International Journal of Advanced Research in Education and TechnologY (IJARETY)

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi crossref

निस्केयर NISCAIR

International Journal of Advanced Research in Education and TechnologY(IJARETY)

| ISSN: 2394-2975 | www.ijarety.in| | A Bi-Monthly, Double-Blind Peer Reviewed & Referred Journal |

|| Volume 7, Issue 6, November-December 2020 ||

| DOI:10.15680/IJARETY.2020.0706018|

Leveraging Deep Learning to Extract Actionable Insights from High-Dimensional Biomedical Data: Opportunities and Challenges

Punitha P[1], Sriraksha S[2], Pushpa S[3], Mala P[4], Manjuala S[5]

Assistant Professors, Department of Computer Science and Engineering, City Engineering College, Bengaluru, Karnataka, India[1,2,3,4,5]

ABSTRACT: Addressing the challenge of extracting actionable insights from complex, high-dimensional biomedical data is crucial for advancing health care. Modern biomedical research generates a variety of data types, including electronic health records, imaging, -omics, sensor data, and textual information, which are often complex, diverse, poorly annotated, and unstructured. Traditional methods of data mining and statistical learning usually involve feature engineering to create effective and robust features, followed by the development of prediction or clustering models. These processes face significant difficulties due to the complexity of the data and the lack of comprehensive domain knowledge. Recent advancements in deep learning offer promising new approaches for creating end-to-end learning models from such complex data. This article reviews recent literature on the application of deep learning in health care and suggests that these approaches could help translate extensive biomedical data into better health outcomes. However, there are limitations and a need for more refined methods, particularly in terms of making these models more understandable to domain experts and citizen scientists. We discuss these challenges and recommend the development of interpretable architectures that can bridge the gap between deep learning models and human understanding.

KEYWORDS: Deep learning, health care, biomedical informatics, translational bioinformatics, genomics, electronic health records

I. INTRODUCTION

Health care is entering a new era where the growing volume of biomedical data is becoming increasingly significant. For instance, precision medicine aims to deliver tailored treatments to individual patients by integrating various data sources such as molecular traits, environmental factors, electronic health records (EHRs), and lifestyle information.

The vast amount of biomedical data presents both opportunities and challenges for health care research. Understanding the relationships between different data components is essential for developing reliable, data-driven medical tools. Previous efforts have focused on linking multiple data sources to create comprehensive knowledge bases for predictive analysis and discovery. Despite promising results from existing models, the adoption of machine learning tools in medicine remains limited. The high-dimensionality, heterogeneity, temporal variability, sparsity, and irregularity of biomedical data present significant challenges. These issues are further complicated by inconsistencies and conflicts among various medical ontologies, such as SNOMED-CT, UMLS, and ICD-9. Additionally, clinical phenotypes can be represented in multiple ways, complicating efforts to standardize and understand these concepts. Traditionally, domain experts manually define phenotypes, which is labour-intensive and may overlook novel patterns. Representation learning methods offer an alternative by automatically discovering the necessary features for prediction from raw data. Deep learning, a subset of representation learning, employs multiple layers of non-linear transformations to derive increasingly abstract representations from raw input. Deep learning has shown impressive results in fields like computer vision, speech recognition, and natural language processing.

Given its success in various domains and ongoing methodological advancements, deep learning holds significant promise for biomedical informatics. Initiatives are already underway to apply deep learning in health care, such as Google DeepMind's plans and Enlitic's use of deep learning for analyzing X-rays and CT scans.

However, deep learning has not yet been extensively tested across the wide range of medical problems that could benefit from its capabilities. The field must address several challenges related to the unique characteristics of health care data,

| ISSN: 2394-2975 | www.ijarety.in| | A Bi-Monthly, Double-Blind Peer Reviewed & Referred Journal |

|| Volume 7, Issue 6, November-December 2020 ||

| DOI:10.15680/IJARETY.2020.0706018|

including its sparse, noisy, heterogeneous, and time-dependent nature. To advance this field, improved methods and tools are needed to integrate deep learning with health care workflows and clinical decision support.

