Deep Science Publishing
https://doi.org/10.70593/978-93-49910-80-5_4

Chapter 4: Applying machine learning to enhance diagnostic accuracy and device functionality

4.1. Introduction

Machine learning is an exciting, emerging field in which computers learn independently from the data they collect and their experience with this data. A subset of artificial intelligence, machine learning has the potential to significantly enhance the functionality of medical devices and any other processes reliant on accurate classification of highly-dimensional sensory data. For classification tasks associated with high-dimensional datasets, suitable tools may not be readily available (Bhelkar & Shedge, 2016; Chalapathi et al., 2019; Analog Devices, 2020).

These tools may range from automated tools such as those utilizing convolutional networks for pixel data classification to customized classifiers using machine learning techniques for both pixel data and non-pixel data. While many medical devices utilize classification for diagnostic purposes, the vast majority of diagnostic devices do not use machine learning. However, we believe that there are substantial advantages in the use of machine learning to enhance the diagnostic capabilities of devices in the medical domain. In this chapter, we will discuss the domain of medical diagnostics and the need and potential for the incorporation of machine learning in diagnostic modalities currently in use.

Over the last few decades, there has been enormous and rapid progress in our ability to generate and collect data and increase improvements of computer algorithms utilizing multivariate techniques to explore the hidden predictive potential of that data. This has occurred at the same time that powerful sensors have been developed to collect medical data associated with the body tissues and fluids. In parallel, there have been increasing advances in computing power and storage as well as in the capabilities of the algorithms for utilizing this computing power to explore the correlated dataset in an efficient and effective manner. As a result, we are amidst a revolution in our ability to utilize data for improvement in diagnosis, prognosis, and prediction of therapeutic response at all levels of evidence (Li et al., 2024; Xu et al., 2024).

4.1.1. Background and Significance

Machine learning has established itself as an essential player in the realm of artificial intelligence, one of the fastest growing fields in technology and science, rapidly transforming business sectors and offering creative solutions to prevailing dilemmas. ML effectively relies on algorithms to assimilate and scrutinize data in novel and sophisticated manners, allowing technology to perform tasks that surpass the efficiency and accuracy of human capabilities. Consequently, ML has become an integral cog in the operation of intelligent systems. The triumphs that intelligent systems have experienced is in large part due to the recent dramatic advancements in the expounding of ML theory as well as powerful enhancements in the computational heft of computers.

Given this backdrop, it is perhaps logical that ML is being progressively embraced by numerous knowledge-based medical disciplines such as radiology, dermatology, pathology, neurology, ophthalmology and cardiology. Indeed, several studies have revealed that ML and, more precisely, deep learning, which performs the modeling task using sparse representations that are learned by convolutional neural networks from the realm of associated data, has successfully outperformed expert medical practitioners in the diagnosis of radiologic chest disease, the prediction of prognosis following image-guided radiotherapy for brain metastases, and sleeplessness categorization with polysomnography data. However, in terms of experimental validation and forthcoming clinical translatability, the grouping of employed neural networks and final datasets remains in its infancy when compared to the MRI biometric community.

4.2. Overview of Machine Learning

Despite the increasing popularity of machine learning (ML), numerous and diverse definitions try to demarcate the subject. One of the most common attributions defines ML as a computer science discipline that develops algorithms that can learn from observed data and make predictions based on these observations. More formally, ML can be defined as a research field aimed at developing computationally feasible procedures that can estimate a mapping function from inputs X to outputs Y based on a database, which has been generated by a probability distribution P(X, Y). When P maps the input/output space, the function can be characterized as error-free. However, since the observations are usually corrupted by noise, the goal of ML is to find an approximation to P. This approximate function, however, is usually defined for a finite number of points in X, whereas it can be used over the entire support of P to define the mapping function.

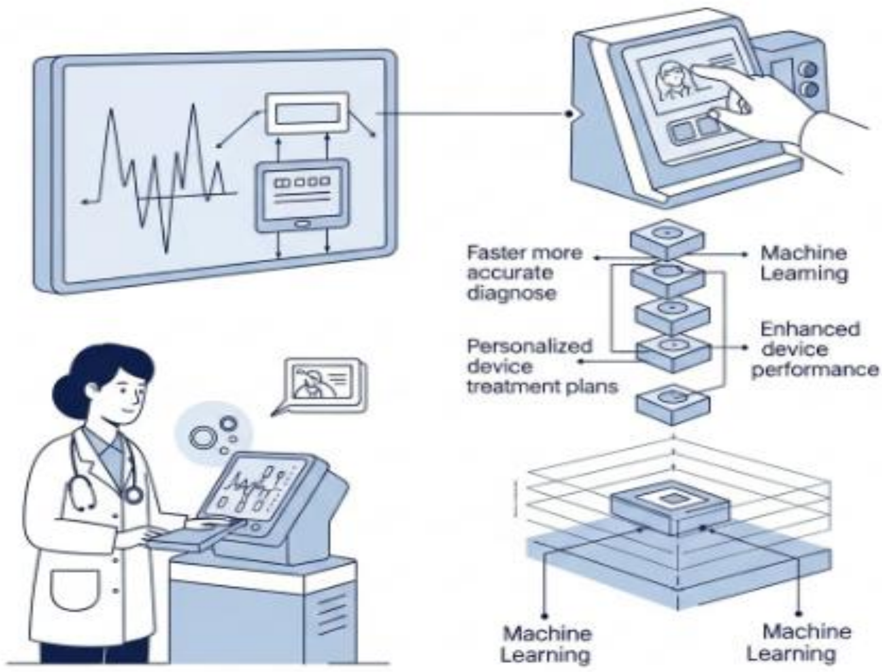

Fig 4.1 : Applying Machine Learning to Enhance Diagnostic Accuracy and Device Functionality.

