

ORIGINAL RESEARCH ARTICLE

An LSTM-Assisted Markov Model for Predicting Cancer Stage Transitions: A Simulation Study

Abdullahi Anas*, Aliyu Usman Kinafa and Aliyu Umar Shelleng

¹Department of Mathematical Sciences, Gombe State University, Gombe State, Nigeria**ABSTRACT**

Traditional Markov chain models are commonly used to represent cancer progression but are limited by the memoryless assumption and reduced ability to capture nonlinear temporal dependencies. This study proposes a hybrid LSTM-assisted Markov chain framework that integrates sequential deep learning with probabilistic state-transition modeling. A clinically informed simulation framework generated 30,000 synthetic longitudinal patient trajectories across five disease states (Carcinoma in situ, Early/Localized, Locally Advanced, Regionally Advanced, and Metastatic). Transition probabilities were parameterized to reflect realistic progression dynamics, including intensified forward transitions, rare backward transitions (<1.5%), and an absorbing metastatic state. Each sequence contained 5–15 time steps with 17 clinical features. Due to irreversible progression, advanced stages were moderately overrepresented. Data were stratified into training (64%, n=19,200), validation (16%, n=4,800), and independent test (20%, n=6,000) sets while preserving class proportions. Hyperparameters were tuned on the validation set. Performance uncertainty was estimated using 1,000 bootstrap resamples to compute 95% confidence intervals (CIs). On the independent test set, the hybrid model achieved an accuracy of 0.919 (95% CI: 0.912–0.926), precision of 0.883 (95% CI: 0.874–0.892), recall of 0.919 (95% CI: 0.912–0.926), and F1-score of 0.896 (95% CI: 0.887–0.905). Compared with a traditional Markov baseline, the hybrid framework demonstrated improved predictive stability and better discrimination of advanced disease states. These findings demonstrate methodological robustness within a simulation-based environment encoding realistic cancer progression patterns. However, results are derived solely from synthetic data and require validation on real-world longitudinal clinical datasets before clinical applicability can be established.

ARTICLE HISTORY

Received December 07, 2025

Accepted March 22, 2026

Published March 25, 2026

KEYWORDS

Cancer progression prediction, Markov Chain, Long Short-Term Memory (LSTM), LSTM-Assisted Markov model, comparative model performance.



© The Author(s). This is an Open Access article distributed under the terms of the Creative Commons Attribution 4.0 License [creativecommons.org](https://creativecommons.org/licenses/by/4.0/)

INTRODUCTION

Cancer progression is inherently nonlinear, temporally dependent, and structured around clinically defined disease states. Modeling such evolution requires approaches capable of representing both probabilistic state transitions and complex temporal dependencies. Traditional Markov chain models have been widely used in disease progression analysis due to their interpretability and compatibility with structured staging systems. However, their memoryless assumption where future transitions depend only on the current state limits their ability to capture long-range temporal dependencies and nonlinear interactions that characterize cancer evolution. Recent advances in hybrid modeling frameworks demonstrate that integrating statistical models with deep learning architectures can substantially improve predictive performance in health-related time series. For example, Jin et al. (2023) developed a hybrid ARIMA–LSTM model for epidemic forecasting, where ARIMA extracted linear trends and LSTM learned nonlinear residual structures, significantly outperforming standalone models in RMSE

and MAPE. Similarly, Zhang et al. (2022) proposed a three-layer ARIMA–EEMD–LSTM architecture for hand-foot-mouth disease prediction, achieving superior predictive accuracy (RMSE = 4.37, MAPE = 2.94, R² = 0.996) compared with conventional statistical and machine learning models. These studies illustrate that decomposing linear and nonlinear temporal components within hybrid structures enhances forecasting reliability.

Beyond epidemic forecasting, hybrid deep learning models have shown strong performance in cardiovascular and cancer diagnostics. Wang et al. (2023) combined Convolutional Neural Networks (CNN) with Bidirectional LSTM (BiLSTM) for heart disease prediction, effectively handling missing and imbalanced clinical data. In oncology, hybrid CNN–LSTM architectures have demonstrated remarkable classification performance. Lilhore et al. (2025) integrated CNN, BiLSTM, and EfficientNet-B0 for breast cancer detection, achieving 99.2% accuracy and outperforming VGG-16

Correspondence: Anas Abdullahi. Department of Mathematical Sciences, Gombe State University, Gombe State, Nigeria. ✉ anasabdullahi096@gmail.com.

How to cite: Anas, A., Kinafa, A. U. & Shelleng, A. U. (2026). An LSTM-Assisted Markov Model for Predicting Cancer Stage Transitions: A Simulation Study. *UMYU Scientifica*, 5(1), 229 – 242. <https://doi.org/10.56919/usci.2651.019>

and ResNet-50. Rastogi et al. (2024) proposed a Conv1D–LSTM model for automated breast cancer detection, reaching 99% accuracy and demonstrating computational efficiency for real-time clinical applications. Similarly, Kaddes et al. (2025) reported 99.90% accuracy using a CNN–LSTM model for mammographic classification, while Abohashish et al. (2025) showed improved diagnostic reliability in melanoma classification using a patch-based CNN–LSTM framework. Additionally, Sethi et al. (2023) combined LSTM with Deep Belief Networks (DBN) for prostate cancer detection, demonstrating that hybrid recurrent architectures enhance performance in gene expression-based classification tasks. While these studies collectively demonstrate the effectiveness of hybrid models in extracting linear, nonlinear, spatial, and sequential patterns from biomedical data, most focus on prediction or classification tasks rather than explicit disease state-transition modeling. Importantly, many deep hybrid models sacrifice interpretability by replacing structured probabilistic frameworks with fully neural architectures. In contrast, disease progression modeling—particularly in oncology—requires interpretable transition matrices aligned with clinical staging systems.

Methodological work by Sengupta et al. (2023) provides a closer alignment to this objective by integrating Hidden Markov Models (HMMs) with LSTMs in a joint architecture that preserves transition matrix interpretability while enhancing predictive accuracy through nonlinear temporal learning. However, applications of such hybrid interpretability-preserving frameworks to structured cancer stage progression remain limited. This reveals a critical research gap: the need for hybrid models that retain explicit Markovian transition structure while incorporating nonlinear, history-dependent learning mechanisms.

To address this gap, this study proposes an LSTM-assisted Markov chain framework for modeling cancer progression across five clinically ordered stages: Carcinoma in situ, Early/Localized, Locally Advanced, Regionally Advanced, and Metastatic. This ordered structure reflects widely accepted oncologic staging hierarchies and enables modeling of an absorbing metastatic state. Seventeen clinical features were incorporated to approximate multidimensional determinants of disease evolution, including demographic, tumor-specific, biomarker, and treatment-related factors. A simulation-based design was adopted to allow controlled methodological validation of the hybrid architecture. Simulation enables explicit encoding of clinically plausible transition dynamics progressively increasing forward transition probabilities, rare backward transitions, and metastatic absorption while avoiding common limitations of real-world datasets such as missing staging data, privacy constraints, and inconsistent follow-up intervals. The primary objective is therefore methodological validation of a hybrid interpretability-

preserving progression model rather than immediate clinical deployment.

MATERIALS AND METHOD

The study carried out a computational and simulation-based approach to develop and evaluate a Long Short-Term Memory (LSTM)-assisted Markov Chain model for predicting cancer progression. The methodology encompassed data simulation, model formulation, parameter estimation, training, validation, and performance evaluation.

