Quantum Computing for Precision Agriculture in Challenging Environments: A Case Study from Northern Morocco

Mohamed Ben Ahmed, Anouar A. Boudhir, and Aziz Mahboub

C3S Laboratory, FSTT, UAE University, Morocco
mbenahmed@uae.ac.ma

Keywords: Quantum Computing, Smart Farming, Medical Cannabis, Remote Sensing, GIS, QAOA, VQLS, VQC

Abstract

The legalization of medical cannabis in Morocco's northern Rif region requires precision agriculture systems capable of supporting highly controlled, traceable and quality-driven cultivation. Medical cannabis is biologically sensitive to micro-variations in soil moisture, vapor pressure deficit (VPD), canopy temperature and nutrient levels, which makes it a demanding testbed for advanced decision-support methods. In this work, we propose and numerically evaluate an end-to-end hybrid quantum–classical framework that combines IoT sensor networks, Sentinel-2 and UAV imagery, GIS integration and quantum-enhanced analytics for regulated medical cannabis cultivation in the Al-Hoceïma region.

The framework instantiates three quantum modules: (i) a variational quantum linear solver (VQLS) for Kriging-based spatial interpolation under sparse sensing, (ii) a variational quantum classifier (VQC) for early stress detection from multi-source features, and (iii) a Quantum Approximate Optimization Algorithm (QAOA) for constrained irrigation scheduling. All experiments are conducted on synthetic yet agro-ecologically calibrated data generated for a 4-hectare virtual plot; no real cannabis-field data or quantum hardware are used. In this controlled simulation setting, the quantum-inspired modules achieve moderate improvements over classical baselines (Kriging, Random Forest, neural networks, MILP), for example reducing interpolation RMSE by about 20% and improving early-stress F1-score by several percentage points.

We explicitly do not claim hardware-level quantum advantage, nor do we provide a formal proof that VQLS or VQC must outperform classical Kriging or machine learning in this regime. Instead, the contribution is a transparent formulation and simulation-based assessment of quantum-compatible workflows for precision agriculture in regulated contexts, together with a critical discussion of their current limitations and the conditions under which they might become competitive in practice.

1. INTRODUCTION

The global medical cannabis industry is expected to grow substantially in the coming years, driven by stringent requirements on product quality, environmental control and traceability (Grand View Research, 2024). In 2021, Morocco legalized medical and industrial cannabis under Law 13-21, authorizing production in selected areas of the Rif mountains including Al-Hoceïma (Kingdom of Morocco, 2021). This shift creates an opportunity for digitally enhanced agriculture in a region characterized by steep topography, heterogeneous soils and strong climatic gradients.

Medical cannabis (*Cannabis sativa L.*) is highly sensitive to its micro-environment. Optimal concentrations of cannabinoids and terpenes require maintaining soil moisture within narrow thresholds (typically 60–80% of field capacity), canopy temperatures in a relatively narrow band (e.g. 22–28°C), and VPD between about 0.8 and 1.2 kPa (Bernstein et al., 2019). Even small spatial variations can significantly impact biochemical profiles and pharmaceutical quality, motivating fine-grained spatial monitoring and control.

Quantum computing, particularly hybrid near-term quantum–classical algorithms, provides new tools for optimization, spatial modelling and classification. Recent advances in variational quantum algorithms (VQAs) demonstrate potential advantages in solving certain linear systems, kernel-based classification tasks and combinatorial optimization problems (Cerezo et al., 2021). These characteristics align conceptually with some computational challenges of precision agriculture, such as data-fusion-based interpolation, detection of subtle stress patterns and resource-constrained planning (Pook et al., 2025). However, current noisy intermediate-scale quantum (NISQ) hardware remains limited in qubit count, coherence and gate fidelity, and practical quantum advantage in real agricultural settings has not yet been demonstrated.

This paper therefore adopts a cautious and simulation-based stance: we propose a hybrid quantum–classical architecture for medical cannabis cultivation and evaluate it on a fully synthetic but agro-ecologically calibrated virtual field. The goal is not to claim superiority of quantum methods today, but to (i) formulate quantum-compatible variants of core precision-agriculture tasks and (ii) quantify, under idealized and noisy simulation, the potential performance gains one might expect on small instances, while explicitly discussing their computational cost and theoretical plausibility.

