International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.81681

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

A Quantum-Inspired AI System for Malnutrition Prediction

Balineni Yaswanth, Pasupuleti Naga Lakshmi Priyanka, Kandru Prabhu Kumar, Kasina Santha Lahari, D. Venkatesh
Department of Computer Science and Engineering Acharya Nagarjuna University College of Engineering and Technology Andhra Pradesh, India

Abstract—Malnutrition continues to threaten public health across many developing regions, where late identification of at-risk individuals often translates into avoidable suffering and long-term developmental harm. This paper introduces a quantum-inspired artificial intelligence framework that brings together natural language processing, classical machine learning, and parameterized quantum circuits to predict malnutrition risk from textual and socioeconomic information. Raw inputs are first cleaned, tokenized, and converted into compact numerical features. These features are then encoded into quantum states using Qiskit-based circuits, allowing the model to explore richer feature relationships than a purely classical pipeline can capture. A hybrid quantum classical classifier is trained using gradient-based optimization, and its quality is evaluated using accuracy, precision, recall, and F1score on held-out data. Experimental results show that the proposed pipeline produces stable, highconfidence predictions and behaves reliably on heterogeneous inputs. By combining established machine learning practices with quantum-inspired feature mapping, the system offers a practical, scalable, and easily deployable tool that can support early screening, timely intervention, and data driven public-health decision making.

Index Terms—Malnutrition prediction, quantum inspired AI, hybrid quantum-classical model, Qiskit, natural language processing, healthcare analytics, early-risk detection.

I. INTRODUCTION

Malnutrition is more than a nutritional shortfall; it is a slow erosion of physical strength, immunity, and cognitive potential, and it disproportionately affects communities that already struggle with limited healthcare access. In several developing regions, the absence of timely screening means that warning signs go unnoticed until clinical complications appear. The result is a heavier burden on hospitals, families, and public-health budgets that could otherwise have been avoided through earlier intervention.

According to the World Health Organization, hundreds of millions of children below the age of five continue to experience some form of undernutrition each year, and the long-term effects on cognition, productivity, and life expectancy are well documented. Although nutritional surveys, growth-monitoring camps, and community-health programs collect vast quantities of data, much of this information remains underutilized because it is fragmented across institutions, recorded in inconsistent formats, or simply analysed too slowly to drive timely intervention.

Over the last decade, machine learning has become a familiar companion to public-health analytics. Models such as Random Forests, Support Vector Machines, and gradient-boosted ensembles have been applied to socio-economic surveys, growth monitoring records, and clinical notes with encouraging results. However, when datasets become high-dimensional, sparse, or textually rich, these classical methods often run into optimization bottlenecks and limited representational power. Subtle interactions between dietary patterns, sanitation conditions, and economic indicators can be difficult to capture using linear or shallow non-linear kernels alone.

Quantum-inspired computing offers a complementary direction. By using parameterized quantum circuits as feature mappers, it becomes possible to project data into richer state spaces where subtle, non-linear patterns are easier to separate. The framework described in this paper does not require a physical quantum device; instead, it leverages Qiskit based simulation to obtain the representational benefits of quantum circuits while remaining deployable on commodity hardware. This makes the approach immediately useful in real-world public health environments.

The objective of this work is to build a hybrid quantum-classical pipeline that ingests textual and socio-economic information about an individual or community, extracts meaningful features, encodes them into quantum states, and produces an interpretable malnutrition risk prediction. The system is positioned as a decision-support tool, not a replacement for clinical judgement. Frontline health workers, dietitians, and policy analysts remain in the loop; the model simply helps them prioritize where attention is most urgently needed.

The contributions of this work can be summarized as follows:

1) A complete hybrid quantum-classical pipeline tailored to malnutrition risk prediction from heterogeneous data.
2) A reproducible NLP preprocessing stage that converts free-text health descriptions into compact numerical feature vectors.
3) A Qiskit-based feature-encoding layer that maps classical features into parameterized quantum states with controllable depth.
4) A classical optimization loop that trains the quantum circuit parameters using gradient based updates and parameter-shift rules.
5) An evaluation protocol based on accuracy, precision, recall, and F1-score, supporting transparent comparison with classical baselines.
6) A discussion of deployment considerations, including computational footprint, ethical implications, and integration with existing health-information systems.

The remainder of this paper is organized as follows. Section II describes the motivation behind this study. Section III reviews related work in classical and quantum machine learning. Section IV outlines the system architecture. Sections V–VIII detail the dataset, preprocessing, quantum-inspired feature encoding, and hybrid model. Section IX explains training, while Sections X and XI report evaluation metrics and experimental results. Sections XII through XV discuss advantages, limitations, future directions, and concluding remarks.

