



Research article

Quantum-enhanced deep learning for severe weather prediction: a 10-qubit QCNN-LSTM for bow echo forecasting

Hélène Canot^{1,*}, Philippe Durand² and Emmanuel Frénod¹

¹ Université de Bretagne-Sud, UMR 6205, Laboratoire de Mathématiques de Bretagne Atlantique, Vannes F-56000, France

² Conservatoire National des Arts et Métiers, Département de Mathématiques et Statistiques (M2N), 292 rue Saint-Martin, Paris Cedex 75141, France

* **Correspondence:** Email: helene.canot@orange.fr.

Abstract: We present a quantum-enhanced deep learning framework for short-term forecasting of extreme convective weather systems, specifically targeting the formation of bow echo structures through prediction clouds. Our approach integrates a quantum convolutional neural network (QCNN) into a classical convolutional neural network-long short-term memory (CNN-LSTM) pipeline, incorporating a 10-qubit variational quantum circuit executed on the *lightning.qubit* simulator. The model is trained using heterogeneous meteorological data (satellite images, convective available potential energy (CAPE) fields, lightning activity) from the extreme storm that hit the Corsican coast on August 18, 2022, preceding the convective event and evaluated during the critical intensification window (0700–0800 h). We compare this hybrid quantum model with classical CNN-LSTM and CNN-LSTM-Transformer architectures using three complementary metrics: Root mean squared error (RMSE), structural similarity index (SSIM), and Wasserstein distance. The quantum model demonstrates superior robustness to input noise and adversarial perturbations fast gradient sign method (FGSM), maintaining spatial coherence and stable prediction under degradation. These findings highlight the capacity of quantum circuits to encode more resilient representations for meteorological inference. This study provides a concrete application of hybrid QCNN architectures in the domain of extreme weather forecasting and lays the groundwork for future integration of topological data analysis to assess the preservation of topological features in quantum predictions.

Keywords: quantum neural networks; extreme weather forecasting; bow echo; QCNN; adversarial robustness; deep learning; entangled representations

1. Introduction

The term *bow echo* refers to the radar signature of a convective structure shaped like an arc. Bow echoes typically form from squall lines and mesoscale convective systems (MCS), especially in environments with moderate to strong wind shear in the lower 2–3 kilometers of the atmosphere. These formations are primarily associated with a descending rear inflow jet (RIJ) that spreads near the surface, leading to strong straight-line winds, heavy precipitation, lightning, and occasionally hail or tornadoes. A distinctive feature of bow echoes is the formation of vortices at both ends of the arc, commonly referred to as bookend vortices. These vortices can intensify the rear jet stream and increase the potential for damaging winds. Numerous studies have investigated the formation and dynamics of bow echoes: [1–4] analyzed their development, while [5–7] proposed various automated detection methods.

A major bow echo event occurred on August 18, 2022, near the Balearic Islands. The convective system moved rapidly northeastward and struck Corsica with wind gusts exceeding 200 km/h, causing five fatalities and widespread destruction. The storm later impacted northern Italy and Austria. This event exhibited the typical characteristics of a *derecho*—a long-lived and intense convective storm complex marked by arc-shaped radar signatures and high-impact wind bursts [8, 9]. The associated radar signatures generally exhibit linear structures, with arcuate or crescent-shaped segments characterized by sharp reflectivity gradients along the convex edge [10–12]. These features often appear as undulations within the main precipitation zone, followed by a bulging pattern associated with meso-depressive systems.

Historically, the identification of bow echoes has relied on Doppler radar reflectivity imagery. Przybylinski [2] described the reflectivity patterns of bow echoes linked to severe weather. Bernatz [13] presented examples of radar signatures, including bow echoes, and discussed their role in identifying intense thunderstorms. Fujita [1] introduced radar-based recognition of bow echoes associated with downbursts. Lee et al. [14] conducted a detailed quantitative study on the evolution and structure of convective bow echoes (CBEs). Parker and Johnson [15] investigated the development of mesoscale convective systems with an emphasis on arc-shaped formations. Klimowski et al. [3] examined the early evolution of bow echoes via radar observations and identified three primary formation modes: Weakly organized cells, squall lines, and supercells. More recently, Patil et al. [7] introduced several algorithmic approaches for detecting and monitoring convective systems. Kamani et al. [5] developed an automatic framework for identifying bow echo patterns in radar imagery using a skeletonization technique and shape-matching algorithms.

Table 1 summarizes prior machine learning/deep learning (ML/DL) work relevant to convective hazards and extreme events, motivating our quantum-enhanced approach.

In parallel, several studies have investigated the potential of quantum convolutional neural networks (QCNNs) for object detection and classification tasks. Notable contributions include the works of Cong et al. [28], Moreira et al. [29], Herrmann et al. [30], Song et al. [31], and Gong et al. [32], who explored various quantum-inspired approaches to pattern recognition and image analysis.

In recent years, quantum machine learning (QML) and quantum neural networks (QNNs) have been applied to meteorological forecasting. Pastori et al. [33] compared multiple QNN architectures with classical neural networks, training them on data derived from high-resolution Icosahedral Non-hydrostatic (ICON) climate model simulations. Their findings indicate that QNNs can achieve

Table 1. Synthesis of ML/DL approaches for meteorology and hydrometeorology (classification, detection, nowcasting, and forecasting).

Study	Method / architecture	Data	Task / phenomenon	Main limitation (reported or typical)
Elhoseiny et al. [16]	Convolutional neural network (CNN) for image classification	Sky/weather images (webcams, satellite, in situ)	Weather labels (sunny/cloudy, etc.)	Label granularity; no explicit temporal dynamics
Trock et al. [17]	ML for attribution / impact quantification	Climate time series + human activity indicators	Anthropogenic impact on extreme heat events	Attribution is sensitive to confounders; interpretability
Liu et al. [18]	Deep CNN for pattern detection	Large-scale climate datasets (satellite, reanalyses)	Tropical cyclones, atmospheric rivers, fronts	Cross-basin generalization; rarity of events
Gagne et al. [19]	CNN	Radar reflectivity + environmental variables	Severe hail classification / diagnosis	Primarily spatial; limited temporal context
Lagerquist et al. [20]	CNN for occurrence forecasting	Numerical Weather Prediction (NWP) fields + radar/observations	Tornado occurrence (nowcasting ~1 h)	Severe class imbalance; false alarms
Ayzel et al. [21]	U-Net / FCN (Fully Convolutional Network) (nowcasting)	Radar sequences	Short-term precipitation nowcasting	Short horizon; sensitivity to radar noise
Price et al. [22]	Foundation/generative artificial intelligence (AI) model (GenCast)	Four decades of meteorological data (multi-source)	Forecasts up to 15 days (including extremes)	Black box; massive data; Out-of-Distribution (OOD) extremes remain difficult
Mounier et al. [23]	U-Net for detection	Application de la Recherche à l'Opérationnel à Méso-Echelle – Ensemble Prediction System (AROME-EPS) high-resolution model outputs (France)	Bow-echo detection	Fine-scale smoothing; no explicit sequence modeling
Gauch et al. [24]	Long short-term memory (LSTM)(time series)	Hydrology + meteorological forcings	Streamflow prediction → precipitation applications	Transferability across basins/time; dependence on exogenous inputs
Kratzert et al. [25]	LSTM (multi-basin learning)	Hydro-meteorological series (e.g., Catchment Attributes and Meteorology for Large-sample Studies (CAMELS))	Hydrological modeling from forcings	Needs robust normalization; data leakage risk
Shi et al. [26]	(Convolutions + memory) (ConvLSTM)	Radar/NWP sequences	Precipitation nowcasting	Computational cost; limited long-range dependencies
Tekin et al. [27]	CNN + LSTM + Attention	Multi-channel radar/NWP sequences	Learned NWP / nowcasting	Complexity; instability on noisy sequences

performance comparable with classical models of similar complexity, while outperforming the standard parameterizations commonly used in climate modeling. In [34], Lachure et al. proposed a lightweight hybrid QCNN that integrates quantum computing with deep learning for temperature predicting. Their hybrid model showed competitive results compared with conventional convolutional neural network (CNN) and long short-term memory (LSTM) architectures. Other studies, including those by Bazgir et al. [35] and Enos et al. [36], explored the potential of quantum computing to enhance traditional machine learning techniques in climate simulations. More recently, Eswaran et al. [37] investigated the use of QNNs for forecasting extreme weather events. Their research focused on designing a quantum neural network capable of analyzing satellite imagery and atmospheric data to detect the precursors of severe storms. The proposed model demonstrated a 30% improvement in forecast accuracy for high-impact weather scenarios, underscoring the potential of quantum-enhanced models in operational meteorology.

Predicting highly dynamic and localized meteorological phenomena remains a significant challenge due to the nonlinear nature of atmospheric processes and their sensitivity to the initial conditions. Classical neural architectures often exhibit limitations in stability and generalization, particularly when faced with noisy or incomplete input data. Within this context, our work investigates the integration of quantum computing into a deep learning pipeline for extreme weather forecasting. Quantum circuits provide a fundamentally different paradigm for information encoding and transformation, exploiting high-dimensional entanglement and non-classical correlations to capture intricate data structures. Building on recent advances in quantum neural networks, we hypothesize that quantum layers can serve as expressive and robust feature spaces, thereby improving generalization capabilities—especially in meteorological applications characterized by spatiotemporal complexity. In this study, we specifically assess the impact of incorporating a variational quantum circuit composed of 10 qubits—simulated via PennyLane’s *lightning.qubit* backend—into a CNN-LSTM framework for short-term bow echo forecasting. Our findings reveal that the resulting hybrid quantum-classical model not only achieves high perceptual accuracy but also exhibits remarkable robustness under both Gaussian noise and fast gradient sign method (FGSM) perturbations. The QCNN lightning 10-qubit model consistently outperforms classical baselines in terms of stability, suggesting that quantum-enhanced representations may offer a promising path toward more resilient extreme weather forecasting systems.

The paper is organized as follows. In Section 2, we introduce the algebraic mathematical concepts that are the basis of quantum computing. Section 3 exhaustively describes the meteorological data used. Section 4 defines the four neural network models and how to evaluate them. Section 5 presents the entire results and their analysis. Finally, Section 6 proposes perspectives such as testing on real hardware and Section 7 concludes with subsequent works that will combine quantum ML and topological data analysis.

2. Mathematical background

2.1. Introduction to Clifford algebras and applications

Quantum neural networks rely on mathematical structures that describe how quantum states evolve and interact. Among these, **Clifford algebras** play a central role because they provide the algebraic framework underlying both quantum mechanics and quantum computation. They offer a compact

formalism to represent spin systems, quantum gates, and the symmetries of unitary evolutions.

In this section, we establish the connection between Clifford algebras, Pauli matrices, and the Special Unitary Group of degree 2 $SU(2)$, and then explain how these mathematical elements are implemented in the quantum layers of our hybrid QCNN-LSTM architecture.

For readers interested in deepening their understanding of the mathematical foundations of quantum computing, particularly the role of Pauli matrices, Clifford algebras, and algebraic structures used in quantum neural networks, we recommend several authoritative references. These works provide rigorous introductions to linear algebra, Lie groups, quantum mechanics, and the central role of Clifford algebras in quantum theory [39–43].

2.2. Quadratic forms and pre-Hilbert space

2.2.1. Bilinear form and signature

Let V be a real (or complex) vector space of dimension n , endowed with a non-degenerate bilinear (or sesquilinear) form $g(\cdot, \cdot)$. Typically, g is symmetric (or Hermitian). A *quadratic form* is given by $Q(v) = g(v, v)$.

Definition 2.1 (Quadratic form). A *quadratic form* on V is a map $Q : V \rightarrow \mathbb{K}$ such that

$$Q(\alpha v) = \alpha^2 Q(v), \quad Q(u + v) + Q(u - v) = 2(Q(u) + Q(v)).$$

The associated bilinear form is

$$g(u, v) = \frac{1}{2}(Q(u + v) - Q(u) - Q(v)).$$

If g is *non-degenerate*, there is a basis of V in which the matrix of g is diagonal, having p positive entries and q negative entries. We speak of a *signature* (p, q) and often use $\mathbb{R}^{p,q}$ to denote such a space of dimension $n = p + q$.

2.3. Definition of Clifford algebra

Definition 2.2 (Clifford algebra). Let (V, g) be a real vector space of dimension n , equipped with a symmetric non-degenerate bilinear form g . The *Clifford algebra* associated with it, denoted by $Cl(V, g)$ or $Cl(V, g)$, is the associative **unital** algebra generated by V with the relation

$$v \cdot v = g(v, v) 1_{Cl(V, g)} \quad \text{for all } v \in V,$$

where $1_{Cl(V, g)}$ is the unit element of the algebra.

Practically, we start from the tensor algebra $T(V) = \bigoplus_{k=0}^{\infty} V^{\otimes k}$ and the quotient by the bilateral ideal generated by

$$\{v \otimes v - g(v, v) 1 : v \in V\}.$$

The resulting quotient is the Clifford algebra.

