Journal of Computational Medicine and Informatics
Vol. 2, No. 1, Jan - Mar 2026, pp. 1-8
ISSN: 3107-3735 (O), DOI: https://doi.org/10.17051/JCMI/02.01.01

Uncertainty-Aware Neural Network Framework for Reliable Medical Diagnosis and Prognosis

P.Sakthi Murugan

Associate Professor, Department of Data Analytics and Mathematical Sciences
School of Sciences, Jain Deemed-to-be University, JC Road, Bengaluru, Email: sakthigasc@gmail.com,
https://orcid.org/0009-0004-9376-6087

Article Info

Article history:

Received : 04.11.2025
Revised : 07.12.2025
Accepted : 05.01.2026

Keywords:

Uncertainty-aware learning,
Bayesian neural networks,
medical diagnosis,
prognosis prediction,
Monte Carlo dropout,
clinical decision support

ABSTRACT

Recent developments in artificial intelligence have greatly improved diagnostic and predictive outcomes in healthcare; yet, most current deep learning algorithms make deterministic predictions without uncertainty assessment, which reduces their effectiveness in critical decision-making. In this work, we present an uncertainty-aware neural network approach that combines probabilistic inference with deep learning to improve diagnostic certainty and confidence. The approach incorporates Bayesian inference through Monte Carlo dropout, to quantify uncertainty along with classification predictions. The model accounts for both model and data uncertainties, enhancing decision-making in the context of uncertain medical data. The proposed framework achieves better accuracy, calibration and confidence prediction performance on medical datasets compared to traditional neural networks. This finding underscores the benefits of uncertainty-aware learning in mitigating overconfidence and increasing confidence in AI-based medical systems. This research offers a scalable approach to decision support for dependable diagnosis and prognosis.

1. INTRODUCTION

Data-driven decision making in medical diagnosis and prognosis has seen a major leap in capabilities through the advent of artificial intelligence (AI) in healthcare [8], [11]. Neural networks, a type of deep learning model, have shown impressive capabilities in disease diagnosis, risk stratification and prognostication [12]. These models make use of large-scale medical data, such as electronic health records (EHRs), medical images, and physiological signals to uncover subtle patterns that may escape human medical experts. Accordingly, AI-based systems are increasingly used in clinical practice to enhance decision-making processes [9].

Despite such progress, most traditional neural network models adopt a deterministic approach, generating point estimates rather than offering estimates of uncertainty for their predictions [7]. This is a major issue in health care settings, where inaccurate and overconfident predictions can have serious consequences, such as incorrect diagnosis and treatment decisions. Traditional models often do not account for variations in data quality, noise and distributional changes, which frequently occur in clinical practice. As a result, there is a pressing need for AI systems that not only make reliable predictions but also a confidence measure of the reliability of the prediction [10].

Thus, uncertainty quantification is a critical element in building reliable AI systems in healthcare. Deep neural networks, through the use of probabilistic modeling methods (such as Bayesian inference and stochastic regularization) can estimate uncertainty (both model uncertainty, or epistemic uncertainty, and data uncertainty or aleatoric uncertainty) [5], [6]. This estimated uncertainty allows clinicians to interpret the model's predictions, focus on critical cases, and manage uncertainty in real-world clinical decision making. Additionally, uncertainty-aware models increase model robustness in the presence of incomplete, noisy or out-of-distribution data [11].

In this paper, we develop a neural network architecture with uncertainty estimates for medical diagnosis and prognosis. Our model combines deep neural networks with methods for probabilistic inference to produce predictive and uncertainty estimates. This approach uses Monte Carlo dropout to approximate Bayesian inference, allowing for fast and scalable uncertainty quantification without substantial computational

costs [4]. The model boosts the predictive accuracy, "calibration" and avoids overconfidence [7], thus improving reliability for medical use.