A related definition views ML as a set of tools to discover patterns in data, enabling data-driven decision-making. This ontology matches the original intention of ML, which is to infer generalizable models from empirical data, a feature that differentiates ML methods from mathematical modeling. This implies two additional characteristics of ML: First, ML is agnostic regarding the model: classical statistical approaches are characterized by the specification of the mathematical form of the model, whereas ML selects or estimates the model from the data, providing a greater number of possible model structures, both linear and non-linear. Second, ML offers flexible approaches for forecasting. While exploratory data analysis focuses on obtaining information from data, using summary statistics and graphs, ML focuses on obtaining knowledge from data by forecasting test data.

ML can be categorized into supervised, unsupervised, and reinforcement learning. In supervised learning, a model is built based on a training set containing features and the labels of each feature vector. When the model is validated and can generalize well in unseen data, it can be used for prediction tasks. In unsupervised learning, only observed data is used, and the learning algorithm is asked to find patterns or groupings. In reinforcement learning, an agent takes actions in an environment to maximize cumulative reward. However, diminished return phenomena appear because of the sequential nature of the tasks.

ML techniques have attracted enormous interest due to their suitability for data-rich environments, such as finance, e-commerce, marketing, social media, and healthcare, among others. Within the healthcare context, applications include predicting diseases, adverse events triggered by therapeutic interventions, and treatment responses from electronic health record data; improving disease surveillance from event monitoring; diagnosing diseases from images, genomic and proteomic data; refining health system management and performance analysis by effectively allocating resources, predicting disease prevalence, and enhancing quality of care.

4.2.1. Definition and Key Concepts

Machine learning (ML), a subset of artificial intelligence (AI) that has emerged over the past few decades, allows computer programs to automatically receive information from data and use it to support decision-making without requiring explicit programming for the particular task. Supervised learning, one of the two major classes of ML algorithms, requires a training set of input-output pairs that allows the program to learn the association of inputs with desired outputs. Once trained, the model can operate on new input data, producing a predicted output. In contrast, the other major class of ML algorithms, unsupervised learning, does not require labels. While supervised learning is more commonly used in image classification tasks, especially in medical imaging, unsupervised learning is often employed in exploratory data analysis problems.

For supervised learning to be useful, a large collection labeled observation pairs must be acquired prior to developing a model for general use. Labeling of medical images for tasks such as disease classification is a labor-intensive process and requires the expertise of medical professionals. Alternatively, unsupervised learning can search through large datasets, locating structure such as clusters or outliers that create a basis for hypotheses or augment the training of models supervised by more limited labeled datasets. Speed and compatibility with parallel processing make so-called "deep" neural networks particularly computationally attractive methods of ML, especially in recent years as massive computer processing capability has become available to researchers.

Technical advances in both ML methodologies and available computational resources have fueled its rapid expansion into many aspects of medicine. Classification tasks, which can encompass the detection and diagnosis of particular diseases, conditions, and physiologic states, represent the most common applied use of ML in the medical domain.

4.2.2. Types of Machine Learning

Machine learning methods can be broadly categorized into supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning. In supervised learning, the input data is annotated with the desired corresponding output, such as class labels in classification applications or numerical values in regression applications. Given the labeled input-output mappings, the goal is to learn a deterministic or stochastic function to be able to reliably predict the output for unseen data. In unsupervised learning, only the input data is given, and the goal is to find hidden structures in the data. Examples include clustering algorithms and generative models. In reinforcement learning, an agent learns to interpret the current state of an environment and then make decisions with the goal of maximizing a reward signal. The learned policy is then employed to guide the agent's behavior in the environment.

Most commonly used machine learning algorithms, such as logistic regression, support vector machines, linear regression, and deep neural networks, are in the supervised or the semi-supervised category. Commonly used unsupervised algorithms include Gaussian mixture models, kernel PCA, dictionaries and auto-encoder based compressive sensing, dictionary learning, and latent variable models such as probabilistic PCA, transfer learning, and variational autoencoders. Some recent attention has also been paid to generative adversarial networks for generating new samples and for semi-supervised learning.

4.2.3. Applications in Healthcare

Machine learning is applied in medicine to predict medical outcomes, derive insight from sensor data, and help with device programming and testing. It solves key problems, such as, diagnosis; precision diagnostics considers a growing number of omics, imaging with deep learning, and physiological measurements to create a multidimensional disease landscape; longitudinal assessment of clinical condition considers integrated pathology from omics, wearable sensors, images, and clinical tests to derive insights on prognosis; and precise prediction of clinical events or therapy response considering large input factors and prediction horizons. Abundant data, computing power, cloud-assisted deployment, model transfer, and clinical unmet needs enable the underlying ecosystem for machine learning in medicine. Healthcare devices increasingly come with integrated embedded sensing and processing capabilities. These devices can benefit from improved users and patient outcomes via enhanced device functionality, device testing and optimization, and better understanding of individual sensor data. Embedding intelligent algorithms within the devices can enhance functional capabilities – allowing, for example, non-contact sleep stage monitoring to enable personalized sleep optimization, and near real-time fetal heart rate monitoring in varied environments to enable effective labor monitoring.

