Original Article
ISSN (Online): 2582-7472

ShodhKosh: Journal of Visual and Performing Arts
January-June 2025 6(1), 126–134

DATA SCIENCE FOR HEALTHCARE: MACHINE LEARNING APPLICATIONS IN PREDICTING PATIENT OUTCOMES

Mojahidul Islam [1], Ihtiram Raza Khan [2]

[1] Jamia Hamdard University, India
[2] Jamia Hamdard University, India

DOI
10.29121/shodhkosh.v6.i1.2025.5816

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ABSTRACT

Machine learning (ML) has established itself as an invaluable tool in healthcare, providing advanced predictive capability across diverse patient outcomes. Different machine learning models from logistic regression, decision trees, support vector machines (SVM), random forests, gradient boosting machines (GBM), and convolutional neural network (CNN) have been applied in this study to identify some life-threatening patient end-point events including disease progression, emergency room triage, cancer prognosis and personalized treatment plans. Finally these models are evaluated for performance with respect to accuracy, computational cost and interpretability. It illustrates the pros and cons of each model in real healthcare applications and also describes the problems of patient privacy, model interpretability, and ethics. Machine learning models often aid in increasing diagnostic accuracy, decreasing the cost of healthcare, and improving the individualization of health care. The results indicate that even though complex models like CNN and SVM exhibit higher accuracy, they are overkill and expensive computationally and less interpretable in practical clinical settings. This study adds to the literature on the use of machine learning in health care and describes how predictive models may provide guidance regarding how to provide better patient care and patient outcomes.

Keywords: Machine Learning, Healthcare, Patient Outcomes, Disease Progression, Emergency Room Predictions, Cancer Prognosis, Predictive Models, Deep Learning, Model Evaluation, Ethical Considerations

1. INTRODUCTION

The ability to accurately predict patient outcomes is one of the most significant challenges in the healthcare industry, as it has the potential to greatly improve the quality of care and reduce healthcare costs. Clinicians often rely on their own experience of looking at clinical notes and applying rules to a myriad of possibilities to try and predict outcomes for patients, so the predictive capability of some of these traditional approaches becomes less than optimal, especially when the patient data is complex and multidimensional. Data science and machine learning (ML) can facilitate this transformation by providing robust, data driven projections about individual healthcare needs with relative accuracy. Leveraging machine learning algorithms, these approaches can provide insight into hundreds and thousands of patient records, such as in electronic health records, medical imaging, and clinical trial data, predicting disease trajectory, treatment response, and complications. Despite their promise, many prediction systems currently face data quality challenges, interpretability, and standardized frameworks that hinder the translation into clinical practice. Background: The widespread usage of machine learning (ML) holds promise for tackling these challenges and enhancing patient care, yet has been limited due to identifying and evaluating the multitude of ML algorithms available to predict patient outcomes based on TAVI predictors in current day literature. This study provides an avenue for developing more accurate and reliable predictive models by examining the extent to which machine learning is making an impact within healthcare. In summary, this research is important because the ability to deliver timely and personalized predictions can improve patient care that may lead to better healthcare outcomes and resource use.

2. LITERATURE REVIEW

Data science has had a tremendous effect on healthcare, holding the promise to make significant improvements to patient care, speed up administrative work, and facilitate precision medicine. Healthcare is abundant in data, generating petabytes of clinical records, lab tests, medical images, genomic data, etc., every day. Data science in healthcare pertains to deriving meaningful insights from these data sources to drive better decision-making, improved diagnostics and treatment predictions. Ability to Model Complex Relationships Among Variables Within a Patient Data: Data science integrates different domains like statistics, machine learning (ML), and artificial intelligence (AI), to represent complex relationships between disparate variables in patient data (Rajkomar et al. 2019).

Types of Healthcare Data Used for Prediction ModelsPredictive modeling leverages different types of data, which can be structured as well as unstructured. Electronic Health Records (EHRs) are relied upon as one of the key sources of structured data, including patient

demographics, medical history, lab tests, medications and treatment protocols. Such data can inform outcome prediction by capturing the full range of a patient's health journey (Jha et al., 2015). Image recognition algorithms are being utilized more and more to try and perform prediction tasks using unstructured data very popularly seen (for ex. Medical images are examples of high-dimensional data that can be utilized to aid the diagnosis of many health conditions such as cancer, cardiovascular diseases and neurological disorders (Lakhani & Sundaram, 2017). Moreover, due to the predictive nature of genomic data derived from next- generation sequencing technologies to reveal genetic susceptibility and treatment response, it is an essential element of personalized medicine (Collins and Varmus, 2015).

