

Federated Learning with Quantum Enhanced LSTM for Applications in High Energy Physics

Abhishek Sawaika¹, Durga Pritam Suggiseti², Udaya Parampalli¹, Rajkumar Buyya¹

¹*qCLOUDS Lab, School of Computing and Information Systems, The University of Melbourne, Australia*

²*Birla Institute of Technology and Science Pilani, Dubai, UAE*

asawaika@student.unimelb.edu.au; f20230242@dubai.bits-pilani.ac.in; (udaya, rbuyya)@unimelb.edu.au

Abstract—Learning with large-scale datasets and information-critical applications, such as in High Energy Physics (HEP), demands highly complex, large-scale models that are both robust and accurate. Even though the computing capacity of current supercomputers is increasing day by day, there is a huge cost associated with the energy consumption of such systems. To tackle this issue and cater to the learning requirements, we envision using a federated learning framework with a quantum-enhanced model. Specifically, we design a hybrid quantum-classical long-shot-term-memory model (QLSTM) for local training at distributed nodes. It combines the representative power of quantum models in understanding complex relationships within the feature space, and an LSTM-based model to learn necessary correlations across data points. Given the computing limitations and unprecedented cost of current stand-alone noisy-intermediate quantum (NISQ) devices, we propose to use a federated learning setup, where the learning load can be distributed to local servers as per design and data availability. We demonstrate the benefits of such a design on a classification task for the Supersymmetry(SUSY) dataset, having 5M rows. Our experiments indicate that the performance of this design is not only better than some of the existing work using variational quantum circuit (VQC) based quantum machine learning (QML) techniques, but is also comparable ($\Delta \sim \pm 1\%$) to that of classical deep-learning benchmarks. An important observation from this study is that the designed framework has <300 parameters and only needs 20K data points to give a comparable performance. Which also turns out to be a $100\times$ improvement than the compared baseline models. This shows an improved learning capability of the proposed framework with minimal data and resource requirements, due to the joint model with an LSTM based architecture and a quantum enhanced VQC.

Index Terms—Federated, SUSY, HEP, LSTM, quantum, representation, distributed, machine learning, deep learning, VQC, encoding, neural network, LHC.

I. INTRODUCTION

In Modern High Energy Physics (HEP) experiments at the Large Hadron Collider (LHC), unprecedented volumes of data are generated, reaching petabyte-scale annually across all major experiments and detectors, such as ATLAS, CMS, LHCb, and ALICE [1]. Extracting rare physical events from overwhelming background processes requires increasingly advanced machinery and machine learning (ML) techniques. Various ML techniques involving deep learning, reinforcement learning, etc., have been successfully used in different tasks such as classification, forecasting, regression analysis, and control for HEP applications [2]–[4]. However, training and deploying such systems with constraints such as data scale,

privacy, and distributed experimental environments is crucial yet challenging [5].

Recent studies in quantum computing have shown the advantages of QML-based algorithms, offering new representational capabilities for learning complex correlations in scientific data [6], [7]. More recently, works exploring QML in HEP using techniques such as variational quantum classifiers, quantum kernels in support vector classifiers, and quantum tensor networks have shown promising results [8]–[11]. This motivates more such explorations in the intersection of QML and HEP.

TABLE I: Quantum Hardware Roadmaps (2025–2029)

Provider	Tech	2025–26	2027–28	2029
IBM [12]	SC	120	1K+	200L
Google [13]	SC	105 (Willow)	EC	FTQC
Atom [14]	NA	1,180	10K	FTQC
Pasqal [15]	NA	140+	200+	100–200L
IonQ [16]	TI	64	10–20K	80KL

Note: SC = Superconducting, NA = Neutral Atoms, TI = Trapped Ions, L = Logical qubits, EC = Error correction, FTQC = Fault-tolerant QC, K = thousands.

Despite rapid progress toward fault-tolerant quantum computing, see Table I, current computers still remain in the NISQ regime [17], constrained by noise and limited error-correction techniques. This demands innovative solutions to address the learning requirements for large-scale problems in HEP. Federated learning (FL) [18] solves this issue by enabling a collaborative approach of learning with data distributed across multiple models (nodes).