However, intelligent healthcare devices need to be validated thoroughly. Generalizing model training with multi-domain, temporal, and subject-specific variations in stress and sleep is important for predictive solutions; this requires extensive modeling. Modeling these variations would usually also require a large dataset, which is not possible for device-related physiological signals, due to the existence of significant inter-subject and intra-subject variations. Hence, transfer learning helps with adaptation models that stress on subject-related factors or exploratory models that use information from multi-subject data to gradually adapt and converge to solution, while generalizing across the variations. Additionally, the processing implementation also requires careful design of the feature extraction, filtering, and sensor data glitch handling aspects along with embedding the models. Further, these require careful lifecycle monitoring for enablement and scheduled recalibration due to variations stemming routine changes, health condition changes, device conditions, and system drift.

4.3. Diagnostic Accuracy in Healthcare

Diagnostic tests are very important in order to classify and possibly assess and characterize a disease or a health problem. Classifying if a patient has cancer, is diabetic or has a coronavirus infection, is only a few examples of some possible diagnostic tests. The probability to detect the disease, if it is really present, is defined as positive predictive value. The probability to detect a common disease, if it is really absent, is defined as a negative predictive value. Both values are a subject of the applied test accuracy and represent the quantitative concept of diagnostic test accuracy. Positive and negative predictive values are functions of the condition prevalence and describe accurate tests to identify a target population that is tested to classify as diseased or non-diseased.

A non-choice will be made in a medical imaging task. In most imaging modalities, a probe will send radiation, light, sound or magnetic sources to the pathology under consideration. A detection device will receive the response signal from the stimulus. The sensitivity of a medical imaging system should be sufficient so that normal patients are not disturbed by interference features in the diagnostic image which could confuse the diagnostics. Also, there should be a high efficient choice probability that abnormal or diseased patients will have the correct expectation-maximization. We reviewed the mutual relationship between diagnostic accuracy and device functionality. Low imaging accuracy can be compensated by advanced diagnostic methods. We will have a look at the performance compromises for some advanced standard methods as well as their possible solutions.

4.3.1. Importance of Accurate Diagnostics

Accurate diagnostics are key to the field of healthcare. Whether for tracking health status, identifying infection or disease, preventing adverse health events, or monitoring recovery to avert complications, diagnostics must be trustworthy and reliable. If a fault occurs in a diagnostic, the implications can be disastrous. Even a simple at-home pregnancy test can cause serious implications for a person's health if its results are not accurate. Consider the case when a patient is misdiagnosed as not being pregnant, and thus does not seek pre-natal care that may diagnose or help her mitigate pregnancy-related health issues, such as gestational diabetes or pre-eclampsia, and later delivers a severely unhealthy infant. Doctors, too, need to be thoroughly trained in how to perform when using and interpreting results from diagnostic devices and tests.

Misdiagnosis can also lead to improper intervention, including medical or surgical procedure errors that carry significant risks, which potentially can worsen health. Consider a tumor that has been diagnosed as benign based on an imaging test to be removed laparoscopically, only to later find that it is confirmed malignant. In this case, it is likely that the patient will need to undergo additional treatment with non-targeted chemotherapy or radiation, which have major systemic effects, or radiation or chemotherapy for a more localized approach. As with pregnancy tests, diagnostic testing may carry the burden of false negatives, but false positives in diagnostics can be equally impactful to patient wellbeing. Misdiagnosis by either result may lead to either unnecessary or lack of treatment, resulting in complications, additional tests and diagnoses, or even increased financial burden.

4.3.2. Current Diagnostic Methods

The diagnostics landscape comprises several traditional diagnostic methods. Physical examination, also known as a clinical exam, is the first step in assessing a patient's healthcare problem. Various forms of imaging modalities, including MRI, CT, and ultrasound, offer visualizations of internal tissues and organs of the human body. Invasive surgical procedures may also be performed if necessary, causing a mixture of reactions, including stress, fear of disease, and increased distrust associated with treatment. In challenging cases, these may be combined with second opinions, but more often through trial-and-error studies, such as guided therapy and diagnostic studies, which add time to the overall process. Defining monitoring and managing disease are part of the feedback during a rational clinical decision-making process. Laboratory investigations contribute towards more precise and quicker diagnostics. This includes microbiology testing, such as bacteria or parasite screening, immunology studies, hormone, blood, and urine analysis, and finally, histopathology investigations.

Nevertheless, these conventional methods have been associated with significant problems, such as limited availability of skilled professionals, sample storage issues, and increasing in-disease prediction costs. Especially for patients with chronic illnesses such as diabetes, asthma, and hypertension, current screening and diagnostic approaches face significant hurdles to achieving patient-centric timely diagnostics due to limited sensitivity. Moreover, samples such as blood and tissue cannot be easily stored and can be difficult to obtain in patients such as infants without severely aggravating conditions. These examples show unmet needs in timely diagnostics in various patient populations, resulting in healthcare services that compromise patient outcomes and experience.