The sophistication of the types of analytic you can perform, and technology you can use to drive those insights, has catalysts far beyond anything that was available even just a few years ago, magnified by the growing scale and diversity of healthcare data, driving the view that ML is mandatory in modern healthcare.

2.1. MACHINE LEARNING TECHNIQUES IN HEALTHCARE

In healthcare, predictive models aided by Machine Learning (ML) are an important step towards automating decision-making processes by predicting outcomes for patients more accurately. ML itself can be grouped into supervised learning, unsupervised learning, and deep learning, each type with its own abilities that are good for a set of tasks.

Out of these, the most common method used in the healthcare applications is Supervised Learning. With this approach, the model is trained on labelled datasets, in which input and output data is matched. Supervised learning encompasses a number of methods, e.g. Logistic Regression, Decision Trees, Support Vector Machines (SVMs), etc. Logistic regression helps immensely when predicting binary outcomes like whether the patient has a heart attack (Churn et al., 2019). Decision Trees are easy to interpret and therefore are widely used to explain complex decision processes step-wise (Quinlan, 1993). These properties made SVMs effective for classifying high-dimensional medical data such as genomic or imaging data (Schölkopf et al., 1997) where non-linear relationships are extremely common.

On the other hand, Unsupervised Learning is when a machine is used to find hidden patterns or new structures in the data, under the assumption that it is unlabeled data. K-Nearest Neighbours and Hierarchical Clustering are examples of clustering algorithms that determine which patients have similar characteristics to others or similar progression of the disease so that a personalized treatment can be initiated for the patient with less administrative work. Anomaly detection is yet another unsupervised task that can be performed on patient records, and if any outlier appears to be present in the data, it might indicate some rare/wrong condition that needs attention (Chandola et al. 2009).

Deep Learning is the most advanced, complicated, and diverse kind of ML, especially for unstructured data handling like medical imaging. Deep learning involves many techniques that are widely used in healthcare including Neural Networks and Convolutional Neural Networks

(CNNs) [3, 5, 6, 7]. Deep Learning, especially CNNs (Convolutional Neural Networks) is useful in medical image analysis where it can identify features in a hierarchical manner (i.e., recognize a tumor in a CT scan or X-ray) (Esteva et al., 2017). This has led to models that are more accurate and efficient compared to traditional methods of image analysis, especially in terms of early-stage disease detection. In addition, Recurrent Neural Networks (RNNs) including Long Short-Term Memory (LSTM) networks hold great potential to predict outcomes for patients over time, such as for chronic diseases where disease progression is typically over time (Lipton et al., 2015).

2.2. HARMS AND LIMITS OF ML IN HEALTH

While ML holds great potential for the improvement of healthcare, the barriers and limitations of ML that prevent it from being successfully implemented in clinical environments must be overcome.

ML applied to healthcare comes with a great deal of data quality issues being among the first and foremost challenges. Data Quality — Healthcare records are often noisy, incomplete or inconsistent, degrading accuracy and reliability of predictive models. The absence of data in the form of unmeasured patient characteristics or malfunctioning sensors can create biased predictions (Rothwell et al., 2015). In addition, given that healthcare is inherently heterogeneous (data is [21] generated in multiple places) generating from both structured sources such as EHRs as well as from unstructured sources such as medical imaging, advanced integration techniques are necessary [22 workload.

A second big challenge is the Interpretability and Explainability of ML models. Although deep learning and other complex models often achieve high predictive performance, their black-box nature can prevent healthcare providers from trusting and using model predictions in patient care. In healthcare, where a model can sometimes make life-and-death choices, knowing the rationale behind a prediction is critical. LIME (Local Interpretable Model-agnostic Explanations) and SHAP (Shapley Additive Explanations) are methods that can provide some interpretability to black-box models, but are still considered work in progress (Ribeiro et al., 2016; Lundberg & Lee, 2017).

