

Explaining Anomalies with Tensor Networks

Hans Hohenfeld^{*†}, Marius Beuerle[‡], and Elie Mounzer^{*}

^{*}German Research Center for Artificial Intelligence, Robotics Innovation Center, 28359 Bremen, Germany

Email: hans.hohenfeld@dfki.de, elie.mounzer@dfki.de

[†]QBITFLOW GmbH, 28217 Bremen, Germany

[‡]University of Bremen, Robotics Research Group, 28359 Bremen, Germany

Abstract—Tensor networks, a class of variational quantum many-body wave functions have attracted considerable research interest across many disciplines, including classical machine learning. Recently, Aizpurua et al. demonstrated explainable anomaly detection with matrix product states on a discrete-valued cyber-security task, using quantum-inspired methods to gain insight into the learned model and detected anomalies. Here, we extend this framework to real-valued data domains. We furthermore introduce tree tensor networks for the task of explainable anomaly detection. We demonstrate these methods with three benchmark problems, show adequate predictive performance compared to several baseline models and both tensor network architectures’ ability to explain anomalous samples. We thereby extend the application of tensor networks to a broader class of potential problems and open a pathway for future extensions to more complex tensor network architectures.

Index Terms—Quantum Inspired Machine Learning, Unsupervised Learning, Anomaly Detection, Explainable AI

I. INTRODUCTION

Anomaly detection [1], [2] is the task of identifying patterns or samples in data that are rare, novel, or deviate from a norm or expectation. It is distinct from, e.g., classification, as the available data for a given problem is usually strongly biased toward regular instances, and samples are (partially) unlabeled. For many applications, such as fraud detection, medical diagnostics, or monitoring of industrial processes, explaining anomalies, rather than just identifying them, can benefit informed decision-making by domain experts [3]. Tensor networks [4], originally a technique to express low-rank approximations of otherwise classically intractable quantum many-body systems, have been proposed for machine learning applications such as compression of neural networks [5]–[8], supervised [9]–[15] and unsupervised [16]–[23] learning. Aizpurua et al. [24] recently proposed an anomaly detection framework based on matrix product states (MPS), demonstrating the identification of anomalous samples in a cyber-security data set and quantum-inspired techniques to extract rich information from the trained model to understand the detected anomalies.

Here, we extend this tensor network based anomaly detection framework to real-valued data, leveraging results on real-valued density estimation with matrix product states [23]. We furthermore demonstrate the applicability of these techniques to tree tensor networks (TTN) and compare their predictive performance and ability to explain anomalies. Additionally, we extend the two-site DMRG-like gradient-based optimization




Fig. 1. Matrix product state (a) and tree tensor network (b). The tensors are contracted with an encoded feature vector. For details we refer to Sec. II.

procedure previously introduced for MPS [17] to tree tensor networks, enabling the adaptation of bond dimensions during training. We evaluate the tensor networks on three anomaly detection benchmark data sets and compare their predictive performance to several baseline models before demonstrating their ability to provide insight into detected anomalies. Our contribution broadens the scope of application for tensor networks in explainable AI and may foster similar research directions for quantum machine learning due to the intricate connection between tensor networks and quantum circuits in the latter field [25]–[27].

The remaining paper is structured as follows: In Sec. II we outline the theoretical background, learning methods, explainability metrics, and experimental setup. Afterward, we present and discuss our findings in Sec. III, before outlining limitations and possible future research directions in Sec. IV.

II. METHOD

A. Tensor Networks

The state vector of a quantum many-body system of size L with local d -dimensional state space $\{\sigma_i\}$ is given by

$$|\psi\rangle := \sum_{\sigma_1, \dots, \sigma_L} c_{\sigma_1, \dots, \sigma_L} |\sigma_1, \dots, \sigma_L\rangle. \quad (1)$$

This general and exact representation does not take known properties of many physical systems of interest, such as locality, correlation length, or criticality, into account [4]. Writing the d^L coefficients $c_{\sigma_1, \dots, \sigma_L} \in \mathbb{C}$ into a tensor $\Psi(\sigma) = c_\sigma$ with $\sigma = (\sigma_1, \dots, \sigma_L)$ enables a compressed re-parameterization of $|\psi\rangle$ as a tensor network. The first class of tensor network