This article reviews recent and upcoming applications of deep learning in medicine, focusing on how these methods could impact health care. We do not aim to provide a detailed technical background or a general overview of deep learning but will concentrate on biomedical data from clinical imaging, EHRs, genomes, and wearable devices. While other data sources like metabolomics and antibodyomics could be valuable, deep learning has not yet made significant strides in these areas. We will introduce the general framework of deep learning, review its medical applications, and discuss the opportunities and challenges related to its use in precision medicine and advanced health care.

II. DEEP LEARNING FRAMEWORK

Machine learning is a broad artificial intelligence technique that extracts patterns from data without needing predefined rules. Its main advantage is creating predictive models without requiring detailed assumptions about underlying processes, which are often not well understood. The typical workflow in machine learning involves four stages: data harmonization, representation learning, model fitting, and evaluation. Historically, building a machine learning system required significant engineering and domain knowledge to convert raw data into a format suitable for learning algorithms, such as classifiers, to identify patterns. Traditional methods often use a single, linear transformation of data, which limits their effectiveness on raw, natural data.

Deep learning differentiates itself from traditional machine learning by its approach to learning from raw data. It employs models with multiple neural network layers to develop data representations with varying levels of abstraction. Unlike traditional artificial neural networks (ANNs), which typically have only a few layers and are optimized for specific tasks, deep learning models use numerous hidden layers to create more complex and generalizable representations. Traditional ANNs usually consist of up to three layers and produce task-specific representations that may not generalize well. In contrast, each layer in a deep learning system generates representations based on the input from the previous layer, optimizing unsupervised criteria. Deep learning's key feature is that these layered representations are learned from the data itself rather than being manually designed.

Figure 1 highlights these distinctions: deep neural networks process inputs through multiple nonlinear layers to initially learn broad patterns, which are then refined in subsequent supervised layers to optimize for specific tasks using backpropagation.

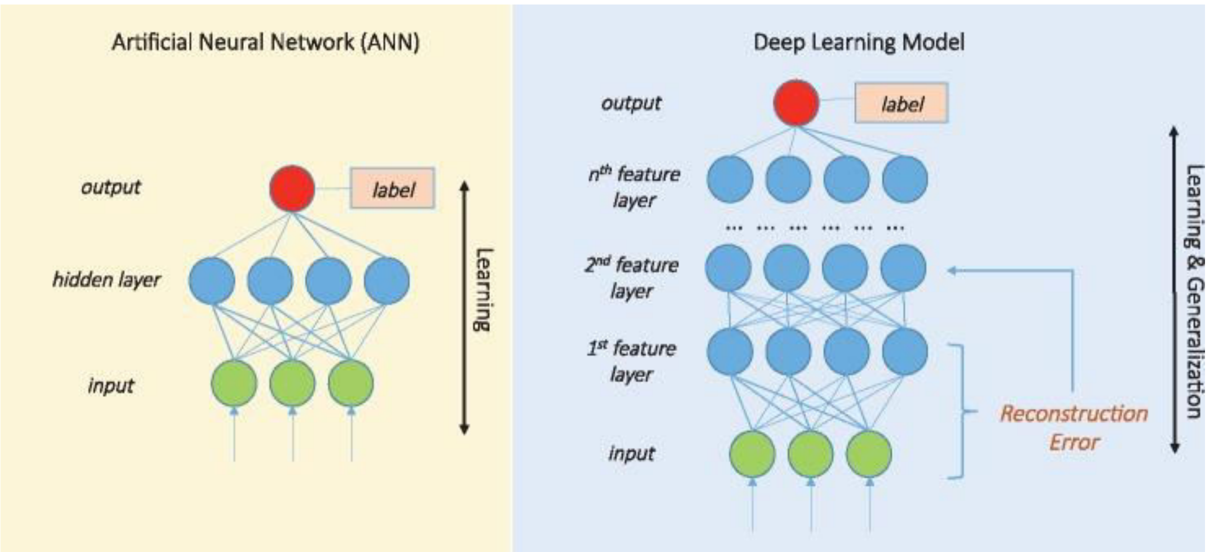

Figure 1. Comparison between ANNs and deep architectures.