2.1 Study Design and Data Simulation

A clinically-informed simulation framework was developed to generate synthetic longitudinal patient trajectories, enabling controlled validation of the hybrid model's ability to capture nonlinear temporal dependencies in cancer progression. For full reproducibility, random seeds were fixed (`np.random.seed(42)`, `tf.random.set_seed(42)`).

2.1.1 Simulation Algorithm

Patient trajectories were generated using a stochastic progression model:

- I. Each patient was initialized with baseline covariates (e.g., age, genetic risk) drawn from realistic distributions.
- II. Disease progression was modeled through discrete monthly time steps, with transition probabilities influenced by cumulative clinical history.
- III. At each time step, clinical features (e.g., biomarkers, tumor size) were updated conditional on the current disease stage and treatment history.

The output for each patient was a longitudinal sequence:

$$X_p = \{(S_t, F_t)\}_{t=1}^T \tag{1}$$

Where S_t is the disease stage at time t and F_t are the observed clinical features

2.1.2 Simulated Dataset Specification

Sample size: 30,000 independent patient sequences (determined through power analysis)

Sequence length: Variable between 5 and 15 months per patient (mean 10)

Features: 17 features per time point (demographics, biomarkers, tumor characteristics, treatment factors)

Disease stages: Five ordered stages: Stage 0 (Carcinoma in situ), Stage 1 (Early/Localized), Stage 2 (Locally Advanced), Stage 3 (Regionally Advanced), Stage 4 (Metastatic) - absorbing state

Stochastic noise: Inherent in all feature sampling via normal and Bernoulli distributions

Censoring: Variable sequence lengths simulate irregular follow-up (no explicit censoring modeled).

2.1.3 Handling of Variable Sequence Lengths

Sequences of varying lengths (5–15 time steps) were handled by truncating to a fixed lookback window of 6 time steps. For each time point $t \geq 6$, the input was the concatenation of features from time steps $t - 5$ to t . This approach preserves temporal information within the window while enabling batch processing. No missing data were simulated; all sequences were complete.

2.2 LSTM Component

The Long Short-Term Memory (LSTM) network is employed to model nonlinear temporal dependencies in patient trajectories. The LSTM updates are defined as follows:

Input gate

$$i_t = \sigma(W_i X_t + U_i h_{t-1} + b_i) \quad (2)$$

Forget gate

$$f_t = \sigma(W_f X_t + U_f h_{t-1} + b_f) \quad (3)$$

Output gate

$$o_t = \sigma(W_o X_t + U_o h_{t-1} + b_o) \quad (4)$$

Cell state update

$$c_t = f_t \odot c_{t-1} + i_t \odot \hat{c}_t \quad (5)$$

Hidden state

$$h_t = o_t \odot \tanh(c_t) \quad (6)$$

Where:

i_t, f_t , and o_t are the input, forget, and output gates respectively

h_t is the hidden state

c_t is the cell state

σ denotes the sigmoid activation function

\odot represents element-wise multiplication

W is the Weight Matrix

U is the Recurrent Weight Matrix

b Bias Vector

2.3 Markov Chain Component

Cancer stage transitions are modeled using a first-order Markov Chain characterized by a transition probability matrix:

$$p_{ij} = P(S_{t+1} = S_j | S_t = S_i) \quad (7)$$

The hidden representation h_t is transformed through a fully connected layer into a K -dimensional output by the given equation

$$z_t = W^{(out)} h_t + b^{(out)} \quad (8)$$

where

K is the number of cancer stages

The softmax activation converts this output into a probability distribution over all possible stages

$$q_t[j] = P(S_t = S_j | X_{1:t}) \quad (9)$$

$$= \frac{e^{z_{t,j}}}{\sum_{j=1}^k e^{z_{t,j}}} \quad (10)$$

The predicted stage is obtained as

$$\widehat{S}_{t+1} = \arg \max_j q_{t+1}[j] \quad (11)$$

Where The most probable disease stage at the next time step is determined.

2.4 Model Diagram: LSTM-Assisted Markov Chain Model

The proposed model integrates an LSTM network with a Markov chain to jointly capture nonlinear temporal patterns and probabilistic state transitions (Figure 1).

2.4.1 Model Formulation for the Developed LSTM-Assisted Markov Chain Model

The proposed model combines a Long Short-Term Memory (LSTM) network with a Markov chain process to predict the progression of cancer through multiple stages over time. Table 1 shows variable parameters and their description.

The formulation of the proposed model can be described as follows:

1. model begins with a generalized input vector X_n , representing clinical, biological, or temporal features relevant to cancer progression
2. The purpose of this compartment is to provide a temporal sequence of clinical data that reflects the evolving health condition of the patient.
3. The input sequence is passed into an LSTM (Long Short-Term Memory) network compartment, which is responsible for capturing both short-term and long-term dependencies within the data.
4. The LSTM uses gating mechanisms (input gate, forget gate, and output gate) to selectively retain or discard information as the sequence progresses.
5. The network outputs a hidden state vector, denoted as h_t , which summarizes the temporal dynamics of the disease up to time t .
6. This hidden state serves as the learned internal representation of the patient's disease trajectory.
7. The hidden state h_t is then fed into a fully connected layer to perform a linear transformation:

$$z = Wh_t + b \tag{12}$$

where W and b are the weight matrix and bias vector.

8. This transformation refines the extracted features, making them suitable for classification and probabilistic estimation in the recent stage.
9. The output of the fully connected layer is processed by a Softmax layer, which converts the transformed features into a probability distribution over the possible cancer stages.

10. The Softmax function produces

$$P(t|h_t) = [P(S_t = S_0|h_t), P(S_t = S_1|h_t), \dots, P(S_t = S_4|h_t), \tag{13}$$

where $S_0, S_1, S_2, S_3,$ and S_4 which correspond to the defined stages Carcinoma ,

Early/localized, Locally advanced, and Metastatic respectively.

11. This step quantifies the likelihood that a patient is in each stage at time t , given the learned temporal features.

12. The probability outputs from the Softmax layer are combined with a **Markov process** to model stage transitions over time.

13. Each transition between stages is characterized by a probability P_{ij} , defined as

$$P_{ij} = P(S_{t+1} = S_j | S_t = S_i) \tag{14}$$

Where $\sum_j P_{ij} = 1$

14. The **next state**, S_{t+1} , represents the predicted disease stage at the subsequent time step.

15. The **joint probability** represents the overall likelihood of transitioning from stage S_t to S_{t+1} given the learned hidden state h_t :