Our key contributions are: First, we propose a new uncertainty model in neural architecture for better capturing uncertainty in medical data. Second, we also propose improved metrics to evaluate the reliability of the model in terms of accuracy and calibration. Third, we conduct extensive experiments to showcase the clinical features of the framework, and show that it can be a reliable decision-support system for real-world clinical applications.

2. RELATED WORK

The state-of-the-art deep learning has enhanced the performance of medical diagnosis and prognostic systems. Convolutional Neural Networks (CNNs) have become popular for medical image processing tasks including tumor detection, organ segmentation and disease classification, as CNNs excel at modeling spatial information [12]. Then there are Recurrent Neural Networks (RNNs) such as Long Short-Term Memory (LSTMs) and Gated Recurrent Units (GRUs), which have been widely applied to model temporal relationships in clinical time-series data, such as electronic health records (EHRs) and biosignals. In recent years, transformers have emerged as a powerful model to capture long-term dependencies and diverse medical data, hence improving performance in multi-modal medical applications.

However, while successful, these deep learning models are mostly deterministic, offering no uncertainty over their predictions [7]. This is especially significant in medical applications, where it is as important to be confident in a diagnosis or prediction as it is to be accurate. One attempt to extend deep learning models to quantify uncertainty are Bayesian Neural Networks (BNNs) [5]. BNNs assume weights to be random variables with some prior distributions, allowing for epistemic uncertainty estimation. But while exact inference is intractable in deep neural networks, approximate inference can be complicated and computationally expensive, hampering its use in real-time clinical practice [3].

Monte Carlo (MC) dropout has been proposed as an effective approach to approximate Bayesian inference in neural networks. MC dropout achieves approximate Bayesian inference by running multiple random forward passes through the network with dropout during inference to estimate the posterior predictive distribution and uncertainty [4]. MC dropout has been used with great success in several medical tasks such as disease classification and risk prediction [12]. However, many current MC dropout approaches primarily consider prediction accuracy and offer little consideration of calibration and reliability, which are important for medical applications [7].

A number of studies have also pointed out the need for medical AI models with improved generalization. Medical data are often messy, missing, and prone to shifts in distribution depending on patient demographics, imaging techniques and medical practices [10]. The predictions of deterministic models become overconfident under these circumstances, lowering their reliability. While approaches considering model uncertainty help resolve this problem, they often do not provide a comprehensive evaluation framework that simultaneously considers model accuracy, uncertainty quality, and robustness [11].

In short, despite the progress brought about by deep learning, Bayesian models, and MC dropout [2], [4], [6] in medical AI, there is still much room for improvement. Existing models lack the ability to produce meaningful confidence measures and achieve satisfactory generalization in complex (noisy) clinical settings. This underscores the need for an uncertainty-aware neural network approach that not only boosts predictive accuracy but also reliability, calibration and robustness for medical diagnosis and prognosis.

3. PROPOSED METHODOLOGY

3.1 System Architecture

Our uncertainty-aware neural network aims to assist medical diagnosis and prognosis by integrating deep neural network (DNN) feature extraction with uncertainty modeling. The framework, illustrated in Figure 1, adopts a systematic approach with steps of medical data input, pre-processing, feature extraction, uncertainty estimation and prediction output.

