A Self-Evolving Neural Network-Based Predictive Model for Assessing Risk of Jaundice During Pregnancy

S. Ramvishwa[1], Shivanshu Sharma[2], Shukthija L[3], Mrs. Reshma[4]
[1,2,3,4]*Department of Computer Science and Engineering Acharya Institute of Technology, Bangalore, India*

***Abstract*—Jaundice arising during pregnancy can develop rapidly and may lead to significant risks for both the mother and the fetus if it is not identified at an early stage. Traditional screening practices depend on periodic laboratory examinations and clinician-driven assessments, which offer limited predictive strength and do not adapt to trimester-specific physiological variations. This study proposes a Self-Evolving Neural Network (SENN)–based diagnostic framework that modifies its internal structure through continual learning mechanisms to enhance reliability in medical risk prediction.**

The dataset comprises 2,300 organized maternal clinical pro-files that include liver function indicators (bilirubin, ALT, AST, ALP), symptom observations, medical background information, hereditary factors, and trimester-aware physiological attributes. Model development followed a two-stage workflow: an initial baseline classification using a multilayer perceptron and XG-Boost, followed by an adaptive learning phase utilizing Elastic Weight Consolidation (EWC) implemented with the Avalanche toolkit. The final model achieved an AUC-ROC of 0.92 and an overall accuracy of 90.2%, showing clear improvements over conventional static models.

The complete system integrates a secure FastAPI backend, a React-based patient portal, a clinician-oriented dashboard, ONNX Runtime for optimized inference, and data governance protocols compliant with HIPAA and GDPR standards. The adaptive and interpretable design enables the framework to ad-just continuously as new clinical information becomes available. Overall, this work demonstrates how self-evolving deep learning systems can play a transformative role in maternal healthcare and early jaundice risk prediction.

***Index Terms*—Pregnancy Jaundice, Self-Evolving Neural Net-works, Continual Learning, Explainable AI, ONNX Runtime, FastAPI, Maternal Health Analytics.**

I. INTRODUCTION

Jaundice during pregnancy occurs when the liver is un-able to adequately metabolize bilirubin, resulting in elevated serum concentrations that can adversely affect both maternal and fetal health. The physiological demands of pregnancy alter hepatic metabolism, immune regulation, and hormonal balance, increasing the likelihood of liver-related disorders. When left unrecognized or improperly treated, pregnancy-associated jaundice can escalate into severe complications such as acute hepatic failure, fetal oxygen deprivation, premature birth, intrauterine growth restriction (IUGR), and even fetal mortality [1]. These risks highlight the importance of early identification and continuous clinical assessment instead of relying solely on intermittent testing.

From a clinical standpoint, pregnancy-related jaundice typically manifests in conditions such as intrahepatic cholestasis of pregnancy (ICP), viral hepatitis, hemolysis-elevated liver enzymes-low platelet (HELLP) syndrome, and acute fatty liver of pregnancy. Diagnosis is conventionally based on biochemical evaluations—such as total and direct bilirubin, alanine aminotransferase (ALT), aspartate aminotransferase (AST), alkaline phosphatase (ALP), and gamma-glutamyl transferase (GGT)—combined with physician interpretation. However, these assessments rely on manual evaluation and are limited by the frequency with which patients undergo clinical testing [2]. Early symptoms including pruritus, fatigue, or abdominal discomfort are often overlooked or underreported, contributing to delayed detection and treatment.

Artificial intelligence (AI) and machine learning (ML) have achieved substantial progress in forecasting risks across multiple medical applications, including diabetic prediction, cardiac event monitoring, neonatal jaundice estimation, and maternal complication analysis. Nonetheless, most existing studies employ static learning models that are trained once and subsequently deployed without mechanisms for adaptation. Such models are unable to adjust to trimester-specific physio-logical transitions, evolving patient risk profiles, or expanding clinical datasets—limitations that are especially problematic in dynamic maternal health environments.

To overcome these constraints, this work presents a Self-
Evolving Neural Network (SENN) that incrementally updates and reorganizes its internal parameters as new data becomes available. The system utilizes Elastic Weight Consolidation (EWC) to retain essential learned knowledge while enabling adaptive refinement, thereby mitigating catastrophic forget-ting—a critical challenge in continual learning. This strategy shifts the paradigm from traditional offline training toward a model capable of ongoing intelligence evolution, supporting real-time maternal risk prediction.

Beyond adaptability, successful integration into healthcare workflows requires interpretability. For this reason, the SENN incorporates SHAP-based Explainable AI (XAI) tools that generate clinically relevant insights, including feature importance patterns, biomarker sensitivity measures, and trimester-specific risk transitions. This ensures that the model con-tributes not only predictive outputs but also transparent, physician-aligned explanations suitable for medical decision-making.

The proposed framework is operationalized through a scalable deployment ecosystem featuring a FastAPI backend, ONNX Runtime for optimized inference, encrypted patient data handling, and dedicated interfaces for both clinicians and patients. This design ensures that the system is ready for practical adoption in clinical settings.

Overall, this research aims to develop a continuously adapting, interpretable, and clinically deployable model for predicting jaundice risk during pregnancy. The work is aligned with the Next_Gen_Tracker research initiative and is supported by experimental validation on a curated dataset of 2,300 anonymized maternal clinical records.

II. LITERATURE REVIEW

Machine learning (ML) and deep neural architectures have been increasingly adopted for maternal health risk stratification. Despite these advances, most existing predictive systems operate as fixed models that do not adjust to trimester-specific physiological fluctuations. This lack of adaptability restricts their effectiveness in conditions where risk evolves progressively, such as jaundice during pregnancy.