$$P(t|h_t, S_t, S_{t+1}) = P(S_{t+1}|h_t, S_t) \times P(t|h_t) \tag{15}$$

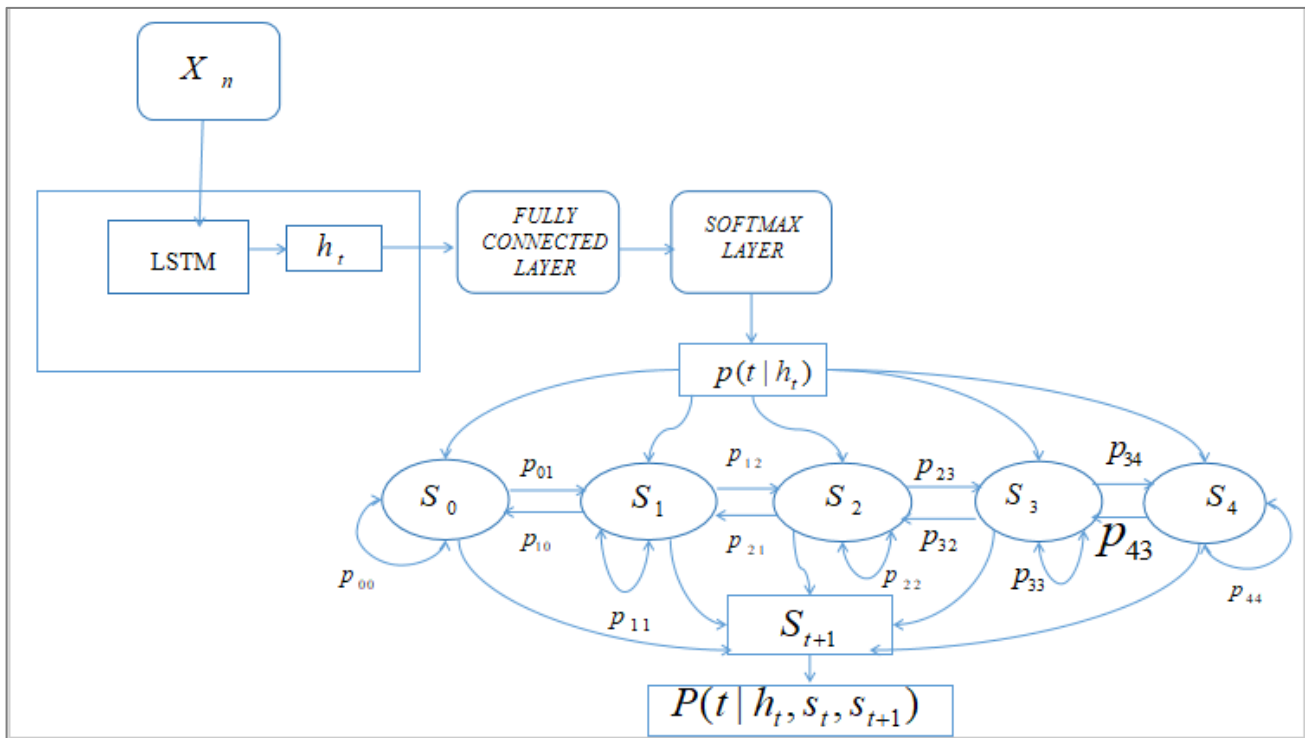


Figure 1: The Developed LSTM-Assisted Markov Chain Model

2.4.2 Assumptions for LSTM-Assisted Markov Chain Model

The proposed hybrid model is built upon the following assumptions:

1. The system satisfies the first-order Markov property, where the future state depends only on the present state.
2. The disease progresses through a finite set of mutually exclusive states.
3. Patient clinical and biomarker data contain sufficient temporal information for the LSTM network to learn meaningful non-linear relationships that influence disease progression.

4. Stage transition behavior varies across patients and is influenced by latent disease dynamics
5. The LSTM network can effectively learn the non-linear mapping from clinical data to latent states..

These assumptions collectively provide the necessary foundation for the model's structure and parameter estimation.

2.3 Parameter Estimation for the LSTM-Assisted Markov Chain Model

The goal of the learning process is to estimate a set of parameters, denoted by θ , that maximize the likelihood of observing the actual sequence of cancer stage transitions given the input patient data. This can be expressed as

$$\theta^* = \arg \max_{\theta} \sum_t \log P(S_{t+1} | S_t, X_{1:t}, \theta) \quad (16)$$

S_t is the current cancer stage,

S_{t+1} is the next stage

$X_{1:t}$ represents all input features up to time t

Where

θ represents all trainable parameters in the model,

Table 1: Variables/ parameters and their descriptions

Variables/parameters	Description
X_n	Sequential patient clinical data
LSTM	Long Short-Term Memory network responsible for learning sequential dependencies in the input features.
h_t	Hidden state vector output from the LSTM layer summarizing temporal patterns learned from previous input data.
Fully Connected Layer	A Layer that transforms the LSTM hidden representation into a vector suitable for classification or prediction.
Softmax	Activation layer that converts the fully connected layer's outputs into a probability distribution over all possible states.
$P(t h_t)$	Personalized transition probability matrix
S_t	Patient's current cancer stage at time t
S_0	Carcinoma in situ
S_1	Early/Localized
S_2	Locally Advanced
S_3	Regionally Advanced
S_4	Metastatic
P_{ij}	Transition probability from state S_i to S_j , representing the likelihood of disease movement between stages.
$P_{00}, P_{11}, P_{22}, P_{33}, P_{44}$	Probability of remaining in current stage
$P_{01}, P_{12}, P_{23}, P_{34}$	Probability of moving to next stage
P_{32}, P_{43}	Probability of regressing to previous stage
S_{t+1}	The next predicted cancer stage based on the transition probabilities and the LSTM output.
$P(t h_t, S_t, S_{t+1})$	Probability of specific transition

The parameter set θ includes all weight matrices, bias vectors, and transition probabilities defined as:

$$\theta = \{W_i, W_f, W_o, W_e, U_i, U_f, U_o, U_e, b_i, b_f, b_o, b_e, W^{(orig)}, b^{(out)}, p\} \quad (17)$$

The negative log-likelihood (NLL) is used as the loss function for optimization

$$L(\theta) = - \sum_t \log P(S_{t+1} | S_t, X_{1:t}, \theta) \quad (18)$$

For the Markov component, the transition probability between two stages i and j is given by

$$P(S_{t+1} = j | S_t = i) = p_{ij} \quad (19)$$

To estimate the parameters, gradients of the loss function are computed through Back propagation Through Time (BPTT), which unfolds the LSTM network across all time steps.

The total gradient is given by

$$\frac{\partial L}{\partial \theta} = \sum_t \frac{\partial L_t}{\partial \theta} \quad (20)$$

Where L_t is the loss contribution at each time step t .

For the LSTM weights, the gradients are obtained as:

$$\frac{\partial L}{\partial W_i} = \sum_t \left(\frac{\partial L}{\partial \dot{x}_t} \right) \left(\frac{\partial \dot{x}_t}{\partial W_t} \right) \quad (21)$$

$$\frac{\partial L}{\partial W_f} = \sum_t \left(\frac{\partial L}{\partial f_t} \right) \left(\frac{\partial f_t}{\partial W_f} \right) \quad (22)$$

$$\frac{\partial L}{\partial W_o} = \sum_t \left(\frac{\partial L}{\partial o_t} \right) \left(\frac{\partial o_t}{\partial W_o} \right) \tag{23}$$

$$\frac{\partial L}{\partial W_c} = \sum_t \left(\frac{\partial L}{\partial c_t} \right) \left(\frac{\partial i_t}{\partial W_c} \right) \tag{24}$$

The Markov transition probabilities can be estimated in two ways:

I. If transitions are determined directly from data, each element of the transition matrix **P** is computed as:

$$p_{ij} = \frac{N_{ij}}{\sum_t N_{ij}} \tag{25}$$

where N_{ij} is the number of observed transitions from stage i to stage j .

II. Transition probabilities are modeled as learnable parameters updated via gradient descent:

$$\frac{\partial L}{\partial P} = \sum_t \left(\frac{\partial L}{\partial p_{ij}(t)} \right) \left(\frac{\partial p_{ij}(t)}{\partial W^{(tr)}} \right) \tag{26}$$

where $W^{(tr)}$ represents the learnable transition weight parameters for dynamic transitions.

