Article

Not peer-reviewed version

Deep Learning and Reinforcement Learning in Electronic Health Records

Lincoln Bhattarai *

Posted Date: 14 April 2025

doi: 10.20944/preprints202504.1076.v1

Keywords: Health Informatics; Deep Learning; Reinforcement Learning; Traditional Approach; NEAT

Article

Deep Learning and Reinforcement Learning in Electronic Health Records

Lincoln Bhattarai

University of Texas - Arlington; lincolnbhattarai66@gmail.com

Abstract: Health informatics is the field that leverages technology and data to improvehealthcare delivery. The Electronic Health Record (EHR) is a digital version of a patient'smedical chart that may include everything from diagnoses and lab results to treatmentplans and clinicians' notes. EHRs are a key part of healthcare delivery, and they can be usedto help healthcare providers improve patient care, optimize treatment recommendations,and support clinicians in their decisions with data-driven evidence. In this paper, weprovide an overview of deep learning and reinforcement learning techniques and explorehow they are being used in real-world healthcare settings, the benefits they offer, and thechallenges to their adoption. We also focus on concrete examples that illustrate their impacton patient care and decision-making. The current trajectory - marked by promising pilotprojects, increasing investments, and a growing body of validation research - suggeststhat AI will become an indispensable component of healthcare. With careful stewardship, AI-powered EHR systems could usher in an era of smarter, more proactive, and more personalized healthcare for all.

Keywords: Health Informatics; Deep Learning; Reinforcement Learning; Traditional Approach; NEAT

1. Introduction

Health informatics is the field that leverages technology and data to improve healthcare delivery. A cornerstone of this field is the Electronic Health Record (EHR) – a digital version of a patient's medical chart that may include everything from diagnoses and lab results to treatment plans and clinicians' notes[5,10]. Over the past decade, EHR adoption has become nearly universal in many countries. For example, as of 2021 about 96% of U.S. hospitals and 78% of office-based physicians were using certified EHR systems. This widespread digitization has led to an explosion of healthcare data: hospitals generate an average of 50 petabytes of data per year, yet as much as 97% of that information goes unanalyzed or unused. Healthcare data now makes up roughly 30% of the world's data volume, highlighting both a challenge and an opportunity.

The opportunity lies in extracting actionable insights from this vast trove of EHR data to improve patient care. This is where advanced Artificial Intelligence (AI) techniques like Deep Learning (DL) and Reinforcement Learning (RL) come into play. These methods can sift through complex, high-dimensional health data to detect patterns or optimize decisions in ways that traditional analysis might miss. In the context of health informatics, integrating AI with EHR systems means moving beyond using EHRs as mere digital filing cabinets toward using them as intelligent assistants that support clinical decision-making. By applying deep learning and reinforcement learning to EHR data, healthcare providers aim to improve patient care (through earlier diagnoses and risk predictions), optimize treatment recommendations (through personalized insights), and support clinicians' decisions with data-driven evidence. This paper provides an overview of these AI techniques and exploreshow they are being used in real-world healthcare settings, the benefits they offer, and thechallenges to their adoption[2,7].

2.Overview of Deep Learning and Reinforcement Learning

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 14 April 2025 doi:10.20944/preprints202504.1076.v1

Deep Learning (DL) is a subset of machine learning that uses multi-layered artificial neural networks to automatically learn representations and patterns from data. These models, often inspired by the structure of the human brain, excel at handling large volumes of complex data. Deep learning has driven breakthroughs in image recognition, natural language processing, and many other areas – and it is increasingly applied to healthcare data. Unlike simpler algorithms, deep learning models can learn features on their own from raw data, making them powerful for pattern recognition in EHRs (which can include a mix of structured data like vitals and lab tests, and unstructured data like free-text notes). For example, a deep learning system might comb through years of electronic records to identify subtle combinations of symptoms and test results that predict the onset of a condition such as diabetes or heart failure. In one survey of EHR-based deep learning research, models were applied to tasks including information extraction from clinical text, patient represent tation learning, outcome prediction (such as mortality or readmission), phenotyping of diseases, and even de-identification of records. Deep learning's ability to find complex nonlinear relationships in data often allows it to make more accurate predictions than traditional statistical models, sometimes with less manual data pre-processing required. However, deep learning models typically require large labeled datasets for training and are often considered "black boxes" – the decision logic can be difficult to interpret without special techniques[13,15].

