

A Review of Entropy-Guided Latent Risk State Discovery for Thyroid Nodule Classification using Deep Learning and Gaussian Hidden Markov Models

Swapnil Banerjee¹, Kakali Karmakar (Sur)²

¹Department of Computer Science and Engineering (AI & ML), Future Institute of Technology
Kolkata-700103, West Bengal, India.

²Department of Mathematics, B. P. Poddar Institute of Management and Technology
Kolkata-700052, West Bengal, India.

Abstract

Conventional thyroid nodule classification studies primarily focus on observable diagnostic states, namely benign and malignant classes. While such approaches provide useful classification outcomes, they offer limited insight into diagnostic uncertainty and hidden structures present within the data. In this work, an entropy-guided latent risk-state discovery framework is proposed by integrating deep learning, Shannon entropy analysis, and Gaussian Hidden Markov Models (HMMs). A convolutional neural network is employed to estimate benign and malignant probabilities from thyroid ultrasound images. Shannon entropy is subsequently computed to quantify prediction uncertainty. The resulting probability and entropy values are combined into observation vectors and analysed using a Gaussian Hidden Markov Model. The proposed framework discovers four latent risk states and reveals a maximum-uncertainty state characterised by nearly equal benign and malignant probabilities and entropy approaching its theoretical maximum. The results demonstrate that entropy-enhanced latent-state modelling provides richer diagnostic interpretation than conventional binary classification and offers a promising direction for uncertainty-aware clinical decision support.

Keywords: Shannon Entropy, Hidden Markov Model, Thyroid Nodule Classification, Deep Learning, Uncertainty Quantification, Latent Risk States

Introduction

Thyroid cancer is among the most common endocrine malignancies worldwide. Early diagnosis plays a crucial role in patient management and treatment planning. Recent advances in deep learning have enabled automated classification of thyroid nodules from ultrasound images. However, most existing approaches focus only on observable diagnostic states, namely benign and malignant categories. Such systems typically provide a final class label without quantifying uncertainty associated with the prediction. In clinical practice, uncertainty is often as important as the prediction itself. Cases with nearly equal benign and malignant likelihoods may require further examination, biopsy, or expert review. Therefore, uncertainty-aware diagnostic systems are highly desirable. To address this limitation, this work incorporates Shannon entropy for uncertainty quantification and Gaussian Hidden Markov Models for latent risk-state discovery.

Literature Review

Deep learning techniques have demonstrated remarkable success in medical image analysis. Convolutional Neural Networks (CNNs) have been widely used for thyroid ultrasound classification. However, most studies report only classification accuracy and related metrics. Information-theoretic measures such as Shannon entropy have recently gained attention for uncertainty quantification. Hidden Markov Models have been successfully applied in healthcare for latent-state analysis and probabilistic modelling. The integration of deep learning, entropy analysis, and HMM-based risk-state discovery remains relatively unexplored in thyroid cancer assessment. [1]

Mathematical Background

CNN Probability Model

The trained CNN outputs the probability of benign and malignant classes:

$$P(B) + P(M) = 1$$

where $P(B)$ denotes the probability of a benign nodule and $P(M)$ denotes the probability of a malignant nodule. [4]

Shannon Entropy

Prediction uncertainty is quantified using Shannon entropy:

$$H(X) = -(P_B \log_2 P_B + P_M \log_2 P_M)$$

Low entropy corresponds to confident predictions, whereas entropy approaching one indicates maximum uncertainty.

Observation Vector

For each image, an observation vector is defined as

$$O_t = [P_B, P_M, H]$$

This vector forms the input to the Hidden Markov Model. [7]

Hidden Markov Model

The HMM consists of hidden states and observation vectors. The transition probability is given by

$$a_{ij} = P(S_{t+1} = j | S_t = i)$$

while the emission probability is

$$b_j(O_t) = P(O_t | S_t = j)$$

The most probable hidden-state sequence is obtained using Viterbi decoding:

$$S^* = \underset{S}{\operatorname{argmax}} P(S, O)$$

Dataset Description

The study utilised a thyroid ultrasound dataset containing benign and malignant thyroid nodule images. Images were preprocessed and resized before being passed to the CNN model. A train-validation-test split was performed to evaluate the proposed framework on unseen samples.[13]

Proposed Methodology

The complete workflow is illustrated in Figure 1.