Though FL is extensively studied in classical settings [19], its integration with quantum models is still a new area of research which needs to be explored further, especially for industrial or large-scale level applications [20]–[22]. Specifically, to the best of our knowledge, it has not been well explored for HEP applications, where one can leverage data across different institutions and countries worldwide and learn with limited resources. In this work, we:

- Design of a hybrid quantum-classical LSTM model (QLSTM) for learning complex correlations with less data.
- Propose a federated framework for learning in HEP applications. This could help in efficient workload distribution and resource sharing.
- Empirically demonstrate the usability and improvement gains of the proposed model for a classification task on

the SUSY dataset [23].

The rest of the paper is organized as follows: We will introduce the concepts of quantum machine learning, the proposed quantum-enhanced LSTM model, and federated learning for HEP in Sections II, III, IV, respectively. We will discuss the results obtained from the experiments conducted in Section V, and we provide a concluding remark in Section VI. The source code for this work can be found at https://github.com/z-ax-qsc/fed_hep.

TABLE II: Similar works comparison in HEP applications

Related Work	Quantum Model	Federated Learning	LSTM
Wu, et al. [8]	✓	✗	✗
Terashi, et al. [9]	✓	✗	✗
Felser, et al. [10]	✓	✗	✗
Tripathi, et al. [11]	✓	✗	✓
Baldi, et al. [23]	✗	✗	✗
This work	✓	✓	✓

II. QUANTUM MACHINE LEARNING

Quantum computing uses quantum-mechanical principles to process information and encode more information (infinite possibilities) in these amplitudes, unlike a classical bit, which can only encode two states [24]. A *qubit* (quantum bit) represents a superposition of basis states $|0\rangle$ and $|1\rangle$. A quantum state of a single qubit system is expressed as:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad (1)$$

where, α, β are complex probability amplitudes equating the normalization condition $|\alpha|^2 + |\beta|^2 = 1$.

This helps in encoding more information (infinite possibilities) in these amplitudes, as opposed to a classical bit, which can only encode two possible values. For a system of n qubits, the quantum state space grows exponentially with dimension 2^n , giving higher representation capabilities for large data encoding in smaller quantum systems. This acts as a motivation for using quantum algorithms such as QML which, in theory, can enable us to study high-dimensional feature spaces while using relatively fewer physical degrees of freedom [25]. This motivates our intuition for exploring a quantum model for HEP use cases.

A QML workflow can be represented as a five-stage process (see Figure 1). This starts with an **encoding layer** to map classical data vector $\mathbf{x} = (x_1, x_2, \dots, x_m)$ into quantum states $\{|\psi_i\rangle\}_{i=1}^n$, via an encoding scheme $G : \mathbf{x} \rightarrow |\psi\rangle$. Once encoded into the quantum states, the data is passed through a **parametrized circuit** $F(\theta)$ for learning. This is analogous to a feed-forward network with trainable parameters. The final state $|\Psi\rangle$ is **measured** to generated classical information $\mathbf{O}$. The measurements are usually linked to the loss of the objective function defined for the chosen learning task, using a **post processing** function $\mathbf{P}$. The whole process is repeated multiple times through a **classical optimizer** for an iterative update of parameters θ .

A variational quantum circuit (VQC) usually comprises the encoding, parametrized circuit, and the measurement layers.

See Figure 3 for a sample VQC. Operations performed in a generic quantum circuit (algorithm) are denoted by a sequence of discrete unitaries ($2^n \times 2^n$) acting on a n qubit system [24].




Fig. 1: QML workflow

A. Data Encoding Schemes

Recent surveys and comparative analyses highlight trade-offs among the encoding methods commonly used in QML [26], [27]. In this section, we will describe two of the most commonly used techniques, namely angle encoding and amplitude encoding.

1) *Angle Encoding (General Technique)*: Angle encoding maps input features into rotation angles of single-qubit operations. This is done by assigning each data feature to a rotation angle on each *qubit*. E.g. for a feature vector $\mathbf{x} = [x_1, x_2, \dots, x_N]$, the encoded quantum state is defined as:

$$|\psi(\mathbf{x})\rangle = \bigotimes_{j=1}^N R_P(x_j) |0\rangle^{\otimes N}, \quad (2)$$

where $R_P(\theta)$ is a Pauli rotation about an axis $\mathbf{P} = \{X, Y, Z\}$ with angle θ .

2) *Amplitude Encoding (Quantum Data Compression)*: Here the classical data vector is mapped into the probability amplitudes of a quantum state, as:

$$|\psi(\mathbf{x})\rangle = \frac{1}{\|\mathbf{x}\|} \sum_{i=1}^{2^n} x_i |i\rangle, \quad (3)$$

where $\|\mathbf{x}\| = \sqrt{\sum_i x_i^2}$, $|i\rangle$ is a n -qubit basis state and $n > \log_2 N$.

