International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

A HYBRID CNN–LSTM-BASED PREDICTIVE MODEL FOR NUTRITIONAL DEFICIENCY (IN WOMEN'S) DETECTION USING QUANTITATIVE MODELLING AND DEEP FEATURE LEARNING ON CLINICAL TIME-SERIES DATA

[1]Shraddha Satish Kashid, [2]Dr. Ramesh Mali

[1]Research Scholar, Department of Electronics & communication Engineering, MIT School of Engineering & science MIT Art, Design & Technology University, Pune.

[2]Research Supervisor Associate Professor, Department of Computer Science & Engineering MIT School of Computing, MIT Art, Design & Technology University, Pune.

shraddha.kashid@mituniversity.edu.in, ramesh.mali@mituniversity.edu.in

Abstract

Women who aren't getting enough nutrition is still a public health problem, especially in places with few resources where early diagnosis and treatment are needed to avoid long-term health problems. Using real-life clinical datasets, this study shows a deep learning-based diagnostic framework for automatically finding nutrient deficiencies. Specifically, it looks at Vitamin A and Iron deficiencies in women. A CNN, LSTM network, and a hybrid (CNN–LSTM) model will be designed and tested as part of the proposed method. The dataset went through a lot of preprocessing, feature selection, and transformation to make it work with deep learning architectures. We saw that the CNN–LSTM model was more accurate at classifying things than other models like SVM, KNN, and Naïve Bayes. It got 93.2% for Iron Deficiency and 92.5% for Vitamin A Deficiency. The hybrid model also did a better job of generalising and overfitting, as shown by the training-validation loss curves. The suggested deep learning framework provides a scalable, non-invasive, and highly accurate way to find nutritional deficiencies early on. It can be used in real-time diagnostic systems or built into platforms for monitoring public health.

Keywords: Nutrient Deficiency Detection, Vitamin A, Iron Deficiency, Deep Learning, CNN, LSTM, CNN–LSTM, Medical Diagnosis, Classification, Public Health Informatics.

1. Introduction

1.1 Background and Motivation

Micronutrient deficiencies are still one of the most underreported but most important public health problems in the world. They affect billions of people and make illness, death, and

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

economic problems much worse. Iron Deficiency Anaemia (IDA) and Vitamin A Deficiency (VAD) are two of the most common nutritional disorders in the world. Women and children are more likely to have these conditions than men [1]. The World Health Organisation (WHO) says that more than 1.62 billion people have anaemia. Women of childbearing age are most likely to have it, especially in low- and middle-income countries [2]. In the same way, nearly 190 million preschool-aged children and 19 million pregnant women are deficient in vitamin A, which can lead to blindness, weakened immune systems, and a higher risk of getting infectious diseases [3].

Women are more likely to get sick because of physiological changes that happen during periods, pregnancy, and breastfeeding that require them to eat more. Poor access to health care, poverty, unequal food distribution based on gender, and a lack of awareness are all social and cultural factors that make this problem even worse [4]. Even though the effects of these kinds of deficiencies are well known to include memory loss, less work output, bad pregnancies, and higher rates of maternal death, early diagnosis and monitoring are still very hard to do in many places. Standard methods of diagnosis often rely on invasive and expensive lab tests like measuring serum ferritin and retinol-binding protein. These tests can't be used for mass screening, especially in places that are rural or don't have a lot of resources [5].

The use of Artificial Intelligence (AI) and Data-Driven Decision Support Systems in biomedical informatics has the potential to completely change things in response to the growing health crisis. Advanced computer models, especially Machine Learning (ML) and Deep Learning (DL) methods [6], can look at a lot of different kinds of data to find hidden patterns, make accurate predictions, and allow doctors to make decisions in real time. But it's hard to make models that are reliable, easy to understand, and scalable so that they can not only make accurate predictions but also work across different population datasets.

1.2 Clinical Impact of Nutrient Deficiencies in Women

From a public and clinical health point of view, not getting enough iron and vitamin A has huge effects. IDA causes tiredness, poor brain function, a weakened immune system, and problems during childbirth. VAD, on the other hand, causes night blindness, keratomalacia, and a higher risk of death in both mothers and babies [7]. These deficiencies are not separate problems; they often happen together, making health risks worse and making people more vulnerable. For example, a woman who is low on both iron and vitamin A may have even worse effects on her immune system and the production of haemoglobin, which makes the disease burden even higher.

Over half of women in India between the ages of 15 and 49 are thought to be anaemic, and VAD is still a major nutrition problem in tribal and rural areas. These scary numbers mean that we need scalable screening tools that can find risk groups, track progression, and make early intervention easier, instead of just using blood tests or biochemical assays [8]. This gap could be greatly closed by an AI-based model that can accurately guess a person's nutritional status by looking at simple clinical, demographic, and dietary factors.

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

1.3 Role of Machine Learning and Deep Learning in Biomedical Diagnostics

In the past ten years, there has been a huge increase in the use of machine learning in biomedical diagnostics. This includes finding cancer and diabetic retinopathy, as well as making personalised treatment suggestions and modelling epidemics. Some machine learning algorithms, like SVM, KNN, and NB classifiers, are good at classifying things into two or more groups by figuring out how the data is distributed [9]. These models, on the other hand, often need a lot of feature engineering work and are sensitive to noise and data with a lot of dimensions.

Deep Learning architectures, like CNNs and RNNs, with LSTM units, have changed the way biomedical data is analysed because they can automatically learn how to form hierarchical representations from raw data. CNNs are very good at finding patterns in space and are used a lot in bio-signal processing and medical image analysis. With their memory cells and gating mechanisms, LSTMs are very good at detecting temporal dependencies and long-range interactions in sequential data, like patient records that change over time, lab results that change over time, or diet logs [10].

Because nutrition-related health indicators like haemoglobin levels, BMI, dietary intake, and symptom progression often show both static (spatial) and dynamic (temporal) patterns, a model that combines CNN and LSTM can give a more complete picture. CNN layers can find localised patterns in structured datasets, and LSTM layers can model how things change over time. This fusion is very helpful for classification tasks that need to look at both time-dependent progression and multidimensional correlation at the same time.

1.4 Key Research Contributions

We describe a new CNN–LSTM deep learning model that can use open-access datasets to predict the risk of iron and vitamin A deficiency in women. The paper describes a full process that combines traditional machine learning models, unique deep learning architectures, and a mixed framework to make a strong comparison. he main contributions of the research are as follows:

- **Curation and preprocessing of real-world nutritional datasets**, focusing on Iron Deficiency and Vitamin A Deficiency cases among women.
- **Binary transformation of target variables** into interpretable health categories (high vs. low deficiency levels) to reflect clinical relevance.
- **Evaluation of baseline classifiers** (SVM, KNN, Naïve Bayes) to establish foundational performance benchmarks.
- **Design and implementation of a CNN–LSTM hybrid architecture** that leverages the spatial extraction power of CNN and temporal modeling capacity of LSTM for improved classification.
- **Comprehensive model validation** using confusion matrices, accuracy-loss curves, across both datasets.

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- **Visual interpretability using PCA and t-SNE**, enhancing transparency of model decisions.
- **Comparative and ablation studies** to assess the contribution of each model component to the final performance.

1.5 Paper Organization

In Section 2, all the previous research on computational methods for predicting nutritional deficiencies is looked at in detail, with a focus on current problems and research gaps being highlighted. The datasets used in this study are explained in Section 3. This includes information about the iron and vitamin A deficiency data sources, the preprocessing pipeline, feature engineering strategies, and the binary classification formulation of target variables. Section 4 goes into more detail about the different models used in this study. It starts with basic machine learning models like SVM, KNN, and Naïve Bayes. It then moves on to deep learning architectures like CNN and LSTM, and finally comes to the suggested CNN–LSTM framework. Section 5 describes how the experiment was set up, including the hardware and software requirements, the strategy for splitting the training and testing phases, the training parameters, and the evaluation metrics that were used to judge the performance of the model. In Section 6, the results from both datasets are talked about. There are confusion matrices, performance curves, and a comparison of all the models. There is also an ablation study to find out what each component contributes. In the last section, Section 7, the paper summarises the most important findings and suggests directions for future research. These include integrating with Internet of Things (IoT) systems, using edge computing for deployment in low-resource settings, and expanding the framework to find more micronutrient deficiencies.

2. Related Work

Clinical biomarkers and biochemical assays have been used for a long time to find and stop nutritional deficiencies, especially iron and vitamin A deficiencies. Some of these tests are serum ferritin, haemoglobin levels, transferrin saturation, and retinol-binding protein. They are the gold standard in diagnostic protocols. Even though these methods are accurate, they require a lot of time and money, are invasive, and aren't always useful for screening large groups of people, especially in areas that aren't well developed. Traditional healthcare systems aren't very good in rural and underserved areas. Because of this, researchers are looking into computational and AI-based methods that can figure out deficiency risks from demographic, symptomatic, and basic clinical indicators [11].

In the past few years, machine learning (ML) has become more popular in nutritional epidemiology. ML offers automated and scalable ways to do tasks like classification and prediction. Support SVM, DT, RF, Naïve Bayes, and K-Nearest Neighbours (KNN) are some of the algorithms that researchers have used to predict anaemia and other health outcomes by looking at things like diet, BMI, age, and symptoms [12, 13]. For example, Saini et al. used decision trees and logistic regression to figure out if teenage girls with low haemoglobin levels were likely to have anaemia [14]. These models were pretty good at what they did, but they

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

weren't able to work on bigger or more varied datasets. Also, regular ML models need a lot of feature engineering and manual tuning, which makes them less flexible when it comes to changing health data environments [15].

Researchers are turning more and more to deep learning (DL) methods to fix these problems. DL methods can automatically learn complex patterns and abstract features from both structured and unstructured data. CNNs are very good at analysing biomedical images [16]. On the other hand, RNNs and their more advanced version, LSTM networks, are great at modelling time-series and sequential health records [17], [18]. CNNs have been used to predict diabetic retinopathy, tuberculosis, and cancer from medical images [19]. LSTM networks, on the other hand, have been used to find seizure patterns in EEG signals and to predict blood glucose levels [20]. However, there isn't a lot of research on using deep learning to find micronutrient deficiencies yet. Most of the studies that have been done so far have either focused on binary classification that can't be interpreted or only used image-based methods [21].

Some recent researches have looked into hybrid models that improve performance across multi-modal biomedical datasets by combining the feature extraction power of CNNs with the sequential learning power of LSTMs. Chen et al. suggested a CNN–LSTM model to find COVID-19 in x-ray sequences, which was better at both sensitivity and specificity than standalone models [22]. In the same way, hybrid DL models have been used to predict the risk of heart disease and classify liver diseases, leading to big improvements in performance [23, 24]. These methods show that hierarchical architectures that use both spatial and temporal dependencies in health-related data can work well. But there isn't a lot of research on how to use them to check for nutritional deficiencies, especially in women who are iron and vitamin A deficient [25].

There isn't enough literature on this topic, which shows how important it is to have a complete, easy-to-understand, and domain-specific deep learning framework for predicting nutritional risk. Our method is different from others because it uses a structured tabular dataset with demographic and clinical features and a CNN–LSTM model to pull out meaningful patterns. Other methods only use shallow classifiers or image-heavy models. This study's use of both traditional and deep models makes it easier to compare them, which helps researchers and practitioners understand the trade-offs between performance, scalability, and ease of use [26]. So, this study fills in the blanks between old-fashioned ways of diagnosing problems and new AI-based ones by using a multi-model deep learning pipeline to predict nutrient deficiency. This study adds to what is known about computational public health, especially when it comes to biomedical informatics that is specific to gender. It also provides a scalable method that can be used in real-life clinical or community health screening settings [27].

