10.36647/IJSEM/12.05.A001
ISSN (Online) 2456 -1304
International Journal of Science, Engineering and Management (IJSEM)
Vol 12, Issue 05, May 2025

Machine Learning Algorithms in Personalized Life Expectancy Prediction: A Systematic Review

[1] Vladyslav Malanin

[1] PhD student, M.D, V.M. Glushkov Institute of Cybernetics of the NAS of Ukraine, Kyiv, Ukraine
Email ID: [1] vladislav.malanin@gmail.com

***Abstract*— Hyperforecasting allows medical professionals and individuals to make decisions about health care, resources, and planning. Yet, actuarial and statistical science methods are generally ineffective in considering various unique, unexpected, and nonlinear components of longevity. New ML algorithms have appeared lately, permitting life expectancy estimates that can adjust to each individual and change as new information comes in. This review tested how Random Forest, Extreme Gradient Boosting, and Neural Networks could forecast each individual's life expectancy in various clinical and demographic settings.**

***The review collects results from peer-reviewed papers published from 2015 to 2025 in PubMed, IEEE Xplore, and Scopus databases. Accuracy, adaptability, how well the machine can be interpreted, and whether it can be easily integrated are measured. Random forests are respected for being sturdy and readable, which is not something XGBoost is known for, but they seem more accurate and performant when using complex data. The greatest potential of deep learning networks is modeling difficult data, but they are less transparent and demand lots of information to function properly.**

ML models are more accurate than many older techniques, yet there are problems with explaining them, ensuring their ethics, and being fair. By reviewing these findings, we can recognize that it's important for AI to be clinically understandable and highlight designs that may improve them in the future.

***Index Terms:* life expectancy prediction, machine learning, Random Forest, XGBoost, neural networks, deep learning, personalized healthcare.**

I. INTRODUCTION

For many years, public health and clinical fields have depended on forecasting how long a person will live. The time a person may live influences decisions on personal risks and national policies. Usually, they base these on typical trends in the population by using methods like linear regression and the Cox proportional hazards model. Still, these tools provide valuable knowledge, but they cannot account for situations when people's health does not follow the same pattern because they assume every hazard affects people similarly.

For this reason, ML methods are being used today to improve the process of life expectancy forecasting. AI Illuminates works best on data with numerous variables, is present in many formats, and can soon change. With help from electronics in healthcare, smart devices, genetic analysis, and technologies for monitoring health, researchers can accumulate data for precise results. Unlike the models used in traditional analytics, machine learning can study the data more effectively and detect complex trends and changes as they occur.

Researchers often use RF, XGBoost, and NNs to examine life expectancy. All these teaching methods have their benefits. Random Forest is praised for being easy to use and understand, as well as being very reliable. It handles planned data well and will probably not develop overfit. Because XGBoost scales and helps reduce over-fitting, it is praised for being fast, accurate, and able to handle missing data. Deep learning models and other Neural Networks are widely applied to understand different relationships in survival analysis.

Nevertheless, applying ML to life expectancy brings about many notable issues. Understanding and trusting how a model in health care operates is very valuable. Additionally, people are concerned about making AI systems that are not fair or not biased. Since death predictions use algorithms, it is essential to look at them from the points of view of both technology and society.

The review examines how well ML applications predict a person's life expectancy. The analysis mainly discusses the differences in how accurate Random Forest, XGBoost, and Neural Networks are, their uses in most situations, and the ease with which their code can be read and put into practice. Since 2015, many books and articles have been written about these methods, and the article mentions what they have and have not achieved and identify areas that still require further study. Therefore, the review expands the understanding of changing from standard risk scores to personal and ethical prediction tools.

Overall, this work ensures ML models are accurate and trustworthy by encouraging people to focus on their usability, fairness, and predictive capabilities. Since machine learning is used for important decisions involving life and death, knowing what it can and cannot do is crucial.