Scope and Positioning

More concretely, this work should be read as a conceptual and numerical feasibility study:

- All data are synthetic and generated from stochastic models calibrated using generic climatic and agronomic ranges consistent with the Al-Hoceïma region; no real cannabis-field measurements are used.
- All quantum algorithms are executed with Qiskit Aer simulators, either under ideal, noiseless conditions or with simple noise models. There is no claim of hardware-level quantum advantage.

- Problem sizes (number of qubits, number of zones, spatial resolution) are intentionally small, so that both quantum and classical baselines are tractable and results are easy to interpret.

Within this strictly defined scope, we investigate three research questions:

- **RQ1:** How can Kriging-based spatial interpolation under sparse IoT sensing be formulated as a VQLS problem, and how do its simulated results compare to classical Kriging on small instances?
- **RQ2:** Can a variational quantum classifier increase sensitivity to early stress conditions, in simulation, compared to competitive classical machine learning models?
- **RQ3:** Can a QAOA formulation of irrigation scheduling capture multi-objective trade-offs (water use, constraint violations, schedule smoothness) in small, illustrative scenarios?

Our main contributions are:

- We design an end-to-end hybrid architecture that couples IoT sensing, UAV/Sentinel-2 remote sensing, GIS databases and quantum modules tailored to regulated medical cannabis in Al-Hoceïma, and we articulate its scope as a simulation framework.
- We formulate Kriging, stress detection and irrigation scheduling as VQLS, VQC and QAOA problems, respectively, and provide explicit mathematical encodings suitable for NISQ devices on small instances.
- We build a synthetic but agro-ecologically calibrated dataset for a 4-hectare virtual plot and benchmark quantum modules against classical baselines (Kriging, Random Forest, neural networks, MILP), reporting illustrative performance differences under idealized and noisy simulation.
- We discuss methodological, computational and contextual limitations (e.g. lack of real data, limited classical baselines, absence of advanced noise modelling) and outline more rigorous experiments, including sensitivity analysis and extended benchmarks, that would be required for stronger conclusions.

2. STUDY AREA AND DATA SOURCES

The study area is defined as a 4-hectare *virtual* plot located near Ait Youssef, Al-Hoceïma (35.25°N, 3.96°W). This location is chosen to align with the legal cannabis-production zones under Law 13-21, but no real field data are accessed. Elevation is assumed to range from 210 to 590 m, with slopes exceeding 30% in some areas, and soils to vary from sandy-loam in lowlands to clay-loam on elevated terraces. The micro-climate is parameterized as semi-humid Mediterranean, with annual precipitation between 600 and 800 mm.

2.1 Remote Sensing Data (Simulated)

We emulate the following remote sensing data streams:

- **Sentinel-2 Level-2A imagery**: top-of-canopy reflectance in visible, NIR and SWIR bands at 10 m resolution, used for NDVI, NDMI and LAI extraction. Acquisition dates are sampled every 5–10 days under assumed clear-sky conditions.
- **UAV multispectral flights**: 8-band camera (550–850 nm), nominal 5 cm GSD, resampled to the common modeling grid.
- **DSM/DTM layers**: simulated from a base DEM with added noise to mimic UAV photogrammetry.

2.2 GIS Layers

GIS layers are represented as raster and vector data structures:

- topographic elevation raster (10 m resolution),
- soil maps (horizon depth, texture, organic matter percentage),
- field parcels and irrigation zones (polygons),
- UAV-derived NDVI rasters resampled to the modeling grid.

2.3 Feature Set for Hybrid Analytics

The different data sources are combined into a unified feature set used by both quantum and classical models. The full list of variables is summarized in Table 1.

3. QUANTUM COMPUTING BACKGROUND

Quantum computing exploits quantum mechanical principles to process information using qubits. Unlike classical bits taking values in $\{0, 1\}$, a qubit encodes

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad |\alpha|^2 + |\beta|^2 = 1, \tag{1}$$

where $(\alpha, \beta) \in \mathbb{C}^2$ define a point on the Bloch sphere.

Applying a Hadamard gate creates superposition:

$$H|0\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}. \tag{2}$$

Two-qubit entanglement is introduced using CNOT:

$$\mathrm{CNOT}\Big(\tfrac{|0\rangle+|1\rangle}{\sqrt{2}} \otimes |0\rangle\Big) = \frac{|00\rangle + |11\rangle}{\sqrt{2}}. \tag{3}$$

We use single-qubit rotations $R_X(\theta)$, $R_Y(\theta)$, $R_Z(\theta)$, entangling gates (CNOT, CRZ) and variational layers with parameter vector $\boldsymbol{\theta}$.