4.3.3. Challenges in Diagnostic Accuracy

Diagnostic tests play a crucial role in confirming the presence of a disease to aid the clinician in clinical decision-making. A diagnostic test must have high sensitivity and specificity for it to be reliable. As medical science advances, it has reached a stage where we have a wide range of sophisticated diagnostic processes available to us, but no diagnostic test today has 100% sensitivity and specificity for every clinical problem. Results are more accurate when more tests are utilized that complement each other. Many diagnostic tests including imaging and biopsies are adjunct to clinical assessment and most clinicians prefer them to augment their clinical impression rather than replace it. Accuracy of diagnostic tests is affected by several factors and recently, attention has been drawn to the overriding nature of the pretest clinical probability. In other words, the ability of a diagnostic modality in predicting the disease state is dependent on the state of the patient before conducting the test and the pretest probability is inferred from prior knowledge based on study of the natural history of diseases.

Diagnostic tests can have suboptimal performance for a variety of reasons. The most fundamental issue is the difficulty in accurately determining the status of cleanly defined disease and nondisease. Most conditions pursued by diagnostic tests are not only defined by signs and symptoms but are also understood to unfold over periods of time. There exists a group of individuals who possess some characteristics of disease or health but do not yet fall clearly into either category, and this category of individuals is often described as diagnostic dilemmas. In many instances, diagnostic dilemmas would resolve over time but in others, it would not, highlighting the fact that the boundaries of diagnostic categories are blurred. Further, the patient's stage of the disease and the presence of confounding comorbidities will also impact the accuracy of tests, especially as it relates to sensitivity. It follows that the error or false results for a diagnostic modality would not be uniform across all the patients being tested for a specific condition, leading to controversies on the sensitivities and specificities stated for most diagnostic tests.

4.4. Machine Learning Techniques for Diagnostics

The availability of increasingly higher-quality electronic health records is opening the door for innovative approaches to provide population-level insights that may improve the understanding of disease risk and treatment response. Because health data is made up of varied modalities that may require differing preprocessing steps and represent various perspectives regarding health risk, many machine learning techniques may fit, and even complement, a research question of interest. Additionally, supervised, self-supervised, and unsupervised techniques alike may prove useful. Below, we provide an overview of several prominent applications of ML methods for predictive modeling, clustering, and risk stratification related to disease diagnosis.

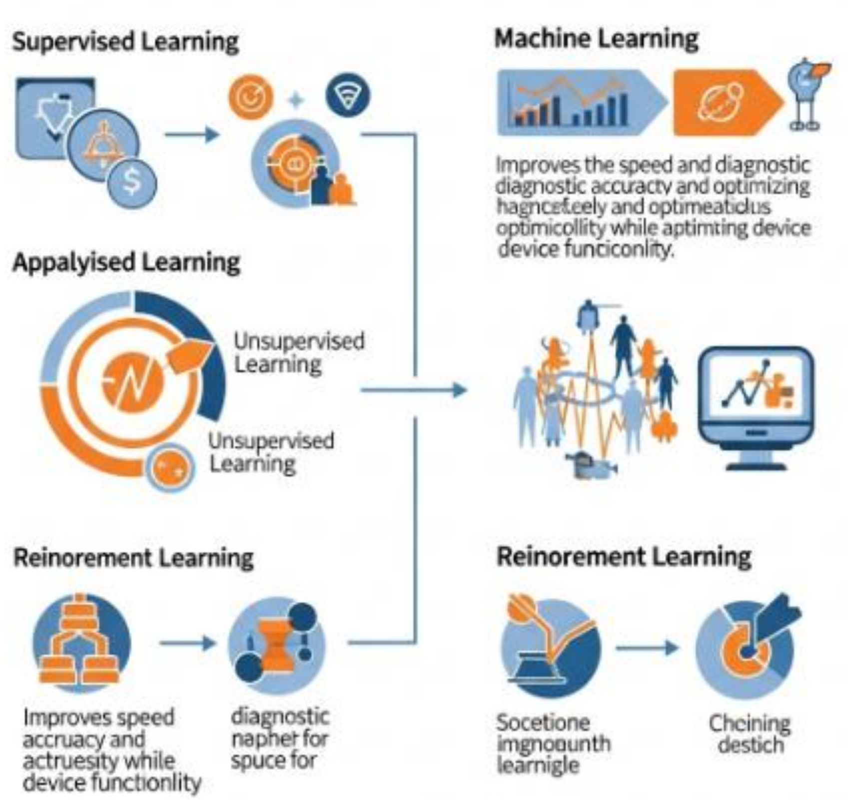

Fig 4.2: Machine Learning Techniques for Diagnostics.

4.4.1. Supervised Learning Approaches Predictive modeling has been a common thread in many ML applications to health data. Knowledge of ground truth labels or outcomes of interest during model training allows the vast amounts of health data to be mined into powerful diagnostic models. Many techniques fall into this category, from traditional supervised regression and classification to more modern deep learning techniques. Internal validation using temporal or external study cohorts has been crucial for establishing the robustness of these predictive models, and much work lies ahead to quantify the clinical impact of these models. Adding structure to the predictive modeling task, hierarchical labels on the disease state can also be incorporated to further complement rare or heterogeneous disease states.

4.4.2. Unsupervised Learning Techniques In the absence of adjudicated event labels on the dataset of interest, unsupervised learning techniques enable deep exploration and embedding of clinical data sources. Self-supervised methods, similar to supervised algorithms, can be trained on large amounts of data to characterize clinically relevant features. Reinforcement learning has gained traction with clinical design help, optimizing its tuning to improve patient and clinician outcomes alike. Traditional unsupervised clustering techniques such as K-means, Expectation-Maximization, and Hierarchical clustering provide interpretable and simplistic disease discovery aids. Topic modeling, probabilistic clustering, predictive modeling, and matrix factorization can also discover latent themes from textual or categorical data. Autoencoders applied to rich and multi-modal datasets may also yield similar dimensionality reductions. Thus, group-level latent features or modes of data variability can be mapped, grouped into disease clusters, or even linked back to ground truth event labels.