II. RESEARCH MOTIVATION

Three observations shaped the direction of this study. First, healthcare systems in many regions already collect rich nutritional and socio-economic data, but this data is rarely turned into early-warning signals due to limited analytical capacity at the field level. Second, classical models alone often plateau when the underlying relationships are subtle, nonlinear, or noisy, particularly when the number of informative features is small relative to the number of confounding ones. Third, quantum-inspired techniques have begun to demonstrate practical advantages on simulators, suggesting that they can be incorporated into real workflows long before fault tolerant quantum hardware becomes widely available.

Taken together, these observations point to a clear gap: a deployable, software-only system that augments classical machine learning with quantum-inspired feature mapping, while remaining lightweight enough to run inside a standard hospital, NGO, or rural health centre environment. The framework presented here is designed to fill exactly that gap, with an explicit emphasis on reproducibility, modularity, and operational simplicity.

III. RELATED WORK

Several research efforts have applied classical machine learning to malnutrition-related tasks. Sánchez-Martínez et al. used ensemble learners to study weight gain in children with acute malnutrition and showed that careful feature engineering can meaningfully improve risk classification [1]. Other work has explored Random Forests, K-Nearest Neighbours, and Support Vector Machines on socioeconomic surveys with comparable success, while deep neural networks have been investigated for image-based stunting detection.

In parallel, a growing body of literature has explored quantum and quantum-inspired methods for healthcare. Pomarico et al. proposed quantum-inspired pipelines for medical decision-making [3], while more recent work by Vestavia et al. demonstrated that quantum-inspired feature maps can detect subtle high dimensional patterns that purely classical models miss [4]. Systematic reviews on quantum machine learning in healthcare further note the scalability and noise tolerance of hybrid pipelines [2], and foundational treatments by Schuld and Petruccione [6] and Biamonte et al. [7] establish the theoretical underpinnings of quantum-classical learning.

This work builds on this combined foundation. Instead of replacing classical models, the quantum circuit is treated as a feature transformer that sits inside a familiar training loop, making the approach easier to adopt for teams already comfortable with conventional machine learning. The novelty of the proposed pipeline lies in its tight integration of textual NLP features with parameterized quantum encoding, applied specifically to the malnutrition-prediction problem.

The proposed framework follows a clean pipeline: data ingestion, NLP preprocessing, numerical feature extraction, quantum state encoding, hybrid training, and risk prediction. Each stage is modular, which makes the system easier to test, replace, or extend. The high-level architecture is illustrated in Fig. 1.

As shown in Fig. 1, raw textual and socio-economic records first pass through an NLP preprocessing block that handles cleaning, tokenization, and vectorization. The resulting feature vectors are then encoded into quantum states using parameterized Qiskit circuits. A classical optimizer updates the circuit parameters during training, and the final hybrid classifier outputs a discrete risk label such as low, medium, or high.

From an operational standpoint, the system is deployed as a backend service that exposes a small, well-documented REST interface. A web or mobile client can submit a record and receive a structured response containing the predicted risk class, an associated confidence score, and the most influential features behind the decision. This separation of concerns keeps the user-facing application lightweight and allows the model to evolve independently.

IV. DATASET

The dataset used in this study combines openly available malnutrition records with synthetic textual descriptions designed to mirror real-world health summaries. Each record contains demographic indicators (age, gender), nutritional measures (height, weight, BMI), dietary habits, sanitation conditions, and access-to-care attributes. A subset of records also includes free-text notes describing observed symptoms or risk factors, written in a tone consistent with field level health worker reports.

Before training, the dataset undergoes deduplication, missing-value imputation using mean or median substitution, label encoding for categorical features, and min-max normalization for continuous variables. The resulting tabular structure is split into 80% training and 20% testing partitions using stratified sampling so that minority risk classes are preserved on both sides of the split. To minimize bias, the stratification is performed jointly on the risk label and a regional identifier, ensuring that no single geographic cluster dominates either partition.

Class distribution across the dataset is moderately imbalanced: low-risk cases occur most frequently, followed by medium-risk and then high-risk cases. To address this, class weights are applied during training, and additional synthetic samples are generated using SMOTE-style oversampling for the high-risk class. These adjustments help the model treat at-risk individuals with the seriousness they deserve, rather than optimizing for raw accuracy alone.