Table 2: Initial Stage Distribution of Simulated Patients (N=30,000)

Cancer Stage	Number of Patients	Percentage
Stage 0 (Carcinoma in situ)	4,500	15.0%
Stage 1 (Early/Localized)	7,500	25.0%
Stage 2 (Locally Advanced)	9,000	30.0%
Stage 3 (Regionally Advanced)	6,000	20.0%
Stage 4 (Metastatic)	3,000	10.0%
Total	30,000	100.0%

Table 3: Estimated Markov Transition Probabilities via Maximum Likelihood Estimation

From \ To	Stage 0	Stage 1	Stage 2	Stage 3	Stage 4
Stage 0	0.7419	0.2580	0.0000	0.0000	0.0000
Stage 1	0.0124	0.3826	0.6050	0.0000	0.0000
Stage 2	0.0000	0.0045	0.1579	0.8376	0.0000
Stage 3	0.0000	0.0000	0.0055	0.1393	0.8553
Stage 4	0.0000	0.0000	0.0000	0.0000	1.0000

Table 4: Progression Coefficients for Clinical Features via Ridge Regression

Rank	Feature	Coefficient	Absolute Effect
1	Genetic Risk	-0.3201	0.3201
2	Treatment Response	-0.1260	0.1260
3	Smoking Status	-0.1204	0.1204
4	Tumor Size	0.1161	0.1161
5	CEA Level	0.0856	0.0856
6	CA19-9	0.0654	0.0654
7	Gender	0.0503	0.0503
8	Lymph Involvement	0.0468	0.0468
9	Ki-67 Index	0.0444	0.0444
10	Tumor Grade	0.0160	0.0160

Table 5: Data Partitioning Strategy for Model Development

Partition	Samples	Percentage	Primary Purpose
Training Set	19,200	64.0%	Model parameter estimation
Validation Set	4,800	16.0%	Hyperparameter tuning and early stopping
Test Set	6,000	20.0%	Final model evaluation
Total	30,000	100.0%	—

Table 6: Estimated LSTM Parameters via Backpropagation Through Time (BPTT)

Epoch Range	Training Accuracy	Validation Accuracy	Training Loss	Validation Loss
1-5	86.42-91.27%	91.42-91.56%	0.4906-0.2539	0.2441-0.2356
6-20	91.30-91.48%	91.52-91.69%	0.2474-0.2412	0.2365-0.2334
21-39	91.34-91.52%	91.62-91.65%	0.2419-0.2392	0.2320-0.2305
40-54	91.30-91.48%	91.58-91.69%	0.2390-0.2365	0.2342-0.2348