Even with higher *qubit* scaling, $O(N)$ vs $O(\log_2 N)$, angle encoding can often yield superior classification accuracy. In empirical studies [28], angle encoding achieved 82% accuracy compared to 75% for amplitude encoding when done on a 4-feature toy dataset. Though amplitude encoding seems to be a viable option for applications and datasets like the LHC, where high-dimensional particle features can be compactly represented for efficient learning, there is no evidence to generalize this fact.

Other compression techniques include a hybrid amplitude encoding, where qubits are divided into blocks and amplitude encoding is applied to each block separately. This balances the qubit overhead with circuit depth [29]. A dense angle-encoding scheme [30] simultaneously encodes two features into the phase and amplitudes of a quantum state, yielding a compression factor of 2 compared to the basic angle encoding mentioned above. One could also explore other compression techniques mentioned in [31], [32].




Fig. 2: Projection of a ground state $|0\rangle$ onto the Bloch sphere using (a) only R_x rotations and (b) with a fully entangled layer and universal rotations using R_y and R_z operations. One can see the difference in representation capabilities of the two techniques, where the latter captures more space in this Hilbert space, and hence could be more powerful for certain learning tasks.

B. Parametrized Quantum Circuits

Parametrized quantum circuits (PQCs) are the core component of variational quantum algorithms (VQAs) and QML models. The choice of PQCs determines how quantum states can be represented, the model's trainability, and the quantum circuit's hardware efficiency [33].

A general framework for a PQC with parameters $\theta = [\theta_1, \theta_2, \dots, \theta_p]$ for an n -qubit system, with L layers (circuit depth), can be written as:

$$U(\theta) = \prod_{l=1}^L U_{\text{ent}} U_{\text{param}}^{(l)}(\theta^{(l)}),$$

where:

- $U_{\text{param}}^{(l)}$ represents a combination of single-qubit rotations (e.g., R_Y , R_Z).

$$U_{\text{param}}^{(l)}(\theta^{(l)}) = \bigotimes_{i=1}^n R(\theta_i^{(l)})$$

- U_{ent} represents entangling gates (e.g., CNOT, CZ) that create correlations between qubits.

such that the total number of parameters $p = O(nL)$ [6].

This generic framework helps in designing any circuit of our choice. One can also design a Hardware-Efficient Ansatz as a generic extension of this framework to minimize gate overhead by using only the native gates available on quantum hardware [34].




Fig. 3: A 3-qubit VQC with angle encoding using R_x rotations and a fully entangled parametrized circuit with R_z and R_y operations. X 's are the input data, whereas θ 's are trainable parameters.

III. QUANTUM ENHANCED LSTM MODEL

Recent studies have shown that integrating variational quantum circuits within a classical recurrent network architecture can be advantageous in general [20], [35]. A traditional QLSTM [36] combines VQCs within an LSTM cell to increase parallelism in terms of feature encoding, capabilities of high-dimensional representations and learning complex correlations. This motivates us to explore and design a quantum enhanced LSTM architecture for HEP applications.

In our case, we additionally add linear layers before VQCs and at the output layer to compress input/latent features, enabling a feasible implementation of large-dimensional use cases on near-term quantum devices (see Figure 4). This is useful in cases such as HEP for learning through compact representations of large-scale data. Furthermore, this integration of classical and quantum models allows us to work on the upcoming near-term models as they run on quantum circuits and classical optimization algorithms.

Initially, the input vector $\mathbf{X}_t$ is projected to the quantum embedding space using a linear layer, s.t. :

$$\mathbf{z}_t = W_c \mathbf{X}_t + \mathbf{b}_c, \quad (4)$$

Then the projected features are encoded into a quantum state using an angle embedding, followed by parameterized




Fig. 4: Quantum Enhanced LSTM model

entangling layers, as shown in Figure 3, and the final outcome is measured in Z basis to give:

$$\mathbf{q}_t = \langle \psi(\mathbf{z}_t, \boldsymbol{\theta}) | \hat{Z} | \psi(\mathbf{z}_t, \boldsymbol{\theta}) \rangle, \quad (5)$$