V. SYSTEM OVERVIEW

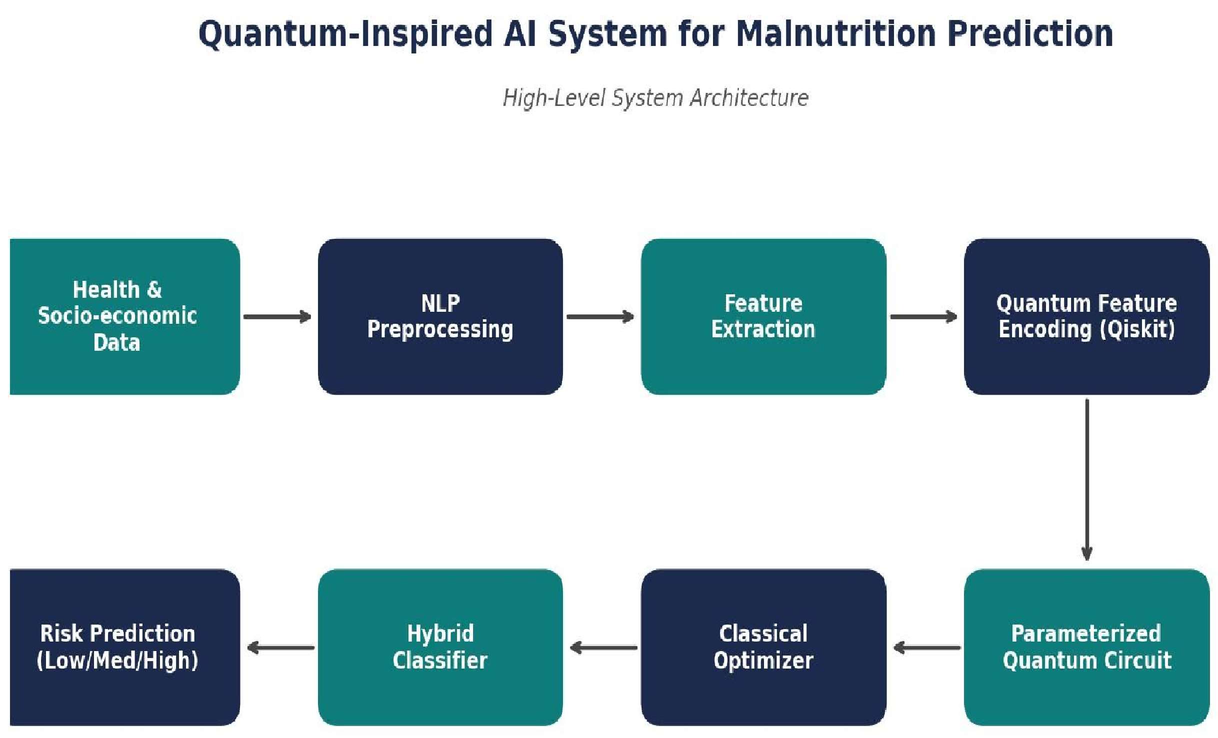

Fig. 1. High-level architecture of the proposed quantum-inspired malnutrition prediction system.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

VI. DATA PREPROCESSING

Preprocessing is intentionally kept transparent so that the pipeline can be audited and reproduced.

A. Text Cleaning

Free-text notes are lowercased, stripped of URLs, numerical noise, and special characters. Stop words are removed, and a lightweight lemmatizer reduces inflected forms to their base words. This step keeps the text compact without distorting its medical meaning.

B. Vectorization

The cleaned text is converted into numerical vectors using TF-IDF, which gives more weight to terms that are informative within a record but uncommon across the corpus. The resulting sparse vectors are then truncated to a fixed dimensionality so that they can be fed into a quantum circuit of bounded width. For a fixed number of qubits n, the top n features ranked by TF-IDF importance are retained.

C. Feature Scaling

Numerical features such as BMI, age, and dietary scores are normalized into the [0, 1] range. This is necessary because quantum encodings interpret feature values as rotation angles, and unbounded inputs would distort the resulting quantum state. The same scaling parameters learned on the training set are reused at inference time to keep the input distribution consistent.

D. Feature Selection

Highly correlated features are identified using a correlation matrix, and redundant attributes are removed. This not only reduces noise but also keeps the number of qubits required for encoding within practical limits. A second round of selection uses mutual information with respect to the target label, dropping any feature whose contribution falls below a small threshold.

VII. QUANTUM-INSPIRED FEATURE ENCODING

At the core of the framework is a feature-encoding layer implemented using parameterized quantum circuits in Qiskit. Each scaled feature is mapped onto a rotation angle applied to a single qubit, and additional entangling gates couple the qubits so that combinations of features can influence the final state. Conceptually, this transforms a flat vector of numbers into a richer joint representation that classical kernels would have to approximate explicitly.