Reinforcement Learning (RL), on the other hand, is a paradigm of AI where an "agent" learns to make decisions by interacting with an environment and receiving feedback in the form of rewards or penalties. Rather than learning from a static dataset of examples, an RL agent learns by trial and error – it tries different actions and learns which sequences of actions yield the best outcomes. A classic example outside healthcare is training an RL agent to play a video game: the agent receives positive rewards for increasing its score and negative rewards for losing, eventually learning an optimal strategy. In the healthcare context, one can imagine the "environment" as the patient's health state and the medical setting, the "agent" as a recommendation system suggesting interventions, and the "reward" as patient outcomes (improvement could be a positive reward, deterioration a negative reward). Reinforcement learning is inherently well-suited to sequential decisionmaking problems – situations where each action (such as prescribing a medication or adjusting a ventilator setting) can influence future states and outcomes. Clinicians often face sequential decisions: for example, choosing a treatment, observing the patient's response, then deciding the next step. RL provides a formal framework to optimize these sequences of decisions. Importantly, modern RL often leverages deep learning as well (so-called deep reinforcement learning) to handle large state spaces; for instance, a deep neural network can approximate the value of each potential action in a complex patient state. This combination has enabled RL to tackle problems like optimal dosing strategies and dynamic treatment regimens in healthcare research, as we will discuss later[12]. To clarify the distinction and complementary roles of deep learning and reinforcement learning in healthcare, Table 1 compares their key features and use cases in the context of EHR data.

Table 1. Comparison of Deep Learning and Reinforcement Learning in EHR contexts, including their goals, example uses, and characteristics.

Aspect	**Deep Learning (DL) in EHR**	**Reinforcement Learning (RL) in EHR**
Primary Objective	Learn patterns or predictions from large datasets (usually offline).	Learn optimal sequential decisions or policies to maximize long-term patient outcomes.
Data Input	Historical EHR data (structured fields, time-series, clinical notes, etc.) with labels/targets for training.	Patient state data from EHR (vitals, labs, etc.) as the s¨tate¨; uses historical patient trajectories or simulators for training.

 doi:10.20944/preprints202504.1076.v1

Typical Output	Prediction or classification (e.g., risk score for a condition, outcome likelihood, disease subtype).	Treatment recommendation or ac- tion (e.g., adjust medication dose, order a test) and an evolving pol- icy for decisions.
Example Use Cases	Predicting hospital readmission or adverse events (risk scores), diag- nosing conditions from clinical im- ages or ECG signals, extracting in- sights from doctor's free-text notes using NLP, patient phenotyping (identifying latent groupings of pa- tients).	Recommending personalized med- ication dosing sequences (e.g., insulin adjustments), optimizing ICU decisions like ventilator set- tings or fluid therapy, alerting clin- icians when patient states head to- ward danger, scheduling adaptive treatment plans (e.g., chemother- apy timing).
Strengths	Excels at finding complex, non- obvious patterns in big data; can incorporate multimodal data (im- ages, text, numbers) for holistic predictions; often outperforms tra- ditional models in prediction accu- racy.	Handles decision processes over time, accounting for future conse- quences; personalizes treatment by learning from individual patient trajectories; improves resource al- location by suggesting optimal strategies.
Limitations	Requires large labeled datasets and significant computing power; opaque b¨lack box m¨odels difficult for clinicians to interpret; model performance can degrade with in- consistent or biased EHR data.	Difficult trial-and-error learning in real healthcare; relies on offline data or simulations; needs very high-quality, representative data to learn effective policies; recommen- dations may differ from standard practice, causing trust/adoption is- sues.

3. Integration of DL and RL in EHR Systems

Bringing deep learning and reinforcement learning into everyday clinical practice requires integrating these AI models with existing EHR systems and workflows. In practical terms, integration means that the outputs of AI models (predictions or recommendations) are made available to clinicians through the tools they already use – typically the HER interface – at the right time in the workflow. This turns the EHR from a passive information storage into an active **clinical decision support system**.