The measurement results are mapped back to the classical space using a linear layer as:

$$\mathbf{s}_t = W_q \mathbf{q}_t + \mathbf{b}_q. \quad (6)$$

This whole process is done in parallel for all four gates, giving rise to $\mathbf{s}_t^f, \mathbf{s}_t^i, \mathbf{s}_t^o, \mathbf{s}_t^c$, such that the memory and hidden state updates are given by:

$$\mathbf{c}_t = \mathbf{s}_t^f \odot \mathbf{c}_{t-1} + \mathbf{s}_t^i \odot \mathbf{s}_t^c, \quad (7)$$

$$\mathbf{h}_t = \mathbf{s}_t^o \odot \tanh(\mathbf{c}_t). \quad (8)$$

Finally, the output is given by

$$\mathbf{Y}_t = W_o \mathbf{c}_t + \mathbf{b}_o. \quad (9)$$

where, $W_c, \mathbf{b}_c, W_q, \mathbf{b}_q, W_o, \mathbf{b}_o$ and $\boldsymbol{\theta}$ are learnable parameters of classical, quantum layers.

IV. FEDERATED LEARNING FOR HIGH ENERGY PHYSICS

FL has been at the forefront as an effective learning paradigm for collaboratively training multiple models across distributed data sources without the need to share data. This, when combined with quantum principles, could also provide privacy guarantees while maintaining computational efficiency [22], [37], [38].

A typical quantum federated learning (QFL) system with a quantum learned model consists of N quantum clients, each maintaining local quantum/classical datasets, and a central server that aggregates global model parameters [22]. In our HEP application use case, this can be represented as shown in Figure 5. Here, each collider represents a client, generating




Fig. 5: Federated learning framework for HEP experiments conducted at different colliders.

data through local experiments, having quantum-classical capabilities to train the local model, and participating in global learning through a central server.

The overall architecture of a federated quantum learning framework is inspired by the work in [18], [20] where the role of individual components in HEP is as follows:

- **Detectors:** These are multiple sensor nodes used to detect and measure different properties of the particles produced in a collider.
- **Data Storage:** It is a storage unit for preprocessing and cleaning of intermittent data collected by detectors.
- **Local Model:** This is the hybrid quantum-classical computing resource to host local models for training. After each round (or decided by participating nodes) of training, the trained weights $\boldsymbol{\theta}$ are shared with a global server.
- **Global server:** The sole task of the global server is to merge the weights received from different nodes and pass them back for synchronization.

This QFL-based framework can be uniquely used to address key issues such as data preservation and resource requirements, given that multiple countries (institutes) participate in such big-data applications for HEP [39]. For computational scalability, and in QFL, each institution trains a quantum classifier on its local data and compute resources, and then aggregates only the trained model parameters through a central server.

When learning with homogeneous data, local models with a common structure can be used across multiple nodes. Whereas, in case of heterogeneous data structures across different detectors, a more complex merging strategy needs to be deployed at the global server, based on the particular learning requirements [40].

V. PERFORMANCE EVALUATION

To evaluate the performance of our framework, we perform experiments with different numbers of nodes, a simple VQC, and the designed QLSTM models, and compare the results with prior work [9], [23]. We use the AUC (area under the curve) of the ROC (receiver operating characteristic) to measure model performance. This gives us an estimate of the model's ability to distinguish between different classes. The terminologies node, model, and server are used interchangeably in this section.

A. The Dataset

For our analysis, we have used a classification problem to identify supersymmetry in data generated by the LHCb experiments [41]. From an operational perspective, the LHCb experiment records pp (proton-proton) collisions delivered in runs of several hours, resulting in data containing high-level reconstructed objects such as tracks, vertices, identified hadrons, and jets, together with derived kinematic and topological variables.

The SUSY dataset [42] is generated from a framework in theoretical physics that extends the Standard Model by introducing a symmetry principle connecting fermions and bosons. At LHC, SUSY searches typically targets final states with multiple jets, leptons, and significant missing transverse momentum from weakly interacting lightest supersymmetric particles. The final state in the detector is two charged leptons (ll), and the missing momentum is carried off by the invisible particle [23].

In our setting, the SUSY dataset serves as the HEP framework, with its features treated as structured inputs to a federated quantum-enhanced LSTM model. Quantum circuits embedded in the recurrent cell learn nonlinear correlations across features and across different detection events. We conduct experiments with two different scenarios, where in the first one we consider all the 18 features, while in the second one we only use the following 7 features, namely 'lepton 1 pT', 'lepton 2 pT', 'missing energy magnitude', 'M_TR_2', 'M_Delta_R', 'lepton 1 eta', 'lepton 2 eta', as identified in [9]. These features are chosen based on their significance in successfully distinguishing between signal and background processes. These also represent the low-level features mentioned in [23].

