

rection below the surface code threshold”. Nature (2024). url: <https://doi.org/10.1038/s41586-024-08449-y>.

[52] Laura Caune, Luka Skoric, Nick S Blunt, Archibald Ruban, Jimmy McDaniel, Joseph A Valery, Andrew D Patterson, Alexander V Gramolin, Joonas Majaniemi, Kenton M Barnes, et al. “Demonstrating real-time and low-latency quantum error correction with superconducting qubits” (2024). url: <https://doi.org/10.48550/arXiv.2410.05202>.

[53] Volodymyr V Sivak, Alec Eickbusch, Baptiste Royer, Shraddha Singh, Ioannis Tsoutsios, Suhas Ganjam, Alessandro Miano, BL Brock, AZ Ding, Luigi Frunzio, et al. “Real-time quantum error correction beyond break-even”. Nature **616**, 50–55 (2023). url: <https://doi.org/10.1038/s41586-023-05782-6>.

[54] Daniel Alcalde Puente and Matteo Rizzi. “Simulations performed to produce figures presented in this work”. Zenodo (2025).

[55] Daniel Alcalde Puente. “Implementation of measurement and feedback based variational circuits”. <https://github.com/danielalcalde/mVQE> (2025).

[56] Matthew Fishman, Steven R. White, and E. Miles Stoudenmire. “The ITensor Software Library for Tensor Network Calculations”. SciPost Phys. Codebases-Page 4 (2022).

[57] Jülich Supercomputing Centre. “Jureca: Data centric and booster modules implementing the modular supercomputing architecture at jülich supercomputing centre”. Journal of large-scale research facilities **7**, A182 (2021).

[58] Adam Bene Watts, Robin Kothari, Luke Schaeffer, and Avishay Tal. “Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits”. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing. Pages 515–526. (2019). url: <https://doi.org/10.1145/3313276.3316404>.

[59] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. “The llama 3 herd of models” (2024). url: <https://doi.org/10.48550/arXiv.2407.21783>.

[60] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. “Learning phrase representations using RNN encoder–decoder for statistical machine translation”. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). Pages 1724–1734. Doha, Qatar (2014). Association for Computational Linguistics.

A The Greenberger–Horne–Zeilinger State

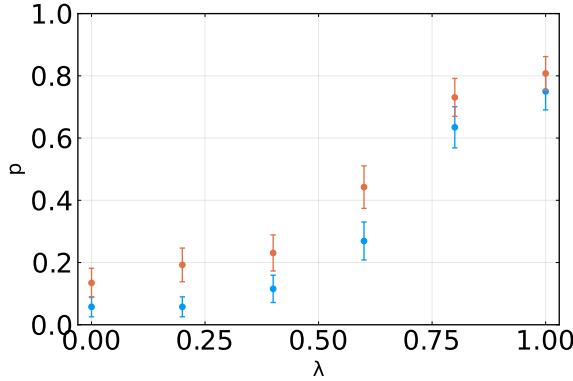


Figure 7: Success probability p as a function of the regularization parameter λ for two different optimization methods for the preparation of the 6 qubit GHZ state. The blue data points correspond to the standard ADAM optimizer, while the orange data points represent the optimizer with ancilla regularization and an update frequency of 5. The error bars indicate the standard deviation over 50 independent optimization runs. As λ increases, the success probability improves for both methods, with the ancilla-regularized approach consistently achieving higher success rates across all values of λ .

To validate our claims regarding local minima, we analyzed the simpler GHZ state, for which an efficient feedback protocol is known [58]. Due to its simplicity, the optimization process was significantly faster compared to the AKLT state, allowing us to perform statistical analysis on the optimization methods presented in this work. The infidelity used to optimize the GHZ state preparation protocol is defined as:

$$\begin{aligned}
 l(\theta; M) &= 1 - F(\theta; M) \\
 &= 1 - \frac{1}{2} |(\langle 00..| + \langle 11..|) |\psi(\theta; M)\rangle|^2 \\
 &= 1 - \frac{1}{2} |\langle 00..|\psi(\theta; M)\rangle|^2 - \frac{1}{2} |\langle 11..|\psi(\theta; M)\rangle|^2 - \text{Re}(\langle 00..|\psi(\theta; M)\rangle \langle \psi(\theta; M)|11..)) .
 \end{aligned} \tag{21}$$

The first two terms in the fidelity introduce two local minima corresponding to the trivial product states $|00..>$ and $|11..>$. The last term reaches its minimum at the GHZ state, $\text{argmin}_{\psi} \text{Re}(\langle 00..|\psi\rangle \langle \psi|11..)) = |\text{GHZ}\rangle$. As discussed in previous sections, measurement-based variational quantum circuits often encounter difficulties due to local minima induced by underutilization of the ancilla qubits. If these two local minima are not suppressed, finding the target state becomes challenging. To address this issue, we introduce a parameter λ to attenuate the impact of the local minima and redefine the loss function as follows:

$$l(\theta; M) = \left(1 - \frac{\lambda}{2}\right) - \frac{1-\lambda}{2} |\langle 00..|\psi(\theta; M)\rangle|^2 - \frac{1-\lambda}{2} |\langle 11..|\psi(\theta; M)\rangle|^2 - \text{Re}(\langle 00..|\psi(\theta; M)\rangle \langle \psi(\theta; M)|11..)) . \tag{22}$$

For $\lambda = 0$, the loss function reduces to the standard infidelity, while for $\lambda = 1$, it exclusively retains the last term. We conducted 50 optimization trials for various values of λ using a small system size of $N_s = 6$, where the optimizations were completed within a few minutes. Both a standard ADAM optimizer and an enhanced version incorporating ancilla regularization with an update frequency of 5 were employed. As expected, higher values of λ resulted in fewer occurrences of local minima, and the proposed local minima avoidance strategies improved convergence by reducing the likelihood of becoming trapped in suboptimal solutions. For $\lambda = 0$ the local minima present are the product states $|00..>$ and $|11..>$ and for $\lambda = 1$ the local minima are defined by having two or more smaller GHZ states that are not entangled with each other.

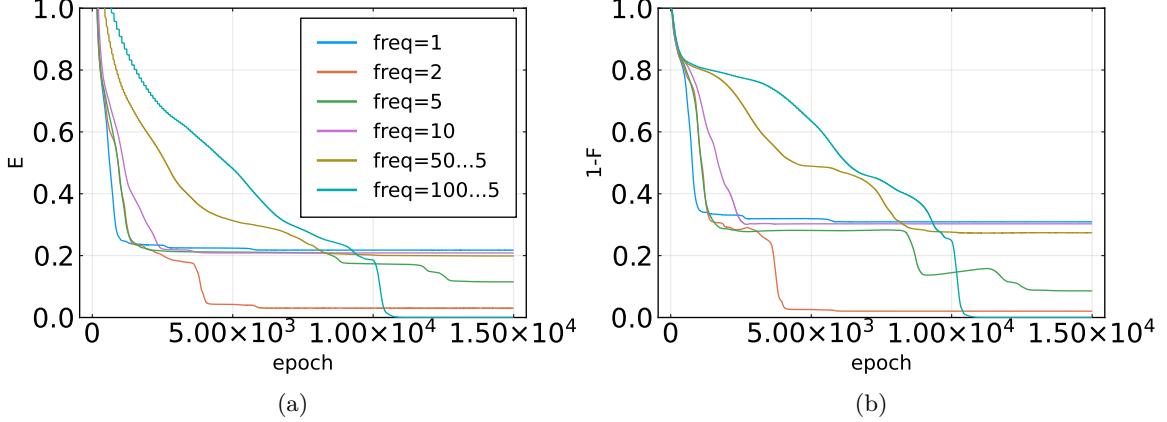


Figure 8: Same optimization as in Fig. 3, but instead of using the infidelity as a cost function the energy is optimized. In panel (a) the energy is plotted against optimization epochs and in panel (b) the fidelity is computed for the same optimization runs.

This confirms the insight we obtained from studying the AKLT. It also confirms our suspicion that the choice of the loss function is an important factor in which local minima are present and how strongly they attract the optimizer.

We note however that the difference between the use of our local minima avoidance minima and standard ADAM is not as large as the one seen for the AKLT. This might be due to the smaller system sizes that were used for this analysis, which are less prone to local minima, but are easier to perform statistics on.

B Optimizing the Energy

In the main text, infidelity was employed as a metric to evaluate how accurately a given protocol approximates the AKLT state. In this appendix, the focus shifts to examining the effects of optimizing the energy instead of the infidelity. The Hamiltonian used for this purpose consists of two components:

$$H = M_{1 \rightarrow 1/2} H_{\text{AKLT}}^1 + \left[1 - M_{1 \rightarrow 1/2} M_{1 \rightarrow 1/2}^\dagger \right] \quad (23)$$

The first term corresponds to the spin-1 AKLT Hamiltonian mapped to a spin-1/2 representation using the operator $M_{1 \rightarrow 1/2}$. The second term is a projector onto the forbidden computational subspace. It enforces that the two spin-1/2 particles form a valid spin-1 state in the ground state.

When energy is used as the objective function, the optimization landscape becomes significantly more complex. As shown in Fig. 8, only one optimization run successfully converged to a local minimum. By contrast, employing fidelity as the objective led to successful convergence in nearly all optimization runs, as illustrated in Fig. 3.

Using infidelity as the figure of merit results in a landscape where only the target state acts as an attractor. In contrast, energy-based optimization renders all eigenstates attractors. This increases the number of local minima and complicates the optimization process. Therefore, when fidelity is available, it is generally preferable to avoid using energy as the objective.

From a classical optimization perspective, fidelity also provides computational advantages. Its gradients are typically less costly to evaluate. The energy gradient takes the form

$$\left\langle \frac{\partial \psi(\theta)}{\partial \theta} \right| H \left| \psi \right\rangle, \quad (24)$$

whereas the gradient of the fidelity is given by

$$\left\langle \psi_{\text{target}} \right| \psi(\theta) \left\langle \frac{\partial \psi(\theta)}{\partial \theta} \right| \left| \psi_{\text{target}} \right\rangle. \quad (25)$$

In most cases, the target state $|\psi_{\text{target}}\rangle$ exhibits much lower entanglement than $H|\psi\rangle$. Consequently, fidelity gradients can be computed more efficiently than energy gradients when using Matrix Product States.

C Recurrent Neural Network

The Recurrent Neural Network (RNN) structure used in this work is illustrated in Fig. 9. This architecture is derived from the Llama 3 model [59], with the self-attention mechanism replaced by an RNN.

The RNN was selected for its ability to capture spatial dependencies necessary for inferring feedback gates from measurement outcomes. Both a bidirectional RNN as well as a unidirectional RNN were used. Note that when using a unidirectional RNN the first layer was switched from a Linear layer to a 1D Convolution with filter size 5, in order to make sure that the RNN could see enough measurement results to its right.

To mitigate the vanishing gradient problem often encountered in RNN training, a Gated Recurrent Unit (GRU)[60] was employed. GRUs address vanishing gradients by controlling how information is updated and forgotten, thus enabling effective learning over extended sequences. Long Short-Term Memory (LSTM) units were also considered, but they demonstrated inferior performance compared to GRUs.

Despite its strengths, the RNN architecture has limitations. Specifically, it struggles with long-range dependencies, reducing its effectiveness for large system sizes compared to transformer-based models. Transformer architectures, which rely on self-attention, may offer better gradient flow, particularly for long sequences. Future work could explore transformer or mamba-based models to potentially enhance performance.

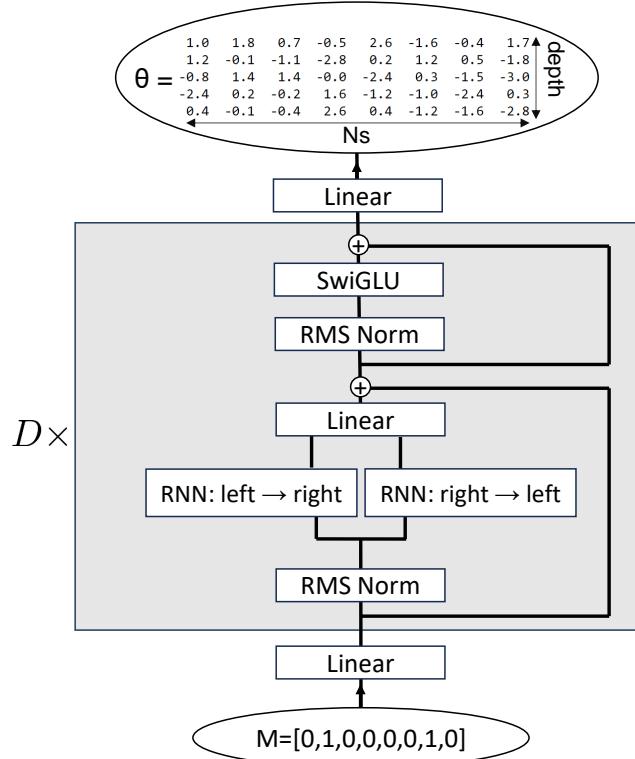


Figure 9: The architecture of the Recurrent Neural Network (RNN) used in this work. The input measurements pass through RMS normalization, a bidirectional RNN, and SwiGLU layers, generating the output angles θ for feedback unitaries. The grey box indicates repeated application of these core layers up to depth D .

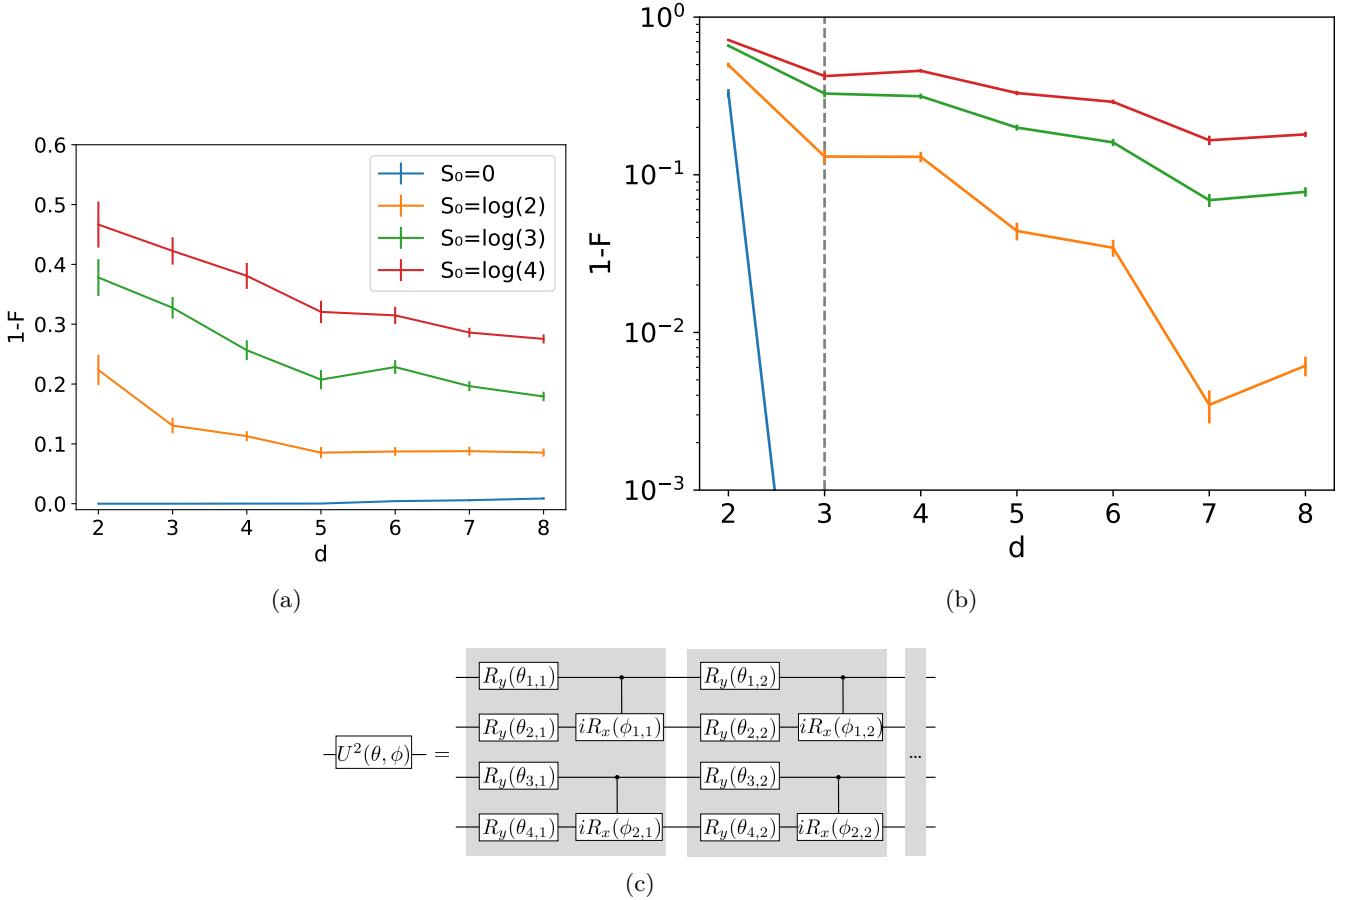


Figure 10: (a) Plot showing the minimally obtained infidelity when trying to learn a target state prepared by a hardware-efficient ansatz as a function of circuit depth d for different initial wave functions with entropies S_0 . The teacher-student approach was used, where the teacher circuit was used to generate the target state, and the student circuit attempted to replicate it. (b) Infidelity as a function of circuit depth with the teacher circuit depth fixed at $d_t = 3$. The plot shows that infidelity decreases with increasing student circuit depth, but at a slow rate. (c) Diagram of the feedback ansatz used, showing the gate structure capable of representing any two-qubit gate at a depth of five.

D Feedback Unitaries

Deciding on the optimal ansatz for the feedback step of the protocol is a nontrivial task. The goal is to use an ansatz that avoids creating long-range entanglement, ensuring that any observed long-range entanglement is entirely due to the measurement process.

To achieve this, a short hardware-efficient ansatz with a structure similar to that depicted in Fig. 1 is ideal. However, short ansätze of this type are prone to numerous local minima [19, 20]. States that are close in Hilbert space may be significantly distant in parameter space, complicating optimization. This becomes specifically difficult when the initial state the variational circuit is applied on carries some entanglement. This was observed when trying to learn the feedback but is laid out here with the help of a simple toy problem.

This can be demonstrated by generating two random parameter sets, θ for the teacher and θ' for the student, and attempting to optimize $\max_{\theta'} \langle \psi_0 | U(\theta)^\dagger U(\theta') | \psi_0 \rangle$ for various depths. Although a minimum exists at $\theta' = \theta$, the ability of the optimization algorithm to locate it depends on the prevalence of local minima in the loss landscape.

Figure 10(a) illustrates this calculation for a hardware-efficient ansatz, where the initial state $|\psi_0\rangle$ is a random MPS with bipartite entropy S_0 and a system size of eight qubits. When the initial entropy

is $S_0 = 0$, the teacher-student infidelity remains close to zero across all circuit depths, indicating an absence of significant local minima in the loss landscape. However, for initial states with nonzero entropy, such as in the feedback protocol where $|\psi_0\rangle_S = \langle M|_A U_1 |0\rangle_{S,A}$, the figure demonstrates that increasing circuit depth helps reduce the number of local minima, but the improvement is not substantial. The initial entropy of the quantum state has a significant impact on the complexity of the loss landscape.

To evaluate the difficulty of reproducing a quantum state prepared at a specific depth, the teacher circuit depth is fixed at $d_t = 3$. As shown in Fig. 10(b), for $S_0 \neq 0$ the infidelity decreases with increasing student circuit depth, but only slowly. These two plots provide insight into the complexity of the loss landscape for shallow circuits with entangled initial states, indicating that such circuits are not suitable if the initial state has some entanglement.

To optimize the circuit effectively, a greater circuit depth would be required, which is undesirable in this context. The aim is to demonstrate that entanglement is primarily propagated by the measurement process, rather than by the unitary gates. Therefore, to avoid creating excessive entanglement, an alternative approach was adopted. Instead of using a standard hardware-efficient ansatz, a sparsely connected ansatz was employed (see Fig. 10). This ansatz utilizes a two-body gate:

$$CiRX(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos\left(\frac{\theta}{2}\right) & \sin\left(\frac{\theta}{2}\right) \\ 0 & 0 & \sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) \end{pmatrix} \quad (26)$$

that can interpolate between the CNOT and the identity gate. The ansatz is capable of representing any two-qubit gate at a depth of five, which was used in all experiments. Notably, the teacher-student infidelity for a depth of five is indeed zero.

E Left correctability of the Protocol

In this section, the method used to determine if a protocol's feedback can be corrected with information from only one direction is presented. This is relevant as for example in the protocol of Smith *et al.* [12] only information from one side is necessary to establish the feedback gates. Using the technique described here, it was confirmed that all protocols developed in this work are indeed left correctable.

The feedback angles are first determined for a specific measurement outcome M . Then, the stability of these angles is assessed by modifying individual bits in the measurement outcome and observing the resulting changes in the feedback angles.

The algorithm proceeds as follows: given a measurement outcome M , the optimal feedback angles θ are determined. Subsequently, M is modified by flipping a specific bit, indicated by ‘index’, resulting in a new measurement outcome M' . The new set of feedback angles θ' is then determined for M' . The difference $\Delta\theta$ between θ and θ' serves as an indicator of whether the feedback mechanism can be corrected using information from a single direction. As an example this procedure was done for the AKLT and plotted in Fig. 11.