4.4.3. Deep Learning Applications The superior performance of deep learning hinges on the availability of sufficient amounts and variety of risks for training. Yet, recent breakthroughs in the application of self-supervised techniques have demonstrated the efficacy of leveraging massive amounts of unlabeled data to learn embeddings. These embeddings can be cleverly paired with supervised learning techniques to maximize the concordance between the supervised and self-supervised studies. Alternatively, pre-trained embeddings can serve as input seeds for downstream supervised tasks with limited amounts of training examples. Various encoding models have been devised to process multi-modal clinical data, from categorical/multi-hot encoders to continuous encoders for health imaging, audio, or time-series data.

4.4.1. Supervised Learning Approaches

The most prominent family of machine learning techniques are the supervised learning approaches. These model the mapping from the input space to the output space, and require labeled training data to develop the model that represents the mapping. The labeled data typically consists of training examples made up of input-output pairs, it is the data from which the algorithm will learn. The input data is usually a nanovector, and the output can include a class label, a class probability, real-valued quantities, or the structured output of objects such as sequences, trees, or graphs. The first two scenarios pertain to classification tasks, while the last belongs to the domain of structured output prediction. This essential requirement of labeled data creates a problem because it is expensive and often difficult to obtain. However, when adequate labeled data is available, supervised learning techniques can attain very high accuracy on unseen data, and are among the most prevalent techniques applied to real-world problems.

The family of supervised learning methods and their variations have seen recent success in diagnostic and prognostic prediction tasks in a variety of application domains. Most of the work follows a similar methodology pattern: construct the main data set from raw data; extract features that inform the diagnostic label, or class, or the prognostic threshold, and which describe its relationship with the input; develop a mapping model that models the relationship between the feature vectors and the diagnostic or prognostic output; and finally validate the model on a test set. These tasks contribute to the majority of the main areas, or classes, of work within the ML community. In the following, we will briefly summarize the most well-established supervised learning methods.

4.4.2. Unsupervised Learning Techniques

Unsupervised learning techniques are a category of machine learning where the model learns from unlabeled data. Unsupervised learning seeks to discover patterns or structures in complex data; Information is learned only through features, kept internal to a model, that are encouraged to extract explanatory and relevant aspects of interest. This property allows fruitful interactions with human subject matter experts, as models can be interrogated for insights and visualizations planned to help users better understand the task domain. The most common tasks performed are data clustering and dimension reduction. Data clustering groups samples according to how similar the samples are to one another. There are numerous algorithms available, such as DBSCAN, hierarchical clustering, and Gaussian mixture models. Dimensions reduction is commonly used to help visualize collections of datapoints in 2D or 3D when their initial feature representation is high-dimensional. Popular algorithms are t-SNE and PCA. However, because images are already low-dimensional objects, dimension reduction has yet to be applied to improve the data description. The most frequently used technique remains Euclidean distance. Other representations can be employed to overcome shortcomings of distance: nearest neighbor graphs; Laplacian eigenmaps, for taking into account the graph structure of the dataset; or unsupervised representation-based classification, which builds a more supervised object representation to perform the classification. Another limitation of the standard Euclidean distance is that it considers samples to be like one another only when they are proximate. To account for this, kernel methods employ kernels to measure the affinity between samples, which consider the influence of the whole dataset and the feature space to compute distance.

4.4.3. Deep Learning Applications

Deep Learning is a subfield of Machine Learning that is based on multiple layers of processing for feature extraction and transformation. Deep learning techniques can be used for detecting salient features in a data set and then for classifying the set, therefore using a hierarchical feature learning structure. Such techniques have been widely popularized after the introduction of Convolutional Neural Networks, which were inspired by biological processes and are widely used for image and video recognition, image classification, medical image analysis, and video analysis. The main advantages of deep learning methods are their capacity of handling high-dimensional data, their high accuracy compared to other methods, and the fact that they require little or no feature selection of the training data.

Though there are more and more applications based on CNN usage and its variants, most of the medical diagnosis deep learning techniques are still in experimental phase, and further research is needed for fully integrating such methods in the medical practice. As an example, a system based on a specific architecture learnt from a large dataset, calculating on patient subtypes measures like sensitivity, specificity, positive predictive value, negative predictive value, F1 score and area under the receiver operating characteristic curve. The system outperformed the radiologists in every biometric, achieving an area under the receiver operating characteristic curve of 0.956 in detecting esophageal cancer, thus behaving on par with expert radiologists.

4.5. Enhancing Device Functionality

Computer applications or services are typically reliant on centralized computing servers located remotely from end-users for computation or data servicing. However, with the rapid advancement of semiconductor technology and an array of sensors within devices, a novel trend is emerging towards the integration of machine learning directly into devices. In contrast to centralized functions, machine learning enables devices to perform highly specific smart tasks, realtime local data processing, or intelligence augmentation directly on the edge with minimal communication with back-end services. While consumer mobile devices have benefited from integration of machine learning capabilities, new powerful and flexible hardware platforms adapted from mobile computers are enabling other traditionally less-intelligent devices ranging from wearables, kitchen appliances, smart home utilities, cars, to industrial sensors to directly host machine learning tasks.