B. Experimental Setup

Given the limitations of quantum simulation on classical machines [43], we use only 20K data points from the chosen dataset and the PennyLane 'lightning-qubit' simulator for our experiments. Simulations are run on an M4 Pro chip with 24GB of RAM and a 16-core, 2.5 GHz processor. We divide the dataset into an 80:20 split for training and testing.

For the purpose of federated learning simulation, we divide the dataset into IID partitions and send them to different copies of QLSTM, representing individual nodes, for training. See Figure 6 for a diagrammatic representation. Individual VQCs in QLSTM are designed using angle encoding and multiple




Fig. 6: Distributed learning setup with a single collider collecting a large amount of data through multiple detectors, and a data distributor which segregates this to multiple nodes (models) for training.

layers of fully-entangled parametrized circuits [6], see Figure 3 for a sample VQC. It has 6 qubits for feature encoding and 4 trainable layers. For classical LSTM we have used fully connected NNs in place of the VQCs. We use the Adam optimizer, the binary cross-entropy loss function, and a learning rate of 0.01 for training. Models are trained for 30 epochs and 5 global rounds. We conduct 10 random experiments and also report the marginal deviations for respective model performances.

C. Results Analysis

To analyze the effect of the federated learning setup, we compare the performance with different numbers of nodes. This is done for an LSTM, QLSTM and a single VQC model. From Figure 7, we can see that the performance decreases slightly with the size of network, which is expected. This is mainly because, in our simulation, as depicted in Figure 6, we divide the entire dataset into multiple chunks; hence, the information reaching each model decreases with the number of nodes. Though the overall decrease is not that significant, with $\Delta < 1\%$ for most of the cases. This provides a fair trade-off for distributed training in large-scale data applications such as this. This decrease is steeper for the simple VQC model, whereas it's almost flat for QLSTM. This indicates that QLSTM can preserve complex relations even with less data. Though, in cases where each node contribute individually with their local data, a fair tradeoff needs to be identified in terms of their participation in the learning process, without hampering the performance of the global model [44], [45].




Fig. 7: Model performance with different number of nodes used in distributed training for federated learning simulation. Each node has a local copy of the model used for training. Results reported here are for the globally merged inference.

In Figure 8, we report the ROC curves for different models trained with all(18) and limited(7) features. We can clearly see that the model with complete features performed better than the model with limited features. This is also evident from the results in Figure 7. Apart from an exception in accuracy, i.e., QLSTM-7 (0.821) > QLSTM-18 (0.800), see Table III. The model trained with all features is also better than that trained with only 7 features in terms of accuracy. This is reversed from the trend reported in [9], where model with lesser features performed better. The difference in ACU values is less ($\approx 0.5\%$) for QLSTM as compared to $\approx 1\%$ for VQC, which indicates that use of less features (statistically significant) is still reasonable for training such hybrid quantum models.

TABLE III: Test accuracy for VQC and QLSTM models trained using all(18) and limited(7) features. The results for federated learning are inferred from a global model with a 3-node setup.

Test Accuracy	$N_{features} = 7$	$N_{features} = 18$
LSTM	0.777 ± 0.007	0.790 ± 0.005
LSTM-Federated (3)	0.775 ± 0.007	0.777 ± 0.005
VQC	0.739 ± 0.006	0.784 ± 0.004
VQC-Federated (3)	0.735 ± 0.006	0.778 ± 0.004
QLSTM	0.821 ± 0.004	0.800 ± 0.003
QLSTM-Federated (3)	0.812 ± 0.004	0.818 ± 0.003


Our experiments also indicate that QLSTM performs better than a single VQC model; see Figures 7 and 8. The best AUC values for QLSTM are (0.88, 0.874), and that for VQC is (0.823, 0.812), under training with (18, 7) features, respectively. This is also true for accuracy values for QLSTM = (0.880, 0.821) and VQC = (0.784, 0.739), see Table III. This could indicate a potential correlation between different data points, which helped enhance learning with the QLSTM-based model. This also reduces the need for a large training dataset.