Algorithm 2 Feedback Correctability Check

```
function LEFT_CORRECTABLE(index)
     $|\psi_1\rangle \leftarrow U_1 |0\rangle$ 
     $M \leftarrow \text{sample}_A(|\psi_1\rangle)$ 
     $\theta \leftarrow \max_{\theta} [\langle \psi_{\text{target}} | U_2(\theta) \langle M | |\psi_1\rangle]$ 
     $M' \leftarrow \text{flipbit}(M, \text{index})$ 
     $\theta' \leftarrow \max_{\theta} [\langle \psi_{\text{target}} | U_2(\theta) \langle M' | |\psi_1\rangle - \lambda \sum_{i,j=\text{index}} (\theta_{i,j} - \theta'_{i,j})^2]$ 
     $\Delta\theta = |\theta - \theta'|$ 
    return  $\Delta\theta$ 
end function
```

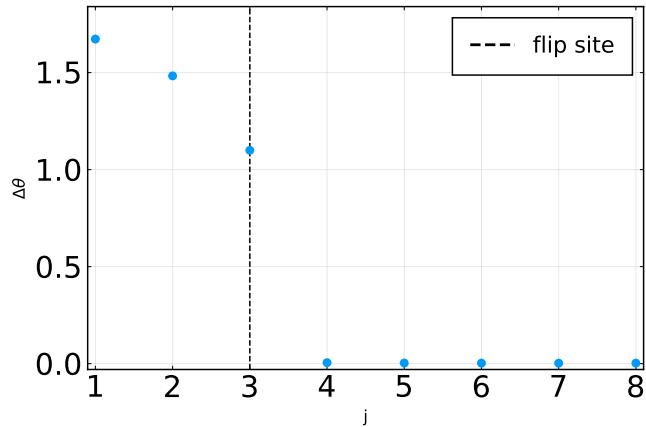


Figure 11: Figure showing the difference in feedback angles, $\Delta\theta$, for each site j after flipping a specific bit in the measurement outcome. The black vertical line represents the site where the bit was flipped.

PUBLICATION III

Efficient Optimization and Conceptual Barriers in Variational Finite Projected Entangled-Pair States

Authors: Daniel Alcalde Puente*, Erik Lennart Weerda*, Konrad Schröder, Matteo Rizzi
Year: 2025
Journal: Physical Review B
Volume/Issue: 111, 195120
DOI: [10.1103/PhysRevB.111.195120](https://doi.org/10.1103/PhysRevB.111.195120)

* Both first authors have contributed equally.

Summary

Projected Entangled Pair States (PEPS) provide a natural tensor network ansatz for representing ground states of local many-body Hamiltonians in two spatial dimensions. By construction, they obey the boundary law scaling of entanglement entropy (Sec. 2.2). Translation-invariant infinite PEPS (iPEPS) have seen substantial methodological progress and are now widely applied to a variety of physical systems [60, 62, 140]. In contrast, a comparably robust and standardized framework is still lacking for finite PEPS [43, 65, 67].

The simulation of finite PEPS presents several distinct computational challenges. Contractions of the double-layer tensor network scale as $\mathcal{O}(D^{10})^1$, making large bond dimen-

¹Although iPEPS calculations suffer from similar unfavorable scaling, the CTMRG algorithm targets a fixed point of the reduced density matrix in the thermodynamic limit. This structure allows for

sions computationally prohibitive. Sampling-based methods using single-layer contractions can reduce the computational cost to $\mathcal{O}(D^7)$ under favorable conditions. Despite this advantage, the broader applicability of finite PEPS is limited not by the expressivity of the ansatz but by the absence of reliable and efficient algorithms for general quantum states. Key aspects such as numerical stability and the cost of evaluating physical observables remain insufficiently characterized. Moreover, the approximations required for tensor network contractions often prevent rigorous variational guarantees for ground state energy estimates. Extensions to systems with long-range interactions have received limited attention and remain an open area of investigation.

This work introduces several methodological advances to address these limitations. A reformulation of the time-dependent variational principle (TDVP), originally developed for neural quantum states [45], is adapted to enhance the stochastic optimization of finite PEPS (Sec. 3.2).

The computational complexity of sampling from finite PEPS via single-layer contractions is analyzed. This analysis reveals an entanglement phase transition in the quantum states² represented by random PEPS. On one side of the transition, random positive tensors exhibit product-state behavior in the thermodynamic limit. On the other side, the PEPS realize the maximal entanglement permitted by the ansatz, and the cost of sampling increases exponentially with system size. In contrast, structured quantum states of physical interest typically exhibit lower contraction complexity than random PEPS. A practical diagnostic, the spectral entropy, is introduced to predict contraction difficulty. This quantity enables targeted acceleration of single-layer contractions. Sec. 2.2.1 demonstrates computational speedups of approximately 50%.

The framework is validated through applications to physically relevant models. Simulations are performed for a gapped Hamiltonian supporting a chiral spin-liquid ground state, which remains a challenging case for PEPS-based methods [141, 142]. The sampling-based approach is also shown to accommodate long-range interactions, as demonstrated by simulations of interacting Rydberg atom arrays.

The variational nature of the PEPS ansatz is preserved through controlled approximations, which support the derivation of rigorous upper bounds on ground state energies.

Future research should aim at a systematic comparison between various approaches for ground state calculations in two-dimensional quantum systems. These include different variants of PEPS, neural quantum states (NQS) [70], tree tensor networks (TTNs) [143–

more flexibility in dynamically adjusting the environment dimensions, which can mitigate some of the associated computational cost.

²Note that in previous works such as Refs. [73, 74], the entanglement phase transition was identified in the boundary MPS used for contractions of random single-layer PEPS, not in the quantum state itself.

145], and augmented TTNs [146]. A clear understanding of the optimal choice of method for different problem classes remains unclear.

Future research should also explore the preparation of two-dimensional quantum states using finite PEPS to represent the evolved state. Previous work presented in this thesis has focused exclusively on one-dimensional systems [I,II], where matrix product state (MPS) representations were employed. With recent advancements in finite PEPS methods, it may become feasible to simulate measurement-based feedback algorithms in two dimensions and to evaluate their effectiveness in preparing nontrivial topological states.

Another promising direction is the application of this framework to time-dependent and out-of-equilibrium quantum dynamics. In Sec. 3.3.2, we propose a method to reduce the number of required Monte Carlo samples to a manageable level, although we do not test it due to a lack of time.

Author’s Contribution

Daniel Alcalde Puente contributed to the conceptual development and the implementation of the computational strategies. He designed the code structure, optimized its performance, and conducted all numerical experiments on the contractibility of both random and physically motivated PEPS. He was one of the two main contributors, alongside Erik Weerda, to the analysis of these results. Although not directly involved in the numerical study of Chiral Spin Liquids or Rydberg atom arrays (Sec. V.b and V.c), he participated in discussions regarding the feasibility of their contractibility. He contributed approximately one-third of the manuscript writing.

Data Availability

The implementation of the algorithms presented in this work is available in the Julia packages `QuantumNaturalfPEPS.jl` [147] and `QuantumNaturalGradient.jl` [148]. These packages provide a flexible interface for implementing arbitrary local Hamiltonians. All scripts used to generate the results in this work, along with the corresponding simulation data, are available on Zenodo [149].

Efficient optimization and conceptual barriers in variational finite Projected Entangled-Pair States

Daniel Alcalde Puente  ^{1,2,*} Erik Lennart Weerda  ^{1,*} Konrad Schröder  ^{1,2} and Matteo Rizzi  ^{1,2}

¹*Institute for Theoretical Physics, University of Cologne, D-50937 Köln, Germany*

²*Forschungszentrum Jülich GmbH, Institute of Quantum Control,*

Peter Grünberg Institut (PGI-8), 52425 Jülich, Germany

(Dated: March 26, 2025)

Projected entangled pair states (PEPS) on finite two-dimensional lattices are a natural ansatz for representing ground states of local many-body Hamiltonians, as they inherently satisfy the boundary law of entanglement entropy. In this paper, we propose the optimization of PEPS via an improved formulation of the time-dependent variational principle (TDVP), namely the minimum-step stochastic-reconfiguration scheme recently introduced for neural quantum states. We further discuss possible numerical issues that might arise in such a sampling-based approach. In this context, we investigate the entanglement properties of random PEPS and find an entanglement phase transition. We note that on one side of this transition, we can identify positive random tensors as product states. To demonstrate the power of the framework described in this paper, we apply the PEPS to study the notoriously challenging chiral spin liquids. Moreover, we exhibit our approach's capability to naturally handle long-range interactions by exploring the phase diagram of Rydberg atom arrays with long-range interactions. We further provide parallelized easy-to-use code, allowing the straightforward application of our method to general Hamiltonians composed of local interaction terms.

I. INTRODUCTION

Understanding the ground state properties of interacting many-body systems remains a challenging topic in condensed matter physics. Tensor network methods [1–3] have emerged as a particularly fruitful approach to address this problem. These methods originate from the development of the *density-matrix renormalization group* [4] and the subsequent analysis using insights from quantum information theory [5–7]. In two spatial dimensions, the most natural tensor network ansatz for ground state studies are the *projected entangled-pair states* (PEPS) [8]. PEPS are particularly suitable as they inherently satisfy a boundary-law scaling of the entanglement entropy.

In recent years the translation-invariant infinite PEPS (iPEPS) have been the subject of substantial method development [9–13], establishing iPEPS as a standard tool for the investigation of ground state properties in two dimensions and allowing for an increasingly wide scope of applicability [14–27]. On finite lattices, several PEPS studies with different algorithms have been performed [28–31], including some utilizing gradient-based approaches [32–34]. However, no standard numerical approach has been established so far due to various numerical challenges. In parallel to these efforts, a completely different approach to the PEPS on finite systems was developed based on the statistical evaluation of expectation values using sampling [35]. Later, gradient-based optimization methods [36, 37] have been used in conjunction with these sampling techniques. These have been applied

successfully to paradigmatic models of frustrated magnetism [38–40]. A particularly promising development for large-scale calculations using the sampling-based approach has been the development of a highly parallelizable method for directly generating samples from the PEPS state vector [41].

In this work, we investigate several aspects of sampling methods for finite PEPS. We emphasize that the most effective gradient-based optimization scheme, which relies on the *time-dependent variational principle* (TDVP) [42–45] can be implemented more efficiently by incorporating recent developments from the neural network literature [46], where it is also referred to as Stochastic Reconfiguration. Moreover, we demonstrate that despite using multiple approximations in numerical computations, the sampling-based PEPS approach can remain variational, as highlighted in [47]. Consequently, it enables the determination of rigorous upper bounds for the ground state energy. We demonstrate that the commonly employed, albeit formally less rigorous, approximation method yields estimates that closely align with the strictly variational results, yet it offers significantly enhanced computational efficiency.

Next, we examine the computational complexity associated with contracting single-layer tensor networks, a crucial step in the sampling-based finite PEPS framework. Our analysis reveals several findings. First, we demonstrate that random PEPS undergo an entanglement phase transition, coinciding with a complexity transition previously identified in random single-layer tensor networks [48]. On one side of this entanglement transition we find the random positive PEPS, in essence, to be product states. Further, the computational difficulty of contractions encountered in physically motivated scenarios is significantly lower than the worst-case scenarios

* Both first authors have contributed equally.

identified in random tensor networks [48]. We find that the entanglement present in the quantum state does not dictate the contractibility of its samples. Additionally, we identify a quantity that effectively serves as a predictor for the contraction complexity of physical PEPS samples.

Finally, we move to the applications of the finite PEPS for challenging physical situations. We show that we can successfully apply the proposed finite PEPS framework to a gapped Hamiltonian hosting a chiral spin-liquid as its ground state, a situation that has been notoriously difficult for PEPS. Further, we point out that the finite PEPS sampling framework can treat very conveniently certain long-range interactions, which can be extremely challenging in alternative approaches. We utilize this to study the phases of a model with long-range interacting Rydberg-atom arrays and point out further possible applications.

The paper is organized as follows: In Sec. II, we present the sampling framework for Projected Entangled Pair States (PEPS) and introduce an improved approach to implementing the time-dependent variational principle (TDVP) within this sampling framework. In Sec. III, we detail the numerical methods employed and address aspects related to variational upper bounds and the approximations involved in PEPS calculations. Sec. IV is dedicated to analyzing the complexity associated with tensor network contractions in sampling-based PEPS computations. Finally, in Sec. V, we apply the developed finite PEPS methods to several models, including chiral spin liquids and Rydberg atom arrays.

II. FINITE PEPS AND SAMPLING

Projected Entangled Pair States (PEPS) provide a powerful ansatz for studying quantum many-body systems on finite lattices of size $L_x \times L_y$. As they satisfy the boundary law of entanglement entropy in two dimensions, they are well-suited for representing ground and low-energy states of local Hamiltonians.

The fundamental building blocks of finite PEPS are local tensors, which are defined on each lattice site

$$T[x, y] = \begin{array}{c} \text{---} \\ \text{---} \end{array} \boxed{T[x, y]} \begin{array}{c} \text{---} \\ \text{---} \end{array} . \quad (1)$$

Every local tensor has a single *physical* index (bold line), labeling a basis of its local physical Hilbert space of dimension d , as well as a set of *virtual* indices (horizontal and vertical lines) of dimension D that connect to the local tensors on neighboring sites. Within the PEPS

ansatz, we express the coefficient tensor of the many-body state as the contraction of all virtual indices, illustrated as connected legs, of the local tensors

$$|\Psi\rangle = \sum_{\{\mathbf{S}\}} \Psi(\mathbf{S}) |\mathbf{S}\rangle, \quad |\mathbf{S}\rangle = |s_{(1,1)} \dots s_{(L_x, L_y)}\rangle$$

$$\Psi(\mathbf{S}) = \begin{array}{c} \text{---} \\ \text{---} \end{array} \boxed{\dots} \begin{array}{c} \text{---} \\ \text{---} \end{array} \quad (2)$$

Here, \mathbf{S} denotes a single many-body configuration, and the red dots signify that we fixed the index at position (x, y) to a certain value $s_{(x,y)}$.

To calculate expectation values with finite PEPS, one in principle has to perform a contraction of a double-layer tensor network, in which one layer corresponds to the *bra*- and the other to the *ket*-vector. However, the exact contraction of a two-dimensional double-layer tensor network is computationally hard, such that schemes to perform the contraction approximately have to be used in practice, see e.g. [30], which can still be challenging.

One alternative to this has been pioneered by Wang *et al.* [35], who demonstrated that one can use sampling techniques to statistically approximate expectation values of finite PEPS. In this framework, only the contraction of single-layer tensor networks is necessary, which can result in substantially cheaper numerical calculations. A very similar sampling-based approach is also employed in methods using the neural quantum state ansatzes [44]. To optimally make use of the statistical approach for ground-state calculations, we only need to be able to perform a few numerical operations. Firstly, in order to evaluate expectation values

$$\langle \hat{O} \rangle = \frac{\langle \Psi | \hat{O} | \Psi \rangle}{\langle \Psi | \Psi \rangle} = \sum_{\mathbf{S}} \underbrace{\frac{|\Psi(\mathbf{S})|^2}{\langle \Psi | \Psi \rangle}}_{p_{\Psi}(\mathbf{S})} \underbrace{\frac{\langle \mathbf{S} | \hat{O} | \Psi \rangle}{\Psi(\mathbf{S})}}_{O_{\text{loc}}(\mathbf{S})} \quad (3)$$

$$\langle \mathbf{S} | \hat{O} | \Psi \rangle = \sum_{\mathbf{S}'} O_{\mathbf{S}, \mathbf{S}'} \Psi(\mathbf{S}'),$$

we need to be able to generate samples \mathbf{S} from the state-vector according to $p_{\Psi}(\mathbf{S})$, and further be able to calculate the coefficient $\Psi(\mathbf{S})$ of the many-body state vector for a given sample. Secondly, to optimize the ansatz (finite PEPS or neural quantum state) one needs the ability to calculate the gradient $\frac{\partial \Psi_{\theta}(\mathbf{S})}{\partial \theta}$ of each state vector coefficient w.r.t. its underlying variational parameters, which we label as θ and suppress in the expressions in general. For the PEPS these variational parameters are

the entries in the local tensors, which results, e.g., in

$$\frac{\partial \Psi_\theta(\mathbf{S})}{\partial T[3, 2]} = \begin{array}{c} \text{Diagram of a 3D PEPS tensor network with 4 layers and 3 columns, representing the derivative of the wavefunction with respect to the third column of the third layer.} \end{array} . \quad (4)$$

In Sec. III we discuss the numerical techniques to perform these necessary tasks and highlight in particular that, for local lattice models, the calculation of (energy) expectation values can be done in a distinctly more efficient way as compared to e.g. neural quantum states.

A. Reformulation of the imaginary-time Schrödinger equation for sampling techniques

To optimize the PEPS, a gradient-based scheme is used. Specifically, the time-dependent variational principle (TDVP) [42–45] is employed. It governs the evolution of parameters θ in a variational wave function $|\Psi(\theta)\rangle$ to ensure that the state vector evolves according to the Schrödinger equation as close as possible. Here, we are specifically interested in the use of the TDVP in the sampling-based approach for the evolution according to the imaginary-time Schrödinger equation. We note that, while we discuss the use for ground state search here, the same approach could be in principle used to simulate real-time evolution. We comment on related difficulties in App. C.

We derive the equations for the evolution of the parameters θ by simple manipulation of the imaginary-time Schrödinger equation

$$|\dot{\Psi}(\theta)\rangle = -H|\Psi(\theta)\rangle \quad (5)$$

$$\frac{\partial |\Psi(\theta)\rangle}{\partial \theta} \dot{\theta} = -H|\Psi(\theta)\rangle \quad (6)$$

$$\sum_i \underbrace{\frac{\partial \langle \mathbf{S} | \Psi(\theta) \rangle}{\partial \theta_i}}_{\tilde{O}_{\mathbf{S},i}} \dot{\theta}_i = -\underbrace{\langle \mathbf{S} | H | \Psi \rangle}_{\tilde{E}_{\mathbf{S}}^{\text{loc}}} \quad \forall \mathbf{S} \quad (7)$$

$$\tilde{O} \dot{\theta} = \tilde{E}_{\mathbf{S}}^{\text{loc}}. \quad (8)$$

Note that if we write Eq. (5) with an explicit normalization for $|\Psi\rangle$ we obtain an additional term.

If the number of parameters is smaller than the dimension of the Hilbert space the above equations will in general not have an exact solution. Instead, we are looking for the best approximation of a solution in the parameter space

$$\min_{\dot{\theta}} \sum_{\mathbf{S}} \left\| \sum_i \tilde{O}_{\mathbf{S},i} \dot{\theta}_i - \tilde{E}_{\mathbf{S}}^{\text{loc}} \right\|^2. \quad (9)$$

Instead of finding the best approximation on the entire Hilbert space, we use Monte Carlo sampling to find a solution on the relevant, sampled part of the Hilbert space whose dimension is the number of samples N_s

$$\min_{\dot{\theta}} \sum_{\mathbf{S}} p_{\Psi}(\mathbf{S}) \left\| \sum_i \underbrace{\frac{\tilde{O}_{\mathbf{S},i}}{\Psi(\mathbf{S})}}_{=:O_{\mathbf{S},i}} \dot{\theta}_i - \underbrace{\frac{\tilde{E}_{\mathbf{S}}^{\text{loc}}}{\Psi(\mathbf{S})}}_{=:E_{\mathbf{S}}^{\text{loc}}} \right\|^2. \quad (10)$$

Conventionally this is solved by defining a matrix G which is approximated on the samples drawn

$$\dot{\theta} = \underbrace{(O^\dagger O)^{-1} O^\dagger}_{=:G} E_{\text{loc}} \quad (11)$$

$$G_{i',i} = \sum_{\mathbf{S}} p_{\Psi}(\mathbf{S}) O_{\mathbf{S},i}^* O_{\mathbf{S},i} \quad (12)$$

Unfortunately, the matrix G has dimensions $N_p \times N_p$, where N_p denotes the number of parameters in the ansatz. For parameter-intensive ansätze such as finite PEPS, where $N_p = L_x L_y D^4 d$, storing G becomes infeasible. Consequently, slow iterative solvers are typically employed. Importantly, the rank of the sampled G matrix is at most the number of samples, N_s .

Recently, in the context of neural quantum states, an alternative method referred to as minimum-step stochastic-reconfiguration (minSR) was proposed by Chen and Heyl [46]. It addresses the issue that arises for ansatzes with large numbers of parameters. Their method works by leveraging the low rank of G in the scenario where there are more parameters than samples

$$\dot{\theta} = O^\dagger \underbrace{(O O^\dagger)^{-1}}_T E_{\text{loc}} \quad (13)$$

$$T_{\mathbf{S},\mathbf{S}'} = \sum_i O_{\mathbf{S},i} O_{\mathbf{S}',i}^*. \quad (14)$$

Here, the matrix T is only considered on the sampled subspace of the Hilbert space and hence has dimensions $N_s \times N_s$. Thus, it is significantly more manageable to store it and compute its pseudoinverse. In Fig. 1 we show that, when we apply this technique to finite PEPS, it substantially increases the efficiency of solving the linear system (8) for $\dot{\theta}$, which is necessary in every optimization step. The fraction of the total compute time of the ground state search with finite PEPS varies substantially for different use cases, but to get some idea of the speedup our code was timed for the J_1 - J_2 model using both methods.

We observe up to a 40% speedup in total compute time, and a 39 times faster solving of the TDVP equation for $L = 16$ PEPS. With higher bond dimensions and large system sizes yielding the most gains. Note that these percentages were computed by comparing a Krylov method to solve Eq. (11) and an eigenvalue decomposition method to solve Eq. (13).

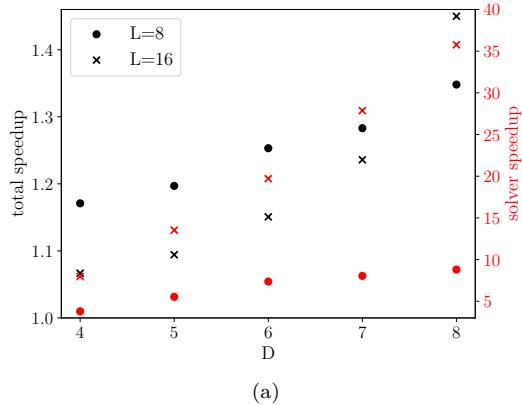


Figure 1. Speedup of employing minSR vs a Krylov-based iterative solver for different bond dimensions for system sizes $L = 8, 16$ with $N_s = 1000$ during the optimization. (black) Is the total time quotient $\frac{t_{\text{total}}[\text{krylov}]}{t_{\text{total}}[\text{minSR}]}$. (red) The quotient time for the solvers $\frac{t_{\text{solver}}[\text{krylov}]}{t_{\text{solver}}[\text{minSR}]}$

III. METHODS

When employing the PEPS on a finite lattice in the sampling-based framework, we use several numerical tools, which are summarized in this section. We further comment on the approximations made in the sampling-based PEPS scheme by comparing energies to true variational upper bounds to the ground state energy, which can be obtained at a higher cost.