2.3.1. Dimension and basis

Proposition 2.1. *The dimension of $\text{Cl}(V, g)$ as a real vector space is 2^n , if $\dim(V) = n$.*

Sketch of proof. Choosing an orthonormal basis (e_1, \dots, e_n) of V , one can build a basis of $\text{Cl}(V, g)$ from the monomials

$$e_{i_1} e_{i_2} \cdots e_{i_k}, \quad 1 \leq i_1 < i_2 < \cdots < i_k \leq n.$$

There are 2^n such monomials (including the unit, corresponding to $k = 0$). □

2.4. Standard Clifford algebras $\mathcal{Cl}_{p,q}$

2.4.1. Definition

When $V = \mathbb{R}^{p,q}$ is endowed with a bilinear form of signature (p, q) , we write

$$\mathcal{Cl}_{p,q} := \text{Cl}(\mathbb{R}^{p,q}, g).$$

If we take the Euclidean form (all positive), we get $\mathcal{Cl}_{n,0}$.

2.4.2. Multiplication relations

Let (e_1, \dots, e_{p+q}) be an orthonormal basis of $\mathbb{R}^{p,q}$, i.e., $g(e_i, e_j) = \delta_{ij} \eta_i$ with $\eta_i = +1$ for $i \leq p$ and $\eta_i = -1$ for $i > p$. Then, in $\mathcal{Cl}_{p,q}$, we have

$$e_i^2 = g(e_i, e_i) = \begin{cases} +1 & \text{if } 1 \leq i \leq p, \\ -1 & \text{if } p+1 \leq i \leq p+q, \end{cases}$$

and

$$e_i e_j + e_j e_i = 2 g(e_i, e_j) 1.$$

2.5. Properties and structure

2.5.1. Grading and involutions

Elements of $\text{Cl}(V, g)$ can be classified by their *grade*, extending the usual exterior algebra decomposition into $\Lambda^k V$. However, the Clifford product modifies the standard wedge-product rules.

There are several classical involutions:

- **Reversion**, which sends $v_1 v_2 \cdots v_k \mapsto v_k \cdots v_2 v_1$;
- **Grade involution** (or parity), which associates the factor $(-1)^k$ with a monomial of grade k .

Such involutions are useful for defining, for example, a *norm* for vectors or to study the internal symmetries of the algebra.

2.5.2. Pin and Spin groups

Within the Clifford algebra, one can define $\text{Pin}(p, q)$ as the group of products of unit (or pseudo-unit) vectors for the quadratic form. It is a double cover of $\text{O}(p, q)$. Similarly, $\text{Spin}(p, q)$ is the subgroup formed by the products of an even number of such vectors, forming a double cover of $\text{SO}(p, q)$.

2.6. Concrete examples and related structures

2.6.1. $\mathcal{Cl}_{0,1}$: Complex numbers

Here, V is one-dimensional with $g(e_1, e_1) = -1$. Hence e_1 satisfies $e_1^2 = -1$. We obtain an isomorphism

$$\mathcal{Cl}_{0,1} \simeq \mathbb{C}$$

by identifying e_1 with the imaginary unit i .

2.6.2. $\mathcal{Cl}_{2,0}$: Quaternions

Consider V of dimension 2 with a Euclidean bilinear form, giving $e_1^2 = e_2^2 = +1$ and $e_1e_2 = -e_2e_1$. Then e_1e_2 behaves as an $((e_1e_2)^2 = -1)$. One can show

$$\mathcal{Cl}_{2,0} \simeq \mathbb{H},$$

which is the algebra of **quaternions**. An example of identification is

$$i \leftrightarrow e_1e_2, \quad j \leftrightarrow e_2, \quad k \leftrightarrow e_1,$$

though conventions may vary.

2.7. Pauli matrices, and the $SU(2)$ group

2.7.1. Pauli matrices

In physics (especially quantum mechanics), the **Pauli matrices** are three 2×2 complex matrices commonly denoted σ_x, σ_y and σ_z :

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

They satisfy the relations

$$\sigma_x\sigma_y = i\sigma_z, \quad \sigma_y\sigma_z = i\sigma_x, \quad \sigma_z\sigma_x = i\sigma_y,$$

and

$$\sigma_x^2 = \sigma_y^2 = \sigma_z^2 = I_2,$$

where I_2 is the 2×2 identity matrix. These mirror the Clifford algebra relations in three dimensions (with Euclidean or mixed signatures). In fact, the Pauli matrices provide a faithful representation of such a Clifford algebra.

2.7.2. Connection with $SU(2)$ and spinorial representations

The group $SU(2)$ is that of 2×2 unitary matrices with a determinant of 1. It is **isomorphic** (as a manifold) to the 3-sphere S^3 and is a double cover of $SO(3)$. Algebraically, it is closely related to the *unit quaternions* \mathbb{H}_1 .

In particular, the usual two-dimensional (2D) complex representation of $SU(2)$ describes the *spin-1/2* behavior of fermions in quantum physics (e.g., electrons).

Hence, the entire family of quantum gates manipulating a qubit (R_X, R_Y, R_Z , Hadamard, etc.) arises directly from the Clifford algebraic structure. This correspondence ensures that the abstract algebraic framework introduced above is not merely formal—it dictates the geometry of qubit transformations on the **Bloch sphere**, the 2D complex Hilbert space representing quantum states.

2.8. From Clifford algebra to quantum circuits

2.8.1. Quantum computing

In quantum computing, the manipulation of qubits (two-level systems) is closely tied to the structure of $SU(2)$ and spinorial representations, specifically

- *Quantum gates* like X, Y, Z or the *Hadamard gate* H are 2×2 complex unitary matrices, often derived from or related to the Pauli matrices.
- *Qubit states* are *spinors* in 2 dimensions (over \mathbb{C}), and the Bloch sphere provides a geometric picture of these states.
- Clifford algebras offer a powerful language for handling larger qubit networks, including higher-dimensional versions (qubits), quantum information encoding, spin correlations, and more.

2.8.2. Qubits and Bloch sphere

A **qubit** is a vector in \mathbb{C}^2 , typically represented as

$$|\psi\rangle = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1.$$

The state space (up to a global phase factor) is identified with the *Bloch sphere*, which is a sphere S^2 . Concretely, any pure qubit state can be written as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\varphi} \sin\left(\frac{\theta}{2}\right)|1\rangle,$$

which corresponds, geometrically, to a point with coordinates (θ, φ) on the sphere.

The link to $SU(2)$ is as follows: Every unitary (detailed) operator acting on a qubit with a determinant of 1 corresponds to a *rotation* on the Bloch sphere. More precisely, there is a surjective group morphism

$$SU(2) \rightarrow SO(3),$$

showing that spin-1/2 (qubit) transformations cover the usual three-dimensional (3D) rotations twice.

This deep duality is directly related to Clifford algebras, which provide a unified framework for describing rotations, spinors, and operations on qubits (quantum logic, quantum gates, etc.).

The following subsections summarize how Clifford algebras and Pauli matrices apply to defining qubits, quantum layers, and hybrid quantum-classical architectures in QNNs.

2.8.3. Quantum layers in QNNs

In a quantum layer, qubits are often initialized in a reference state (commonly $|0\rangle^{\otimes n}$), then undergo a sequence of parametric unitary operators:

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle^{\otimes n}.$$




Figure 1. Bloch sphere.

Quantum gates are typically constructed from Pauli rotations of the form

$$R_x(\theta) = e^{-\frac{i}{2}\theta\sigma_x}, \quad R_y(\theta), \quad R_z(\theta).$$

Hence the entire parametric circuit is heavily based on Clifford-algebraic notions, with a topology adapted to the data (classification, generation, etc.).

2.8.4. Integration into the QCNN-LSTM architecture

The hybrid quantum convolutional neural network with LSTM (QCNN-LSTM) developed in this work uses these algebraic principles to construct a trainable quantum layer embedded within a classical deep learning pipeline.

- 1) **Classical feature extraction (CNN + LSTM):** Meteorological inputs (satellite, convective available potential energy (CAPE), lightning) are processed through convolutional and temporal layers to extract latent spatio-temporal features.
- 2) **Quantum encoding and transformation:** The latent vector is normalized and encoded into a 10-qubit quantum state via angle embedding, each feature controlling a Pauli rotation $R_y(x_i)$. The qubits are then entangled through a chain of (Controlled-NOT) CNOT gates—realizations of Clifford operations that generate a correlated quantum representation of the input data.
- 3) **Measurement and fusion:** Expectation values $\langle Z_i \rangle$ of the Pauli-Z operator are measured and merged with the LSTM temporal embedding to produce the final forecast image of the bow echo cloud at time $t + 1$.

This hybrid approach leverages the representational strength of CNNs along with the high-dimensional Hilbert space explorations of quantum circuits, aiming to provide improvements in expressivity and generalization.

2.8.5. Classical data encoding

In quantum neural networks, classical data $x \in \mathbb{R}^n$ must be mapped into quantum states. A common technique is *angle embedding* as follow:

$$R_y(x) = e^{-\frac{i}{2} x \sigma_y} = \begin{pmatrix} \cos(\frac{x}{2}) & -\sin(\frac{x}{2}) \\ \sin(\frac{x}{2}) & \cos(\frac{x}{2}) \end{pmatrix}.$$

For n features, one may construct

$$U(\mathbf{x}) = \bigotimes_{i=1}^n R_y(x_i),$$

thereby embedding the classical vector \mathbf{x} into a multi-qubit state. This sets the stage for further entangling operations and parametric gates.

2.8.6. Entanglement and measurement

Entanglement. Entanglement is a key resource in QNNs, enabling non-classical correlations among qubits. In practice, one often applies single-qubit rotations followed by multi-qubit gates (e.g. CNOT) that *entangle* the qubits. A typical pattern for n qubits might be:

$$U_{\text{entangle}}(\boldsymbol{\theta}) = \prod_{j=1}^{n-1} (\text{CNOT}_{j,j+1} R_y(\theta_j)),$$

where $R_Y(\theta_j)$ is a single-qubit rotation about the y -axis on qubit j , and $\text{CNOT}_{j,j+1}$ is a two-qubit controlled-NOT gate acting on qubit j (control) and qubit $j + 1$ (target). The product over j indicates that these gates are applied across the qubit chain (for $j = 1, 2, \dots, n - 1$), entangling each qubit with its neighbor. The CNOT gate flips the target qubit only when the control qubit is in state $|1\rangle$, and acts as the identity otherwise. In the two-qubit computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ (the control qubit is listed first), the CNOT operator is represented by the matrix

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix},$$

which clearly swaps the $|10\rangle$ and $|11\rangle$ states (flipping the target qubit when the control is $|1\rangle$). Through these entangling operations, the qubits become correlated in a way that cannot be described by any product state, thereby increasing the expressiveness of the QNN.

$$U_{\text{entangle}}(\boldsymbol{\theta}) = \prod_j \text{CNOT}_{j,j+1} \cdot R_y(\theta_j) \quad (2.1)$$

Measurement. After the state preparation and entangling operations, the final step in a QNN circuit is to measure the qubits to obtain a classical output. We typically measure each qubit on the computational (Z) basis, which corresponds to the eigenstates $|0\rangle$ and $|1\rangle$ of the Pauli- Z operator. The




Figure 2. Simple 4-qubit quantum circuit containing an R_Y gate on each qubit, then a chain of CNOT gates connecting adjacent qubits.

measurement observable for qubit i can be represented by the Pauli-Z operator Z_i , and the expected measurement outcome is given by its quantum expectation value

$$\langle Z_i \rangle = \langle \psi | Z_i | \psi \rangle$$

where $|\psi\rangle$ is the state of the system before measurement. This expectation value $\langle Z_i \rangle$ will be +1 if qubit i is in the state $|0\rangle$ (since $Z|0\rangle = +1|0\rangle$), and -1 if it is in the state $|1\rangle$ ($Z|1\rangle = -1|1\rangle$). For a qubit in a superposition state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, the expectation evaluates to $\langle Z_i \rangle = |\alpha|^2 - |\beta|^2$, which is the difference between the probabilities of obtaining 0 vs. 1 upon measurement. In other words, $\langle Z_i \rangle$ ranges between -1 and $+1$, indicating the bias towards $|0\rangle$ (positive values) or $|1\rangle$ (negative values) for that qubit. Collecting the results for all n qubits, we obtain a vector of the expectation values:

$$\mathbf{z} = (\langle Z_1 \rangle, \langle Z_2 \rangle, \dots, \langle Z_n \rangle)$$

which represents the measured output of the QNN. Each component $\langle Z_i \rangle$ in this vector is a classical number summarizing the measurement outcome of qubit i . Together, they can be used for further classical processing or evaluation of the QNN's performance.

Through this integration, the Clifford algebraic structure is not an abstract prelude but the mathematical justification for the design of the quantum layer. It governs the transformations applied to encoded meteorological features and ensures that the quantum circuit explores a non-commutative, high-dimensional space capable of capturing complex atmospheric correlations. The resulting hybrid model inherits the stability and robustness properties of Clifford-based unitary transformations, which contribute to the observed resilience of the QCNN-LSTM against noise and adversarial perturbations.

This direct lineage, from algebraic theory to operational architecture, establishes a rigorous bridge between mathematical physics and data-driven prediction, forming the theoretical foundation for the quantum models analyzed in the following sections.