Fig. 2. In each training step, the canonical center (dark blue) is contracted with its neighbor in the traversal direction (a). The gradient $\nabla_C \Psi(\mathbf{x})$ is computed by contracting the network around C and removing C . The outer product of the remaining vectors is the gradient tensor (b). After each update step, the contracted tensor undergoes a truncated singular value decomposition. The singular values are contracted in the direction of traversal to move the canonical center (c). To compute the reduced density matrix of a sub-system, the tensor network is contracted with its transpose, leaving the physical bonds of the sites of the sub-system open (d). To compute the reduced density matrix for a conditional marginal distribution, the sites the distribution is conditioned on are fixed to $|x_i = a_i\rangle\langle x_i = a_i|$, $a_i \in \mathcal{I}$ in the contraction (e). The areas with gray backgrounds contract to the identity. The traversal order for the two-site update on TTNs from the canonical center (right-most leaf) is akin to a depth-first traversal with back-tracking, where all but the last leaf are only traversed upwards (f). All computations in (a)-(e) are shown for matrix product states but work for TTNs analogously.

states we consider in this work are *matrix product states* (MPS), decomposing $\Psi(\sigma)$ into [28]

$$\Psi(\sigma_1, \dots, \sigma_L) := \sum_{\{\alpha\}} A_{\alpha_1}^{\sigma_1} A_{\alpha_1 \alpha_2}^{\sigma_2} \cdots A_{\alpha_{L-1}}^{\sigma_L}. \quad (2)$$

Each A^{σ_i} is a $D_{i-1} \times D_i$ matrix with $D_0 = D_L = 1$, graphically illustrated in Fig. 1a. An exact decomposition leads to $D_i \leq d^{\frac{L}{2}}$, for many applications, such as the ground state computation of local Hamiltonians with gapped eigenspectrum on a 1-dimensional lattice, a rank-reduced approximation suffices, as those states follow an area law of entanglement entropy [29]. For any bi-partition of $|\psi\rangle$ into sub-systems $\mathcal{X}$ and $\mathcal{Y}$ at bond j , the entanglement entropy S is bound by $S(\rho_{\mathcal{X}}) \leq \log D_j$, enabling faithful representation of area-law states with $D_i \ll d^{\frac{L}{2}}$ [30], typically $D_i \sim \mathcal{O}(100)$.

Tree tensor networks (TTN) were introduced to study strongly correlated systems with non-local interactions [31]. For a balanced binary tree, as shown in Fig. 1b, $\Psi(\sigma)$ is given by

$$\Psi(\sigma) := \sum_{\{\alpha\}} A_{\alpha_1 \alpha_2} \prod_{n=1}^{L/2-2} A_{\alpha_n \alpha_{2n+1} \alpha_{2n+2}} \prod_{n=1}^{L/2} A_{\alpha_{n+L/2-2}}^{\sigma_{2n-1} \sigma_{2n}}. \quad (3)$$

The binary tree construction is no necessity. TTNs can be expressed by arbitrary trees and their structure even be subject to optimization [32].

B. Unsupervised Learning with Tensor Networks

A tensor network state can be interpreted as a probability distribution of d^L basis states. Let $\mathbf{x} \in \{0, 1\}^L$ and $d = 2$ then

$$P(\mathbf{x}) = |\langle \psi | \mathbf{x} \rangle|^2 = |\Psi(\mathbf{x})|^2, \quad (4)$$

with $\mathbf{x} := |x_1\rangle \otimes |x_2\rangle \otimes \cdots \otimes |x_L\rangle$. For a data set $\mathcal{D} \subseteq \{0, 1\}^L$, we can optimize the tensor network to represent the probability distribution over $\mathcal{D}$ by minimizing the negative log-likelihood loss [17], [19]

$$\mathcal{L} := -\frac{1}{|\mathcal{D}|} \sum_{\mathbf{x} \in \mathcal{D}} \log P(\mathbf{x}). \quad (5)$$