ISSN (Online) 2456 -1304
International Journal of Science, Engineering and Management (IJSEM)
Vol 12, Issue 05, May 2025

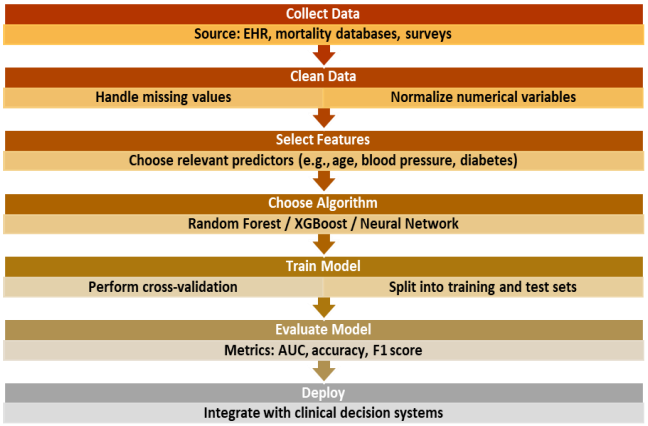


Fig. 1. Machine Learning Workflow for Life Expectancy Prediction

II. RANDOM FOREST IN LIFE EXPECTANCY PREDICTION

Before, Random Forest was the main tool for machine learning tasks in biomedicine and clinical prediction. Breiman developed the Random Forest algorithm in 2001 by making several decision trees during training and selecting the most common class (classification) or average prediction (regression) from the trees to output the overall result. Its ability to work with unclear and nonsimple patterns, along with its usefulness and meaningfulness, makes it the preferred option for calculating life expectancy for each person.

Random Forest has been effective in mortality and life expectancy because it handles high-dimensional and diverse data, often seen in modern healthcare. Missing values, variable priority, and complex relationships among the variables do not trouble the algorithm. This area is significant in health settings, given that available data are incomplete, noisy, and a mixture of categories involving many variables like demographics, biomarkers, basic illnesses, genes, and symptoms.

Manuscripts have established that Random Forest can be used to forecast life expectancy in various medical scenarios and among different groups of people. Scientists in this study developed RF models that aimed to predict survival for cancer patients and obtained accuracy above 0.8, proving that their predictions were almost the same as the actual results. Other investigations have used RF to evaluate the risk of heart disease death, the weakness commonly seen in the elderly, and the life spans of people living with diabetes, kidney disorders, and HIV/AIDS. In general, RF performed as well as or even better than Cox proportional hazards models and logistic regression, often in data with nonlinear curves and many connections between variables.

Random Forest has an important advantage in this field because of its variable importance measurements. With the algorithm, it is clear which features are most important for prediction, explaining to doctors how results are achieved. In the case of older individuals, RF algorithms have repeatedly pointed out that age, number of diseases, certain biomarkers (albumin and white blood cells), and health conditions rank among the main determinants of life expectancy. They can strengthen decision-making and identify risk factors that may cause people to change.

Still, there are some drawbacks to the effectiveness of Random Forest. RFs are normally simpler to explain than neural networks, though they act as mysterious ensembles when many trees are included. If the variables are correlated, the significance of each can look different from reality. Moreover, interpreting partial dependence plots always needs proper care. In addition, if the outcome we are interested in is heavily concealed due to censoring or where there is lopsided data, the RF performance may fall significantly.

Examining whether radiofrequency models can be used with different groups is necessary. However, the algorithm's ability to prevent overfitting comes from being built as an ensemble, yet non-representative and biased data can still affect its prediction results. If the data used to train a Random Forest model comes from urban hospitals, it may lead the model to provide accurate estimates only for those areas. It is necessary to use extra data and look into groups of patients to validate such models and ensure equal results for everyone.

Experts in RF have recently tried to improve how RF can be understood and how it works in survival analysis. Random Survival Forests adjust splitting and prediction to include the information from time-to-event data and handle censoring. Numerous works have shown that survival functions, hazard functions, and average survival times can be measured by RSF better than by using standard RF when predicting mortality rates. In addition, designers have developed methods that merge RF with new algorithms for selecting key features, imputing missing information, and building blended ensembles, including RF and other popular models.

All in all, Random Forest is relied on as an effective and useful approach for predicting life expectancy in medical science and research. Managing nonlinear interactions and avoiding much preprocessing attracts it to the personalized prognosis field. Additionally, the method uses its methods to evaluate which features are important, helping doctors infer meaningful predictions from the system. Still, fairness, validation, and explainability must be considered to reach its full benefits. As NLP-generated Big Data increases, Random Forest is expected to contribute to developing personalized life expectancy models, though only under careful thought about the model's strengths, accuracy, and usefulness for doctors.