3. Meteorological data

3.1. Data description

The dataset used in this study is based on a combination of heterogeneous sources from different European meteorological services, covering the main dynamic factors related to the formation and


evolution of bow echo phenomena. These data, primarily in the form of images, constitute the inputs and outputs of the hybrid neural networks models.

The satellite images come from the EUMETSAT service [44] and provide a spatio-temporal visualization of the cloud cover associated with convective systems. These data are provided at a 15-minute cadence, allowing us to monitor the dynamics of storm cells with a relatively fine temporal resolution. They represent the cloud structure associated with the bow echo phenomenon.


The CAPE data are sourced from Météo France [45] and represent the atmospheric instability available to fuel convection. They are initially provided at an hourly resolution, in the form of scalar fields that we converted into spatialized maps for integration into the model's convolutional pipeline.

The lightning data comes from the Swiss meteorological service Meteologix [46]. These radar images synthesize the intensity of electrical activity on the ground, made available every 15 minutes. These data complement the multi-source approach by providing an indirect measure of the severity of convective phenomena.


Finally, the bow echo radar images, used as model targets, come from Météo France's high-resolution radar system. These images are available at a frequency of 5 minutes, thus providing a temporal granularity that is particularly useful for analyzing and predicting the fine structure of the arcus.




EUMETSAT satellite image



CAPE map (Météo France)



Lightning radar (Meteologix)



Bow echo radar (Météo France)

Figure 3. Representative examples of the four datasets used in our hybrid models: Satellite imagery (left top), convective energy fields (right top), lightning radar (left bottom), and bow echo radar (right bottom).

To ensure temporal consistency across all data sources, we applied linear interpolation to each input dataset. This harmonization step enabled us to reproject all images onto a regular 5-minute

interval sequence, aligned with the highest temporal resolution of the target dataset (Météo France radar images). As a result, all input data are uniformly synchronized and temporally sequenced, which facilitates temporal pattern learning by the LSTM and improves the consistency of predicted storm cell trajectories.

EUMETSAT. Among the EUMETSAT satellite images used in this study, some clearly reveal the presence of atmospheric gravity waves (see Figure 3), forming near the developing bow echo system. These waves appear as regularly spaced, parallel cloud bands, indicating vertical oscillatory motion within the troposphere. Gravity waves arise when an air parcel is displaced from its equilibrium position in a stably stratified atmosphere. As the atmosphere attempts to restore equilibrium, oscillations in vertical velocity, pressure, and temperature propagate outward from the disturbance. In convective environments, such waves can be triggered by downdrafts, gust fronts, or vertical wind shear associated with intense thunderstorm activity. Their visual signature in EUMETSAT imagery offers valuable insight into local atmospheric instability and the interactions between convective updrafts and stable stratification layers aloft. These wave patterns act as indirect indicators of the developing storm's severity and dynamics. Furthermore, they constitute a rich source of spatial information that CNNs are particularly well-suited to exploit. Their presence reinforces the relevance of EUMETSAT satellite imagery as a core input to our model, complementing dynamic predictors such as CAPE and lightning activity. This justifies the deliberate inclusion of EUMETSAT images in our dataset, not only for their usefulness in tracking convective systems, but also for the physical processes they capture, which are both meteorologically significant and computationally exploitable by deep learning architectures.

CAPE. One of the key input variables in our model is the CAPE, a fundamental thermodynamic parameter used to assess the potential intensity of convective phenomena. CAPE quantifies the amount of buoyant energy available to an air parcel for vertical ascent and thus serves as a direct indicator of atmospheric instability. Mathematically, CAPE is defined as

$$\text{CAPE} = \int_{z_{\text{LFC}}}^{z_{\text{EL}}} g \left(\frac{T_{v,p} - T_{v,e}}{T_{v,e}} \right) dz$$

where: z_{LFC} is the level of free convection, z_{EL} is the equilibrium level, g is the gravitational acceleration, $T_{v,p}$ is the virtual temperature of the rising air parcel and $T_{v,e}$ is the virtual temperature of the environment.

This integral quantifies the positive buoyancy experienced by a parcel between the level of free convection and the equilibrium level regions where the parcel is warmer (and thus less dense) than its surroundings. CAPE is expressed in joules per kilogram (J/kg), with values exceeding 2000 J/kg typically indicating strong potential for deep convection, and values above **3000 J/kg** as observed in our dataset suggesting conditions that are favorable for *extreme convective events* such as bow echoes or supercells. In this study, CAPE values were interpolated in time and transformed into gridded spatial maps. These maps are fed into the CNN as image-like inputs, enabling the model to capture both the magnitude and spatial gradients of instability fields. This representation allows the CNN to exploit fine-scale spatial patterns and detect precursors to convective development, making CAPE a physically meaningful and computationally effective input in our hybrid prediction architecture.

Radar-based precipitation and lightning imagery. In addition to satellite imagery and CAPE maps, our model incorporates radar imagery capturing both lightning activity and precipitation

structures. The lightning data, sourced from Meteorlogix [46], provides high-resolution information on the electrical activity associated with convective systems. These images serve as dynamic indicators of a storm's intensity and are updated every 15 minutes. The target variables consist of radar images of bow echo systems, provided by Météo France at a 5-minute temporal resolution. These reflectivity maps capture the fine-scale morphology of arc-shaped convective systems, particularly those approaching the Corsican coastline—a region where sudden and intense storm surges pose significant meteorological risks. By integrating these radar datasets, the model is equipped to detect both the electrical and hydrometeorological signatures of severe convective events, thereby improving its ability to anticipate the onset of hazardous weather conditions in real time.

3.2. Dataset structuring

The complete dataset corresponds to a mesoscale convective system observed between 0500 h and 0800 h local time. After preprocessing and linear interpolation, all data sources were temporally resampled to an interval of 5 minutes, resulting in 37 synchronized images per variable (inputs and targets). All images were resized to a spatial resolution of 256×256 pixels. This sequence captures the progressive intensification of the system, with a particular focus on a key transformation: Around 0700 h local time, the initial squall line abruptly reorganized into a bow echo, which is clearly visible in the radar reflectivity imagery. This structural shift, associated with enhanced convection and wind severity, occurred just as the system approached the Corsican coastline. The proximity to land plays a pivotal role in this reorganization. By 0700 h, the trajectory and dynamics of the storm system made its coastal impact virtually certain. As a result, we evaluate the performance of our four neural architectures specifically over the final hour of the sequence (0700 h to 0800 h local time), which represents a high-stakes forecasting window where early predictions are crucial for alert systems and risk mitigation.

4. Methodology

We designed and implemented four architectures of increasing complexity and expressivity, all trained on the meteorological data described in the paragraph above. In the first part of our study, we consider three distinct model configurations: A CNN-LSTM model, combining spatial feature extraction with temporal modeling via LSTM layers, and a CNN-LSTM-Transformer, where the LSTM component is replaced by self-attention mechanisms to capture long-range temporal dependencies. A first version of the QCNN-LSTM, where a variational quantum circuit is embedded as a learnable layer in the architecture using PennyLane KerasLayer and the default.qubit simulator with 8 qubits. These three models share a common input–output structure and are trained under the same meteorological time windows. The goal is to assess their respective strengths in terms of spatial coherence, stability, and noise resilience.

In the second part of the study, we introduce a more advanced quantum model: A QCNN-LSTM using a custom 10-qubit quantum layer with the high-performance lightning.qubit backend. Unlike the earlier KerasLayer-based model, this architecture is implemented through a fully customized quantum layer (PennyLaneQuantumLayer), enabling greater control over the quantum operations, training dynamics, and gradient flow.

This version not only increases the quantum capacity by moving from 8 to 10 qubits, but also

serves to evaluate whether deeper quantum encoding enhances robustness. In addition, it allows us to apply adversarial robustness testing (FGSM) on the quantum feature extractor, assessing the network's resilience to input perturbations.

4.1. The CNN-LSTM architecture

The first architecture combines CNNs with LSTM layers to capture both spatial and temporal dependencies in the data; see Figure 4. The CNN branch processes sequences of meteorological images (e.g., satellite, CAPE, lightning) through a shared convolutional stack, using five convolutional blocks with increasing depth. Each image in the input sequence is processed individually via time-distributed layers, producing a compact spatial feature vector. This vector is then flattened and concatenated with an LSTM output that encodes the temporal dynamics from auxiliary inputs (e.g., time series, positional encodings). The final output is a predicted red–green–blue (RGB) image of the bow echo cloud at time $t+1$.

CNN-LSTM-transformer architecture. To improve the modeling of long-range temporal interactions, the second architecture replaces the LSTM with a Transformer encoder. After CNN-based feature extraction, temporal sequences are fed into a multi-head self-attention module. The Transformer component aims to enhance the model's ability to selectively focus on the most relevant temporal dependencies across the input sequence.

QCNN-LSTM with eight qubits (KerasLayer + default.qubit). The third model introduces quantum computation by replacing the penultimate dense layer with a quantum variational circuit, implemented via the KerasLayer wrapper from PennyLane. The quantum layer consists of 8 qubits, initialized with the single-qubit rotation R_X gates and entangled using BasicEntanglerLayers. Inputs to the quantum layer are projected onto a compact feature space via a classical dense layer of size eight with tanh activation. The output of the quantum circuit a vector of Pauli-Z expectation values is then fused with the LSTM stream to generate the predicted image. To build our hybrid CNN + quantum LSTM architecture, we employed PennyLane, an open-source software framework designed for differentiable quantum programming. PennyLane provides seamless integration with TensorFlow, allowing us to embed a trainable quantum circuit directly within a classical deep learning pipeline. This quantum circuit is wrapped into a KerasLayer, enabling joint optimization of both classical and quantum parameters via gradient descent. The differentiable nature of the quantum node ensures compatibility with backpropagation, making PennyLane an ideal tool for quantum-enhanced learning on structured meteorological data.

QCNN-LSTM with 10 qubits (Custom layer + lightning.qubit). The final model, and the most advanced, integrates a custom quantum layer based on the lightning.qubit backend; see Figure 4. The variational circuit uses 10 qubits and a depth-2 entanglement pattern. Unlike the KerasLayer version, this model uses a fully customized `tf.keras.layers.Layer` implementation, enabling direct control over quantum computation and seamless integration with the TensorFlow backpropagation engine. A targeted FGSM adversarial attack was conducted to evaluate its resilience.

Role of dense layers in the four models. In the four architectures studied, dense layers were introduced after feature extraction. This choice was motivated by the small size of the dataset, which makes training models that are too deep or too parametric difficult without overfitting. Dense layers allow us to reduce the dimensionality of convolutional representations (flatten); project the extracted features into a regular latent space, acting as a generalization kernel between heterogeneous

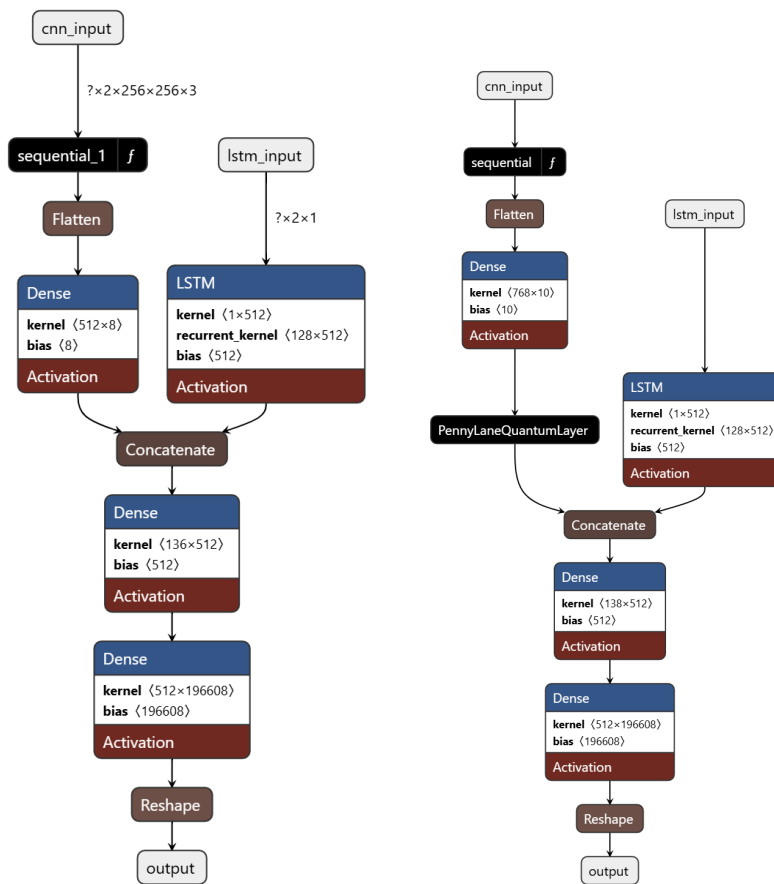


Figure 4. CNN LSTM (left) and quantum CNN lightning 10 qubits (right) architectures.

representations (weather images + time); and promote learning convergence, thanks to their fully connected structure, by efficiently disseminating information between channels. In the case of CNN-LSTM, the dense layer between the CNN and the fusion acts as a complexity reducer. In the Transformer, the dense layers in the attention block allow resampling of the information after the attention phase. Finally, in the quantum model (QCNN), the dense layer is essential to adapt the size of the latent space to that of the quantum circuit (number of qubits).