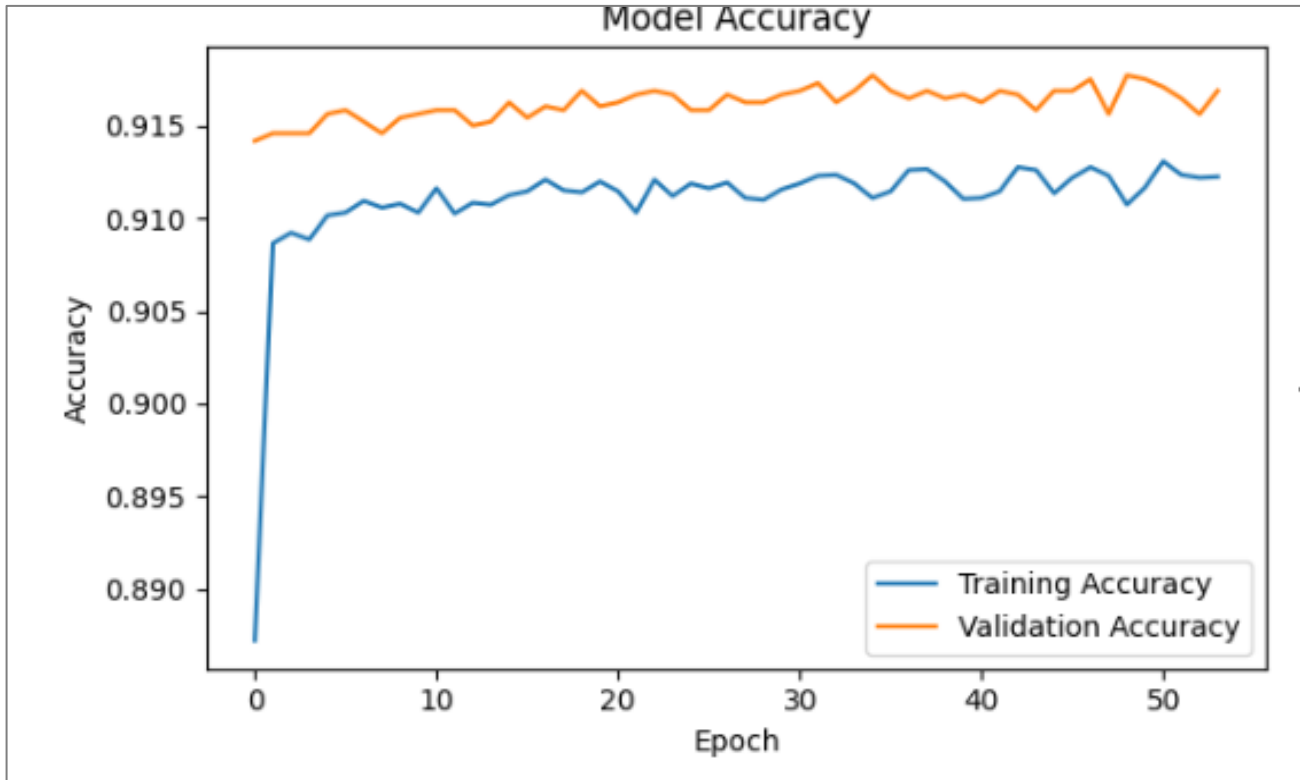


Figure 2: Model Accuracy




Figure 3: Model Loss

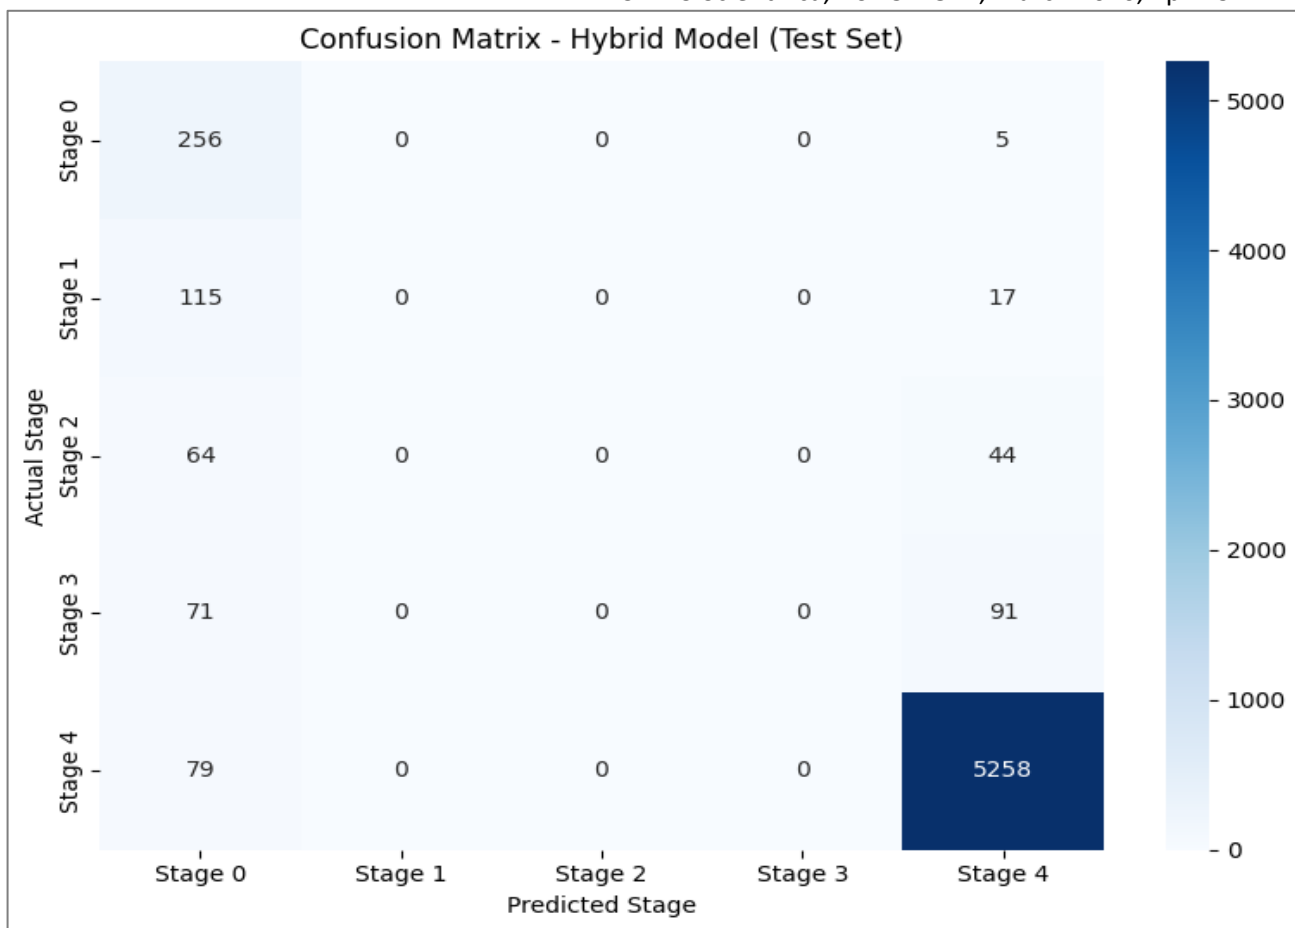


Figure 4: Confusion Matrix for LSTM-Assisted Markov Chain Model

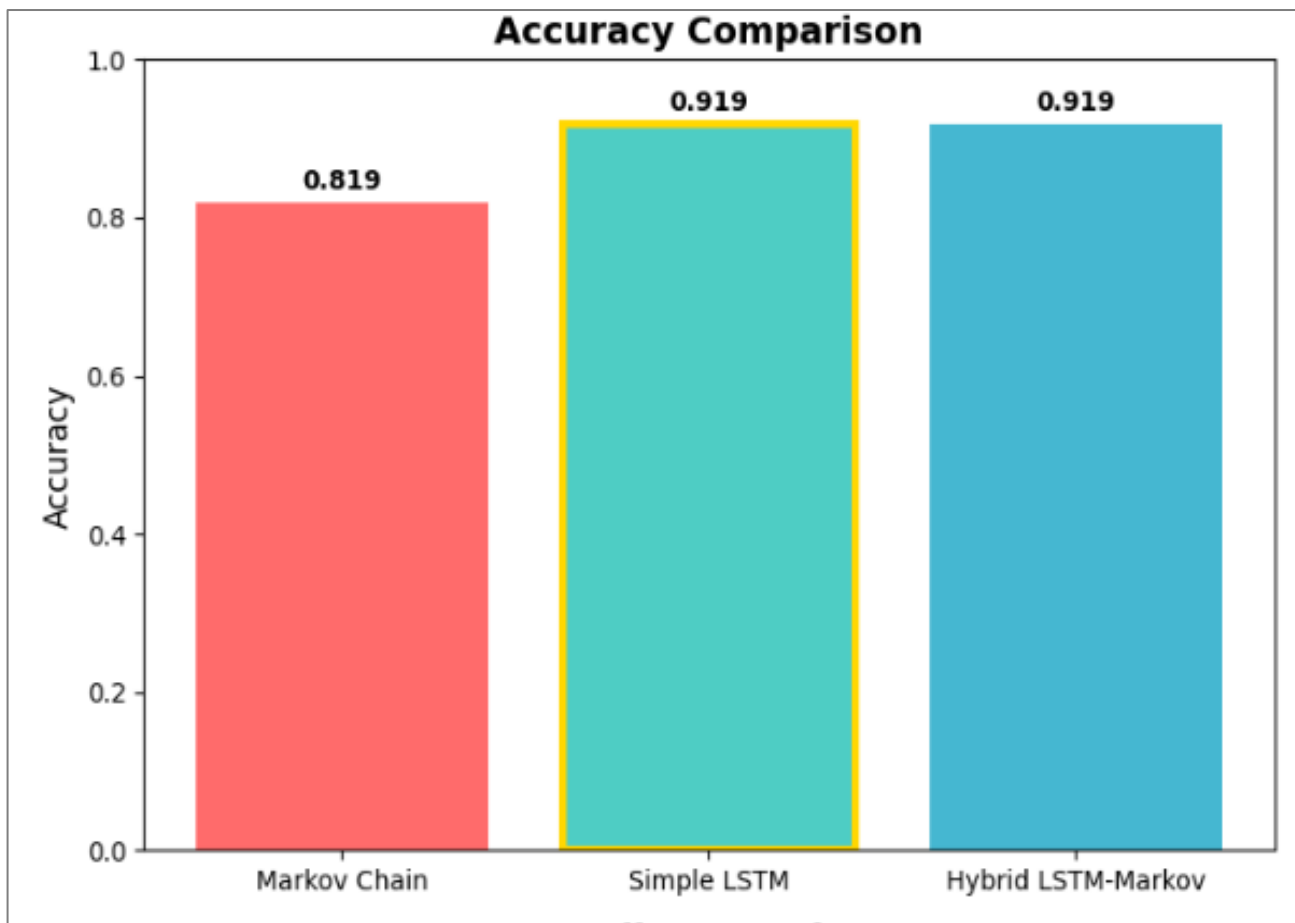


Figure 5: Accuracy Chart




Figure 6: Precision Chart




Figure 7: Recall Chart

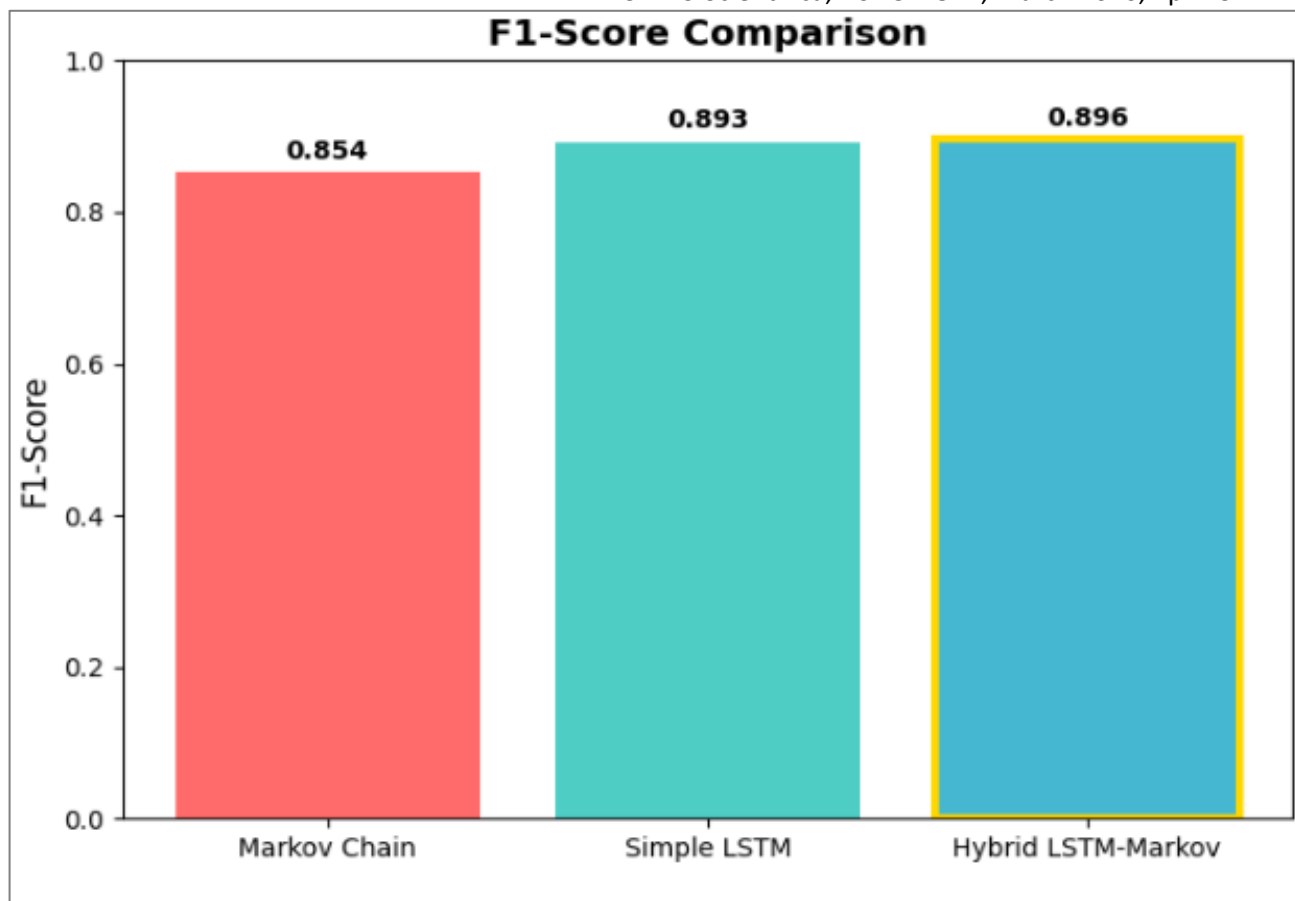


Figure 8: F1-score chart

Table 7: Per-Stage Performance of the LSTM-Assisted Markov Chain Model (Test Set, N=6,000)

Cancer Stage	Precision	Recall	F1-Score	Support (N)
Stage 0	0.4376	0.9808	0.6052	261
Stage 1	0.0000	0.0000	0.0000	132
Stage 2	0.0000	0.0000	0.0000	108
Stage 3	0.0000	0.0000	0.0000	162
Stage 4	0.9710	0.9852	0.9781	5,337
Overall Accuracy	—	—	0.9190	6,000
Macro Average	0.2817	0.3932	0.3167	6,000
Weighted Average	0.8827	0.9190	0.8963	6,000

Table 8: Model Performance Comparison on Test Data

Model	Accuracy	Precision	Recall	F1-Score
Mackov Chain	0.8185 (81.85%)	0.9106 (91.06%)	0.8185 (81.85%)	0.8536 (85.36%)
Simple LSTM	0.9193 (91.93%)	0.8725 (87.25%)	0.9193 (91.93%)	0.8930 (89.30%)
LSTM-Assisted Markov	0.9190 (91.90%)	0.8827 (88.27%)	0.9190 (91.90%)	0.8963 (89.63%)

2.4 Hyperparameters and Training Details

The LSTM architecture comprised three stacked LSTM layers with 128, 64, and 32 units respectively, each with dropout (0.2) and recurrent dropout (0.2). The output was fed into three dense layers (128, 64, 32 units) with ReLU activation and dropout rates of 0.3, 0.3, and 0.2. Training used the Adam optimizer (learning rate = 0.001, clipnorm = 1.0) with batch size 32 for up to 100 epochs. Early stopping with patience of 15 epochs, monitored on validation loss, restored the best weights from epoch 39. All hyperparameters were fixed based on initial experimentation prior to final test set evaluation.

2.5 Evaluation Metrics

Model performance was assessed using: Accuracy, Precision, Recall, F1-Score and Confusion Matrix. Comparison with baseline models: Pure Markov Chain and standalone LSTM.

2.6 Software and Hardware