To train the tensor network, we start with one site as canonical center [28] and traverse the network. In each training step, the current canonical center is contracted with its neighboring site in direction of traversal, as shown in Fig. 2a. The contracted site C is then updated with one or several steps of (stochastic) gradient descent with the gradient

$$\nabla_C \mathcal{L} = -\frac{2}{|\mathcal{D}|} \sum_{\mathbf{x} \in \mathcal{D}} \frac{\nabla_C \Psi(\mathbf{x})}{\Psi(\mathbf{x})}, \quad (6)$$

updating C by $C \leftarrow C - \eta \nabla_C \mathcal{L}$, with a learning rate $0 < \eta \ll 1$. The procedure to compute $\nabla_C \Psi(\mathbf{x})$ is shown in Fig. 2b. After updating, C is decomposed using a truncated SVD, where we control the bond dimension by discarding singular values below a threshold:

$$C \approx U \Sigma V^\dagger. \quad (7)$$

As shown in Fig. 2c, the singular values are then contracted into the next site in the direction of traversal, making it the new canonical center. For matrix product states, the traversal is a sweeping motion from left to right and back. Tree tensor networks have been trained with single site updates in previous works [19], here we adapt the two-site learning procedure to a depth-first traversal with backtracking, illustrated in Fig. 2f.

C. Encoding Real-Valued Data

To encode real-valued data into a feature layer as shown in Figure 1, we leverage the framework introduced in Ref. [23]. Given a sample $\mathbf{x} \in \mathcal{D}$ with $\mathbf{x} = (x_1, \dots, x_L)$, each feature $x_i \in \mathcal{I}$ is encoded with an isometric feature map $\xi: \mathcal{I} \rightarrow \mathbb{R}^N$ composed of N feature functions

$$\xi(x_i) := |x_i\rangle = \begin{bmatrix} g_1(x_i) \\ g_2(x_i) \\ \vdots \\ g_N(x_i) \end{bmatrix}. \quad (8)$$

The authors of Ref. [23] discuss several feature functions and their induced priors. We found the best results by rescaling all features to $\mathcal{I} = [0, 1]$ and encode with

$$g_n(x_i) := \sqrt{2n+1} \tilde{P}_n(x_i). \quad (9)$$

$\tilde{P}_n : \mathbb{R} \rightarrow \mathbb{R}$ is the n -th shifted Legendre polynomial

$$\tilde{P}_n(x_i) := \frac{1}{n!} \frac{d^n}{dx_i^n} (x_i^2 - x_i)^n, \quad (10)$$

fulfilling the orthogonality relation

$$\int_0^1 \tilde{P}_m(x) \tilde{P}_n(x) dx = \frac{1}{2n+1} \delta_{mn}. \quad (11)$$

The factor $\sqrt{2n+1}$ scales the polynomials to an orthonormal basis.

D. Extracting Information

Tensor networks are a representation of a class of quantum states. Thus, we can apply quantum information theoretic methods to analyze the trained model and data samples of interest. In Ref. [24], the authors established the reduced density matrix corresponding to a subset of features as the primary tool for this purpose. Conceptually, the reduced density matrix corresponds to the marginal distribution the model assigns to the sub-system. Fig. 2d shows the contraction of an MPS with its conjugate transpose, leaving the physical indices of the sub-system open, which, after re-shaping, computes the reduced density matrix. Here, we operate on real-valued data, in principle, requiring integrating over the variables being marginalized out. Given the isometric encoding map, we have

$$\int_{\mathcal{I}} |x_i\rangle\langle x_i| = \mathbb{I}, \quad (12)$$

and this step can be omitted. The conditional reduced density matrix, corresponding to the conditional marginal distribution of a sub-systems of features, is computed by the contraction shown in Fig. 2e. The variables to condition the distribution on are fixed by introducing $|x_i = a_i\rangle\langle x_i = a_i|$, $a_i \in \mathcal{I}$ into the contraction. If we make one of the sites within the boundaries of the sub-system the canonical center before either of these contractions, the areas shaded in gray in Fig. 2d and 2e contract to the identity [28]. When working with discrete-valued data encoded into basis states, the multinomial distribution over the features of the sub-system can be directly read from the diagonal of the reduced density matrix. This information is not readily available for real-valued data encoded in an isometric feature layer, but the reduced density matrix induces a quasi-probability density function