A key benefit of on-device machine learning is that it allows for real-time data processing. Data collected from sensors, such as cameras or microphones, can be processed locally without incurring network latency. Thus, tasks can be performed rapidly, improving user experiences. For instance, on-device machine learning on mobile cameras is used for computational photography applications that augment, enhance, or even change the captured images or video in real-time, boosting the capabilities of traditional cameras significantly. More generally, on-device machine learning enhances the capabilities of traditional accelerometers, gyroscopes, compasses, and GPS sensors, which are often not accurate meaningfully for practical applications including health tracking, navigation, and motion sensing. The convenience of having an ability to rapidly process data and provide inference on-device for various applications enables better interactions and improves the experience for users.

4.5.1. Integration of Machine Learning in Devices

Machine learning increases developmental capabilities to create intelligent devices that meet current diagnosis needs. This section describes the hardware methods for model integration, activities implemented by models running on the device, and the edgeML model. Device integration is necessary for a device to be released into the market or to be within a use case that requires, for getting a model from a laboratory environment to real use, hardware and software development and validation, device and application, considering necessary confidentiality and security. The main activity performed is inference, running a model to predict classes or output device diagnostic or correction values. While the inference activity is on the device, it is possible to run other activities from different people within the device application architecture, deployment model, and processing type. Device integration for inference is about the speed and accuracy of device configuration for model execution and the amount of inference runs done by the model on the existing resources. This section addresses only the basics of device inference for internal enhancement for its use cases, diagnostics, corrections, and content extraction.

The simplest way to integrate a model flavor into device architecture is to run inference on its chip or make a model onboard the application OS. Because this solution blends different environments, the inference resources must be enough to have a reliable model output latency with adequate accuracy for all observed situations, and the safety of the application function must not be affected by the execution of the inference. EdgeML defines the methodologies and techniques structures to allow ML model designs that run onboard the devices. This device integration approach is essential for devices aimed at unsupported use cases or with high security and confidentiality requirements. For example, supported devices, where sending the input data or the model output is detectable, and the action generated is risky, such as industrial machines.

4.5.2. Real-time Data Processing

Self-learning algorithms that analyze data in real-time offer the potential to elevate the specificity and sensitivity of devices and the data they provide caregivers. This enables devices to make recommendations while conducting the intervention, rather than waiting for the user to input information. This is appealing not only to the time- and attention-deprived caregiver or parent, but also to all researchers and practitioners concerned with the integrity of data captured by devices, as it removes the reliance on user input for accurate insights. This change could be particularly advantageous for devices capturing complex data and must make recommendations while the task is being executed, for instance, when a device provides auditory scaffolding during therapy and assesses quality in real-time.

Of note is the gradual transition from remote post-processing of data to live feedback and recommendations. At the core of this transition are strides in the development of real-time, efficient data analysis methods. These strides have been focused on both data and the type of machine learning algorithm. Remarkable contributions have been made in facial recognition, which is now at the core of many devices assessing behavioral health. Low-power CNNs enable real-time recognition, as do re-engineered architectures created for the analysis of streaming video. Several facial recognition learning approaches have been deployed by research groups, as open-source implementations jointly developed with companies, or as for-profit solutions. Other domains are now following suit, with potential benefits for data integration and quality, including breath analysis for telehealth applied to infections, lung ventilation function assessment for enhanced telemonitoring of COPD, and several post-stroke applications involving imaging during rehabilitation therapies.

4.5.3. User Experience Improvements

End-user experience is a key differentiator in how a device is accepted and continues to be used after initial purchase. This experience goes well beyond the general purpose functions that a device may offer. Typically devices are designed or used to look for a specific event. A key user interaction is therefore to present, mark, or annotate such events after measurement, and this task is typically completed using dedicated but relatively generic software after data capture. These definitive event markers provide the foundation for any further time-series analysis and potentially the data upon which diagnosis or therapy decisions are made. However, the task of marking events in wearable data can be time-consuming and arduous. For example, lay users may have to spend hours using software to find the minute number or message time of the actual occurrence of a particular event, and this level of user involvement can deter even an interested user from engaging with the software, leading to lost useful data.

Improvement in the user experience provided by devices is increasingly being discussed as a driver for consumer use and engagement with these technologies. Therefore, the development of device systems that can automatically identify event timings present obvious advantages including reducing user workload and confusion, allowing rapid data-linking during event-rich measuring, and potentially facilitating therapy decisions based on device data combined with pre-defined clinical rules or algorithms. The combination of event annotation and use of these annotated events may result in a feedback system loop. Regular reports to users could help them adjust their behavior to improve the outcome.

4.6. Case Studies

Specific applications of ML show computer vision tasks as image denoising and labeling, and general ML focus on predicting outcomes using tabular data. Some ML applications in healthcare demonstrate the variety of improvements that ML can enable in healthcare delivery, since a significant proportion of modern healthcare is based on health records analysis using standard statistical prediction approaches.

Machine Learning in Radiology

Radiology is undergoing an unprecedented transformation with the advent of ML techniques for image acquisition, enhancement, and analysis. AI is currently reshaping not only how images are analyzed but also how they are enhanced and generated. ML has already been able to efficiently denoise images acquired at low doses to decrease the exposure to harmful ionizing radiation. Fully autonomous deep learning algorithms have been or are projected to be released shortly to perform screening tasks, capable of detecting early breast cancer by mammography, lung cancer by CT, or diabetic retinopathy by fundus imaging. ML computer vision libraries allow virtually any hospital or medical center with a computer to build on previously documented propriety and non-proprietary algorithms to undertake the research to validate their performance using local data.