ISSN (Online) 2456 -1304

International Journal of Science, Engineering and Management (IJSEM)

Vol 12, Issue 05, May 2025

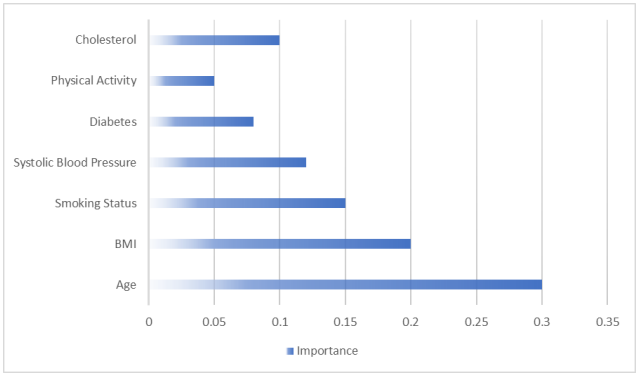


Fig. 2. Random Forest Feature Importances

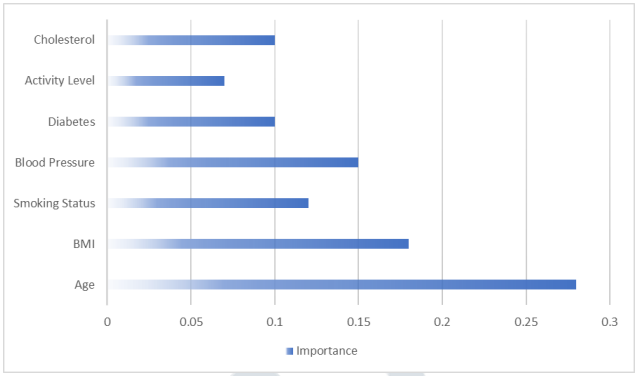


Fig. 3. XGBoost Feature Importances

III. XGBOOST IN LIFE EXPECTANCY PREDICTION

Many experts now consider XGBoost, which is short for Extreme Gradient Boosting, to be a leading algorithm for supervised machine learning, mainly for tasks such as predicting health or survival. In 2016, XGBoost was introduced by Tianqi Chen and Carlos Guestrin as a boosted-gradient algorithm that uses regularization, parallelization, and efficient calculations. Because of these characteristics, it performs best on structured data based on speed, scalability, and high performance. According to studies, XGBoost performs better than similar methods when predicting human life expectancy.

Its main strength is that it takes a step-by-step approach to making decisions. When similar trees are added to correct past mistakes, the model improves loss functions and achieves fewer trees than usual Random Forests. Also, XGBoost uses techniques called L1 and L2 penalties to avoid problems related to overfitting, as this issue often comes up in predicting medical outcomes using data with lots of features yet are sometimes sparse and may have skewed class distribution. In life expectancy estimation, it is especially important since censorship, non-normal survival patterns, and diverse risk levels are usual.

Many researchers have relied on XGBoost to estimate how long an individual may live, depending on various clinical and demographic facts. In oncology, experts have tested the algorithm to guess the survival of those with lung, breast, and colorectal cancers using EHRs, aspects from scans, and genetics. According to AUC and C-index, XGBoost has been shown to perform better than the Cox regression and logistic models. In this branch of medicine, doctors have successfully used XGBoost to predict deaths from heart failure. Such applications explain how an algorithm gathers various health details and produces useful information on a person's risk factors.

In addition, XGBoost is beneficial for life expectancy modeling since it handles missing data well and allows using many types of input data. Unlike some artificial intelligence (AI) algorithms that depend on extensive handling of missing values, XGBoost learns how to fill in those lost items during tree building by default. The value of this feature increases in clinical data because such data often has missing information due to differences in how information is recorded or patient noncompliance. On top of that, the framework makes it clear which aspects of the data have the most impact on what the model predicts.

Still, while it performs well, XGBoost cannot model the time events occur, which prevents its direct use in survival analysis. As a result, the algorithm has been changed to support survival loss functions or used together with the Cox regression model. To illustrate, some survival-XGBoost models improve how Cox's model is fitted without sacrificing the advantages of the boosting approach. This enables XGBoost to analyze censored records and provides a life expectancy estimate and the curves or functions for survival distribution. Many current methods, such as XGBSE (XGBoost Survival Ensembles), have Python libraries designed for them and are gaining ground in clinical ML.