$$q(\mathbf{x}) := \text{tr}(|\mathbf{x}\rangle\langle \mathbf{x}| \rho_{\mathcal{X}}), \quad (13)$$

which after normalization by its integral can be used to compute expected value, expected variance, and the expected covariance matrix for the sub-system $\mathcal{X}$. From the reduced density matrix two further measures of interests can be computed. That is, the von Neumann entropy

$$S(\rho_{\mathcal{X}}) := -\text{tr}(\rho_{\mathcal{X}} \log \rho_{\mathcal{X}}), \quad (14)$$

TABLE I
PROPERTIES OF THE DATA SETS USED IN THIS STUDY.

Data set	Ref.	L^a	Samples	Anomalies	N_{MPS}^b	N_{TTN}^c
ECG5000	[33]	140	4,998	2,079	4	4
Satellite	[34]	36	6,435	2,036	5	5
Spambase	[35]	57	4,601	1,813	6	5

^a Number of features

^b Number of encoding feature functions for MPS

^c Number of encoding feature functions for TTN

and the mutual information I of two sub-systems $\mathcal{X}$ and $\mathcal{Y}$

$$I(\mathcal{X}; \mathcal{Y}) := S(\rho_{\mathcal{Y}}) + S(\rho_{\mathcal{X}}) - S(\rho_{\mathcal{X}, \mathcal{Y}}). \quad (15)$$

With these, we can analyze the learned correlations and dependencies between individual and groups of features.

E. Learning Setup

We use three different benchmark data sets for this study. Their main properties are summarized in Table I. The ECG5000 [33] data set contains about 5,000 heartbeats from an ECG recording of a patient with congestive heart failure. Each heartbeat is normalized in time to equal length and described by 140 features. Satellite [34] contains 6,435 multi-spectral satellite images, each sample being a 3×3 patch recorded in four different bands. The third set, Spambase [35], is a collection of 4,601 email samples, described by the word count of 57 different key characters or words used to identify spam emails. N_{MPS} and N_{TTN} in Table I refer to the number of encoding feature functions for each data set and tensor network type. The selected data sets provide a variety of application domains and are commonly used in the anomaly detection literature. Thus, baseline performance comparisons are readily available. Yet, while being instructive for the techniques discussed here, they are relatively small in scale. This choice is a consequence of the considerable overhead in computational resources required for the tensor network models, which we discuss in Sec. III-E.

All models are trained unsupervised. To this end, we prepare each data set to contain 95% regular samples and 5% anomalies, either unlabeled. The anomalies are composed in equal parts of those provided with the data set and generated anomalies, for which we use the techniques described in Ref. [2]. The prepared data sets are then split into a stratified 10-fold. Each model is trained and evaluated ten times, with one of the folds as a hold-out set. We perform two evaluation tasks. First, we assess the model's ability to separate regular from anomalous samples in the nine training folds. Afterward, we conduct an inductive evaluation with the hold-out set.

To assess the anomaly detection capability of both tensor networks, before demonstrating their ability to explain anomalous samples, we compare their performance to three baseline models. We use a one-class support vector machine (SVM) [36], an isolation forest (IFO) [37] and an autoencoder (AE) [38] with two hidden layers in the encoder and decoder part. Before experiments, the hyper-parameters of all models

TABLE II
MEAN AUCROC SCORE AND STD. DEVIATION ON BOTH TASKS.

Task	Model	ECG5000		Satellite		Spambase	
		Mean	Std.	Mean	Std.	Mean	Std.
Separation	MPS	0.92	0.02	0.90	0.01	0.83	0.01
	TTN	0.96	0.00	0.89	0.01	0.82	0.01
	SVM	0.95	0.00	0.80	0.01	0.77	0.01
	IFO	0.93	0.00	0.90	0.01	0.88	0.01
	AE	0.93	0.01	0.87	0.03	0.81	0.01
Inductive	MPS	0.92	0.03	0.89	0.05	0.83	0.06
	TTN	0.96	0.01	0.89	0.05	0.82	0.05
	SVM	0.95	0.02	0.80	0.07	0.77	0.06
	IFO	0.93	0.02	0.90	0.05	0.88	0.06
	AE	0.94	0.03	0.86	0.06	0.81	0.05

Mean and standard deviation for the area under the ROC curve.
Best result per task/data set marked with green background.

and for all data sets were tuned extensively with Optuna [39]. We implemented the experimental setup using Jax [40] and opt_einsum [41], code is available under an open-source license at Ref. [42].