A similar trend was also observed with a classical LSTM-based model, where LSTM performed better than the VQC. This suggests that a recurrent model architecture can be more efficient in learning with limited data. However, we observed that the LSTM model didn't outperform QLSTM, with $\sim 3\%$ degradation in AUC and $\sim 8\%$ in accuracy, suggesting that a quantum model can improve learning through its complex non-linear correlation mapping using entanglement. An important trade-off to mention here is that the runtime for training such a hybrid QLSTM model on a classical machine was around $\sim 3\times$ higher than LSTM and $\sim 2\times$ higher than the VQC model. However, such tradeoff can be reduced to some extent by using a quantum computer for direct quantum simulation of large-scale models.


While comparing the results from existing works, we observe that our best model using QLSTM performs better than most of the cases in [9], and is comparable from classical benchmarks in [23]. Best AUC for QCL [9] (~ 0.825) < QLSTM-7 (~ 0.874). AUC values of the QLSTM model trained here $\in [0.86, 0.88]$, which lies only within $\pm 1\%$ range of the best models (neural-network, deep-learning) reported in [23], with AUC $\in [0.87, 0.89]$. Whereas, for LSTM, we observed more variation than QLSTM when compared against results in [23]. AUC values for the LSTM model $\in [0.83, 0.86]$, leading to a degradation of $\sim 3\%$.

It is important to note that the models used in this work have < 300 parameters. This is significantly lower, even for the worst-performing model in [23], with $\sim 300K$. Moreover, we have used only 20K data points for training, which is again lower than the 5M used in [23]. This gives us an indication of high representational capabilities and complex interaction learning with hybrid quantum-enhanced QLSTM models. Hence, this $\sim 1\%$ difference in AUC values can be overpowered by the $100\times$ reduction in data size and trainable

parameters achieved for the proposed design. Even the model under federated/distributed learning has similar results, which justifies the use of this framework for large-scale applications, such as in HEP.



(a) Number of features = 18



(b) Number of features = 7

Fig. 8: ROC curve comparing performance on test dataset for models trained on different numbers of features, i.e., 18 and 7 for sub-figures (a) and (b), respectively. Here, the results for federated learning are inferred from a global model with a 3-node setup.

VI. CONCLUSIONS AND FUTURE WORK

We engineered a hybrid quantum-enhanced LSTM model and demonstrated the feasibility of a federated learning framework for large-scale HEP applications. Our results on the SUSY dataset demonstrated that a hybrid QLSTM outperforms other related works using basic quantum models. While it showed only $\sim 1\%$ degradation compared to large classical deep learning benchmarks, there was a significant trade-off ($\sim 100\times$) in data size and the number of model parameters used for learning. This gives us a hint on high representation capabilities and complex interaction learning in a hybrid quantum-classical model, such as a QLSTM, even with limited data and model size. The federated-learned QLSTM also performed similar to a stand-alone model, justifying the use of such framework in constrained environments. For demonstration, we used a single data source and distributed it across multiple nodes to simulate a homogeneous-IID case of federated learning. Though, it is to be noted that this can be easily extended to learning with multiple data sources simulating non-IID distributions. As part of future work, we will explore the impact of other compact encoding techniques in QML models and other heterogeneous learning scenarios in high-energy/particle physics applications. We also plan to analyze this model on a quantum hardware and study the effect of various noise characteristics.

ACKNOWLEDGMENTS

This work is supported by the University of Melbourne and Maitri scholarships from the Department of Foreign Affairs and Trade, Government of Australia.