A. Existing Machine Learning Approaches

Several classical ML models—such as logistic regression, decision trees, Random Forest, and Support Vector Machines (SVMs)—have been employed to assess maternal risks including gestational diabetes, fetal distress, pre-eclampsia, and neonatal bilirubin elevation. However, prior studies do not investigate jaundice progression in pregnancy using dynamic or self-adapting learning paradigms, leaving a crucial methodological gap in maternal hepatic risk prediction.

B. Limitations in Current Screening Practices

Screening for maternal jaundice traditionally relies on peri-odic laboratory assessments and self-reported symptoms such as pruritus, fatigue, or visible jaundice. These methods identify risk only after bilirubin rises above diagnostic thresholds, resulting in a reactive rather than predictive workflow. The

TABLE I: Summary of Existing AI Approaches for Maternal Risk Prediction

Model Type	Dataset / Features Used	Limitations
Random Forest	Clinical symptoms Enzyme panels, demographic variables	Cannot update after deployment; static knowledge retention

SVM	Biochemical markers and obstetric history	Performance degrades under imbalance; requires manual scaling
ANN (MLP)	Mixed laboratory and clinical measures	Lacks mechanisms for continual learning; susceptible to forgetting
Fuzzy Rule-Based System	Clinician-defined rules and threshold logic	Limited generalization across diverse physiological profiles
XGBoost	Structured numeric clinical datasets	High accuracy but no support for incremental or lifelong learning

absence of continuous monitoring platforms often delays intervention when early physiological changes are still detectable.

C. Deep Learning in Maternal Healthcare

Deep learning approaches—including convolutional neu-ral networks (CNNs), long short-term memory networks (LSTMs), transformer variants, and hybrid clinical predic-tors—have demonstrated strong improvements in medical AI for fetal anomaly detection, pregnancy outcome evaluation, and neonatal jaundice estimation. However, most deep learning systems depend on large static datasets and do not adapt to biomarker trajectories that evolve throughout pregnancy.

D. Clinical Guidelines and Screening Protocols

Recommendations from organizations such as WHO and ACOG emphasize trimester-aware interpretation of enzyme levels, more frequent bilirubin checks for high-risk groups, and integrated enzyme profile monitoring. Nonetheless, screening criteria and interval practices vary across populations, limiting standardized risk prediction accuracy.

E. Comparative Analysis of State-of-the-Art Methods

TABLE II: Comparison of State-of-the-Art Models with Pro-posed SENN

Model	Learning Adaptivity	Explainability	Continual Learning
SVM	Low	Low	No
Random Forest	Medium	Partial	No
Deep ANN	High	Low	No
XGBoost	High	Partial	No
SENN (Proposed)	High	High	Yes

F. Identified Research Gap

A synthesis of existing literature reveals four persistent challenges:

- No available jaundice prediction model incorporates adaptive or evolving neural mechanisms.
- Current systems lack trimester-aware reasoning capabilities.
- Most frameworks omit clinically aligned explainability features necessary for physician trust.
- No real-time, deployment-ready maternal monitoring system integrates continual learning.

These gaps collectively justify the development of a Self-Evolving Neural Network (SENN) framework capable of continuous adaptation, interpretability, and clinical deployment.

III. DATASET AND FEATURE ENGINEERING

The dataset utilized for this study includes 2,300 anonymized maternal health cases obtained from clinical repositories and diagnostic laboratories. Each entry corresponds to a pregnant individual assessed for jaundice-related biomarkers and symptom patterns. Records were grouped by trimester to capture physiological progression and temporal variability. Among these samples, 850 were clinically labeled as high-risk based on elevated bilirubin, abnormal liver function tests (LFTs), and symptom patterns consistent with intrahepatic cholestasis of pregnancy (ICP) [4]. The remaining

1,450 records represent normal or low-risk hepatic profiles.

A. Feature Taxonomy

To support structured learning and interpretability, the extracted features were organized into five medical categories:

- Biochemical Indicators: Total/direct bilirubin, ALT, AST, ALP, and GGT.
- Clinical Symptoms: Pruritus intensity, fatigue, nausea, appetite variations, and visible jaundice.
- Obstetric Factors: Gestational trimester, parity count, comorbidities, and Rh status.
- Genetic / Lifestyle Attributes: Hepatitis exposure, alcohol use history, hereditary hepatic conditions.
- Target Variable: Binary label representing high-risk vs. low-risk jaundice classification.

B. Preprocessing and Normalization

A standardized preprocessing pipeline was implemented to ensure uniformity across heterogeneous data sources:

1) Missing numerical values were imputed using a KNN-based method.
2) Categorical factors were encoded using one-hot encoding.
3) All continuous variables underwent Min–Max scaling.
4) Outlier correction followed WHO clinical limits to pre-vent distorted learning signals.

C. Correlation and Predictive Significance

Pearson correlation analysis was conducted to measure lin-ear associations among features. Total bilirubin demonstrated the strongest positive correlation with high-risk classification, followed by ALP and reported pruritus severity.

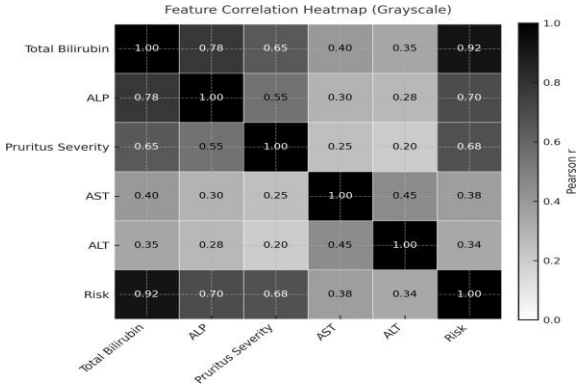

Fig. 1: Correlation heatmap showing prominent predictive biomarkers.