Figure 1: Overall methodology.

The CNN generates probability estimates for benign and malignant classes. Shannon entropy is then computed to quantify uncertainty. The probability values and entropy are combined into observation vectors and analysed using a Gaussian Hidden Markov Model. The HMM automatically discovers latent diagnostic states and enables risk-state interpretation.

Experimental Results

Classification Results

Table 1: Classification Performance

Metric	Value
Accuracy	75.07%
Precision	72.95%
Recall	79.47%
F1 Score	76.07%

The obtained performance metrics indicate that the proposed CNN model achieves reliable classification of thyroid nodules. The balanced values of accuracy, precision, recall, and F1-score suggest effective discrimination between benign and malignant cases. These results validate the use of the model for generating probability estimates that are subsequently utilized for entropy computation and Hidden Markov Model-based risk-state analysis.[5]




Figure 2: Confusion Matrix of CNN predictions.

The confusion matrix demonstrates the classification capability of the proposed FibonacciNet model on the test dataset. A large number of samples were correctly classified in both benign and malignant categories, indicating that the model is able to effectively distinguish between the two classes. The relatively smaller number of misclassifications suggests good overall predictive performance.




Figure 3: Accuracy and Loss Plots

The training accuracy and loss curves further illustrate the learning behaviour of the model. Training accuracy shows a consistent upward trend, while training loss decreases steadily over successive epochs, indicating effective feature learning. Although some fluctuations are observed in the validation curves, the overall performance demonstrates satisfactory convergence and provides a reliable foundation for the subsequent entropy-based uncertainty analysis and Hidden Markov Model risk-state discovery. [9]

Entropy Analysis




Figure 4: Entropy distribution across discovered risk states.

The entropy analysis demonstrates that uncertainty is not uniformly distributed across samples. Certain latent states exhibit significantly higher entropy values, indicating diagnostically ambiguous cases. This observation confirms that Shannon entropy effectively captures the confidence level of the model and provides additional information beyond the final benign or malignant prediction. [7]

Overall, the entropy analysis confirms that Shannon entropy serves as an effective uncertainty quantification measure for thyroid nodule classification. The observed relationship between entropy and latent risk states demonstrates that uncertainty-aware modelling can reveal hidden diagnostic structures that remain inaccessible in conventional binary classification approaches. This finding highlights the importance of integrating information-theoretic measures with probabilistic state modelling for enhanced medical decision support. [6, 7]

Hidden State Discovery




Figure 5: Discovered latent risk states using Gaussian HMM.

Table 2: Discovered Latent Risk States

State	P_B	P_M	Entropy	Risk State
S1	0.89	0.11	0.47	Low Risk
S2	0.51	0.49	0.98	Maximum Uncertainty
S3	0.26	0.74	0.80	Very High Risk
S4	0.72	0.28	0.85	Intermediate Risk

Maximum Uncertainty State

The Gaussian Hidden Markov Model identified a latent state characterised by nearly equal benign and malignant probabilities ($P_B = 0.51$, $P_M = 0.49$) and entropy approaching the theoretical maximum ($H = 0.98$). This state represents diagnostically ambiguous cases where the classifier exhibits maximum uncertainty. Such cases cannot be distinguished through conventional binary classification alone and demonstrate the importance of entropy-guided latent-state modelling. [7]

Sample Prediction Analysis




Figure 6: Representative test image with prediction, entropy and risk-state assignment.

Table 3: Representative Sample Predictions

	GroundTruth	Predicted_Label	P_Benign	P_Malignant	Entropy	HiddenState	RiskState
0	Benign	Malignant	0.292000	0.708000	0.871000	2	High Risk
1	Benign	Benign	0.851000	0.149000	0.607000	0	Low Risk
2	Malignant	Malignant	0.183000	0.817000	0.686000	2	High Risk
3	Benign	Benign	0.574000	0.426000	0.984000	1	Intermediate Risk
4	Malignant	Malignant	0.319000	0.681000	0.904000	2	High Risk

Discussion

Unlike traditional thyroid classification studies that focus solely on observable benign and malignant states, the proposed framework investigates uncertainty through Shannon entropy and discovers hidden latent risk states using Gaussian Hidden Markov Models. The entropy analysis revealed meaningful differences between discovered states. In particular, the maximum-uncertainty state corresponded to nearly equal benign and malignant probabilities, indicating cases where the model was unable to make a confident decision.[3]

The discovered latent states provide a richer interpretation than standard binary classification. Rather than producing only benign or malignant labels, the framework identifies confidence levels, uncertainty regions, and hidden risk structures within the data.