REFERENCES

- [1] CERN, “Cern data centre passes the 200-petabyte milestone,” <https://home.cern/news/news/computing/cern-data-centre-passes-200-petabyte-milestone>, 2017. Accessed January 2026.
- [2] G. Karagiorgi, G. Kasieczka, S. Kravitz, B. Nachman, and D. Shih, “Machine learning in the search for new fundamental physics,” *Nature Reviews Physics* 2022 4:6, vol. 4, pp. 399–412, 5 2022.
- [3] W. B. He, Y. G. Ma, L. G. Pang, H. C. Song, and K. Zhou, “High-energy nuclear physics meets machine learning,” *Nuclear Science and Techniques* 2023 34:6, vol. 34, pp. 88–, 6 2023.
- [4] D. Guest, K. Cranmer, and D. Whiteson, “Deep learning and its application to LHC physics,” *Annual Review of Nuclear and Particle Science*, vol. 68, pp. 161–181, 10 2018.
- [5] P. W. Hatfield, J. A. Gaffney, G. J. Anderson, S. Ali, L. Antonelli, S. Başgeçmez du Pree, J. Citrin, M. Fajardo, P. Knapp, B. Kettle, B. Kustowski, M. J. MacDonald, D. Mariscal, M. E. Martin, T. Nagayama, C. A. Palmer, J. L. Peterson, S. Rose, J. J. Ruby, C. Shneider, M. J. Streeter, W. Trickey, and B. Williams, “The data-driven future of high-energy-density physics,” *Nature* 2021 593:7859, vol. 593, pp. 351–361, 5 2021.
- [6] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, “Supervised learning with quantum-enhanced feature spaces,” *Nature* 2019 567:7747, vol. 567, pp. 209–212, 3 2019.
- [7] M. Sajjan, J. Li, R. Selvarajan, S. H. Sureshbabu, S. S. Kale, R. Gupta, V. Singh, and S. Kais, “Quantum machine learning for chemistry and physics,” *Chem. Soc. Rev.*, vol. 51, pp. 6475–6573, 2022.
- [8] S. L. Wu, J. Chan, W. Guan, S. Sun, A. Z. Wang, C. Zhou, D. Livescu, and M. Carena, “Application of quantum machine learning using the quantum kernel algorithm to high energy physics analysis at the lhc using ibm simulators and quantum hardware,” *Physical Review Research*, vol. 3, p. 033221, 2021.