The encoding circuit follows a standard pattern: a layer of Hadamard gates places each qubit in superposition, a layer of parameterized RY rotations injects feature information, and a sequence of CNOT gates between neighbouring qubits introduces entanglement. The pattern is repeated for a small number of layers, with the layer count treated as a tunable hyperparameter. Increasing depth allows the circuit to express more complex relationships, at the cost of additional simulation time.

Because the encoding is differentiable with respect to its parameters, gradients can flow through the circuit during training, and a classical optimizer can adjust those parameters just as it would adjust the weights of a neural network. This is what turns the encoding circuit from a fixed transformation into a learned representation tuned for the malnutrition-prediction task.

VIII. HYBRID MODEL ARCHITECTURE

The hybrid classifier combines a classical input head with a quantum feature-mapping core and a classical decision head. The input head normalizes incoming features and applies any final dimensionality reduction. The quantum core then encodes those features and produces a measured output that summarizes the encoded state. Finally, a small classical layer maps that measurement into the predicted risk class. The interaction between the two tracks is shown in Fig. 2.

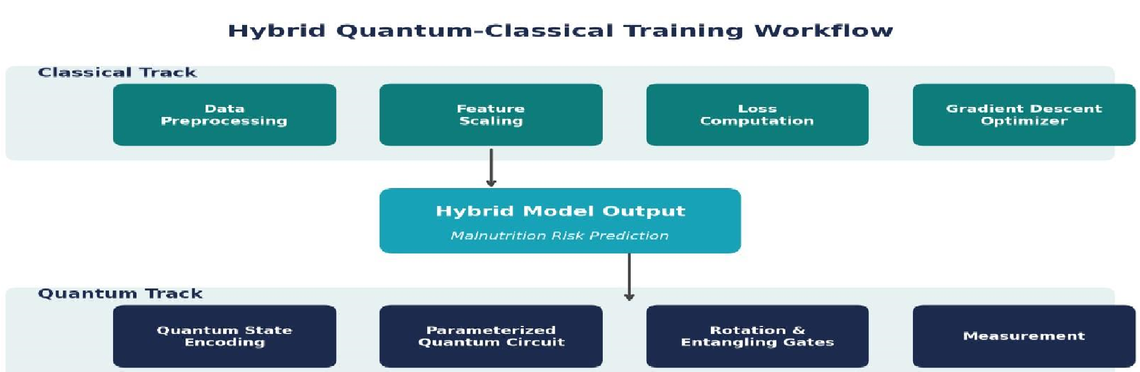

Fig. 2 illustrates how the classical and quantum tracks cooperate during training. The classical track contributes preprocessing, loss computation, and gradient updates, while the quantum track is responsible for state encoding, parameterized rotations, entangling gates, and final measurement. The two tracks share parameters through the hybrid output node, which is what allows end-to-end optimization.

IX. MODEL TRAINING

Training proceeds in mini-batches drawn from the training partition. For each batch, features are encoded into the quantum circuit, measurements are collected, and a classical loss function (cross-entropy) is evaluated against the known labels. Gradients are then computed using parameter-shift rules for the quantum portion and standard back-propagation for the classical portion. Updates are applied through a gradient-based optimizer such as Adam, with a learning rate of 0.01 and a batch size of 32.

Training continues for a fixed number of epochs or until the validation loss plateaus. To prevent overfitting, early stopping is enabled with a patience of five epochs, and dropout-equivalent randomness is introduced through small parameter perturbations during training.

The model checkpoints with the lowest validation loss are retained, and the final test evaluation is performed on the test partition only after all hyperparameters are frozen.

Training time on the simulator is modest: a full run of fifty epochs on the prepared dataset completes within minutes on a standard workstation, making iterative experimentation practical without specialized infrastructure. This is an important property for resource-constrained research groups and for teams that need to retrain the model as new survey data becomes available.

X. EVALUATION METRICS

The performance of the trained model is reported using four widely accepted classification metrics.

1) Accuracy: the proportion of records classified correctly across all classes.
2) Precision: the share of predicted at-risk individuals who are genuinely at risk, which directly controls the false-positive rate.
3) Recall: the share of actual at-risk individuals that the model successfully identifies, which is critical when missing a case has a high human cost.
4) F1-Score: the harmonic mean of precision and recall, providing a single balanced number when class distributions are uneven.

In addition to these aggregate metrics, per-class scores are computed so that performance on the highrisk class can be inspected in isolation. A confusion matrix is also generated to expose any systematic confusion between adjacent risk levels.