A. On the boundary MPS method

In order to contract finite two-dimensional tensor networks, we use the boundary-MPS method [8, 30]. This method can be applied to single-layer tensor networks, as they appear in Eq. (3), as well as to double-layer network contraction, which is more costly, as we discuss below. We label the local tensors on a finite square lattice of size $L_x \times L_y$ by their positions in the lattice $T[x, y]$ and assume that they have an identical bond dimension D for each of their virtual indices. The entries in the local tensors are the parameters of the PEPS ansatz $\theta = \{T[1, 1], \dots, T[L_x, L_y]\}$. The central objects of this contraction algorithm are the boundary-MPS environments, which approximate the contraction of a set of rows of the network starting from a boundary. We will, for completeness, discuss here an example of an approximate contraction using boundary-MPS for the uppermost rows. The boundary-MPS environment can be recursively defined so that it approximates the contraction of the i uppermost rows of the tensor network as an MPS

of bond dimension D_c :

In the second row, the approximation indicates that we employ an MPS–MPO multiplication to truncate the bond dimension of the resulting MPS from $D_c D$ back to D_c . For this task, several algorithms have been developed in the past [49, 50]. Here, we employ the density matrix algorithm [50]. For the initial boundary MPS on the top, the uppermost row of local tensors is used. Of course, a boundary-MPS can be defined completely analogously for the approximate contraction for the lowermost rows up to row j , which we refer to as $E^l[j]$. These boundary-MPS can be combined to approximate the contraction of the network from both sides

$$\text{tTr} \left[\prod_{x,y} T[x,y] \right] \approx E^u[i] \cdot E^l[i+1], \forall i \in [1, L_x - 1]. \quad (16)$$

where $t\text{Tr}$ is the tensor trace over all indices not shown explicitly. As mentioned above, the boundary-MPS algorithm can be used both for single- as well as for double-layer tensor networks with different computational costs. Specifically, for the application in the double layer case the cost scales as $\mathcal{O}(D_c^3 D^4) + \mathcal{O}(d D_c^2 D^6)$ [30], while for a single layer the dominant scaling is $\mathcal{O}(D_c^3 D^3) + \mathcal{O}(D_c^2 D^4)$. Let us note that the bond dimension of the boundary-MPS D_c used in practical calculations should be increased whenever more accurate results are needed. Specifically, one can decide on D_c dynamically by setting a fixed value for the maximal truncated weight, cf. App. A for an analysis of the impact of different choices of the maximal truncated weight. Due to this analysis, in the following, a maximal truncated weight of 10^{-4} was used. We refer to Sec. IV for a discussion of the possible problems and usability of the boundary-MPS algorithm for different single-layer situations. Lastly, we mention that the boundary-MPS are also used for the calculations of gradients, which are necessary for optimization.

B. On reusing environments and the convenient treatment of long-range interactions

When using the boundary-MPS algorithm to compute the components of Eq. (3), it is often possible to reduce the computational cost by avoiding the repeated calculation of boundary-MPS environments. For local operators that act non-trivially only on a limited number of sites, the sum in

$$O_{\text{loc}}(\mathbf{S}) = \sum_{\mathbf{S}'} O_{\mathbf{S}'\mathbf{S}} \frac{\Psi(\mathbf{S}')}{\Psi(\mathbf{S})}, \quad (17)$$

is restricted to configurations \mathbf{S}' that differ from \mathbf{S} only by a few modified elements within the support of \hat{O} . Consequently, a significant computational advantage of PEPS, compared to neural quantum ansätze, lies in the ability to reuse precomputed environments for $\Psi(\mathbf{S})$ on lattice sites outside the support of \hat{O} .

If the operator \hat{O} acts non-locally, reusing boundary-MPS environments becomes less effective. However, if \hat{O} is diagonal in the computational basis, $O_{\text{loc}}(\mathbf{S})$ in Eq. (17) becomes independent of $\Psi(\mathbf{S})$. In this case, no additional computations are required. This property is particularly advantageous in Sec. V C, where it facilitates the treatment of long-range interactions in Rydberg atom arrays.

This approach can also be applied when the operator is not entirely diagonal. For instance, consider a Hamiltonian consisting of long-range Heisenberg interactions. The terms of this Hamiltonian can be divided into three sets: one diagonal in the X -basis, one in the Y -basis, and one in the Z -basis. To evaluate the expectation values of the operators in each set, one can transform the PEPS into the corresponding diagonal basis by a global spin rotation. In a diagonal basis, the calculations become straightforward. Note, that one has to sample independently for each of the transformed PEPS. Hence, this procedure becomes less effective if the Hamiltonian consists of more sets of operators that do not share a common diagonalizing basis.

C. On Numerical Stability

The method used to contract single-layer samples of finite PEPS, $\Psi(\mathbf{S})$, involves numerous MPS-MPO multiplications. To ensure numerical stability during these operations, each tensor $(E^{u/l}[i])_j$ of the environment is normalized. This is achieved by dividing the tensor by a factor f_{ij} such that its largest entry equals one. This normalization prevents the elements of the tensors from becoming out of range for machine precision under successive contractions. Alongside the normalized environment, we store $f_i = \sum_j \ln(f_{ij})$. The complete environment can then be reconstructed by multiplying by $\exp(f_i)$ at the end of the process.

Numerical stability is also critical when evaluating expressions of the form $\frac{\Psi(\mathbf{S}')}{\Psi(\mathbf{S})}$, which frequently appear in

Eq. (3). In this case, we improve stability by computing $\exp(\ln(\Psi(\mathbf{S}')) - \ln(\Psi(\mathbf{S})))$ instead, thereby avoiding potential numerical inaccuracies in direct division.

D. On Direct Sampling

The evaluation of Eq. (3) additionally requires the generation of samples that follow the probability distribution of the PEPS state $p_\Psi(\mathbf{S}) = \frac{|\Psi(\mathbf{S})|^2}{\langle \Psi | \Psi \rangle}$. We employ here a scheme that directly generates samples from the PEPS state vector [41]. This direct-sampling scheme has several substantial benefits with respect to Markov-chain-based ones [36]. Firstly, it avoids autocorrelation problems and the necessity for problem-dependent update schemes, as the samples are sampled independently of each other. Secondly, for the same reason, the sample generation can be parallelized.

It is worth noting that a double-layer tensor network needs to be contracted when employing the direct sampling scheme. However, this can be done with a small contract bond dimension $D_c \sim D$, while for expectation values with double-layers a larger contract bond dimension $D_c \sim D^2$ needs to be chosen. Thus, for small D the double-layer contractions in the direct sampling algorithm are not a dominant contributor to the total computational time due to the other computations necessary for finite PEPS optimizations. In the case of large D , the double-layer environments can be computed asynchronously. This will make them slightly outdated, but since these errors are corrected using importance sampling, their accuracy can be easily monitored by making sure that the statistical error for the energy is small enough. We further comment on this and summarize the direct-sampling scheme in App. B.

E. On Variational Upper Bound to the Ground State Energy

The estimator of Eq. (3) defines a variational upper bound for the ground state energy, provided that a consistent numerical value is assigned to $\Psi(\mathbf{S})$. However, for efficiency in numerical calculations (to reuse boundary-MPS), we may use $L - 1$ different approximations for $\Psi(\mathbf{S})$, depending on which boundary-MPS we use. We denote these different approximations $\Psi(\mathbf{S})_i := E^u[i] \cdot E^l[i + 1]$. Consequently, the energy expectation value obtained in this manner does not strictly serve as a variational upper bound.

Nevertheless, as long as these different approximations are very close to each other, we can expect to obtain a good approximation of the energy expectation value that is not far from the upper bound. It is possible to obtain a variational bound on the energy by restricting ourselves to a single approximation $\Psi(\mathbf{S})_i$ throughout the calculation of the expectation value, which is computationally significantly more demanding. To evaluate the accuracy

of the efficient, non-variational scheme, we compare the expectation value obtained in this way to the one obtained to give a variational bound. This method was also used in Ref. [47] to establish variational upper bounds for PEPS.

As an example, the variational energy for the ground state of the J_1 - J_2 model with $J_2 = 0.58$ and $L = 10$ was calculated using a bond dimension $D = 6$ PEPS and $\Psi(\mathbf{S})_{i=5}$, with a contraction cutoff of 10^{-4} (which translates to $D_c \approx 12$) and 10^5 samples. The same samples were used for both methods. The fast, inexact method yielded $\langle H \rangle = -187.856 \pm 0.036$, while the upper bound calculation produced $\langle H \rangle_{\text{bound}} = -187.805 \pm 0.027$.

To estimate the error between these two values, the elementwise difference of the sampled E^{loc} was computed as $\langle E_{\text{bound}}^{\text{loc}}(\mathbf{S}) - E^{\text{loc}}(\mathbf{S}) \rangle_{\mathbf{S} \in p_\Psi} = 0.051 \pm 0.013$. This result suggests that during optimization, the ansatz exploited variations in $\Psi(\mathbf{S})_i$ to lower the energy. However, the error is minor, affecting only the fifth significant digit. A more detailed analysis is provided in Appendix A.

IV. CONTRACTION COMPLEXITY ISSUES IN SAMPLING PEPS

A central mathematical task in the sampling-based schemes for finite PEPS is the contraction of single-layer tensor networks. These single-layer networks emerge in the calculation of the wave-function amplitude of a single many-body configuration $\Psi(\mathbf{S}) = \langle \mathbf{S} | \Psi \rangle$ as well as in the evaluation of local estimators $\hat{O}_{\text{loc}}(\mathbf{S}) = \langle \mathbf{S} | \hat{O} | \Psi \rangle / \Psi(\mathbf{S})$. However, it was recently discussed in the context of random tensor networks that the contraction of these single-layer tensor networks with the iterative MPS-MPO approach described in Sec. II, can be a task of drastically varying complexity [48].

In this section, we are going to analyze the situation first for random PEPS and then focus on the use case for finite PEPS calculations, which involve non-random tensors. We first find that the complexity transition found by Schuch *et al.* [48] corresponds to an entanglement transition of the random PEPS. We then demonstrate that the complexity of contracting the single-layer tensor networks that appear in non-random situations, such as the ground state search of a physical model, is generically less problematic than the ones encountered for random PEPS. We point out that, therefore, the naive random PEPS as a starting point of a PEPS optimization should be avoided.

Note that similar analyses have been done in the context of random quantum circuits [51].

A. Tools for the analysis of contraction complexity

1. Contraction complexity indicators

When contracting a two-dimensional single-layer tensor network, as necessary to obtain $\Psi(\mathbf{S})$, with the boundary-MPS method described in Sec. III A, we can define several quantities that are helpful in the analysis. Firstly, we can fix a maximal truncation error ϵ_{trunc} that we allow within the truncation of the MPS-MPO product. We define the boundary-MPS bond dimension $D_c(\epsilon_{\text{trunc}})$, which is necessary to achieve this truncation-accuracy ϵ_{trunc} in every step. This quantity can be used to indicate the computational difficulty of the single-layer contraction.

We further want to define a quantity that can illustrate the convergence of the scalar value ψ of the contraction of the single-layer tensor network as a function of the cutoff boundary-MPS bond dimension D_c . For this, we define

$$\Delta\psi(D_c, D_c^{\max}) = \frac{\psi(D_c^{\max}) - \psi(D_c)}{\psi(D_c^{\max})}, \quad (18)$$

where D_c^{\max} is chosen to be substantially larger than D_c ($D_c^{\max} \gg D_c$). This quantity $\Delta\psi(D_c, D_c^{\max})$, which we refer to as the *relative contraction error* of ψ , can be used to establish suitable contract bond dimension D_c , to be used in practical calculations. It can additionally be employed to visualize the convergence behavior of different single-layer tensor networks.

2. Geometric Entanglement

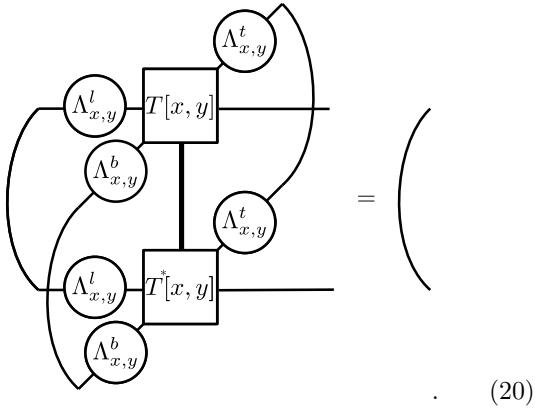
The geometric entanglement per site, denoted as S_G , quantifies the entanglement in a quantum many-body system on a per-site basis. It has previously been used in the MPS context as a tool to visualize phase transitions in [52]. It is defined using the maximum fidelity Λ_{\max} between a given pure normalized quantum state $|\psi\rangle$ and the set of all product states $|\phi\rangle$:

$$\Lambda_{\max} = \max_{|\phi\rangle} |\langle \phi | \psi \rangle|^2, \\ S_G(|\psi\rangle) = -\frac{1}{N} \log_2 \Lambda_{\max}. \quad (19)$$

This maximum fidelity measures how close the state $|\psi\rangle$ is to being fully separable, and the per-site geometric entanglement S_G provides a straightforward to compute measure of the entanglement in the quantum state which is zero in the case of a product state and one in the case of a maximally entangled state. An efficient scheme to compute the geometric entanglement is described in Sec. D.

3. Spectral properties in Vidal gauge

Generically, tensor networks enjoy so-called *gauge freedom*, from the fact that for every index of the tensor network that is contracted, we can insert a pair of invertible matrices $M \cdot M^{-1} = \mathbb{I}$ without changing the result of the contraction. In loop-free tensor networks, this gauge freedom plays a crucial role, but even in the context of PEPS a particular gauge choice, which we refer to as *Vidal gauge* is often used in the context of studies using the simple-update technique for optimization[53, 54]. This gauge choice can be achieved numerically on PEPS with the use of several algorithms [53, 55]. In the Vidal gauge the PEPS is represented in a form such that in addition to the local tensors, we have diagonal, non-negative matrices Λ on all bonds of the network. The tensors are gauged in such a way as to fulfill a *isometry condition* if the adjacent bond matrices are absorbed into the local tensors



This condition holds in all directions. Inspired by the fact that for an MPS in the Vidal gauge, the bond matrices carry the Schmidt coefficients [56] we choose to study the properties of the bond matrices of our PEPS in the Vidal gauge. Specifically, we will use the *average spectral entropy*

$$H_{\text{spec}} = \frac{1}{N_b} \sum_{i \in \text{bonds}} H_{\text{spec}}^i, \quad (21)$$

$$H_{\text{spec}}^i = \sum_k |S_k|^2 \log(|S_k|^2),$$

where S_k are the entries of the bond matrices of the PEPS in the Vidal gauge. We will see in the following sections that this average spectral entropy can be related to the difficulty of contracting the samples of a PEPS.

B. Random PEPS

We begin our analysis from random PEPS since their wave function amplitudes $\Psi(\mathbf{S})$ for any sample \mathbf{S} correspond to the random two-dimensional tensor networks

investigated in [48]. A central result of that study is that the complexity of the contraction of a Haar-random two-dimensional tensor network with the boundary-MPS method can be changed when shifting the mean value of the random tensors λ towards positive entries. At a mean value of $\lambda = 1/D$, where D denotes the bond dimension of the local tensors in the two-dimensional network, the complexity of the approximate contraction changes drastically. The change in complexity is measured e.g. with the bipartite entropy of the boundary MPS used for the contraction or equivalently the boundary MPS bond dimension necessary.

We are now investigating the properties of random PEPS to see if we can find a physical quantity that serves as a predictor for the complexity of contracting its samples. As a first step, we are going to again shift the mean of the entries of the random PEPS by λ and examine its entanglement via the geometric entanglement, cf. Sec. IV A 2. The results are shown in Fig. 2(a). We notice that at large shifts λ , the geometric entanglement becomes very small and the random PEPS represent essentially a product state. By extrapolation in system size, we show that in infinite systems this seems to be the case. If we now decrease the shift λ we notice an increase in the geometric entanglement. At the value of $\lambda = 1/D$, we find a transition in the geometric entanglement, when scaling the size of the system under consideration. At shift values below this transition, we find that the geometric entanglement plateaus to a value set by the system size. The value of the plateau converges to 1 (the maximum geometrical entanglement) as the system size grows.

This suggests that random PEPS go through an entanglement phase transition (as measured here by the geometric entanglement) as a function of the value of the mean-value shift λ which coincides with a complexity transition for the task of contracting the corresponding samples as established in [48].

This drastic change in the computational difficulty of contracting the random-single-layer networks suggests an important question. Is such a transition in the difficulty of the contraction a generic feature, when changing the tensors of the network? This is an important, practical consideration for sampling-based PEPS calculations because we will continuously manipulate the PEPS tensors during the ground state search. In this ground state search procedure, should one expect to generically hit a complexity phase transition, at which point the calculation of observables becomes very hard and the method stops working in practice?

We investigate this question in several steps. First, we are again considering random PEPS, which we now manipulate in a different way. Instead of shifting the mean value of the entries, as done previously, we multiply the virtual legs of every local tensor in the PEPS with diagonal matrices $S(\alpha)_{i,j} = \delta_{i,j} 1/i^\alpha$, whose entries decay as

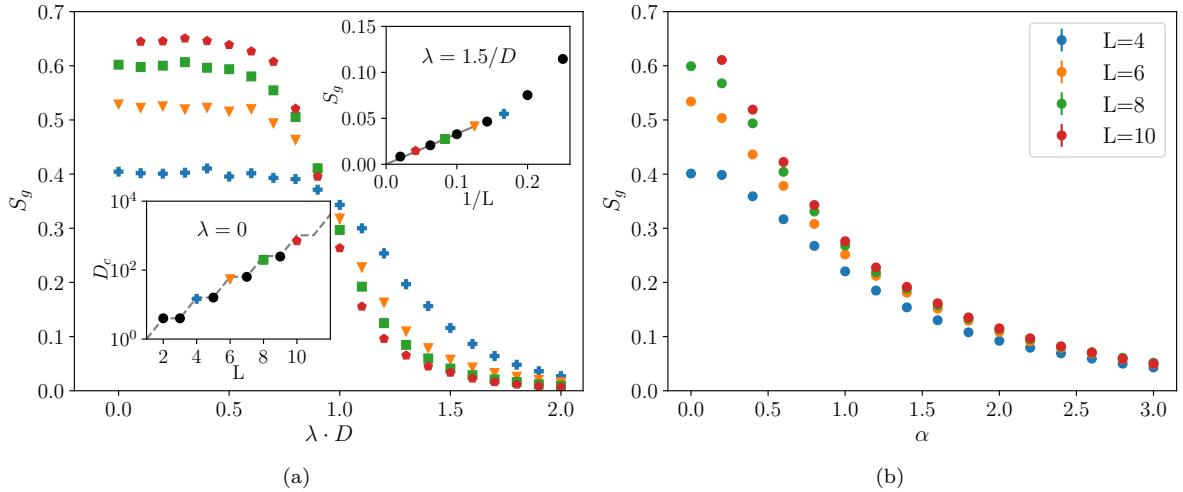


Figure 2. Geometrical entanglement S_g for random finite PEPS with bond dimension $D = 4$ and different system sizes for (a) a normally distributed tensor shifted by a constant λ and (b) a normally distributed tensor with algebraically decaying diagonal matrices $S_{i,j} = \delta_{i,j} 1/i^\alpha$ multiplied on each nonphysical link. The bottom left inset in (a) shows the dynamically determined contracted bond dimension D_c of the boundary MPS with a cutoff of 10^{-4} at different system sizes for $\lambda = 0$, where the dashed line represents the maximal possible bond dimension for an MPS of size L . The top right inset in (a) presents S_g versus system size, demonstrating that the geometrical entanglement approaches zero in the infinite system size limit, indicating that the PEPS collapses into a product state.

a function of α

$$T(\alpha) = (S(\alpha)^{\otimes 4})T ; T(\alpha) = -\begin{array}{c} \text{S}(\alpha) \\ \text{S}(\alpha) \end{array} \text{---} \boxed{T} \text{---} \begin{array}{c} \text{S}(\alpha) \\ \text{S}(\alpha) \end{array}. \quad (22)$$

We show the geometric entanglement of the resulting PEPS as a function of the decay-parameter α in Fig. 2(b). We observe a smooth increase of the geometric entanglement as the decay-parameter α decreases. This suggests a smooth entanglement-crossover as a function of α and stands in stark contrast to the entanglement transition shown in Fig. 2(a). We take this as first, tentative evidence that the drastic complexity transition observed for random tensors with a shifted mean λ is not generic.

Furthermore, we point out, that the contraction complexity of samples of a PEPS state vector cannot be entirely due to its entanglement properties. To illustrate this, we consider a PEPS $|\phi\rangle$, made up of a single-layer network to which we attach physical legs, without con-

necting them to the single layer.

$$|\phi\rangle = \begin{array}{c} \text{---} \text{---} \text{---} \text{---} \\ | \quad | \quad | \quad | \\ \text{---} \text{---} \text{---} \text{---} \end{array} \quad (23)$$

In this construction, every basis vector of the resulting many-body state vector has the same coefficient. Such a situation represents a product state, such that there is no entanglement for any bipartition. However, if we choose a single layer in the above construction that is hard to contract, every sample of this PEPS is hard to contract even though the PEPS represents a product state.

C. Towards physical PEPS

We now move to investigate PEPS that are obtained during the ground state search of generic physical models. We aim to find a way to characterize and predict the contraction complexity of the samples of the PEPS obtained in these scenarios.