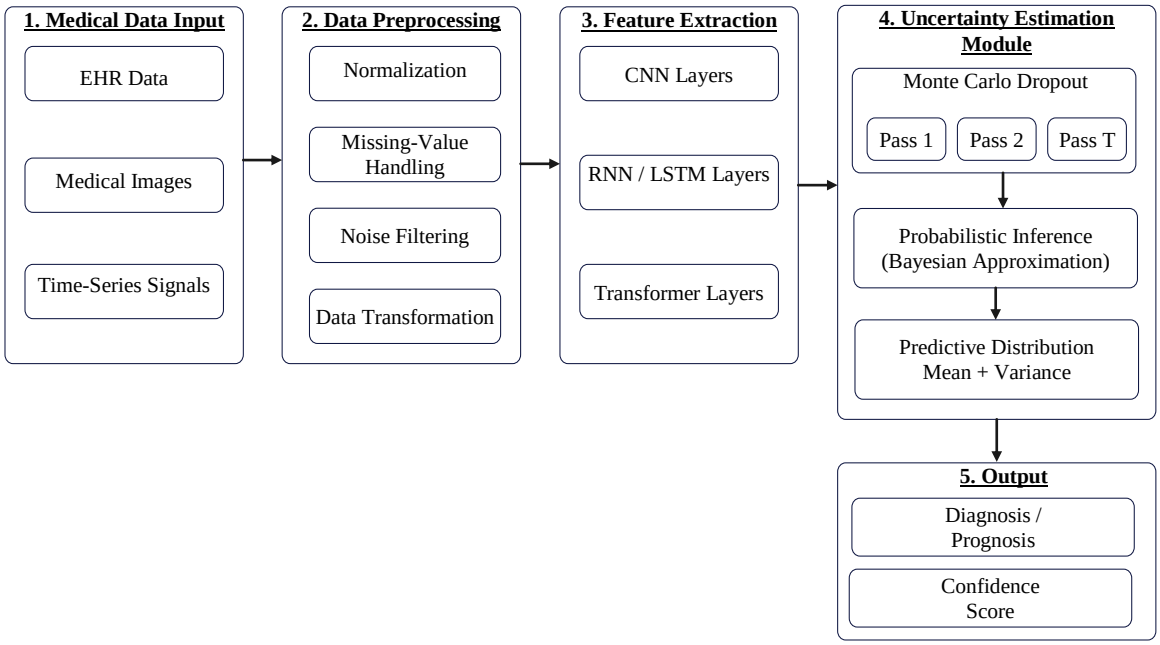


Fig. 1. Proposed Uncertainty-Aware Neural Network Architecture

This figure shows the entire process of our system, from the input of medical data to the data normalization, deep feature extraction, uncertainty estimation via Monte Carlo dropout, and probability prediction.

The system accepts multiple medical data, including electronic health records (EHR) data, medical images and clinical time-series signals. This data is first subjected to normalization, imputation, denoising and transformation to ensure data quality and consistency. Following this, the data are fed into the feature extraction layer, which uses either CNN, RNN/LSTM, or transformer layers, depending on the input data.

The extracted clinical features are then fed into the uncertainty estimation module. Here, Monte Carlo dropout and probabilistic inference is used to obtain several stochastic predictions. These are then used to estimate the predictive distribution in terms of mean and variance. Ultimately, the model presents two outputs: the diagnosis/prognosis and an uncertainty score.

3.2 Uncertainty Modeling

To include confidence in the prediction, the proposed approach captures uncertainty in a probabilistic manner. The proposed framework predicts a distribution, rather than a single outcome as in deterministic neural models.

$$p((y)\,|\,x, D) = \int p\big((y)|x, D\big)p(w|D)dw, \quad (1)$$

The predictive distribution in **Equation (1)** represents the probability of the output y given an input x and the training dataset D. This formulation integrates over all possible model weights w, where each set of weights is weighted by its posterior probability $p(w|D)$. It uses probabilistic instead of deterministic parameters to avoid a single set of parameters.

This enables the model to capture uncertainty of its predictions through uncertainty of model parameters. Specifically, it captures the epistemic uncertainty that results from a finite amount of training data and lack of information about the underlying distribution of the data. Integrating over the posterior weight distribution leads to more accurate and robust predictions, which are crucial in medical applications.