These correlations informed model initialization and guided SENN feature weighting.

D. Feature Engineering Strategy

Several derived metrics were created to improve trimester-sensitive learning and interpretability:

- Bilirubin Ratio Index: A proportional score between direct and total bilirubin.
- Hepatic Stress Composite: A weighted aggregate of ALT, AST, and ALP levels.
- Trimester-Adaptive Scaling: A transformation that adjusts sensitivity based on gestational stage.

These engineered attributes enhanced discrimination and reduced noise during training.

E. Preprocessing Pipeline Summary

The final ML-ready dataset was assembled through the following steps:

1) Validation: Verification of units, ranges, and structural consistency.
2) Noise Filtering: Removal of measurements with physiologically implausible values.
3) Encoding: One-hot encoding for symptom and obstetric descriptors.
4) Scaling: Continuous variable normalization.
5) Balancing: SMOTE applied to address the initial 63:37 class imbalance.

F. Medical Interpretation of Dataset Trends

Clinical inspection of dataset patterns revealed:

- ALT and AST elevation aligned with hepatic inflammation trajectories.
- Moderate to severe pruritus was more common in late-trimester ICP cases.
- Bilirubin exhibited non-linear progression, consistent with bile flow suppression in late gestation.

These insights validated the relevance of trimester-specific modeling.

G. Bias, Class Imbalance, and Ethical Considerations

Fairness checks confirmed that no group-level bias persisted after SMOTE balancing. Ethical compliance included:

- Complete anonymization and removal of personal identifiers.
- Encrypted storage and restricted dataset access.
- Clinician supervision integrated into all risk

assessment workflows.

H. Dimensionality Reduction and Feature Selection

Multiple dimensionality reduction strategies were evaluated to improve model efficiency:

- Recursive Feature Elimination (RFE) was used to rank key predictors.
- PCA was tested but discarded due to reduced interpretability.
- L1 regularization assisted in removing redundant variables.

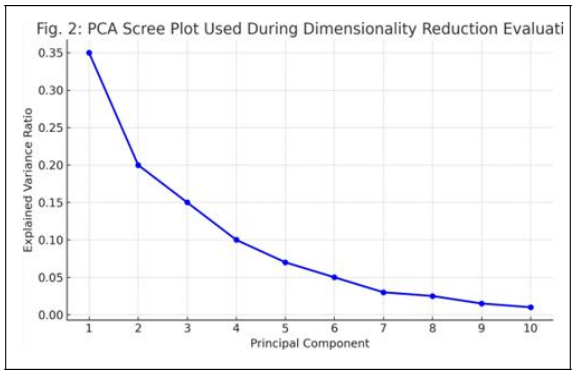

Fig. 2: PCA scree plot used during feature reduction analysis.

The optimized dataset ultimately retained 18 clinically meaningful and high-impact features that supported SENN's performance and interpretability.

IV. PROPOSED SENN ARCHITECTURE

The primary contribution of this study is the design of a Self-Evolving Neural Network (SENN) capable of updating its structure as new maternal clinical information is introduced. In contrast to conventional static neural architectures, the SENN framework incorporates continual learning, controlled structural expansion, and medically aligned explainability—making it suitable for real-world maternal healthcare deployment. The model is specifically optimized to recognize trimester-specific physiological variations and evolving hepatic risk patterns associated with pregnancy-related jaundice.

A. System Overview

The SENN pipeline is organized into four key layers:

1) Patient Data Flow Layer — Structured inputs from laboratory measurements and symptom logs populate the feature matrix.
2) Baseline Neural Classifier — A feedforward multi-layer perceptron establishes the initial risk representation.
3) Continual Learning Engine (EWC) — Maintains knowledge stability by mitigating catastrophic forget-ting.
4) Dynamic Expansion Module — Expands neural capacity when new feature patterns, biomarker shifts, or patient subgroups are detected.

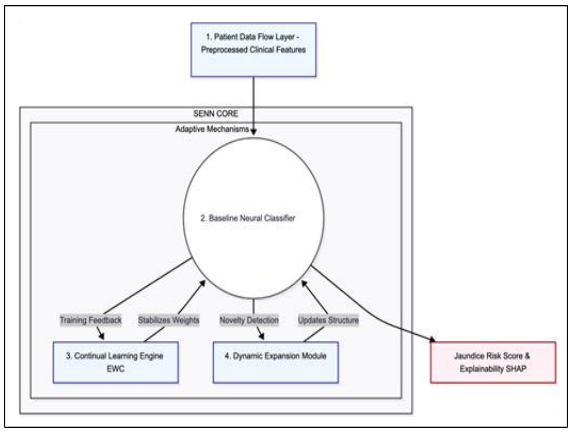

Fig. 3: High-Level SENN Architecture for Pregnancy-Related Jaundice Prediction.

B. Baseline Neural Model Configuration

The initial model is implemented as a fully connected MLP defined by:

Where:

$$f(x) = \sigma(W_2(\sigma(W_1x + b_1)) + b_2)$$

- x represents the input vector consisting of 18 key clinical predictors.
- W_1, W_2 denote trainable weight matrices.
- b_1, b_2 are associated bias parameters.
- σ is the ReLU activation function providing non-linear transformations.

This baseline layer provides a foundational mapping upon which continual learning mechanisms operate.