4.2. Training and prediction of bow echo images

Each of the three input datasets (satellite imagery, lightning density maps, CAPE fields, and bow echo output) consists of 37 temporally ordered RGB images of size 256×256 . The images are sampled hourly from 0500 to 0800 h, capturing the development cycle of a mesoscale convective system.

For training, we use data from 0500 to 0655 h, corresponding to the pre-convective and initiation phases. The models learn to predict the next image in the bow echo sequence \hat{Y}_{t+1} from the corresponding meteorological inputs at time t .

The testing phase focuses on the critical forecast window between 0700 and 0800 h, during which the bow echo cloud system forms and intensifies. This allows us to assess the model's ability to generalize to unseen and rapidly evolving patterns.

To ensure temporal coherence, the data are chronologically split without shuffling. This choice reflects a realistic operational scenario where the models must extrapolate the storm's dynamics based on prior observations.

In this work, we aim to forecast the evolution of severe convective structures specifically, bow echoes at future time step $t + 1$, based on a sequence of meteorological observations at time t . These include satellite imagery of convection, radar-based lightning data, and CAPE maps. Let us denote the input triplet at time t as:

$$\mathbf{X}_t = \{\mathbf{I}_t^{\text{sat}}, \mathbf{I}_t^{\text{cape}}, \mathbf{I}_t^{\text{lightning}}\}, \quad \mathbf{X}_t \in \mathbb{R}^{3 \times 256 \times 256 \times 3}$$

and the target image (bow echo cloud) at time $t + 1$ as:

$$\mathbf{Y}_{t+1} = \mathbf{I}_{t+1}^{\text{bow echo}} \in \mathbb{R}^{256 \times 256 \times 3}.$$

The model processes \mathbf{X}_t through a time-distributed CNN to extract the spatial features from each modality independently. These features are then passed through a fully connected projection layer:

$$\mathbf{z}_{\text{cnn}} = \tanh(W_{\text{dense}} \cdot \text{Flatten}(\text{CNN}(\mathbf{X}_t))),$$

where $W_{\text{dense}} \in \mathbb{R}^{n_q \times d}$ projects the features into a space that is compatible with the quantum circuit of $n_q = 8$ qubits.

The projected vector \mathbf{z}_{cnn} is encoded into a parameterized quantum circuit $U(\mathbf{z}_{\text{cnn}}, \boldsymbol{\theta})$, composed of R_X rotations followed by entanglement gates (basic entangler layers). The circuit outputs the expectation value of the Pauli-Z observable on each qubit:

$$\mathbf{z}_{\text{quantum}} = [\langle Z_0 \rangle, \langle Z_1 \rangle, \dots, \langle Z_7 \rangle].$$

In parallel, a simple LSTM network processes a normalized abstract time sequence $\mathbf{T}_t \in \mathbb{R}^{3 \times 1}$, capturing the latent temporal dependencies. Its output $\mathbf{z}_{\text{lstm}} \in \mathbb{R}^{128}$ is concatenated with the quantum vector:

$$\mathbf{z}_{\text{fused}} = [\mathbf{z}_{\text{quantum}}, \mathbf{z}_{\text{lstm}}]$$

This fused latent representation is passed through dense decoding layers to produce a predicted RGB image of the arcus cloud:

$$\hat{\mathbf{Y}}_{t+1} = \text{Reshape}(\sigma(W_{\text{out}} \cdot \text{ReLU}(W_{\text{fused}} \cdot \mathbf{z}_{\text{fused}})))$$

The model is trained using the MSE loss:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N \|\hat{\mathbf{Y}}_{t+1}^{(i)} - \mathbf{Y}_{t+1}^{(i)}\|^2$$

Training was conducted with early stopping (patience = 10), using a batch size of 1 due to the memory overhead of the quantum simulation. During testing, the last two available samples are used as input to generate $\hat{\mathbf{Y}}_{t+1}$, the predicted image of the bow echo. Quantitative evaluation metrics include the structural similarity index (SSIM), Wasserstein distance, and root mean squared error (RMSE).

4.3. Metrics and evaluation protocol

To objectively assess and compare the performance of the four neural network architectures, we employ three complementary evaluation metrics.

4.3.1. Evaluation metrics

(1) RMSE. RMSE evaluates the pixel-wise accuracy between the predicted and true bow echo image by comparing their flattened pixel vectors:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{Y}_i - Y_i)^2},$$

where \hat{Y}_i and Y_i denote the predicted and actual pixel intensities respectively. RMSE is sensitive to large deviations and is computed after rescaling the image range to $[0, 255]$.

(2) SSIM. SSIM quantifies the perceptual similarity between two images by evaluating their luminance, contrast, and structural alignment:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)},$$

where μ_x, μ_y are local means, σ_x, σ_y are standard deviations, and σ_{xy} is the covariance between patches x and y . SSIM values are computed using the full RGB image with a dynamic range of 255.

(3) Wasserstein distance. To measure the spatial coherence between predicted and real images, we compute the Wasserstein distance over the flattened grayscale intensities as follows:

$$\mathcal{W}_1(\hat{Y}, Y) = \inf_{\gamma \in \Pi(\hat{Y}, Y)} \int_{\mathbb{R}^2 \times \mathbb{R}^2} \|x - y\| d\gamma(x, y).$$

This metric measures the distance between the predicted and real distributions of pixel values, interpreted as the minimal cost to transform one distribution into the other. It provides complementary insights to RMSE by focusing on the global statistical distribution rather than local error.

For each architecture, we report the RMSE, SSIM, and Wasserstein scores with both clean and noisy data. This dual evaluation allows us to quantify the robustness of the models to observational perturbations, a critical aspect in real-world prediction of severe weather.

4.3.2. Evaluation scenarios.

Each model is evaluated on the following:

- Clean data: Standard input data without perturbation.
- Noisy data, corrupted with Gaussian noise to simulate degraded observational conditions. Gaussian noise is defined as:

$$\tilde{\mathbf{X}} = \text{clip}(\mathbf{X} + \mathcal{N}(\mu, \sigma^2), 0, 1),$$

where $\mu = 0$, $\sigma = 0.1$, and the clipping ensures that the noisy images remain within the $[0, 1]$ interval. This noise simulates realistic disturbances that could affect image sensors or data

transmission in operational meteorological systems. We applied Gaussian noise to satellite imagery.

- Adversarial data for the quantum CNN lightning 10q model. To rigorously evaluate the adversarial robustness of the hybrid QCNN model, we implemented the FGSM on the classical convolutional input pathway, following the principle introduced by Goodfellow et al. [38] The FGSM attack perturbs the input x in the direction of the loss gradient to produce an adversarial sample x^{adv} :

$$x^{\text{adv}} = x + \epsilon \cdot \text{sign}(\nabla_x \mathcal{L}(x, y)).$$

In our implementation, we extracted the subnetwork leading up to the quantum interface layer (*reduce_to_quantum*), isolating the classical convolutional component of the QCNN. A dummy loss was computed between this intermediate output and a null target to induce a gradient direction without propagating through the quantum layer, which remained frozen for stability and computational efficiency.

For all models and conditions, we computed the time series of RMSE, SSIM, and Wasserstein distance over each test timestep; the distributional statistics (mean, standard deviation) across the series; and boxplots and comparative line plots to visualize the models' stability, variability, and perceptual accuracy.

4.3.3. Model groups

To reflect architectural differences, we split the analysis into the following two groups:

Classical and hybrid models: CNN-LSTM, CNN-LSTM-Transformer, and QCNN-LSTM (8 qubits, KerasLayer);

Quantum-enhanced model: QCNN-LSTM (10 qubits, lightning.qubit backend).

This separation allows for a focused discussion on quantum enhancement effects under standard and adversarial conditions.

5. Results

5.1. Classical and hybrid models

5.1.1. Comparative performance analysis on clean data

The performance of the three architectures CNN-LSTM, CNN-LSTM-Transformer, and Quantum CNN-LSTM was evaluated using three complementary metrics: RMSE, SSIM, and Wasserstein distance. The temporal evolution of these metrics throughout the prediction window (0700 to 0800 h) is shown in Figure 5.

RMSE. The CNN-LSTM model achieves the lowest and most stable RMSE values throughout the sequence, with optimal predictions around 0725 and 0750 h. The QCNN-LSTM model reports slightly higher RMSEs but demonstrates consistent performance ($\approx 27\text{--}33$), suggesting robust behavior. In contrast, the CNN-LSTM-Transformer exhibits significant variability, including a major spike at 0735 h (RMSE ≈ 62), indicating higher sensitivity to spatial patterns.

SSIM. From a perceptual perspective, CNN-LSTM clearly outperforms the other models, reaching SSIM values up to 0.98, indicating excellent structural preservation of arcus formations. QCNN-LSTM



Figure 5. RMSE, SSIM, and Wasserstein distance evolution over time for the three models on clean input data.

maintains intermediate values (around 0.6) with low variance. The Transformer-based model shows frequent perceptual drops, with SSIM plunging to 0.3 between 0730 and 0740 h.

Wasserstein distance. This metric further highlights the stability of the quantum model: QCNN-LSTM maintains the lowest and most consistent Wasserstein distances (4.7), suggesting accurate spatial alignment between prediction and reality. In contrast, the CNN-LSTM and Transformer models show frequent peaks exceeding 15 or 25, reflecting instability in capturing the spatial mass transport of convective systems.

Overall, CNN-LSTM excels in perceptual and numerical accuracy but is less stable in terms of

spatial transport. QCNN-LSTM stands out for its consistent robustness across all metrics, particularly in Wasserstein distance a potentially critical property for rare and extreme event forecasting. Meanwhile, CNN-LSTM-Transformer, although theoretically more expressive, demonstrates reduced operational reliability under clean input conditions.

5.1.2. Visual analysis of predicted arcus images

Figure 6 presents a visual comparison between the ground truth arcus cloud images and the predictions generated by CNN-LSTM, CNN-LSTM-Transformer, and *QCNN-LSTM*. The selected timestamps (0705, 0725, 0735, and 0755 h) capture different phases in the evolution of the convective system.

0705 h. At this early stage, all three models produce reasonably accurate predictions. CNN-LSTM preserves the bow echo structure with sharp contours and consistent color rendering. CNN-LSTM-Transformer introduces minor distortions, particularly around the peripheral regions. QCNN-LSTM generates smoother outputs with less detail, yet retains the global spatial coherence.

0725 h. This timestamp corresponds to one of the best-performing cases for all models. CNN-LSTM delivers high fidelity in both the structure and texture. CNN-LSTM-Transformer performs acceptably, though geometric inconsistencies remain. QCNN-LSTM, while slightly less sharp, captures the arcus shape and positioning robustly.

0735 h. A challenging and discriminating scenario. CNN-LSTM-Transformer output degrades notably, showing displaced and fragmented storm structures. CNN-LSTM retains the arcus morphology but with reduced clarity. QCNN-LSTM, though more diffuse, maintains spatial alignment, aligning with its low Wasserstein distance at this step.

0755 h. In this final case, CNN-LSTM produces a highly accurate prediction, preserving both the structure and texture of the arcus cloud. QCNN-LSTM retains the main shape with slight blurring. CNN-LSTM-Transformer continues to exhibit instability in localizing the convective system.

In summary, CNN-LSTM provides the sharpest and most perceptually accurate forecasts, in line with its superior RMSE and SSIM scores. QCNN-LSTM stands out for its robustness and spatial consistency, delivering reliable predictions even in complex scenes. In contrast, CNN-LSTM-Transformer, though theoretically powerful, demonstrates sensitivity and temporal instability in this operational context. This visual analysis reinforces the quantitative results and highlights the specific strengths and limitations of each architecture.

5.1.3. Robustness analysis: Performance with and without Gaussian noise

This section evaluates the robustness of the three models CNN-LSTM, CNN-LSTM-Transformer, and QCNN-LSTM under Gaussian noise. The input satellite images were perturbed with zero-mean Gaussian noise ($\mu = 0$, $\sigma = 0.1$) and evaluated using the metrics RMSE, SSIM, and Wasserstein distance. The goal is to assess the model resilience to degraded input quality, as illustrated by the six figures below.