Variational quantum algorithms (VQAs) exploit NISQ hardware by combining parameterized quantum circuits with classical optimizers. In this work we employ:

- VQLS for linear system solving (Kriging),
- VQC for multi-class stress classification,
- QAOA for combinatorial irrigation optimization.

Table 1: Overview of input variables used in the hybrid quantum–classical analytics pipeline (all variables are simulated but agro-ecologically calibrated).

Source	Variable	Symbol	Spatial Resolution	Temporal Resolution	Used in
IoT sensors	Soil moisture (vol. content)	Z_{sm}	Point (36 nodes)	Hourly / daily aggregates	VQLS, VQC, QAOA
IoT sensors	Canopy temperature	Z_{ct}	Point (36 nodes)	Hourly / daily	VQLS, VQC
IoT sensors	Vapor pressure deficit (VPD)	Z_{vpd}	Point (36 nodes)	Hourly / daily	VQLS, VQC
Sentinel-2 L2A	NDVI, NDMI	$Z_{\mathrm{ndvi}}, Z_{\mathrm{ndmi}}$	10 m (upsampled to 8 m grid)	5–10 days (clear-sky)	VQLS, VQC
UAV multispectral	High-resolution NDVI	–	5 cm (resampled to 8 m)	Campaign-based (weekly)	VQLS, VQC
UAV DTM/DSM	Elevation, slope, aspect	–	10 m (derived)	Static	VQLS, VQC, QAOA
Soil maps	Texture, horizon depth, OM%	–	Polygon / rasterized at 8 m	Static	VQLS, VQC
Irrigation system data	Zone layout, valve capacity	–	Field-level polygons	Static	QAOA
Meteorological forcing	Temperature, radiation, ET_0	–	Gridded / station-based	Daily	QAOA, stress labelling

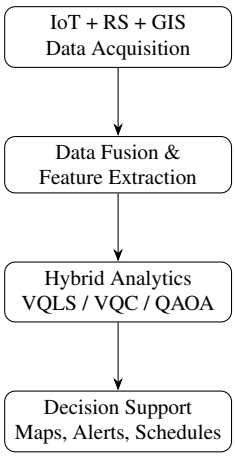

Figure 1: Compact representation of the hybrid quantum–classical architecture, highlighting the main logical layers.

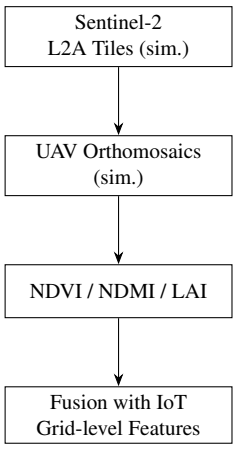

Figure 2: Compact remote sensing pipeline from simulated imagery to fused RS–IoT features.

4. HYBRID QUANTUM–CLASSICAL ARCHITECTURE

The proposed system integrates IoT sensing, remote sensing, GIS layers and quantum-enhanced analytics. The design follows a four-layer conceptual architecture: (i) data acquisition, (ii) preprocessing and feature extraction, (iii) hybrid analytics (classical + quantum), and (iv) decision support. Given the focus on numerical feasibility, we represent the architecture in a compact, single-column schematic (Figure 1) to avoid layout issues.

5. REMOTE SENSING WORKFLOW

Remote sensing is used to complement IoT sensors for monitoring vegetation vigor, water stress and soil conditions. The workflow integrates both satellite and UAV imagery.

From Sentinel-2, vegetation indices are computed as:

$$\mathrm{NDVI} = \frac{NIR - RED}{NIR + RED}, \qquad \mathrm{NDMI} = \frac{NIR - SWIR}{NIR + SWIR}. \tag{4}$$

UAV imagery provides fine-grained NDVI maps (5 cm) that are resampled to the modeling grid to capture intra-parcel variability.

A compact remote sensing pipeline is shown in Figure 2.