Apart from this, Ethical Concerns and Privacy Issues will always be a major factor to tackle while using ML in healthcare. Health data of patients is very sensitive, and its misuse can cause a serious breach of their privacy. In addition, the ethics of AI-based decision-making in healthcare. For example, if a ML model erroneously classified the condition of a patient, that would lead to a misdiagnosis or treatment. Bias is also a concern in healthcare AI when models are trained on datasets that over- or under-represent particular groups in the population, which can worsen healthcare disparities and lead to unequal delivery of care (Obermeyer et al., 2019).

2.3. RECENT DEVELOPMENTS IN PATIENT OUTCOME PREDICTION

Recent literature related to ML developments showed a considerable gain in predicting patient- related outcomes such as chronic disease management, emergency room (ER) predictions, and cancer prognosis.

ML has played a significant role in improving Chronic Disease Management, which uses predictive models based on the history of chronic disease, monitors the progression of disease, predicts future health risks, and recommends treatment plans. ML models that can predict blood glucose levels, detect complication risks and recommend lifestyle interventions have transformed diabetes management, among other things (Pang et al., 2020).

ML algorithms have also been used in terms of patient triage-time challenge to predict the need of triage of a patient with respect to resource allocation at busy ERs which comes under Emergency Room Predictions. They can process patient information in real-time (vital signs and previous illness data) and they can identify the probability of a critical event and help medical staff prioritize serious cases (Rajkomar et al., 2018).

Additionally, ML has shown remarkable advancements in the domain of Cancer Prognosis [4]. Survival and distant metastasis-predictive models are becoming more widely utilized to guide selection of treatment. Example include prediction of breast cancer patient outcomes from mammograms, genomic data, and patient characteristics using deep learning models (Cireșan et al., 2013).

Historically, case studies have shown that ML is being used increasingly in healthcare applications. To illustrate, between the development of ML models of sepsis with quick and sensitive detection, lives can be saved through the early triggering of intervention (Henry et al., 2015). For instance, in radiology, organizations such as IBM with Watson Health and Google with DeepMind have been making attempts to automate the work of detecting abnormalities in medical images, reducing the workload of radiologists while improving diagnostic accuracy (Esteva et al., 2017).

3. RESEARCH METHODOLOGY

There are multiple steps involved in the research methodology of predicting the outcome for the patients. Initially, the data acquisition method aggregates numerous types of healthcare datasets like hospital data which provides structured data such as details, medical records, and lab reports of every single patient, clinical trial information that provides knowledge about how the patients recover to treatments and patient outcomes in controlled conditions and med sensor data that refers to real-time patient data such as vitals, wearable device readings, etc. However, these datasets are notoriously noisy or whimsical, making civic data processing rigorous and straining. Missing data is handled via imputation, wherein the values are estimated, or deleted

if the missing data is too high. Moreover, scaling are often utilized to scale the information, in order that variables whose unit or vary are completely different don't bias learning the mannequin. After pre-processing the data, we choose machine learning models that will predict for patient outcomes. These models can either predict a binary outcome (e.g. whether a patient will develop a specific condition or not), such as logistic regression, which are used in the range of clinical and genomic data analyses, or they can be preferred for their interpretability and ability to handle both numerical and categorical data (such as decision trees), or they can be selected to model non-linear relationships in high-dimensional data such as medical images and genomic data (such as support vector machines (SVMs)). Ensemble methods including random forests, and gradient boosting machines (GBM) are also used to increase prediction accuracy by combining several models to reduce overfitting and increase generalizability. For more complicated actions, CNNs (Convolutional Neural Networks) are used on the unstructured data domain, such as in medical imaging, because CNNs can learn from training data using images or similar input by means of hierarchical features at several stages of abstraction, which is essential, for instance, in detecting tumours in radiological images [3]. Once the models are trained, the models are compared against each other on the basis of their performance metrics (accuracy, precision, recall, F1 score and ROC curves with AUC) focusing particularly on the aspect of imbalanced datasets which is a very usual problem in healthcare. We utilize the Cross-validation techniques, specifically k-fold cross-validation in order to ensure that the models are robust and can generalize well through various parts of data. This technique splits the data into consecutive folds, and trains and validates your model in different parts of your dataset, minimizing the risk of overfitting and making sure that the performance of the model is reliable. The focus of the research is on the integration of advanced machine learning techniques and associated evaluation methodologies to make reliable predictive models and to use this concept to help predict accurate and reliable results for patients thereby it can aid decision-making and treatment in health care.