Additionally, XGBoost allows for the convenient validation of models and testing their strength. Cross-validation and sampling approaches allow practitioners to adjust the model and assess its performance when used on different people. As a result of XGBoost's efficiency, these steps can be applied to big datasets or to many other groups, such as people from various age groups and ethnic backgrounds or suffering from several types of diseases.

However, understanding how they function can be challenging, as with other ensemble methods. Even though XGBoost is more straightforward than neural networks, the technique can be difficult to know when many trees and features interact. Researchers have started using SHAP and LIME to provide a better picture of how a model makes each prediction. In a survival model, SHAP values help analyze which features increase or decrease the expected survival for each patient.

ISSN (Online) 2456 -1304
International Journal of Science, Engineering and Management (IJSEM)
Vol 12, Issue 05, May 2025

However, XGBoost is unsuited for fields where most data comes in unstructured forms, such as images, texts, or time-related recordings. Some studies have shown that using XGBoost with NLP to handle clinical notes or CNNs to extract valuable information from medical images has worked well. Thanks to these methods, the model can better predict life expectancy based on various data collected about a person.

XGBoost is a highly flexible option among machine learning methods for making life expectancy predictions. Telemedicine is very useful because it is speedy, accurate, and works well under pressure. When XGBoost is fitted for survival analysis and paired with some explainability methods, it becomes a model that performs well and can be understood. Since healthcare prioritizes tailored and data-based care, XGBoost is expected to play a major role in creating helpful and personalized tools for physicians.

Table I: Summary of Machine Learning Techniques in Life Expectancy Prediction

Algorithm	Key Features	Strengths	Weaknesses	Typical Use Cases
Random Forest	Ensemble of decision trees	Robust to overfitting, interpretable	Less effective with sparse data	Structured clinical datasets
XGBoost	Gradient-boosted decision trees	High accuracy, handles missing data	Hyperparameter tuning complexity	Mortality prediction, comorbidity modeling
Neural Networks	Layers of interconnected nodes	Captures non-linear relationships	Requires large datasets, less interpretable	EHRs, genomics, imaging
Deep Learning	Deep multi-layered architectures	Learns from unstructured/multi-modal data	"Black-box" nature, high compute demand	Time-series EHRs, NLP, imaging analytics

IV. NEURAL NETWORKS AND DEEP LEARNING IN LIFE EXPECTANCY PREDICTION

Thanks to deep learning neural networks, new ways of predictive modeling have seen widespread use in several areas, such as healthcare. Since they can represent difficult, curved links between many factors without using explicit features, they work well to predict life expectancy, including the body, behavior, environment, and health care. Since Random Forest or XGBoost only employs simple decision rules, neural networks can still learn complex patterns using multiple transforming layers. As a result, discoveries may be made about hidden patterns normally overlooked.

The basis of neural networks is that artificial neurons are organized in layers called input, hidden, and output; each type computes the weighted sums of its inputs and then uses a nonlinear function. Having larger hidden layers in deep learning models allows them to learn from simple to more complex aspects gradually. Working with EHRs, imaging data, genomics, and continuous monitoring signals in assessing life expectancy is very beneficial, thanks to the hierarchical structure described above.

Over the past few years, different types of neural networks have been used in life expectancy models based on the input data. FNNs have been used where inputs include demographics, tests, and medical histories. On average, the examples cited surpass the results seen with linear regression or Cox proportional hazards models when forecasting mortality for older populations and individuals with chronic diseases. It has been shown that, by applying FNNs, the prediction of short- and long-term death in ICU patients achieved AUC values typically better than 0.90.

Among the neural network types, RNNs have performed well, with LSTM and GRUs serving as effective choices when dealing with a series of health measurements over time. They are especially helpful when predicting life expectancy based on blood pressure charts, how often medication is followed, or sensor readings. RNNs can identify how recent changes in health indicators can affect the likelihood of patients living longer. In a few cases, attention mechanisms are added on top of RNNs to pay more attention to influential previous data and improve how accurately they can interpret and predict events.

CNNs, commonly linked to analyzing images, have lately been applied to predicting life expectancy, mainly in radiology. Researchers have trained CNNs using X-ray and CT images, allowing them to estimate the chances of death based on identifying features such as abnormal tissues, irregular lungs, and artery calcification imaged in the scan. Typically, images are combined with other medical records using different models to estimate risk better. Research has shown that CNNs can be used to study ECG data and similar signals, highlighting fine details important to understanding a patient's health outcome.