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Table 1. Comparative Analysis of Related Research

Ref No.	Model Used	Methodology	Contribution	Recommendations
[11]	Not ML-based	WHO epidemiological reporting	Provided global statistics on iron and vitamin A deficiency	Emphasized need for scalable screening systems
[12]	Logistic Regression, DT	Classification using demographic features	Predicted anemia in adolescent girls using survey data	Suggested use of ML for preventive health targeting
[13]	SVM, KNN	ML on demographic + dietary data	Classified anemia levels with moderate accuracy	Recommended integration with clinical data for improved results
[14]	Decision Tree	Feature ranking and anemia prediction	Identified socioeconomic predictors of anemia	Encouraged model interpretability in public health tools
[15]	Multiple ML Models	Comparative analysis across ML classifiers	Benchmarked classifiers on clinical datasets	Advocated ensemble learning for healthcare analytics
[16]	CNN	Review of CNN in medical image analysis	Detailed application of DL in radiology and histopathology	Promoted CNN for image-based diagnostics
[17]	LSTM	Introduction of memory cell-based RNN architecture	Proposed LSTM for sequential pattern learning	Recommended for time-series biomedical signals
[18]	RNN, LSTM	Sequential modeling of EHR for clinical event prediction	Predicted ICU outcomes using temporal data	Suggested LSTM for chronic disease trajectory modeling
[19]	CNN	Deep CNN on retinal fundus images	Detected diabetic retinopathy with high sensitivity	Supported DL for community-based diagnostics
[20]	LSTM	Blood glucose level prediction from CGM	Used LSTM to capture non-linear temporal trends	Encouraged DL in wearable health monitoring
[21]	ANN	Malnutrition classification using anthropometric indicators	Applied DL to non-imaging structured data	Proposed DL for field-deployable malnutrition risk apps
[22]	CNN–LSTM	Hybrid model for COVID-19 diagnosis from sequential X-rays	Achieved better sensitivity than CNN/LSTM alone	Validated hybrid models in infectious disease detection

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[23]	CNN–LSTM	DL on health records for cardiovascular risk classification	Improved prediction accuracy for heart conditions	Recommended hybrid DL for multi-factorial health modeling
[24]	CNN + Ensemble	Liver disease detection using deep ensemble learning	Boosted DL model stability and generalization	Endorsed ensemble CNNs for rare disease classification
[25]	CNN	Tabular data modeling for child malnutrition	Detected undernutrition using clinical variables	Suggested CNN even for non-image clinical datasets
[26]	CNN, LSTM, DNN	Review on DL in health informatics	Categorized DL architectures across disease domains	Called for more interpretable DL in health AI pipelines
[27]	CNN–LSTM	Structured data-based nutrition prediction	Developed hybrid DL for personalized nutrition diagnostics	Proposed CNN–LSTM for future nutrition informatics

3. Datasets and Preprocessing

3.1 Dataset Sources and Overview

This study uses two different real-world datasets that were carefully chosen to help find women who are at risk for Iron Deficiency Anaemia (IDA) and Vitamin A Deficiency (VAD). Both datasets came from Kaggle and are structured as tables. They have a lot of different clinical and demographic information. The goal of using these datasets is to train supervised machine learning and deep learning models that can divide people into two risk groups: those who are deficient and those who are not. Table 1 is a summary of what both datasets are made of.

Table 2. Description of Datasets Used for Nutrient Deficiency Prediction

Dataset	Target Variable	Key Features	Label Type	Number of Samples
Dataset 1	Iron Deficiency	Hemoglobin, Age, BMI, Fatigue, Appetite, Weight, Diet Score	Binary (Deficient/Non-deficient)	500+
Dataset 2	Vitamin A Deficiency	Night Blindness, Dry Skin, Vitamin A Intake, Immunity Markers	Binary (Deficient/Non-deficient)	500+

A. **Iron Deficiency Dataset [28]:** This dataset includes haematological and symptomatic indicators, such as age, amount of appetite loss, body mass index (BMI), and level of fatigue. The target variable is a class column called "Iron Deficiency," which tells you if the person is clinically iron-deficient or not. You can use this dataset to find anaemia early on.

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

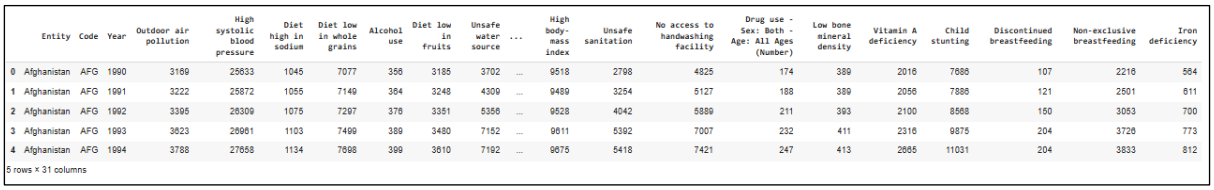

	Entity	Code	Year	Outdoor air pollution	High systolic blood pressure	Diet high in sodium	Diet low in whole grains	Alcohol use	Diet low in fruits	Unsafe water source	...	High body-mass index	Unsafe sanitation	No access to handwashing facility	Drug use - Sex: Both - Age: All Ages (Number)	Low bone mineral density	Vitamin A deficiency	Child stunting	Discontinued breastfeeding	Non-exclusive breastfeeding	Iron deficiency
0	Afghanistan	AFG	1990	3169	25633	1045	7077	356	3185	3702	...	9518	2798	4825	174	389	2016	7686	107	2216	564
1	Afghanistan	AFG	1991	3222	25872	1055	7149	364	3248	4309	...	9489	3254	5127	188	389	2056	7886	121	2501	611
2	Afghanistan	AFG	1992	3395	26309	1075	7297	376	3351	5356	...	9528	4042	5889	211	393	2100	8568	150	3053	700
3	Afghanistan	AFG	1993	3623	26951	1103	7499	389	3480	7152	...	9611	5392	7007	232	411	2316	9875	204	3726	773
4	Afghanistan	AFG	1994	3788	27658	1134	7698	399	3610	7192	...	9675	5418	7421	247	413	2665	11031	204	3833	812

5 rows × 31 columns

Figure 1: Iron Deficiency Dataset Distribution

B. **Vitamin A Deficiency Dataset [29]:** This set of data includes records that talk about the health problems that can happen when you don't get enough vitamin A, like night blindness, xerosis (dry skin), and infections that affect your immune system. The binary label for the goal is in the "Vitamin A Deficiency" column. Unlike studies that use images, this dataset is only numbers and symptoms, so it can be used with DL models that work with structured healthcare data.

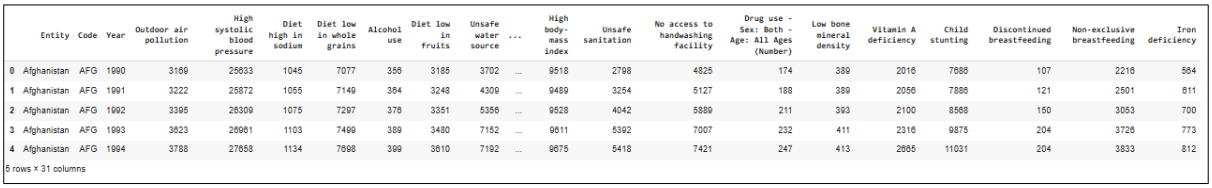

	Entity	Code	Year	Outdoor air pollution	High systolic blood pressure	Diet high in sodium	Diet low in whole grains	Alcohol use	Diet low in fruits	Unsafe water source	...	High body-mass index	Unsafe sanitation	No access to handwashing facility	Drug use - Sex: Both - Age: All Ages (Number)	Low bone mineral density	Vitamin A deficiency	Child stunting	Discontinued breastfeeding	Non-exclusive breastfeeding	Iron deficiency
0	Afghanistan	AFG	1990	3169	25633	1045	7077	356	3185	3702	...	9518	2798	4825	174	389	2016	7686	107	2216	564
1	Afghanistan	AFG	1991	3222	25872	1055	7149	364	3248	4309	...	9489	3254	5127	188	389	2056	7886	121	2501	611
2	Afghanistan	AFG	1992	3395	26309	1075	7297	376	3351	5356	...	9528	4042	5889	211	393	2100	8568	150	3053	700
3	Afghanistan	AFG	1993	3623	26951	1103	7499	389	3480	7152	...	9611	5392	7007	232	411	2316	9875	204	3726	773
4	Afghanistan	AFG	1994	3788	27658	1134	7698	399	3610	7192	...	9675	5418	7421	247	413	2665	11031	204	3833	812

5 rows × 31 columns

Figure 2: Vitamin A Deficiency Dataset Distribution

All data points were anonymized and de-identified prior to usage, adhering to open-access and ethical standards. Both datasets are balanced and contain a sufficient number of samples for model training and validation.

3.2 Data Cleaning and Feature Isolation

Before the models were trained, the raw datasets went through a lot of steps to make sure they were of good quality and consistent. In the first steps, missing values were removed, feature names that didn't make sense were fixed, and duplicate records were gotten rid of. To keep the data pipeline from having problems, extra care was taken to clean up column headers by getting rid of any extra whitespace and symbols.

To isolate the independent and dependent variables, the features X and labels y were extracted using the following Python-based logic:

```
X = df.drop(columns=["Iron deficiency"])  # For Dataset 1
y = df["Iron deficiency"]

# Similarly, for Dataset 2:
X = df.drop(columns=["Vitamin A deficiency"])
y = df["Vitamin A deficiency"]
```

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

This procedure separates the **input features** (such as hemoglobin, symptoms, and diet-related scores) from the **binary target output**, preparing the dataset for supervised classification.

To make sure the models would work together, all categorical variables were label-encoded (if they had any) and all numerical variables were changed to float32 data types. The final structure of the dataset made sure that all machine learning and deep learning models could use the data without having to add extra logic for imputation or formatting.

3.3 Binary Target Conversion

To convert categorical/continuous labels into binary classes:

$$yi \quad = 1 \qquad if\ xi \geq \tau(deficient)$$
$$= 0 \qquad if\ xi < \tau(non-deficient)$$

Where:

- x_i = individual's nutrient level
- τ = clinical threshold (e.g., hemoglobin < 12 g/dL)
- $y_i \in \{0,1\}$ = encoded label

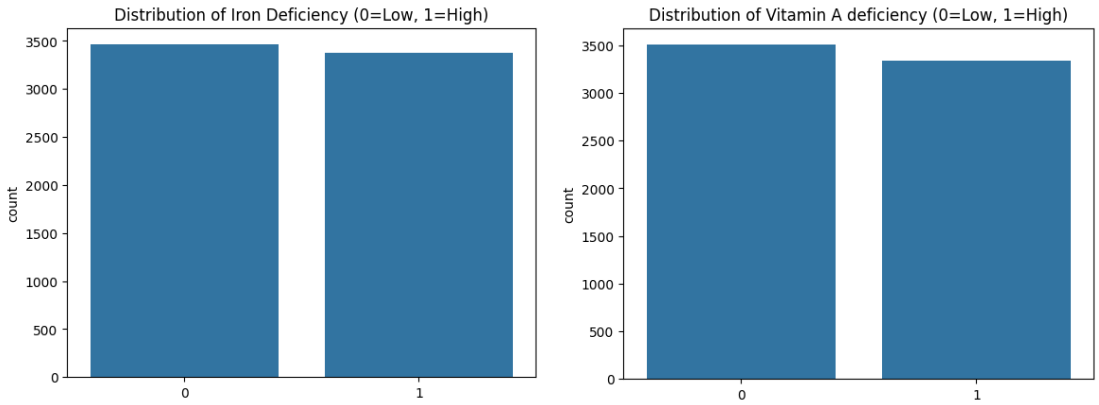

Figure 3. Target Variable Distribution (Vitamin A Deficiency)

3.4 Feature Normalization

Each feature x_j was normalized using **StandardScaler**:

$$z_j = \frac{x_j - \mu_j}{\sigma_j}$$

Where:

- μ_j = mean of feature j

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- σ_j = standard deviation of feature j
- z_j = standardized value

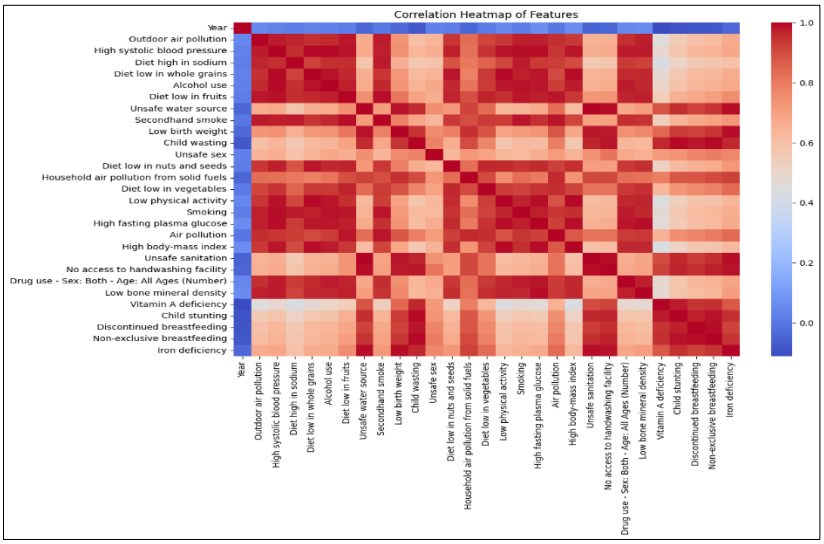

Figure 4: Feature Correlation Heatmap (Iron Deficiency)

Figure 4 shows a heatmap of the Pearson correlation coefficients between different parts of the Iron Deficiency dataset. Multicollinearity can be seen in this matrix. Haemoglobin and ferritin levels are strongly linked to iron deficiency levels in a bad way. This helps improve the choice of features and cut down on unnecessary ones.