While ANNs are usually composed by three layers and one transformation toward the final outputs, deep learning architectures are constituted by several layers of neural networks. Layer-wise unsupervised pre-training allows deep

| ISSN: 2394-2975 | www.ijarety.in| | A Bi-Monthly, Double-Blind Peer Reviewed & Referred Journal |

|| Volume 7, Issue 6, November-December 2020 ||

| DOI:10.15680/IJARETY.2020.0706018|

networks to be tuned efficiently and to extract deep structure from inputs to serve as higher-level features that are used to obtain better predictions. Breakthroughs in unsupervised pre-training, advancements in techniques to prevent overfitting, the use of general-purpose graphic processing units to accelerate computations, and the development of high-level frameworks for building neural networks (such as Theano, Caffe, and TensorFlow) have established deep learning as a leading technology in various fields. Deep learning excels at uncovering complex patterns in high-dimensional data and has achieved outstanding results in object detection in images, speech recognition, and natural language understanding and translation. In healthcare, deep learning has also shown promising results, including the detection of diabetic retinopathy from retinal fundus images, skin cancer classification, and predicting DNA- and RNA-binding protein sequence specificities. These successes mark the beginning of a new era in medical care, driven by intelligent tools and deep learning technologies.

III. CLINICAL IMAGING

Following their success in computer vision, deep learning techniques were first applied to clinical data, particularly for image processing. Initial efforts focused on analyzing brain MRI scans to predict Alzheimer's disease and its variations. Convolutional neural networks (CNNs) were also employed in other medical areas, such as analyzing low-field knee MRI scans to automatically segment cartilage and assess the risk of osteoarthritis. This method, despite using 2D images, outperformed a state-of-the-art technique that relied on manually selected 3D multi-scale features.

Deep learning has been used for various other medical imaging tasks, including segmenting multiple sclerosis lesions in multi-channel 3D MRI and differentiating between benign and malignant breast nodules in ultrasound images. More recently, Gulshan et al. applied CNNs to detect diabetic retinopathy in retinal fundus photographs, achieving high sensitivity and specificity across around 10,000 test images, compared to certified ophthalmologist annotations. Additionally, CNNs demonstrated performance comparable to 21 board-certified dermatologists in classifying skin cancer types from biopsy-proven clinical images, with a dataset of 130,000 images, including 1,942 biopsy-labelled test images.

IV. ELECTRONIC HEALTH RECORDS

Recently, deep learning has been applied to process aggregated electronic health records (EHRs), incorporating both structured data (such as diagnoses, medications, and laboratory tests) and unstructured data (such as free-text clinical notes). Most of this research focuses on using deep learning for specific, typically supervised, predictive clinical tasks, often demonstrating that deep learning outperforms conventional machine learning models on various metrics like Area under the Receiver Operating Characteristic Curve, accuracy, and F-score. While many studies use end-to-end supervised networks, some also explore unsupervised models to derive latent patient representations, which are then analyzed using simpler classifiers like random forests or logistic regression.

Several studies have applied deep learning to predict diseases based on patient clinical status. Liu et al. used a four-layer CNN to predict conditions such as congestive heart failure and chronic obstructive pulmonary disease, showing notable improvements over baseline methods. DeepCare, an end-to-end deep dynamic network, employed RNNs with long short-term memory (LSTM) units, pooling, and word embedding to infer illness states and predict future outcomes. This model also adjusted for irregularly timed events typical in longitudinal EHRs and incorporated medical interventions for dynamic prediction. DeepCare was evaluated for disease progression modelling, intervention recommendations, and risk prediction in diabetes and mental health patients. Choi et al. developed Doctor AI, an end-to-end RNN model with gated recurrent units (GRUs) for predicting diagnoses and medications, demonstrating significantly higher recall compared to shallow baselines and good adaptability across institutions. Miotto et al. proposed a three-layer Stacked Denoising Autoencoder (SDA) to learn deep patient representations from EHRs for disease risk prediction, showing improved results over raw EHRs and traditional methods. Liang et al. used Restricted Boltzmann Machines (RBMs) to uncover novel concepts in EHRs, achieving better prediction accuracy for various diseases.