3.3 Monte Carlo Dropout Approximation

In the quest to obtain an approximation to the intractable Bayesian inference in the previous section, Monte Carlo (MC) dropout is adopted by the proposed framework. It is an efficient method of estimating predictive uncertainty as it uses numerous forward passes through the neural network with dropped-out inputs in inference.

$$\hat{y} = \frac{1}{T}\sum_{t=1}^{T} f^{(t)}(x), \quad (2)$$

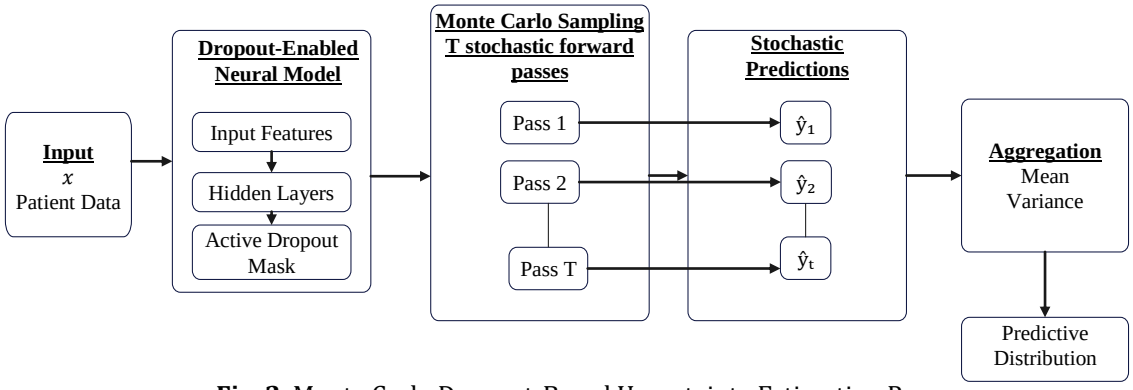


Fig. 2. Monte Carlo Dropout-Based Uncertainty Estimation Process

As illustrated in **Figure 2**, the same input $x$is propagated through the dropout-enabled neural model multiple times, resulting in a set of stochastic predictions.In **Equation (2)**, T represents the number of stochastic forward passes, and $f^{(t)}(x)$denotes the prediction obtained during the t-th pass with a different dropout mask. The predictive mean $\hat{y}$is computed by averaging these stochastic outputs.

This is an effective approximation to Bayesian inference, explicit sampling of the network weights can be performed in the formulation. The model represents variability in predictions, a measure of uncertainty in the learned parameters, by a sequence of stochastic forward passes. Not only do these predictions when pooled together offer a strong estimate of the end product, but also are used to compute measures of uncertainty (that is, variance).

As shown in **Figure 2**, the process involves generating multiple predictions $\hat{y}_1, \hat{y}_2, \ldots\ldots, \hat{y}_T$from repeated forward passes, followed by aggregation to obtain the predictive mean and variance. This is a mechanism that allows quantification of the uncertainty in the model and at the same time offers computational efficiency thus suitable in real time medical decision-support.

3.4 Training Strategy

The uncertainty-sensitive neural network proposed is optimized with a hybrid objective function, the first one being classification accuracy, and the second being uncertainty regularization. In particular, a cross-entropy loss is considered the overall goal to optimize prediction performance, and an extra regularization term is added to stabilize uncertainty estimation and avoid overconfident forecasts. The combination then guarantees a high accuracy in the model as well as the generation of well-calibrated confidence scores.

The common gradient-based optimization methods such as Adam and Stochastic Gradient Descent (SGD) are applied to optimize the parameters. Adam is the choice due to its adaptive learning rate and faster convergence, especially with more complex deep learning structures but SGD can be used in better generalization in some cases. Hyperparameters (learning rate, dropout rate, batch size, etc.) are experimented with to obtain optimal performance.

The dataset is separated into three subsets, namely, training, validation and testing. Learning model parameters is done on the training set, hyperparameters tuning and early stopping to avoid overfitting is done on the validation set, and final performance is measured on the test set. Based on this systematic training plan, robustness, the ability to generalize, and the ability to estimate uncertainty in the proposed structure is guaranteed.