To start we consider the ground state PEPS approximation obtained for the $J_1 - J_2$ model, cf. benchmarks in Sec. V. We will now describe a way to create from

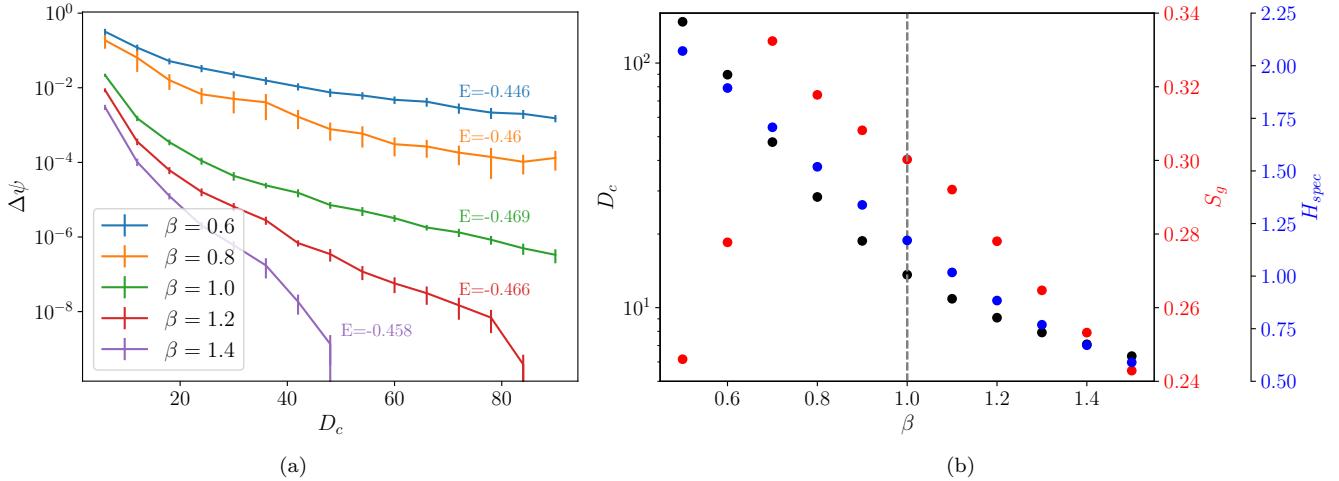


Figure 3. (a) The relative contraction error, $\Delta\psi(D_c, D_c^{\max} = 200)$ is plotted against the contraction dimension D_c for various values of β , with which the Vidal spectra were modified as $\{S\}_{\text{bonds}} \mapsto \{S^\beta\}_{\text{bonds}}$. The energy of the resulting states is also recorded. (b) The contraction dimension D_c , geometrical entanglement S_g , and average spectral entropy H_{spec} are plotted as functions of β . The results show that H_{spec} is correlated with D_c , whereas S_g does not exhibit a direct correlation with D_c .

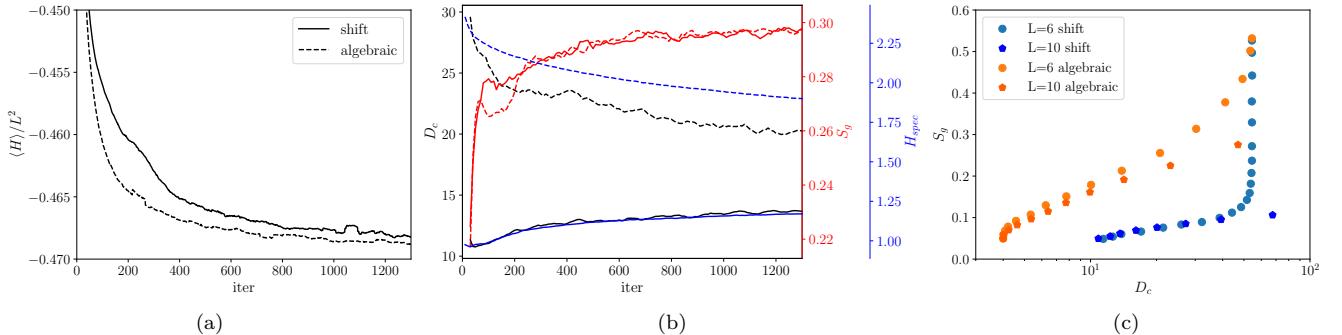


Figure 4. Different quantities calculated during the optimization of a PEPS initialized (solid line) with a shift $\lambda = 2$ and (dashed line) an algebraically decaying spectrum $\alpha = 1.5$. (a) Energy per site $\langle H \rangle / L^2$ versus optimization iterations. (b) Evolution of the contraction bond dimension D_c (black), geometrical entanglement S_g (red), and spectral entanglement H_{spec} (blue) during the optimization. (c) Contract environment Dimension D_c for randomly sampled finite PEPS that was either shifted or multiplied times an algebraic decaying spectra plotted against its geometrical entanglement.

this ground-state PEPS a family of low-energy PEPS, which is well suited for our study of the contraction complexity. We choose to represent the ground state PEPS in the Vidal gauge, cf. Sec. IV A 1, which allows us to generate from it a set of new PEPS by the following procedure. We can take the set of spectra on the bonds of the ground state PEPS in the Vidal gauge and take them to the power of β ,

$$\{S\}_{\text{bonds}} \mapsto \{S^\beta\}_{\text{bonds}}, \quad (24)$$

such that for every choice of β we obtain a unique new PEPS. For the family of PEPS that we have obtained in this way we now investigate the contraction complexity with the use of the relative contraction error $\Delta\psi(D_c, D_c^{\max})$. In Fig. 3(a) we show for several choices of β , that such a manipulation yields PEPS with different

contraction complexities for their samples. In fact it is quite clear from Fig. 3(a), that a manipulation that flattens the spectra ($\beta < 1$) yields a PEPS whose samples are more difficult to contract. Conversely, a steepening of the spectrum ($\beta > 1$) yields PEPS whose samples are cheaper to contract. This suggests that the spectrum in the Vidal gauge might be related to the contraction complexity of physical PEPS. Remarkably, all of the PEPS that we generate in this way are energetically close to the ground state, even if the contractability of their samples varies drastically.

For this reason, we choose this family of low-energy PEPS, parameterized by the single parameter β , and try to find a quantity that can predict how difficult it is to contract the samples of a given PEPS. In Fig. 3(b) we show several candidate quantities as a function of β

as well as the cutoff boundary-MPS bond dimension D_c which is necessary to maintain an error of $\epsilon_{\text{trunc}} = 10^{-4}$ during the contraction procedure. We first point out that the geometric entanglement does not seem to be a good predictor of the cost of the single-layer contractions, as might have been conjectured based on the investigations of random PEPS in the previous section. In fact for small β the cutoff bond dimension D_c increases, while the geometric entanglement decreases. This should further be taken as evidence that random PEPS do not generically share properties with physical PEPS. However, motivated by the fact that a manipulation of the Vidal-spectra of the PEPS seems to have a large impact on the contractability of their samples, cf. Fig. 3(a), we also consider the average spectral entropy H_{spec} of the PEPS at different β in the Vidal gauge. Fig. 3(b) shows that this quantity seems to be correlated with the difficulty of contraction.

In order to further substantiate the above finding that the spectral properties in the Vidal gauge are strongly correlated with the cost of contracting samples of a physical PEPS, we look at another family of PEPS, namely those obtained during the ground state search with imaginary time evolution, cf. Sec. II A. To this end, we have performed two optimization runs with different initial tensors. We have chosen a set of mean-shifted random tensors ($\lambda = 2$) as a starting PEPS, as well as a set of random PEPS that were manipulated according to Eq. (22) with $\alpha = 1.5$. Fig. 4 shows the results of these two ground-state searches. Fig. 4(a) shows the energy of the PEPS during the optimization. We find that both optimizations converge to a similar energy. However, the random tensors manipulated according to Eq. (22) serve as a more effective initial choice for the PEPS. This approach leads to convergence in fewer iteration steps. Additionally, each iteration step is faster to compute due to the lower contraction dimension D_c .

This is a direct consequence of the relation between geometrical entanglement S_g and contract bond dimension D_c for the shifted tensors. In Fig. 4(c) these two quantities are plotted against each other for different values of shift λ and algebraic decay α . One can see that to obtain a state distinct from the product state for the shifted initialization, a large contract bond dimension is needed, while for the algebraic decaying random PEPS, this is not the case.

More importantly, we find again that the average spectral entropy H_{spec} is strongly correlated with the contraction cost of the samples of the PEPS during the optimization. This is true for both choices of initial tensors. We close this section by concluding that, firstly, the complexity/entanglement phase transition as they can occur in random tensors-network/PEPS does not seem to be a generic feature of low energy PEPS of physical models, which can be considered as a validation of the applicability of the sampling PEPS methods. Secondly, we find that the spectral properties of the PEPS in the Vidal gauge are correlated with the contraction difficulty of its

samples.

V. APPLICATIONS

Now that we have set contraction complexity issues on firm(er) grounds, we proceed to use the finite PEPS framework discussed in Sec. II to investigate several challenging situations. We start by considering the well-studied $J_1 - J_2$ model [40, 57, 58] and use this as a starting point to compare the finite PEPS results to those of tree-tensor-network [3, 59–61] in the light of the different entanglement scaling of the two ansätze. We then move to an investigation of a Hamiltonian hosting a chiral spin liquid as its ground state; a notoriously difficult quantum state for PEPS [62, 63]. This has so far not been investigated with finite PEPS, but recently experienced advances with infinite PEPS optimized variationally with the help of automatic differentiation [21]. Finally, we show how powerful the described method can be for systems with long-range interactions by applying it to describe ground states of Rydberg atom arrays [64, 65], a very prominent platform for quantum simulation.

A. Entanglement scaling and comparison to Tree Tensor Networks: $J_1 - J_2$ -model

This section aims to demonstrate that finite PEPS optimized from scratch using sampling TDVP can represent the ground states of the J_1 - J_2 model on the square lattice

$$\hat{H} = J_1 \sum_{\langle i,j \rangle} \vec{S}_i \vec{S}_j + J_2 \sum_{\langle\langle i,j \rangle\rangle} \vec{S}_i \vec{S}_j. \quad (25)$$

Previous studies have extensively analyzed the ability of finite PEPS to capture the different phases of this model [38–40]. Here, we verify certain points of the established phase diagram and benchmark a specific point against TTNs (tree tensor networks) computed using [66].

We find the expected patterns of the local magnetization, as shown in Fig. 5(a). In the Néel antiferromagnetic phase, we find a checkerboard pattern, while at large values of J_2/J_1 we find a stripe-ordered pattern. In the intermediate region, we find vanishing local magnetization. Let us note that the magnitude of the local magnetization is smaller at the edges of the system due to the impact of a smaller number of neighbors and, hence, larger quantum fluctuations dressing the order.

To assess how the accuracy of finite PEPS scales with system size and bond dimension, we refer to Fig. 5(b). In this figure, the variance per site is plotted against the expectation value of the energy for both tree tensor networks and finite PEPS as is often done in the Neural Quantum state literature [67]. Linear extrapolation to the zero-variance limit shows the estimated energy value for trees and PEPS match. We note that to reach the

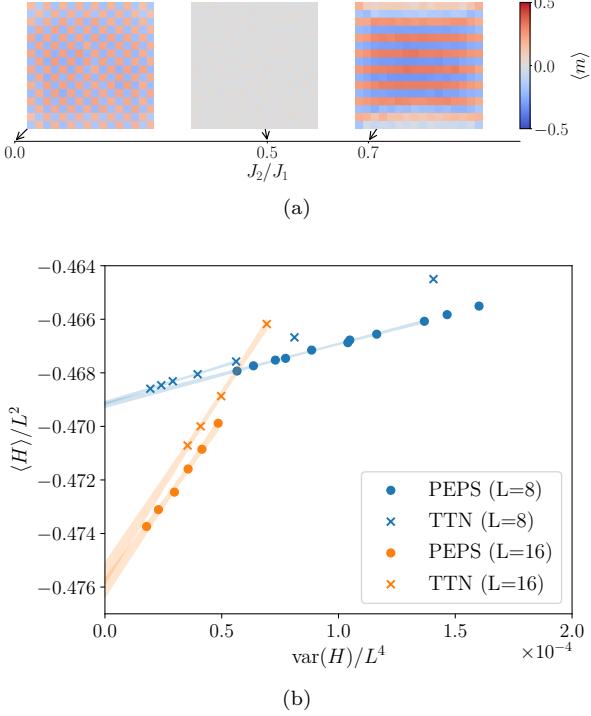


Figure 5. (a) Average magnetization per site for three distinct phases of the J_1 - J_2 phase diagram. (b) Variational energies plotted against energy variances at $J_2/J_1 = 0.58$ for system sizes $L = 8$ and $L = 16$. The data points shown as circles correspond to PEPS states with bond dimension $D = 8$, optimized using a cutoff of 10^{-4} , and subsequently evaluated with a more stringent cutoff of 10^{-6} . Each point represents a PEPS state at specific iterations during optimization: iterations $\{350, 410, 490, 650, 790, 870, 1250, 1710, 1950, 3000, 4000\}$ for $L = 8$, and $\{830, 1100, 1500, 2420, 4000, 9020\}$ for $L = 16$. A detailed analysis determining the appropriate cutoff value is provided in App. A. Data points represented by crosses indicate Tree Tensor Network (TTN) results with bond dimensions $D = \{50, 100, 150, 200, 250, 300, 350\}$ for $L = 8$ and $D = \{50, 100, 150, 200\}$ for $L = 16$. Dashed lines indicate linear fits employed for extrapolating energies to the zero-variance limit.

regime of truly linear relations between energy and variance, larger bond dimensions might need to be considered, especially for the tree tensor networks on larger systems.

PEPS fulfill the boundary law of entanglement entropy and are thus a good ansatz for any system size. In contrast, TTN, which does not fulfill the boundary law, requires a drastically increasing bond dimension to capture the same amount of entanglement as a PEPS at larger system sizes. However, they can still be an effective ansatz in smaller systems.

As expected, at a small system size of $L = 8$, TTNs with bond dimension 300 achieve 45% lower energies than PEPS with bond dimension 8, taking the extrapolated energy E_{\min} as a reference value $\frac{\langle H \rangle_{\text{TTN}300} - E_{\min}}{\langle H \rangle_{\text{PEPS}8} - E_{\min}} = 0.45$.

However, at a larger system size of $L = 16$, this trend reverses due to the more favorable entanglement scaling of the PEPS ansatz. In this case, the PEPS ansatz attains 39% lower energy, with $\frac{\langle H \rangle_{\text{PEPS}8} - E_{\min}}{\langle H \rangle_{\text{TTN}200} - E_{\min}} = 0.39$.

It is important to note that the bond dimension for the tree tensor networks with $L = 16$ was limited to $D = 200$ due to the memory constraints of the V100 GPU with 32 GB of memory.

B. Chiral spin liquid

Several recent experiments have made strides toward the realization of chiral topological many-body quantum states on quantum simulation platforms [68–70]. These pioneering works involve very few atoms and substantial efforts are being employed towards realizations with increasing numbers of constituents. For this effort, a faithful simulation of the experimental situation on classical computers - the generation of so-called *digital twins* - is crucial for benchmarking purposes. The finite PEPS are the natural candidate for this task, as they share the expected entanglement scaling of the chiral topological states [71], that are targeted in the experiments.

As a proof of principle, we examine the Hamiltonian proposed by Nielsen *et al.* [72],

$$\hat{H} = J_1 \sum_{\langle i,j \rangle} \vec{S}_i \vec{S}_j + J_2 \sum_{\langle\langle i,j \rangle\rangle} \vec{S}_i \vec{S}_j + i\lambda \sum_{\square} (P_{ijkl} - P_{ijkl}^{-1}), \quad (26)$$

which has a chiral spin liquid ground state at $J_1 = 2 \cos(0.06\pi) \cos(0.14\pi)$, $J_2 = 2 \cos(0.06\pi) \sin(0.14\pi)$ and $\lambda = 2 \sin(0.06\pi)$. The operator P_{ijkl} is defined for all plaquettes of the system and acts as cyclic permutations on the local Hilbert spaces. The corresponding terms in the Hamiltonian explicitly break time-reversal symmetry.

Recent works using an iPEPS ansatz have shown that in the thermodynamic limit, chiral topological ground states of local Hamiltonians [21, 25] can be successfully approximated using the iPEPS ansatz if variational optimization is employed [11–13, 73]. For finite system sizes, as relevant for the next generation of mesoscopic cold-atom experiments, however, such a numerical demonstration is to this date lacking. We show that with the approach for the finite PEPS detailed in this paper, we can indeed find chiral topological states as ground states.

To start, we show in Fig. 6(a) the extrapolation of the energy density against the variance, cf. Sec. V A, for a 16×16 square lattice with open boundary conditions. The estimated energy expectation value of $\langle H \rangle = -0.9697 \pm 0.0005$ is quite close to the result obtained for infinite systems using translationally invariant infinite PEPS where it is found numerically to be close to -1 [21]. One can get even closer to the expectation value found in the thermodynamic limit by extracting the energy expectation value for the bulk of the finite PEPS. For the terms of the Hamiltonian corresponding to the center 8×8 lattice sites of our 16×16 lattice, we obtain an energy of

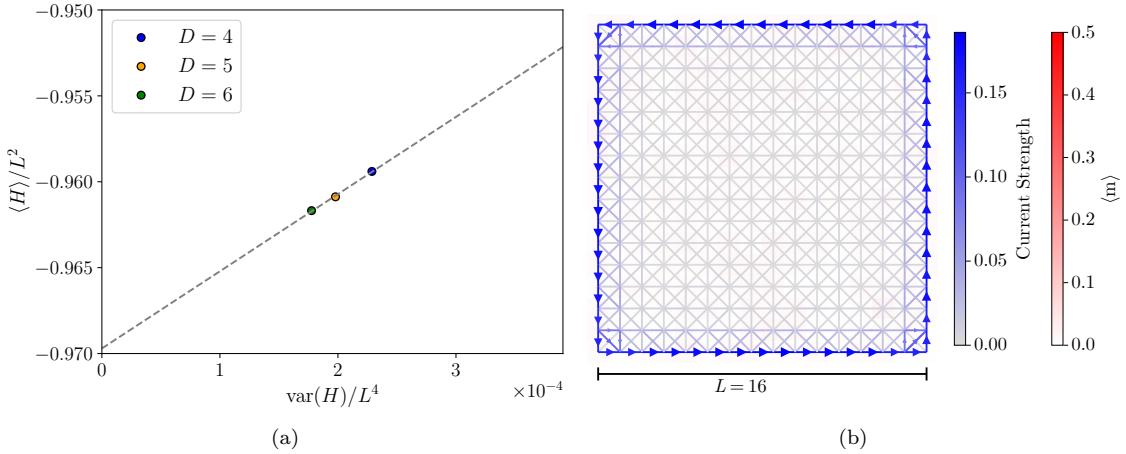


Figure 6. Results for the chiral spin liquid. (a) Energy extrapolation using the variance of the results for different bond dimensions. (b) Illustration of observables on a 16×16 lattice. On the vertices, we show the local magnetization, which is homogeneously very small ($m^2 \approx 7 \cdot 10^{-5}$). On the edges, we show the current. We find a current on the edge of the lattice with preferred chirality.

$E_{\text{bulk}} = -0.99335 \pm 0.00439$. Next, we focus on local observables that could act as local order parameters. We plot $m_i = |\langle \vec{S}_i \rangle|$ for all sites of the lattice in Fig. 6(b). We find that the local magnetization becomes very small $m^2 \approx 7 \cdot 10^{-5}$, suggesting vanishing magnetization in the state.

To briefly investigate the chiral nature of the state, we focus on its edge. For chiral topological states on finite systems with open boundary conditions, we expect chiral behavior at the edges, as demonstrated e.g. in [61, 74]. In [75], it was shown that for Mott insulators, loop currents around triangles can be shown to have the form

$$\mathcal{I}_{ij,k} \sim \frac{\mathbf{r}_{ij}}{|\mathbf{r}_{ij}|} \vec{S}_i \cdot (\vec{S}_j \times \vec{S}_k), \quad (27)$$

where $\mathcal{I}_{ij,k}$ denotes the current contribution of the loop current around the triangle consisting of sites i, j and k along the edge connecting site i and j . This quantity is promising for our investigation as the triple-product is invariant under $SU(2)$ spin rotations, which is required for its expectation value not to vanish for the chiral spin liquid. The Hamiltonian in Eq. (26) contains next-nearest neighbor spin interactions (which can be interpreted as perturbatively arising from next-nearest neighbor hopping terms in an underlying fermionic model). Thus, in order to investigate the contribution of the loop current in our quantum state, we sum over all triangles of our system, consisting of two pairs of nearest neighbors and one pair of next-nearest neighbors. The results are shown in Fig. 6(b). We find that the circulating currents in the bulk cancel out, such that no net current is found, while at the edges, such a cancellation does not happen, such that a chiral current around the edge emerges, as expected.

C. Rydberg atom arrays and long-range interactions

Long-range interactions are present in many different scenarios of interest in modern condensed matter- or atomic physics. Prominent examples include Coulomb-interactions ($U_C \sim \frac{1}{r}$) in ab initio electron systems, dipolar-interactions in cold quantum gases ($U_d \sim \frac{1}{r^3}$) [76] as well as van-der-Waals interaction in Rydberg atom arrays ($U_{vdW} \sim \frac{1}{r^6}$) [64, 65].

Traditionally, the treatment of such long-range interactions in the context of tensor-network methods (specifically in two dimensions) has been challenging in non-sampling-based approaches. This is because, for a system of n sites, the evaluation of every single term necessitates an evaluation of the order of n^2 terms, each one of which might be hard to evaluate depending on the specific method used. Therefore, in almost all cases, interactions have been truncated to a few neighbors. One can make substantial progress by employing projected entangled-pair operator methods [77, 78], which come at the cost of approximating the long-range interaction.