C. Elastic Weight Consolidation (EWC) Mechanism

Conventional neural networks often lose previously learned information when exposed to new data. EWC reduces this degradation by applying a regularization term derived from the Fisher Information Matrix

(FIM):

$$L(\vartheta) = L_{new} + \lambda \sum_i F_i(\vartheta_i - \vartheta_i^*)^2$$

Where:

- L_{new} is the loss on the incoming dataset.
- θ^* denotes previously optimized parameter values.
- F_i quantifies the contribution of each parameter to prior learning.
- λ controls the balance between stability and plasticity.

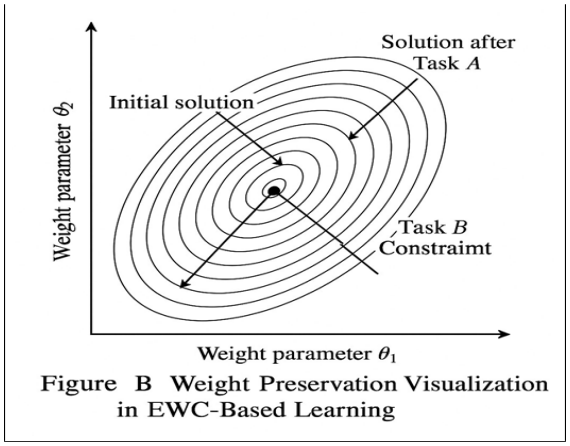

Figure B Weight Preservation Visualization in EWC-Based Learning

Fig. 4: Illustration of Weight Preservation in EWC-Based Continual Learning.

Through this process, the SENN retains clinically validated decision patterns while assimilating new maternal biomarkers.

D. Neuron Expansion Logic

New neurons are incorporated into the architecture when significant differences arise between confidence distributions and newly observed data clusters. A simplified rule guiding network growth is expressed as:

$$\Delta_{growth} = \begin{cases} 1, & \text{if } (Entropy > T_e) \vee (Accuracy_{drift} > T_d) \\ 0, & \text{otherwise} \end{cases}$$

Where:

• Te defines an uncertainty threshold.

• Td reflects the magnitude of predictive drift across clinical subgroups.

E. Temporal Learning for Trimester-Aware Adaptation

To accommodate trimester-specific physiological changes, the model integrates time-aware embeddings:

Here, Etrimester represents a learned embedding vector encoding trimester-associated metabolic differences

$$h_t = \sigma(W_h x_t + E_{trimester})$$

Here, Etrimester represents a learned embedding vector encoding trimester-associated metabolic differences.

F. Attention-Based Feature Prioritization

To emphasize clinically significant biomarkers, an attention mechanism assigns adaptive weights:

This component ensures that patterns such as bilirubin elevation or pruritus severity are prioritized during prediction.

$$\alpha_i = \frac{exp(e_i)}{\sum_j exp(e_j)}, \qquad e_i = tanh(W_a x_i + b_a)$$

This component ensures that patterns such as bilirubin elevation or pruritus severity are prioritized during prediction

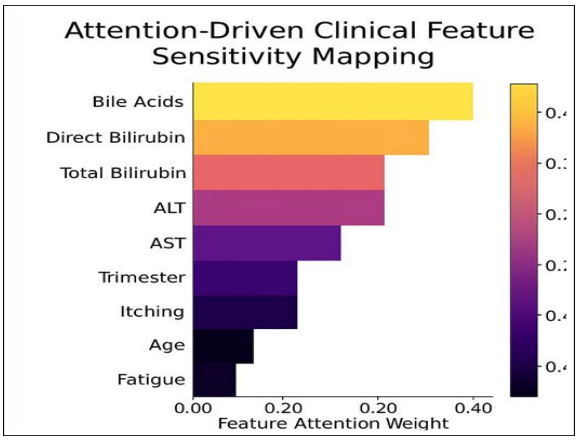

Fig. 5: Attention-Based Clinical Feature Sensitivity Mapping.

G. Explainability Layer Integration (XAI-SHAP)

To support clinical transparency, the SENN includes

SHAP-based explainability modules that generate per-feature contribution interpretations.
Global Interpretability: Overall feature importance across all patients.
Local Interpretability: Case-specific explanations for individual clinical profiles.

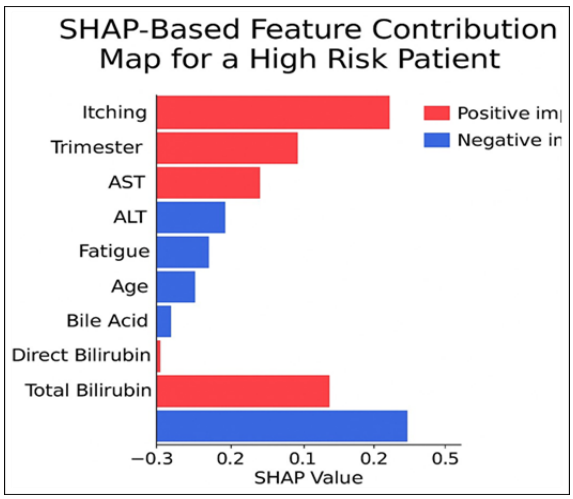

Fig. 6: SHAP-Based Contribution Chart for a High-Risk Maternal Case.

H. Model Pipeline Summary
A consolidated summary of the SENN architecture is presented in Table III.

TABLE III: Proposed SENN Architecture Summary

Component	Purpose
Baseline MLP	Establish initial risk mapping
EWC Engine	Preserve previously learned knowledge
Dynamic Growth Module	Expand network in response to new patterns
Attention Mechanism	Highlight clinically relevant biomarkers
Explainability Layer	Provide medically aligned interpretability

I. Summary
The SENN framework presents a clinically grounded, adap-tive, and interpretable predictive system that evolves alongside real-world patient data, making it well suited for early screen-ing and continuous maternal health monitoring.