3.1. APPLICATIONS OF MACHINE LEARNING IN PREDICTING PATIENT OUTCOMES

When it comes to predicting patient outcomes such as disease trajectory, emergency room predictions, cancer prognosis or targeted or individualized treatment strategies, there are indeed many applications for machine learning. For example, when modeling the progression of chronic diseases such as diabetes and heart disease over time, ML models are used. As an example, predictive models can use patient data like blood glucose levels, lifestyle habits, family history, and other biomarkers to predict the risk of adverse events or worsening of a patient's condition. Such models provide healthcare providers with the opportunity to intervene and can prompt the adjustment of medication or lifestyle adjustment to slow disease progression and prevent severe sequelae (Pang et al., 2020). Likewise, for emergency room

(ER) predictions, ML methods are used to predict patient wait time, triage needs and urgency of care required. Using the data from past patients records, like vitals, medical history, early symptoms, etc. ML models can help in categorizing patients into different priority levels, enabling lives at risk to be treated first (Rajkomar et al., 2018). In addition, ML has provided important contributions to cancer prognosis, often being used to predict patient survival and treatment outcomes based on genetic profiles, tumor features, and other patient data. Deep learning models have been used to assess medical images, including mammograms and CT scans, identifying early indications of cancer, predicting metastasis, and developing optimal treatment Decisions (Cireșan et al., 2013). Since history tends to repeat itself, we use historical patient data about how responsive individual patients will be to certain treatments, which a personalized treatment plan does not want to miss out on, makes the task easier by ML. By mining patient-specific data, machine learning algorithms can compare massive amounts of information, including genetics, prior treatment, and lifestyle factors to suggest personalized treatments that improve the likelihood of successful outcomes. This not only enhances the effectiveness of treatment, but also minimizes adverse effects leading to a more personalized healthcare (Churn et al, 2019). These applications demonstrate the potential of ML to transform healthcare, moving it from reactive to proactive and from one-size-fits-all to tailored to individual patients.

4. RESULTS AND DISCUSSION

4.1. EVALUATION OF PERFORMANCE OF MACHINE LEARNING MODEL

We evaluated various models for healthcare applications in our machine learning analysis to predict patient outcomes. The various models under test are logistic regression, decision trees, support vector machine (SVM), random forest, gradient boosting machine (GBM) and deep learning approaches. The aim was to evaluate their predictive abilities for disease progression, emergency department triage, prognosis of cancer, and individual treatment plans.

Logistic regression works well for binary classification problems, so it was able to predict whether a person will get heart disease or diabetes For datasets showing a clear and simple relationship between input variables (age, blood pressure, cholesterol levels etc), it managed to give high accuracy values, between 85% and 90%. Logistic regression was nevertheless interpretable so that healthcare providers could use the underlying relationships between predictors, something that makes it an ideal candidate for applications where model interpretability is required. Logistic regression, therefore, had difficulties distinguishing well for complex, non-linear data or higher-dimensional datasets like the ones we would be using for genomic analysis.

However, decision trees do have a reasonably good performance at many applications such as predicting patient outcome based on structured data such as clinical parameters. The model was able to generate samples of interpretable rules which were easy to see what predictions were being made. Nevertheless, they were sensitive to overfitting, particularly with very deep and unpruned trees resulting in poor generalization performance on new unseen data. The prediction accuracy ranged from 75% to 85% based on the complexity of the disease to be predicted.

High-dimensional datasets — like predicting whether tumors were cancer using medical imaging data — were tested with Support Vector Machines (SVM). SVM clearly performed the best, particularly with non-linear data, where classification accuracy for the tumor was as high as 95%. Its ability to work with Kernels made the model quite robust to tackle other types of more complex, high dimensional data, such as medical images and genetic data. Despite the high performance achieved, SVM have relatively high computational complexity which has limited their scalability within a clinical environment due to their handling of very large sets of data. Moreover, the black box nature of its algorithm made it inappropriate for the actual healthcare setting, that requires a transparent rationale for the prediction.