In stark contrast to these challenges, within the sampling-based approach used in this paper, the evaluation of long-range interactions becomes trivial, as long as the interaction is diagonal in the computational basis. As already briefly noted in Sec. III A, the evaluation of the expectation value of any operator \hat{O}_{diag} diagonal in the computational basis, requires no additional numerical calculations beyond the wave function amplitude for every sample

$$\langle \mathbf{S} | \hat{O}_{\text{diag}} | \Psi \rangle = O_{\mathbf{S}\mathbf{S}} \Psi(\mathbf{S}). \quad (28)$$

This is true irrespective of how local the operator in question acts, which is why we can treat these long-range

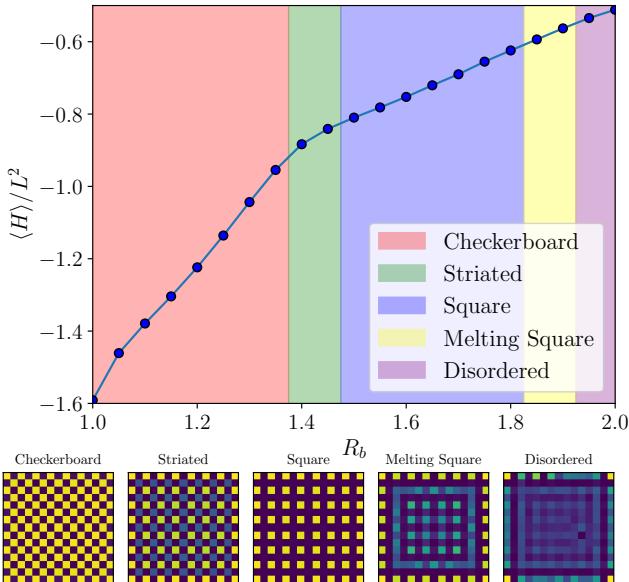


Figure 7. Results on the Rydberg atom arrays. We show the energy of the configurations we found for the fixed detuning $\Delta = 3$ and various values of R_b . We highlight the local Rydberg-density patterns for the different phases below.

interactions cheaply, which we exploit in the following. Additionally, the direct sampling procedure does not rely on a specific update scheme for generating the relevant samples of a state vector. The approach used in this paper is ideally suited for the investigation of models with long-range interactions in the finite PEPS framework.

To demonstrate this, and keeping with the theme of digital twins of current cold atom experiments, we are investigating the phases of a Rydberg atom array on a square lattice. Such a system can be described by

$$\hat{H} = \frac{1}{2} \sum_i \sigma_i^x - \Delta \sum_i n_i + \sum_{i \neq j} \frac{1}{((\mathbf{r}_i - \mathbf{r}_j)/R_b)^6} n_i n_j, \quad (29)$$

where the local Hilbert spaces on the square lattice are spanned by $\{|g\rangle, |r\rangle\}$, where $|g\rangle$ denotes a local atom in the ground state, while $|r\rangle$ indicates it being in an excited (Rydberg-)state. We use the Rabi-frequency Ω , the coupling constant of the first term, to fix the energy scale by setting it to unity together with the lattice distance.

This situation has recently been explored experimentally [65] as well as in numerical studies [33, 79]. One of the central conclusions of one of the numerical studies has been the importance of the treatment of long-range tails of the Rydberg interactions [33]. This finding strongly suggests a reconsideration of studies on Rydberg atom arrays on various lattice geometries in with methods that can treat these long-range interactions faithfully.

We investigate a slice of the phase diagram at fixed detuning $\Delta = 3$ whilst varying the Rydberg-blockade radius R_b . In Fig. 7, we show that several charge-density type phases, as well as disordered configurations, emerge.

At small R_b we find a checkerboard-state, that transitions into a striated phase ($R_b \approx 1.4 - 1.45$), in which the density of the checkerboard pattern is modulated. At $R_b = 1.5$ a homogeneous square phase emerges that eventually melts (via an intermediary phase) into a quantum disordered phase at $R_b = 1.9$. These results are consistent with the phase diagram obtained in [33]. In addition to the local Rydberg densities, we show the expectation values of the energy for the different ground states of the Rydberg atom array model. As expected, the slope of the energy density becomes flatter as the overall density of Rydberg excitations becomes smaller at a larger blockade radius R_b .

These results demonstrate that sampling-based finite PEPS calculation can be used comparatively easily to study particular models of long-range interacting Rydberg atom arrays. This makes it possible to investigate the phase diagram for these models on other lattice geometries [80] and aspect ratios as well as to study the influence of finite-size effects on the stabilization of phases. Particularly interesting for these investigations are the situations in which the Rydberg models are proposed to host topologically ordered ground states [81, 82].

VI. CONCLUSIONS AND DISCUSSION

In this work, we have discussed a framework for making use of sampling methods for finite PEPS calculations. To this end, we have pointed out how to solve the equations for the optimization more efficiently using an approach from the neural network community called minSR [46]. We further have pointed out that variational bounds can be obtained using the finite PEPS in the sampling approach - which we then proceeded to use to verify that the standard approximations made are highly accurate. We then moved on to discuss complexity issues that might arise in this approach and have given a tentative quantity of the physical PEPS, that is correlated with the contraction complexity of its samples and have introduced a new initialization strategy based on these considerations. Finally, we applied the finite PEPS in the context of a chiral spin liquid and of long-range interacting Rydberg atom arrays. We have pointed out that certain long-range interactions can be treated trivially in the sampling PEPS approach.

This demonstration on Rydberg atom arrays opens the door to the accurate numerical study of a multitude of related situations from frustrated geometries and topological order in Rydberg atoms [81, 82], to phases of dipolar gases in cold atom platforms [76], and Nitrogen-vacancy centers [83, 84]. One crucial application of this is the generation of digital twins of state-of-the-art quantum simulation platforms, which plays an important role in benchmarking these powerful experiments.

We note that a technique recently proposed in the context of iPEPS, for treating *bra*- and *ket*-layer separately [85] might be useful to reduce computational cost

in the double-layer boundary-MPS during the direct sampling procedure.

An important work for the future is an extensive comparison between the different methods, like different flavours of PEPS, NQS [44], TTNs [3, 59, 60], augmented TTNs [86], etc., for ground states calculations of two dimensional quantum systems.

Another area worth exploring further is how the single-layer contractability relates to the entanglement in the state represented by the finite PEPS. We suspect that given an entanglement entropy S , there exists a finite PEPS that will minimize the hardness of contracting its samples.

Algorithm and open source code. An implementation of the algorithms discussed here is available as open source libraries [87, 88]. Parts of them are making use of the ITensor library [89, 90]. All scripts used to generate the data for this analysis, along with the resulting simulation data, are available on Zenodo [91].

VII. ACKNOWLEDGEMENTS

The authors thank Niklas Tausendpfund for generating the tree-tensor-network data used for comparison with the package TTN.jl [66].

This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) via Project-ID 277101999 – CRC network TRR 183 (“Entangled states of matter”) and under Germany’s Excellence Strategy – Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC 2004/1 – 390534769, and by the Horizon Europe programme HORIZON-CL4-2023-DIGITAL-EMERGING-01-CNECT via the project 101135699 (SPINUS) and Horizon Europe programme HORIZON-CL4-2022-QUANTUM-02-SGA via the project 101113690 (PASQuanS2.1). E. L. W. thanks the Studienstiftung des deutschen Volkes for support. D.A. acknowledge funding by the German Federal Ministry of Education and Research (BMBF) for support under the thematic programme “Quantum technologies – from the basics to the market”, project number 13N16202 “Noise in Quantum Algorithms (NiQ)”.

The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS [92] (Grant NeTeNeSyQuMa) and the FZ Jülich for computing time on JURECA [93] (institute project PGI-8) at Jülich Supercomputing Centre (JSC).

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

[2] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Rev. Mod. Phys.* **93**, 045003 (2021).

[3] P. Silvi, F. Tschirsich, M. Gerster, J. Jünemann, D. Jaschke, M. Rizzi, and S. Montangero, The Tensor Networks Anthology: Simulation techniques for many-body quantum lattice systems, *SciPost Phys. Lect. Notes*, 8 (2019).

[4] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).

[5] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Annals of Physics* **326**, 96 (2011), january 2011 Special Issue.

[6] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Rev. Mod. Phys.* **82**, 277 (2010).

[7] N. Laflorencie, Quantum entanglement in condensed matter systems, *Physics Reports* **646**, 1 (2016), quantum entanglement in condensed matter systems.

[8] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions (2004), arXiv:cond-mat/0407066 [cond-mat.str-el].

[9] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical simulation of infinite-size quantum lattice systems in two spatial dimensions, *Phys. Rev. Lett.* **101**, 250602 (2008).

[10] P. Corboz, T. M. Rice, and M. Troyer, Competing states in the t - j model: Uniform d -wave state versus stripe state, *Phys. Rev. Lett.* **113**, 046402 (2014).

[11] P. Corboz, Variational optimization with infinite projected entangled-pair states, *Phys. Rev. B* **94**, 035133 (2016).

[12] L. Vanderstraeten, J. Haegeman, P. Corboz, and F. Verstraete, Gradient methods for variational optimization of projected entangled-pair states, *Phys. Rev. B* **94**, 155123 (2016).

[13] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, Differentiable programming tensor networks, *Phys. Rev. X* **9**, 031041 (2019).

[14] H. J. Liao, Z. Y. Xie, J. Chen, Z. Y. Liu, H. D. Xie, R. Z. Huang, B. Normand, and T. Xiang, Gapless spin-liquid ground state in the $s = 1/2$ kagome antiferromagnet, *Phys. Rev. Lett.* **118**, 137202 (2017).

[15] I. Niesen and P. Corboz, Emergent haldane phase in the $s = 1$ bilinear-biquadratic heisenberg model on the square lattice, *Phys. Rev. B* **95**, 180404 (2017).

[16] J.-Y. Chen, L. Vanderstraeten, S. Capponi, and D. Poilblanc, Non-abelian chiral spin liquid in a quantum antiferromagnet revealed by an ipeps study, *Phys. Rev. B* **98**, 184409 (2018).

[17] S. S. Chung and P. Corboz, Su(3) fermions on the honeycomb lattice at $\frac{1}{3}$ filling, *Phys. Rev. B* **100**, 035134 (2019).

[18] B. Ponsioen, S. S. Chung, and P. Corboz, Period 4 stripe in the extended two-dimensional hubbard model, *Phys. Rev. B* **100**, 195141 (2019).

[19] H.-Y. Lee, R. Kaneko, L. E. Chern, T. Okubo, Y. Yamaji,

N. Kawashima, and Y. B. Kim, Magnetic field induced quantum phases in a tensor network study of kitaev magnets, *Nature communications* **11**, 1639 (2020).

[20] O. Gauthé, S. Capponi, M. Mambrini, and D. Poilblanc, Quantum spin liquid phases in the bilinear-biquadratic two-su(4)-fermion hamiltonian on the square lattice, *Phys. Rev. B* **101**, 205144 (2020).

[21] J. Hasik, M. Van Damme, D. Poilblanc, and L. Vanderstraeten, Simulating chiral spin liquids with projected entangled-pair states, *Phys. Rev. Lett.* **129**, 177201 (2022).

[22] M. Peschke, B. Ponsioen, and P. Corboz, Competing states in the two-dimensional frustrated kondo-necklace model, *Phys. Rev. B* **106**, 205140 (2022).

[23] B. Ponsioen, S. S. Chung, and P. Corboz, Superconducting stripes in the hole-doped three-band hubbard model, *Phys. Rev. B* **108**, 205154 (2023).

[24] B. Ponsioen, S. S. Chung, and P. Corboz, Superconducting stripes in the hole-doped three-band hubbard model, *Phys. Rev. B* **108**, 205154 (2023).

[25] E. L. Weerda and M. Rizzi, Fractional quantum hall states with variational projected entangled-pair states: A study of the bosonic harper-hofstadter model, *Phys. Rev. B* **109**, L241117 (2024).

[26] J. Hasik and P. Corboz, Incommensurate order with translationally invariant projected entangled-pair states: Spiral states and quantum spin liquid on the anisotropic triangular lattice, *Phys. Rev. Lett.* **133**, 176502 (2024).

[27] P. Schmoll, J. Naumann, J. Eisert, and Y. Iqbal, Bathing in a sea of candidate quantum spin liquids: From the gapless ruby to the gapped maple-leaf lattice (2024), arXiv:2407.07145 [cond-mat.str-el].

[28] V. Murg, F. Verstraete, and J. I. Cirac, Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states, *Phys. Rev. A* **75**, 033605 (2007).

[29] V. Murg, F. Verstraete, and J. I. Cirac, Exploring frustrated spin systems using projected entangled pair states, *Phys. Rev. B* **79**, 195119 (2009).

[30] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Unifying projected entangled pair state contractions, *New Journal of Physics* **16**, 033014 (2014).

[31] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Algorithms for finite projected entangled pair states, *Phys. Rev. B* **90**, 064425 (2014).

[32] M. Scheb and R. M. Noack, Finite projected entangled pair states for the hubbard model, *Phys. Rev. B* **107**, 165112 (2023).

[33] M. J. O'Rourke and G. K.-L. Chan, Entanglement in the quantum phases of an unfrustrated rydberg atom array, *Nature Communications* **14**, 5397 (2023).

[34] J. Gray, quimb: A python package for quantum information and many-body calculations, *Journal of Open Source Software* **3**, 819 (2018).

[35] L. Wang, I. Pižorn, and F. Verstraete, Monte carlo simulation with tensor network states, *Phys. Rev. B* **83**, 134421 (2011).

[36] W.-Y. Liu, S.-J. Dong, Y.-J. Han, G.-C. Guo, and L. He, Gradient optimization of finite projected entangled pair states, *Phys. Rev. B* **95**, 195154 (2017).

[37] W.-Y. Liu, H. Zhai, R. Peng, Z.-C. Gu, and G. K. Chan, Accurate simulation of the hubbard model with finite fermionic projected entangled pair states, arXiv preprint arXiv:2502.13454 (2025).

[38] W.-Y. Liu, S. Dong, C. Wang, Y. Han, H. An, G.-C. Guo, and L. He, Gapless spin liquid ground state of the spin- $\frac{1}{2}$ $J_1 - J_2$ heisenberg model on square lattices, *Phys. Rev. B* **98**, 241109 (2018).

[39] W.-Y. Liu, Y.-Z. Huang, S.-S. Gong, and Z.-C. Gu, Accurate simulation for finite projected entangled pair states in two dimensions, *Phys. Rev. B* **103**, 235155 (2021).

[40] W.-Y. Liu, S.-S. Gong, Y.-B. Li, D. Poilblanc, W.-Q. Chen, and Z.-C. Gu, Gapless quantum spin liquid and global phase diagram of the spin-1/2 j_1 - j_2 square antiferromagnetic heisenberg model, *Science Bulletin* **67**, 1034 (2022).

[41] T. Vieijra, J. Haegeman, F. Verstraete, and L. Vanderstraeten, Direct sampling of projected entangled-pair states, *Phys. Rev. B* **104**, 235141 (2021).

[42] S. Sorella, M. Casula, and D. Rocca, Weak binding between two aromatic rings: Feeling the van der waals attraction by quantum monte carlo methods, *The Journal of chemical physics* **127**, 014105 (2007).

[43] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-dependent variational principle for quantum lattices, *Physical review letters* **107**, 070601 (2011).

[44] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).

[45] L. Vanderstraeten, J. Haegeman, and F. Verstraete, Tangent-space methods for uniform matrix product states, *SciPost Phys. Lect. Notes* **7** (2019).

[46] A. Chen and M. Heyl, Empowering deep neural quantum states through efficient optimization, *Nature Physics* **20**, 1476 (2024).

[47] W.-Y. Liu, S.-J. Du, R. Peng, J. Gray, and G. K.-L. Chan, Tensor network computations that capture strict variationality, volume law behavior, and the efficient representation of neural network states, *Physical Review Letters* **133**, 260404 (2024).

[48] J. Chen, J. Jiang, D. Hangleiter, and N. Schuch, Sign problem in tensor network contraction (2024), arXiv:2404.19023 [quant-ph].

[49] E. M. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, *New Journal of Physics* **12**, 055026 (2010).

[50] I. P. McCulloch, From density-matrix renormalization group to matrix product states, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P10014 (2007).

[51] M. McGinley, W. W. Ho, and D. Malz, Measurement-induced entanglement and complexity in random constant-depth 2d quantum circuits (2024), arXiv:2410.23248.

[52] R. Orús and T.-C. Wei, Visualizing elusive phase transitions with geometric entanglement, *Physical Review B* **82**, 155120 (2010).

[53] S.-J. Ran, W. Li, B. Xi, Z. Zhang, and G. Su, Optimized decimation of tensor networks with super-orthogonalization for two-dimensional quantum lattice models, *Phys. Rev. B* **86**, 134429 (2012).

[54] S. S. Jahromi and R. Orús, Universal tensor-network algorithm for any infinite lattice, *Phys. Rev. B* **99**, 195105 (2019).

[55] J. Tindall and M. Fishman, Gauging tensor networks with belief propagation, *SciPost Phys.* **15**, 222 (2023).

[56] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, *Phys. Rev. Lett.* **91**, 147902

(2003).

[57] P. Chandra and B. Doucot, Possible spin-liquid state at large s for the frustrated square heisenberg lattice, *Phys. Rev. B* **38**, 9335 (1988).

[58] E. Dagotto and A. Moreo, Phase diagram of the frustrated spin-1/2 heisenberg antiferromagnet in 2 dimensions, *Phys. Rev. Lett.* **63**, 2148 (1989).

[59] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).

[60] L. Tagliacozzo, G. Evenbly, and G. Vidal, Simulation of two-dimensional quantum systems using a tree tensor network that exploits the entropic area law, *Phys. Rev. B* **80**, 235127 (2009).

[61] M. Gerster, M. Rizzi, P. Silvi, M. Dalmonte, and S. Montangero, Fractional quantum hall effect in the interacting hofstadter model via tensor networks, *Phys. Rev. B* **96**, 195123 (2017).

[62] T. B. Wahl, H.-H. Tu, N. Schuch, and J. I. Cirac, Projected entangled-pair states can describe chiral topological states, *Phys. Rev. Lett.* **111**, 236805 (2013).

[63] J. Dubail and N. Read, Tensor network trial states for chiral topological phases in two dimensions and a no-go theorem in any dimension, *Phys. Rev. B* **92**, 205307 (2015).

[64] H. Labuhn, D. Barredo, S. Ravets, S. De Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, Tunable two-dimensional arrays of single rydberg atoms for realizing quantum ising models, *Nature* **534**, 667 (2016).

[65] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature* **595**, 227 (2021).

[66] N. Tausendpfund, W. Krinitzin, M. Schmitt, and M. Rizzi, TTN.jl – A tree tensor network library for calculating groundstates and solving time evolution (0.1), Zenodo.

[67] Y. Nomura, Helping restricted boltzmann machines with quantum-state representation by restoring symmetry, *Journal of Physics: Condensed Matter* **33**, 174003 (2021).

[68] L. W. Clark, N. Schine, C. Baum, N. Jia, and J. Simon, Observation of Laughlin states made of light, *Nature* **582**, 41 (2020).

[69] J. Léonard, S. Kim, J. Kwan, P. Segura, F. Grusdt, C. Repellin, N. Goldman, and M. Greiner, Realization of a fractional quantum hall state with ultracold atoms, *Nature* **619**, 495 (2023).

[70] P. Lunt, P. Hill, J. Reiter, P. M. Preiss, M. Gałka, and S. Jochim, Realization of a Laughlin state of two rapidly rotating fermions, *Phys. Rev. Lett.* **133**, 253401 (2024).

[71] M. Haque, O. Zozulya, and K. Schoutens, Entanglement entropy in fermionic Laughlin states, *Phys. Rev. Lett.* **98**, 060401 (2007).

[72] A. E. Nielsen, G. Sierra, and J. I. Cirac, Local models of fractional quantum hall states in lattices and physical implementation, *Nature Communications* **4**, 2864 (2013).

[73] J. Naumann, E. L. Weerda, M. Rizzi, J. Eisert, and P. Schmoll, An introduction to infinite projected entangled-pair state methods for variational ground state simulations using automatic differentiation, *SciPost Phys. Lect. Notes*, 86 (2024).

[74] S. Banerjee, W. Zhu, and S.-Z. Lin, Electromagnetic signatures of a chiral quantum spin liquid, *npj Quantum Materials* **8**, 63 (2023).

[75] L. N. Bulaevskii, C. D. Batista, M. V. Mostovoy, and D. I. Khomskii, Electronic orbital currents and polarization in mott insulators, *Phys. Rev. B* **78**, 024402 (2008).

[76] T. Lahaye, C. Menotti, L. Santos, M. Lewenstein, and T. Pfau, The physics of dipolar bosonic quantum gases, *Reports on Progress in Physics* **72**, 126401 (2009).

[77] M. J. O'Rourke, Z. Li, and G. K.-L. Chan, Efficient representation of long-range interactions in tensor network algorithms, *Phys. Rev. B* **98**, 205127 (2018).

[78] M. J. O'Rourke and G. K.-L. Chan, Simplified and improved approach to tensor network operators in two dimensions, *Phys. Rev. B* **101**, 205142 (2020).

[79] R. Samajdar, W. W. Ho, H. Pichler, M. D. Lukin, and S. Sachdev, Complex density wave orders and quantum phase transitions in a model of square-lattice rydberg atom arrays, *Phys. Rev. Lett.* **124**, 103601 (2020).

[80] R. Samajdar, W. W. Ho, H. Pichler, M. D. Lukin, and S. Sachdev, Quantum phases of rydberg atoms on a kagome lattice, *Proceedings of the National Academy of Sciences* **118**, e2015785118 (2021).

[81] R. Verresen, M. D. Lukin, and A. Vishwanath, Prediction of toric code topological order from rydberg blockade, *Phys. Rev. X* **11**, 031005 (2021).

[82] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, *Science* **374**, 1242 (2021).

[83] J. Randall, C. Bradley, F. Van Der Gronden, A. Galicia, M. Abobeih, M. Markham, D. Twitchen, F. Machado, N. Yao, and T. Taminiau, Many-body-localized discrete time crystal with a programmable spin-based quantum simulator, *Science* **374**, 1474 (2021).