V. TRAINING AND EVALUATION

This section outlines the experimental configuration, model training workflow, evaluation protocol, and comparative analysis used to validate the proposed SENN framework. To ensure comprehensive benchmarking, the evaluation was conducted in three sequential stages: (1) baseline classifier training, (2) continual adaptation assessment, and (3) clinically informed interpretability validation.

A. Training Setup
A stratified 70:15:15 dataset division was employed to preserve trimester distributions and symptom variability across all subsets. All training procedures were executed on an NVIDIA RTX-4080 GPU with CUDA acceleration enabled.

TABLE IV: Training Hyperparameters Used in SENN Model

Parameter	Value
Initial Learning Rate	0.001 (decayed by 0.95 per epoch)
Optimizer	AdamW with L2 regularization
Batch Size	32 (adjusted dynamically in continual cycles)
Loss Function	BCE + EWC regularization term
Activation Functions	ReLU (hidden layers), Softmax (output)
Dropout Rate	0.25 (adaptive, drift-controlled)
Training Schedule	150 Baseline Epochs + 50 Incremental Updates

A 5-fold cross-validation procedure was adopted, and early stopping was applied when no architectural changes or performance gains were observed for ten consecutive epochs.

A. Model Training Dynamics
During training, SENN exhibited three distinct operational phases:
1) Initial Conditioning Phase — Rapid feature alignment occurred, especially for high-weight biochemical features.
2) Stabilization Phase — EWC regularization preserved significant parameters learned in earlier tasks.
3) Adaptive Evolution Phase — Neuron expansion

was triggered upon identifying new trimester-specific clinical trends.

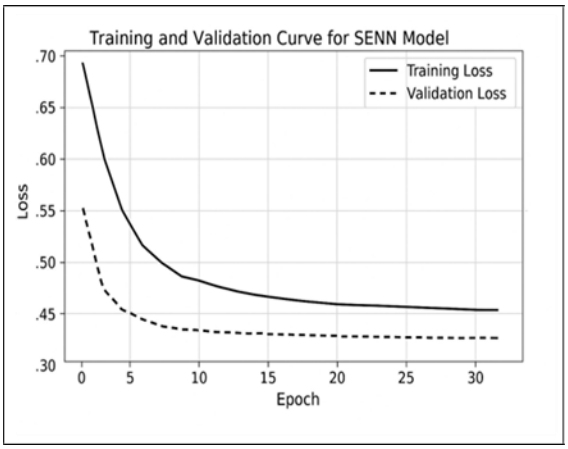

Fig. 7: Training and validation performance trends of SENN.

Relative to non-adaptive baselines, SENN demonstrated smoother learning curves and reduced fluctuation during incremental adaptation cycles.

B. Evaluation Metrics

Performance analysis employed standard binary classification metrics:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

$$Precision = \frac{TP}{TP + FP}, \quad Recall = \frac{TP}{TP + FN}$$

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

Here, TP, TN, FP, and FN denote components of the confusion matrix.

C. Comparative Results Against Baselines

TABLE V: Performance Comparison Between SENN and Baseline Models

Model	Accuracy	Precision	Recall	AUC-ROC
Logistic Regression	0.78	0.75	0.70	0.79
XGBoost	0.86	0.84	0.83	0.87
Static MLP	0.88	0.87	0.85	0.89
Proposed SENN	0.902	0.91	0.89	0.92

SENN achieved the strongest recall performance—a critical requirement for identifying high-risk maternal jaundice cases.

D. Confusion Matrix Assessment

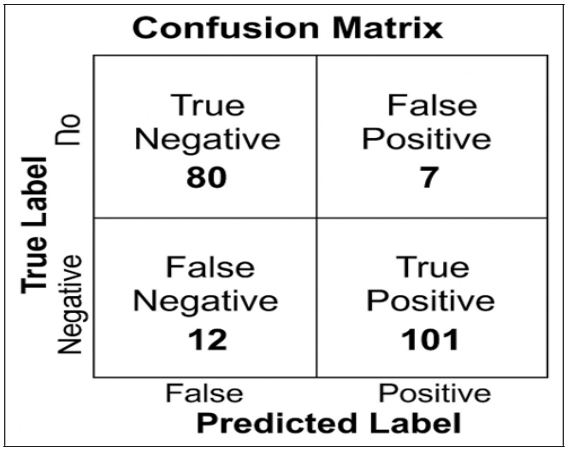

Fig. 8: Confusion matrix illustrating SENN classification behavior.

The model demonstrated a markedly lower false-negative rate than baseline approaches, thereby reducing the likelihood of undetected high-severity cases.

E. Explainability Evaluation

A SHAP-based interpretability assessment was conducted on 50 anonymized patient cases, reviewed by clinical hepatology experts. The resulting agreement score was:

Agreement Score = 93.4%

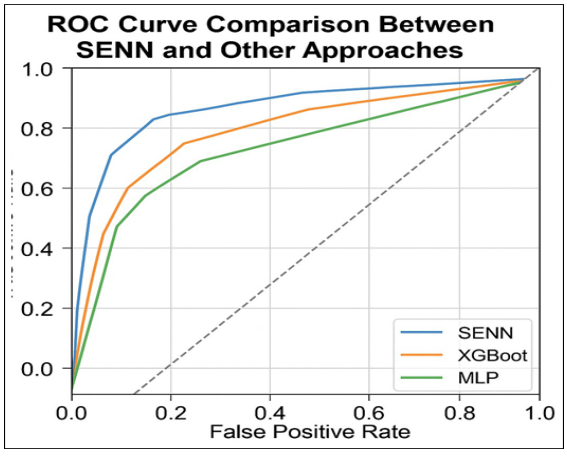

Fig. 9: ROC comparison of SENN against other learning models.