RMSE. CNN-LSTM model is the most sensitive to noise, with a noticeable increase in variance and outliers. QCNN-LSTM remains very stable, showing almost no variation in the error distribution between clean and noisy conditions. Interestingly, CNN-LSTM Transformer shows a slight improvement in the median RMSE with noise, suggesting a possible regularization effect or

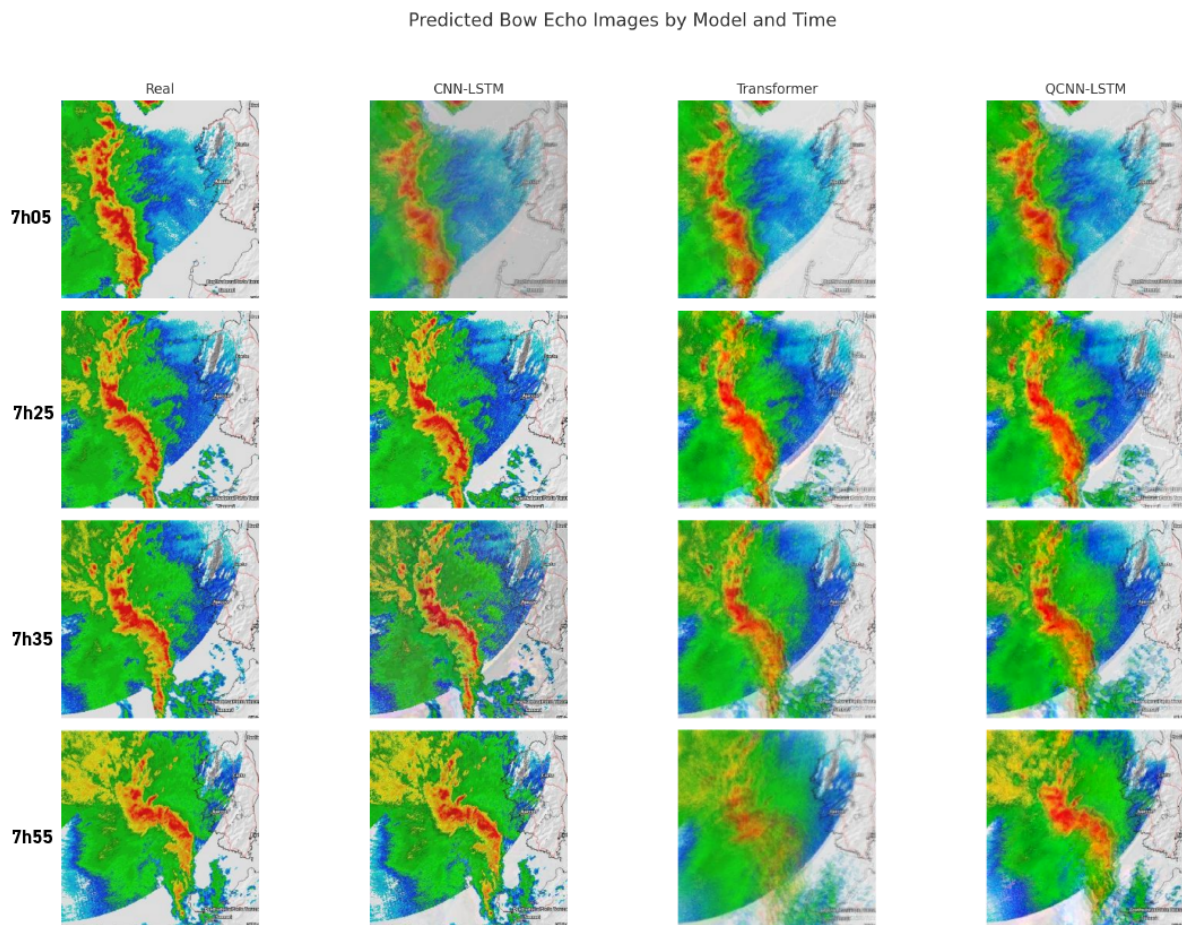


Figure 6. Comparison of predicted bow echo (arcus) images at different times (0705, 0725, 0735, and 0755 h). The top row shows the real images, while the subsequent columns display the predictions generated by the three models: CNN-LSTM, CNN-LSTM-Transformer, and QCNN-LSTM. Visually, CNN-LSTM's predictions appear more accurate in terms of structural detail, while QCNN-LSTM predictions are more stable in texture and spatial consistency. Transformer-based predictions suffer from occasional structural distortions and displacements. These visual trends are consistent with the quantitative metrics reported in the previous section.

robustness to perturbation.

SSIM. CNN-LSTM retains the highest SSIM scores and even benefits slightly from noise, suggesting improved generalization. QCNN-LSTM displays moderate but steady values. CNN-LSTM Transformer model suffers from overall lower SSIM, though its predictions stabilize slightly under noisy conditions.

Wasserstein distance. QCNN-LSTM maintains low and stable Wasserstein distances regardless of noise, confirming its spatial consistency. CNN-LSTM shows fluctuations but no systematic degradation. CNN-LSTM Transformer again appears to be inconsistent, with large variations and

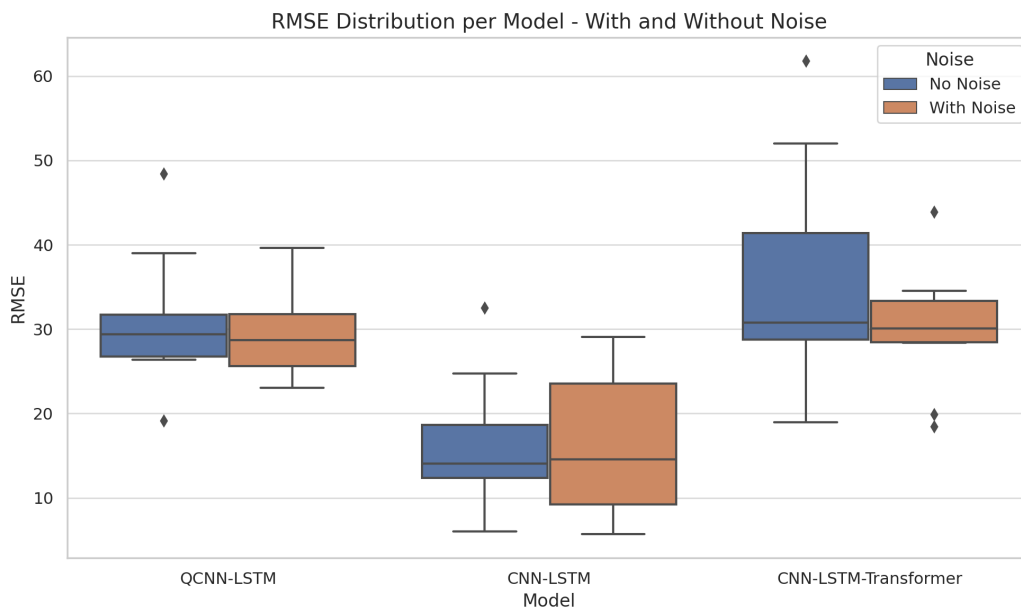


Figure 7. Distribution of RMSE for the models with and without Gaussian noise.

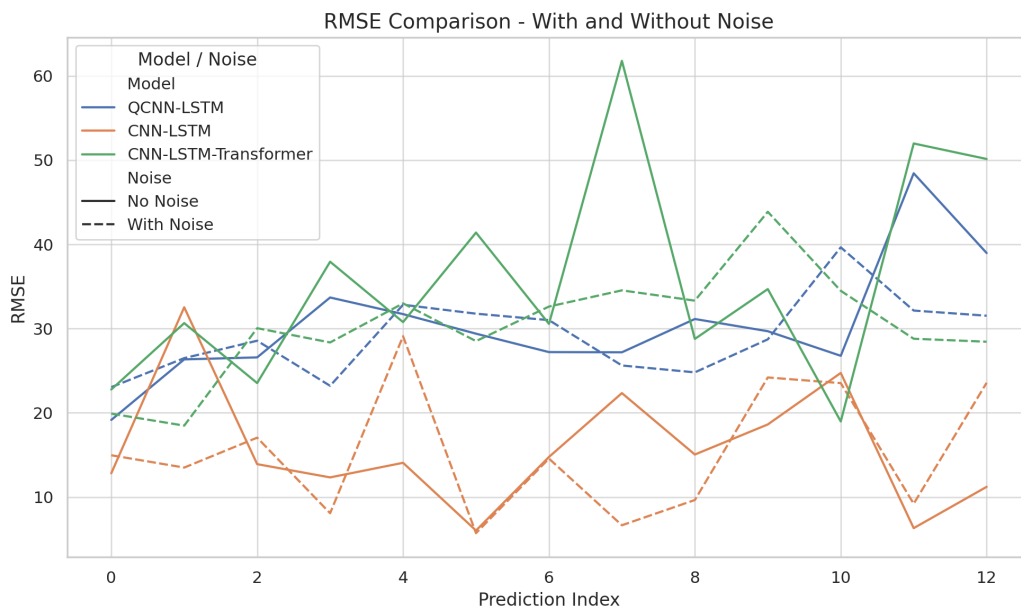


Figure 8. Evolution of RMSE across predictions: Noisy vs. clean images.

multiple peaks under both conditions.

Table 2 summarizes the observed behaviors across metrics.

These results highlight QCNN-LSTM as the most robust model against Gaussian perturbations, offering consistent predictions across all metrics. While CNN-LSTM remains excellent under clean inputs, its sensitivity to noise may limit operational reliability. CNN-LSTM Transformer shows unstable performance, suggesting that it requires further calibration or architecture tuning for this type of application.

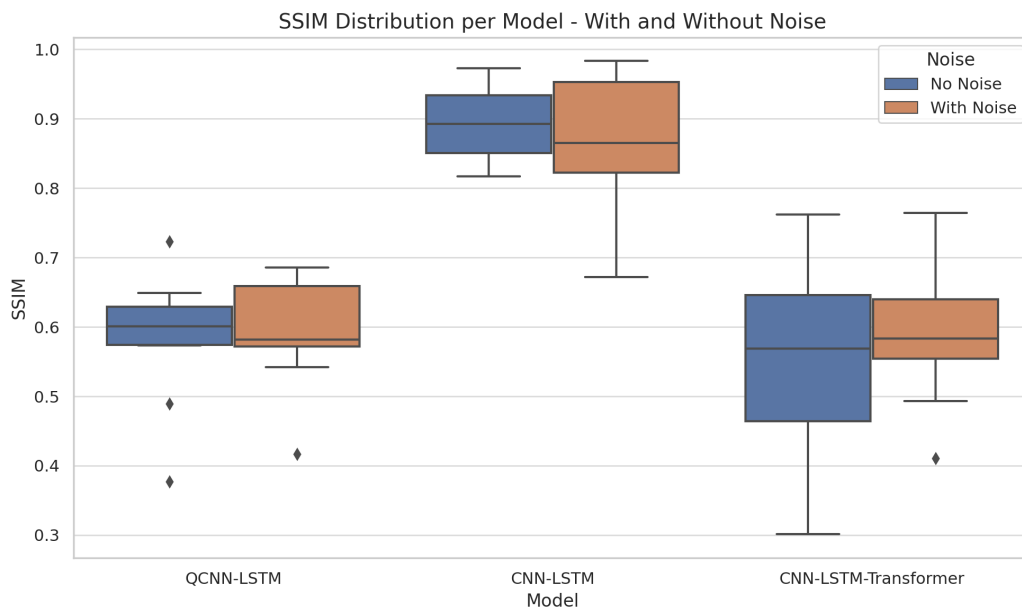


Figure 9. Distribution of SSIM for the models with and without Gaussian noise.

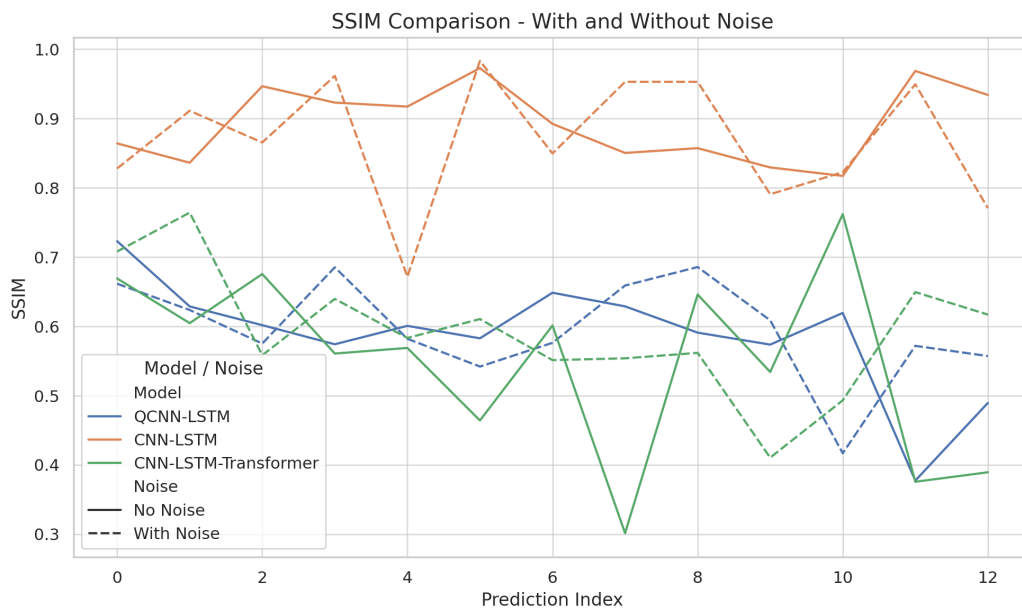


Figure 10. Evolution of SSIM across predictions: Noisy vs. clean images.

Table 2. Metrics of the three models.