[84] J. Cai, A. Retzker, F. Jelezko, and M. B. Plenio, A large-scale quantum simulator on a diamond surface at room temperature, *Nature Physics* **9**, 168 (2013).

[85] J. Naumann, E. L. Weerda, J. Eisert, M. Rizzi, and P. Schmoll, Variationally optimizing infinite projected entangled-pair states at large bond dimensions: A split-ctmrg approach (2025), arXiv:2502.10298.

[86] T. Felser, S. Notarnicola, and S. Montangero, Efficient tensor network ansatz for high-dimensional quantum many-body problems, *Phys. Rev. Lett.* **126**, 170603 (2021).

[87] D. Alcalde Puente, E. L. Weerda, and K. Schröder, Quantum natural finite peps, Will be available after publication and before that, upon request. (2025).

[88] D. Alcalde Puente, Quantum natural gradient, <https://github.com/danielalcalde/QuantumNaturalGradient.jl> (2025).

[89] M. Fishman, S. R. White, and E. M. Stoudenmire, The ITensor Software Library for Tensor Network Calculations, *SciPost Phys. Codebases*, 4 (2022).

[90] M. Fishman, S. R. White, and E. M. Stoudenmire, Codebase release 0.3 for ITensor, *SciPost Phys. Codebases*, 4 (2022).

[91] D. Alcalde Puente, E. L. Weerda, K. Schröder, and M. Rizzi, Data for "efficient optimization and conceptual barriers in variational finite projected entangled-pair states", 10.5281/zenodo.15046369 (2025).

[92] Jülich Supercomputing Centre, JUWELS Cluster and Booster: Exascale pathfinder with modular supercom-

puting architecture at JSC, *J. Large-Scale Rec. Fac.* **7**, A183 (2021).

[93] Jülich Supercomputing Centre, JURECA: Data centric and booster modules implementing the modular supercomputing architecture at JSC, *J. Large-Scale Res. Fac.* **7**, A182 (2021).

[94] A. J. Ferris and G. Vidal, Perfect sampling with unitary tensor networks, *Phys. Rev. B* **85**, 165146 (2012).

[95] W. Krinitzin, N. Tausendpfund, M. Rizzi, M. Heyl, and M. Schmitt, Roughening dynamics of interfaces in two-dimensional quantum matter (2024), arXiv:2412.10145 [quant-ph].

[96] M. Schmitt and M. Heyl, Quantum many-body dynamics in two dimensions with artificial neural networks, *Physical Review Letters* **125**, 100503 (2020).

Appendix A: Sampling-based PEPS schemes can yield a variational upper bound on the energy

To further investigate the effect of different approximations of the wave function amplitude $\Psi(\mathbf{S})$ on the variational upper bound of the ground state energy, we compare the energy estimates obtained using two different computational approaches. As described in the main text, the faster method reuses the environments computed during boundary-MPS contraction to evaluate the local estimator $E_{\text{loc}}(\mathbf{S})$, whereas the slower, variational approach enforces a consistent wave function approximation, ensuring a strict upper bound for the ground state energy.

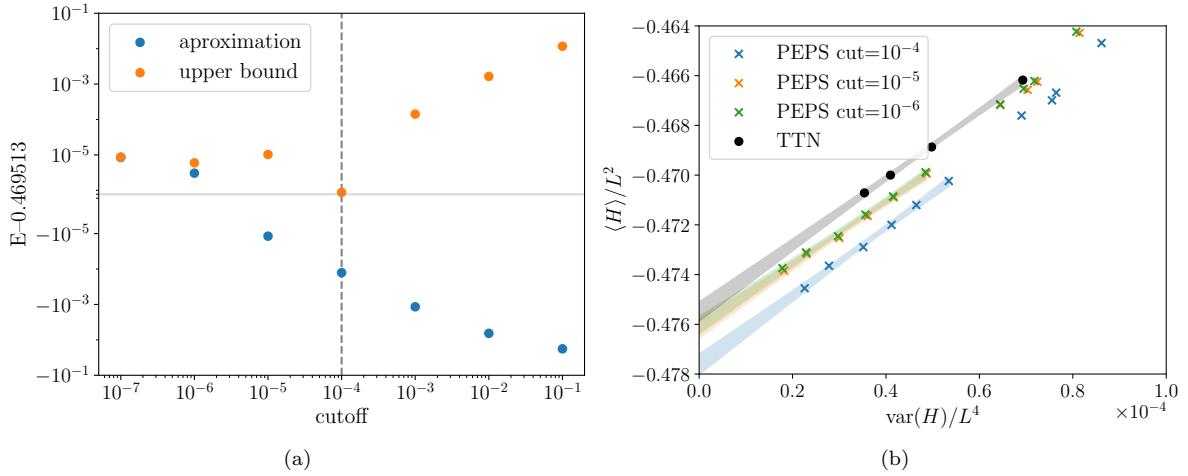


Figure 8. (a) For a set of 10^4 samples at $L = 10$ the energy was computed for different cutoffs in two different ways: (blue) the faster way, which takes advantage of the environments to compute the $\Psi(\mathbf{S})$ needed to compute $E_{\text{loc}}(\mathbf{S})$, and (orange) the slower way which always uses the same environments to compute an upper bound for the ground state energy. Note that the same samples were used for all calculations. The lowest variational energy value was used as the offset for the plot. (b) The energy vs variance plot for $L = 16$ is displayed in Fig. 5(b) but calculated to a different accuracy using the different cutoffs for 10^5 samples using the fast approximate method.

Figure 8(a) presents the results for a set of 10^4 samples, showing the energy deviation from the lowest obtained variational energy as a function of the contraction cutoff. For very small cutoffs (10^{-7} to 10^{-5}), both methods yield nearly identical energy estimates, suggesting that the approximation error is negligible in this regime. As the cutoff increases, deviations increase polynomially in the cutoff, with the fast method producing slightly lower energy estimates than the upper-bound method. This indicates that the ansatz indeed exploits inconsistencies in the wave function approximations to achieve artificially lower energy values.

Notably, at a cutoff of 10^{-4} , the same value used during optimization, the strict variational calculation achieves its lowest value. Since $\langle S | \Psi(\text{cutoff}) \rangle = E^u[i](\text{cutoff}) \cdot E^l[i+1](\text{cutoff})$ was indirectly optimized using the cut environments E it makes sense that its energy would perform best.

These results highlight the importance of choosing an appropriate contraction cutoff: while excessively tight cutoffs increase computational cost without significant accuracy gains, too loose a cutoff compromises the variational nature of the ansatz. The optimal choice depends on balancing these factors to ensure efficient yet reliable energy estimates.

An alternative approach for estimating an upper bound to the ground state energy involves a different, unbiased estimator that is computationally less demanding. Specifically, the method proposed at the end of Sec. III B can be utilized. In this approach, the expectation values of the Hamiltonian terms are evaluated on a rotated basis where the Hamiltonian terms become diagonal. As a result, the computation of $E_{\text{loc}}(\mathbf{S})$ becomes independent of the wave function, eliminating the need for different approximations of $\Psi(S)$, and the errors that come with them.

This estimator is unbiased because the finite PEPS was not optimized using this method, ensuring that the optimization process could not exploit discrepancies between the different wave functions. The unbiased estimator yields $\langle H \rangle = -187.907 \pm 0.042$, which is in agreement with the other estimates. This consistency suggests that the approximations employed are sufficiently accurate to prevent significant deviations between different approximation schemes. Consequently, the use of additional methods appears unnecessary unless there is a reason to question the validity of the primary method.

More practical than generating expensive variational upper bounds is to vary the cutoff and see if energy and variance converge. For example, the PEPS states presented in Fig. 8(b) were optimized using a cutoff value of 10^{-4} . However, evaluating the energies obtained with this cutoff reveals discrepancies with the Tree Tensor Network (TTN) results, indicating inaccurate ground-state energies. This discrepancy disappears once the accuracy is increased. This observation underscores the necessity for meticulous care in PEPS calculations. In particular, ensuring sufficient accuracy in environment approximations is crucial to prevent the ansatz from exploiting numerical errors introduced during truncation.

Appendix B: Direct sampling of PEPS

The direct-sampling scheme generates a many-body configuration \mathbf{S} with a probability $p(\mathbf{S})$, which serves as an approximation of $p_\Psi(\mathbf{S}) = \frac{|\Psi(\mathbf{S})|^2}{\langle \Psi | \Psi \rangle}$. The discrepancy between $p(\mathbf{S})$ and $p_\Psi(\mathbf{S})$ can be corrected using importance sampling, as will be discussed later.

To proceed, we introduce the shorthand notation $\mathbf{S}_i := \{s_{i,1}, \dots, s_{i,L_y}\}$ to represent the collection of local configurations on the i -th row. This allows us to express the probability of a configuration as

$$p(\mathbf{S}) = p(\mathbf{S}_1) \prod_{i=1}^{L_y-1} p(\mathbf{S}_{i+1} | \mathbf{S}_{<i+1}), \quad (\text{B1})$$

where $p(\mathbf{S}_i | \mathbf{S}_{<i})$ represents the conditional probability given the configurations on the uppermost $i-1$ rows, denoted by $\mathbf{S}_{<i} = \{\mathbf{S}_1, \dots, \mathbf{S}_{i-1}\}$.

We generate each of the probabilities in Eq. (B1) sequentially, beginning with $p(\mathbf{S}_1)$. This probability corresponds to the reduced density matrix of the topmost row of the PEPS, denoted as $\rho[1]$. As this step involves tracing out all degrees of freedom except those on the first row, an approximate contraction of the double-layer tensor PEPS network for all rows except the first is required.

To achieve this, we define $D^l[i]$ as the double-layer contraction of the i lowermost rows, computed using the boundary-MPS method described above (cf. Fig. 9). In practice, we truncate the bond dimension of $D^l[i]$ to D_c^{double} . Additional comments on this double-layer boundary can be found at the end of this section.

An approximation of the reduced density matrix for the first row can thus be expressed as

$$\begin{aligned} \rho[1] &= T^u[1] \cdot D^d[2] \cdot (T^u)^*[1] \\ &= \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} T[1, 1] \quad T[1, 2] \quad \dots \quad T[1, L_y] \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \\ &= \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} D^l[2]_1 \quad D^l[2]_2 \quad \dots \quad D^l[2]_{L_y} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \\ &= \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} T^*[1, 1] \quad T^*[1, 2] \quad \dots \quad T^*[1, L_y] \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \end{aligned} \quad (\text{B2})$$

where $T^u[i]$ represents the product of all local PEPS tensors in the i -th uppermost row, as illustrated in the equation above.

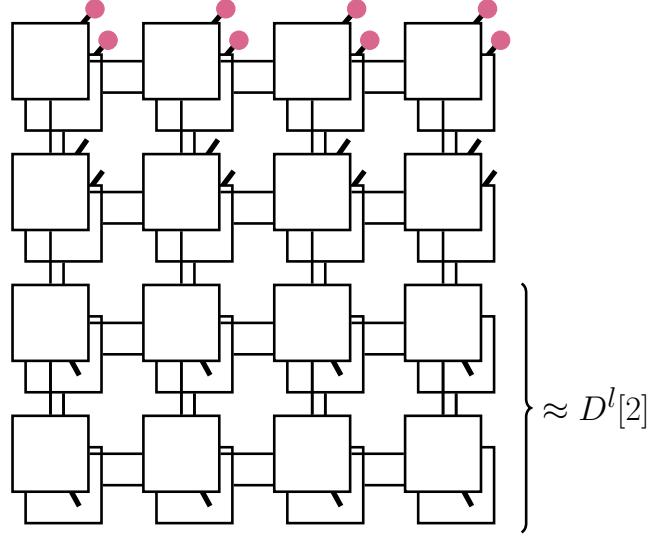


Figure 9. Illustration of the conditional reduced density matrix for the second row $\rho_{S_1}[2]$ used in the direct sampling scheme. The spin configurations of the first row are fixed to configurations S_1 , which is illustrated by the colored dots. The physical legs of the second row are left open (on the support of $\rho_{S_1}[2]$), while all physical legs on the lower rows are traced out. The contraction of the two lowermost rows, on which the physical indices are traced out, are in our calculation approximated by a boundary-MPS $D^l[2]$.

It is noteworthy that this reduced-density matrix possesses a one-dimensional structure. Consequently, efficient methods for sampling MPS wavefunctions can be directly applied to generate samples S_1 from the probability distribution $p(S_1)$ [94]. This established MPS technique resembles the direct-sampling approach for PEPS discussed here, as the probability $p(S_1)$ is similarly expressed as a product of conditional probabilities, analogous to Eq. (B1), which are evaluated sequentially from left to right (or vice versa).

We now proceed to the generation of the conditional probabilities $p(S_i|S_{<i})$, which are derived from the conditional reduced density matrix $\rho_{S_{<i}}[i]$ of the i -th row, as illustrated in Fig. 9. The support of $\rho_{S_{<i}}[i]$ consists of the degrees of freedom on the i -th row, while all degrees of freedom on the rows above are fixed to the configurations $S_{<i}$.

To construct this conditional reduced density matrix, we combine the i uppermost rows using the boundary-MPS method, resulting in

$$\begin{aligned}
 T_{S_{<i}}^u[i] &:= \left(T_{S_{<i}}[i]_1 \xrightarrow{} T_{S_{<i}}[i]_2 \xrightarrow{} \dots \xrightarrow{} T_{S_{<i}}[i]_{L_y} \right) \\
 &\approx \left(\begin{array}{c} T[1, 1] \xrightarrow{} T[1, 2] \xrightarrow{} \dots \xrightarrow{} T[1, L_y] \\ \vdots \quad \vdots \quad \vdots \quad \vdots \\ T[i-1, 1] \xrightarrow{} T[i-1, 2] \xrightarrow{} \dots \xrightarrow{} T[i-1, L_y] \\ \vdots \quad \vdots \quad \vdots \quad \vdots \\ T[i, 1] \xrightarrow{} T[i, 2] \xrightarrow{} \dots \xrightarrow{} T[i, L_y] \end{array} \right), \tag{B3}
 \end{aligned}$$

where the bond dimension is truncated to D_s to ensure computational efficiency. Note that in practice, D_s can be chosen to be significantly smaller than D_c . All truncations are performed using the algorithm described in [50].

All rows below the i -th one are traced out, which is represented using the double-layer boundary-MPS $D^l[L_x - i]$. This approach allows for an efficient and stable computation of the conditional probabilities required for the sampling procedure. With these objects, we can express the conditional reduced density matrix $\rho_{\mathbf{S}_{*}}[i]*$ as

$$\rho_{\mathbf{S}_{*}}[i] = T_{\mathbf{S}_{*}}^u[i] \cdot D^l[i+1] \cdot (T_{\mathbf{S}_{*}}^u)^*[i], \quad (\text{B4})***$$

which has the same form as Eq. (B2) such that we can again use the sampling algorithms for MPS to obtain the conditional probability $p(\mathbf{S}_i | \mathbf{S}_{*})*$ from it. This is summarized in *Alg. 1*.

Algorithm 1 Sampling algorithm

```

function SAMPLE( $T, D^l; D_c, D_s$ )
   $E^u[0] \leftarrow 1$ 
  for  $i = 1, \dots, L$  do
     $T_{\mathbf{S}_{}}^u[i] \leftarrow \text{mul}(T[i, :], E^u[i-1]; D_s)$ 
     $S_i \leftarrow \text{sample}(T_{\mathbf{S}_{}}^u[i], D^l[i+1])$ 
     $T_{\text{proj}}[i] \leftarrow \text{proj}(T[i, :], S_i)$ 
     $E^u[i] \leftarrow \text{mul}(E^u[i-1], T_{\text{proj}}[i]; D_c)$ 
  end for
  return  $\mathbf{S}, E^u$ 
end function

```

We close this summary of the direct sampling procedure with a few comments. Firstly, the direct sampling procedure involves the calculation of boundary-MPS approximation for double layers of PEPS. However, as was noted in [41], we can get away with taking small values for the environment bond dimension $D_c^{\text{double}} \sim D$ of this double-layer boundary-MPS. Additionally, these double-layer environments have to be calculated only once and can then be reused to generate an arbitrary number of samples for the corresponding PEPS. Due to this fact, the calculation of the double-layer boundary-MPS for small enough D only accounts for a small fraction of the computational time used.

Secondly, as mentioned at the beginning of this section, the fact that we obtain any sample \mathbf{S} with a probability $p(\mathbf{S})$ which is an approximation of the actual probability $p_\Psi(\mathbf{S})$ can be corrected for with an additional factor of $\frac{p_\Psi(\mathbf{S})}{p(\mathbf{S})}$ in Eq. (3). Note that once we have a sample \mathbf{S} we can obtain $\Psi(\mathbf{S})$ needed for $p_\Psi(\mathbf{S})$ accurately with a single layer contraction.

The computation of double-layer environments becomes increasingly demanding for larger bond bond dimensions, taking up a sizable share of the computational time for $D \geq 7$. This issue can be mitigated by computing the double-layer environments asynchronously. This approach generates samples using previously computed double-layer environments while updated environments are simultaneously computed. Although asynchronous generation introduces slight inaccuracies due to outdated environments, these deviations are corrected through importance sampling. Moreover, the magnitude of these errors can be continuously monitored by evaluating the statistical error of the energy, ensuring that statistical biases remain within acceptable limits. In practice for a $L = 16$ $D = 8$ PEPS with 2000 samples, the double layer environments will lag 5 optimization steps behind while not causing any significant change in the error metrics.

Appendix C: Real-Time Evolution

The equations introduced in Sec. II A can, in principle, be applied to real-time evolution. However, two major challenges prevent achieving high-fidelity results in practice.

The first challenge arises from the typically linear growth of entanglement entropy during real-time evolution. PEPS inherently obey an area-law scaling of entanglement entropy and are limited to modest bond dimensions, generally ranging from 1 to 10, due to computational constraints. Consequently, accurate simulations over long time intervals are usually infeasible. Exceptions occur in special cases where entanglement entropy grows sublinearly, such as the domain-wall dynamics recently examined in Ref. [95].

The second challenge relates to the number of samples required for accurately solving the TDVP equations. In optimization scenarios, the fidelity of individual evolution steps is less critical, provided convergence to the correct state is eventually achieved. In contrast, real-time evolution is highly sensitive to fidelity losses, which accumulate exponentially over multiple steps. Ensuring high-fidelity evolution steps thus requires a substantial number of samples, often several orders of magnitude larger than the number of variational parameters. Schmidt *et al.* [96], for instance,

used around 10^6 samples for an ansatz with roughly 10^3 parameters. A PEPS with moderate system size ($L = 10$) and bond dimension ($D = 6$) contains substantially more parameters ($N_p \approx 10^5$). Therefore, meeting the sampling requirements in practical scenarios is expected to be a significant computational challenge.

Appendix D: Efficient computation of geometric entanglement

Several methods exist to compute the geometric entanglement for finite PEPS. In this study, a two-step approach is employed. The first step provides an approximate product state that maximizes the overlap

$$\begin{aligned}\Lambda_{\max} &= \max_{|\phi\rangle} |\langle\phi|\Psi\rangle|^2, \\ S_G(|\psi\rangle) &= -\frac{1}{N} \log_2 \Lambda_{\max},\end{aligned}\tag{D1}$$

while the second step refines this initial approximation.

To efficiently approximate the optimal state $|\phi\rangle = \bigotimes_{i=1}^{L^2} |\phi_i\rangle$, a sampling-based strategy is utilized. Rather than sampling directly from $|\Psi\rangle$ in the z -basis, the same algorithm is used, but now the spins are sequentially optimized to maximize the overlap with the conditional reduced density matrix. This selection is performed iteratively by minimizing

$$\min_{|\phi_j\rangle} \langle\Psi| \left(\bigotimes_{i=1}^j |\phi_i\rangle\langle\phi_i| \right) |\Psi\rangle. \tag{D2}$$

After obtaining this initial approximation, the overlap $|\langle\phi|\Psi\rangle|^2$ can be further optimized through a sweep across all $|\phi_i\rangle$. During this sweep, each $|\phi_i\rangle$ is individually optimized by maximizing its overlap, allowing for efficient reuse of previously computed environments.

CHAPTER 7

Conclusion

This thesis has investigated how quantum measurements can be utilized for quantum state preparation and ground state search. By exploiting non-unitary dynamics and incorporating adaptive feedback, the developed methods address key challenges associated with the preparation of complex quantum states, particularly within the constraints of noisy intermediate-scale quantum (NISQ) devices. These contributions enhance the prospects for scalable quantum state preparation by providing new algorithmic tools for quantum information processing.

This chapter presents a summary of the main contributions of the thesis and outlines several directions for future research.

A central contribution, presented in Publication [I], is the development and theoretical analysis of a dissipative protocol for preparing the one-dimensional Affleck-Kennedy-Lieb-Tasaki (AKLT) state. By incorporating engineered ancilla measurement and reset, the protocol operates beyond the idealized Markovian regime. It identifies an optimal finite measurement and reset interval that balances entanglement generation and convergence rate. The protocol also exhibits robustness against realistic dephasing noise, indicating its potential for experimental implementation.

Another key result, detailed in Publication [II], is the introduction of a self-learning framework that integrates projective measurements and adaptive classical feedback into variational quantum circuits (VQCs). This approach autonomously discovers efficient constant-depth preparation protocols for specific AKLT edge states, including new deterministic solutions. The study also identifies a new class of local minima specific to feedback-based VQCs and proposes strategies to mitigate the challenges they present to optimization.

Publication [III] introduces methodological advances for ground state search using finite PEPS by analyzing the complexity of performing projective measurements on both

random and physically motivated quantum states. This analysis informs the design of improved initialization strategies for variational optimization. In particular, the spectral entropy H_c is proposed as a diagnostic tool to quantify the difficulty of tensor contractions. This thesis continues the analysis in Sec. 2.2.1, demonstrating that insights into H_c enable targeted acceleration of single-layer contractions required for sampling.

The results presented in this thesis suggest several promising directions for future research. The dissipative protocol introduced in Publication [I], initially developed for Hamiltonian dynamics, could be adapted to gate-based architectures. Additionally, it may benefit from optimization techniques such as the variational quantum eigensolver (VQE). A natural extension of this work involves the preparation of two-dimensional states using the finite PEPS ansatz from Publication [III].