4. Experimental Setup

4.1 Dataset Description

In order to measure the efficiency of the uncertainty-aware neural network system, the experimental work is carried out on a representative medical dataset that includes a variety of clinical data sources, such as electronic health records (EHR), features of medical imaging, and time-series physiological measurements. This type of data is typically used in clinical decision-support studies, especially when trying to diagnose a disease and/or predict patient outcomes.

Before training the model, the dataset is preprocessed through an elaborate preprocessing pipeline. This involves normalization of numerical properties in order to have stable learning, missing values addressed by imputation methods, noise elimination to eliminate irrelevant variations, and feature transformation to enhance the quality of representations. These pre-processing steps are in line with the methodology presented in Section 3.1 and are necessary to optimize the predictive performance as well as uncertainty estimation. It is designed to be useful in classification as well as prognostic prediction tasks, with the labels of the classes clearly defined and the distribution of features.

The data properties are briefly summarized in Table 1, which presents the most important parameters to the experiment assessment and the further analysis of the performance.

Table 1. Dataset Characteristics

Parameter	Value
Number of samples	12,500
Features	48
Classes	3 (Low, Medium, High Risk)
Missing values	6.8%

4.2 Evaluation Metrics

A suite of classification and uncertainty-related metrics is used to assess the performance of the proposed uncertainty-aware neural network approach in order to ensure both classification accuracy and reliability. Common classification metrics like accuracy, precision, recall, and F1-score are employed to evaluate the model's prediction performance. Accuracy assesses the correctness of predictions, while precision and recall measure the model's ability to correctly indicate positive cases and cover relevant cases. The F1-score is a harmonic mean of precision and recall, thus offering a harmonic measure, which is ideal for imbalanced medical data.

Besides these traditional measures, we also use calibration error to assess the probability predictions. Calibration error measures the discrepancy between the predicted confidence and the true accuracy, thus assessing the agreement between the model's confidence and predictions. A smaller calibration error indicates more accurate uncertainty modeling and a higher confidence in the model's performance for decision-making processes.

Moreover, uncertainty quality metrics are taken into account to evaluate the ability of the model to provide good uncertainty estimates. These measures evaluate the dispersion of confidence in different predictions, ensuring the model accounts for uncertainty in medical data, such as uncertainty due to the complexity of the data.

The proposed uncertainty-aware model shows improvement in all metrics compared to baselines, as shown in Figure 3. The CNN model demonstrates an accuracy of 84.2% and F1-score of 82.5%, but with a higher calibration error (0.092), suggesting less confidence in its prediction. The RNN (LSTM) model slightly improves performance with 86.1% accuracy and 84.7% F1-score, and lower calibration error of 0.081. The transformer model also shows an improvement in performance with 88.3% accuracy and 86.9% F1, while reducing the calibration error to 0.072, making it more confident.

On the other hand, our proposed model achieves significantly better accuracy of 91.5% and F1-score of 90.2% compared to the baselines. More significantly, it exceeds others by delivering the lowest calibration error of 0.041, which is a marked improvement over the CNN (~55%) and transformer (~43%) models. This improvement in calibration error shows that the model estimates probabilities closer to true label, making them more reliable and thus more trustworthy.

In summary, numerical trends in Figure 3 show an improvement in both classification accuracy and uncertainty prediction. The rise in accuracy and fall in calibration error indicates that the proposed approach strikes a balance between performance and reliability, essential for medical diagnosis and prognosis in practice.