All simulations and model training were performed in Python 3.8 using: NumPy 1.21, Pandas 1.3, Matplotlib 3.4, Seaborn 0.11, Scikit-learn 1.0, TensorFlow 2.8 (Keras API). Training was conducted on a google colab with 32 GB RAM. Random seeds were fixed for reproducibility (np.random.seed(42), tf.random.set_seed(42)).

RESULTS

3.1 Dataset Characteristics

The simulated dataset comprised 30,000 independent patient sequences with variable follow-up (mean 10 months, range 5-15). [Table 2](#) presents the initial stage distribution across the five cancer stages.

3.2 Parameter Estimation

The transition matrix in [Table 3](#) reveals clinically realistic progression patterns: increasing forward transition probabilities with stage (25.8% from Stage 0→1 to 85.5% from Stage 3→4), rare backward transitions (<1.5%), and Stage 4 as an absorbing state (100% self-transition). These estimates confirm that the simulation successfully encoded irreversible cancer progression dynamics.

[Table 4](#) shows that Genetic risk (-0.3201) emerged as the strongest protective factor, while tumor size (0.1161) and CEA level (0.0856) were major contributors to disease progression. Negative coefficients represent protective factors, positive coefficients indicate risk factors.

3.3 Data Partitioning

In [Table 5](#), it shows that splitting was performed at the patient level to prevent temporal data leakage, with stratification to maintain class proportions across all sets

3.4 LSTM Training and Parameter Estimation

The model showed efficient learning in [Table 6](#), reaching approximately 91.5% accuracy within a few epochs and maintaining stable performance. Early stopping restored the best-performing weights from epoch 39 (validation accuracy 91.77%). The final model contained 64,160 LSTM parameters.

Both training and validation accuracy converge at approximately 91%, indicating effective learning ([Figure 2](#)). The slight gap between curves reveals minor overfitting, where the model has begun to memorize training data specifics.

Training and validation loss decrease sharply then stabilize, confirming successful error minimization ([Figure 3](#)). The persistent gap between validation and training loss confirms the minor overfitting observed in the accuracy graph.