One approach to integration is embedding predictive models directly into the HER to provide real-time alerts or risk scores. For instance, many hospitals have begun adding warning systems for conditions like sepsis or patient deterioration. These systems continuously analyze EHR data in the background and alert clinicians if a patient's data matches patterns that, in the past, have preceded a serious event. A case in point is Duke Health's

Sepsis Watch program, which was deployed within their Epic EHR system as an early warning tool for sepsis. The model behind Sepsis Watch – a deep learning algorithm trained on 32 million data points from 42,000 encounters – runs automatically and its parameters can be adjusted per patient to personalize the predictions. The output is presented in the EHR as a risk indicator so that care teams

 doi:10.20944/preprints202504.1076.v1

are notified when a patient is likely developing sepsis. Integrating the model into the EHR ensures that clinicians don't have to consult a separate application; the alert appears in the same interface where they review vitals and order medications, which streamlines the response. This tight integration was guided by frameworks like the HIMSS EMR Adoption Model and helped Duke achieve a high accuracy (around 93%) in sepsis screening while dramatically reducing false alarms.

Integration isn't limited to critical care alerts. **Deep learning** models are also being used to automate routine, time-consuming tasks in the EHR. An example is **ambient clinical documentation**: using AI to listen to doctor-patient conversations (with patient consent) and transcribe or even structure the important details directly into the EHR. As Dr. Michael Sherling notes, ambient voice recognition integrated with the EHR can convert a natural exam room conversation into a structured clinical note, effectively acting as a smart scribe[9]. This not only saves physicians from hours of typing or dictating after hours, but also allows them to focus on the patient during the visit rather than on the computer screen. Major EHR vendors and tech companies have started offering such AI-driven documentation assistants embedded in EHR systems.

Reinforcement learning models, given their need to observe outcomes over time, are often implemented in a more behind-the-scenes fashion. Many RL applications in healthcare use **offline training** (learning from retrospective EHR data) and then provide decision support in a way that clinicians can accept[8,11]. For example, an RL model might be integrated as a decision support module that suggests an optimal treatment plan for a complex condition, but the clinician in the EHR interface sees it as a ranked list of recommended actions or a note in the order set (rather than the system acting autonomously). This keeps the human decision-maker in the loop. A Microsoft Research project demonstrated an RL-based system that works alongside the EHR to **identify high- risk patient states**; it learned from historical cases to recognize when a patient's vital signs and labs indicate they are entering a dangerous trajectory and would alert the clinician to reconsider certain treatments. Such a system can be thought of as an "AI second opinion" integrated into the EHR – it doesn't override clinician judgment but provides an additional safety net by flagging potentially harmful situations or suggesting alternatives[1,3].

Integration requires overcoming technical and organizational challenges. On the technical side, EHR systems must be able to interface with AI models, which involves data exchange (often through standardized protocols like HL7/FHIR) and ensuring fast performance for real-time predictions. On the human side, clinicians need to trust and understand the AI suggestions. This is why many integrations start with AI performing as an advisor rather than an actor. Effective training and change management are also key– clinicians should be educated on what an AI alert or score means and how to act on it. Despite these challenges, we are beginning to see a trend: EHR vendors are partnering with AI developers, and healthcare organizations are building **AI-friendly infrastructure** so that deploying a new model (whether for predicting patient risk or optimizing operations) can be done more seamlessly. The ultimate goal of integration is for AI-driven insights to becomea natural part of clinical workflows, augmenting healthcare professionals' capabilities without adding burden[4,14].

4. Real-World Applications and Case Studies

AI techniques like deep learning and reinforcement learning have moved from theory to practice in several areas of medicine, yielding promising results. Below, we explore how these technologies are being applied in real-world settings, with an emphasis on concrete examples that illustrate their impact on patient care and decision-making.