Particular attention should be given to topological states such as the toric code. These states could potentially be prepared using both passive protocols, as developed in Publication [I], and active feedback-based methods described in Publication [II]. The passive protocol is expected to prepare the desired state accurately, although it lacks control over the encoded logical qubits. In contrast, the feedback-based approach may offer a strategy for accelerating the selection of a specific logical qubit state, similar to how it enabled the deterministic selection of an edge mode in the AKLT case.

However, the preparation of topological states presents a greater challenge. In the AKLT case, the presence of short-range entanglement simplifies the selection of edge modes. For topological states, where long-range entanglement is inherent, such selection is expected to be significantly more difficult. Furthermore, the applicability of shallow-circuit active learning techniques in two dimensions remains an open question due to the optimization difficulties already observed in one-dimensional systems.

Another avenue for future research is the real-time evolution of finite PEPS for the characterization of out-of-equilibrium dynamics. In Sec. 3.3, we showed that the number of samples needed to compute accurate gradients scales linearly with the number of parameters in the ansatz as well as with the entanglement of the quantum state. It remains to be clarified whether this unfavorable scaling with entanglement is specific to sampling-based PEPS or also affects Neural Quantum States (NQS), where time evolution is implemented in a similar manner and has demonstrated significant success in simulating out-of-equilibrium dynamics in two dimensions [89]. It would be important to carry out a comparable analysis for NQS to determine whether they exhibit the same sensitivity to entanglement. This would clarify whether the observed scaling is an inherent feature of the underlying poly-linear structure of tensor networks or a more general limitation of variational sampling approaches.

A further avenue for future research is the proposed hybrid TEBD-sampling approach

for real-time PEPS evolution (Sec. 3.3.2), as it could greatly reduce the sampling overhead. In addition, translationally invariant PEPS offer a promising ansatz for homogeneous systems and may reduce computational cost. Lastly, future research should aim at a systematic comparison between various approaches for ground state search in two-dimensional quantum systems. These include different variants of PEPS, neural quantum states (NQS) [70], tree tensor networks (TTNs) [143–145], and augmented TTNs [146]. A clear understanding of the optimal choice of method for different Hamiltonians remains unclear and should be investigated.

In summary, this thesis has introduced new protocols for quantum state preparation, algorithmic improvements for PEPS optimization, and conceptual insights into the role of measurements and feedback in variational quantum algorithms. Together, these contributions provide a foundation for continued progress toward the efficient preparation and simulation of complex quantum systems.

Bibliography

- [1] John Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (Aug. 2018), p. 79. ISSN: 2521-327X. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79). URL: <https://doi.org/10.22331/q-2018-08-06-79>.
- [2] Jens Eisert. “Entangling power and quantum circuit complexity”. In: *Physical Review Letters* 127.2 (2021), p. 020501. URL: <https://doi.org/10.1103/physrevlett.127.020501>.
- [3] Tzu-Chieh Wei. *Measurement-Based Quantum Computation*. Mar. 2021. DOI: [10.1093/acrefore/9780190871994.013.31](https://doi.org/10.1093/acrefore/9780190871994.013.31). URL: <https://oxfordre.com/physics/view/10.1093/acrefore/9780190871994.001.0001/acrefore-9780190871994-e-31>.
- [4] Youwei Zhao et al. “Realization of an error-correcting surface code with superconducting qubits”. In: *Physical Review Letters* 129.3 (2022), p. 030501.
- [5] Eric Chitambar and Gilad Gour. “Quantum resource theories”. In: *Reviews of modern physics* 91.2 (2019), p. 025001.
- [6] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf. “Sequential Generation of Entangled Multiqubit States”. In: *Phys. Rev. Lett.* 95 (11 Sept. 2005), p. 110503. DOI: [10.1103/PhysRevLett.95.110503](https://doi.org/10.1103/PhysRevLett.95.110503). URL: <https://doi.org/10.1103/PhysRevLett.95.110503>.
- [7] Marco Cerezo et al. “Variational quantum algorithms”. In: *Nature Reviews Physics* 3.9 (2021), pp. 625–644. URL: <https://doi.org/10.1038/s42254-021-00348-9>.
- [8] Eric R Anschuetz and Bobak T Kiani. “Quantum variational algorithms are swamped with traps”. In: *Nature Communications* 13.1 (2022), p. 7760. URL: <https://doi.org/10.1038/s41467-022-35364-5>.
- [9] Tameem Albash and Daniel A. Lidar. “Adiabatic quantum computation”. In: *Rev. Mod. Phys.* 90 (1 Jan. 2018), p. 015002. DOI: [10.1103/RevModPhys.90.015002](https://doi.org/10.1103/RevModPhys.90.015002). URL: <https://doi.org/10.1103/RevModPhys.90.015002>.
- [10] Pimonpan Sompet et al. “Realizing the symmetry-protected Haldane phase in Fermi–Hubbard ladders”. In: *Nature* (2022), pp. 1–5. URL: <https://doi.org/10.1038/s41586-022-04688-z>.

[11] Zhi-Yuan Wei, Daniel Malz, and J. Ignacio Cirac. “Efficient adiabatic preparation of tensor network states”. In: *Physical Review Research* 5.2 (May 2023). ISSN: 2643-1564. DOI: [10.1103/physrevresearch.5.1022037](https://doi.org/10.1103/physrevresearch.5.1022037). URL: <https://doi.org/10.1103/physrevresearch.5.1022037>.

[12] Barbara Kraus et al. “Preparation of entangled states by quantum Markov processes”. In: *Physical Review A—Atomic, Molecular, and Optical Physics* 78.4 (2008), p. 042307.

[13] Sebastian Diehl et al. “Quantum states and phases in driven open quantum systems with cold atoms”. In: *Nature Physics* 4.11 (2008), pp. 878–883.

[14] Frank Verstraete, Michael M Wolf, and J Ignacio Cirac. “Quantum computation and quantum-state engineering driven by dissipation”. In: *Nature physics* 5.9 (2009), pp. 633–636.

[15] Anne Matthies, Mark Rudner, Achim Rosch, and Erez Berg. “Programmable adiabatic demagnetization for systems with trivial and topological excitations”. In: *Quantum* 8 (2024), p. 1505.

[16] Hefeng Wang. “Quantum algorithm for preparing the ground state of a system via resonance transition”. In: *Scientific Reports* 7.1 (2017), p. 16342.

[17] Jia-Jin Feng, Biao Wu, and Frank Wilczek. “Quantum computing by coherent cooling”. In: *Physical Review A* 105.5 (2022), p. 052601.

[18] Stefano Polla, Yaroslav Herasymenko, and Thomas E O’Brien. “Quantum digital cooling”. In: *Physical Review A* 104.1 (2021), p. 012414.

[19] Leo Zhou, Soonwon Choi, and Mikhail D Lukin. “Symmetry-protected dissipative preparation of matrix product states”. In: *Physical Review A* 104.3 (2021), p. 032418.

[20] Yunzhao Wang, Kyrylo Snizhko, Alessandro Romito, Yuval Gefen, and Kater Murch. “Dissipative preparation and stabilization of many-body quantum states in a superconducting qutrit array”. In: *Physical Review A* 108.1 (2023), p. 013712.

[21] Sthitadhi Roy, JT Chalker, IV Gornyi, and Yuval Gefen. “Measurement-induced steering of quantum systems”. In: *Physical Review Research* 2.3 (2020), p. 033347. URL: <https://doi.org/10.1103/physrevresearch.2.033347>.

[22] Nathanan Tantivasadakarn, Ryan Thorngren, Ashvin Vishwanath, and Ruben Verresen. “Long-Range Entanglement from Measuring Symmetry-Protected Topological Phases”. In: *Phys. Rev. X* 14 (2 June 2024), p. 021040. DOI: [10.1103/PhysRevX.14.021040](https://doi.org/10.1103/PhysRevX.14.021040). URL: <https://doi.org/10.1103/PhysRevX.14.021040>.

[23] Tsung-Cheng Lu, Leonardo A. Lessa, Isaac H. Kim, and Timothy H. Hsieh. “Measurement as a Shortcut to Long-Range Entangled Quantum Matter”. In: *PRX Quantum* 3 (4 Dec. 2022), p. 040337. DOI: [10.1103/PRXQuantum.3.040337](https://doi.org/10.1103/PRXQuantum.3.040337). URL: <https://doi.org/10.1103/PRXQuantum.3.040337>.

[24] Ruben Verresen, Nathanan Tantivasadakarn, and Ashvin Vishwanath. “Efficiently preparing Schrödinger’s cat, fractons and non-Abelian topological order in quan-

tum devices,(2021)”. In: *arXiv preprint arXiv:2112.03061* (2021). URL: <https://doi.org/10.48550/arXiv.2112.03061>.

[25] Lorenzo Piroli, Georgios Styliaris, and J Ignacio Cirac. “Quantum circuits assisted by local operations and classical communication: Transformations and phases of matter”. In: *Physical Review Letters* 127.22 (2021), p. 220503. URL: <https://doi.org/10.1103/physrevlett.127.220503>.

[26] Nathanan Tantivasadakarn, Ruben Verresen, and Ashvin Vishwanath. “Shortest Route to Non-Abelian Topological Order on a Quantum Processor”. In: *Phys. Rev. Lett.* 131 (6 Aug. 2023), p. 060405. DOI: [10.1103/PhysRevLett.131.060405](https://doi.org/10.1103/PhysRevLett.131.060405). URL: <https://doi.org/10.1103/PhysRevLett.131.060405>.

[27] Kevin C. Smith, Eleanor Crane, Nathan Wiebe, and S.M. Girvin. “Deterministic Constant-Depth Preparation of the AKLT State on a Quantum Processor Using Fusion Measurements”. In: *PRX Quantum* 4 (2 Apr. 2023), p. 020315. DOI: [10.1103/PRXQuantum.4.020315](https://doi.org/10.1103/PRXQuantum.4.020315). URL: <https://doi.org/10.1103/PRXQuantum.4.020315>.

[28] David T. Stephen and Oliver Hart. “Preparing matrix product states via fusion: Constraints and extensions”. In: *Phys. Rev. B* (Apr. 2025), pp. –. DOI: [10.1103/cv3q-518w](https://doi.org/10.1103/cv3q-518w). URL: <https://doi.org/10.1103/cv3q-518w>.

[29] Laurens Lootens, Clement Delcamp, Dominic Williamson, and Frank Verstraete. “Low-Depth Unitary Quantum Circuits for Dualities in One-Dimensional Quantum Lattice Models”. In: *Phys. Rev. Lett.* 134 (13 Apr. 2025), p. 130403. DOI: [10.1103/PhysRevLett.134.130403](https://doi.org/10.1103/PhysRevLett.134.130403). URL: <https://doi.org/10.1103/PhysRevLett.134.130403>.

[30] Rahul Sahay and Ruben Verresen. “Classifying One-Dimensional Quantum States Prepared by a Single Round of Measurements”. In: *PRX Quantum* 6 (1 Feb. 2025), p. 010329. DOI: [10.1103/PRXQuantum.6.010329](https://doi.org/10.1103/PRXQuantum.6.010329). URL: <https://doi.org/10.1103/PRXQuantum.6.010329>.

[31] Kevin C. Smith, Abid Khan, Bryan K. Clark, S.M. Girvin, and Tzu-Chieh Wei. “Constant-Depth Preparation of Matrix Product States with Adaptive Quantum Circuits”. In: *PRX Quantum* 5 (3 Sept. 2024), p. 030344. DOI: [10.1103/PRXQuantum.5.030344](https://doi.org/10.1103/PRXQuantum.5.030344). URL: <https://doi.org/10.1103/PRXQuantum.5.030344>.

[32] Daniel Malz, Georgios Styliaris, Zhi-Yuan Wei, and J Ignacio Cirac. “Preparation of matrix product states with log-depth quantum circuits”. In: *Physical Review Letters* 132.4 (2024), p. 040404.

[33] Lorenzo Piroli, Georgios Styliaris, and J. Ignacio Cirac. “Approximating Many-Body Quantum States with Quantum Circuits and Measurements”. In: *Phys. Rev. Lett.* 133 (23 Dec. 2024), p. 230401. DOI: [10.1103/PhysRevLett.133.230401](https://doi.org/10.1103/PhysRevLett.133.230401). URL: <https://doi.org/10.1103/PhysRevLett.133.230401>.

[34] Yaroslav Herasymenko, Igor Gornyi, and Yuval Gefen. “Measurement-Driven Navigation in Many-Body Hilbert Space: Active-Decision Steering”. In: *PRX Quantum* 4 (2 June 2023), p. 020347. DOI: [10.1103/PRXQuantum.4.020347](https://doi.org/10.1103/PRXQuantum.4.020347). URL: <https://doi.org/10.1103/PRXQuantum.4.020347>.

[35] Samuel Morales, Yuval Gefen, Igor Gornyi, Alex Zazunov, and Reinhold Egger. “Engineering unsteerable quantum states with active feedback”. In: *Phys. Rev. Res.* 6 (1 Mar. 2024), p. 013244. DOI: [10.1103/PhysRevResearch.6.013244](https://doi.org/10.1103/PhysRevResearch.6.013244). URL: <https://doi.org/10.1103/PhysRevResearch.6.013244>.

[36] Zhikang T. Wang, Yuto Ashida, and Masahito Ueda. “Deep Reinforcement Learning Control of Quantum Cartpoles”. In: *Phys. Rev. Lett.* 125 (10 Sept. 2020), p. 100401. DOI: [10.1103/PhysRevLett.125.100401](https://doi.org/10.1103/PhysRevLett.125.100401). URL: <https://doi.org/10.1103/PhysRevLett.125.100401>.

[37] Sangkha Borah, Bijita Sarma, Michael Kewming, Gerard J. Milburn, and Jason Twamley. “Measurement-Based Feedback Quantum Control with Deep Reinforcement Learning for a Double-Well Nonlinear Potential”. In: *Phys. Rev. Lett.* 127 (19 Nov. 2021), p. 190403. DOI: [10.1103/PhysRevLett.127.190403](https://doi.org/10.1103/PhysRevLett.127.190403). URL: <https://doi.org/10.1103/PhysRevLett.127.190403>.

[38] V. V. Sivak et al. “Model-Free Quantum Control with Reinforcement Learning”. In: *Phys. Rev. X* 12 (1 Mar. 2022), p. 011059. DOI: [10.1103/PhysRevX.12.011059](https://doi.org/10.1103/PhysRevX.12.011059). URL: <https://doi.org/10.1103/PhysRevX.12.011059>.

[39] Matteo Puviani, Sangkha Borah, Remmy Zen, Jan Olle, and Florian Marquardt. “Non-Markovian Feedback for Optimized Quantum Error Correction”. In: *Phys. Rev. Lett.* 134 (2 Jan. 2025), p. 020601. DOI: [10.1103/PhysRevLett.134.020601](https://doi.org/10.1103/PhysRevLett.134.020601). URL: <https://doi.org/10.1103/PhysRevLett.134.020601>.

[40] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. “Matrix product states and projected entangled pair states: Concepts, symmetries, theorems”. In: *Rev. Mod. Phys.* 93 (4 Dec. 2021), p. 045003. DOI: [10.1103/RevModPhys.93.045003](https://doi.org/10.1103/RevModPhys.93.045003). URL: <https://link.aps.org/doi/10.1103/RevModPhys.93.045003>.

[41] Román Orús. “A practical introduction to tensor networks: Matrix product states and projected entangled pair states”. In: *Annals of Physics* 349 (2014), pp. 117–158. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2014.06.013>. URL: <https://www.sciencedirect.com/science/article/pii/S0003491614001596>.

[42] F. Verstraete and J. I. Cirac. *Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions*. 2004. arXiv: [cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066) [cond-mat.str-el]. URL: <https://arxiv.org/abs/cond-mat/0407066>.

[43] Wen-Yuan Liu, Shao-Jun Dong, Yong-Jian Han, Guang-Can Guo, and Lixin He. “Gradient optimization of finite projected entangled pair states”. In: *Phys. Rev. B* 95 (19 May 2017), p. 195154. DOI: [10.1103/PhysRevB.95.195154](https://doi.org/10.1103/PhysRevB.95.195154). URL: <https://link.aps.org/doi/10.1103/PhysRevB.95.195154>.

[44] Tom Vieijra, Jutho Haegeman, Frank Verstraete, and Laurens Vanderstraeten. “Direct sampling of projected entangled-pair states”. In: *Phys. Rev. B* 104 (23 Dec. 2021), p. 235141. DOI: [10.1103/PhysRevB.104.235141](https://doi.org/10.1103/PhysRevB.104.235141). URL: <https://link.aps.org/doi/10.1103/PhysRevB.104.235141>.

[45] Ao Chen and Markus Heyl. “Empowering deep neural quantum states through efficient optimization”. In: *Nature Physics* 20.9 (2024), pp. 1476–1481. DOI: <https://doi.org/10.1038/s41567-024-02566-1>. URL: <https://doi.org/10.1038/s41567-024-02566-1>.

[46] Daniel Alcalde Puente. *Implementation of Measurement and Feedback Based Variational Circuits*. <https://github.com/danielalcalde/mVQE>. 2025.

[47] David Perez-Garcia, Frank Verstraete, Michael M Wolf, and J Ignacio Cirac. “Matrix product state representations”. In: *arXiv preprint quant-ph/0608197* (2006).

[48] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac. *Matrix Product State Representations*. 2007. arXiv: [quant-ph/0608197](https://arxiv.org/abs/quant-ph/0608197) [quant-ph]. URL: <https://arxiv.org/abs/quant-ph/0608197>.

[49] J Ignacio Cirac, David Perez-Garcia, Norbert Schuch, and Frank Verstraete. “Matrix product states and projected entangled pair states: Concepts, symmetries, theorems”. In: *Reviews of Modern Physics* 93.4 (2021), p. 045003.

[50] Taku Matsui. “A Characterization of Pure Finitely Correlated States”. In: *Infinite Dimensional Analysis, Quantum Probability and Related Topics* 01.04 (1998), pp. 647–661. DOI: [10.1142/S0219025798000351](https://doi.org/10.1142/S0219025798000351). eprint: <https://doi.org/10.1142/S0219025798000351>. URL: <https://doi.org/10.1142/S0219025798000351>.

[51] Norbert Schuch, Ignacio Cirac, and David Pérez-García. “PEPS as ground states: Degeneracy and topology”. In: *Annals of Physics* 325.10 (2010), pp. 2153–2192. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2010.05.008>. URL: <https://www.sciencedirect.com/science/article/pii/S0003491610000990>.

[52] Frank Pollmann, Erez Berg, Ari M Turner, and Masaki Oshikawa. “Symmetry protection of topological phases in one-dimensional quantum spin systems”. In: *Physical Review B—Condensed Matter and Materials Physics* 85.7 (2012), p. 075125.

[53] Frank Pollmann and Ari M Turner. “Detection of symmetry-protected topological phases in one dimension”. In: *Physical Review B—Condensed Matter and Materials Physics* 86.12 (2012), p. 125441.

[54] Ian Affleck, Tom Kennedy, Elliott H Lieb, and Hal Tasaki. “Rigorous results on valence-bond ground states in antiferromagnets”. In: *Physical review letters* 59.7 (1987), p. 799.

[55] Xie Chen, Zheng-Cheng Gu, and Xiao-Gang Wen. “Classification of gapped symmetric phases in one-dimensional spin systems”. In: *Phys. Rev. B* 83 (3 Jan. 2011), p. 035107. DOI: [10.1103/PhysRevB.83.035107](https://doi.org/10.1103/PhysRevB.83.035107). URL: <https://link.aps.org/doi/10.1103/PhysRevB.83.035107>.

[56] Tzu-Chieh Wei, Ian Affleck, and Robert Raussendorf. “Affleck-Kennedy-Lieb-Tasaki state on a honeycomb lattice is a universal quantum computational resource”. In: *Physical review letters* 106.7 (2011), p. 070501.

[57] J. Eisert, M. Cramer, and M. B. Plenio. “Colloquium: Area laws for the entanglement entropy”. In: *Rev. Mod. Phys.* 82 (1 Feb. 2010), pp. 277–306. DOI: [10.1103/RevModPhys.82.277](https://doi.org/10.1103/RevModPhys.82.277). URL: <https://link.aps.org/doi/10.1103/RevModPhys.82.277>.

[58] Nicolas Laflorencie. “Quantum entanglement in condensed matter systems”. In: *Physics Reports* 646 (2016). Quantum entanglement in condensed matter systems, pp. 1–59. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2016.06.008>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157316301582>.

[59] Ling Wang, Iztok Pižorn, and Frank Verstraete. “Monte Carlo simulation with tensor network states”. In: *Phys. Rev. B* 83 (13 Apr. 2011), p. 134421. DOI: [10.1103/PhysRevB.83.134421](https://doi.org/10.1103/PhysRevB.83.134421). URL: <https://link.aps.org/doi/10.1103/PhysRevB.83.134421>.

[60] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac. “Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions”. In: *Phys. Rev. Lett.* 101 (25 Dec. 2008), p. 250602. DOI: [10.1103/PhysRevLett.101.250602](https://doi.org/10.1103/PhysRevLett.101.250602). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.101.250602>.

[61] Philippe Corboz, T. M. Rice, and Matthias Troyer. “Competing States in the t - J Model: Uniform d -Wave State versus Stripe State”. In: *Phys. Rev. Lett.* 113 (4 July 2014), p. 046402. DOI: [10.1103/PhysRevLett.113.046402](https://doi.org/10.1103/PhysRevLett.113.046402). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.113.046402>.

[62] Philippe Corboz. “Variational optimization with infinite projected entangled-pair states”. In: *Phys. Rev. B* 94 (3 July 2016), p. 035133. DOI: [10.1103/PhysRevB.94.035133](https://doi.org/10.1103/PhysRevB.94.035133). URL: <https://link.aps.org/doi/10.1103/PhysRevB.94.035133>.