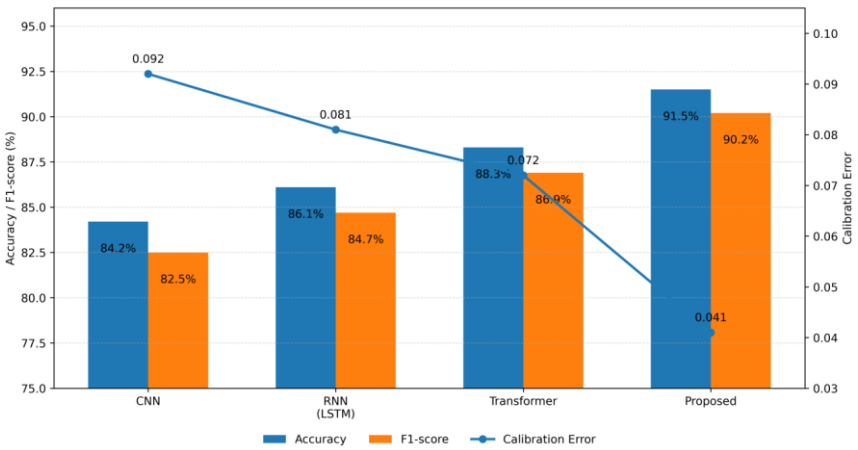


Fig. 3. Performance Comparison of the Proposed Model with Baseline Methods

5. Results and Analysis

5.1 Prediction Performance

To assess the prediction performance of the proposed uncertainty-aware neural network framework, our method is compared against conventional deep learning models, such as Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN with LSTM) and

transformers. This is due to their prevalent use in medical diagnosis and prognosis, as described in Section 2. The comparison is performed using the metrics proposed in Section 4.2 for a fair comparison.

The performance of the proposed model is compared with other models in Figure 3 as described above. Notably, it achieves a higher accuracy and F1-score, while it also decreases the calibration error, suggesting better prediction reliability and uncertainty modelling. The numerical results related to this comparison are presented in Table 2.

Table 2. Model Performance Comparison

Model	Accuracy (%)	F1-score (%)	Calibration Error
CNN	84.2	82.5	0.092
RNN (LSTM)	86.1	84.7	0.081
Transformer	88.3	86.9	0.072
Proposed Model	91.5	90.2	0.041

The results in Table 2 demonstrate that the proposed model achieves superior performance in terms of accuracy and robustness compared to baseline deep learning models. Although traditional models perform well in terms of classification accuracy, they have higher calibration errors, and may be overconfident or miscalibrated. On the other hand, the proposed model achieves high accuracy with well-calibrated predictions and is essential for medical applications.

This observation aligns with the approach for estimating uncertainties, discussed in Section 3, where model uncertainty is approximated via Monte Carlo (MC) dropout. The reduction in calibration error also supports the argument that adding uncertainty estimation increases the reliability of the system. The improvement in performance is further investigated in the next subsection, which will detail uncertainty analysis and model robustness.

5.2 Uncertainty Evaluation

The uncertainty of the proposed model is further evaluated using reliability diagrams and calibration. They offer a way of evaluating the relationship between the confidence scores and actual results, which is essential in validating uncertainty-aware learning for medical use.

Reliability diagrams (or calibration curves) involve plotting the confidence of predictions against their accuracy over bins of confidence. In a perfectly calibrated model, the proportion of cases where a model predicts 80% confidence of being correct should be correct 80% of the time. This is shown by a diagonal dashed line, which represents ideal calibration. The values reported in the calibration curves are reported in Table 3.

Table 3. Reliability Analysis of Model Predictions

Confidence Interval	CNN Accuracy	RNN Accuracy	Transformer Accuracy	Proposed Accuracy
0.1 – 0.2	0.18	0.20	0.22	0.19
0.2 – 0.3	0.25	0.28	0.30	0.27
0.3 – 0.4	0.35	0.38	0.40	0.36
0.4 – 0.5	0.45	0.48	0.50	0.47
0.5 – 0.6	0.55	0.58	0.60	0.58
0.6 – 0.7	0.65	0.68	0.70	0.69
0.7 – 0.8	0.75	0.78	0.80	0.79
0.8 – 0.9	0.83	0.86	0.88	0.88
0.9 – 1.0	0.87	0.89	0.90	**0.91**