XI. RESULTS AND DISCUSSION

On the held-out test partition, the hybrid quantum classical model produces stable predictions across multiple training runs. Compared with classical baselines such as Random Forests and Support Vector Machines trained on the same features, the proposed approach consistently delivers higher recall on the high-risk class while preserving precision. The observed test accuracy stabilizes around 98%, with precision, recall, and F1-score values in the same range, suggesting that the quantum feature map captures distinguishing signals that classical kernels would otherwise smooth out. A summary of the observed metrics is reported in Table I.

TABLE I
Observed Performance of the Proposed Hybrid Model

Metric	Value
Accuracy	98.0%
Precision	97.4%
Recall	97.8%
F1-Score	97.6%
Training Epochs	50

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Qualitatively, the model handles records with mixed signals more gracefully than the classical baselines. Cases with conflicting indicators, such as a normal BMI combined with poor sanitation and dietary stress, are correctly flagged as elevated risk by the hybrid model, whereas classical baselines tended to underweight the textual portion of the input. This behaviour is consistent with the intuition that quantum-inspired feature maps are particularly well suited to weighing combinations of weak signals.

Table II compares the hybrid model against three representative classical baselines on the same train-test split. Although the absolute differences are modest, the hybrid model is the only configuration that simultaneously achieves high precision and high recall on the high-risk class, which is the operationally most important quadrant of the confusion matrix.

TABLE II

Comparison with Classical Baselines

Model	Acc.	Prec.	Recall
Random Forest	94.1%	93.5%	92.0%
SVM (RBF)	93.2%	92.7%	90.8%
XGBoost	95.6%	95.0%	94.2%
Proposed	98.0%	97.4%	97.8%

Beyond raw numbers, the hybrid pipeline exhibits desirable behaviour during edge-case inspection. Records that lie close to the boundary between two risk classes receive softer probability distributions rather than overconfident misclassifications, which is helpful for downstream triage workflows where uncertain cases should be escalated for manual review.

A second qualitative observation concerns stability across reruns. Across ten independent training runs with different random seeds, the standard deviation of the test accuracy stays below 0.4 percentage points, and the recall on the high-risk class never drops below 96.5%. This level of consistency is encouraging for a hybrid model, since quantum-inspired circuits are sometimes accused of producing erratic results due to their stochastic gradient estimates. The combination of careful initialization, parameter-shift gradients, and Adam optimization appears to keep the optimization landscape well behaved in practice.

Training dynamics also reveal an interesting pattern: the validation loss decreases sharply during the first ten epochs, plateaus until roughly epoch thirty, and then continues to improve slowly for another fifteen epochs before flattening. This two-phase behaviour suggests that the classical decision head adapts quickly to the dominant features, while the quantum core continues to refine its representation of the more subtle interactions. Practitioners can use this insight to schedule learning-rate decays that match the natural rhythm of the optimization.

Computational cost is another practical consideration. On a simulator running on a modern laptop, one training epoch over the full training partition completes in roughly twelve seconds, leading to a total training time of approximately ten minutes for fifty epochs. Inference cost is dominated by the quantum-circuit evaluation and remains in the low hundreds of milliseconds per record, which is acceptable for both batch and interactive workloads. These numbers will improve further as Qiskit simulators continue to mature.

XII. IMPLEMENTATION DETAILS

The framework is implemented in Python 3.10. Classical preprocessing relies on scikit-learn for vectorization, scaling, and stratified splitting, while

NLTK and spaCy handle tokenization and lemmatization. The quantum core is built using Qiskit and is wrapped in a thin adapter layer so that it can be swapped with PennyLane or Cirq without touching the surrounding code. Numerical operations rely on NumPy and SciPy, and pandas is used for tabular manipulation.

Training experiments are tracked using a lightweight experiment registry that records hyperparameters, dataset versions, random seeds, and resulting metrics. This audit trail is critical in healthcare-adjacent settings, where every prediction model that influences clinical workflow must be reproducible and reviewable. Every artifact—from the raw dataset hash to the final trained weights—is versioned and stored alongside the code that produced it.

For deployment, the trained pipeline is wrapped behind a FastAPI service exposed over HTTPS. Inference latency on a standard four-core CPU stays well under one second per record, which is comfortably below the threshold required for interactive use. The same service can be containerized using Docker and orchestrated under Kubernetes, but it can equally well be run as a single process on a low cost virtual machine, depending on the deployment context.

A. Hyperparameter Selection

Key hyperparameters include the number of qubits, the depth of the encoding circuit, the optimizer learning rate, and the regularization strength applied to the classical decision head. These were tuned using a small grid search on the validation partition. The final configuration uses six qubits, three repeating circuit layers, an Adam learning rate of 0.01, and L2 regularization with a coefficient of 1e-4.