Deep learning has also been utilized to model continuous time signals, such as laboratory results, to identify specific phenotypes. Lipton et al. employed RNNs with LSTM units to classify diagnoses based on multivariate time series of clinical measurements from pediatric intensive care units, showing significant improvements over traditional models. Che et al. used SDAs, regularized with prior knowledge based on ICD-9 codes, to detect physiological patterns in clinical time series. Lasko et al. applied a two-layer stacked Autoencoder (without regularization) to model serum uric acid measurements for distinguishing between gout and acute leukemia. Razavian et al. evaluated CNNs and RNNs with

| ISSN: 2394-2975 | www.ijarety.in| | A Bi-Monthly, Double-Blind Peer Reviewed & Referred Journal |

|| Volume 7, Issue 6, November-December 2020 ||

| DOI:10.15680/IJARETY.2020.0706018|

LSTM units for predicting disease onset from laboratory tests alone, outperforming logistic regression with hand-engineered features.

Neural language models have also been used on EHRs to learn embedded representations of medical concepts like diseases, medications, and laboratory tests for analysis and prediction. For example, Tran et al. used RBMs to abstract ICD-10 codes for predicting suicide risk in a mental health cohort. Additionally, deep architectures based on RNNs have shown promise in removing protected health information from clinical notes, facilitating the automatic de-identification of free-text patient summaries.

The prediction of unplanned patient readmissions has also gained attention. Nguyen et al. proposed Deepr, an end-to-end CNN-based architecture that detects and combines clinical motifs in longitudinal EHRs to stratify medical risks. Deepr demonstrated strong performance in predicting readmissions within six months and identified meaningful and interpretable clinical patterns.

V. GENOMICS

Deep learning is increasingly applied in high-throughput biology to manage and interpret vast, high-dimensional datasets, such as those from DNA sequencing and RNA measurements. These deep models help uncover high-level features, enhance performance over traditional approaches, and provide deeper insights into the structure of biological data. For a comprehensive overview of these developments, refer to more detailed reviews [93–96].

Early applications of neural networks in genomics involved replacing conventional machine learning models with deep architectures, while maintaining the same input features. For instance, Xiong et al. used a fully connected feed-forward neural network to predict the splicing activity of individual exons, employing over 1,000 predefined features from the exons and adjacent introns. This approach improved prediction accuracy compared to simpler methods and identified rare mutations linked to splicing misregulation.

More recent advancements involve using convolutional neural networks (CNNs) directly on raw DNA sequences, eliminating the need for predefined features. CNNs, by applying convolutional operations on small input regions and sharing parameters across these regions, handle larger sequence windows more effectively. For example, Alipanahi et al. developed Deep Bind, a CNN-based architecture that predicts DNA- and RNA-binding protein specificities. Deep Bind could identify both known and novel motifs, assess sequence alterations, and detect functional single nucleotide variations (SNVs). Similarly, Zhou and Troyanskaya employed CNNs to predict chromatin marks from DNA sequences, while Kelley et al. created Basset, an open-source framework for predicting DNase I hypersensitivity and assessing the impact of SNVs on chromatin accessibility. CNNs were also used by Anger Mueller et al. to predict DNA methylation states in single-cell bisulfide sequencing and by Koh et al. to denoise genome-wide chromatin immunoprecipitation sequencing data, improving the accuracy of chromatin mark prevalence estimates.

In addition to CNNs, other deep learning architectures have been explored in genomics. Sparse auto encoders (AEs) have been used for tasks such as classifying cancer cases from gene expression profiles and predicting protein backbones. Moreover, deep neural networks have significantly advanced drug discovery processes in genomic medicine.

VI. MOBILE

Sensor-equipped smartphones and wearable are revolutionizing mobile health monitoring applications. The boundary between consumer health wearable's and medical devices is blurring, enabling a single device to monitor a wide array of medical risk factors. These devices hold the potential to provide patients with personal analytics that could enhance health, support preventive care, and assist in managing chronic conditions. Deep learning is anticipated to play a crucial role in analyzing data from these devices. However, the application of deep learning to health care sensing is still emerging, primarily due to hardware constraints. Efficiently running deep learning models on mobile devices to process noisy and complex sensor data remains challenging and resource-intensive.