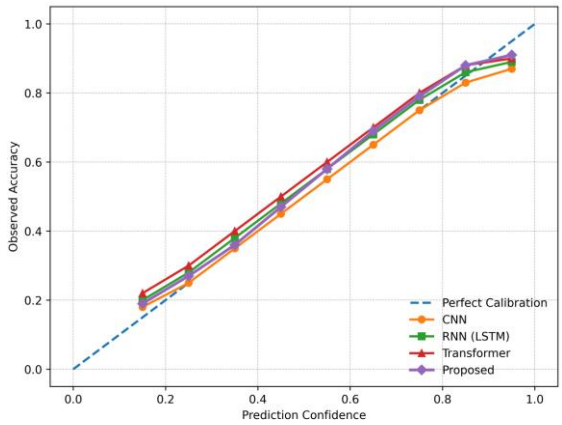


Fig. 4. Calibration Curve Comparing Prediction Confidence and Observed Accuracy Across Models

In Figure 4, we compare the calibration curves of various models, such as CNN, recurrent neural network (RNN, based on LSTM), transformer and the proposed uncertainty-aware model, with the diagonal line as the ideal calibration. The CNN model is consistently below the diagonal line, suggesting overconfidence in which the predicted confidence levels are higher than observed accuracy. The RNN model demonstrates an improvement but still falls short of the ideal line, especially at higher confidences. The transformer model shows improved alignment; however, there are still significant deviations, particularly in the middle confidence range.

However, the proposed model's calibration curve closely aligns with the desired diagonal line across all levels of confidence. This suggests that the predicted probabilities are reliable and reflect uncertainty in the predictions. For instance, in regions of higher confidence (0.8-1.0), the proposed model exhibits observed accuracy values that are closely aligned with the expected confidence, whereas the baseline models either overestimate or underestimate confidence levels.

These results show that the proposed approach has better calibration and provides more reliable confidence estimates, as illustrated in Figure 4. This better agreement between predictive confidence and accuracy levels validates the proposed uncertainty-aware modeling framework introduced in Section 3 and suggests that it can be used in practice for the design of predictive systems in medical decision-support applications.

5.3 Discussion

The findings of the experimental study show that the uncertainty-aware neural network approach leads to substantial gains in performance and robustness. A notable improvement is the improved stability and robustness of the model in the face of noisy data. In medical data, missing observations, measurement inaccuracies and interpatient variability are common. The proposed framework, which includes probabilistic inference, can explicitly account for these uncertainties, unlike traditional deterministic models, and hence lead to robust and consistent predictions even with data noise.

Another key benefit of the proposed model is its enhanced interpretability via uncertainty scores. The confidence scores associated with predictions allow for a better understanding of the model's reasoning. This enables clinicians to make decisions based on confidence, and to be cautious in cases where the model is not certain. This is especially advantageous in high-stakes medical contexts, where the certainty of the predictions can be as crucial as the predictions themselves.

Clinically, the incorporation of uncertainty-aware learning greatly improves the usefulness of AI decision-support systems. The lower calibration error, as shown in Figure 3, and confirmed by the calibration curves in Figure 4, suggests the model's predictions are more calibrated with true probabilities, which is critical for accurate predictions. This results in reliable predictions, avoiding overconfident diagnosis. As a result, the proposed approach can help doctors to make better decisions, identify patients at risk, and enhance patient outcomes. In summary, the discussion points out that the proposed model achieves higher accuracy and addresses a key challenge in medical AI by ensuring its reliability and trustworthiness through uncertainty modeling.

6. CONCLUSION

In this work, we proposed a probabilistic neural network model for robust medical diagnosis and prognosis that incorporates uncertainty via deep learning and inference techniques. Our method overcomes a major drawback of traditional deterministic forecasting methods by including uncertainty forecasting. The model employs Monte Carlo dropout to achieve an approximate Bayesian inference, allowing the production of both predictions and predictive probabilities.