6. MODELING AND QUANTUM ALGORITHMS

6.1 Spatio-Temporal Environmental Field

We denote the environmental vector at location $\mathbf{s}$ and time t as

$$\mathbf{X}(\mathbf{s}, t) = \left[Z_{\mathrm{sm}}, Z_{\mathrm{ct}}, Z_{\mathrm{vpd}}, Z_{\mathrm{ndvi}}\right]^{\top}. \tag{5}$$

Spatial reconstruction aims to estimate

$$\hat{\mathbf{X}}(\mathbf{s}, t) = \sum_{i=1}^{n} \lambda_i(t)\, \mathbf{X}(\mathbf{s}_i, t), \tag{6}$$

where $\mathbf{s}_i$ denote sensor locations and $\lambda_i(t)$ are Kriging weights.

6.2 Quantum Kriging via VQLS

Classical Kriging solves a linear system

$$A\mathbf{y} = \mathbf{b}. \tag{7}$$

VQLS approximates $A^{-1}\mathbf{b}$ by minimizing

$$\mathcal{L}_{\mathrm{VQLS}}(\boldsymbol{\theta}) = \big\| A|\psi(\boldsymbol{\theta})\rangle - |\mathbf{b}\rangle \big\|^2. \tag{8}$$

A generic ansatz has the form

$$|\psi(\boldsymbol{\theta})\rangle = \prod_{l=1}^{L} \left[\bigotimes_{i=1}^{n} R_Y(\theta_{l,i})\, R_Z(\theta_{l,i+n}) \right] U_{\mathrm{ent}}^{(l)} |0\rangle^{\otimes n}, \tag{9}$$

where $U_{\mathrm{ent}}^{(l)}$ encodes entanglement patterns.

6.3 Quantum Classifier for Stress Detection

Feature vectors (soil moisture, canopy temperature, VPD, NDVI, NDMI and topographic covariates) are standardized and reduced by PCA to a low-dimensional representation compatible with the available number of qubits. They are then encoded using amplitude encoding:

$$|\mathbf{x}\rangle = \frac{1}{\|\mathbf{x}\|} \sum_j x_j |j\rangle. \tag{10}$$

A parameterized circuit $U(\boldsymbol{\phi})$ is applied and a subset of qubits is measured to obtain class labels:

$$U(\boldsymbol{\phi})|\mathbf{x}\rangle \to |y\rangle, \tag{11}$$

where $y \in \{0, 1, 2\}$ corresponds to Normal, Early Stress, and Severe Stress in our simulated dataset.

6.4 QAOA for Irrigation Optimization

We consider irrigation decisions $u_{z,k} \in \{0, 1\}$ for zone z at time k, with irrigation dose

$$I_{z,k} = d_z\, u_{z,k}. \tag{12}$$

Soil moisture dynamics use a simple bucket model:

$$\theta_{z,k+1} = \theta_{z,k} + \alpha_z I_{z,k} - ET_{z,k} - D_{z,k}, \tag{13}$$

with $ET_{z,k}$ evapotranspiration and $D_{z,k}$ drainage.

The multi-objective cost is

$$\begin{aligned} J(\mathbf{u}) = \lambda_{\mathrm{w}} \sum_{z,k} I_{z,k} \\ + \lambda_{\mathrm{dev}} \sum_{z,k} \mathrm{dev}\big(\theta_{z,k}\big) \\ + \lambda_{\mathrm{sm}} \sum_{z,k} \big| u_{z,k} - u_{z,k-1} \big|, \end{aligned} \tag{14}$$

with $\mathrm{dev}(\theta_{z,k})$ penalizing deviations outside a target interval $[\theta^{\min}, \theta^{\max}]$.

We express $J(\mathbf{u})$ as a QUBO

$$\min_{\mathbf{x} \in \{0,1\}^N} \mathbf{x}^\top Q \mathbf{x}, \tag{15}$$

where $\mathbf{x}$ stacks all $u_{z,k}$ and Q encodes both objectives and soft constraints. The QAOA cost Hamiltonian is

$$H_C = \sum_{i \le j} Q_{ij} Z_i Z_j, \tag{16}$$

and the mixer is

$$H_M = \sum_i X_i. \tag{17}$$

The QAOA state of depth P is

$$|\boldsymbol{\gamma}, \boldsymbol{\beta}\rangle = \prod_{p=1}^{P} e^{-i\beta_p H_M} e^{-i\gamma_p H_C} |+\rangle^{\otimes N}. \tag{18}$$

7. THEORETICAL CONSIDERATIONS ON VQLS VS CLASSICAL KRIGING

From a theoretical standpoint, there is currently *no general proof* that VQLS-based solvers must outperform classical linear solvers for Kriging problems of the size considered here (36 sensors, moderate grid). Classical algorithms based on Cholesky factorization or conjugate gradients are highly optimized and numerically stable for such dimensions.