Recent studies have explored ways to address these hardware limitations. For instance, Lane and Georgiev introduced a low-power deep neural network inference engine that leverages both the Central Processing Unit (CPU) and Digital Signal Processor (DSP) of mobile devices without significantly burdening the hardware. They also developed Deepx, a

| ISSN: 2394-2975 | www.ijarety.in| | A Bi-Monthly, Double-Blind Peer Reviewed & Referred Journal |

|| Volume 7, Issue 6, November-December 2020 ||

| DOI:10.15680/IJARETY.2020.0706018|

software accelerator designed to reduce the resource demands of deep learning on mobile devices, making it possible to perform large-scale deep learning more efficiently and outperforming cloud-based solutions.

While we did not find studies applying deep learning directly to commercial wearable devices for health monitoring, there has been research on data from phones and medical monitors. One notable area is Human Activity Recognition (HAR), which, though not focused exclusively on medical applications, holds promise for clinical use. For example, Hammerla et al. used CNNs and RNNs with LSTM units to predict freezing of gait in Parkinson's disease patients. This condition involves difficulty in initiating movement, such as walking. Their study, which used accelerometer data from various body parts of 10 patients, found that RNNs significantly outperformed other models, including CNNs, highlighting the potential of deep learning for clinical activity recognition.

Another study by Zhu et al. used CNNs to predict Energy Expenditure (EE) from triaxial accelerometer and heart rate sensor data during physical activities. Accurate EE prediction is crucial for tracking personal activity and preventing chronic diseases. Their deep learning approach yielded significantly better results compared to regression and shallow neural networks.

In other clinical areas, deep learning has enhanced the analysis of portable neurophysiological signals such as Electroencephalogram (EEG), Local Field Potentials (LFPs), and Photoplethysmography (PPG). For instance, Sathyanarayana et al. applied deep learning to predict sleep quality using actigraphy data, achieving 46% better performance in specificity and sensitivity compared to logistic regression. This research underscores the growing potential of deep learning in analyzing and improving clinical outcomes based on sensor data.

VII. CHALLENGES AND OPPORTUNITIES

Despite the promising advancements in deep learning for health care, several key challenges remain unresolved. Addressing these issues is crucial for advancing the application of deep learning in this domain. Here, we outline the primary challenges and potential research directions to overcome them:

Key Challenges

- Data Volume: Deep learning models require vast amounts of data to function effectively. For instance, fully connected neural networks have numerous parameters that need extensive training data. In fields like computer vision or speech recognition, large datasets are available, which contributes to the success of deep learning. However, health care data is more limited due to the relatively small global patient population and the complexity of diseases. Training robust deep learning models in health care demands more extensive data than is typically available.
- Data Quality: Health care data is often heterogeneous, noisy, and incomplete, posing significant challenges for training deep learning models. Unlike more structured fields, medical data may suffer from issues such as sparsity, redundancy, and missing values, making it difficult to develop accurate models.
- Temporality: Diseases evolve over time in unpredictable ways, yet many deep learning models are designed for static inputs. Developing models that can effectively handle temporal data, such as longitudinal patient records or evolving health conditions, is a key area needing innovation.
- Domain Complexity: The complexity of medical problems, including the heterogeneity of diseases and the incomplete understanding of their progression, makes deep learning applications more challenging compared to more straightforward domains like image or speech analysis. The limited number of patients and the complexity of medical conditions add to these difficulties.
- Interpretability: Deep learning models are often criticized as "black boxes," especially in health care where understanding the reasoning behind predictions is crucial. Clinicians need to understand how models arrive at their conclusions to trust and act on them. Enhancing model interpretability is essential for adoption in clinical practice.

VIII. FUTURE RESEARCH DIRECTIONS

a) Feature Enrichment: To address data limitations, integrating diverse data sources (e.g., EHRs, social media, wearable devices, and genomic data) can provide a more comprehensive view of patient health. Research could focus on methods to effectively combine these heterogeneous data sources and develop deep learning models that can leverage a broad array of features.