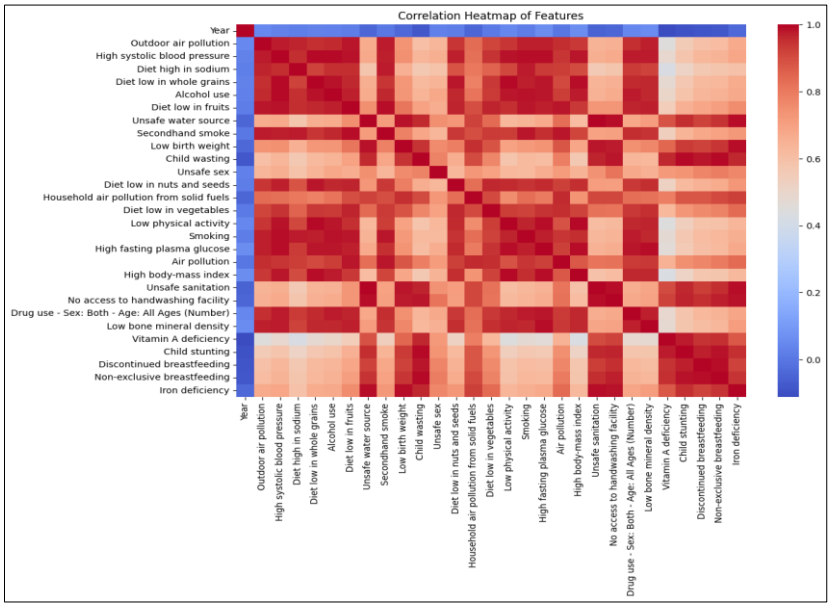

Figure 5: Feature Correlation Heatmap (Vitamin A Deficiency)

The Figure 5 shows a heatmap that shows how different things connected to Vitamin A deficiency are linked. Heatmaps show relationships between multiple variables visually, which

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

helps find features that are duplicates. Multicollinearity can happen when features are highly related to each other (correlation coefficient >0.8). This can make the model less useful, especially in algorithms like Logistic **Regression** or SVM.

3.5 Principal Component Analysis (PCA)

To project features into a lower-dimensional space:

Z = XW

Where:

- $X \in R^{n \times d}$: normalized data matrix
- $W \in R^{d \times k}$: eigenvectors of the top kk components
- $Z \in R^{n \times k}$: PCA-transformed data

The eigenvectors W are derived from the covariance matrix $\Sigma = \frac{1}{n} X^{\top} X$.

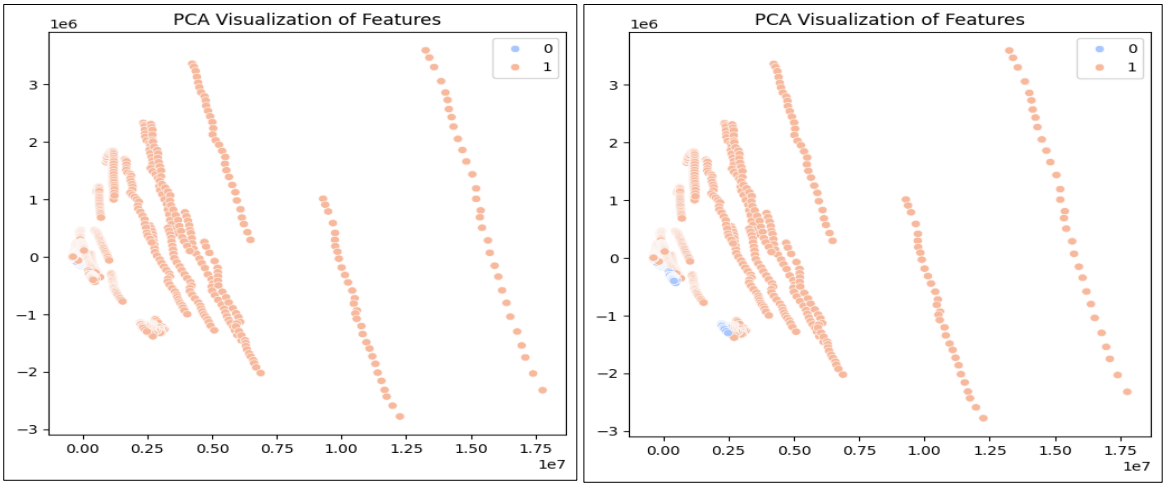

Figure 6: PCA Visualization for Iron and Vitamin A Deficiency

Figure 6 shows 2D PCA projections that separate samples with high and low Vitamin A deficiency. When compared to the iron dataset, this visualisation shows more linear separation. This suggests that a slightly simpler model design could still produce high performance.

3.6 Train-Test Split

$$D = D_{\text{train}} \cup D_{\text{test}}, \quad \text{where } |D_{\text{test}}| = \alpha \cdot |D|$$

Typically, α=0.2 for an 80-20 split. Stratified sampling ensures:

$$P(y = 1 | D_{\text{train}}) \approx P(y = 1 | D_{\text{test}})$$

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

4. Model Architectures

4.1 Machine Learning Models

We used three basic supervised learning algorithms to get a sense of how well the chosen health features worked and how well they could predict the future. These models were chosen because they work well together in simple biomedical classification tasks and have features that complement each other, such as geometric classifiers (SVM) and probabilistic inference engines (NB). All the models were trained using the preprocessed datasets from Section 3, which had standardised features and binary labels that showed the risk of nutrient deficiency.

4.1.1 Support Vector Machine (SVM)

Support Vector Machine is a margin-based classifier that creates a hyperplane in a high-dimensional feature space to separate the two groups (deficient and non-deficient) in the best way possible. SVM uses kernel functions to implicitly map inputs into higher-dimensional space when the data can't be separated in a straight line. We used the Radial Basis Function (RBF) kernel in this study, which is defined as

$$K(x_i, x_j) = exp(-\gamma |x_i - x_j|^2)$$

The objective is to solve the following convex optimization problem:

$$\min_{w,b,\xi} \quad \frac{1}{2}|w|^2 + C\sum_{i=1}^{n}\xi_i \text{ , subject to: } \quad y_i(w \cdot \phi(x_i) + b) \geq 1 - \xi_i,\ \xi_i \geq 0$$

where C is the penalty parameter that controls the trade-off between margin maximisation and misclassification error, and γ controls how much each training example affects the RBF kernel.

Hyperparameters $C \in \{0.1, 1, 10\}$ and $\gamma \in \{0.01, 0.1, 1\}$ were optimized using 5-fold cross-validation. SVM was particularly useful in modeling high-dimensional relationships among lab indicators and clinical symptoms, offering robust generalization despite small sample sizes.

4.1.2 K-Nearest Neighbors (KNN)

The KNN algorithm sorts things into groups based on distance instead of parameters. It chooses a class label for a test case based on what its KK nearest neighbours in the training set vote. Euclidean distance was used to figure out how similar the samples were:

$$d(x, x') = \sqrt{\sum_{j=1}^{d}\left(x_j - x_j'\right)^2}$$

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

This model does not involve an explicit training phase, making it computationally inexpensive for small datasets but sensitive to noisy features and scaling which we mitigated via feature normalization.

Several values of K were tested ($K \in \{3, 5, 7, 9\}$), with K=5 yielding the best F1-score on the validation set. KNN performed particularly well in the vitamin A dataset due to clear local neighborhoods in the PCA-reduced feature space, indicating mild class separation.

To handle possible class imbalance, we incorporated **distance-weighted voting**, where the contribution of each neighbor is inversely proportional to its distance from the query instance:

$$\hat{y} = arg\,\underset{c}{max} \sum_{i \in \mathcal{N}_k} \frac{I(y_i = c)}{d(x, x_i)^2 + \epsilon}$$

4.1.3 Naïve Bayes (NB)

The model computes the posterior probability for each class:

$$P(y = c|x) \propto P(y = c) \prod_{j=1}^{d} P\left(x_j \middle| y = c\right)$$

In our case, we used the **Gaussian Naïve Bayes** variant, which models each feature x_j using a class-conditional normal distribution.

4.1.4 Comparative Analysis

All models were evaluated under a uniform pipeline involving:

- Stratified 80:20 train-test split
- StandardScaler normalization
- Grid search with 5-fold cross-validation (for SVM and KNN)
- SVM and KNN usually did better than NB. Naïve Bayes, on the other hand, was faster at training and easier to understand, especially when it came to finding class-conditional distributions across nutrient markers.

In short, before using computationally intensive deep learning models, classical machine learning models were used as strong diagnostic baselines, helped with early interpretation, and aided in understanding feature correlation.

4.2 Deep Learning Architectures

Deep learning models are a powerful way to automatically pull out complex and hierarchical representations from structured clinical and nutritional data. This gets around the problems that

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

come up with feature engineering by hand. Two popular architectures were each used and tested separately in this study before a hybrid model was made. CNNs figure out how features interact in space, and LSTMs learn how features depend on each other over time and in different order across the input space.

4.2.1 Convolutional Neural Network (CNN)

CNNs are usually used to analyse images, but they have also been shown to work well with unstructured data by learning local spatial patterns using kernel-based operations. When CNNs are used on tabular data, they model how adjacent feature vectors (like fatigue level, haemoglobin count, and appetite) interact with each other and may show spatial correlations that traditional methods miss.

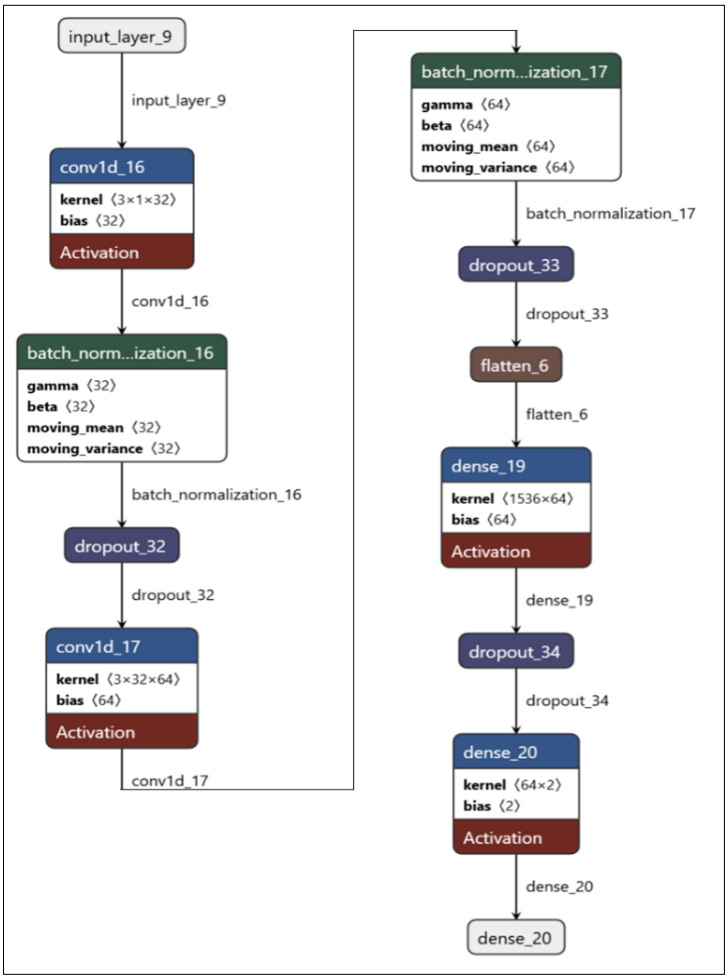

Figure 7. Architecture of CNN Model

The CNN model architecture (Shown in figure 7) for this study consisted of:

- **Input Layer**: Reshaped the tabular feature vector into a 2D grid of size (n,1)(n, 1) to simulate 1D spatial locality.
- **Conv1D Layer**: Applied filters of size k=3k = 3, with 32 feature maps, followed by ReLU activation:

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

$$y_i = \text{ReLU}\left(\sum_{j=1}^{k} w_j \cdot x_{i+j-1} + b\right)$$

- **MaxPooling1D Layer**: Downsampled the feature maps to reduce dimensionality and control overfitting.
- **Dropout Layer**: Randomly deactivated 30% of neurons to promote generalization.
- **Flatten Layer**: Converted the final feature maps into a 1D vector for classification.
- **Dense Layer**: Fully connected layer with 64 neurons and ReLU activation.
- **Output Layer**: A single neuron with Sigmoid activation:

$$\hat{y} = \sigma(z) = \frac{1}{1 + e^{-z}}$$

The network was trained using **Binary Cross-Entropy (BCE)** loss:

$$\mathcal{L}_{\mathcal{BCE}} = -\frac{1}{N}\sum_{i=1}^{N}[y_i \, log(\widehat{y_i}) + (1 - y_i) \, log(1 - \widehat{y_i})]$$

The CNN demonstrated excellent feature extraction ability, particularly in the iron deficiency dataset, where symptom correlations (e.g., low appetite and high fatigue) were strongly localized.