III. RESULTS

A. Anomaly Detection Performance

In Table II, we report the mean area under the curve for the receiver operating characteristic (AUCROC) and its standard deviation for all models, data sets over the ten repetitions of the separation and inductive anomaly detection tasks. In general, all models perform adequately and in line with previous results reported on these or similar data sets in the literature [2] for unsupervised methods. Notably, the mean performance remains almost unchanged, going from the separation task to the evaluation on the hold-out set, with only the standard deviation increasing slightly. Setting the threshold for flagging a sample anomalous is a trade-off between sensitivity and specificity, which regularly is domain and task-dependent. Since our primary focus here is explainability and not task-specific predictive performance, we set the threshold to the intersection between true-positive and true-negative rates, which provided sufficient results for this study.

B. Identifying Anomalous Features

After detecting anomalous samples, understanding what features of the sample are outside a defined norm benefits diagnostics and informs decision-making by domain experts. This information is readily available from the tensor network models with the techniques discussed in Sec. II-D. As an example, we show the per-feature distribution learned by a matrix product state and by a tree tensor network on the ECG5000 data set in Fig. 3a. After normalizing the quasi-probability density defined in Eq. (13) for each feature, we compute its expected value and expected standard deviation, depicted by the green line and the area around it in Fig. 3a. Additionally, a normal sample (blue) and ECG recording flagged as anomalous (yellow) are shown. The anomalous sample has several ranges of features outside the expected

standard deviation, indicated by the red-shaded areas. Here, the sharp rise and following inversion of the amplitude at the end of the sample stand out, indicating a type of arrhythmia known as premature ventricular contraction [43]. Setting the appropriate range relative to the per-feature distribution to consider a feature anomalous is a decision to be made by domain experts on a task-dependent basis. We additionally find that the TTN learned a narrower distribution with a sharper distinction of features. Consequently, the anomalous features stand out more pronounced compared to the model learned by the MPS.

C. Inspecting Model Structure

The trained tensor networks also provide insight into the structure captured by the model. Fig. 3b shows the all-to-all feature mutual information computed from an MPS and a TTN trained on the *Satellite* data set. Next to it, the same measure of feature dependencies is shown for the entire polluted data set, as well as separately for the regular and anomalous samples. We determined the latter three with a low-bias histogram-based estimator [44]. All values have been scaled to the interval $[0,1]$ and the diagonal set to 0, for better visibility. Both tensor networks reveal a similar, structured pattern that corresponds to the structure of the data set discussed in Sec. II-E. The same pattern can be recognized, albeit less sharp, in the mutual information computed for the regular samples, but not for the polluted data set, we used to train the models. For a real-world task, such details may not be known upfront, and the trained tensor network model can support analyzing and understanding structure and feature dependence in the data, separating regular from anomalous samples during training.

D. Conditional Sampling

For some applications, it can be helpful to understand how anomalous features should be different, given the context of the remaining features of a sample. Fig. 3c shows such an evaluation for the *Spambase* data set. Several features, which are individual characters or words present in an email, are flagged, as their relative frequency is outside or at the margins of the per-feature expected standard deviation computed from a trained matrix product state. As described in Sec. II-D, we can construct a conditional reduced density matrix, fixing the features considered normal, and use it to compute the expected value of the resulting conditional marginal distribution. We note that feature 57, after computing its expected value from the conditional marginal distribution, is still outside the expected standard deviation of the same feature's marginal distribution, emphasizing that neither of the presented techniques should be considered in isolation nor that one is preferable over the other. Their applicability and usefulness rather has to be assessed in combination and in the context of the task, to provide the required understanding of detected anomalies.