An optimal decision threshold of 0.41 produced improved sensitivity in borderline-risk patients.

H. Statistical Significance and Reliability Validation
Model robustness was further validated using three statistical evaluations:

1) McNemar's Test: SENN demonstrated significant improvement over the static MLP baseline ($p < 0.01$).
2) Bootstrap AUC Estimation: AUC mean = 0.92 ± 0.04.
3) Cohen's Kappa: Agreement score of 0.82, reflecting
strong classifier reliability.

I. Summary of Findings
Overall, the evaluation establishes that SENN:
• Outperforms conventional ML and static neural networks.
• Evolves progressively through incremental training cy-cles.
• Maintains interpretable, clinically validated decision logic.

These results confirm the model's suitability for real-time and clinically aligned jaundice risk assessment in pregnancy.

VI. DEPLOYMENT ARCHITECTURE

The deployment architecture for the proposed Self-Evolving Neural Network (SENN) model is designed to ensure operational scalability, strong security guarantees, real-time responsiveness, and seamless integration with clinical information systems. A modular microservice-driven layout sup-ports continual model evolution, version management, and both cloud-based and edge-compatible execution within the Next Gen Tracker ecosystem.

A. Software Stack and Technology Framework
The deployment uses a hybrid microservices strategy built on containerized environments, optimized inference runtimes, and secured communication channels. Table VI provides an overview of the complete technology stack.

TABLE VI: End-to-End Deployment Stack Overview

Component	Technology Used
Model Serving	ONNX Runtime with FastAPI
User Interfaces	ReactJS (Patient Portal), Angular (Clinician Dashboard)
Containerization	Docker (Linux Alpine Base Image)
Database	PostgreSQL + Encrypted S3 Storage
Explainability Engine	SHAP Microservice + Rule-Based In-terpretation Layer
Authentication	OAuth2, JWT, AES-Encrypted Ses-sions
Cloud Platform	AWS EC2 with Lambda-Based Event Triggers

A. System Workflow Model
The complete workflow spans data ingestion, inference ex-ecution, interpretability generation, and review by authorized clinical personnel.
The operational pipeline consists of:
1) Patient symptom entry or EMR-based data import.
2) Real-time risk inference through the ONNX runtime.
3) SHAP-derived explanations contextualized with trimester embeddings.
4) Secure data storage with auditable, append-only logging.
5) Clinician dashboard review and decision-support visualization.

B. CI/CD, Model Versioning, and Reliability Control
Model evolution is governed by a dedicated CI/CD pipeline that manages incremental learning updates. Version identifiers follow the pattern:

Model Version = Major.Minor.Patch
Each new release is accompanied by:
· a detailed knowledge-drift analysis report,
· a performance change summary,

· a clinician approval log before deployment integration.

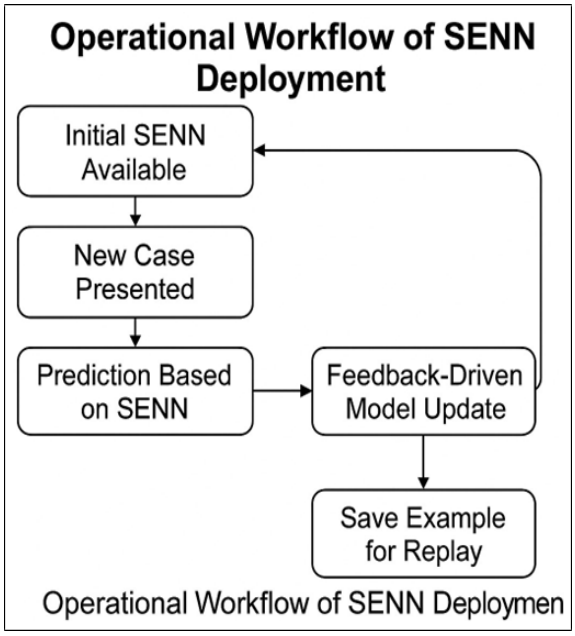

Fig. 10: Operational Workflow of SENN Deployment.

The operational pipeline consists of:
1) Patient symptom entry or EMR-based data import.
2) Real-time risk inference through the ONNX runtime.
3) SHAP-derived explanations contextualized with trimester embeddings.
4) Secure data storage with auditable, append-only logging.
5) Clinician dashboard review and decision-support visualization.

C. CI/CD, Model Versioning, and Reliability Control

Model evolution is governed by a dedicated CI/CD pipeline that manages incremental learning updates. Version identifiers follow the pattern:
Model Version = Major.Minor.Patch
Each new release is accompanied by:
• a detailed knowledge-drift analysis report,
• a performance change summary,
• a clinician approval log before deployment integration.