Potential advantages of VQLS are expected, if at all, in regimes where:

- the Kriging system matrix A becomes very large and structured (e.g. block Toeplitz, low-rank perturbations),
- efficient quantum oracles for preparing $|b\rangle$ and implementing A as a block-encoding are available,
- and the cost of classical dense or sparse linear algebra becomes prohibitive.

In contrast, our synthetic scenario remains in a small-scale regime where classical solvers are extremely efficient. The VQLS results reported here should therefore be interpreted as evidence that Kriging can be *embedded* into a VQLS formulation, not as a claim that VQLS is theoretically or practically superior to Kriging in this setting. A more rigorous analysis of condition numbers, block-encodings and asymptotic scaling would be required to argue for any asymptotic quantum advantage, which is beyond the scope of this simulation paper.

8. EXPERIMENTAL SETUP

8.1 Synthetic Scenario and Data Generation

A 4-hectare cannabis cultivation plot is represented as a 120×120 regular grid (approximately 8 m resolution). Environmental fields (soil moisture, canopy temperature, VPD, NDVI) are generated as Gaussian processes with exponential spatial covariance, parameterized to produce realistic spatial correlation lengths (tens of meters) and temporal autocorrelation. Elevation and slope fields are generated from a smoothed random surface consistent with the assumed topographic range.

A total of 36 virtual IoT sensors are positioned using stratified sampling over the field. Their time series are obtained by sampling the underlying Gaussian processes plus additive

noise. Synthetic Sentinel-2 and UAV reflectance data are created by applying simple forward models to the environmental fields (e.g. NDVI as a nonlinear transform of simulated reflectances), then downsampling or resampling to the desired spatial resolutions.

Stress labels (Normal, Early, Severe) are derived from thresholds on relative soil moisture, NDVI anomaly and VPD deviation around a cannabis-specific target range. The final dataset is mildly imbalanced (approximately 58% Normal, 27% Early, 15% Severe).

8.2 Quantum and Classical Models

Quantum algorithms are executed with Qiskit Aer simulators under ideal, noiseless conditions and, for a subset of experiments, with simple noise models:

- VQLS: 6 qubits, 3 layers, statevector backend, L-BFGS-B optimizer (noiseless only);
- VQC: 4 qubits, 4 variational layers, sampler backend, evaluated both noiseless and with a depolarizing noise model calibrated to typical single- and two-qubit error rates;
- QAOA: depth $P = 4$, 8 qubits, statevector backend, COBYLA optimizer (noiseless only).

Classical baselines are:

- Ordinary Kriging using PyKrige with exponential variogram model;
- Random Forest and a shallow neural network for stress classification;
- Mixed-Integer Linear Programming (MILP) for irrigation scheduling.

We acknowledge that these baselines do *not* exhaust the state of the art: advanced Gaussian Process models with more flexible kernels, deep learning architectures for spatio-temporal data (e.g. ConvLSTM, U-Net) and modern optimizers (CMA-ES, Bayesian optimization frameworks such as Optuna or Scikit-Optimize) could provide stronger classical performance. Evaluating quantum-inspired methods against such baselines is left as important future work.

A summary of key hyperparameters is provided in Table 2. All comparative results are based on repeated runs (50 Monte Carlo simulations) to estimate means and standard deviations. The focus is on relative behavior in this controlled setup, not on absolute performance in real fields.

8.3 Noise-aware Quantum Simulation (VQC)

To obtain a first indication of the impact of realistic noise on the VQC classifier, we run an additional experiment with a simple depolarizing noise model in Qiskit Aer. Single-qubit and two-qubit depolarizing channels are parameterized with error probabilities of $p_1 = 1 \times 10^{-3}$ and $p_2 = 5 \times 10^{-3}$, respectively, which are broadly representative of current NISQ devices. All other settings are kept identical to the noiseless VQC configuration.