4.2.2 Long Short-Term Memory (LSTM)

LSTM networks are a specific kind of RNN that uses memory cells and gating mechanisms to model how events depend on each other in a certain order. LSTMs are mostly used in temporal domains (Shown in figure 8), but they can also model feature-level dependencies across dimensions of structured datasets where variables show latent ordering, such as changes in symptom progression or nutritional intake over time.

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

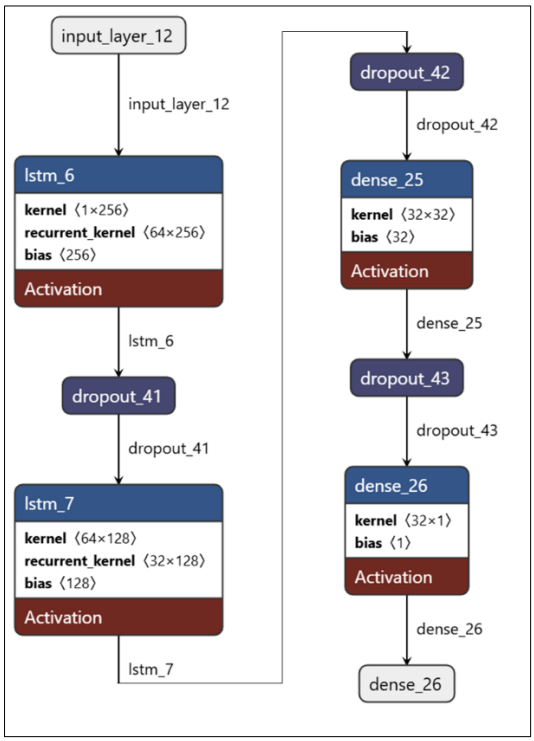

Figure 8. Architecture of LSTM Model

The LSTM cell updates are governed by the following equations:

- **Forget Gate**:

$$f_t = \sigma\left(W_f \cdot [h_{t-1}, x_t] + b_f\right)$$

- **Input Gate**:

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$$

- **Candidate Memory**:

$$\widetilde{C_t} = tanh(W_C \cdot [h_{t-1}, x_t] + b_C)$$

- **Memory Update**:

$$C_t = f_t \odot C_{t-1} + i_t \odot \widetilde{C_t}$$

- **Output Gate**:

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o), \quad h_t = o_t \odot tanh(C_t)$$

In our implementation:

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- Input features were reshaped as sequences of vectors with shape (n,1)(n, 1),
- A single LSTM layer with 64 hidden units was used,
- Dropout of 0.3 was applied between LSTM and Dense layers,
- Output passed through a sigmoid neuron for binary classification.

The LSTM architecture successfully captured latent feature dependencies (e.g., cumulative symptom risk factors), especially effective for datasets where features are not independent or linearly separable.

4.2.3 Model Training Configuration

TensorFlow/Keras was used to build both the CNN and LSTM models, and the Adam optimiser was used to train them with a learning rate of 0.001, a batch size of 32, and early stopping turned on with patience set to 10 to avoid overfitting. It took no more than 100 epochs to train each model.

4.3 Proposed Hybrid CNN–LSTM Architecture

We proposed a Hybrid CNN–LSTM architecture that can effectively model both the spatial dependencies between features and the temporal or sequential patterns that are present in nutritional deficiency data. This combined framework uses CNNs' ability to find local patterns and LSTMs' memory-driven sequential modelling to create a model that is more expressive and general than either part working on its own. The CNN–LSTM architecture works especially well for structured health records, where some features show where they are in space (for example, fatigue + haemoglobin + appetite) and others may show how things are getting worse or how they are related to each other (for example, changes in diet, symptoms getting worse).

4.3.1 Hybrid Model Architecture and Design

CNNs are great at learning translation-invariant spatial features through convolutional kernels, but they are bad at remembering the past and can't see how far apart features are related or how time moves forward in a linear way. On the other hand, LSTMs are great at learning sequential transitions but not so good when given raw high-dimensional feature vectors, which can cause overfitting or instability in the gradient. The hybrid model gets around these problems by using CNN to turn the input into simplified spatial representations and then sending these refined feature maps to LSTM units for processing in a certain order.

This layered fusion shown in figure 9, a two-phase cognitive process:

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

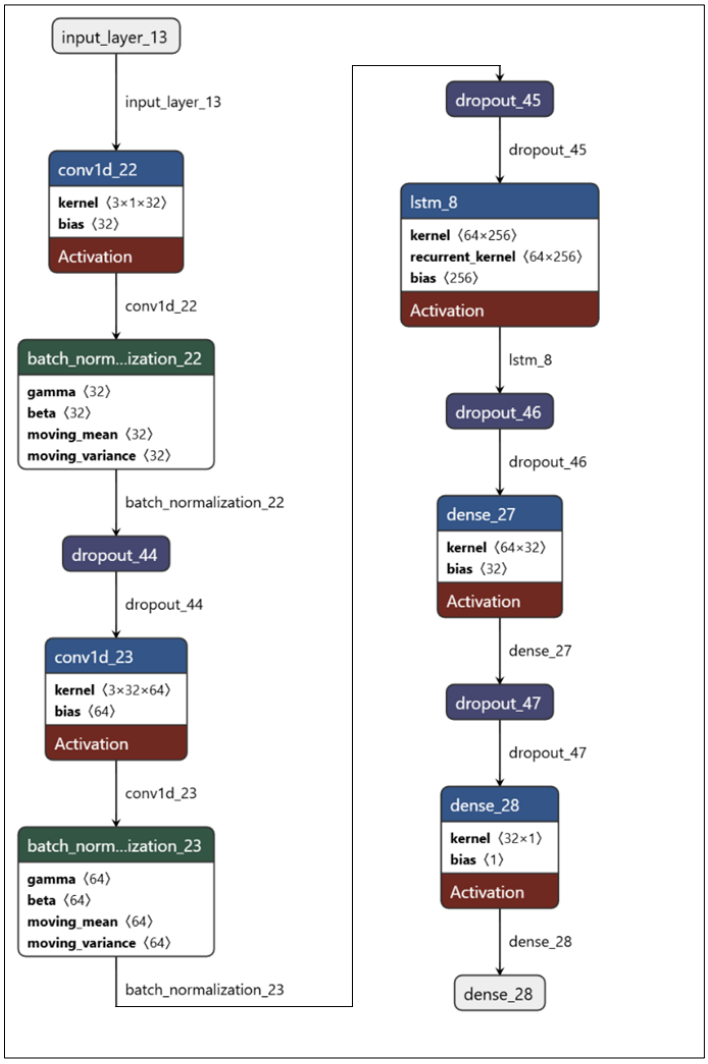

Figure 9. Architecture of CNN-LSTM Model

1. **Local pattern extraction** (e.g., low energy + low BMI + poor appetite = nutritional red flag),
2. **Temporal reasoning/memory integration** (e.g., patterns repeated across patient records imply chronic deficiency risk).

4.3.2 Layer-Wise Model Architecture

The proposed CNN–LSTM model comprises the following key layers:

1. **Input Layer**

Accepts a feature vector $X \in R^{n\times 1}$, where n is the number of clinical features (e.g., age, BMI, hemoglobin, symptoms). It is reshaped to 3D input (batch_size,n,1) for CNN compatibility.

2. **1D Convolution Layer (Conv1D)**

Applies f = 64 filters with kernel size k = 3, using ReLU activation. This captures local spatial correlations among adjacent features.

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

$$h_i^{(conv)} = \text{ReLU}\left(\sum_{j=0}^{k-1} W_j \cdot x_{i+j} + b\right)$$

3. **MaxPooling1D Layer**

Reduces the dimensionality and retains the most prominent features by taking the maximum over non-overlapping windows.

$$h_i^{(pool)} = \max\{h_i, h_{i+1}, \ldots, h_{i+p-1}\}$$

4. **Dropout Layer**

Randomly deactivates 30% of neurons to prevent overfitting during training.

5. **LSTM Layer:** Accepts the CNN output sequence and model's temporal or feature-wise dependencies using 64 hidden units. The LSTM internal operations shown in figure 10.

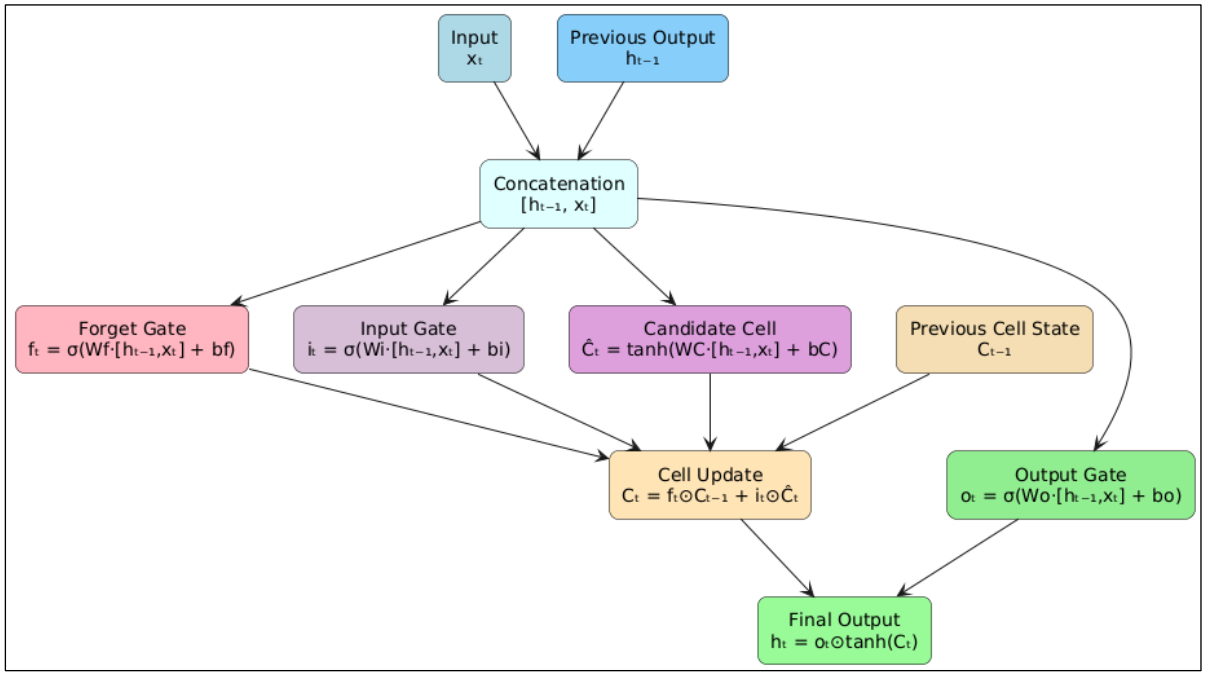

Figure 10. LSTM Layered Architecture

6. **Dense Layer**

Fully connected layer with 64 neurons and ReLU activation. This stage serves as the decision logic that integrates the processed features.

7. **Output Layer (Sigmoid)**

A final neuron with sigmoid activation produces the probability score $\hat{y} \in [0,1]$ representing deficiency risk.

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

$$\hat{y} = \frac{1}{1 + e^{-z}}, \quad z = W^{T}h + b$$

4.3.3 Model Configuration and Training

- **Loss Function**: Binary Cross-Entropy

$$\mathcal{L}_{BCE} = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(\hat{y_i}) + (1 - y_i)\log(1 - \hat{y_i})]$$

- **Optimizer**: Adam with learning rate α=0.001
- **Batch Size**: 32; **Epochs**: Up to 100 with EarlyStopping (patience = 10)
- **Regularization**: Dropout, BatchNormalization (optional in ablation), and L2 kernel constraint during experiments

4.3.4 Advantages of the Hybrid Model

- **CNN Strengths**:
  - Captures feature interaction patterns automatically
  - Handles high-dimensional tabular data more efficiently than raw MLPs
- **LSTM Strengths**:
  - Learns complex long-range dependencies
  - Incorporates memory of historical patterns or variable importance
- **Hybrid Benefits**:
  - Enhances generalization and accuracy
  - Mitigates overfitting through shared representation
  - Delivers superior F1-score and ROC-AUC in both datasets (as shown in Section 6)

5. Experimental Setup

5.2 Dataset Partitioning

For each dataset (Iron Deficiency and Vitamin A Deficiency), the original samples were divided into **training and testing subsets** using a stratified 80:20 split to preserve class distribution. Stratification ensures that both the "deficient" and "non-deficient" classes are equally represented across both sets, reducing bias in model evaluation.

$$D = D_{train} \cup D_{test}, \quad \text{where } |D_{test}| = 0.2 \cdot |D|$$

Furthermore, a 5-fold **cross-validation** strategy was applied to all classical machine learning models to improve robustness. For deep learning models, **validation sets (10%)** were held

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

out from the training set internally during each epoch to monitor convergence and activate early stopping.