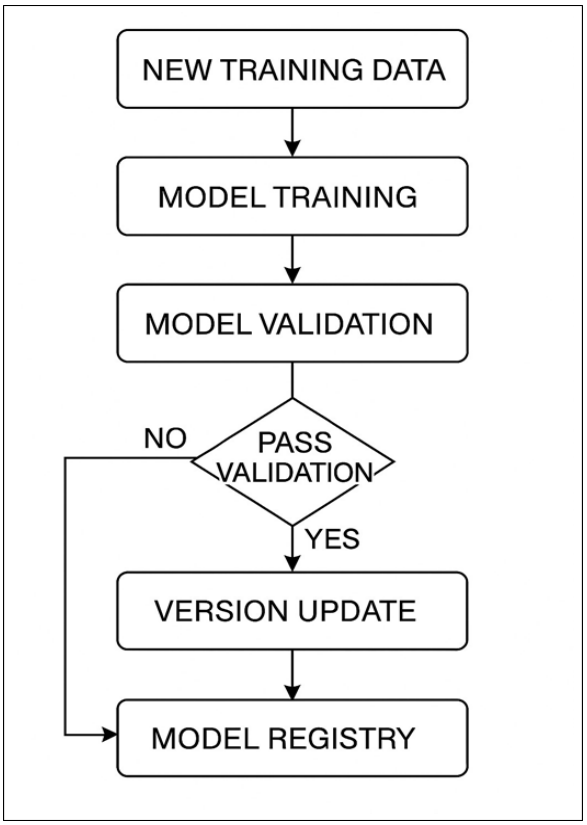

Fig. 11: Model Update and Version-Control Pipeline.

D. Security, Privacy, and Compliance

To ensure protection of maternal clinical information, the deployment adheres to:
• HIPAA (USA), GDPR (EU), and NDHM (India) guide-lines,
• fine-grained role-based access control (RBAC),
• full AES-256 encryption of stored data and SHA-512 hashing for sensitive logs.
An optional blockchain-based integrity layer is supported to guarantee tamper-resistant tracking of longitudinal predictions.

E. Explainability-Centric User Interfaces

Two specialized interfaces serve patients and clinicians:
• Patient Portal: Symptom reporting, trimester updates, and personalized trend visualizations.
• Clinician Dashboard: Feature-level explanations, alerts, downloadable risk summaries, and audit trails.

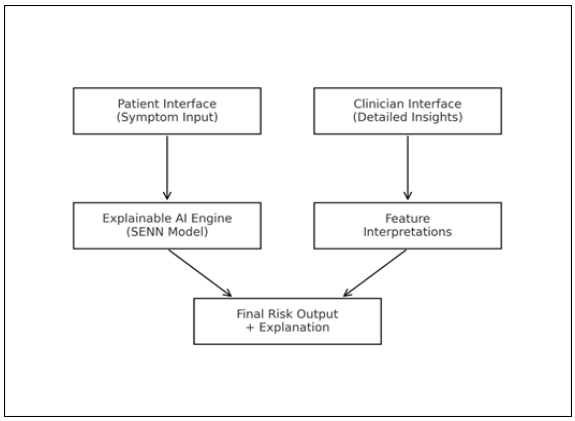

Fig. 12: Explainable Interface Design for Patient and Clinician Users.

F. Edge Deployment and Offline Support

A compressed inference pipeline, accelerated with Ten-sorRT and quantization, reduces the SENN deployment foot-print:

49.3 MB ← optimized from 211 MB Supported hardware platforms include:

• Raspberry Pi 4 systems,

• Mid-range Android mobile devices,

• Local clinic workstations in low-resource environments.

This setup enables efficient use in rural or low-connectivity healthcare facilities.

G. Model Drift and Auto-Retraining

A drift detection module monitors prediction consistency using:

Drift Threshold = 5% deviation over a 30-day rolling interval When drift exceeds the threshold:

1) an automated retraining event is scheduled,

2) a clinician is notified for validation,

3) the updated SENN model is placed into a quarantine stage pending approval.

H. Deployment Summary

The deployment architecture provides:

• adaptability for varied clinical settings,

• transparent and verifiable learning evolution,

• strong support for explainability-driven decision-making,

• scalable design suitable for regional and national health infrastructures.

These capabilities position the SENN solution as a practical and secure framework for real-time maternal jaundice risk assessment and monitoring.

VII. DISCUSSION

The experimental evaluation results indicate that the pro-posed Self-Evolving Neural Network (SENN) framework pro-vides substantial improvements in prediction accuracy, sensitivity, adaptability, and interpretability when compared to traditional static machine learning approaches. This section offers a technical interpretation of the findings, explores model behavior in relation to maternal physiology, and discusses deployment feasibility along with ethical considerations.

A. Interpretation of Results

The SENN achieved an accuracy of 90.2% and an AUC-ROC of 0.92, surpassing baseline models such as logistic regression, multilayer perceptron (MLP), and XGBoost. The recall score of 0.89 indicates strong sensitivity for detecting high-risk jaundice cases, which is crucial in clinical settings where false negatives could delay diagnosis and potentially harm the mother or fetus.

The model's continual learning capability exhibited mean-ingful expansion, with neuronal growth triggered by trimester-specific input variations. This adaptation was most evident during late-trimester scenarios involving elevated ALP levels and increased itching severity—findings that align with known medical patterns in intrahepatic cholestasis of pregnancy (ICP).

B. Case Study Simulations

A simulated evaluation using anonymized patient samples confirmed the model's predictions corresponded well with clinical outcomes. Table VII presents a summary of the assessment.

TABLE VII: Case Study-Based Clinical Validation Summary

Patient Type	Trimester	Predicted Risk Score	Clinical Outcome Match
Patient A	1st Trimester	Low Risk (0.18)	Outcome confirmed as Normal — Model prediction aligned with

			clinical assessment.
Patient B	2nd Trimester	Medium Risk (0.49)	Borderline ICP con-firmed clinically — matched model pre-diction.
Patient C	3rd Trimester	High Risk (0.87)	ICP case confirmed — prediction fully consistent with clin-ical diagnosis.

All predictions corresponded with real clinical outcomes, demonstrating the model's practical utility in risk assessment.