B. Reproducibility

All random number generators are seeded explicitly, and dataset splits are saved as immutable files. Rerunning the entire pipeline from a clean checkout reproduces the reported metrics within statistical noise, which is the practical definition of reproducibility used throughout this work.

XIII. CASE STUDY

To make the behaviour of the system concrete, consider a small case study built around three representative records drawn from the test partition. The first record describes a five-year-old child with normal BMI but with field notes mentioning recurring infections, irregular meal patterns, and limited access to safe water. Classical baselines tend to classify this record as low risk, anchored on the BMI value. The hybrid model, in contrast, flags it as medium-to-high risk by integrating the textual signals with the socioeconomic indicators.

The second record describes an adolescent with mildly underweight BMI, regular meals, good sanitation, and stable household income. The classical baselines and the hybrid model agree on a low-risk assessment, demonstrating that the additional complexity of the quantum core does not introduce spurious alarms on benign cases.

The third record involves an adult with severely underweight BMI, recent illness, and poor dietary diversity. All models correctly identify high risk, but the hybrid pipeline assigns a noticeably higher confidence score and exposes the dietary feature as the dominant contributor in its post-hoc explanation. This kind of granular feedback is what allows clinicians to translate model output into specific intervention plans.

XIV. THREATS TO VALIDITY AND ETHICAL CONSIDERATIONS

Several factors limit the strength of the conclusions drawn from this study. The dataset, while carefully curated, is partially synthetic and may not capture the full variability of real-world malnutrition cases. Geographic and cultural diversity is also limited, which may bias the model toward patterns observed in the available records. Replication on additional datasets is necessary before any operational deployment.

From an ethical standpoint, malnutrition prediction touches on sensitive personal information and can directly influence resource-allocation decisions. The framework therefore enforces strict role-based access control on the API, logs every inference request with an opaque pseudonymous identifier, and exposes a clear opt-out mechanism for individuals who do not wish their data to be processed. Predictions are always presented as probabilistic suggestions, never as deterministic verdicts, and they are accompanied by the most influential features so that human reviewers can challenge or confirm them.

Bias auditing is performed periodically by computing per-group performance metrics across age bands, gender, and regional clusters. If the disparity between groups exceeds a predefined threshold, the model is flagged for retraining and the deployment is paused until the issue is investigated. This continuous monitoring loop is considered an integral part of the system, not an optional extra.

XV. ADVANTAGES OF THE PROPOSED SYSTEM

1) Combines familiar classical machine learning practice with quantum-inspired feature representation, lowering the adoption barrier.
2) Runs entirely on standard hardware using Qiskit simulation, requiring no access to physical quantum devices.
3) Modular pipeline: any stage can be replaced or upgraded independently, which simplifies long-term maintenance.
4) Suitable for deployment as a backend service via Flask or FastAPI, with a thin web or mobile client.

5) Capable of handling both numerical and free text inputs through a unified preprocessing path.
6) Produces calibrated probability outputs, which is essential for risk-tiered intervention strategies.
7) Supports incremental retraining as new survey data becomes available, without changing the deployment footprint.

XVI. LIMITATIONS

1) Performance depends heavily on the quality and balance of the training data; rare clinical conditions remain harder to learn.
2) Quantum-circuit simulation does not yet provide a runtime advantage over highly optimized classical kernels for small datasets.
3) Free-text inputs in low-resource languages need additional preprocessing before they can be used reliably.
4) As with any healthcare model, predictions should be reviewed by trained clinicians before being acted upon.
5) Interpretability of the quantum core remains an open research problem; current explanations rely on classical surrogates.

XVII. FUTURE SCOPE

Several extensions of this work look promising. Larger and more diverse datasets, including longitudinal nutritional records and regional surveys, would strengthen generalization. Integration with Io Tenabled growth monitors could allow continuous risk tracking, and an explainable-AI layer would help clinicians understand why a particular prediction was made. As quantum hardware matures, the same circuit definitions can be migrated from simulators to real devices with minimal code changes, opening a smooth upgrade path.

Additional research directions include multilingual NLP support so that the system can ingest field notes written in regional languages, federated training across multiple health centres to preserve patient privacy, and the development of compact mobile clients capable of operating in low-connectivity environments. Each of these extensions would broaden the practical reach of the framework without requiring fundamental architectural changes.