8.4 Classical Pipeline Validation on Non-cannabis Greenhouse Data

To check that the classical part of the pipeline (data fusion and baseline models) behaves reasonably on real measurements, we perform a small, separate experiment on a greenhouse dataset for tomatoes. The dataset contains soil moisture and canopy temperature from 16 sensors, combined with simple vegetation indices derived from RGB imagery. We use only classical models (Kriging for interpolation and Random Forest for stress labelling), without any quantum modules. This experiment is not meant to approximate cannabis cultivation, but to validate that the basic classical workflow is applicable to real greenhouse data.

9. RESULTS (SIMULATED AND CLASSICAL)

We emphasize that all quantum-related results in this section are obtained on synthetic data and ideal or noisy quantum simulators. They should be interpreted as *illustrative* of possible behaviors in small-scale scenarios, not as definitive evidence of quantum advantage.

9.1 Spatial Interpolation (VQLS vs Classical Kriging)

In this particular synthetic setup, VQLS-based Kriging achieves a lower RMSE than classical Kriging (Table 3), corresponding to a relative reduction of about 20–25%. However, this comes at significantly higher computational cost on the simulator: the simulated runtime is almost one order of magnitude larger than that of classical Kriging, even at this small scale. A full cost–benefit analysis would need to consider wall-clock time, energy usage and access cost for real quantum hardware, which we do not attempt here. Given these facts, there is currently no practical justification for preferring VQLS over classical Kriging for problems of this size; the main value of this experiment is conceptual.

9.2 Stress Detection (VQC vs Classical ML)

In noiseless simulation, the VQC model yields a modest but consistent improvement in overall accuracy and F1-score, particularly for the Early Stress class (Table 4). This is encouraging, but remains a toy instance: it assumes low-dimensional feature encodings and ideal quantum operations. It also does not compare against deep spatio-temporal models (e.g. ConvLSTM, U-Net) that could potentially capture richer structure in real data.

9.3 Effect of Simulated Quantum Noise on VQC

When a simple depolarizing noise model is introduced (Table 5), VQC performance degrades by a few percentage points in accuracy and F1-score, especially for the Early Stress class. This is consistent with expectations for NISQ devices: shallow circuits with limited noise can still provide reasonable performance, but the gap between ideal and noisy operation is non-negligible. More realistic device-specific noise modelling and error-mitigation strategies would be needed to further bridge this gap.

9.4 Irrigation Optimization (QAOA vs MILP)

The QAOA-based controller achieves lower water use and fewer constraint violations than a simple rule-based strategy in this

Table 2: Hyperparameters of quantum and classical models used for interpolation, stress detection and irrigation optimization in the synthetic scenario.

Model	Qubits / Params	Depth / Layers	Optimizer	Backend / Solver	Task
VQLS	6 qubits	3 layers	L-BFGS-B	Qiskit Aer (statevector)	Kriging interpolation
VQC	4 qubits	4 variational layers	COBYLA	Qiskit Aer (sampler)	Stress classification
QAOA	8 qubits	$P = 4$	COBYLA	Qiskit Aer (statevector)	Irrigation optimization
Ordinary Kriging	–	–	–	PyKrige (exponential variogram)	Interpolation baseline
Random Forest	–	–	–	scikit-learn (300 trees)	Stress classification
Neural Network	–	2 hidden layers	Adam (10^{-3})	PyTorch (128–64 neurons)	Stress classification
MILP (Gurobi)	–	–	Branch & Cut	Gurobi Engine	Irrigation baseline

Table 3: Spatial interpolation performance (50 Monte Carlo runs, mean ± std) in the synthetic field.

Method	RMSE	MAE	Time (s)
IDW	0.071 ± 0.005	0.055 ± 0.004	2.3 ± 0.3
Ordinary Kriging	0.053 ± 0.003	0.041 ± 0.003	15.7 ± 2.1
Quantum VQLS	$\mathbf{0.041 \pm 0.002}$	$\mathbf{0.032 \pm 0.002}$	128.4 ± 18.6

Table 4: Multi-class (Normal, Early, Severe Stress) classification (mean ± std) on the synthetic dataset (noiseless simulation).

Model	Acc.	F1(N)	F1(E)	F1(S)
Random Forest	0.882 ± 0.012	0.91	0.79	0.88
Neural Network	0.894 ± 0.010	0.92	0.81	0.89
VQC (ideal)	$\mathbf{0.921 \pm 0.008}$	**0.94**	**0.86**	**0.92**

small synthetic example (Table 6), with performance comparable to the MILP baseline. Given the limited problem size (8 qubits) and ideal simulation, these results are best interpreted as evidence that QAOA can *encode* relevant irrigation objectives, rather than as proof of superiority over classical optimization methods, especially when more advanced solvers (e.g. nonlinear programming, CMA-ES, Bayesian optimization) are considered.