Something very powerful about deep learning is its ability to learn and predict results straight from raw data without requiring software users to analyze the data before training. By doing this, the model can identify hidden risks from clinical notes, voice transcripts, or DNA data. Researchers have found that using transformer architectures (like BERT, GPT, BioBERT) in NLP helps identify predictive factors from physician notes, discharge notes, and general descriptions of social circumstances. Changing the text to high-dimensional lists helps the models discover patterns in

ISSN (Online) 2456 -1304
International Journal of Science, Engineering and Management (IJSEM)
Vol 12, Issue 05, May 2025

information that strongly impact life expectancy yet can't be easily described using structured rules.

There are difficulties associated with using neural networks in healthcare. AI companies need to be more transparent. Doctors prefer not to employ deep learning in places where understanding the outcome is as significant as obtaining accurate results because they cannot easily explain the reasoning. Even if LSP, SHAP, and saliency maps show how some predictions are made, the explanations can still be vague or vary depending on the method. Relying on these algorithms could lead to trust issues, affecting treatment, palliative care, or allocating medical resources.

Sometimes, getting the needed data is a challenge. Usually, deep learning works best when fed with a lot of data that has been accurately annotated. Estimating life expectancy cannot be easy as healthcare institutions are divided, privacy concerns exist, and data is inconsistently logged. Still, transfer and federated learning are not easy to use since it is complicated for different healthcare systems to link together. Moreover, models trained with deep learning could end up carrying the biases shown in the training data. If data about certain people, such as those in cities, some races, or those who go to the doctor more often, is the main focus, then people who are not similar to these may end up affected negatively by its results. Issues mentioned above should be managed by regular monitoring, properly reducing bias, and paying attention to ethics while the technology is in use.

Yet, neural networks have the potential to be very effective in predicting how long patients will live. It has been observed that using data from deep learning and clinicians simultaneously can lead to better predictions than using only one or the other. Reconciling what is precise in models and physicians' clinical judgment can be achieved by merging model forecasts with doctors' intuition. Besides, the continuous growth of XAI, model distillation, and rule-based neural-symbolic systems might one day fill the interpretability space that prevents deep learning from being used in critical settings.

All in all, both neural networks and deep learning are changing how personalized life expectancy is estimated. They excel in analyzing hard-to-work-with datasets, making them the top candidates for AI-based health predictions. With progress in understanding deep learning models, more people will likely rely on them to make equitable, individualized, and well-informed decisions about what lies ahead for patients.

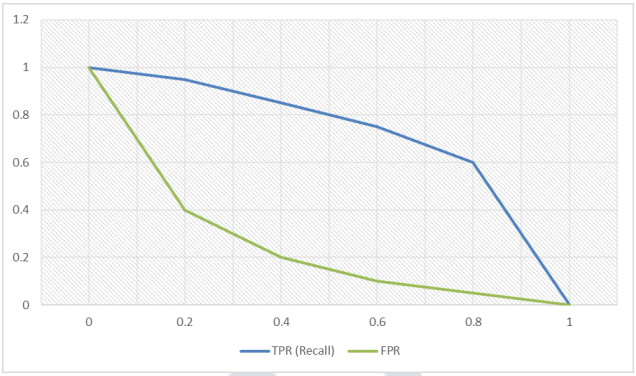


Fig. 4. ROC Curves for ML Models on Mortality Prediction

V. COMPARATIVE ANALYSIS AND SYNTHESIS

Improving ML methods has changed the way life expectancy can be predicted for individuals. RF, XGBoost, and neural networks all showcase different advantages and disadvantages, influenced by how they are constructed, processed data, and used in healthcare. Reviewing and comparing the models reveals each medical approach and its unique and matched advantages for specific groups of people. Combining multiple decision trees allows Random Forest to process data that fits into a structure, has various features, and has some missing entries, but it requires little preprocessing. It handles different health-related datasets successfully and is commonly preferred for its interpretation and stabilization ability. Clinicians find the model useful since it allows them to recognize the key features that strongly impact the output. Unfortunately, understanding high-dimensional or slightly complicated data, such as images, genes, or series of events, may not be possible with RF.