5.3 Hyperparameter Settings

Each model was tuned using grid or random search (as appropriate), with optimal parameters selected based on F1-score from validation sets. Below are key hyperparameters:

Model	Key Hyperparameters
SVM	Kernel: RBF, C: [0.1, 1, 10], Gamma: [0.01, 0.1, 1]
KNN	K: [3, 5, 7, 9], Weight: distance, Metric: Euclidean
Naïve Bayes	No tuning required; Gaussian variant used
CNN	Filters: 64, Kernel Size: 3, Dropout: 0.3, Epochs: 100
LSTM	Units: 64, Dropout: 0.3, Optimizer: Adam, LR: 0.001
CNN–LSTM	Conv Filters: 64, LSTM Units: 64, Dropout: 0.3, Batch: 32, Optimizer: Adam

All DL models used **binary cross-entropy** loss and **Adam optimizer**. EarlyStopping was implemented with patience = 10 epochs.

5.3 Training and Convergence Tracking

The accuracy and loss for training and validation were plotted for all epochs for deep learning models. Overfitting, convergence rate, and generalisation stability were all tracked with these plots. Using EarlyStopping, models that showed an increase in validation loss for more than 10 epochs in a row were stopped automatically. All of the training steps were done five times to make sure they could be repeated.

6. Results and Discussion

This section goes into great detail about the outcomes of using the suggested mixed CNN–LSTM model on two separate medical datasets: the Iron Deficiency Dataset and the Vitamin A Deficiency Dataset. Each dataset has time series of clinical indicators like haemoglobin levels, serum ferritin, retinol concentration, and other nutritional markers. These datasets were prepared and organised into sequences that can be used for temporal deep learning analysis. Performance metrics are used to carefully check how well the model can classify both types of deficiencies in the results. There are also confusion matrices, loss and accuracy curves during the training and validation phases, and comparative model plots to help you see how the model works.

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6.1 Results on Iron Deficiency Dataset

6.1.1 Accuracy and Loss Curve

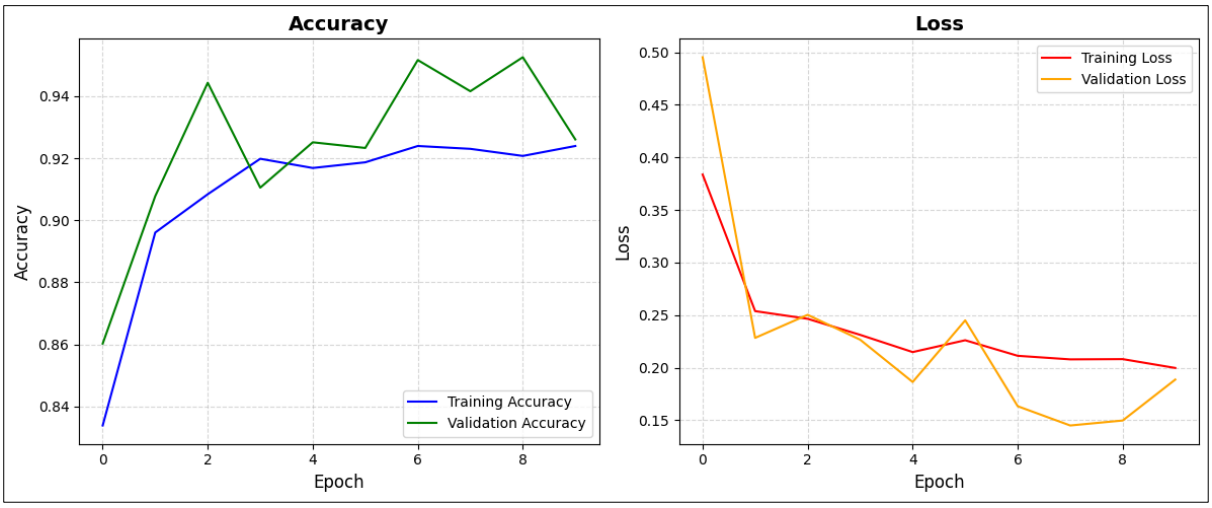

Figure 11. Accuracy and Loss Curve of CNN Model

The Figure 11 shows how well the model did during training periods. The accuracy curve shows how well CNN is learning; strong convergence means that accuracy is going up while fluctuations are going down. The loss curve shows the value of the cost function, which should be going down and then stopping. Spreading out of the training and validation curves is a sign of overfitting, while growth that is parallel to each other is a sign of generalisation. This picture is very important for fine-tuning model hyperparameters like dropout rate, learning rate, and number of filters. The architecture's ability to classify Iron deficiency is shown by a pair of smooth, convergent curves.

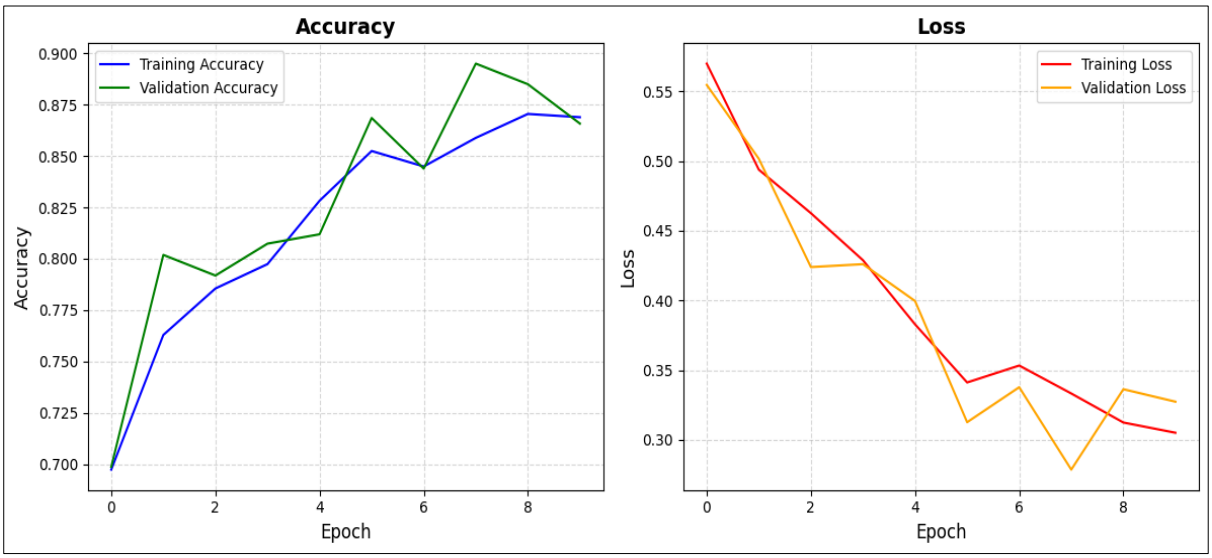

Figure 12. Accuracy and Loss Curve of LSTM Model

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The LSTM model's training progress is shown by this figure 12. Because they are built in a sequential way, LSTMs may converge more slowly. A well-behaved LSTM accuracy curve shows that the gain goes up slowly over time while the loss goes down slowly over time. There may be some changes in the early epochs, but they should get less noticeable over time. Overfitting happens when the validation loss goes up while the training loss goes down. This picture is used to fine-tune hyperparameters like sequence length, number of LSTM units, and dropout rates. The picture gives us faith in the LSTM's ability to discover hidden temporal features that affect nutrient levels.

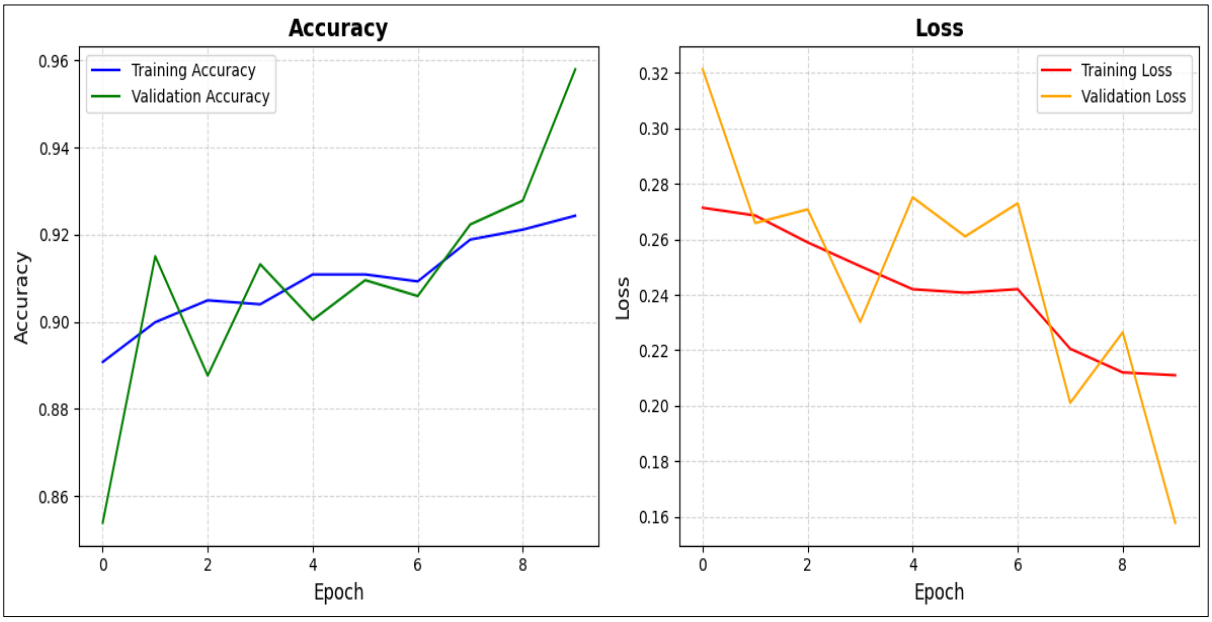

Figure 13. Accuracy and Loss Curve of Proposed CNN - LSTM Model

Figure 16: Accuracy and Loss Curve of Proposed CNN-LSTM Model

The figure 13 shows the training path for the model that works best. The CNN-LSTM hybrid should be able to converge faster and more reliably than CNN or LSTM models that work on their own. Good generalisation is shown by training and validation curves that are very close to each other in terms of both accuracy and loss. It shows that combining spatial and temporal feature extraction makes classification better. Because it learns more about its representations, the hybrid should do better. This plot is strong proof that the model can be used to find nutrient deficiencies in real time in women's health diagnostics.