[63] Román Orús and Guifré Vidal. “Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction”. In: *Physical Review B—Condensed Matter and Materials Physics* 80.9 (2009), p. 094403.

[64] Román Orús. “Exploring corner transfer matrices and corner tensors for the classical simulation of quantum lattice systems”. In: *Physical Review B—Condensed Matter and Materials Physics* 85.20 (2012), p. 205117.

[65] V. Murg, F. Verstraete, and J. I. Cirac. “Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states”. In: *Phys. Rev. A* 75 (3 Mar. 2007), p. 033605. DOI: [10.1103/PhysRevA.75.033605](https://doi.org/10.1103/PhysRevA.75.033605). URL: <https://link.aps.org/doi/10.1103/PhysRevA.75.033605>.

[66] Michael Lubasch, J. Ignacio Cirac, and Mari-Carmen Bañuls. “Algorithms for finite projected entangled pair states”. In: *Phys. Rev. B* 90 (6 Aug. 2014), p. 064425. DOI: [10.1103/PhysRevB.90.064425](https://doi.org/10.1103/PhysRevB.90.064425). URL: <https://link.aps.org/doi/10.1103/PhysRevB.90.064425>.

[67] Michael Lubasch, J Ignacio Cirac, and Mari-Carmen Bañuls. “Unifying projected entangled pair state contractions”. In: *New Journal of Physics* 16.3 (Mar. 2014), p. 033014. DOI: [10.1088/1367-2630/16/3/033014](https://doi.org/10.1088/1367-2630/16/3/033014). URL: <https://dx.doi.org/10.1088/1367-2630/16/3/033014>.

[68] Sandro Sorella, Michele Casula, and Dario Rocca. “Weak binding between two aromatic rings: Feeling the van der Waals attraction by quantum Monte Carlo methods”. In: *The Journal of chemical physics* 127.1 (2007), p. 014105. DOI: <https://doi.org/10.1063/1.2746035>. URL: <https://doi.org/10.1063/1.2746035>.

[69] Jutho Haegeman et al. “Time-dependent variational principle for quantum lattices”. In: *Physical review letters* 107.7 (2011), p. 070601. URL: https://doi.org/10.1103/10.1007/3-540-10579-4_20.

[70] Giuseppe Carleo and Matthias Troyer. “Solving the quantum many-body problem with artificial neural networks”. In: *Science* 355.6325 (2017), pp. 602–606. DOI: <https://doi.org/10.1126/science.aag2302>. URL: <https://doi.org/10.1126/science.aag2302>.

[71] Laurens Vanderstraeten, Jutho Haegeman, and Frank Verstraete. “Tangent-space methods for uniform matrix product states”. In: *SciPost Phys. Lect. Notes* (2019), p. 7. DOI: [10.21468/SciPostPhysLectNotes.7](https://doi.org/10.21468/SciPostPhysLectNotes.7). URL: <https://scipost.org/10.21468/SciPostPhysLectNotes.7>.

[72] Wen-Yuan Liu, Yi-Zhen Huang, Shou-Shu Gong, and Zheng-Cheng Gu. “Accurate simulation for finite projected entangled pair states in two dimensions”. In: *Phys. Rev. B* 103 (23 June 2021), p. 235155. DOI: [10.1103/PhysRevB.103.235155](https://doi.org/10.1103/PhysRevB.103.235155). URL: <https://link.aps.org/doi/10.1103/PhysRevB.103.235155>.

[73] Sofía González-García et al. “Random insights into the complexity of two-dimensional tensor network calculations”. In: *Physical Review B* 109.23 (2024), p. 235102.

[74] Jielun Chen, Jiaqing Jiang, Dominik Hangleiter, and Norbert Schuch. *Sign problem in tensor network contraction*. 2024. arXiv: [2404.19023 \[quant-ph\]](https://arxiv.org/abs/2404.19023). URL: <https://arxiv.org/abs/2404.19023>.

[75] Shi-Ju Ran, Wei Li, Bin Xi, Zhe Zhang, and Gang Su. “Optimized decimation of tensor networks with super-orthogonalization for two-dimensional quantum lattice models”. In: *Phys. Rev. B* 86 (13 Oct. 2012), p. 134429. DOI: [10.1103/PhysRevB.86.134429](https://doi.org/10.1103/PhysRevB.86.134429). URL: <https://link.aps.org/doi/10.1103/PhysRevB.86.134429>.

[76] Saeed S. Jahromi and Román Orús. “Universal tensor-network algorithm for any infinite lattice”. In: *Phys. Rev. B* 99 (19 May 2019), p. 195105. DOI: [10.1103/PhysRevB.99.195105](https://doi.org/10.1103/PhysRevB.99.195105).

PhysRevB. 99. 195105. URL: <https://link.aps.org/doi/10.1103/PhysRevB.99.195105>.

- [77] Joseph Tindall and Matt Fishman. “Gauging tensor networks with belief propagation”. In: *SciPost Phys.* 15 (2023), p. 222. DOI: [10.21468/SciPostPhys.15.6.222](https://doi.org/10.21468/SciPostPhys.15.6.222). URL: <https://scipost.org/10.21468/SciPostPhys.15.6.222>.
- [78] Jutho Haegeman, Christian Lubich, Ivan Oseledets, Bart Vandereycken, and Frank Verstraete. “Unifying time evolution and optimization with matrix product states”. In: *Physical Review B* 94.16 (2016), p. 165116.
- [79] Michael Lubasch, J Ignacio Cirac, and Mari-Carmen Banuls. “Algorithms for finite projected entangled pair states”. In: *Physical Review B* 90.6 (2014), p. 064425.
- [80] Piotr Czarnik, Jacek Dziarmaga, and Philippe Corboz. “Time evolution of an infinite projected entangled pair state: An efficient algorithm”. In: *Physical Review B* 99.3 (2019), p. 035115.
- [81] Ho N Phien, Johann A Bengua, Hoang D Tuan, Philippe Corboz, and Román Orús. “Infinite projected entangled pair states algorithm improved: Fast full update and gauge fixing”. In: *Physical Review B* 92.3 (2015), p. 035142.
- [82] Ravi Teja Ponnaganti, Matthieu Mambrini, and Didier Poilblanc. “Real-time dynamics of a critical resonating valence bond spin liquid”. In: *Physical Review B* 106.19 (2022), p. 195132.
- [83] Ravi Teja Ponnaganti, Matthieu Mambrini, and Didier Poilblanc. “Tensor network variational optimizations for real-time dynamics: Application to the time-evolution of spin liquids”. In: *SciPost Physics* 15.4 (2023), p. 158.
- [84] Augustine Kshetrimayum, Hendrik Weimer, and Román Orús. “A simple tensor network algorithm for two-dimensional steady states”. In: *Nature communications* 8.1 (2017), p. 1291.
- [85] Ryui Kaneko and Ippei Danshita. “Tensor-network study of correlation-spreading dynamics in the two-dimensional Bose-Hubbard model”. In: *Communications Physics* 5.1 (2022), p. 65.
- [86] Claudius Hubig and J Ignacio Cirac. “Time-dependent study of disordered models with infinite projected entangled pair states”. In: *SciPost Physics* 6.3 (2019), p. 031.
- [87] Augustine Kshetrimayum, Marcel Goihl, and Jens Eisert. “Time evolution of many-body localized systems in two spatial dimensions”. In: *Physical Review B* 102.23 (2020), p. 235132.
- [88] Valentin Murg, Frank Verstraete, and J Ignacio Cirac. “Variational study of hardcore bosons in a two-dimensional optical lattice using projected entangled pair states”. In: *Physical Review A—Atomic, Molecular, and Optical Physics* 75.3 (2007), p. 033605.
- [89] Markus Schmitt and Markus Heyl. “Quantum many-body dynamics in two dimensions with artificial neural networks”. In: *Physical Review Letters* 125.10

(2020), p. 100503. DOI: <https://doi.org/10.1103/PhysRevLett.125.100503>. URL: <https://doi.org/10.1103/PhysRevLett.125.100503>.

[90] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf. “Sequential Generation of Entangled Multiqubit States”. In: *Phys. Rev. Lett.* 95 (2005), p. 110503. DOI: [10.1103/PhysRevLett.95.110503](https://doi.org/10.1103/PhysRevLett.95.110503).

[91] F. G. S. L. Brandão and M. J. Kastoryano. “Finite Correlation Length Implies Efficient Preparation of Quantum Thermal States”. In: *Commun. Math. Phys.* 365 (2019), p. 1. DOI: [10.1007/s00220-018-3150-8](https://doi.org/10.1007/s00220-018-3150-8).

[92] L. Piroli, G. Styliaris, and J. I. Cirac. “Quantum Circuits Assisted by Local Operations and Classical Communication: Transformations and Phases of Matter”. In: *Phys. Rev. Lett.* 127 (2021), p. 220503. DOI: [10.1103/PhysRevLett.127.220503](https://doi.org/10.1103/PhysRevLett.127.220503).

[93] Y. Ge, A. Molnár, and J. I. Cirac. “Rapid adiabatic preparation of injective projected entangled pair states and Gibbs states”. In: *Phys. Rev. Lett.* 116 (2016), p. 080503. DOI: [10.1103/physrevlett.116.080503](https://doi.org/10.1103/physrevlett.116.080503).

[94] N. Schuch, D. Pérez-García, and I. Cirac. “Classifying quantum phases using matrix product states and projected entangled pair states”. In: *Phys. Rev. B* 84 (2011), p. 165139. DOI: [10.1103/PhysRevB.84.165139](https://doi.org/10.1103/PhysRevB.84.165139).

[95] S. Bachmann, W. D. Roeck, and M. Fraas. “The adiabatic theorem and linear response theory for extended quantum systems”. In: *Communications in Mathematical Physics* 361 (2018), p. 997. DOI: [10.1007/s00220-018-3117-9](https://doi.org/10.1007/s00220-018-3117-9).

[96] Sean D Barrett and Pieter Kok. “Efficient high-fidelity quantum computation using matter qubits and linear optics”. In: *Physical Review A—Atomic, Molecular, and Optical Physics* 71.6 (2005), p. 060310.

[97] Zhi Zhao et al. “Experimental demonstration of five-photon entanglement and open-destination teleportation”. In: *Nature* 430.6995 (2004), pp. 54–58.

[98] Daniel E Browne and Terry Rudolph. “Resource-efficient linear optical quantum computation”. In: *Physical Review Letters* 95.1 (2005), p. 010501.

[99] Serge Rosenblum et al. “Fault-tolerant detection of a quantum error”. In: *Science* 361.6399 (2018), pp. 266–270.

[100] Julio T Barreiro et al. “An open-system quantum simulator with trapped ions”. In: *Nature* 470.7335 (2011), pp. 486–491.

[101] Yiheng Lin et al. “Dissipative production of a maximally entangled steady state of two quantum bits”. In: *Nature* 504.7480 (2013), pp. 415–418.

[102] Shyam Shankar et al. “Autonomously stabilized entanglement between two superconducting quantum bits”. In: *Nature* 504.7480 (2013), pp. 419–422.

[103] Fernando Pastawski and Beni Yoshida. “Fault-tolerant logical gates in quantum error-correcting codes”. In: *Physical Review A* 91.1 (2015), p. 012305.

[104] John Dengis, Robert König, and Fernando Pastawski. “An optimal dissipative encoder for the toric code”. In: *New Journal of Physics* 16.1 (2014), p. 013023.

- [105] Mario Motta et al. “Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution”. In: *Nature Physics* 16.2 (2020), pp. 205–210.
- [106] Rajeev Acharya et al. “Quantum error correction below the surface code threshold”. In: *Nature* (2024). URL: <https://doi.org/10.1038/s41586-024-08449-y>.
- [107] Laura Caune et al. “Demonstrating real-time and low-latency quantum error correction with superconducting qubits”. In: *arXiv preprint arXiv:2410.05202* (2024). URL: <https://doi.org/10.48550/arXiv.2410.05202>.
- [108] Volodymyr V Sivak et al. “Real-time quantum error correction beyond break-even”. In: *Nature* 616.7955 (2023), pp. 50–55. URL: <https://doi.org/10.1038/s41586-023-05782-6>.
- [109] Riccardo Porotti, Antoine Essig, Benjamin Huard, and Florian Marquardt. “Deep reinforcement learning for quantum state preparation with weak nonlinear measurements”. In: *Quantum* 6 (2022), p. 747.
- [110] Pranav Vaidhyanathan, Florian Marquardt, Mark T Mitchison, and Natalia Ares. “Quantum feedback control with a transformer neural network architecture”. In: *arXiv preprint arXiv:2411.19253* (2024).
- [111] Faisal Alam and Bryan K Clark. “Learning dynamic quantum circuits for efficient state preparation”. In: *arXiv preprint arXiv:2410.09030* (2024). URL: <https://doi.org/10.48550/arXiv.2410.09030>.
- [112] Yuxuan Yan, Muzhou Ma, You Zhou, and Xiongfeng Ma. “Variational LOCC-Assisted Quantum Circuits for Long-Range Entangled States”. In: *Phys. Rev. Lett.* 134 (17 Apr. 2025), p. 170601. DOI: [10.1103/PhysRevLett.134.170601](https://doi.org/10.1103/PhysRevLett.134.170601). URL: <https://doi.org/10.1103/PhysRevLett.134.170601>.
- [113] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. “Barren plateaus in quantum neural network training landscapes”. In: *Nature communications* 9.1 (2018), p. 4812.
- [114] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. “Cost function dependent barren plateaus in shallow parametrized quantum circuits”. In: *Nature communications* 12.1 (2021), p. 1791. URL: <https://doi.org/10.1038/s41467-021-21728-w>.
- [115] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. “Cost function dependent barren plateaus in shallow parametrized quantum circuits”. In: *Nature communications* 12.1 (2021), p. 1791.
- [116] Andrea Skolik, Jarrod R McClean, Masoud Mohseni, Patrick Van Der Smagt, and Martin Leib. “Layerwise learning for quantum neural networks”. In: *Quantum Machine Intelligence* 3 (2021), pp. 1–11.
- [117] Zoë Holmes, Kunal Sharma, Marco Cerezo, and Patrick J Coles. “Connecting ansatz expressibility to gradient magnitudes and barren plateaus”. In: *PRX quantum* 3.1 (2022), p. 010313.

[118] Eric Ricardo Anschuetz. “Critical Points in Quantum Generative Models”. In: *International Conference on Learning Representations*. 2022. URL: <https://openreview.net/forum?id=2f1z55GVQN>.

[119] Xuchen You and Xiaodi Wu. “Exponentially many local minima in quantum neural networks”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 12144–12155.

[120] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. “Evaluating analytic gradients on quantum hardware”. In: *Physical Review A* 99.3 (2019), p. 032331.

[121] Gian Giacomo Guerreschi and Mikhail Smelyanskiy. “Practical optimization for hybrid quantum-classical algorithms”. In: *arXiv preprint arXiv:1701.01450* (2017).

[122] Daniel Faílde, José Daniel Viqueira, Mariamo Mussa Juane, and Andrés Gómez. “Using differential evolution to avoid local minima in variational quantum algorithms”. In: *Scientific Reports* 13.1 (2023), p. 16230.

[123] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. “Quantum circuit learning”. In: *Physical Review A* 98.3 (2018), p. 032309.

[124] Matthew Fishman, Steven R. White, and E. Miles Stoudenmire. “The ITensor Software Library for Tensor Network Calculations”. In: *SciPost Phys. Codebases* (2022), p. 4. DOI: [10.21468/SciPostPhysCodeb.4](https://doi.org/10.21468/SciPostPhysCodeb.4). URL: <https://scipost.org/10.21468/SciPostPhysCodeb.4>.

[125] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. “Multilayer feedforward networks are universal approximators”. In: *Neural Networks* 2.5 (1989), pp. 359–366. ISSN: 0893-6080. DOI: [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8). URL: <https://www.sciencedirect.com/science/article/pii/0893608089900208>.

[126] George Cybenko. “Approximation by superpositions of a sigmoidal function”. In: *Mathematics of control, signals and systems* 2.4 (1989), pp. 303–314.

[127] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. “Deep feedforward networks”. In: *Deep learning* 1 (2016), pp. 161–217.

[128] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.

[129] Vinod Nair and Geoffrey E Hinton. “Rectified linear units improve restricted boltzmann machines”. In: *Proceedings of the 27th international conference on machine learning (ICML-10)*. 2010, pp. 807–814.

[130] Noam Shazeer. “Glu variants improve transformer”. In: *arXiv preprint arXiv:2002.05202* (2020).

[131] Biao Zhang and Rico Sennrich. “Root mean square layer normalization”. In: *Advances in Neural Information Processing Systems* 32 (2019).

[132] Jeffrey L Elman. “Finding structure in time”. In: *Cognitive science* 14.2 (1990), pp. 179–211.

[133] Dian Wu, Riccardo Rossi, Filippo Vicentini, and Giuseppe Carleo. “From tensor-network quantum states to tensorial recurrent neural networks”. In: *Phys. Rev. Res.* 5 (3 July 2023), p. L032001. DOI: [10.1103/PhysRevResearch.5.L032001](https://doi.org/10.1103/PhysRevResearch.5.L032001). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.5.L032001>.

[134] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. “Learning long-term dependencies with gradient descent is difficult”. In: *IEEE transactions on neural networks* 5.2 (1994), pp. 157–166.

[135] Kyunghyun Cho et al. “Learning phrase representations using RNN encoder-decoder for statistical machine translation”. In: *arXiv preprint arXiv:1406.1078* (2014).

[136] Abhimanyu Dubey et al. “The llama 3 herd of models”. In: *arXiv preprint arXiv:2407.21783* (2024). URL: <https://doi.org/10.48550/arXiv.2407.21783>.

[137] Ashish Vaswani et al. “Attention is all you need”. In: *Advances in neural information processing systems* 30 (2017).

[138] Albert Gu and Tri Dao. “Mamba: Linear-time sequence modeling with selective state spaces”. In: *arXiv preprint arXiv:2312.00752* (2023).

[139] Daniel Alcalde Puente and Matteo Rizzi. *Simulations performed to produce figures presented in Learning Feedback Mechanisms for Measurement-Based Variational Quantum State Preparation*. 2025. DOI: [10.5281/zenodo.14654836](https://doi.org/10.5281/zenodo.14654836). URL: <https://doi.org/10.5281/zenodo.14654836>.

[140] Laurens Vanderstraeten, Jutho Haegeman, Philippe Corboz, and Frank Verstraete. “Gradient methods for variational optimization of projected entangled-pair states”. In: *Phys. Rev. B* 94 (15 Oct. 2016), p. 155123. DOI: [10.1103/PhysRevB.94.155123](https://doi.org/10.1103/PhysRevB.94.155123). URL: <https://link.aps.org/doi/10.1103/PhysRevB.94.155123>.

[141] Juraj Hasik, Maarten Van Damme, Didier Poilblanc, and Laurens Vanderstraeten. “Simulating Chiral Spin Liquids with Projected Entangled-Pair States”. In: *Phys. Rev. Lett.* 129 (17 Oct. 2022), p. 177201. DOI: [10.1103/PhysRevLett.129.177201](https://doi.org/10.1103/PhysRevLett.129.177201). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.177201>.

[142] Wen-Yuan Liu et al. “Gapless spin liquid ground state of the spin- $\frac{1}{2}$ $J_1 - J_2$ Heisenberg model on square lattices”. In: *Phys. Rev. B* 98 (24 Dec. 2018), p. 241109. DOI: [10.1103/PhysRevB.98.241109](https://doi.org/10.1103/PhysRevB.98.241109). URL: <https://link.aps.org/doi/10.1103/PhysRevB.98.241109>.

[143] Y.-Y. Shi, L.-M. Duan, and G. Vidal. “Classical simulation of quantum many-body systems with a tree tensor network”. In: *Phys. Rev. A* 74 (2 Aug. 2006), p. 022320. DOI: [10.1103/PhysRevA.74.022320](https://doi.org/10.1103/PhysRevA.74.022320). URL: <https://link.aps.org/doi/10.1103/PhysRevA.74.022320>.

- [144] L. Tagliacozzo, G. Evenbly, and G. Vidal. “Simulation of two-dimensional quantum systems using a tree tensor network that exploits the entropic area law”. In: *Phys. Rev. B* 80 (23 Dec. 2009), p. 235127. DOI: [10.1103/PhysRevB.80.235127](https://doi.org/10.1103/PhysRevB.80.235127). URL: <https://link.aps.org/doi/10.1103/PhysRevB.80.235127>.
- [145] Pietro Silvi et al. “The Tensor Networks Anthology: Simulation techniques for many-body quantum lattice systems”. In: *SciPost Phys. Lect. Notes* (2019), p. 8. DOI: [10.21468/SciPostPhysLectNotes.8](https://doi.org/10.21468/SciPostPhysLectNotes.8). URL: <https://scipost.org/10.21468/SciPostPhysLectNotes.8>.
- [146] Timo Felser, Simone Notarnicola, and Simone Montangero. “Efficient Tensor Network Ansatz for High-Dimensional Quantum Many-Body Problems”. In: *Phys. Rev. Lett.* 126 (17 Apr. 2021), p. 170603. DOI: [10.1103/PhysRevLett.126.170603](https://doi.org/10.1103/PhysRevLett.126.170603). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.170603>.
- [147] Daniel Alcalde Puente, Erik Lennart Weerda, and Konrad Schröder. *Quantum Natural Finite PEPS*. github/Quantum-NaturalfPEPS.jl. 2025. URL: <https://github.com/KonradSchroeder/QuantumNaturalfPEPS.jl>.
- [148] Daniel Alcalde Puente. *Quantum Natural Gradient*. github/QuantumNaturalGradient.jl. 2025. URL: <https://github.com/danielalcalde/QuantumNaturalGradient.jl>.
- [149] Daniel Alcalde Puente, Erik Lennart Weerda, Konrad Schröder, and Matteo Rizzi. *Data for "Efficient optimization and conceptual barriers in variational finite Projected Entangled-Pair States"*. Zenodo, Mar. 2025. DOI: [10.5281/zenodo.15046369](https://doi.org/10.5281/zenodo.15046369). URL: <https://doi.org/10.5281/zenodo.15046369>.