Model	Noise Impact	RMSE	SSIM	Wasserstein
CNN-LSTM	Sensitive	↑ (degraded)	↑ (improved)	~ (moderate)
QCNN-LSTM	Robust	≈ (stable)	≈ (stable)	≈ (best and stable)
CNN-LSTM-Transformer	Inconsistent	↓ (slight)	↑ (less variance)	↑ (still unstable)

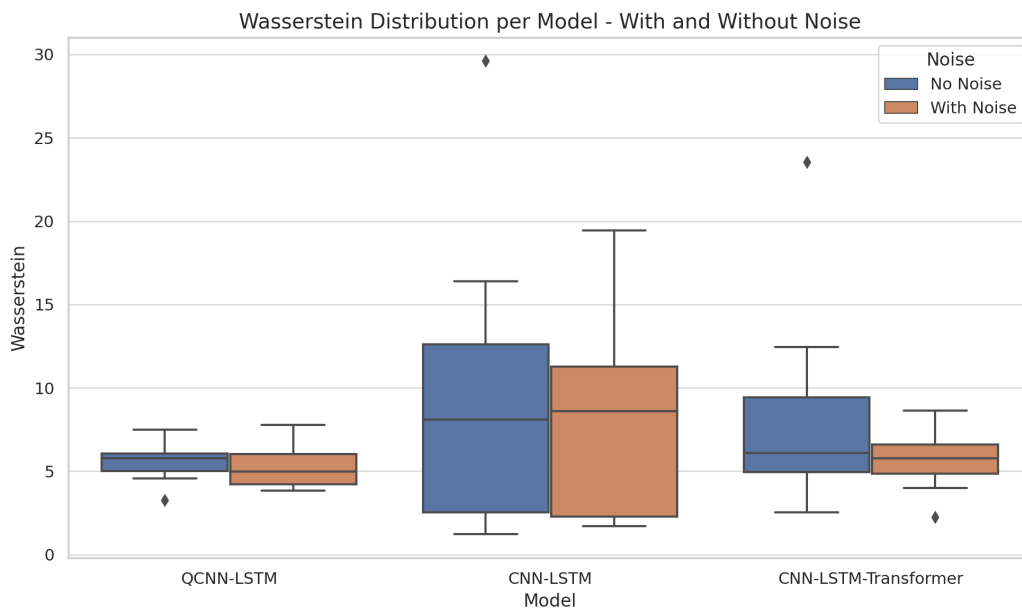


Figure 11. Wasserstein distance of the models with and without Gaussian noise.




Figure 12. Evolution of Wasserstein distance: Noisy vs. clean images.

Finally, CNN-LSTM excels on noise-free data, particularly in terms of visual fineness (SSIM) and raw accuracy (RMSE), but is very sensitive to noise, which can be problematic in real-world conditions. QCNN-LSTM is the most robust, with stable performance regardless of the noise level. It provides spatially reliable predictions (Wasserstein) and appreciable resilience, making it a good candidate for forecasting real-world extreme events. CNN-LSTM Transformer remains the most unstable model, showing little consistency and sensitivity to the input conditions, despite its theoretical complexity.

5.2. Enhanced quantum resilience with lightning.qubit backend

In this subsection, we present the results obtained with the quantum-enhanced model using a 10-qubit variational quantum circuit implemented via the lightning.qubit backend from PennyLane. This architecture represents the most advanced and scalable configuration tested in this study, designed to explore the full potential of quantum circuits for feature encoding in meteorological forecasting. Unlike the 8-qubit QCNN integrated via KerasLayer, this model uses a fully custom quantum layer embedded directly in the TensorFlow pipeline. It enables greater flexibility in its architectural design and enhanced numerical stability through direct quantum differentiation. This setup allows us to evaluate the impact of increased quantum expressivity on predictive accuracy and resilience. The following results assess the model's performance on clean, noisy, and adversarial (FGSM) inputs. Special attention is given to its robustness in maintaining high prediction quality under perturbations, which is a critical requirement for operational deployment in extreme weather forecasting systems.

The model's architecture remains hybrid, combining a time-distributed CNN block for spatial feature extraction, LSTM for temporal context encoding, and a variational quantum circuit applied to a condensed representation of the CNN features. The quantum layer is composed of 10 qubits and two layers of entangling gates (basic entangler layers), with backpropagation enabled through differentiable quantum simulation; see Figure 4.

5.2.1. Performance and robustness

The QCNN lightning was evaluated on the same meteorological image prediction task, using the last 13 time steps of the dataset for inference. The results, presented in Figures 13 and 14, show that this model achieves comparable accuracy to the standard QCNN, but with an even more stable profile across metrics.

To further assess the resilience of the lightning-based QCNN, we subjected it to FGSM adversarial perturbations ($\epsilon = 0.15$). The attack had virtually no effect on its prediction quality, as evidenced by minimal variations in RMSE, SSIM, and Wasserstein distance (see Table 3). This highlights the intrinsic robustness of quantum-enhanced architectures and supports their potential for use in safety-critical forecasting applications.

5.3. Adversarial evaluation using FGSM

To assess the robustness of the model, adversarial inputs x^{adv} were generated using FGSM and fed to the full hybrid QCNN model. The resulting predictions were compared with those obtained from clean inputs x^{clean} using three key metrics: RMSE, SSIM, and Wasserstein distance.

Despite the application of strong perturbations ($\epsilon = 0.15$), the QCNN lightning 10-qubit model exhibited remarkable resilience. The differences between clean and adversarial predictions were negligible across all evaluation metrics. These findings underscore the inherent stability of the quantum-enhanced representation and its robustness against adversarial attacks targeting early convolutional layers.

According to the comparative table of the performance of the 10-qubit QCNN lightning model, with and without FGSM attack, negligible variations are observed, which confirms the exceptional robustness of the architecture.

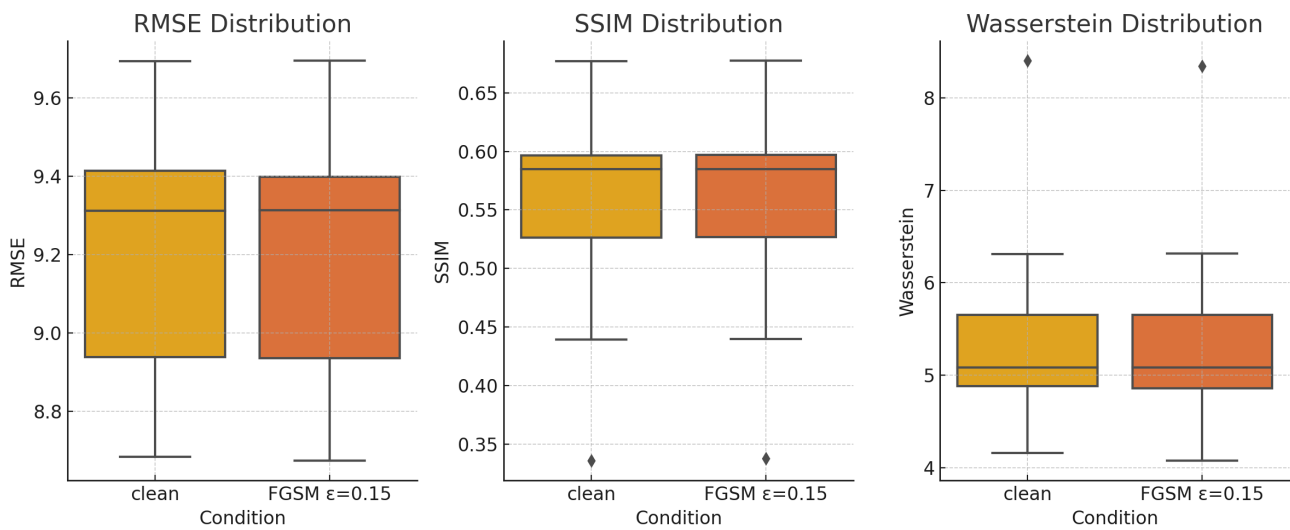


Figure 13. Boxplot comparison of the QCNN lightning (10 qubits) model’s performance on clean inputs versus FGSM adversarial perturbations ($\epsilon = 0.15$). The RMSE, SSIM, and Wasserstein distance distributions remain nearly identical, highlighting the model’s strong robustness.

Table 3. Performance statistics of the QCNN lightning model (10 qubits) on the original data and under FGSM attack ($\epsilon = 0.15$). The relative variations (Δ) between the two conditions are negligible, which confirms the stability of the model.

Metric	Mean (clean)	Std (clean)	Mean (FGSM)	Std (FGSM)	Δ (%)
RMSE	9.270	0.280	9.270	0.280	0.00
SSIM	0.578	0.077	0.577	0.078	-0.17
Wasserstein	5.080	0.600	5.070	0.600	-0.20

Table 3 highlights the remarkable stability of the QCNN lightning (10 qubits) model under a FGSM adversarial attack with $\epsilon = 0.15$. The differences in performance between clean and perturbed predictions are negligible, with a variation of only -0.17% in SSIM and -0.20% in Wasserstein distance. The RMSE remains strictly unchanged. These results confirm that the hybrid quantum architecture simulated with lightning.qubit exhibits an inherent robustness to localized perturbations, which is a major asset for critical applications in climatology or safety-sensitive forecasting. This quantum resilience could be further enhanced with explicit defense strategies and represents a promising direction toward achieving practical quantum advantage.

These observations validate the robustness of the QCNN under adversarial conditions. To further investigate the role of meteorological inputs in stabilizing predictions, we conducted an ablation study by removing CAPE from the inputs.

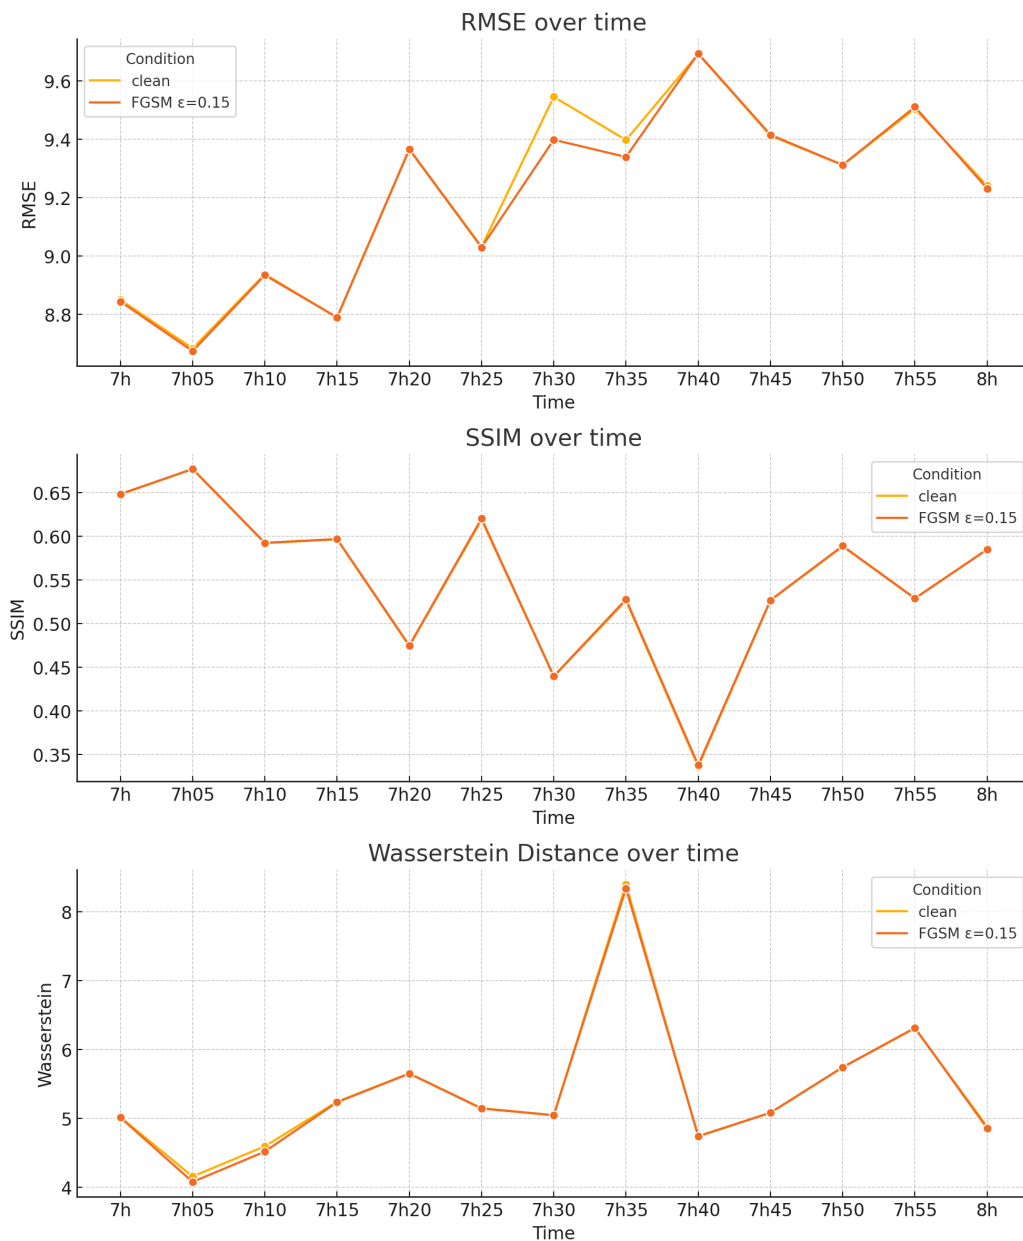


Figure 14. Evolution of RMSE, SSIM, and Wasserstein distance over time for the QCNN lightning (10 qubits) model under clean and adversarial (FGSM $\epsilon = 0.15$) inputs. The nearly overlapping curves confirm the exceptional resistance of the model to perturbations.