Predictive Analytics in Patient Monitoring

AR, the real-time overlay of virtual information on a real-world interface at the point of care, has been termed pregnancy's killer app. When used to disclose information on the state of the patient or fetus to the provider and patient that otherwise remained hidden, it has shown promise for decreasing premature birth rates. With the enhancement of predictive dashboards using ML models, providers at risk for burnout could have an interactive tool at their disposal to prioritize their care to patients with the greatest need at that moment.

AI in Pathology

AI has enabled the real-time digital analysis of key tissue features in biopsies, thus mobilizing the long-grounded theory that tissue structures could serve as an early disease warning niche. Biopsy adverse events, Wegener's granulomatosis, and other interstitial pulmonary diseases have been identified using tissue architecture in images that are, on their own, insufficiently effective for pathologists' qualitative assessment.

4.6.1. Machine Learning in Radiology

Advanced machine learning (ML) techniques, which are used commonly in a multitude of applications, including radiology, often do not examine what is learned, especially in conditions of "black-box" techniques such as deep learning. They require massive amounts of data for training, yet achieve superhuman performance on certain tasks. Other forms of ML rely more on permissible algorithms based on statistical theory, and expert knowledge of the domain of application, and commonly highlight exploratory analytics to augment other approaches at drawing inferences from data. These approaches indeed have an important role in radiology as well, especially in exploring discovered patterns, which could be useful in setting future hypothesis-driven work. A survey purported the large-scale availability of data among the downsides mentioned in the radiology space, which would lead to better optimization of deep convolutional neural networks (DCNNs), as well as limited greater accessibility and transparency for lesser regulated specialties such as radiology, since radiology deals with a considerable amount of sensitive data. Irrespective of the debate, deep learning operations, especially in radiologic applications concerning image categorization, localization, and detection, as well as semantic segmentation achieving superhuman performance, the radiology community is now cautiously optimistic about the potential use of AI and ML in clinical practice. The major aspects relate to detection, segmentation, and classification tasks, while a radiology application/environment includes images accompanying text reports, or attempts to learn directly, essentially leveraging multimodal data.

4.6.2. Predictive Analytics in Patient Monitoring

Traditional statistical techniques have long been utilized in health care settings for predictive capabilities and forecasting future events. Logistic regression and general linear models can uncover a variety of predictors related to the event of interest. However, these traditional supervised learning approaches typically only deal with a small number of predictor variables, as results become quickly overfit due to the high dimensionality of the data. The growth of important data sources offers a unique opportunity to provide better, more informative, and timely predictions. For example, digitally monitored telemetric data can offer the opportunity to inform clinical decision-making earlier than was previously considered possible. Experts have estimated that up to 40% of in-hospital cardiac arrests could be appropriate to predict. Early warning systems based on traditional statistics have existed for certain conditions for a while, including early warning scores. These early warning systems betray their heritage by relying primarily on physiologic signs, such as body temperature, heart rate, blood pressure, and urine output.

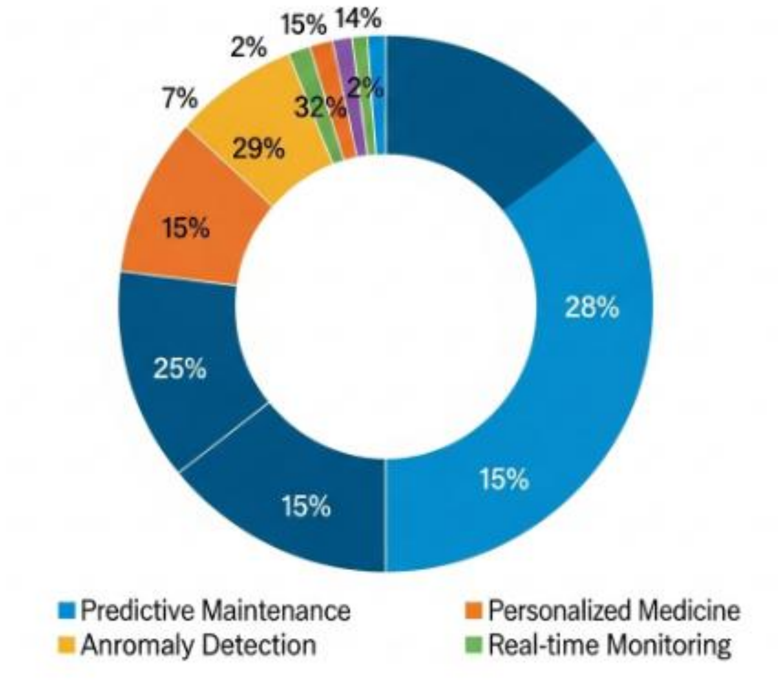

Fig : Applying Machine Learning to Enhance Diagnostic Accuracy and Device Functionality.

4.6.3. AI in Pathology

Pathology is a research and diagnostic tool that deals with the study of diseases, changes due to genetic and cellular mutation, and the reactions and interactions of all these changes. The digital transformation of pathology started with the introduction of whole-slide scanners. Such systems scan microscope slides and produce high-resolution digitized images of the whole tissue section, allowing remote access and image analysis. Digital pathology, unlike conventional light microscopy, increases the productivity of pathologists and allows the embedding of advanced quantitative analysis.