XGBoost is better than RF because it boosts the performance of trees by constructing them step by step and using advanced regularization. Overall, it gives better results than RF when the situation is not very noisy, and the data is not too complex. Due to its guard against overfitting, it is an excellent choice for programs working with real clinical data that is sometimes unreliable and poor in quality. XGBoost is known for working well, yet it cannot automatically learn features like neural networks and relies on features engineered by humans.

Neural networks and deep learning models can handle complex, not linear, information of different types better than other models. With deep learning, features that machines can understand, such as bits of text, data, and sensor feeds, are processed more thoroughly. Their edge is considerably greater when the analysis concerns extensive data or involves observing changes through time. Nonetheless, using deep learning leads to some serious problems: it eats up a lot of data, requires significant computing abilities, and can be hard to interpret, which can raise issues about how people will trust it.

ISSN (Online) 2456 -1304
International Journal of Science, Engineering and Management (IJSEM)
Vol 12, Issue 05, May 2025

Deep learning is often shown to be the most accurate, followed fairly closely by XGBoost and then RF. Under difficult conditions, RF and XGBoost could work more effectively since they demand less data and are faster to train. In addition, RF and XGBoost perform well on organized data, whereas neural networks perform notably in situations with many variables.

Another way to compare how effectively a therapy can be implemented and used in real-world settings is by using the following: Since RF and XGBoost are simple to understand and work with, they are more suitable for integration into healthcare systems. Even though deep learning models are very effective, they generally require specific platforms and strict testing before being used in key areas such as care for patients at the end of their lives. The model picked depends on the complicated process, how much data there is, and the particular concern being studied.

Notably, hybrid strategies are coming to use several algorithms in combination. RF, XGBoost, and deep learning outcomes were integrated, which helped accuracy while allowing the models to be used in other settings. They mean that using several techniques, each chosen according to where people come from could improve how we estimate differences in life expectancy.

In general, the best approach depends on how your data is structured if you require explanations of the results and the setting of your clinical application. More studies are needed to compare the results of AI methods across various groups of patients and ensure that their decisions are always made transparently. For healthcare to be fair and reliable, machine learning should adapt to every patient.

Table II: Summary of Machine Learning Techniques in Life Expectancy Prediction

Algorithm	Key Features	Strengths	Weaknesses	Typical Use Cases
Random Forest	Ensemble of decision trees	Robust to overfitting, interpretable	Less effective with sparse data	Structured clinical datasets
XGBoost	Gradient-boosted decision trees	High accuracy, handles missing data	Hyperparameter tuning complexity	Mortality prediction, comorbidity modeling
Neural Networks	Layers of interconnected nodes	Captures non-linear relationships	Requires large datasets, less interpretable	EHRs, genomics, imaging
Deep Learning	Deep multi-layered architectures	Learns from unstructured/multi-modal data	"Black-box" nature, high compute demand	Time-series EHRs, NLP, imaging analytics

VI. ETHICAL, LEGAL, AND PRACTICAL CONSIDERATIONS

Predicting an individual's life span with machine learning tools involves ethical, legal, and practical difficulties. As it becomes easier to prepare accurate prognostic forecasts with technology, the role of these forecasts in our lives goes beyond the latest computers. They discuss the most important problems in people's lives, including personal freedom, justice, privacy, trust, and the idea that there is a code of right and wrong.

When thinking about ethics, the biggest concern is the impact of life expectancy predictions on medical decisions and a patient's freedom. Telling patients or clinicians how much time they can expect improves decisions on various treatments, end-of-life care, and readiness to cope with death. When we can't explain the reasons behind the predictions or if they fail, this increases the risk that our decisions will fulfill them. As a result, it is important that AI-related technologies aid informed decision-making and do not impede it.

Machine learning models can also be biased, making that another important ethical issue. Historical healthcare data often causes AI models to reflect the lack of equal access, treatment, and outcomes between people of different races, socioeconomic groups, and genders. If the problem is not corrected, the biases can be put into predictive tools, leading to unequal estimates of life expectancy that might harm people from marginalized groups. Promoting fairness involves carefully monitoring, adding different datasets, and applying algorithms that prevent people from being judged differently. Managing AI ethically requires regularly paying attention and using any new data that is more inclusive.

Since ML models predict life expectancy, people wonder about responsibility and consent when using them. Should a negative result from an AI prediction occur, it is not always easy to decide if the developer, the organization deploying it, or the clinician who used it is to blame. Currently, there are not enough guidelines for regulatory bodies to ensure patients who experience algorithmic issues in healthcare can hold anyone accountable. Besides explaining the diagnosis, doctors must also clarify how the diagnosis was reached through an algorithm.