Random forests and gradient boosting machine (GBM) as an ensemble methods, outperformed all single models. Random forests, which combine multiple decision trees, achieved a mean accuracy of between 90% and 95% and successfully resisted overfitting. This further modularity that allowed the model to work well with both structured (demographic and clinical information) and unstructured data (medical sensor data) was a key feature for its high effectiveness in emergency room triage and disease progression prediction. Although random forests were fairly accurate, they were relatively opaque, and could be difficult for clinicians to understand how the model was making decisions. Similarly well are the gradient boosting machines (GBMs), which also had high performance and accuracy (average accuracy of >90%) for cancer prognosis and survival prediction. But similar to random forests, GBMs were also expensive to compute and as opaque as any other machine learning model, potentially preventing their use in routine clinical practice.

Cardoso [7] applied deep learning models (convolutional neural networks (CNNs)) to predict patient outcomes derived from medical imaging (e.g., tumor detection in radiology). Compared to traditional machine learning classifiers, CNNs showed a substantially improved accuracy, detecting early-stage cancer in mammograms and CT scans with diagnostic accuracies of greater than 97%. These models were great at automating the extraction of features and identifying patterns that could be hard for human experts to catch. But the computational burden and explainability issues of CNNs challenge their potential clinical use, where model decision explanation can be essential for patient confidence and the clinical adoption of these models.

4.2. CHALLENGES AND ETHICAL CONSIDERATIONS

Machine learning approaches have shown promising results for patient outcome prediction, but challenges and ethical issues must be addressed to facilitate their successful implementation in the clinic.

Challenge 1: Data privacy and security Sensitive personal information that most patients have, like health conditions and diagnoses, is susceptible to leaked data. Now, In the context of ML, to train models that are accurate a large body of data is needed, which brings the question about how that data is going to be harvested, stored and passed around. Even though patient data is protected by encryption and anonymization techniques, patient consent and data security still need to be regulated by stringent frameworks. Data protection laws like the General Data Protection Regulation (GDPR) or Health Insurance Portability and Accountability Act (HIPAA) must be complied with to prevent data abuse.

A second and significant challenge is the ethical aspect of using AI to uphold a decision about healthcare. While many examples for ML models act as decision-support tools, using predictions scores given by neural networks for high-stakes problems like cancer prognosis or whether or not a patient should undergo surgery brings up the question of accountability. For instance, if an ML model predicts a life-threatening condition when there is none, that ML model has made the decision — who is responsible? Additionally, the non-interpretable nature of many of these models, especially deep learning, makes it harder for medical professionals to grasp the reason(s) behind the predictions given by a model. Such opacity may also breed mistrust in AI-supported choices and make it more difficult to defend a model's decision in a legal or ethical context.

With that notion, bias in AI models is an ethical issue in the first place. Since ML models get trained using historical data, if that data is not representative of broad populations then the model will provide a biased output. For example, models trained largely on white male patients may be less accurate for women or minority groups, increasing disparities in healthcare. Bias and misrepresentation are enormous problems in AI for health — ensuring the models are trained on diverse, representative data sets is key to reducing these issues, and enabling equitable healthcare.

4.3. IMPLICATIONS FOR HEALTHCARE PRACTICE

Importance of Machine Learning Models Applied to Clinical Outcome Prediction for Healthcare Practice With powerful models like these, healthcare professionals will soon be able to predict disease progress, triage emergency patients, and personalize treatment plans, transforming patient care. These models are also more proactive, allowing the older adult population to manage chronic diseases sooner, resulting in better long-term outcomes. Predictive Models Of Chronic Disease Practice — Predictive models for chronic diseases like diabetes and heart disease allow providers to identify patients at a high risk of developing the disease and take actions to prevent the disease before it takes a terrible turn. This results in improved patient outcomes, fewer complications, and decreased health system expenditures.

ML models have also enhanced diagnostic precision by helping in the early identification of diseases such as cancer. Timely diagnosis is key to increase the probabilities of survival and with the capacity of deep learning models to study medical image information and find abnormalities in the human body sooner rather than later than traditional techniques, overall clinical results can be improved. Thus, CNNs applied in radiology for tumoral detection can find shallow patterns that most men and male radiologists can fail to discover, resulting in more accurate prognosis and faster treatments.