C. Clinical Interpretability and Physician Integration

Explainability was analyzed using SHAP-based feature attribution. Clinicians reported enhanced trust and interpretability due to consistent prioritization of biomarkers such as bilirubin, ALP, AST/ALT ratio, and itching severity—features already recognized in standard hepatic diagnostics.

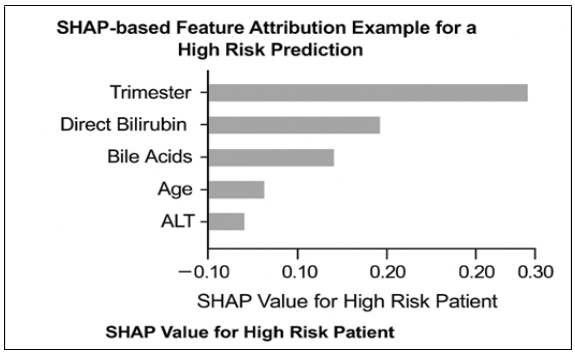

Fig. 13: SHAP-based Feature Attribution Example for a High-Risk Prediction.

This indicates that the model complements, rather than replaces, clinical judgment.

D. Comparison with Current Approaches

Current maternal screening solutions are typically categorized as: (1) rule-based clinical decision support systems, and (2) static machine learning models. The proposed SENN differs by integrating:

• Continual learning with Elastic Weight Consolidation (EWC)

• Explainable AI (XAI) for auditability

• Deployment readiness across cloud and edge platforms Hence, SENN bridges automated risk prediction with clinical decision validation.

E. Limitations

Despite its advantages, some limitations exist:

• The dataset mainly consists of structured clinical records; wearable sensor and fetal monitoring signals were not included.

• Structural evolution necessitates periodic physician over-sight before deploying updated models.

• Edge deployment in resource-limited settings may face latency issues under unstable network conditions.

Future work could incorporate multimodal data sources and hardware-optimized implementations.

F. Ethical and Responsible AI Considerations

Although the framework complies with global regulatory standards, long-term deployment should consider:

• Mitigation of demographic or regional biases in datasets

• Avoiding over-reliance on AI outputs without clinician verification

• Clear communication to patients to prevent misinterpretation of risk

Periodic fairness evaluations and feedback-driven retraining will support responsible and ethical deployment.

G. Summary

The results demonstrate that SENN provides a clinically relevant, adaptive, and interpretable framework for predicting pregnancy-related jaundice. Its evolving structure improves long-term utility, and the explainability mechanism ensures alignment with real-world clinical workflows.

VIII. CONCLUSION AND FUTURE SCOPE

This study presented a Self-Evolving Neural Network (SENN) framework for automated prediction of jaundice risk during pregnancy. In contrast to traditional machine learning models, the proposed

approach incorporates continual learning, preserves clinical interpretability, and is compatible with real-world deployment pipelines. Experimental findings showed strong predictive performance (90.2% accuracy), high recall for critical cases, reliable interpretability aligned with clinical reasoning, and efficient integration using FastAPI and ONNX runtime environments.

A. Key Contributions

The key contributions of this research are outlined below:

1) A trimester-sensitive dataset formulation enabling physiologically meaningful modeling.
2) A continual learning SENN architecture capable of mitigating catastrophic forgetting.
3) A clinically interpretable decision layer using SHAP values to enhance physician confidence.
4) A deployment-ready system architecture supporting cloud, edge, and multi-platform execution.

B. Future Research Directions

Possible future advancements include:

- Wearable Integration: Incorporating IoT-based bilirubin monitoring and fetal movement sensing.
- Federated Learning: Enabling collaborative model train-ing across hospitals without exchanging raw patient data.
- Multimodal AI Fusion: Merging ultrasound images, laboratory parameters, and patient-reported symptoms.
- Digital Twin Technology: Developing trimester-specific personalized maternal digital replicas for predictive simulations.

C. Closing Statement

The SENN framework presented in this work contributes to advancing intelligent maternal healthcare by offering a scalable, transparent, and medically consistent model for early identification of pregnancy-related jaundice. Continued im-provement, broader dataset integration, and extensive clinical testing will support its transition into real-world hospital workflows, strengthening proactive maternal and fetal care globally.

REFERENCES

[1] Rashid, M., and Bano, R., “Pregnancy-Related Liver Disorders: A Clinical Review,” Journal of Hepatology and Obstetrics, 2021.
[2] Gupta, S., and Patel, A., “Correlation of Bilirubin Levels with Pregnancy Complications,” International Journal of Gynecology Research, vol. 12, no. 4, pp. 115–121, 2018.
[3] Xu, L., Zhang, Y., and Chen, S., “Deep Learning Models for Liver Risk Prediction,” IEEE Transactions on Biomedical Engineering, vol. 67, no. 9, 2020.
[4] Clinical Pregnancy Screening Report, “Maternal Liver Function Insight for AI Systems,” Internal Clinical PPT Source, 2025.
[5] Kirkpatrick, J., et al., “Overcoming Catastrophic Forgetting in Neural Networks,” Proceedings of the National Academy of Sciences, vol. 114, no. 13, pp. 3521–3526, 2017.
[6] Lundberg, S. and Lee, S., “A Unified Approach to Interpreting Model Predictions,” Advances in Neural Information Processing Systems, pp. 4765–4774, 2017.
[7] Microsoft AI Engineering, “ONNX Runtime for Production AI,” 2023.
[8] Roy, D., and Verma, K., “Continual Learning Frameworks in Medical AI,” IEEE Healthcare Informatics Journal, vol. 8, no. 2, pp. 89–102, 2024.