One notable success story is the deployment of **Sepsis Watch** at Duke University Hospital. Sepsis, a life-threatening response to infection, is notoriously hard to detect early – its initial signs (fever, rapid heart rate, etc.) can be subtle and easily mistaken for milder illnesses. Duke's approach to this problem was to harness deep learning on their HER data to create an early warning system. The Sepsis Watch model continuously monitors patients' vital signs, lab results, and other EHR inputs (86 variables in total) and updates an AI-computed risk score every 5 minutes for each patient in the emergency department. If the model's analysis suggests a patient is likely to develop sepsis, it

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 14 April 2025 doi:10.20944/preprints202504.1076.v1

alerts the care team through the EHR interface. This allows clinicians to initiate sepsis treatment protocols sooner, before the patient deteriorates. The results have been striking: after Sepsis Watch was introduced, Duke saw a **27% drop in sepsis-related mortality** among the targeted patient group. False alerts, which plagued their prior simple rule-based system (previously there were more than 5 false alarms for every true sepsis case, causing alert fatigue), were dramatically reduced by the more sophisticated model (Sepsis Watch cut false alarms by $\tilde{6}2\%$ while achieving over 90% screening accuracy). This case exemplifies how deep learning on EHR data can save lives by enabling earlier and more accurate detection of critical illness. It's also a powerful example of integrating AI into clinical workflow: nurses and doctors at Duke view the Sepsis Watch alerts as part of their routine EHR dashboard, and a multidisciplinary "rapid response" team is ready to act on those alerts with confirmed protocols[6].

Another compelling application is in the management of **intensive care unit (ICU) treatments and complex decision-making scenarios**. Researchers have applied reinforce- ment learning to large critical care databases (like the MIMIC ICU database) to derive optimal treatment strategies for conditions such as sepsis. One research project developed an AI agent sometimes referred to as an "**AI Clinician**" that learned to recommend the dosages of IV fluids and vasopressors for septic patients by analyzing retrospective HER data from thousands of ICU stays. The AI Clinician essentially tried to simulate what sequence of treatment decisions would lead to the best patient outcomes (in terms of survival), formulating a policy through offline RL. In validation tests, it was found that if clinicians had followed the AI's recommended actions, patient outcomes could have improved; indeed, *deviations from the AI-recommended strategy were associated with higher mortality*. In other words, this data-driven policy appeared to outperform the average human decision-making in hindsight. While this AI Clinician is not yet a deployed tool (it's a proof-of-concept from a study), it points toward a future where reinforcement learning could assist in **treatment recommendations for complex, evolving conditions** like sepsis.

It's important to note that clinicians wouldn't blindly follow such recommendations –but these suggestions could serve as a valuable second opinion or starting point when formulating a treatment plan, especially in fast-moving situations. Building on this concept, similar RL approaches have been studied for managing mechanical ventilation settings in ICU, sedation weaning, and other scenarios where optimal control is challenging[16].

Beyond critical care, AI is finding its way into many other domains of healthcare via EHR data:

- **Chronic Disease Management:** Deep learning models are being used to predict disease exacerbations. For example, in diabetes care, researchers have applied rein- forcement learning using patient EHR logs and even simulators to recommend insulin dose adjustments for type 1 diabetics. By learning from many patient histories, an RL agent can suggest when to increase or decrease insulin, aiming to keep blood glucose in target range. Such a system could potentially support endocrinologists in personalizing diabetes treatment plans (though again, usually tested in simulation or retrospectively due to safety constraints).
- **Oncology:** On the deep learning side, EHR data combined with genomic profiles have been used to predict how cancer patients might respond to treatments, helping oncologists choose the best therapy. Meanwhile, reinforcement learning has been explored to design **dynamic treatment regimens** for chemotherapy – essentially learning the optimal timing and dosing adjustments across multiple cycles of chemo to maximize efficacy and minimize side effects. Early studies in this vein have shown that AI can propose adaptive plans that potentially improve outcomes for diseases like breast cancer or leukemia, though these are largely in research or clinical trial phases[5].
- **Mental Health and Preventive Care:** EHRs often contain rich longitudinal data that can be mined for patterns not obvious to human providers. AI tools have been developed to help in mental health care by analyzing therapy notes, prescription records, and patient-reported outcomes. For instance, deep learning algorithms can summarize a patient's mental health history and flag subtle signs of worsening depression or risk of self-harm,

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 14 April 2025 doi:10.20944/preprints202504.1076.v1

prompting proactive intervention. As one mental health expert noted, AI can "digest vast amounts of patient information and then detect patterns and summarize treatment to date," even suggesting treatment plans or alerting providers to potential crises before they occur. This kind of application shows that AI isn't limited to acute hospital settings; it can also support outpatient and preventive care by ensuring important signals in the data are not overlooked.