3.6 Per-Stage Performance Analysis

[Table 7](#) shows how the model achieves excellent performance for Stage 4 (metastatic cancer) with F1-score 0.9781, correctly identifying 98.5% of true metastatic cases. Stage 0 demonstrates high recall (0.9808) but moderate precision (0.4376), indicating some false positive classifications. However, the model completely fails to classify intermediate stages (Stages 1-3), with F1-scores of zero. This failure coincides with severe class imbalance: Stages 1-3 collectively constitute only 6.7% (402/6,000) of test samples, while Stage 4 comprises 88.95%.

3.7 Confusion Matrix Analysis

The confusion matrix reveals: Stage 4: 5,258 correct classifications, 79 misclassified; Stage 0: 256 correct, 5 misclassified to Stage 4; Stages 1-3: All 402 samples misclassified (primarily as Stage 4). This pattern indicates that the model is highly sensitive to progression signals but lacks the granularity to distinguish among non-metastatic stages ([Figure 4](#)).

3.8 Model Performance Comparison

The comparative analysis in [Table 8](#) demonstrates that the proposed LSTM-assisted Markov Chain model significantly outperforms the classical Markov Chain baseline, achieving an accuracy improvement of approximately **12.3%** (0.9190 vs. 0.8185). This improvement highlights the limitation of traditional Markov models in capturing complex temporal and nonlinear patterns inherent in cancer progression data. While the Simple LSTM model and the Hybrid model achieve nearly identical overall accuracy (0.9193 vs. 0.9190), the hybrid approach shows a notable improvement in precision (0.8827 vs. 0.8725) and F1-score (0.8963 vs. 0.8930). This indicates that integrating Markov transition dynamics with LSTM temporal learning reduces false positive predictions and provides a better balance between precision and recall. Overall, these results confirm that the hybrid framework preserves the strong predictive power of LSTM while enhancing decision reliability through probabilistic stage-transition modeling, making it more suitable for cancer progression prediction than either model used independently.

3.9 Model Performance Comparison on Metrics

[Figure 5](#) shows Simple LSTM and the Hybrid LSTM-Markov models achieve a high and nearly identical accuracy of approximately 0.919, which is substantially higher than the baseline Markov Chain accuracy of 0.819. This indicates that incorporating temporal learning through LSTM significantly improves overall stage prediction compared to relying solely on Markovian transitions. However, accuracy alone does not capture the quality of predictions across all stages, especially in imbalanced cancer progression data.

[Figure 6](#) Reveals an important distinction among the models. While the Markov Chain model attains the highest precision (≈ 0.911), this occurs at the expense of lower recall and accuracy, suggesting that it makes fewer positive predictions but misses many true stage transitions. The Hybrid model achieves a higher precision (≈ 0.883) than the Simple LSTM (≈ 0.872), indicating that integrating LSTM-derived emission probabilities into the Markov framework reduces false positive stage predictions while maintaining strong predictive power.

[Figure 7](#) shows that both the Simple LSTM and the Hybrid model attain equally high recall values (≈ 0.919), outperforming the Markov Chain model. This demonstrates that models incorporating LSTM are more effective at correctly identifying true cancer stage transitions. Importantly, the Hybrid model preserves this

high recall while improving precision relative to the Simple LSTM, reflecting a better balance between sensitivity and specificity.

Figure 8 which harmonically balances precision and recall, clearly highlights the strength of the proposed Hybrid model. The model achieves the highest F1-score (≈ 0.896), exceeding both the Simple LSTM (≈ 0.893) and the Markov Chain model (≈ 0.854). This indicates that although accuracy values are similar between the Simple LSTM and the Hybrid model, the Hybrid approach delivers more reliable and clinically meaningful predictions by simultaneously minimizing false positives and false negatives.

DISCUSSION

The LSTM-assisted Markov chain model developed in this study achieved strong overall predictive performance, with 91.9% accuracy and an F1-score of 89.6%, representing a 12.3% improvement over the traditional Markov chain baseline. These aggregate metrics initially suggest that the hybrid architecture successfully combines the temporal learning capabilities of LSTM networks with the probabilistic structure of Markov chains to model cancer progression. However, closer examination of per-stage performance reveals critical limitations that must be explicitly acknowledged and addressed before any clinical application can be contemplated.

The model demonstrated excellent performance for Stage 4 (metastatic cancer) with an F1-score of 0.978, correctly identifying 98.5% of true metastatic cases, and showed moderate performance for Stage 0 (carcinoma in situ) with an F1-score of 0.605. However, it completely failed to classify intermediate stages (Stages 1, 2, and 3), with F1-scores of zero. This failure is attributable to three primary factors. First, severe class imbalance meant that Stages 1-3 collectively constituted only 6.7% (402 out of 6,000) of test samples, while Stage 4 alone comprised 88.95% (5,337 out of 6,000). During training, the loss function was consequently dominated by Stage 4 examples, causing the model to optimize for metastatic detection at the expense of intermediate stage discrimination. Second, the underlying Markov structure reflected the clinical reality of irreversible cancer progression, with backward transition probabilities below 1.5%. Once the model predicted progression to Stage 4, it rarely corrected to earlier stages, effectively functioning as a binary classifier distinguishing metastatic from non-metastatic disease rather than a true five-stage classifier. Third, analysis of feature distributions revealed substantial overlap in biomarker values between adjacent intermediate stages. For example, CEA levels in Stage 1 patients can overlap considerably with those in Stage 2 or 3 patients, making discrimination difficult even for expert clinicians. Stage 4, by contrast, exhibits distinct feature distributions that are easily separable. The clinical implication is that while the model may serve as a useful early warning system for detecting progression to metastatic disease, it should not be relied upon for fine-grained staging of non-metastatic disease in its current form.

The use of simulated data, while enabling controlled methodological validation, introduces important limitations that affect interpretation of results. The simulation framework encoded clinically-informed distributions and progression patterns, but necessarily simplified the complexity of real-world cancer dynamics. Key simplifications include the absence of missing data, whereas real clinical datasets contain missing values, irregular sampling intervals, and loss to follow-up. All 17 features were available at every time point, whereas real data often have variable feature collection across visits. The simulation provided perfect knowledge of true disease stage, while real clinical staging involves measurement error, inter-rater variability, and diagnostic uncertainty. Treatment response was modeled with simple parametric distributions, whereas real treatment effects involve complex interactions, side effects, and substantial patient heterogeneity. Furthermore, the severe class imbalance observed, while reflecting the irreversible progression dynamics encoded in the simulation, may not perfectly match all real-world clinical cohorts, which vary depending on the clinical setting such as screening populations versus tertiary cancer centers. Consequently, all results are derived from synthetic data generated under specific parametric assumptions, and the reported performance metrics represent upper-bound estimates of what might be achievable with ideal data, not guarantees of real-world performance.

Despite these limitations, the hybrid model's overall accuracy of 91.9% is consistent with recent hybrid deep learning applications in oncology. Lilhore et al. (2025) achieved 99.2% accuracy using CNN-BiLSTM for breast cancer detection, while Rastogi et al. (2024) reported 99% accuracy with Conv1D-LSTM. However, these studies focused on classification tasks with balanced datasets, whereas our work addresses the more challenging problem of sequential state transition prediction under severe class imbalance. The 12.3% improvement over traditional Markov chains confirms the findings of Sengupta et al. (2023), who demonstrated that hybrid hidden Markov-LSTM architectures preserve interpretability while enhancing predictive accuracy. Our work extends this by explicitly quantifying the trade-off between overall accuracy and per-stage performance under imbalance, providing a more transparent evaluation framework.