Fig. 3. Three examples of information extraction from the trained tensor network models. Both, MPS and TTN correctly identified an anomalous heartbeat in the ECG5000 data set (a). Beyond the detection, both models can identify the parts of the ECG, that explain the anomaly, in this example the sharp spike and reversal of amplitude at the end. This information can support a domain expert to judge the presented data. Notably, the TTN learned a sharper and narrower distribution, than the MPS. The all-to-all feature mutual information extracted from the MPS and TTN (b) shows a structural pattern, that can also be seen in the normal samples from the *Satellite* data set, but not in the polluted data, the models were trained on. For better visibility, values are scaled to the interval $[0, 1]$ and the diagonal set to 0. The tensor networks not only enable identifying anomalous features of a sample (c) but also computing their expected values, conditioned on the normal features of the sample, providing additional insight into a detected anomaly.

E. Model Complexity & Numerical Precision

Both tensor network models provide valuable insight to understand anomalous samples as well as the structure of the data set used for training. This additional utility, does not come without a cost. The tensor network models are notably larger compared to the autoencoder we used as a baseline model, with between 1.5 and 2.5 times the trainable parameters. Combined with the inherently sequential learning algorithms, optimizing two sites at a time, make their training computationally demanding and potentially time and resource prohibitive for high-dimensional data sets. Another aspect contributing to this limitation is numeric stability and precision. A tensor network is an approximation of a quantum state, assigning a probability to each state encoded in the feature layer described in Sec. III-B. As a consequence, for each given basis $|\sigma\rangle$, the total probability has to be distributed among d^L basis states, which for sufficiently large d and L regularly leads to small probabilities during initialization and training, requiring double-precision floating point parameters for the relatively small data sets considered in this study.

IV. CONCLUSION

Here, we extended the framework for explainable anomaly detection with matrix product states, introduced in Ref. [24] to real-valued data and tree tensor networks. We showed that both tensor network architectures perform adequately in unsupervised anomaly detection compared to several baseline models on three different benchmark data sets. Furthermore, we demonstrated how quantum-inspired techniques can be used in the real-valued domain, to extract information about anomalous features, gain insight into the learned data set

structure and feature dependencies, as well the computation of expected values for individual features conditioned on parts of a sample. Thereby, we extend the scope of applicability for explainable anomaly detection with tensor networks to a broader class of potential problems. Given their roots in the study of quantum many-body systems and their connection to quantum machine learning [25]–[27], tensor network based techniques in machine learning can bridge the gap between the current small-scale quantum machine learning in the NISQ era and potential future learning methods on large-scale, fault-tolerant quantum computers. One avenue of future research to this end could be methods to compute measures for explainable machine learning as presented here, directly from a quantum circuit based tensor network model as well as the efficient transferability of classically (pre-)trained tensor networks onto quantum circuits. A further natural extension of our work would be the evaluation of the presented techniques to other types of tensor networks such as PEPS [45] and MERA [46] and in-depth comparison of their capabilities in a wider variety of anomaly detection tasks. Furthermore, we intend to address the computational constraints in terms of model size and training time discussed in Sec. III-E in future work.

ACKNOWLEDGMENT

This work was funded by the German Ministry of Economic Affairs and Climate Action (BMWK) and the German Aerospace Center (DLR) in project QuDA-KI (grant no. 50RA2206). We thank Borja Aizpurua for the helpful discussion and Uğur Bolat, Lukas Groß, and Gunnar Schönhoff for their feedback on the manuscript.