Conclusion

This study proposed an entropy-guided latent risk-state discovery framework for thyroid nodule assessment by integrating deep learning, Shannon entropy analysis, and Gaussian Hidden Markov Models. Shannon entropy successfully quantified diagnostic uncertainty, while the Gaussian Hidden Markov Model automatically discovered four meaningful latent risk states. A maximum-uncertainty state characterised by nearly equal benign and malignant probabilities was identified, revealing diagnostically ambiguous cases that cannot be adequately represented through conventional binary classification. The results demonstrate that entropy-enhanced latent-state modelling provides richer interpretability and deeper diagnostic insight than traditional classification systems, making it a promising approach for uncertainty-aware medical image analysis and risk-aware clinical decision support.

References

1. Ajabani, Deep Himmatbhai. *Predicting Alzheimer's Progression in Mild Cognitive Impairment: Longitudinal MRI with HMMs and SVM Classifiers*. International Journal of Advanced Computer Science and Applications (IJACSA), 14(12), 2023.
2. Chattopadhyay, Subhagata. *Towards Predicting Recurrence Risk of Differentiated Thyroid Cancer with a Hybrid Machine Learning Model*. MedInformatics, 3(2), 2026.
3. Kwok, H.; Coult, J.; Blackwood, J.; Sotoodehnia, N.; Kudenchuk, P.; Rea, T. *A Method for Continuous Rhythm Classification and Early Detection of Ventricular Fibrillation During CPR*. Resuscitation, pp. 90–97, 2022.
4. Liu, Yuxi; Zhang, Zhenhao; Qin, Shaowen. *NeuralHMM: A Deep Markov Network for Health Risk Prediction using Electronic Health Records*. IEEE, 2024.
5. Nilsson, F.; Stridh, M.; Sörnmo, L. *Frequency Tracking of Atrial Fibrillation using Hidden Markov Models*. Conference Proceedings of the IEEE Engineering in Medicine and Biology Society, pp. 1406–1409, 2006.
6. Niu, Shuai; Ma, Jing; Yin, Qing; Bai, Liang; Li, Chen; Yang, Xian. *A Deep Clustering-Based State-Space Model for Improved Disease Risk Prediction in Personalized Healthcare*. Original Research, 2024.
7. Núñez Mejía, Sebastián. *Hidden Markov Models for Early Detection of Cardiovascular Diseases*. Revista Ingeniería Solidaria, 20(1), 2024.
8. Oliveira, J.; Sousa, C.; Coimbra, M. *Coupled Hidden Markov Model for Automatic ECG and PCG Segmentation*. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1023–1027, 2017.
9. Pamulaparthivenkata, Saigurudatta; Sharma, Jitendra; Vishwanath, Manish; Avacharmal, Rajiv; Mulukuntla, Sarika; Vahidbabanu. *Utilizing EHR in Machine Learning-Based Systems for Early Heart Disease Prediction in Healthcare Applications*. IEEE, 2024.

10. Petrutiu, S.; Sahakian, A. V.; Swiryn, S. *Abrupt Changes in Fibrillatory Wave Characteristics at the Termination of Paroxysmal Atrial Fibrillation in Humans*. *Europace*, 9(7):466–470, 2007.
11. Pimentel, M. A. F.; Santos, M. D.; Springer, D. B.; Clifford, G. D. *Heart Beat Detection in Multimodal Physiological Data Using a Hidden Semi-Markov Model and Signal Quality Indices*. *Physiological Measurement*, 36(8):1717–1727, 2015.
12. Sangaiah, A. K.; Arumugam, M.; Bin Bian, G. *An Intelligent Learning Approach for Improving ECG Signal Classification and Arrhythmia Analysis*. *Artificial Intelligence in Medicine*, 103:1–14, 2020.
13. Sowmya Abirami. *Thyroid Ultrasound Dataset*. Kaggle. Available at: <https://www.kaggle.com/datasets/sowmyaabirami/thyroid-ultrasound-dataset>