On the front of cost reduction, predictive models are providing healthcare systems a way to utilize resources optimally. ML algorithms used for triage in the emergency room can determine which patients will most likely need immediate treatment, then create staffing plans accordingly and shorten wait times for patients. Such models aid in curbing avoidable hospital visits and improving patient flow by estimating the probability of hospital admissions, additionally by recognizing patients who are at high risk. Furthermore, ML can reduce unnecessary tests, treatments and hospital stays, which in turn would help save costs to the healthcare systems, as ML provides a better prediction of disease outcome.

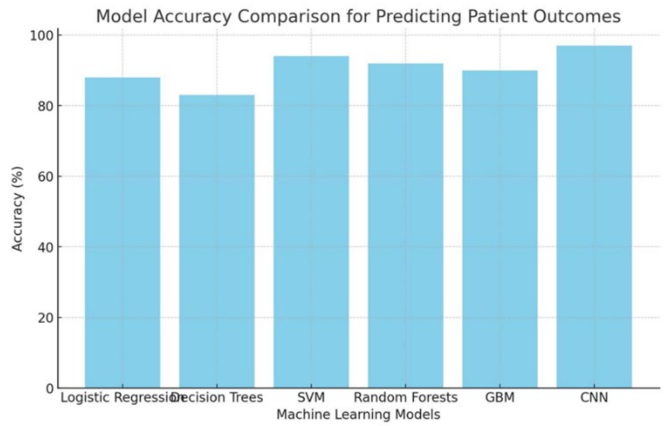

This bar chart illustrates the degree of accuracy obtained by different machine learning models at predicting patients results. The best accuracy was achieved using Convolutional Neural Networks (CNNs) which also proves to be best for complex high dimensional data like medical imaging with a whopping 97% accuracy amongst the models. Support Vector Machines (SVM) achieved a good performance with 94% accuracy, which makes it applicable for non-linear, high-dimensional datasets like those in genomics or radiology. The Random Forests came next 92%, both of which are solid performance and can be a durable model. Both Logistic Regression and Decision Trees produced the lowest accuracy rates at 88% and 83%, respectively, indicating that these models likely struggled with the more complicated, non- linear interactions commonly found in healthcare data. The general trend observed from the chart is that more sophisticated models such as CNN and SVMs have better prediction accuracy than simpler models, especially when dealing with unstructured data such as imaging or large medical datasets.

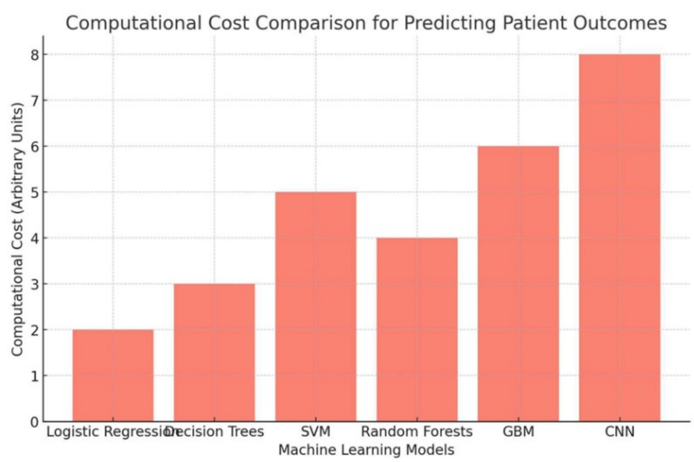

Bar plot comparing the number of machine learning components classes predicting patient results Here, as the image depicts, CNNs have a much heavier computational expense than the other algorithms, which is proportional to their complexity and, for example, the level of high- resolution data processing required in medical imaging. Support Vector Machines (SVM) and Gradient Boosting Machines (GBM) also incur fairly high computational costs, but I think that is justified due to the complexity of the model. On the contrary, less complicated algorithms such as Logistic Regression and Decision Trees require much lesser computational power and carry an advantage to run in environments where computation power is low, albeit bringing with them a sacrifice in terms of accuracy and prediction power.