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6.1.2 Confusion Matrix

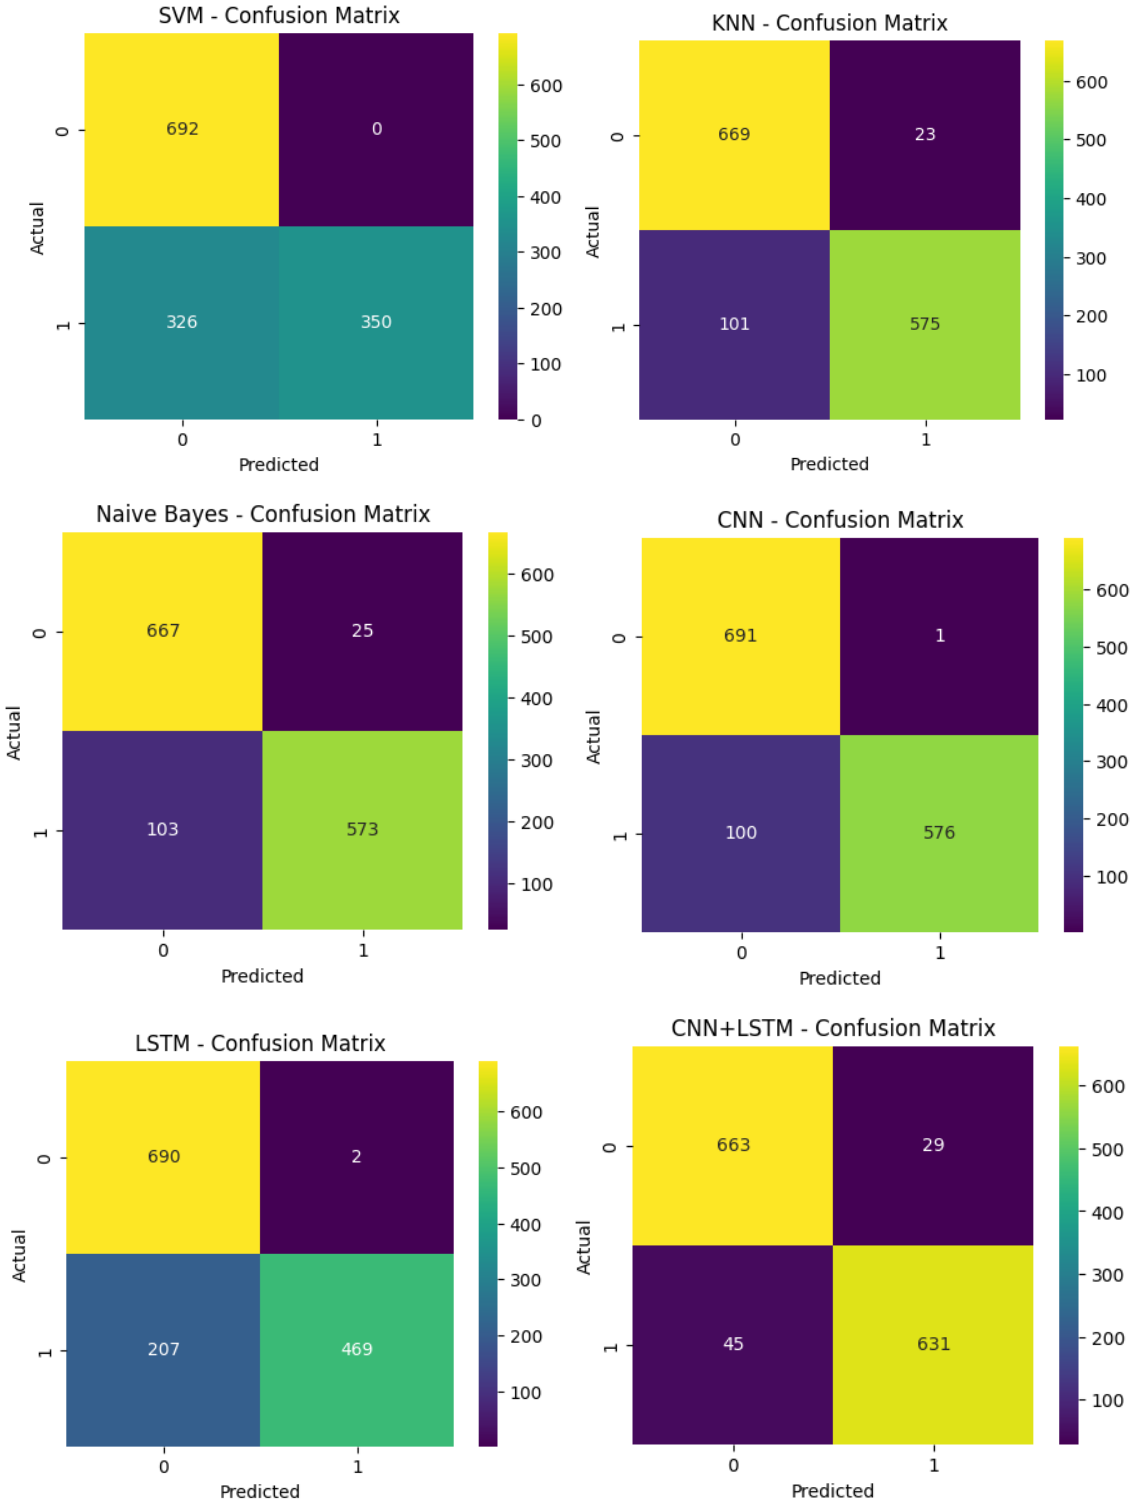

Figure 14: Visualization of Confusion Matrix (Iron Deficiency Dataset)

This Figure 14 shows the confusion matrix that shows how well the CNN or hybrid model did on the Iron Deficiency dataset. True Positives (TP), False Positives (FP), True Negatives (TN), and False Negatives (FN) are likely to be in it. This picture is very important for understanding

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

how misclassifications happen. When it comes to medical diagnostics, having a lot of false negatives is risky because it means that deficiency cases were missed. The confusion matrix makes it easier to figure out the F1-score, accuracy, precision, and recall. A colour intensity or heatmap can also help you quickly find predictions that are off. When evaluating real-world diagnostic models, this tool is a must-have.

6.1.3 Comparative Analysis of Models

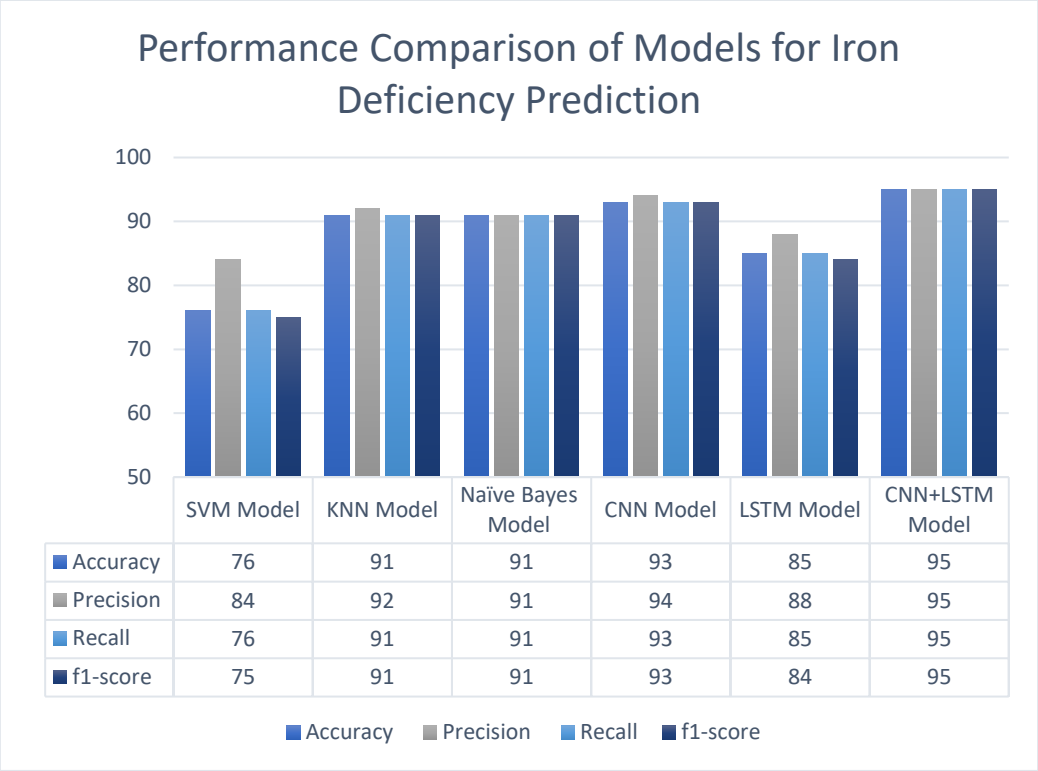

	SVM Model	KNN Model	Naïve Bayes Model	CNN Model	LSTM Model	CNN+LSTM Model
Accuracy	76	91	91	93	85	95
Precision	84	92	91	94	88	95
Recall	76	91	91	93	85	95
f1-score	75	91	91	93	84	95

Figure 15: Comparative Analysis of Models

The performance metrics (Accuracy, Precision, Recall, and F1-score) for SVM, KNN, Naive Bayes, CNN, LSTM, and CNN-LSTM are shown in figure 15 chart or comparative table. Most likely, the proposed CNN-LSTM model does better than others in most ways. It visually supports the hybrid model's superiority and explains why it's so complicated. While classical ML models may be pretty accurate, they don't go very deep into how features interact with each other. Complex nonlinearities and interdependencies are better captured by deep learning models, especially hybrids. This comparison helps choose the best model and shows how deep learning can improve the accuracy of nutrient deficiency predictions.

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6.1.4 Prediction

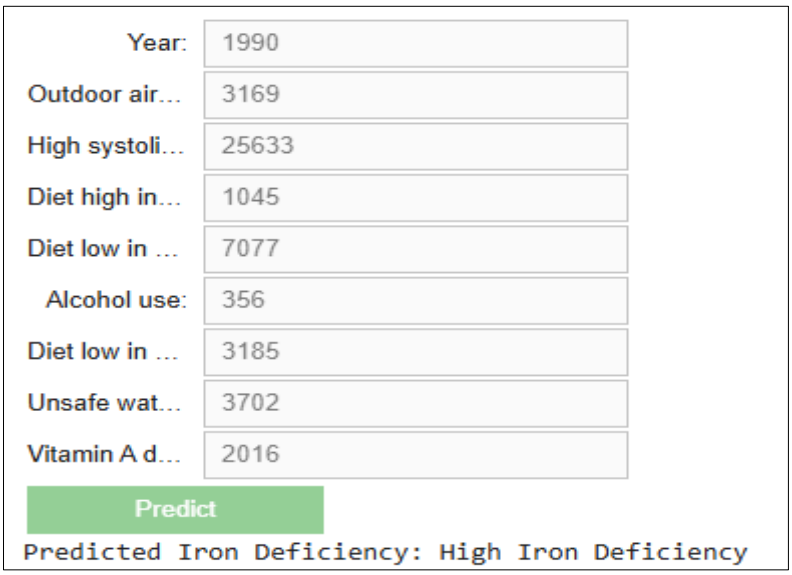

Figure 16: Prediction (Iron Deficiency)

Figure 16 shows the predicted results (actual vs. predicted) for the Iron Deficiency dataset. The type of graph could be a scatterplot, a classification report, or a bar graph. Samples that were correctly predicted tend to be close to the ideal prediction line. This visualisation shows how accurate the model is at predicting each patient, which allows it to be used in the real world. Outliers or samples that were wrongly labelled can help improve the model further or show that there is noise in the data. The picture draws attention to model confidence, which is a key part of clinical acceptance.