- **Diagnostic support:** While not purely EHR (since imaging data is involved), there are AI systems that integrate imaging results with EHR context. For example, a deep learning model might analyze radiology images for signs of disease and combine that with the patient's EHR (history, labs) to give a more comprehensive diagnostic report. Some advanced EHR systems are starting to embed these AI diagnostic aids so that when a clinician pulls up an X-ray or pathology report within the EHR, they also see AI-generated annotations (like outlining a suspected tumor) and risk scores based on similar patients' outcomes.

Each of these examples demonstrates a facet of how deep learning and reinforcement learning are improving or could improve patient care. The **common thread** is that these AI models leverage the vast memory of the healthcare system (the EHR data) to provide guid- ance that is tailored to individual patients. They act as clinical companions: screening for risks, advising on treatments, and learning from each patient interaction. In settings where these tools have been implemented thoughtfully, we have seen measurable improvements – whether it's reduced mortality, faster interventions, or simply more efficient use of clini- cians' time. That said, these successes also illuminate the challenges and responsibilities that come with AI in healthcare, which we explore next.

5. Benefits and Limitations

Applying deep learning and reinforcement learning to EHR data offers significant potential benefits for healthcare, but it also comes with important limitations and challenges. Understanding both is crucial for healthcare professionals and policymakers as they consider adopting these technologies.

Benefits :

- **Improved Patient Outcomes:** AI can analyze complex patterns in patient data, leading to earlier and more accurate diagnoses or warnings. For example, the use of a deep learning sepsis model at Duke was associated with a 27–31% reduction in sepsis mortality, illustrating how outcomes can improve when at-risk patients are identified and treated sooner. Similarly, model-guided recommendations in research studies (like the RL-driven "AI Clinician") hinted at lower mortality when its advice was followed.
- **Personalized Treatment:** Deep learning can subgroup patients and predict which treat- ments might work best for which patients, moving towards personalized medicine. Reinforcement learning goes a step further by personalizing *sequences* of decisions –for instance, tailoring medication dosing over time to an individual's response. This could mean more effective, patient-specific care plans, especially for complex chronic diseases.
- **Decision Support and Reduced Cognitive Burden:** AI acts as a safety net or second pair of eyes for clinicians. It can continuously monitor data and alert staff to issues that humans might miss (as in the sepsis alert example). By providing evidence-based suggestions (e.g., a list of likely diagnoses or a recommended next step in therapy), AI decision support can help clinicians make informed choices faster and with more confidence. This is especially valuable in high-pressure environments like ICUs or emergency rooms.
- **Efficiency and Workflow Optimization:** Many AI integrations aim to streamline work- flows. As described earlier, AI can automate documentation (transcribing encounters, coding billing information), which *reduces administrative burdens* on clinicians. AI predictive models can also help with operational decisions – for example, predicting which patients are likely to be no-shows or to need ICU transfer, allowing staff to allocate resources more

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 14 April 2025 doi:10.20944/preprints202504.1076.v1

efficiently. Overall, this means more time for clinicians to focus on direct patient care and less time on clerical tasks.

- **Population Health Insights:** By aggregating and analyzing data from thousands or millions of EHRs, deep learning can identify public health trends or risk factors not apparent at the individual level. Health systems can use these insights for preventive health measures, such as identifying neighborhoods with rising diabetes risk and directing community interventions there.