REFERENCES

- [1] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, pp. 15:1–15:58, 2009.
- [2] S. Han, X. Hu, H. Huang, M. Jiang, and Y. Zhao, "ADBBench: Anomaly Detection Benchmark," in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 32 142–32 159.
- [3] Z. Li, Y. Zhu, and M. Van Leeuwen, "A Survey on Explainable Anomaly Detection," *ACM Trans. Knowl. Discov. Data*, vol. 18, no. 1, pp. 23:1–23:54, 2023.
- [4] R. Orús, "A practical introduction to tensor networks: Matrix product states and projected entangled pair states," *Annals of Physics*, vol. 349, pp. 117–158, 2014.
- [5] G. S. Novikov, M. E. Panov, and I. V. Oseledets, "Tensor-train density estimation," in *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, vol. 161. PMLR, 2021, pp. 1321–1331.
- [6] J. Kossaifi, A. Bulat, G. Tzimiropoulos, and M. Pantic, "T-Net: Parametrizing Fully Convolutional Nets With a Single High-Order Tensor," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 2019, pp. 7814–7823.
- [7] T. Nguyen, D. Lyakh, E. Dumitrescu, D. Clark, J. Larkin, and A. McCaskey, "Tensor Network Quantum Virtual Machine for Simulating Quantum Circuits at Exascale," *ACM Transactions on Quantum Computing*, vol. 4, no. 1, pp. 6:1–6:21, 2022.
- [8] A. Tomut, S. S. Jahromi, A. Sarkar, U. Kurt, S. Singh, F. Ishtiaq, C. Muñoz, P. S. Bajaj, A. Elborady, G. del Bimbo, M. Alizadeh, D. Montero, P. Martin-Ramiro, M. Ibrahim, O. T. Alaoui, J. Malcolm, S. Mugel, and R. Orus, "CompactifAI: Extreme Compression of Large Language Models using Quantum-Inspired Tensor Networks," 2024.
- [9] E. Stoudenmire and D. J. Schwab, "Supervised Learning with Tensor Networks," in *Advances in Neural Information Processing Systems*, vol. 29. Curran Associates, Inc., 2016.
- [10] E. M. Stoudenmire, "Learning relevant features of data with multi-scale tensor networks," *Quantum Science and Technology*, vol. 3, no. 3, p. 034003, 2018.
- [11] D. Liu, S.-J. Ran, P. Wittek, C. Peng, R. B. García, G. Su, and M. Lewenstein, "Machine learning by unitary tensor network of hierarchical tree structure," *New Journal of Physics*, vol. 21, no. 7, p. 073059, 2019.
- [12] Z.-Z. Sun, C. Peng, D. Liu, S.-J. Ran, and G. Su, "Generative tensor network classification model for supervised machine learning," *Physical Review B*, vol. 101, no. 7, p. 075135, 2020.
- [13] I. Glasser, N. Pancotti, and J. I. Cirac, "From Probabilistic Graphical Models to Generalized Tensor Networks for Supervised Learning," *IEEE Access*, vol. 8, pp. 68 169–68 182, 2020.
- [14] J. A. Reyes and E. M. Stoudenmire, "Multi-scale tensor network architecture for machine learning," *Machine Learning: Science and Technology*, vol. 2, no. 3, p. 035036, 2021.
- [15] M. L. Wall and G. D'Aguzzo, "Tree-tensor-network classifiers for machine learning: From quantum inspired to quantum assisted," *Physical Review A*, vol. 104, no. 4, p. 042408, 2021.
- [16] A. Novikov, D. Podoprikin, A. Osokin, and D. P. Vetrov, "Tensorizing Neural Networks," in *Advances in Neural Information Processing Systems*, vol. 28. Curran Associates, Inc., 2015.
- [17] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, "Unsupervised Generative Modeling Using Matrix Product States," *Physical Review X*, vol. 8, no. 3, p. 031012, 2018.
- [18] Z. Li and P. Zhang, "Shortcut Matrix Product States and its applications," 2018.
- [19] S. Cheng, L. Wang, T. Xiang, and P. Zhang, "Tree tensor networks for generative modeling," *Physical Review B*, vol. 99, no. 15, p. 155131, 2019.
- [20] S. Dolgov, K. Anaya-Izquierdo, C. Fox, and R. Scheichl, "Approximation and sampling of multivariate probability distributions in the tensor train decomposition," *Statistics and Computing*, vol. 30, no. 3, pp. 603–625, 2020.
- [21] Y. Hur, J. G. Hoskins, M. Lindsey, E. M. Stoudenmire, and Y. Khoo, "Generative modeling via tensor train sketching," *Applied and Computational Harmonic Analysis*, vol. 67, p. 101575, 2023.