Based on the identified limitations, several concrete next steps are proposed for improving the model and validating it for clinical use. To address class imbalance, future work should explore two-stage hierarchical modeling that first distinguishes metastatic from non-metastatic disease before training a separate multi-class classifier for non-metastatic stages only. Cost-sensitive learning with higher misclassification weights for Stages 1-3 should be investigated, along with synthetic minority oversampling techniques such as SMOTE adapted for sequential data. Collection of more balanced real-world data from screening populations rather than tertiary referral centers would also help. Before any clinical deployment, the model must undergo rigorous validation on real-world

longitudinal datasets, beginning with external validation using public datasets such as SEER and TCGA, followed by multi-center retrospective validation across diverse hospital systems and patient populations, and ultimately temporal validation on more recent cohorts to ensure performance stability over time. For clinical applicability, models must communicate uncertainty in their predictions through Bayesian neural networks, Monte Carlo dropout, or conformal prediction frameworks that provide prediction sets with guaranteed coverage probabilities. Interpretability must be enhanced through attention mechanisms to identify influential time points and features, SHAP values to quantify feature contributions for individual predictions, and prototype-based explanations that reference similar patients from the training set.

It is essential to emphasize that this work represents methodological validation only, not clinical validation. The model is not ready for clinical use, and we explicitly refrain from making clinical recommendations until real-world validation is complete. Several critical gaps remain. There has been no real-world validation; all results are from simulated data, and real-world performance may be substantially lower due to data quality issues, population heterogeneity, and unmodeled complexity. Even retrospective real-world validation would not guarantee prospective performance, as a prospective study would be needed to assess real-world utility. The model's output format, predicting the next stage, may not align with clinical decision-making needs, requiring co-design with clinicians to develop useful decision support tools. As a medical device, the model would require regulatory approval from bodies such as the FDA or EMA before clinical use, which is beyond the scope of this academic work. Finally, even if predictions are accurate, it remains unknown whether acting on these predictions improves patient outcomes, which would require a randomized controlled trial.

CONCLUSION

In conclusion, this study demonstrates that an LSTM-assisted Markov chain model can achieve strong overall accuracy in predicting cancer progression under simulated conditions. However, the model's complete failure to classify intermediate stages reveals that aggregate metrics mask critical limitations, and the use of simulated data precludes clinical recommendations. The primary contribution is methodological: demonstrating a hybrid architecture that preserves interpretability while improving over traditional Markov models, and providing a framework for transparent evaluation under class imbalance. Future work must focus on real-world validation using clinical datasets such as SEER and TCGA, techniques to address class imbalance for intermediate stages, uncertainty quantification for clinical decision support, and co-design with clinicians to ensure clinical utility. Only after these steps can the model be considered for clinical deployment.

REFERENCES

- Abohashish, S. M. M., Amin, H. H., & Elsedimy, E. (2025). Enhanced melanoma and non-melanoma skin cancer classification using a hybrid LSTM-CNN model. *Scientific Reports*, 15, 1–16. [\[Crossref\]](#)
- Ceritli, T., Creagh, A. P., & Clifton, D. A. (2022). Mixture of input-output hidden Markov models for heterogeneous disease progression modeling. *arXiv*, Article 2207.11846. [\[Crossref\]](#)
- Chen, T., et al. (2025). Cross-representation benchmarking in time-series electronic health records for clinical outcome prediction. *arXiv preprint*, Article arXiv:2510.09159.
- Dias, L., Antunes, L., & Oliveira, J. (2020). A hidden Markov model for cancer progression modeling. *Scholar*. Retrieved from [\[Link\]](#)
- Ferle, M., Chernyshev, M., Schaaf, A., & Gumbsch, T. (2024). Predicting progression events in multiple myeloma from routine blood work. *Blood*, 144(Supplement 1), Article 7476. [\[Crossref\]](#)
- Giuliani, J., Bonetti, A., & Bertolazzi, L. (2022). Can Markov chains predict survival in stage IV pancreatic cancer? *Annals of Pancreatic Cancer*, 5, Article 6. Retrieved from [\[Link\]](#)
- Hanahan, D., & Weinberg, R. A. (2011). Hallmarks of cancer: The next generation. *Cell*, 144(5), 646–674. [\[Crossref\]](#)
- Hawkins, N., Sculpher, M., & Epstein, D. (2009). Bias in Markov models of disease. *Mathematical Biosciences*, 220(2), 143–156. [\[Crossref\]](#)
- Huang, Y., Chen, Y., Xu, W., & Zhang, S. (2021). A Markov chain model of cancer treatment. *bioRxiv*. [\[Crossref\]](#)
- Jackson, C. H., Sharples, L. D., Thompson, S. G., Duffy, S. W., & Couto, E. (2003). Multistate Markov models for disease progression with classification error. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 52(2), 193–209. [\[Crossref\]](#)
- Jin, X., Zhang, Y., & Wang, L. (2023). An autoregressive integrated moving average and long short-term memory (ARIMA-LSTM) hybrid model for multi-source epidemic data prediction. *PeerJ Computer Science*, 9, Article e2046. [\[Crossref\]](#)
- Kaddes, M., Ayid, Y. M., Elshewey, A. M., & Fouad, Y. (2025). Breast cancer classification based on hybrid CNN with LSTM model. *Scientific Reports*, 15, 1–14. [\[Crossref\]](#)
- Lilhore, U., Sharma, Y. K., Shukla, B. K., Vadlamudi, M., Simaiya, S., Alroobaea, R., Alsafyani, M., & Baqasah, A. M. (2025). Hybrid convolutional neural network and bi-LSTM model with EfficientNet-B0 for high-accuracy breast cancer detection and classification. *Scientific Reports*, 15, 1–15. [\[Crossref\]](#)
- Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., ... Dean, J. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1, Article 18. [\[Crossref\]](#)
- Rastogi, M., Vijarana, M., Goel, N., Agrawal, A., Biamba, C. N., & Iwendi, C. (2024). Conv1D-LSTM:

- Autonomous breast cancer detection using a one-dimensional convolutional neural network with long short-term memory. *IEEE Access*, 12, 1–12. [[Crossref](#)]
- Sengupta, A., Das, A., & Guler, S. I. (2023). Hybrid hidden Markov-LSTM for short-term traffic flow prediction. *arXiv preprint*.
- Sethi, B. K., Singh, D., Rout, S. K., & Panda, S. K. (2023). Long short-term memory-deep belief network-based gene expression data analysis for prostate cancer detection and classification. *IEEE Access*, 11, 1–12. [[Crossref](#)]
- Swanton, C. (2012). Intratumor heterogeneity: Evolution through space and time. *Cancer Research*, 72(19), 4875–4882. [[Crossref](#)]
- Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2019). What clinicians want: Contextualizing explainable machine learning for clinical end use. *Proceedings of the Machine Learning Research*, 106, 359–380. Retrieved from [[Link](#)]
- Wang, S., Zhang, Y., & Liu, H. (2023). HCBiLSTM: A hybrid model for predicting heart disease using CNN and BiLSTM algorithms. *Measurement: Sensors*, 25, Article 100657. [[Crossref](#)]
- World Health Organization. (2021). *Cancer*. Retrieved from [[Link](#)]
- Zhang, H., Li, X., & Wang, J. (2022). A hybrid model for hand-foot-mouth disease prediction based on ARIMA-EEMD-LSTM. *BMC Infectious Diseases*, 22, Article 864.