As all image-based medical specialties, pathology is also experiencing a rapid introduction of Machine Learning in its workflow. Pathology, amongst image-based specialties, is probably the one that has the largest room for disruption due to the clear limitations in accuracy suffered by human observers and the expected boost in performance by using Machine Learning technology in assisting pathologists' decisions. It is indeed commonly thought that digital image analysis will be instrumental in creating a more efficient and effective practice of pathology in the near future. The first use cases in aiding pathologists' decision-making are already acquiring regulatory approval and subsequently coming into the market, for example in the diagnosis of breast, prostate, and lung cancer.

In the last decade, the push to increase productivity and the advances in computer resources, especially deep learning, have rendered artificial intelligence (AI) a probability-like optimized technology for some of the most difficult daily tasks in the field. Medical image analysis is rapidly moving towards the transfer of responsibilities on the shoulders of algorithms, who are taking over many low-level tasks, such as the quantification of cellular markers from immunohistochemical images or the recognition of areas of interest in histological sections, such as tumoral areas. Long-term strategies seek to accomplish higher-level tasks using high predictive performance automatic procedures, starting from disease diagnosis.

4.7. Conclusion

With a fast-growing aging population together with the increasing incidence of cardiovascular and neurological diseases, there is a clear demand for healthcare systems that can offer smarter and automated decision-making supported by leading transformative tools such as artificial intelligence and machine learning methods. The widespread adoption of wearables, telemedicine technologies, cardiovascular and neurological diagnostic and bio-micro/nano-sensing devices promise to address this growing demand and to remedy several shortcomings in current healthcare practices. AI and ML methods can be applied to improve the diagnostic accuracy of several health issues such as mortality from cardiovascular diseases and neurological distress, as well add-on value to the functionality of such diagnostic and therapeutic devices. In this regard, the recent breakthrough with large language models and self-supervised learning paradigm inspires the prospect of a universal AI model for several tasks.

The future progress of using AI/ML in the space of health technology is towards the development of a one-stop comprehensive platform that can allow universal diagnostics anytime and anyplace, creating more equitable and sustainable access to healthcare and dispelling the huge burden for primary healthcare professionals. However, models must be responsibly developed and their misuse or malfunctions should be hence prevented. This comes with the necessity of collaboration between different fields of work: When it comes to medical diagnosis, AI cannot work in isolation — it must work with the help of and hand-in-hand with healthcare professionals. Moreover, any devices designed for the advancement in this field should be co-created with health end-users for safe and trustworthy use.

Developing diagnostics services and devices using AI/ML offers the possibility to build smarter and more successful solutions for digitizing, automating, and improving the accuracy of diagnosis processes. Moreover, in the realm of diagnostic services, the entire process could be shortened and made less complex as the solutions would be easier to transmit and understand, therefore making testing easily available for the general public and widening the intervention gap.

4.7.1. Future Trends

In the coming years, the bottom line focus of device development and deployment needs to be making devices capable of doing the jobs they were meant to do in as effective and accurate a way as possible. The constant drain on the attention resources of healthcare providers caused by device misuse or misunderstanding to the extent that it becomes common practice may risk lives, especially when it comes to at-risk patients. New generations of smart devices coupled with easier analytics pipelines and streamlined capabilities should enable informed users to utilize multi-modality sensing potential alongside intelligent analysis more effectively. Cross-device diagnostic capability enhanced by multiple modalities is a capability that is further still to be fully realized. It is certainly true that interesting results have begun to emerge in the cross-signaling space, but that is about it for the time being. Systems which fuse inputs from multiple modalities and/or devices into a single more complete view of the physiological makeup of the patient are obviously a holy grail solution. Untapped periods of interest cueing driven by the sharp rise and fall detection capability could stimulate ever-more responsiveness herein.

Careful thought must be given to the take up and scope of use of novel devices intended to supplement and enhance existing healthcare protocols, especially within non-traditional user groups. This is, however, a consideration which cuts both ways. Development of improved feedback methods on devices expanding capability has the potential to significantly reduce the attentional burden of healthcare providers while simultaneously increasing efficiency. Devices which trap engagement context information, and then present it in a form which is usable and of value to providers post-hoc generate novel possibilities for engaging users at risk but not previously included in regular health tracking.

References

Chalapathi, G. S. S., Chamola, V., Vaish, A., & Buyya, R. (2019). Industrial Internet of Things (IIoT) applications of edge and fog computing: A review and future directions. arXiv Preprints. arXiv

Li, H., Wang, X., Feng, Y., Qi, Y., & Tian, J. (2024). Driving intelligent IoT monitoring and control through cloud computing and machine learning. arXiv Preprints. arXiv

Xu, J., Wan, W., Pan, L., Sun, W., & Liu, Y. (2024). The fusion of deep reinforcement learning and edge computing for real-time monitoring and control optimization in IoT environments. arXiv Preprints. arXiv

Analog Devices. (2020). Embedded sensor platform with AI algorithms—Locally from big data to smart data. Analog Devices Technical Articles. Analog Devices

Bhelkar, V., & Shedge, D. (2016). Different types of wearable sensors and health monitoring systems: A survey. In Proceedings of the 2016 2nd International Conference on Applied and Theoretical Computing and Communication Technology (iCATccT) (pp. 43–48). IEEE.MDPI