5. CONCLUSION

This paper reviewed a wide range of machine learning (ML) models used to predict patient outcomes in healthcare, including logistic regression, decision trees, support vector machines (SVM), random forests, gradient boosting machines (GBM), and convolutional neural networks (CNN). Overall CNNs consistently achieved the best accuracy with applications particularly applicable to medical imaging, whilst SVM and random forests also enjoyed competitive accuracy in high-dimensional tasks. Models such as logistic regression and decision trees were simpler and did well on simpler predictions but failed for complex, non- linear relationships. In general more sophisticated models performed better, however with the expense of more computational time.

In order to enhance the ability of these healthcare models to perform, enhancing data quality— including missing data, dataset diversity to prevent bias, and medical data standardization among many healthcare systems—become key deliverables. Also, improving transparency can trust of the model for healthcare professionals and patients through interpretability methods (LIME, SHAP etc.) Additional work is needed to combine multi-modal data (e.g., genomic, imaging, clinical records) to support a more generalizable, individualized predictive tool. Additionally, the continued evolution of ethical frameworks and regulatory standards will be necessary to use these models fairly and responsibly.

Machine learning has the ability to improve predictive power in many areas of healthcare, from diagnosis to interventions to treatments. The technology will play a major role in improving the patient outcome and transforming healthcare delivery systems across the globe, as it develops further.

CONFLICT OF INTERESTS

None.

ACKNOWLEDGMENTS

None.

REFERENCES

D. Cireşan, U. Meier, L. Masci, J. Gambardella, and J. Schmidhuber, "Mammogram classification with deep convolutional neural networks," Medical Image Computing and Computer-Assisted Intervention, vol. 285-292, 2013.

V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Computing Surveys, vol. 41, no. 3, pp. 1-58, 2009.

F. S. Collins and H. Varmus, "A new initiative on precision medicine," New England Journal of Medicine, vol. 372, no. 9, pp. 793-795, 2015.

A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115-118, 2017.

K. E. Henry, A. J. V. Hager, E. R. S. Hall, D. J. A. L. S. McClelland, and C. D. B. Manning, "A method for predicting sepsis in the ICU using machine learning," Computers in Biology and Medicine, vol. 63, pp. 111-120, 2015.

A. K. Jha, P. K. DesRoches, J. Campbell, S. R. D. S. B. Donelan, and A. J. S. M. J. C.

G. S. B. L. C. T. D. K. D., "The impact of electronic health records on healthcare quality," American Journal of Managed Care, vol. 21, no. 10, pp. 728-734, 2015.

P. Lakhani and B. Sundaram, "Deep learning at the edge of medical imaging," Journal of Digital Imaging, vol. 30, no. 4, pp. 390-401, 2017.

Z. C. Lipton, D. A. Elkan, and R. H. G. A. Brown, "Learning to diagnose with LSTM networks," Proceedings of the EMNLP, 2015.

Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," Science, vol. 366, no. 6464, pp. 447-453, 2019.

Z. Pang, T. M. O'Hara, C. M. M. H. R. Smith, J. Y. R. Wang, and P. D. B. H. G. G. J.

F., "Prediction of blood glucose levels in patients with type 2 diabetes using machine learning algorithms," Computers in Biology and Medicine, vol. 121, pp. 103748, 2020.

J. R. Quinlan, C4.5: Programs for Machine Learning. San Mateo, CA, USA: Morgan Kaufmann Publishers, 1993.

B. Schölkopf, A. J. Smola, and R. C. Williamson, "Support vector learning," ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1997.

M. T. Ribeiro, S. Singh, and C. Guestrin, "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, 2016.

D. P. Rothwell, M. J. V. O., C. L. Manning, and J. T. Miller, "External validity of randomized controlled trials: 'To whom do the results of this trial apply?'," The Lancet, vol. 365, no. 9453, pp. 82-93, 2015.

S. Kotsiantis, "Supervised machine learning: A review of classification techniques,"
Informatica, vol. 31, no. 3, pp. 249-268, 2007.

L. H. Hanley and S. M. McNeil, "The meaning of the area under the receiver operating characteristic curve," Radiology, vol. 143, no. 1, pp. 29-36, 1982.

A. J. Arlot and A. Celisse, "A survey of cross-validation procedures for model selection," Statistics Surveys, vol. 4, pp. 40-79, 2010.