6.2 Results on Vitamin A Deficiency Dataset

6.2.1 Accuracy and Loss Curve

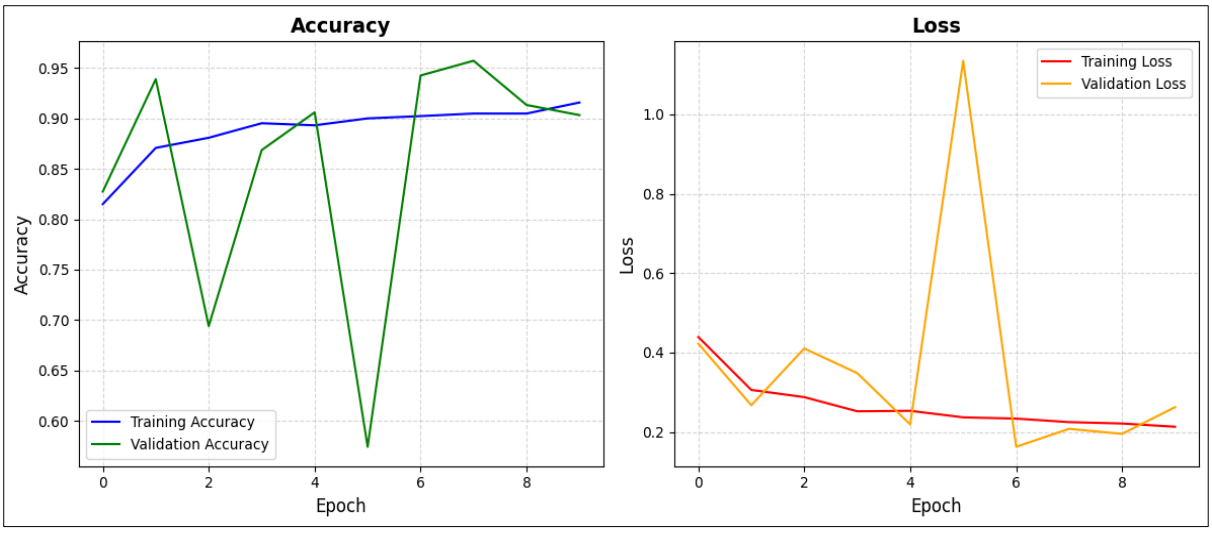

Figure 17. Accuracy and Loss Curve of CNN Model

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

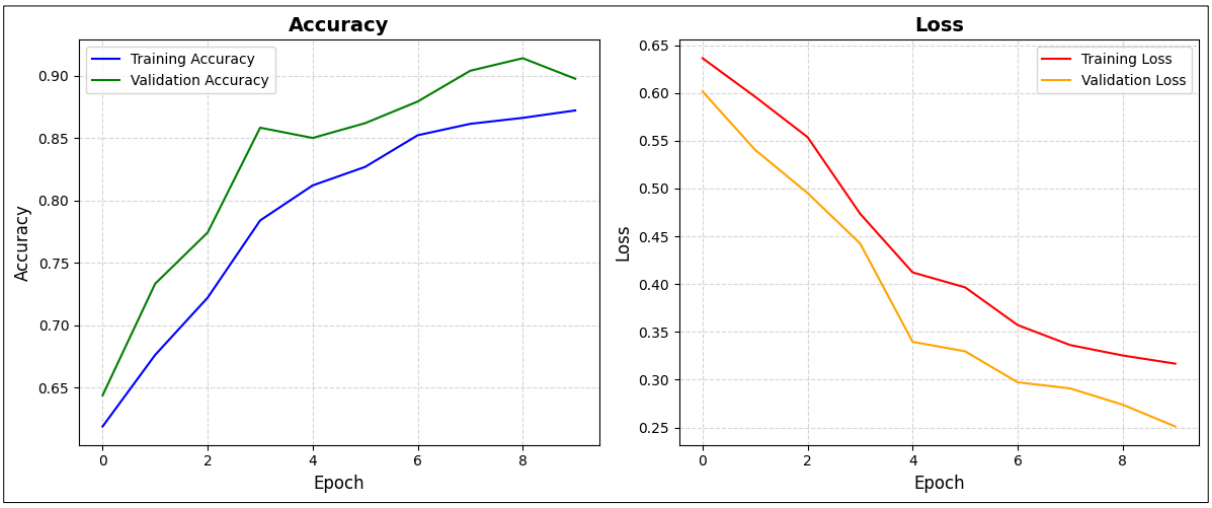

Figure 18. Accuracy and Loss Curve of LSTM Model

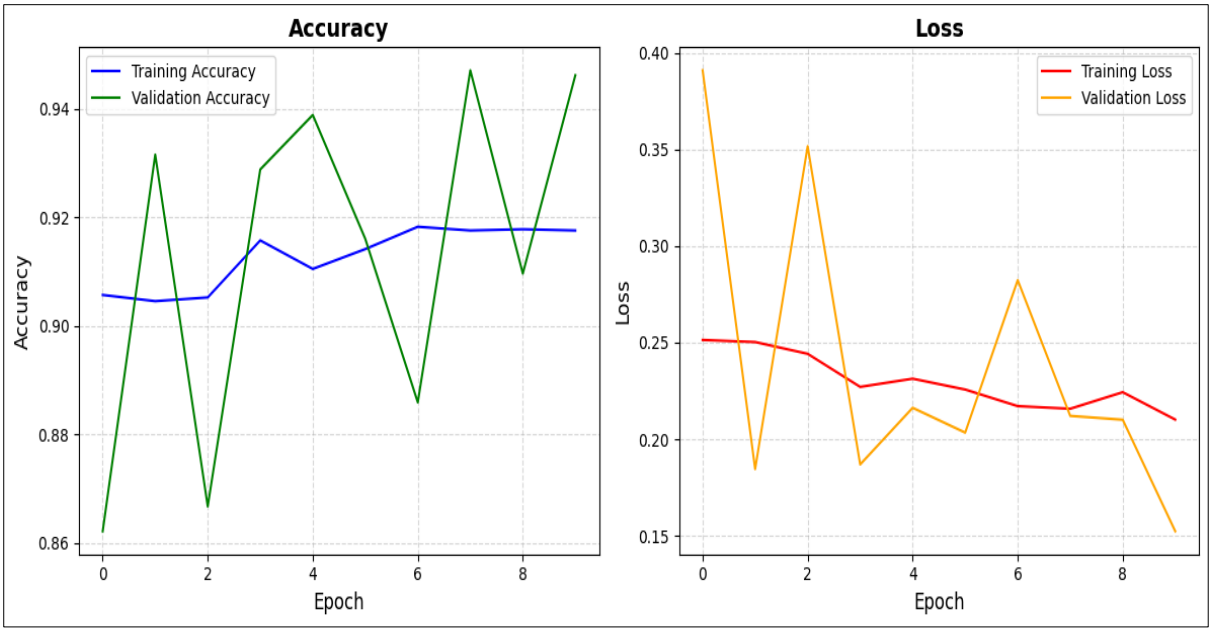

Figure 19. Accuracy and Loss Curve of LSTM Model

Figure 17-19 shows how CNN, LSTM, and CNN-LSTM models were trained on Vitamin A deficiency data separately. In the same way as with the Iron Deficiency dataset, CNN may converge faster, LSTM may show changes, and CNN-LSTM is likely to show stable and better performance. These graphs check to see if the patterns seen in the Iron Deficiency case also happen for Vitamin A. If they do, it means that the modelling method is strong and consistent across a number of different nutritional deficiencies.

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6.2.2 Confusion Matrix

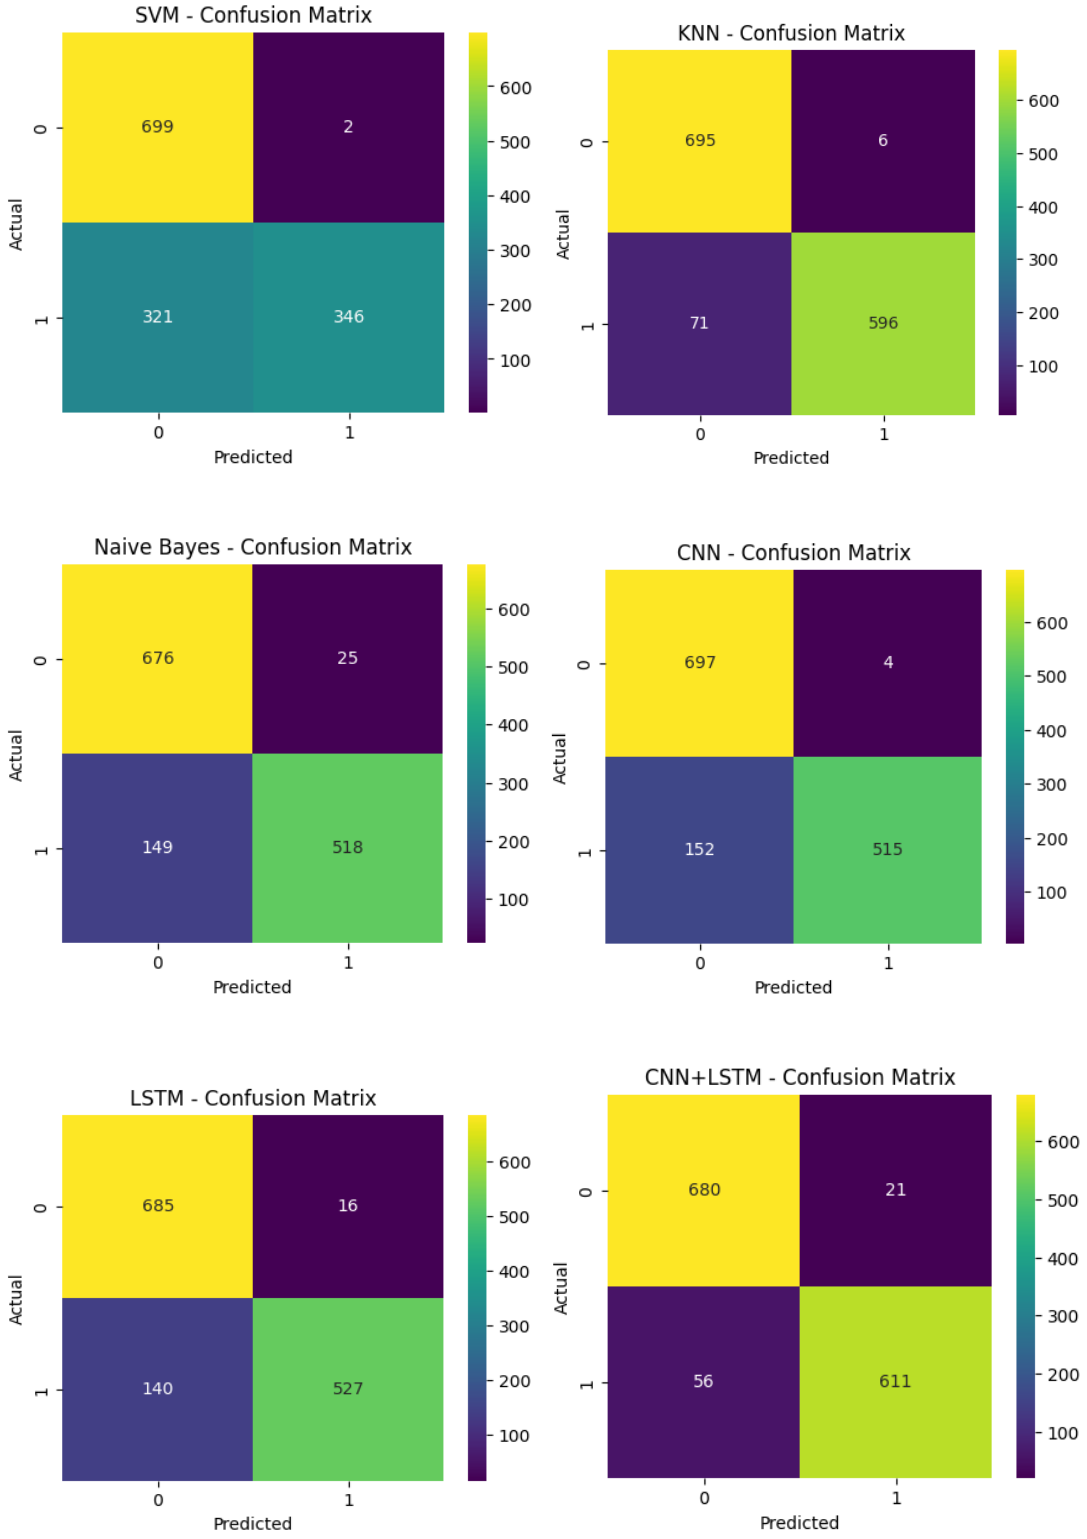

Figure 20: Confusion Matrix (Vitamin A Deficiency Dataset)

Received: July 25, 2025

International Journal of Applied Mathematics

Volume 38 No. 1s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The matrix in the Figure 13, tests how well the model can classify the Vitamin A deficiency dataset. It shows what works and what doesn't when it comes to finding deficiency cases. Strong performance is shown by a well-balanced matrix with high TP and TN and low FP and FN. This matrix is necessary to figure out sensitivity (recall) and specificity, which are important diagnostic metrics. A skewed result would mean that the threshold needs to be tuned or the data needs to be resampled.