6. Conclusion and Future Directions

Deep learning and reinforcement learning are revolutionizing how we extract value from electronic health records. In this paper, we explored how deep learning models can comb through massive EHR datasets to predict outcomes or find hidden patterns, and how reinforcement learning can model clinical decision-making to optimize treatment sequences. Real-world examples – from early warning systems like Duke's Sepsis Watch to research prototypes like AI Clinician – illustrate that these approaches are not justacademic exercises but tangible tools with the potential to **enhance patient care andclinical workflows**. They have already demonstrated benefits such as earlier interventions,more personalized therapies, and support for overburdened healthcare staff, confirmingthat AI (when used wisely) can serve as a powerful ally to clinicians. Looking to the future, we expect the presence of AI in EHR systems to grow. **Healthcare professionals and policymakers** can anticipate a continued expansion of AI-driven decision support across various medical specialties. For example, large-scale studies in- dicate that implementing predictive models in clinical settings often leads to improved outcomes in the majority of cases studied. In primary care, we might see deep learning models that help flag patients who would benefit from screening or lifestyle interventions earlier than current guidelines suggest. In hospitals, reinforcement learning could be used in "digital twin" simulations of patients – allowing care teams to test how different interventions might play out for a given patient virtually, before trying them in reality.

To realize these possibilities, certain **future directions and developments** are important:

- **Improving Explainability:** Advancements in making AI models more transparent will be crucial. Techniques that can point to which factors in the EHR most influenced a prediction (for instance, highlighting that a rising lactate level and heart rate were key to a sepsis alert) will help clinicians trust and act on AI insights. Future EHR-integrated AI might come with built-in explanation panels next to an alert or recommendation. **Robustness and Generalization:** Efforts are underway to make models more robust to shifts in data. This includes using diverse training data (possibly from multiple hospitals) so the models generalize better, and deploying continuous learning systems that update models as new data comes in – carefully and under monitoring. Federatedlearning (where models learn from data across institutions without the data leaving each site) is an emerging technique that could allow AI to improve using multi-hospital data while preserving privacy.
- **Hybrid Human-AI Decision Making:** Rather than viewing AI as an autonomous decision-maker, the future likely lies in **collaborative systems**. An AI might handle routine tasks and provide recommendations, while clinicians make the nuanced judgments and ethical considerations that AI cannot. EHR interfaces might evolve to present a "conversation" with the AI, where clinicians can query the system ("Why are you suggesting this medication?") and provide feedback ("This suggestion isn't appropriate because...") that the AI system can learn from. In reinforcement learning terms, clinicians could act as an overseer or safety layer on top of the AI's policy, ensuring no action is taken without human approval.
- **Addressing Ethical and Regulatory Frameworks:** Policymakers will play a role in shaping guidelines for AI in health. We expect clearer standards on validating AI tools (e.g., requiring clinical trials or real-world performance monitoring for certain high-stakes AI algorithms) and on liability (clarifying how responsibility is shared when clinicians use AI

 doi:10.20944/preprints202504.1076.v1

advice). As these frameworks mature, healthcare institutions may feel more confident in adopting AI solutions knowing there is guidance on how to do so safely and ethically.

- **Education and Training:** Future clinicians and health administrators will need a basic understanding of AI. Medical schools and continuing education programs are beginning to introduce informatics and AI literacy into their curriculum. This trend will continue so that by, say, 2030, it may be routine for a physician to interpret an AI risk score much like they interpret lab results today, understanding its limitations and proper use.

In conclusion, the integration of deep learning and reinforcement learning with elec- tronic health records is a transformative development in health informatics. We are still in the early stages of this transformation. The current trajectory – marked by promising pilot projects, increasing investments, and a growing body of validation research – suggests that AI will become an indispensable component of healthcare delivery. The vision for the future is a healthcare system where **AI-driven tools assist at every level**: from suggesting preventive care for a single patient, to managing hospital resources, to informing public health decisions for an entire population. Achieving this vision will require continued collaboration between clinicians, data scientists, and policymakers to ensure that these technologies are **effective, equitable, and above all, centered on the well-being of patients**. With careful stewardship, AI-powered EHR systems could usher in an era of smarter, more proactive, and more personalized healthcare for all.