9.5 Classical Pipeline Validation on Tomato Greenhouse Data

On the tomato greenhouse dataset, the classical pipeline attains reasonable interpolation and stress-prediction performance (Table 7), with error magnitudes and classification metrics comparable to those reported in the literature for similar setups. This limited experiment does not directly validate the cannabis-specific scenario, but it shows that the classical part of the proposed architecture can operate on real greenhouse data without major modifications.

10. DISCUSSION

10.1 Interpretation of Simulation Results

Across all three tasks, quantum-inspired models show improved or comparable performance relative to classical baselines in the chosen synthetic configurations. This is consistent with the expressive power of VQAs when given enough circuit depth and properly tuned hyperparameters. However, these results are obtained:

Table 5: Impact of a simple depolarizing noise model on VQC performance (mean ± std over 30 runs).

VQC configuration	Accuracy	F1(Early)	F1(Severe)
Ideal (noiseless)	0.921 ± 0.008	0.86 ± 0.010	0.92 ± 0.007
Depolarizing noise	0.902 ± 0.011	0.82 ± 0.014	0.89 ± 0.010

Table 6: Irrigation scheduling over a short 7-day horizon (mean ± std) in a small synthetic scenario.

Strategy	Water	Dev. RMSE	Violations	Smooth.
Rule-based	154.2 ± 8.3	0.117 ± 0.015	12.4 ± 2.1	0.21
MILP	140.1 ± 6.2	0.095 ± 0.011	8.2 ± 1.3	0.18
QAOA	$\mathbf{129.8 \pm 5.7}$	$\mathbf{0.088 \pm 0.009}$	$\mathbf{6.8 \pm 1.1}$	**0.14**

- on a fully synthetic dataset with relatively simple stochastic structure,
- on small problem sizes (few qubits, limited number of zones),
- under ideal or mildly noisy quantum simulation,
- and against classical baselines that do not exhaust the current state of the art in spatial statistics or spatio-temporal deep learning.

As such, the improvements should be viewed as indicating that the proposed encodings are *capable* of representing useful solutions, not as evidence of practical advantage over well-engineered classical techniques in real fields.

10.2 Limitations

This work has several important limitations:

- **Synthetic data for cannabis.** All cannabis-related inputs and labels are generated from simplified stochastic models, calibrated only at the level of typical ranges and correlation patterns. Real cannabis fields exhibit more complex, non-stationary dynamics and management practices. Even though the tomato greenhouse experiment partially validates the classical pipeline, it does not remove the need for cannabis-specific field trials.
- **No real quantum hardware.** All quantum experiments use Qiskit Aer simulators. While we add a first noise-aware experiment for VQC, current NISQ devices suffer

Table 7: Classical interpolation and stress classification on a real tomato greenhouse dataset (mean $\pm$ std over 10-fold cross-validation).

Model	Metric	Value	Comment
Kriging (soil moisture)	RMSE	0.048 ± 0.006	Interpolation across 16 sensors
Random Forest (stress)	Accuracy	0.873 ± 0.015	Binary stress label (low/high)
Random Forest (stress)	F1(stress)	0.84 ± 0.020	Minority stressed class

from decoherence, gate errors and limited qubit counts, which would likely degrade performance and restrict problem size. More realistic device-specific noise modelling and error mitigation are still missing here.

- **Small-scale toy instances.** The numbers of qubits (4–8) and sensors (36) are small. Classical methods easily handle these sizes with excellent accuracy and speed; there is no computational bottleneck in this regime.

- **Limited classical baselines.** We restrict attention to ordinary Kriging, Random Forest, a shallow neural network and MILP. More advanced baselines, such as Gaussian Processes with adaptive kernels, deep spatial architectures (U-Net, ConvLSTM) and modern black-box optimizers (CMA-ES, Bayesian optimization via Optuna or Scikit-Optimize), are not considered.

- **Limited statistical analysis.** We report means and standard deviations over repeated runs, but do not provide bootstrap confidence intervals or multiple-testing corrections. A more rigorous statistical treatment would improve the robustness of the conclusions.

- **No explicit cost model.** Beyond reporting simulated runtimes, we do not quantify total cost in terms of wall-clock time, energy consumption and potential access costs to quantum hardware compared to classical high-performance computing.

- **Contextual constraints.** Due to regulatory, logistical and privacy constraints, no real data from licensed cannabis producers in Morocco are used. The link to the Moroccan context is therefore conceptual rather than empirical.

10.3 Future Directions

Future work should address these limitations by:

- **Hybrid validation on real crops.** Extending the classical and quantum pipelines to multiple real greenhouse datasets (tomatoes, peppers, etc.) with more diverse conditions, and progressively integrating cannabis data as it becomes accessible.

- **Extended classical benchmarks.** Comparing quantum-inspired methods against stronger baselines, including Gaussian Processes with advanced kernels, deep spatio-temporal models (U-Net, ConvLSTM, graph neural networks) and modern optimization frameworks such as CMA-ES, Optuna or Scikit-Optimize.

- **Noise-aware quantum simulations.** Incorporating realistic device-specific noise models in Qiskit Aer (e.g. from IBM or IonQ backends) and, when possible, executing small instances on real QPUs to quantify the gap between ideal, simulated-noisy and hardware performance.

- **Sensitivity and uncertainty analysis.** Performing systematic experiments varying noise levels, circuit depth, number of qubits and hyperparameters, computing bootstrap confidence intervals and robustness metrics to distinguish algorithmic effects from tuning artefacts.

- **Theoretical analysis.** Developing a more formal understanding of when VQLS-based solvers or QAOA formulations could theoretically offer asymptotic advantages over classical methods for spatial interpolation or irrigation optimization, leveraging tools from quantum linear algebra and optimization complexity.

- **Economic and deployment analysis.** Evaluating whether, under plausible future hardware and cloud-access scenarios, quantum-enhanced workflows could be economically justified for high-value regulated crops, taking into account hardware, energy and integration costs.

11. CONCLUSION

We have proposed a simulation-based hybrid quantum–classical smart farming architecture for regulated medical cannabis in the Al-Hoceïma region. Using synthetic IoT, UAV and Sentinel-2 data for a virtual 4-hectare field, we formulated Kriging, stress detection and irrigation scheduling as VQLS, VQC and QAOA problems, respectively, and compared their performance to classical baselines under ideal and noisy quantum simulation. A small tomato greenhouse experiment further validated the classical part of the pipeline on real data.

The numerical results show that, in small synthetic scenarios, quantum-inspired modules can match or moderately outperform standard classical methods on key metrics such as interpolation error, early-stress F1-score and irrigation-water usage. These findings do not establish quantum advantage, nor do they imply that VQLS or QAOA should be preferred to well-tuned classical solvers at the scales considered; rather, they demonstrate that agricultural decision-support tasks can be expressed in quantum-compatible forms and achieve reasonable performance under idealized conditions.

We argue that, given the strong regulatory and quality constraints of medical cannabis, such quantum-enhanced workflows are worth investigating now at the conceptual level, so that they can be deployed and evaluated empirically once quantum hardware, agricultural data infrastructures and classical baselines have been jointly considered in a more comprehensive way.

ACKNOWLEDGMENTS

This research was supported by the Moroccan National Center for Scientific and Technical Research (CNRST) under grant AgriQ-2024/07. The authors thank the Al-Hoceïma Provincial Directorate of Agriculture for providing general information on regional climate and agronomic conditions.

References

Bernstein, N., Gorelick, J., Koch, S., 2019. Cannabis horticulture: Perspectives and applications. S. Chandra, H. Lata, M. El-Sohly (eds), *Cannabis sativa L. - Botany and Biotechnology*, Springer, Cham.

Cerezo, M., Arrasmith, A., Babbush, R. et al., 2021. Variational quantum algorithms. *Nature Reviews Physics*, 3(9), 625–644.

Grand View Research, 2024. Medical cannabis market size, share & trends analysis report 2024–2030. Industry Market Research Report. Grand View Research.

Kingdom of Morocco, 2021. Law 13-21 on the lawful uses of cannabis. Official Bulletin No. 7000. In French: Loi 13-21 relative aux usages licites du cannabis.

Pook, T. et al., 2025. Assessing the potential of quantum computing in agricultural optimization. *Nature Food*, 6(2), 145–158.