It is also not easy to implement because of other practical issues. Data and compatibility problems may prevent life expectancy models from being used in various healthcare systems. Because each institution documents patients differently, a model designed on one institution's data often fails to work for many other populations. Furthermore, since building the necessary computational system to run advanced learning models is impossible for some, it can increase digital health inequality.

ISSN (Online) 2456 -1304
International Journal of Science, Engineering and Management (IJSEM)
Vol 12, Issue 05, May 2025

For these technologies to be effectively used in care, patients and providers must trust one another. It depends on the systems appearing precise and treating people fairly and transparently. For this trust to occur, those who work in data science, ethics, clinical practice, and making policies should cooperate to ensure that machine learning in life expectancy prediction is effective and fair.

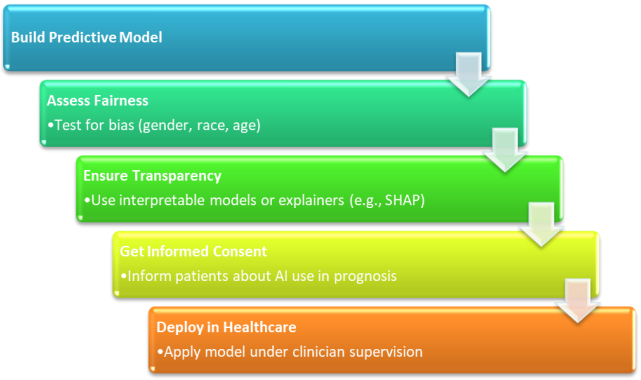


Fig. 5. Ethical and Legal Considerations in ML-based Life Prediction

VII. FUTURE DIRECTIONS

While machine learning develops, using it to predict personal life expectancy is expected to provide new chances for advancement, but it must also account for new challenges. Future studies are expected to look into important aspects that will raise both the technology and the benefits of such predictive models.

Integrating data from amniotic-AV robotics with other systems is becoming a significant trend. Especially genetic variables, physiological recordings from wearable devices, test imaging, and how the patient reports their health can complete a clinical model and show the whole picture of a person's health changes. Improving the fusion of various data using effective architectures in deep learning will lead to much higher accuracy and better personal features for users.

Exploring and establishing explainable AI that works well for life expectancy models is also important. Being more transparent will help doctors and patients participate more effectively in managing their health. Approaches such as counterfactual reasoning and causal inference may help explain ML decisions and make them trusted by professionals.

Federated and privacy-preserving learning approaches for healthcare help train effective models without revealing confidential patient information. By overcoming sharing and compatibility issues, they could enable many more societies and health systems to get better estimations for life expectancy.

It is important to check the results over time to see how they are used in practice. Analyzing these models before use allows us to highlight what might go wrong, what issues could arise, and where improvements are required. To avoid making healthcare disparities worse, researchers will rely on the views of both patients and clinicians when introducing new technologies.

Overall, the future in this field depends on advances in technology, a focus on ethics, and working with doctors and medical professionals. Work in this field can progress if multimodal tools, transparent explanations, protection of privacy, and testing in the real world are all improved.

VIII. CONCLUSION

With machine learning, it is now possible to accurately predict health outcomes for each individual. Random Forest, XGBoost, and neural networks have strengths that suit various problems. In healthcare, Random Forest makes predictions strong and straightforward, whereas XGBoost helps improve the accuracy of predictions. Unlike other methods, neural networks and deep learning can model complicated and multidimensional information.

Even so, it is not common for clinicians to use it in practice. Because of possible biases and the psychological effects of the results, it is necessary to create and manage any healthcare decision model carefully. Laws and policies should be updated to ensure the responsible use of algorithms in making decisions. Other issues, for example, poor data, difficulty with different systems' interaction, and the power of the computing systems required, can influence whether it is possible to apply these models in healthcare settings.

With time, developing systems that combine multiple data sources, using explainable AI, and ensuring privacy should improve both the use and performance of machine learning models. In addition, it is important to confirm that these tools work in practice and help make care fair to every patient.

While machine learning can predict life expectancy well, it needs experts from various fields to ensure it remains responsible and safe. Handling these multiple issues may allow this field to develop more accurate, widely trusted, and useful methods for physicians.