5.4. Ablation analysis: role of CAPE in the 10-qubit QCNN

To further assess the impact of meteorological inputs on the quantum model's performance, we conducted an ablation experiment comparing the 10-qubit QCNN trained with three inputs (satellite imagery, lightning radar, and CAPE maps) versus a reduced version with only two inputs (satellite imagery and lightning radar). All other components of the model and training procedure remained unchanged.

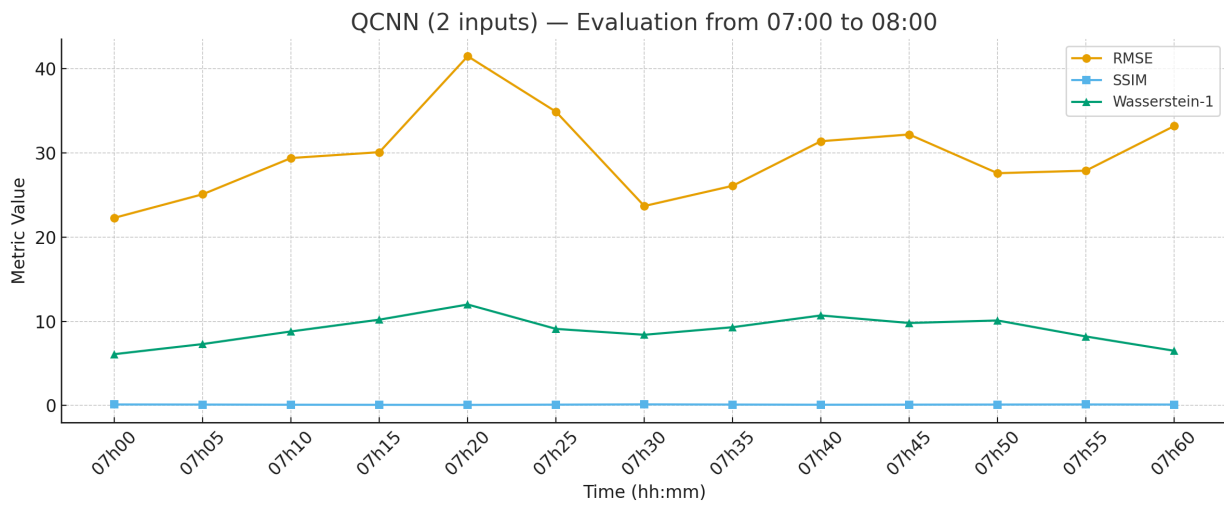


Figure 15. QCNN-10 model with two inputs (satellite + lightning). RMSE, SSIM, and Wasserstein-1 metrics were computed every 5 minutes from 0700 to 0800 h.

The comparison (see Figure 15) reveals that the reduced-input model achieves slightly better RMSE values during the early phase of the event (0700–0730 h), when convective activity is still moderate. However, it exhibits a significant degradation in all metrics during the peak convective dynamics (0730–0800 h), with RMSE spikes exceeding 30, reduced SSIM, and large Wasserstein distances.

In contrast, the full model with CAPE maintains stable and lower RMSE values, superior SSIM (~ 0.6), and significantly better Wasserstein-1 scores across the entire hour. Furthermore, under FGSM perturbations ($\epsilon = 0.15$), the three-input model shows greater robustness with minimal deviation between clean and perturbed predictions.

These results suggest that thermodynamic inputs like CAPE provide a stabilizing effect during a storm’s intensification, enhancing both the accuracy and the resilience of quantum-enhanced architectures under dynamic and potentially adversarial conditions.

5.5. Final comparative analysis

The comparative results, illustrated in Figure 16, provide a clear overview of each model’s behavior across all evaluation metrics. The classical CNN-LSTM and CNN-LSTM+Transformer models exhibit high variability in both RMSE and Wasserstein distance, particularly when subjected to Gaussian noise. Although the Transformer variant slightly improves the SSIM in certain sequences, it remains highly sensitive to perturbations.

The QCNN model, based on eight qubits and implemented via default.qubit with *KerasLayer*, demonstrates consistent performance and superior robustness. However, the most striking result emerges from the implementation of QCNN lightning (10 qubits). Not only does this model achieve lower error rates, but it also maintains invariant metrics under FGSM adversarial attacks, as previously demonstrated. The ablation results highlight the added value of CAPE as input in enhancing the QCNN’s temporal robustness and predictive sharpness.

These findings suggest that hybrid quantum architectures, particularly those disentangled from

Global Boxplots Comparison of Model Performance

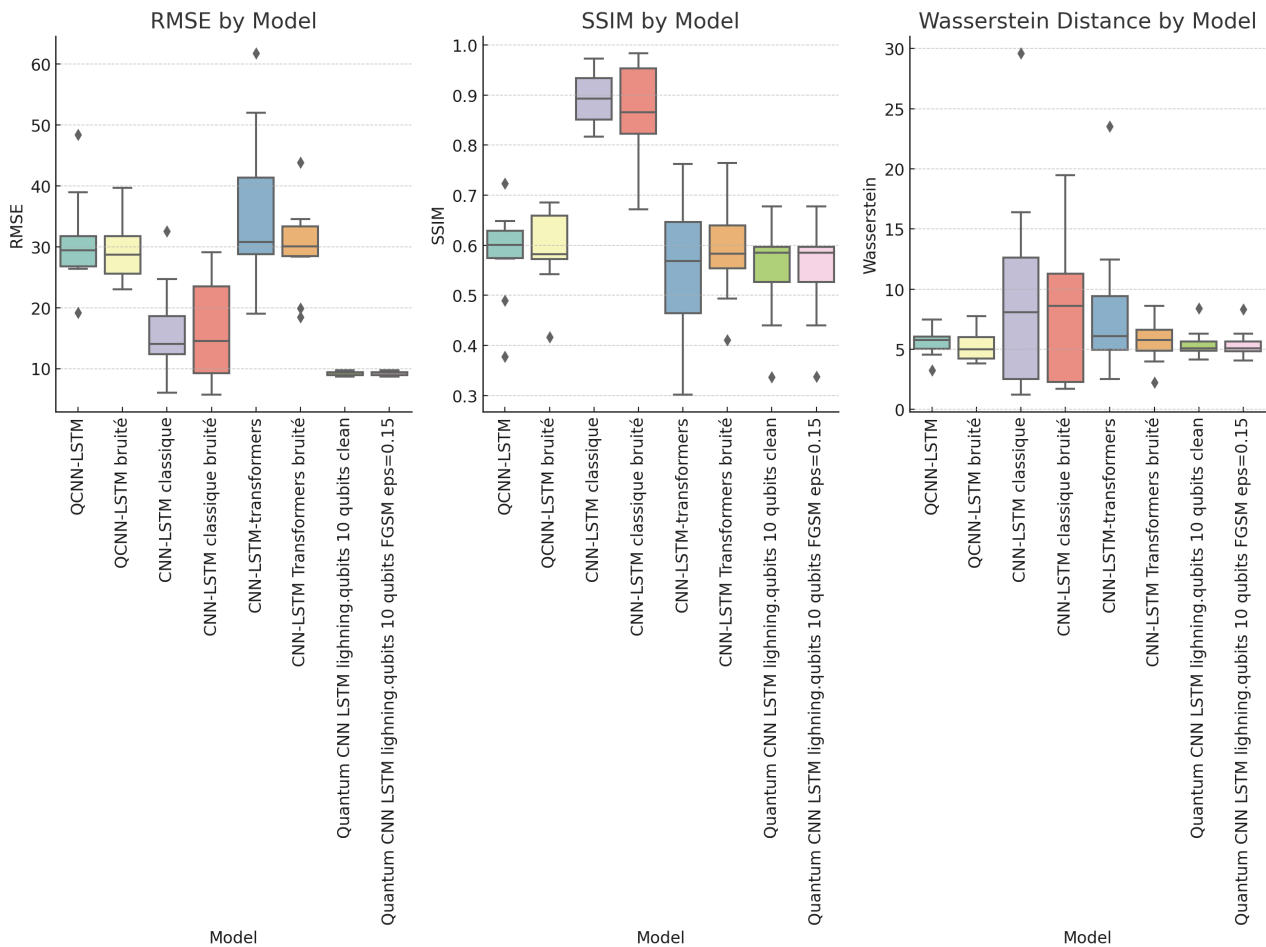


Figure 16. Global boxplots comparison of all models in terms of RMSE, SSIM, and Wasserstein metrics under clean and noisy conditions.

Keras abstraction layers and executed on performant simulators like lightning.qubit, offer compelling advantages in terms of both accuracy and adversarial robustness. This opens up a promising avenue for future deployments on real quantum hardware and establishes a strong precedent for practical quantum machine learning in complex real-world applications such as forecasting extreme weather.

6. Perspectives for real quantum hardware deployment

This study demonstrates the effectiveness and robustness of hybrid quantum–classical models for the short-term prediction of extreme weather events.

6.1. Limitations

While the proposed quantum-enhanced framework demonstrates stable and promising results, several limitations should be acknowledged to contextualize the current findings.

First, the present study focuses on a single meteorological event—a three-hour bow echo system

comprising 37 satellite–radar images—used as a proof of concept to evaluate the hybrid QCNN architecture. No cross-validation was performed across independent storms or seasonal datasets, which limits the generalization of the model. Future work will expand the dataset to multiple convective systems with a similar topology in order to assess robustness across time and space.

Second, in the robustness analysis, the FGSM adversarial attack was applied only up to the quantum interface layer. Because gradients were not propagated through the quantum circuit, the attack corresponds to a partial white-box scenario. This constraint reflects the current limitations in differentiating through quantum nodes within hybrid TensorFlow–PennyLane frameworks. In future implementations, fully end-to-end differentiable attacks will be explored to assess the intrinsic robustness of the quantum component.

Third, the training process was performed using the lightning.qubit simulator, which scales exponentially with the number of qubits. Although our 10-qubit QCNN with depth 2 remains computationally tractable, the training time per epoch is non-negligible (70 minutes for 150 epochs). Further work will evaluate deployment on real noisy intermediate-scale quantum (NISQ) devices, where additional factors such as sampling noise, decoherence, and readout error mitigation must be considered.

Overall, these limitations are consistent with the exploratory nature of this study and point toward concrete research directions. The forthcoming stages of our work aim to combine larger-scale data assimilation with execution on photonic or superconducting quantum platforms, thereby closing the loop between simulation and experimental quantum learning.

These encouraging results open several promising directions for future research and real-world deployment. Although our experiments were conducted using state-of-the-art simulators, the ultimate goal is to run the QCNN architecture on real quantum hardware. We are currently exploring deployment on platforms such as Quandela MosaiQ (photonics-based) [47], Amazon Braket [48], and OVHcloud [49] in collaboration with Quandela. These efforts aim to assess the feasibility of executing end-to-end pipelines on actual quantum processors, starting with inference tasks and eventually moving to full training loops. To adapt QCNN architectures to real quantum devices, future work will focus on optimizing quantum circuits to meet hardware constraints such as limited qubit counts, connectivity, and noise levels. Techniques such as circuit transpilation, qubit mapping, and parameter freezing during inference could significantly reduce the computational costs and enhance stability.

A long-term objective is the integration of quantum-enhanced models in edge-based or cloud-assisted weather forecasting systems, particularly for early warning of convective hazards. By offloading key inference components to quantum processors, we expect faster decision-making pipelines in meteorological monitoring, especially when latency is critical.

The demonstrated robustness of quantum-enhanced learning suggests potential applications beyond meteorology. Similar architectures could be transferred to domains such as climate modeling, earthquake prediction, or flood forecasting, where noisy, high-dimensional, and dynamic data are the norm.

In summary, the deployment of QCNNs on quantum hardware represents a natural next step toward operationalizing quantum machine learning for environmental science. While technical challenges remain, the convergence of robust hybrid models and advancing quantum hardware offers an exciting horizon for future innovation.

7. Conclusions and future work

This study introduced an innovative hybrid approach for short-term forecasting of a rare extreme weather phenomenon represented here by bow echo cloud formations. We compared three deep learning architectures, CNN-LSTM, CNN-LSTM-Transformer, and quantum CNN-LSTM (QCNN), using multi-source meteorological data (satellite imagery, CAPE fields, and lightning activity) from the storm of August 18, 2022. The models were trained on data collected from 0500 to 0655 h, and tested on the critical prediction window from 0700 to 0800 h, when the convective system formed and intensified. Among the key findings, the CNN-LSTM model achieved excellent perceptual performance on clean inputs, with SSIM values reaching up to 0.98. However, it was highly sensitive to Gaussian noise and adversarial perturbations. CNN-LSTM-Transformer, while theoretically well-suited for temporal modeling, exhibited instability and strong variability, particularly in terms of spatial coherence and Wasserstein distance. In contrast, QCNN-LSTM, integrating an 8-qubit variational quantum circuit, demonstrated excellent stability across all metrics (RMSE, SSIM, Wasserstein distance) under noisy conditions. More significantly, its enhanced version, the QCNN lightning (10 qubits) based on the PennyLane lightning.qubit backend maintained not only high visual and metric performance but also remarkable resilience to adversarial attacks. FGSM-based perturbations with $\epsilon = 0.15$ produced no noticeable degradation, suggesting the architecture's potential for adversarially robust forecasting. These results underscore the promise of quantum-enhanced neural networks in capturing compact, stable, and geometrically coherent representations of convective storm systems, even under degraded conditions.