- [22] Y. Peng, Y. Chen, E. M. Stoudenmire, and Y. Khoo, "Generative Modeling via Hierarchical Tensor Sketching," 2023.
- [23] A. Meiburg, J. Chen, J. Miller, R. Tihon, G. Rabusseau, and A. Perdomo-Ortiz, "Generative learning of continuous data by tensor networks," *SciPost Physics*, vol. 18, no. 3, p. 096, 2025.
- [24] B. Aizpurua, S. Palmer, and R. Orus, "Tensor networks for explainable machine learning in cybersecurity," *Neurocomputing*, vol. 639, p. 130211, 2025.
- [25] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, "Towards quantum machine learning with tensor networks," *Quantum Science and Technology*, vol. 4, no. 2, p. 024001, 2019.
- [26] H.-M. Rieser, F. Köster, and A. P. Raulf, "Tensor networks for quantum machine learning," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 479, no. 2275, p. 20230218, 2023.
- [27] B. Aizpurua, S. S. Jahromi, S. Singh, and R. Orus, "Quantum Large Language Models via Tensor Network Disentangled," 2024.
- [28] U. Schollwöck, "The density-matrix renormalization group in the age of matrix product states," *Annals of Physics*, vol. 326, no. 1, pp. 96–192, 2011.
- [29] M. B. Hastings, "Locality in Quantum and Markov Dynamics on Lattices and Networks," *Physical Review Letters*, vol. 93, no. 14, p. 140402, 2004.
- [30] F. Verstraete and J. I. Cirac, "Matrix product states represent ground states faithfully," *Physical Review B*, vol. 73, no. 9, p. 094423, 2006.
- [31] V. Murg, F. Verstraete, Ö. Legeza, and R. M. Noack, "Simulating strongly correlated quantum systems with tree tensor networks," *Physical Review B*, vol. 82, no. 20, p. 205105, 2010.
- [32] T. Hikihara, H. Ueda, K. Okunishi, K. Harada, and T. Nishino, "Automatic structural optimization of tree tensor networks," *Physical Review Research*, vol. 5, no. 1, p. 013031, 2023.
- [33] Y. Chen, Y. Hao, T. Rakthanmanon, J. Zakaria, B. Hu, and E. Keogh, "A general framework for never-ending learning from time series streams," *Data Mining and Knowledge Discovery*, vol. 29, no. 6, pp. 1622–1664, 2015.
- [34] A. Srinivasan, "Statlog (Landsat Satellite)," 1993.
- [35] M. Hopkins, E. Reeber, G. Forman, and J. Suermondt, "Spambase," 1999.
- [36] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the Support of a High-Dimensional Distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
- [37] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in *2008 Eighth IEEE International Conference on Data Mining*, 2008, pp. 413–422.
- [38] G. E. Hinton and R. R. Salakhutdinov, "Reducing the Dimensionality of Data with Neural Networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [39] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A Next-generation Hyperparameter Optimization Framework," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD '19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 2623–2631.
- [40] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: Composable transformations of Python+NumPy programs," 2018.
- [41] D. G. a Smith and J. Gray, "Opt_einsum - A Python package for optimizing contraction order for einsum-like expressions," *Journal of Open Source Software*, vol. 3, no. 26, p. 753, 2018.
- [42] H. Hohenfeld, M. Beuerle, and E. Mounzer, "dfki-ric-quantum/tn_explainability: Paper release," Aug. 2025. [Online]. Available: <https://doi.org/10.5281/zenodo.16738077>
- [43] E. P. Gerstenfeld and T. De Marco, "Premature Ventricular Contractions," *Circulation*, vol. 140, no. 8, pp. 624–626, 2019.
- [44] A. Hacine-Gharbi, P. Ravier, R. Harba, and T. Mohamadi, "Low bias histogram-based estimation of mutual information for feature selection," *Pattern Recognition Letters*, vol. 33, no. 10, pp. 1302–1308, 2012.
- [45] F. Verstraete and J. I. Cirac, "Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions," 2004.
- [46] G. Vidal, "Class of Quantum Many-Body States That Can Be Efficiently Simulated," *Physical Review Letters*, vol. 101, no. 11, p. 110501, 2008.