REFERENCES

[1] Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

[2] Chen, T., and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. https://doi.org/10.1145/2939672.2939785

[3] Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G. S., Thrun, S., and Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine, 25*(1), 24–29. https://doi.org/10.1038/s41591-018-0316-z

[4] Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT Press.

[5] Huang, S., Cai, N., Pacheco, P. P., Narrandes, S., Wang, Y., and Xu, W. (2018). Applications of support vector machine

(SVM) learning in cancer genomics. *Cancer Genomics and Proteomics*, *15*(1), 41–51. https://doi.org/10.21873/cgp.20063

[6] Katuwal, G. J., and Chen, X. (2016). Machine learning model interpretation in biomedical research: Application to protein classification. *Methods*, *111*, 55–65. https://doi.org/10.1016/j.ymeth.2016.05.014

[7] Lipton, Z. C. (2016). The mythos of model interpretability. *Communications of the ACM*, *61*(10), 36–43. https://doi.org/10.1145/3233231

[8] Miotto, R., Li, L., Kidd, B. A., and Dudley, J. T. (2016). Deep patient: An unsupervised representation to predict patients' future from the electronic health records. *Scientific Reports*, *6*, 26094. https://doi.org/10.1038/srep26094

[9] Miotto, R., Wang, F., Wang, S., Jiang, X., and Dudley, J. T. (2018). Deep learning for healthcare: Review, opportunities, and challenges. *Briefings in Bioinformatics*, *19*(6), 1236–1246. https://doi.org/10.1093/bib/bbx044

[10] Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, *366*(6464), 447–453. https://doi.org/10.1126/science.aax2342

[11] Rajkomar, A., Dean, J., and Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, *380*(14), 1347–1358. https://doi.org/10.1056/NEJMra1814259

[12] Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144. https://doi.org/10.1145/2939672.2939778

[13] Schulam, P., and Saria, S. (2015). Reliable decision support using counterfactual models. *Advances in Neural Information Processing Systems*, *28*, 1697–1705.

[14] Shickel, B., Tighe, P. J., Bihorac, A., and Rashidi, P. (2018). Deep EHR: A survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. *IEEE Journal of Biomedical and Health Informatics*, *22*(5), 1589–1604. https://doi.org/10.1109/JBHI.2017.2767063

[15] Shortreed, S. M., Laber, E., Lizotte, D. J., Stroup, T. S., Pineau, J., and Murphy, S. A. (2011). Informing sequential clinical decision-making through reinforcement learning: An empirical study. *Machine Learning*, *84*(1-2), 109–136. https://doi.org/10.1007/s10994-010-5239-3

[16] Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, *25*(1), 44–56. https://doi.org/10.1038/s41591-018-0300-7

[17] Vayena, E., Blasimme, A., and Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine*, *15*(11), e1002689. https://doi.org/10.1371/journal.pmed.1002689

[18] Wang, F., Casalino, L. P., and Khullar, D. (2019). Deep learning in medicine — promise, progress, and challenges. *JAMA Internal Medicine*, *179*(3), 293–294. https://doi.org/10.1001/jamainternmed.2018.7117

[19] Wang, X., Yu, K., Li, W., and Ma, J. (2020). A review of machine learning in life expectancy prediction. *IEEE Access*, *8*, 56962–56971. https://doi.org/10.1109/ACCESS.2020.2984033

[20] Wang, Y., Kung, L., and Byrd, T. A. (2018). Big data analytics: Understanding its capabilities and potential benefits for healthcare organizations. *Technological Forecasting and Social Change*, *126*, 3–13. https://doi.org/10.1016/j.techfore.2015.12.019

[21] Yang, J., and Leskovec, J. (2012). Modeling information diffusion in implicit networks. *Proceedings of the 2012 IEEE International Conference on Data Mining*, 599–608. https://doi.org/10.1109/ICDM.2012.142

[22] Zhang, Y., Milinovich, A., and Feng, L. (2019). Using electronic health records, a deep learning framework for predicting mortality in intensive care units. *Journal of Biomedical Informatics*, *98*, 103271. https://doi.org/10.1016/j.jbi.2019.103271

[23] Zhou, Z., Siddiquee, M. M. R., Tajbakhsh, N., and Liang, J. (2019). UNet++: A nested U-Net architecture for medical image segmentation. *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, 3–11. https://doi.org/10.1007/978-3-030-00889-5_1