Future research directions

Two main research directions are proposed to extend this work:

- 1) **Topological data analysis (TDA) on real bow echo cloud imagery from 0700 to 0800 h.** TDA will be used to extract robust topological descriptors (e.g., persistence diagrams, Betti numbers) reflecting the geometry and dynamics of bow echo structures.
- 2) **TDA of QCNN predictions.** Building on the observed stability of the QCNN architectures, particularly the lightning 10-qubit version, we aim to determine whether their predictions preserve the topological invariants that are characteristic of convective systems. This will assess the model's ability not only to reproduce visible cloud formations but also to encode the deeper dynamical and topological features crucial for anticipating extreme meteorological phenomena.

Use of AI tools declaration

The authors declare they have not used artificial intelligence (AI) tools in the creation of this article.

Conflict of interest

The authors declare that there is no conflict of interest.

References

1. Fujita TT, (1978) *Manual of downburst identification for project NIMROD* (No. NASA-CR-156953).
2. Przybylinski RW, (1995) The bow echo: Observations, numerical simulations, and severe weather detection methods. *Weather Forecast* 10: 203–218. [https://doi.org/10.1175/1520-0434\(1995\)010%3C0203:TBEONS%3E2.0.CO;2](https://doi.org/10.1175/1520-0434(1995)010%3C0203:TBEONS%3E2.0.CO;2)
3. Klimowsky BA, Hjelmfelt MR, Bunkers MJ, (2004) Radar observations of the early evolution of bow echoes. *Weather Forecast* 19: 727–734. [https://doi.org/10.1175/1520-0434\(2004\)019%3C0727:ROOTEE%3E2.0.CO;2](https://doi.org/10.1175/1520-0434(2004)019%3C0727:ROOTEE%3E2.0.CO;2)
4. Davis C, Atkins N, Bartels D, Bosart L, Bosart L, Coniglio M, Bryan G, et al. (2004) The bow echo and MCV experiment: Observations and opportunities: observations and opportunities. *Bull Am Meteorol Soc* 85: 1075–1093. <https://doi.org/10.1175/BAMS-85-8-1075>
5. Kamani MM, Farhat F, Wistar S, Wang JZ, (2016) Shape matching using skeleton context for automated bow echo detection, In: *2016 IEEE International Conference on Big Data (Big Data)*, 901–908. <https://doi.org/10.1109/BigData.2016.7840685>
6. Zeng F, Hou J, (2018) Automatic detection of bow echoes from weather radar images, In: *2018 13th World Congress on Intelligent Control and Automation (WCICA)*, 503–508. <https://doi.org/10.1109/WCICA.2018.8630468>
7. Patil V, Das SK, Phadke A, (2019) Methods for mesoscale convective systems detection and tracking: A survey, In: *2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, 1–7. <https://doi.org/10.1109/ICCCNT45670.2019.8944656>
8. Johns RH, Hirt WD, (1987) Derechos: Widespread convectively induced windstorms. *Weather Forecast* 2: 32–49. [https://doi.org/10.1175/1520-0434\(1987\)002%3C0032:DWCIW%3E2.0.CO;2](https://doi.org/10.1175/1520-0434(1987)002%3C0032:DWCIW%3E2.0.CO;2)
9. Corfidi SF, Coniglio MC, Cohen AE, Mead CM, (2016) A proposed revision to the definition of “derecho”. *Bull Am Meteorol Soc* 97: 935–949. <https://doi.org/10.1175/BAMS-D-14-00254.1>
10. Canot H, Durand P, Frénod E, (2025) Bow echo alarm system using topological data analysis. *Appl Math Mod Challenges* 3: 44–63. <https://doi.org/10.3934/ammc.2025002>
11. Půčik T, The Derecho and Hailstorms of 18 August 2022, European Severe Storms Laboratory, 2022. Available from: <https://www.essl.org/cms/the-derecho-and-hailstorms-of-18-august-2022/>.
12. Fery L, Faranda D, (2023) Changes in synoptic circulations associated with documented derechos over France in the past 70 years. *Weather Clim Dyn* 2023. <https://doi.org/10.5194/wcd-2023-8>
13. Bernatz R, Bow Echo and Strong Rear Inflow Jet, 2023. Available from: https://www.faculty.luther.edu/bernatzr/Courses/Sci123/comet/radar/severe_signatures/navmenu.php_printname_print_bow_echo.htm_page_1.0.0.htm.
14. Lee WC, Wakimoto RM, Carbone RE, (1992) The evolution and structure of a “bow-echo-microburst” event. Part II: The bow echo. *Mon Weather Rev* 120: 2211–2225. [https://doi.org/10.1175/1520-0493\(1992\)120%3C2211:TEASOA%3E2.0.CO;2](https://doi.org/10.1175/1520-0493(1992)120%3C2211:TEASOA%3E2.0.CO;2)

15. Parker MD, Johnson RH, (2000) Organizational modes of midlatitude mesoscale convective systems. *Mon Weather Rev* 128: 3413–3436. [https://doi.org/10.1175/1520-0493\(2001\)129%3C3413:OMOMMC%3E2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129%3C3413:OMOMMC%3E2.0.CO;2)
16. Elhoseiny M, Huang S, Elgammal A, (2015) Weather classification with deep convolutional neural networks, In: *2015 IEEE International Conference on Image Processing (ICIP)*, 3349–3353. <https://doi.org/10.1109/ICIP.2015.7351424>
17. Trok JT, Barnes EA, Davenport FV, Diffenbaugh NS, (2024) Machine learning—Based extreme event attribution. *Sci Adv* 10: eadl3242. <https://doi.org/10.1126/sciadv.adl3242>
18. Liu Y, Racah E, Correa J, Khosrowshahi A, Lavers D, Kunkel K, et al. (2016) Application of deep convolutional neural networks for detecting extreme weather in climate datasets, preprint, arXiv:1605.01156. <https://doi.org/10.48550/arXiv.1605.01156>
19. Gagne II DJ, Haupt SE, Nychka DW, Thompson G, (2019) Interpretable deep learning for spatial analysis of severe hailstorms. *Mon Weather Rev* 147: 2827–2845. <https://doi.org/10.1175/MWR-D-18-0316.1>
20. Lagerquist R, McGovern A, Homeyer CR, Gagne II DJ, Smith T, (2020) Deep learning on three-dimensional multiscale data for next-hour tornado prediction. *Mon Weather Rev* 148: 2837–2861. <https://doi.org/10.1175/MWR-D-19-0372.1>
21. G Ayzel, M Heistermann, A Sorokin, O Nikitin, O Lukyanova, (2019) All convolutional neural networks for radar-based precipitation nowcasting. *Proc Comput Sci* 150: 186–192. <https://doi.org/10.1016/j.procs.2019.02.036>
22. Price I, Sanchez-Gonzalez A, Alet F, Andersson TR, El-Kadi A, Masters D, et al. (2025) Probabilistic weather forecasting with machine learning. *Nature* 637: 84–90. <https://doi.org/10.1038/s41586-024-08252-9>
23. Mounier A, Raynaud L, Rottner L, Plu M, Arbogast P, Kreitz M, et al. (2022) Detection of bow echoes in kilometer-scale forecasts using a convolutional neural network. *Artif Intell Earth Syst* 1: e210010. <https://doi.org/10.1175/AIES-D-21-0010.1>
24. Gauch M, Kratzert F, Klotz D, Nearing G, Lin J, Hochreiter S, (2021) Rainfall-runoff prediction at multiple timescales with a single long short-term memory network. *Hydrol Earth Syst Sci* 25: 2045–2062. <https://doi.org/10.5194/hess-25-2045-2021>
25. Kratzert F, Klotz D, Shalev G, Klambauer G, Hochreiter S, Nearing G, (2019) Towards learning universal, regional, and local hydrological behaviors via machine learning applied to large-sample datasets. *Hydrol Earth Syst Sci* 23: 5089–5110. <https://doi.org/10.5194/hess-23-5089-2019>
26. Shi X, Chen Z, Wang H, Yeung DY, Wong WK, Woo WC, (2015) Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Adv Neural Inf Process Syst* 2015: 28.
27. Tekin SF, Fazla A, Kozat SS, (2024) Numerical weather forecasting using convolutional-LSTM with attention and context matcher mechanisms. *IEEE Trans Geosci Remote Sens* 62: 1–13. <https://doi.org/10.1109/TGRS.2024.3409084>
28. Cong I, Choi S, Lukin MD, (2019) Quantum convolutional neural networks. *Nat Phys* 15: 1273–1278. <https://doi.org/10.1038/s41567-019-0648-8>

29. Moreira MS, Guerreschi GG, Vlothuizen W, Marques JF, van Straten J, Premaratne SP, et al. (2023) Realization of a quantum neural network using repeat-until-success circuits in a superconducting quantum processor. *NPJ Quantum Inf* 9:118. <https://doi.org/10.1038/s41534-023-00779-5>
30. Herrmann J, Llima SM, Remm A, Zapletal P, McMahon NA, Scarato C, et al. (2022) Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases. *Nat Commun* 13: 4144. <https://doi.org/10.1038/s41467-022-31679-5>
31. Song Y, Li J, Wu Y, Qin S, Wen Q, Gao F, (2024) A resource-efficient quantum convolutional neural network. *Front Phys* 12: 1362690. <https://doi.org/10.3389/fphy.2024.1362690>
32. Gong LH, Pei JJ, Zhang TF, Zhou NR, (2024) Quantum convolutional neural network based on variational quantum circuits. *Opt Commun* 550: 129993. <https://doi.org/10.1016/j.optcom.2023.129993>
33. Pastori L, Grundner A, Eyring V, Schwabe M, (2025) Quantum neural networks for cloud cover parameterizations in climate model, preprint, arXiv:2502.10131. <https://doi.org/10.48550/arXiv.2502.10131>
34. Lachure S, Damahe L, Lachure J, Sawarkar A, Bhati SS, Chhabra R, et al. (2024) A lightweight hybrid quantum convolution neural network for temperature forecasting, In: *AIP Conference Proceedings* 3188: 090008. <https://doi.org/10.1063/5.0240207>
35. Bazgir A, Zhang Y, (2024) QESM: A leap towards quantum-enhanced ML emulation framework for earth and climate modeling, preprint, arXiv:2410.01551. <https://doi.org/10.48550/arXiv.2410.01551>
36. Enos GR, Reagor MJ, Henderson MP, Young C, Horton K, Birch M, et al. (2021) Synthetic weather radar using hybrid quantum-classical machine learning, preprint, arXiv:2111.15605. <https://doi.org/10.48550/arXiv.2111.15605>
37. Eswaran U, Eswaran V, Murali K, Eswaran V, (2024) Quantum-based predictive modeling for extreme weather events, In: *The Rise of Quantum Computing in Industry 6.0 Towards Sustainability: Revolutionizing Smart Disaster Management*, Cham: Springer Nature Switzerland, 123–140. https://doi.org/10.1007/978-3-031-73350-5_8
38. Goodfellow IJ, Shlens J, Szegedy C, (2014) Explaining and harnessing adversarial examples, preprint, arXiv:1412.6572. <https://doi.org/10.48550/arXiv.1412.6572>
39. Griffiths DJ, Schroeter DF, (2018) *Introduction to Quantum Mechanics*, Cambridge University Press. <https://doi.org/10.1017/9781316995433>
40. Hestenes D, (1999) *New Foundations for Classical Mechanics*, Dordrecht: Springer Netherlands. <https://doi.org/10.1007/0-306-47122-1>
41. Lounesto P, (2001) Clifford algebras and spinors, In: *Clifford Algebras and Their Applications in Mathematical Physics*, Dordrecht: Springer Netherlands, 25–37. https://doi.org/10.1007/978-94-009-4728-3_2
42. Nakahara M, (2018) *Geometry, Topology and Physics*, CRC Press. <https://doi.org/10.1201/9781315275826>

43. Nielsen MA, Chuang IL, (2010) *Quantum Computation and Quantum Information*, Cambridge University Press.
44. EUMETSAT, Données Satellitaires de l'Orage en Corse du 18 Août 2022, 2022. Available from: <https://www.eumetsat.int/>.
45. Météo-France, Données Météorologiques de l'Orage en Corse du 18 Août 2022, 2022. Available from: <https://donneespubliques.meteofrance.fr>
46. Météorologix, Données Météorologiques de l'orage en Corse du 18 Août 2022, 2022. Available from: <https://www.meteorologix.com>.
47. Quandela, MosaiQ-Quandela's Photonic Quantum Cloud Platform, 2023. Available from: <https://www.quandela.com/mosaiq>.
48. Amazon Web Services, Amazon Braket: Explore and Experiment with Quantum Computing, 2023. Available from: <https://aws.amazon.com/braket/>.
49. OVHcloud, OVHcloud Quantum Project, 2023. Available from: <https://www.ovhcloud.com/en/quantum/>.



AIMS Press

© 2025 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)