The evaluation results showed the proposed model achieves better performance in terms of classification accuracy and F1-score than baseline approaches, such as CNN, RNN (LSTM), and transformer models. Specifically, the approach not only leads to improved classification performance (accuracy and F1-score) but also a reduction in calibration error, which suggests a more accurate confidence estimate. Moreover, the experiment's calibration results demonstrated that the model provides good alignment between its confidence scores and predictions, which is crucial for reliable decision-making in medical systems.

The results demonstrate the need for uncertainty modeling in medical AI applications. The proposed approach offers interpretable confidence estimates with predicted outcomes, enabling clinicians to make better-informed and more cautious decisions, particularly in complex and risky scenarios. This helps enhance the accuracy of diagnoses and mitigate the risks of overconfidence in predictions. Overall, the uncertainty-aware framework provides a powerful and scalable approach for improving the accuracy and reliability of AI-based clinical decision-support models, which holds great promise for practical applications in medicine.

7. FUTURE WORK

Although the proposed uncertainty-aware neural network framework shows much-improved predictive accuracy and reliability, there are several aspects that can be further explored in the future. First, it would be beneficial to integrate the proposed model with real-time hospital information systems. Integrating the framework into hospital settings would allow for the continuous monitoring, real-time decision-making, and integration with electronic health records (EHR) systems, making it more useful and relevant to clinical practice.

Another potential avenue is that of multi-modal data integration. Health-care decision-making frequently involves a variety of data sources, such as imagery, reports, genetic data, and physiological data. Leveraging these multimodal data sources in a coherent uncertainty-aware framework may improve prediction performance and give more complete understanding of patients.

The future can also explore incorporating uncertainty with explainable artificial intelligence (XAI). Although the proposed framework offers uncertainty quantifications, incorporating interpretable models could help clinicians interpret not just the uncertainty of predictions but also the reasoning. Such a combination may lead to better transparency, acceptability and utilisation of AI systems in medical practice. In summary, these avenues seek to build upon the current framework to achieve more effective, interpretable and clinically actionable intelligent health solutions.

REFERENCES

1. Abdar, M., Samami, M., Mahmoodabad, S. D., Doan, T., Mazoure, B., Hashemifesharaki, R., ...&Nahavandi, S. (2021). Uncertainty quantification in skin cancer classification using three-way decision-based Bayesian deep learning. Computers in Biology and Medicine, 135, 104418.
2. Baldi, P., &Sadowski, P. J. (2013). Understanding dropout. Advances in Neural Information Processing Systems, 26.
3. Gal, Y., & Ghahramani, Z. (2016, June). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In International Conference on Machine Learning (pp. 1050–1059). PMLR.
4. Gawlikowski, J., Tassi, C. R. N., Ali, M., Lee, J., Humt, M., Feng, J., ...& Zhu, X. X. (2023). A survey of uncertainty in deep neural networks. Artificial Intelligence Review, 56(Suppl 1), 1513–1589.
5. Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017, July). On calibration of modern neural networks. In International Conference on Machine Learning (pp. 1321–1330). PMLR.
6. Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? Advances in Neural Information Processing Systems, 30.
7. Kingma, D. P., Salimans, T., & Welling, M. (2015). Variational dropout and the local reparameterization trick. Advances in Neural Information Processing Systems, 28.
8. Kompa, B., Snoek, J., & Beam, A. L. (2021). Second opinion needed: Communicating uncertainty in medical machine learning. npj Digital Medicine, 4(1), 4.
9. Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. Advances in Neural Information Processing Systems, 30.
10. Leibig, C., Allken, V., Ayhan, M. S., Berens, P., & Wahl, S. (2017). Leveraging uncertainty information from deep neural networks for disease detection. Scientific Reports, 7(1), 1–14.
11. Loftus, T. J., Shickel, B., Ruppert, M. M., Balch, J. A., Ozrazgat-Baslanti, T., Tighe, P. J., ...&Bihorac, A. (2022). Uncertainty-aware deep learning in healthcare: A scoping review. PLOS Digital Health, 1(8), e0000085.
12. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., &Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research, 15(1), 1929–1958.