- [9] K. Terashi, M. Kaneda, T. Kishimoto, M. Saito, R. Sawada, and J. Tanaka, "Event classification with quantum machine learning in high-energy physics," *Computing and Software for Big Science*, vol. 5, no. 1, p. 2, 2021.
- [10] T. Felser, M. Trenti, L. Sestini, A. Gianelle, D. Zuliani, D. Lucchesi, and S. Montangero, "Quantum-inspired machine learning on high-energy physics data," *npj Quantum Information* 2021 7:1, vol. 7, pp. 111–, 7 2021.
- [11] S. Tripathi, H. Upadhyay, and J. Soni, "A quantum machine learning-based predictive analysis of CERN collision events," *Scientific Reports* 2025 16:1, vol. 16, pp. 682–, 12 2025.
- [12] IBM Quantum Team, "Ibm quantum roadmap 2025: Practical quantum computing era," <https://www.ibm.com/roadmaps/quantum/>, 2025.
- [13] H. Neven, "The quantum echoes algorithm breakthrough," <https://blog.google/innovation-and-ai/technology/research/quantum-echoes-willow-verifiable-quantum-advantage/>, 2025. Google Research Blog, October 2025.
- [14] Atom Computing, "Highly scalable quantum computing with neutral atoms," <https://atom-computing.com/wp-content/uploads/2025/01/Atom-Computing-Whitepaper-2025.pdf>, 2025. Whitepaper, January 2025.
- [15] Pasqal, "Pasqal releases 2025 roadmap showcasing upgradable architecture toward fault-tolerant quantum computing," <https://www.pasqal.com/wp-content/uploads/2025/10/Pasqal-Roadmap-2025.pdf>, 2025. Roadmap document, June 2025.
- [16] IonQ, "Ionq hits aq 64 milestone ahead of schedule," <https://www.ionq.com/blog/ionq-hits-aq-64-milestone-ahead-of-schedule-and-sets-its-sights-even-higher>, 2025. IonQ Blog, October 2025.
- [17] J. Preskill, "Quantum computing in the nisc era and beyond," *Quantum*, vol. 2, p. 79, 2018.
- [18] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 54, pp. 1273–1282, 2017.
- [19] P. Kairouz, H. B. McMahan, B. Avenet, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al., "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [20] A. Sawaika, S. Krishna, T. Tomar, D. P. Suggisetti, A. Lal, T. Shrivastav, N. Innan, and M. Shafique, "A privacy-preserving federated framework with hybrid quantum-enhanced learning for financial fraud detection," in *Proceedings of the 2025 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 01, pp. 1578–1588, 2025.
- [21] N. Innan, M. A.-Z. Khan, A. Marchisio, M. Shafique, and M. Bennai, "Fedqnn: Federated learning using quantum neural networks," *arXiv preprint arXiv:2403.10861*, 2024.
- [22] S. Y.-S. Chen and S. Yoo, "Federated quantum machine learning," *Entropy (Basel)*, vol. 23, p. 460, Apr 2021.
- [23] P. Baldi, P. Sadowski, and D. Whiteson, "Searching for exotic particles in high-energy physics with deep learning," *Nature Communications*, vol. 5, 2014.
- [24] B. Zohuri, "What is quantum computing and how it works," *Journal of Material Sciences & Manufacturing Research*, vol. 3, pp. 3–5, 2020.
- [25] J. Biamonte et al., "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, 2017.
- [26] M. B. Pande, "A comprehensive review of data encoding techniques for quantum machine learning problems," in *Proceedings of the Second International Conference on Emerging Trends in Information Technology and Engineering (ICETITE)*, pp. 1–7, 2024.
- [27] M. Rath and H. Date, "Quantum data encoding: a comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy," *EPJ Quantum Technology*, vol. 11, p. 72, 12 2024.
- [28] S. Kumar and A. Kumar, "Comparative study of amplitude versus angle encoding in variational quantum classifiers," *Physics Journal*, vol. 7, no. 2, pp. 49–193, 2025.
- [29] Z. Li, X. Fu, L. Meng, and R. Du, "A repetitive amplitude encoding method for enhancing the mapping ability of quantum neural networks," *Nature Scientific Reports*, vol. 15, p. 17651, August 2025.
- [30] E. Ovalle-Magallanes et al., "Quantum angle encoding with learnable rotation applied to quantum machine learning," *Neural Computing and Applications*, vol. 35, pp. 12345–12358, 2023.
- [31] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, "Data re-uploading for a universal quantum classifier," *Quantum*, vol. 4, p. 226, 2 2020.
- [32] R. Larose and B. Coyle, "Robust data encodings for quantum classifiers," *Physical Review A*, vol. 102, p. 032420, 9 2020.
- [33] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, "Expressive power of parametrized quantum circuits," *Physical Review Research*, vol. 2, p. 033125, 2020.
- [34] L. Leone, J. Gibbs, L. Di Ruscio, N. Sangouard, and M. Cerezo, "On the practical usefulness of the hardware efficient ansatz," *Quantum*, vol. 8, p. 1395, 2024.
- [35] S. Z. Khan et al., "Quantum long short-term memory (qlstm) vs. classical lstm: A comparative analysis for solar power forecasting," *Frontiers in Physics*, vol. 12, p. 1439180, 2024.
- [36] S. Y.-C. Chen, S. Yoo, and Y.-L. L. Fang, "Quantum long short-term memory," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8622–8626, 2022.
- [37] W. Li, S. Lu, and D. L. Deng, "Quantum federated learning through blind quantum computing," *Science China Physics, Mechanics & Astronomy* 2021 64:10, vol. 64, pp. 100312–, 9 2021.
- [38] M. Chehimi and W. Saad, "Quantum federated learning with quantum data," *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 2022-May, pp. 8617–8621, 2022.
- [39] T. Basaglia, M. Bellis, J. Blomer, J. Boyd, C. Bozzi, D. Britzger, S. Campana, C. Cartaro, G. Chen, B. Couturier, et al., "Data preservation in high energy physics," *The European Physical Journal C*, vol. 83, no. 9, p. 795, 2023.
- [40] M. Ye, X. Fang, B. Du, P. C. Yuen, and D. Tao, "Heterogeneous Federated Learning: State-of-the-art and Research Challenges," *ACM Computing Surveys*, vol. 56, 3 2024.
- [41] C. Brust, A. Katz, S. Lawrence, and R. Sundrum, "SUSY, the Third Generation and the LHC," *Journal of High Energy Physics* 2012 2012:3, vol. 2012, pp. 103–, 3 2012.
- [42] D. Whiteson, "SUSY." UCI Machine Learning Repository, 2014. DOI: <https://doi.org/10.24432/C54606>.
- [43] K. Young, M. Scese, and A. Ebnenasir, "Simulating quantum computations on classical machines: A survey," *arXiv preprint arXiv:2311.16505*, 2023.
- [44] L. Fu, H. Zhang, G. Gao, M. Zhang, and X. Liu, "Client Selection in Federated Learning: Principles, Challenges, and Opportunities," *IEEE Internet of Things Journal*, vol. 10, pp. 21811–21819, 12 2023.
- [45] J. Zhao, X. Zhu, J. Wang, and J. Xiao, "Efficient client contribution evaluation for horizontal federated learning," *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 2021-June, pp. 3060–3064, 2021.