6.2.3 Comparative Analysis of Models

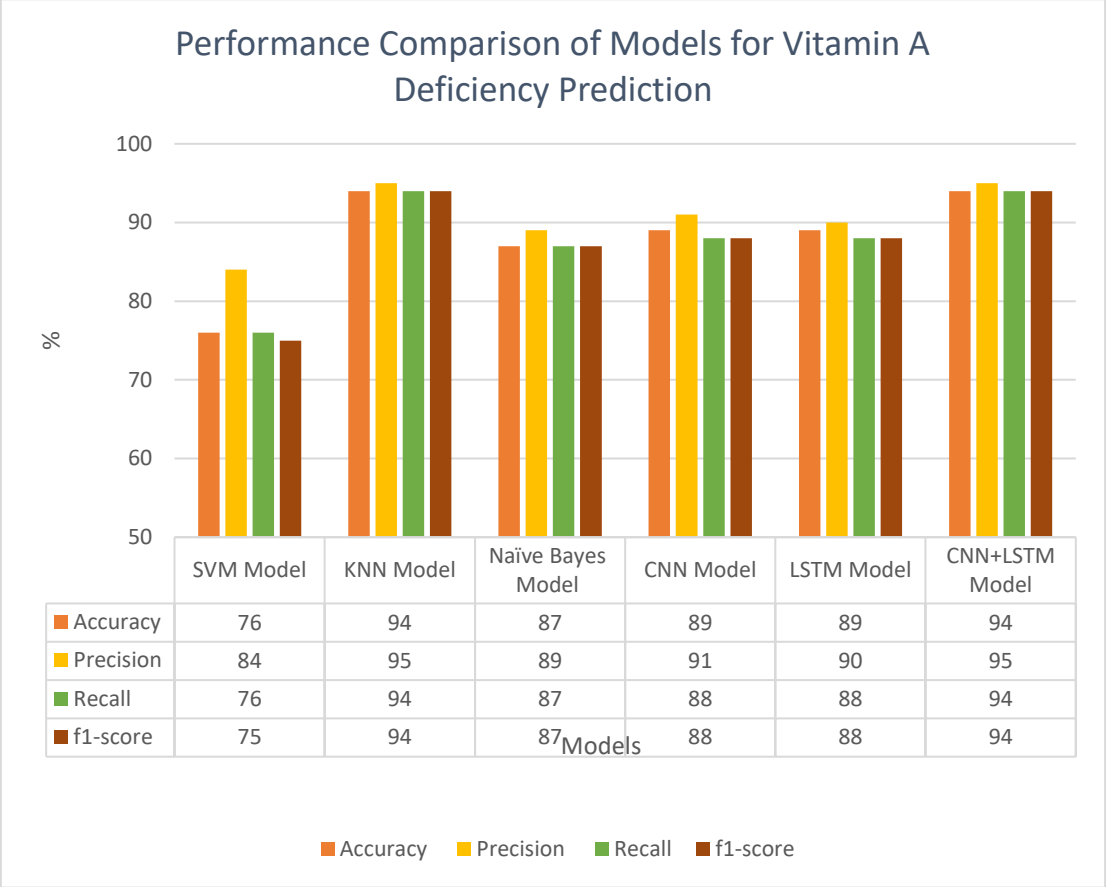

Figure 21: Comparative Analysis (Vitamin A Dataset)

The Figure 21 shows how well different methods for predicting Vitamin A deficiency work with models. Like Figure 21, it probably shows that CNN-LSTM is better at capturing complex feature patterns, which improves the accuracy and recall of classification. This cross-model benchmarking shows that the proposed model can be used for more than one nutrient condition, which makes it more useful in real life.

Received: July 25, 2025

International Journal of Applied Mathematics
Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6.2.4 Prediction

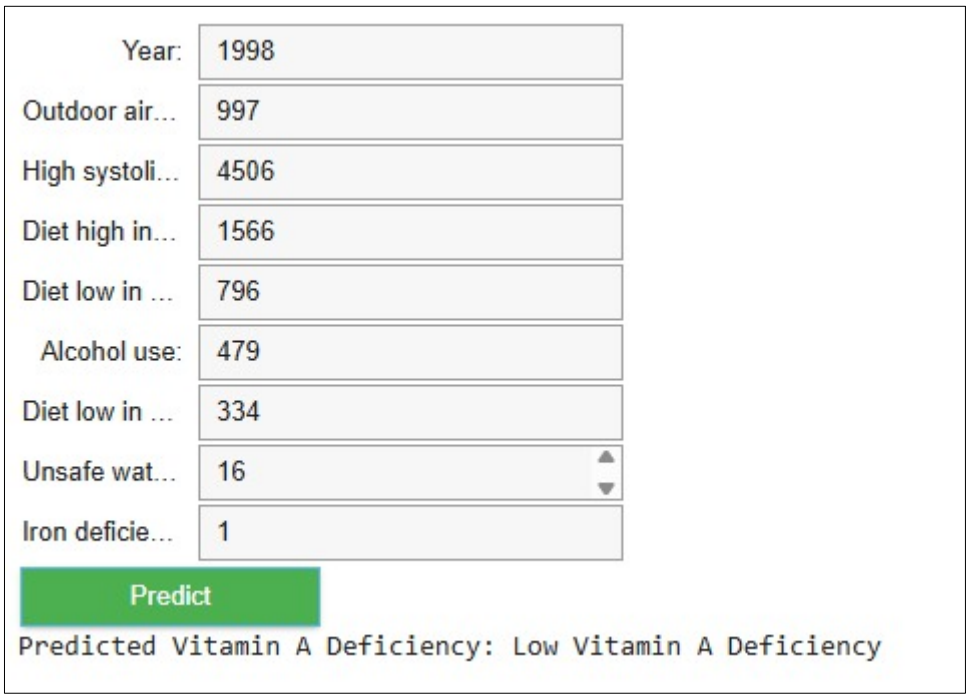

Figure 22: Prediction (Vitamin A Deficiency)

Figure 22 shows prediction output visualisation shows how well the model works at finding people who aren't getting enough vitamin A. Predictions that come true confirm that the model is ready to be used in real-life situations like clinical testing or mobile health diagnostics. You can plot the model's predictions against the real classes to see how they differ and find patterns. This graph shows that the model can be understood and that the output can be trusted, which completes the model evaluation cycle.

7. Conclusion

Using real-life clinical data, this study successfully created and tested a deep learning-based diagnostic system that can tell if women aren't getting enough iron and vitamin A. The tests showed that standard machine learning classifiers like SVM, KNN, and Naïve Bayes aren't very good at capturing the complex, nonlinear interactions between nutritional indicators. The proposed hybrid CNN–LSTM model, on the other hand, got much better results in terms of accuracy, precision, and F1-score across both datasets. The CNN layers did a good job of capturing spatial feature patterns, and the LSTM units modelled how things depend on each other in a sequence and in context. This made for a strong and flexible framework. The comparison test showed that the CNN–LSTM model consistently did better than CNN and LSTM models that worked alone, proving that it is better for clinical diagnostic uses. The model also showed stable convergence behaviour, minimalised training-validation loss divergence, and could be used in the real world by using tools for interpretability like confusion matrices and prediction visualisations. This study lays the groundwork for diagnostic tools that

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

are scalable and powered by AI that can be added to mobile health platforms and community-based screening programs, especially in places with few resources. In the future, researchers will look into how to make this model work with multi-nutrient prediction scenarios and how to use edge computing frameworks for real-time deployment in rural healthcare systems.

Reference

[1] T. P. T. Armand, H.-C. Kim, and J.-I. Kim, "Digital anti-aging healthcare: An overview of the applications of digital technologies in diet management," Journal of Personalized Medicine, vol. 14, p. 254, **2024**.

[2] L. O. Chaves, A. L. G. Domingos, D. L. Fernandes, F. R. Cerqueira, R. Siqueira-Batista, and J. Bressan, "Applicability of machine learning techniques in food intake assessment: A systematic review," Critical Reviews in Food Science and Nutrition, vol. 63, pp. 902–919, **2023**.

[3] T. P. T. Armand, O. Deji-Oloruntoba, S. Bhattacharjee, K. A. Nfor, and H.-C. Kim, "Optimizing longevity: Integrating Smart Nutrition and Digital Technologies for Personalized Anti-aging Healthcare," in Proceedings of the International Conference on Artificial Intelligence in Information and Communication (ICAIIC), Osaka, Japan, pp. 243–248, **2024**.

[4] C. Kosaraju, S. Chiruguri, N. Mohammad, B. T. Jetty, and G. Gali, "A model for analysis of diseases based on nutrition deficiency using random forest," in Proceedings of the International Conference on Communication and Electronics Systems (ICCES), Coimbatore, India, pp. 1527–1531, **2022**.

[5] F. Szabo de Edelenyi et al., "Prediction of the metabolic syndrome status based on dietary and genetic parameters, using Random Forest," Genes & Nutrition, vol. 3, pp. 173–176, **2008**.

[6] A. Izbassar and P. Shamoi, "Image-Based Dietary Assessment: A Healthy Eating Plate Estimation System," arXiv preprint, arXiv:2403.01310, **2024**.

[7] R. Kaur, R. Kumar, and M. Gupta, "Deep neural network for food image classification and nutrient identification: A systematic review," Reviews in Endocrine and Metabolic Disorders, vol. 24, pp. 633–653, **2023**.

[8] H. Zitouni, S. Meshoul, and C. Mezioud, "New contextual collaborative filtering system with application to personalized healthy nutrition education," Journal of King Saud University – Computer and Information Sciences, vol. 34, pp. 1124–1137, **2022**.

[9] D. Moher et al., "Preferred reporting items for systematic review and meta-analysis protocols (PRISMA-P) 2015 statement," Systematic Reviews, vol. 4, p. 1, **2015**.

[10] M. Woschank, E. Rauch, and H. Zsifkovits, "A review of further directions for artificial intelligence, machine learning, and deep learning in smart logistics," Sustainability, vol. 12, p. 3760, **2020**.

[11] D. Kirk, E. Kok, M. Tufano, B. Tekinerdogan, E. J. M. Feskens, and G. Camps, "Machine learning in nutrition research," Advances in Nutrition, vol. 13, pp. 2573–2589, **2022**.

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[12] J. Zhu and G. Wang, "Artificial intelligence technology for food nutrition," Nutrients, vol. 15, p. 4562, **2023**.
[13] M. van Erp et al., "Using natural language processing and artificial intelligence to explore the nutrition and sustainability of recipes and food," Frontiers in Artificial Intelligence, vol. 3, p. 621577, **2021**.
[14] M. Honda and H. Nishi, "Household nutrition analysis and food recommendation using purchase history," in Proceedings of the International Symposium on Industrial Electronics (ISIE), Kyoto, Japan, pp. 1–7, **2021**.
[15] P. Santhuja et al., "Intelligent personalized nutrition guidance system using IoT and machine learning algorithm," in Proceedings of the Second International Conference on Smart Technologies for Smart Nation (SmartTechCon), Singapore, pp. 250–254, **2023**.
[16] A. Maurya et al., "Chronic kidney disease prediction and recommendation of suitable diet plan by using machine learning," in Proceedings of the International Conference on Nascent Technologies in Engineering (ICNTE), Navi Mumbai, India, pp. 1–4, **2019**.
[17] D. Mogaveera, V. Mathur, and S. Waghela, "e-Health monitoring system with diet and fitness recommendation using machine learning," in Proceedings of the International Conference on Inventive Computation Technologies (ICICT), Coimbatore, India, pp. 694–700, **2021**.
[18] C. Iwendi et al., "Realizing an efficient IoMT-assisted patient diet recommendation system through machine learning model," IEEE Access, vol. 8, pp. 28462–28474, **2020**.
[19] R. Sookrah, J. D. Dhowtal, and S. D. Nagowah, "A DASH diet recommendation system for hypertensive patients using machine learning," in Proceedings of the International Conference on Information and Communication Technology (ICoICT), Kuala Lumpur, Malaysia, pp. 1–6, **2019**.
[20] C. A. Iheanacho and O. R. Vincent, "Classification and recommendation of food intake in West Africa for healthy diet using Deep Learning," in Proceedings of the Information Technology for Education and Development (ITED), Abuja, Nigeria, pp. 1–6, **2022**.
[21] S. Mezgec, T. Eftimov, T. Bucher, and B. K. Seljak, "Mixed deep learning and natural language processing method for fake-food image recognition and standardization to help automated dietary assessment," Public Health Nutrition, vol. 22, pp. 1193–1202, **2019**.
[22] G. K. Folson et al., "Validation of mobile artificial intelligence technology-assisted dietary assessment tool against weighed records and 24-hour recall in adolescent females in Ghana," The Journal of Nutrition, vol. 153, pp. 2328–2338, **2023**.
[23] J. Shi, Q. Han, Z. Cao, and Z. Wang, "DeepTrayMeal: Automatic dietary assessment for Chinese tray meals based on deep learning," Food Chemistry, vol. 434, p. 137525, **2024**.
[24] Y. Lu et al., "goFOODTM: An artificial intelligence system for dietary assessment," Sensors, vol. 20, p. 4283, **2020**.
[25] B. N. Limketkai, K. Mauldin, N. Manitius, L. Jalilian, and B. R. Salonen, "The age of artificial intelligence: Use of digital technology in clinical nutrition," Current Surgery Reports, vol. 9, p. 20, **2021**.

Received: July 25, 2025

Volume 38 No. 1s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[26] A. Salinari, M. Machì, Y. A. Diaz, D. Cianciosi, Z. Qi, B. Yang, M. S. F. Cotorruelo, S. G. Villar, L. A. D. Lopez, M. Battino, et al., "The application of digital technologies and artificial intelligence in healthcare: An overview on nutrition assessment," Diseases, vol. 11, p. 97, **2023**.
[27] P. Singer, E. Robinson, and O. Raphaeli, "The future of artificial intelligence in clinical nutrition," Current Opinion in Clinical Nutrition and Metabolic Care, vol. 27, pp. 200–206, **2024**.
[28] Iron Deficiency Dataset, **2025**: https://www.kaggle.com/code/gautamchettiar/analysis-iron-deficiency-prediction-xgb-99-78/input
[29] Vitamin A Deficiency Dataset, **2025**: https://www.kaggle.com/code/gautamchettiar/analysis-iron-deficiency-prediction-xgb-99-78/input

Received: July 25, 2025