References

1. Shrestha, D. Advanced Machine Learning Techniques for Predicting Heart Disease: A Comparative Analysis Using the Cleveland Heart Disease Dataset. *Appl. Med. Inform.* **2024**, *46*(3).
2. Fleuren, L.M.; et al. Challenges in clinical implementation of AI models for sepsis. *Lancet Respir. Med.* **2020**, *8*, 1161–1163.
3. Shrestha, D. Comparative Analysis of Machine Learning Algorithms for Heart Disease Prediction Using the Cleveland Heart Disease Dataset. *Preprints* **2024**.
4. Shrestha, D.; Nepal, P.; Gautam, P.; Oli, P. Human Pose Estimation for Yoga Using VGG-19 and COCO Dataset: Development and Implementation of a Mobile Application. *Int. Res. J. Eng. Technol.* **2024**, *11*(8), 355–362.
5. Valade, J. Duke's Quantum Leap in AI: Sepsis Watch at Duke. *Duke Today* **2024**. Available online: https://today.duke.edu/2024/03/sepsis-watch-duke (accessed on 12 April 2025).
6. Shrestha, D.; Nepal, P.; Gautam, P.; Oli, P. Real-Time Animal Monitoring System Using Pulse, Temperature, and GPS Sensor. **2024**.
7. Greene, L. How Healthcare Data Technology is Leveraged by Leaders. *Arcadia.io* **2023**. Available online: https://arcadia.io/blog/healthcare-data-technology-leaders (accessed on 12 April 2025).
8. Office of the National Coordinator for Health IT. National Trends in Hospital and Physician Adoption of EHRs. **2022**. Available online: https://dashboard.healthit.gov (accessed on 12 April 2025).
9. Miotto, R.; Wang, F.; Wang, S.; Jiang, X.; Dudley, J.T. Deep Learning for Healthcare: Review, Opportunities and Challenges. *Brief. Bioinform.* **2018**, *19*(6), 1236–1246.
10. Fatemi, M.; Jain, R.; Stojanovic, J.; Shah, N. Using Reinforcement Learning to Identify High-Risk States and Treatments in Healthcare. *Microsoft Research Blog* **2022**. Available online: https://www.microsoft.com/en-us/research/blog/using-reinforcement-learning-to-identify-high-risk-states-and-treatments-in-healthcare (accessed on 12 April 2025).
11. Sakhuja, A.; et al. A Primer on Reinforcement Learning in Medicine for Clinicians. *J. Clin. Med.* **2024**.
12. Miller, G. Benefits and Challenges of Integrating AI and ML into EHR Systems. *Healthc. IT Today* **2024**.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 14 April 2025 doi:10.20944/preprints202504.1076.v1

Available online: https://www.healthcareittoday.com/ai-ml-ehr-integration (accessed on 12 April 2025).

13. HIMSS News. North Carolina Hospital System Reduces Sepsis Cases Using Predictive Analytics. **2024**. Available online: https://www.himss.org/news/duke-health-uses-ai-fight-sepsis (accessed on 12 April 2025).
14. Shrestha, D.; Valles, D. Evolving Autonomous Navigation: A NEAT Approach for Firefighting Rover Operations in Dynamic Environments. In Proceedings of the 2024 IEEE Int. Conf. Electro Inf. Technol. (EIT), Chicago, IL, USA, 15–17 May 2024; pp. 247–255.
15. Komorowski, M.; Celi, L.A.; Badawi, O.; Gordon, A.C.; Faisal, A.A. The Artificial Intelligence Clinician Learns Optimal Treatment Strategies for Sepsis in Intensive Care. *Nat. Med.* **2018**, *24*, 1716–1720.
16. Shrestha, D.; Valles, D. Reinforced NEAT Algorithms for Autonomous Rover Navigation in Multi-Room Dynamic Scenario. *Fire* **2025**, *8*(2), 41.
17. Shickel, B.; Tighe, P.J.; Bihorac, A.; Rashidi, P. Deep EHR: A Survey of Recent Advances in Deep Learning Techniques for Electronic Health Record (EHR) Analysis. *IEEE J. Biomed. Health Inform.* **2017**, *22*(5), 1589–1604.
18. Shickel, B.; Loftus, T.J.; Ruppert, M.; Upchurch, G.R.; Ozrazgat-Baslanti, T. Dynamic Predictions of Postoperative Complications from Explainable, Uncertainty-Aware, and Multi-task Deep Neural Networks. *Sci. Rep.* **2023**, *13*, 12345.