IJARETY

| ISSN: 2394-2975 | www.ijarety.in| | A Bi-Monthly, Double-Blind Peer Reviewed & Referred Journal |

|| Volume 7, Issue 6, November-December 2020 ||

| DOI:10.15680/IJARETY.2020.0706018|

b) Federated Inference: Building deep learning models across different clinical institutions while preserving patient privacy is crucial. Federated learning approaches, which enable model training across distributed data sources without sharing sensitive information, could be a promising solution. This area intersects with cryptography and privacy research, such as homomorphic encryption and secure multiparty computation.
c) Model Privacy: Ensuring privacy in deep learning models, especially when scaled through cloud computing, is critical. Techniques like differential privacy, which protects individual data points, need to be adapted for deep learning models. Ongoing research is required to develop effective privacy-preserving methods that do not compromise model performance.
d) Incorporating Expert Knowledge: Integrating existing medical knowledge into deep learning models can enhance their effectiveness. This could involve using medical literature and databases to inform model development or applying semi-supervised learning techniques to leverage both labeled and unlabelled data.
e) Temporal Modelling: Developing models that account for temporal aspects of health data is essential. Research should focus on temporal deep learning models, including those that use RNNs, memory networks, and attention mechanisms, to better capture and predict changes in patient health over time.
f) Interpretable Modelling: Improving the interpretability of deep learning models is crucial for clinical adoption. Research should focus on methods to make model predictions more understandable, such as developing algorithms that explain model decisions or integrating existing explanatory tools into deep learning frameworks.

By addressing these challenges and exploring these research directions, the field of deep learning in health care can advance towards more effective, interpretable, and privacy-preserving solutions that significantly benefit patient care and clinical practice.

IX. APPLICATIONS

Deep learning methods have indeed proven to be powerful tools in enhancing traditional machine learning approaches, particularly in fields like computer vision and natural language processing. Their application in health care is no exception, offering significant potential to transform data analysis and improve clinical outcomes. Here's a summary of the key points regarding the role and future of deep learning in health care:

a) Complex and Heterogeneous Data: Biomedical research increasingly relies on complex, heterogeneous, and often poorly annotated data sources, such as EHRs, imaging data, -omics profiles, and sensor data. These data types are generally unstructured, making them challenging to analyze with conventional methods.
b) Early Successes with Deep Learning: Initial applications of deep learning in health care have demonstrated the ability to model and learn from these diverse data sources effectively. For example, deep learning has improved predictive accuracy in various clinical applications and has shown potential in integrating different types of data.
c) Challenges and Areas for Improvement: Combining and processing diverse data types (e.g., EHRs, genomics, wearable data) remains a challenge. Improved methods for data integration are needed to leverage the full potential of deep learning. Deep learning models are often criticized for being "black boxes." Enhancing the interpretability of these models is crucial for clinical acceptance and to provide transparency in decision-making. Ensuring the privacy and security of health data is paramount. Advances in privacy-preserving techniques, such as federated learning and differential privacy, are necessary to protect sensitive patient information. Many health conditions evolve over time, and deep learning models must effectively handle temporal data to capture disease progression and changes in patient status.
d) Future Directions: Deep learning has the potential to scale up to handle large volumes of patient records and to create comprehensive patient representations. Such systems could integrate data from various sources into a unified model, improving clinical support and decision-making. Ideally, deep learning models would provide a single, integrated patient representation that includes data from EHRs, genomics, environmental factors, and wearable devices. This approach could lead to more personalized and accurate health care solutions. Deep learning can facilitate both hypothesis-driven and exploratory research in clinical domains. It can aid in clustering, visualizing patient cohorts, and stratifying disease populations, thus supporting more nuanced and effective research and data analysis.
e) Real-World Applications: The success of deep learning models, such as Wang et al.'s work on Parkinson's disease subtyping, highlights the potential of these approaches in real-world health care scenarios. By capturing disease progression patterns more accurately, deep learning can enhance disease subtyping and treatment strategies.