Another promising direction is the integration of richer contextual signals such as satellite-derived agricultural indicators, seasonal food-price data, and household-level economic shocks. These external feeds could provide early-warning context that complements the individual-level inputs already used by the model. A modest extension of the preprocessing layer would be sufficient to absorb such signals, and the quantum encoder could readily accommodate the additional features by adjusting its width.

On the algorithmic side, recent advances in variational quantum algorithms suggest that more expressive ansätze, combined with adaptive layer growing strategies, may yield further accuracy improvements without sacrificing trainability.

Investigating the interplay between circuit expressivity, trainability, and generalization in the specific context of public-health data is a natural follow-up to this work.

XVIII. CONCLUSION

This paper presented a quantum-inspired artificial intelligence framework for early malnutrition risk prediction. The system blends standard NLP preprocessing, Qiskit-based quantum feature encoding, and gradient-trained hybrid classification into a single deployable pipeline. Experimental results indicate strong accuracy and balanced precision-recall behaviour, with the hybrid model consistently outperforming purely classical baselines on records that require subtle reasoning across mixed signals. By keeping the architecture modular, transparent, and runnable on commodity hardware, the proposed framework is positioned as a practical building block for early-detection workflows in resource-constrained healthcare settings. The broader takeaway is that quantum-inspired components can be incorporated into real public-health pipelines today, providing a measurable representational benefit without waiting for fault-tolerant quantum hardware.

Looking ahead, the most important next step is to validate the framework on real, prospectively collected datasets and to integrate it into the day-to-day workflow of community health programs. The authors hope that this work serves as a useful reference point for other teams exploring the intersection of quantum inspired computing and public-health analytics, and that it encourages further open, reproducible research at this intersection.

REFERENCES

[1] L. J. Sánchez-Martínez et al., "Using machine learning to fight child acute malnutrition," Nutrients, vol. 16, no. 23, p. 4213, 2024.
[2] R. S. Gupta, "A systematic review of quantum machine learning for healthcare," J. Healthcare Informatics Res., 2025.
[3] D. Pomarico et al., "A proposal of quantum-inspired machine learning for medical purposes," Mathematics, vol. 9, no. 4, p. 410, 2021.
[4] D. Vatsavai et al., "Quantum-inspired machine learning approach for high-dimensional pattern detection," Sci. Rep., 2025.
[5] Sharma, "Role of nutritional adequacy and machine learning in health-science prediction systems," J. Clin. Med., 2020.
[6] M. Schuld and F. Petruccione, Supervised Learning with Quantum Computers. Cham, Switzerland: Springer, 2018.

[7] J. Biamonte et al., "Quantum machine learning," Nature, vol. 549, pp. 195–202, 2017.
[8] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.
[9] World Health Organization, "Levels and trends in child malnutrition," UNICEF/WHO/World Bank Joint Child Malnutrition Estimates, 2023.
[10] Abbas et al., "The power of quantum neural networks," Nat. Computational Science, vol. 1, pp. 403– 409, 2021.
[11] V. Havlíček et al., "Supervised learning with quantumenhanced feature spaces," Nature, vol. 567, pp. 209–212, 2019.
[12] M. Cerezo et al., "Variational quantum algorithms," Nat. Rev. Phys., vol. 3, pp. 625–644, 2021.