Deep learning holds great promise for advancing health care data analysis and clinical decision support. While early applications have shown significant improvements in predictive power and integration capabilities, ongoing challenges related to data integration, interpretability, security, and temporal modelling must be addressed. With continued

| ISSN: 2394-2975 | www.ijarety.in| | A Bi-Monthly, Double-Blind Peer Reviewed & Referred Journal |

|| Volume 7, Issue 6, November-December 2020 ||

| DOI:10.15680/IJARETY.2020.0706018|

advancements and refinement, deep learning can pave the way for the next generation of predictive health care systems, capable of handling vast amounts of data and providing comprehensive patient insights.

X. CONCLUSIONS

The rapid evolution of deep learning offers unprecedented opportunities to tackle the challenges posed by complex, high-dimensional biomedical data. Biomedical datasets, including electronic health records, medical imaging, -omics data, and sensor inputs, are vast, heterogeneous, and often unstructured. Traditional data mining and statistical methods have struggled to cope with the intricacies and scale of these data types, largely due to limitations in feature engineering and the need for substantial domain knowledge. Deep learning, with its ability to automatically extract meaningful patterns and build predictive models, presents a powerful solution. End-to-end learning models, such as convolutional neural networks (CNNs) for medical imaging and recurrent neural networks (RNNs) for time-series health data, are already demonstrating their potential in biomedical research. These models can uncover hidden relationships in the data, offering insights that were previously unattainable with conventional approaches. From predicting disease onset to aiding in precision medicine, deep learning's ability to process and interpret biomedical data is a significant step toward improving healthcare outcomes. A primary concern is the black-box nature of many deep learning models, which makes it difficult for domain experts, such as clinicians, to interpret how these models arrive at their conclusions. The opacity of these models poses significant barriers to trust and usability in clinical settings. As healthcare relies heavily on explainability and interpretability for patient safety and legal accountability, there is an urgent need for interpretable deep learning models. Developing these interpretable models could enhance the acceptance of deep learning in healthcare, allowing professionals to verify and trust the outcomes produced by these systems. Techniques such as attention mechanisms, saliency maps, and model simplifications are steps toward achieving this goal. Future research must focus on creating hybrid systems that balance predictive power with transparency, ensuring that actionable insights derived from biomedical data are both accurate and understandable. In conclusion, deep learning holds tremendous promise for transforming how we analyze biomedical data. By addressing the challenges of interpretability and trust, deep learning can be fully harnessed to improve medical decision-making and patient care, ultimately advancing the healthcare field.

REFERENCES

1. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115-118. https://doi.org/10.1038/nature21056
2. Ravi, D., Wong, C., Deligianni, F., Berthelot, M., Andreu-Perez, J., Lo, B., & Yang, G. Z. (2017). Deep learning for health informatics. IEEE Journal of Biomedical and Health Informatics, 21(1), 4-21. https://doi.org/10.1109/JBHI.2016.2636665
3. Miotto, R., Wang, F., Wang, S., Jiang, X., & 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
4. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., & van Ginneken, B. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60-88. https://doi.org/10.1016/j.media.2017.07.005
5. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444. https://doi.org/10.1038/nature14539
6. Zhou, Y., Hou, J., Lu, Y., & Shen, L. (2020). Deep learning and its applications in biomedical image segmentation. IEEE Reviews in Biomedical Engineering, 13, 97-110. https://doi.org/10.1109/RBME.2020.2987262
7. Jiang, Z., Dong, X., Li, Y., & Tian, Y. (2022). Explainable deep learning: A survey of methods and applications in healthcare. Journal of Biomedical Informatics, 124, 103961. https://doi.org/10.1016/j.jbi.2021.103961
8. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347-1358. https://doi.org/10.1056/NEJMra1814259
9. Xing, L., Wu, Q., & Liang, Y. (2020). Interpretable deep learning in healthcare: Potential applications and challenges. Journal of Medical Systems, 44(8), 1-10. https://doi.org/10.1007/s10916-020-01618-7
10. 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.

IJARETY

International Journal of Advanced Research in Education and Technology

ISSN: 2394-2975

www.ijarety.in editor.ijarety@gmail.com