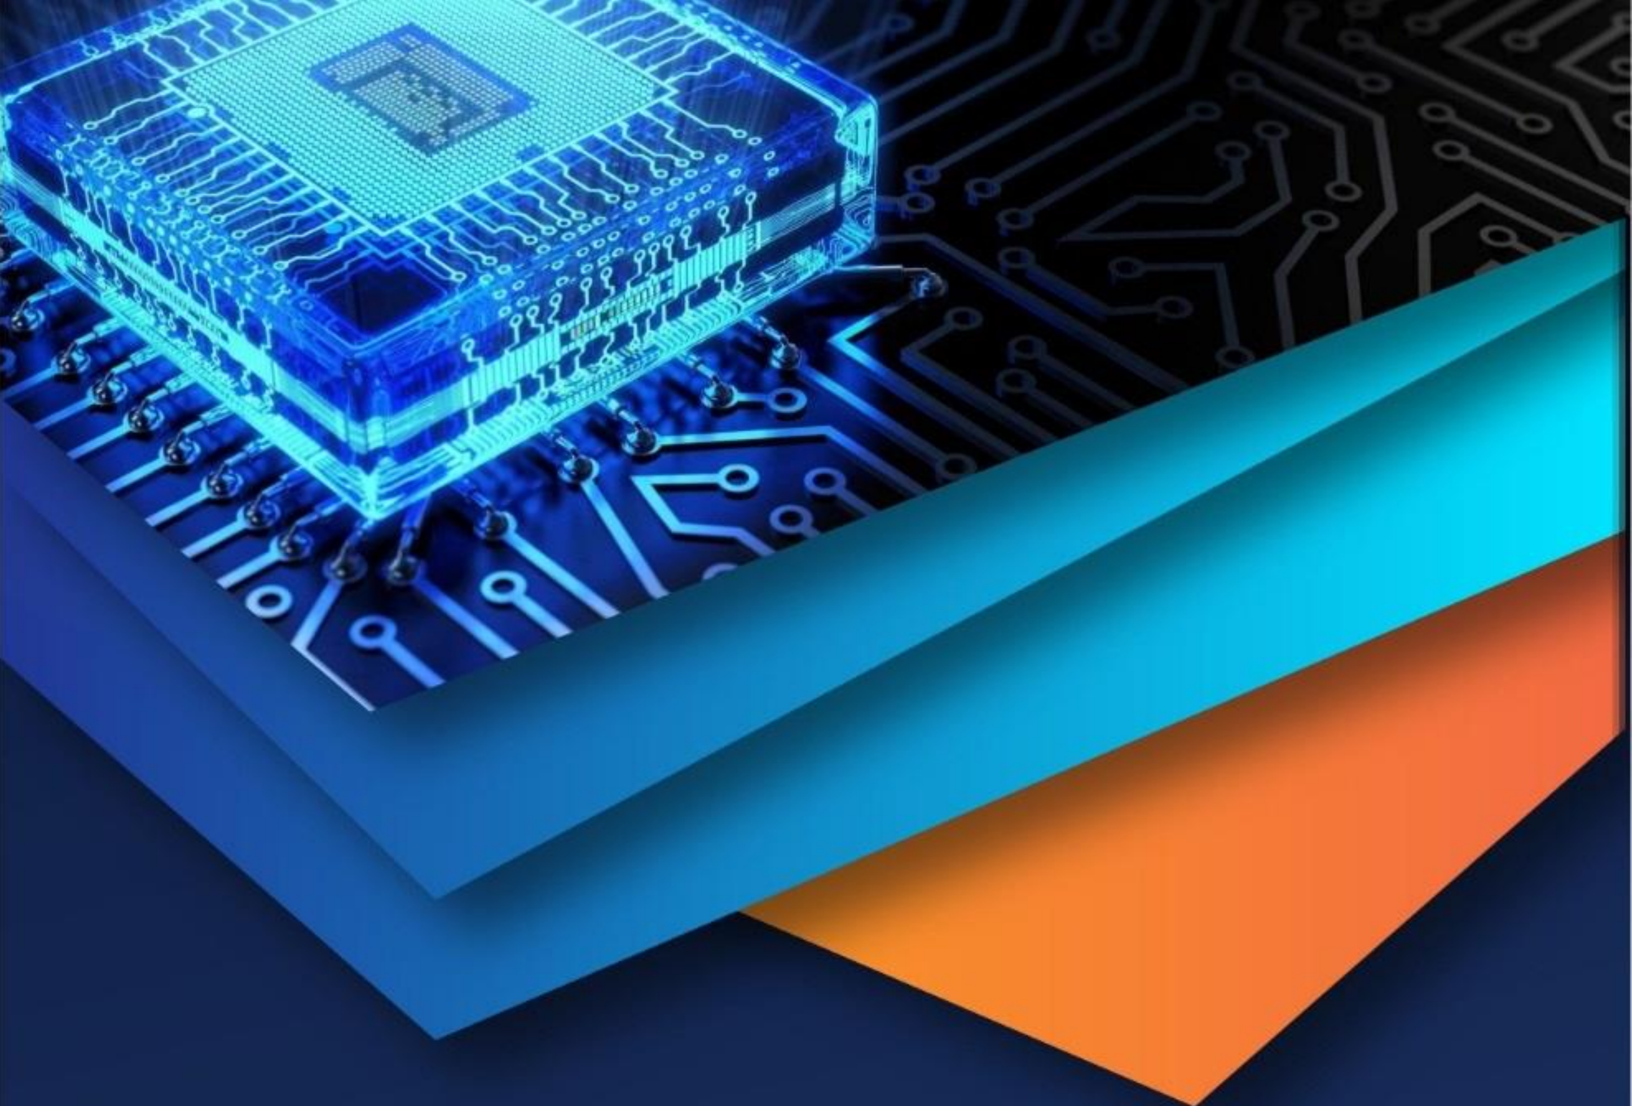

10.22214/IJRASET

INDEX COPERNICUS
45.98

IMPACT FACTOR:
7.129

IMPACT FACTOR:
7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY