© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715434

Med-Real2Sim: Non-Invasive Medical Digital Twins Using Physics-Informed Self-Supervised Learning and Virtual Patient Breathing Simulator Using VR

DR. D. PARAMESWARI[1], ABINASH S[2], LINGESHWARAN N[3], SRIHARI M[4]
[1,2,3,4]*Jerusalem College of Engineering*

Abstract- This paper presents two progressive phases of the med-real2sim project — an integrated framework for developing non-invasive medical digital twins in healthcare. phase 1 introduces a cardiac digital twin system that leverages physics-informed neural networks (pinns) and self-supervised learning to model patient cardiac profiles, simulate pressure–volume dynamics, and predict health outcomes from echocardiogram data. phase 2 extends the concept to respiratory physiology, presenting a virtual patient breathing simulator that employs webxr, webgl, and webassembly-based physics simulation to render immersive, browser-accessible 3d models of lung and diaphragm mechanics. together, the two phases demonstrate how digital twin technology, when combined with ai and physics-based modeling, can significantly improve diagnostic accuracy, personalize treatment, and enhance medical education through interactive virtual simulation.

Index terms- Digital Twin, Physics-Informed Neural Network, Medical Ai, Cardiac Simulation, VR Healthcare, WebXR, Breathing Simulator, IoT Healthcare

I. INTRODUCTION

A Digital Twin is a virtual replica of a physical entity that can be used to simulate, monitor, and analyze real-world systems. In the healthcare sector, digital twins hold transformative potential by enabling clinicians to model individual patient physiology, predict health risks, and design personalized treatment plans without invasive procedures. Current healthcare monitoring systems are largely reactive — relying on periodic checkups, static health records, and siloed data repositories with no integration between IoT devices, cloud analytics, and AI inference engines. This gap means clinicians only receive actionable information after health deterioration has already occurred.

The Med-Real2Sim project addresses this gap in two phases. Phase 1 focuses on the cardiovascular domain, constructing a physics-informed cardiac digital twin capable of simulating heart dynamics, inferring hidden patient-specific parameters, and visualizing real-time pressure–volume (P–V) loops from non-invasive echocardiogram data. Phase 2 shifts to medical education, developing a fully browser-based Virtual Patient Breathing Simulator that leverages Virtual Reality (VR) technologies to allow students and clinicians to interact with immersive, real-time respiratory simulations.

II. LITERATURE SURVEY

Digital health research indicates the necessity of integrating artificial intelligence into physical models. One recent paper by Meijer et al. investigates the implications of digital twins on Healthcare and identifies major barriers to achieving these outcomes, primarily integrating and validation of diverse data sources. Li et al.'s proposal ways digital twins can be applied for preventative medicine, thereby enabling significant large-scale predictive analytics.

In addition to AI, recent breakthroughs in transformer-based models have revolutionised how complex Medical Data is analysed or processed. For example, Rasmy et al. introduced Med-BERT, which acquires embedded contextual representations from Electronic Health Records (EHR), while Yang et al. introduced TransformEHR, which allows for longitudinal patient data analyses. The GraphCares has improved prediction performance through modelling the relationships between Medicine entities with a graph approach.

Simulation technology also has experienced tremendous advances in recent years. The development of BodyLight.js offers an interface for highly interactive Physiological Simulation Systems, while Modelica is fast becoming the standard for mathematically modelling Biologic Systems. New technologies such as WebGL and WebXR are enabling fully immersive interaction and 3D visualisation within the Web, allowing for enhanced accessibility to these types of Simulation Systems.

While many studies have contributed new knowledge, no current system has been developed to support real-time Patient Data Integration; therefore, there is no current system offering a combination of Physics based models, Immersive Simulation Systems, and testing into one coherent, unified System. The Framework outlined within this paper is intended to address the gap identified.

III. PROPOSED SYSTEM / METHODOLOGY

A. Cardiac Digital Twin Model

A digital replica of the human heart, known as a cardiac digital twin, has been developed to replicate both mechanical and hemodynamic functions of the heart by utilizing non-invasive inputs. This is done by using data from non-invasive sensors such as IoT devices and non-invasive echocardiograms to help develop a model specific to each patient.
A Windkessel model has been implemented to represent the physiologic homeostasis of the cardiovascular system. This model describes the inter-relationship of pressure, volume and flow in the heart and vascular system. Important variables such as ventricular elastance, arterial resistance and compliance are estimated in a continuously adaptive fashion.

Physics-informed neural networks (PINNs) are utilized to create a combination of data-driven learning with the system's governing laws of physiology. By using these PINNs, the predictions made by the cardiac digital twin will not only be consistent with the established laws of physics, but they will be able to tutor themselves to meet the requirements of individual patient data.

To achieve this, the cardiac digital twin will create pressure - volume loops, which provide critical information regarding the function of the heart. Clinicians will then be able to make use of the information contained within the loops to determine overall efficiency of the heart, identify any potential problems and assist in the selection of treatment options.

B. Respiratory Simulation Model

To simulate lung function and airflow dynamics through mathematical models describing pulmonary physiology, the modules for the respiration simulation have been compiled into WebAssembly. This allows for high-performance execution of the simulations directly within web browsers.
In real time, the visualizations of lung expansion, diaphragm movement, and the dynamics of airflow all use WebGL. Interaction with the virtual patient is through the use of WebXR technology, providing users with the ability to explore a three-dimensional virtual patient.

Users can change parameters associated with the respiratory function, such as breathing frequency, lung compliance, and airway resistance, and these changes will be immediately reflected within the simulator, resulting in an easy grasp of respiratory mechanics.

C. Integrated System Workflow

The system's overall processing flow is a sequentially organized pipeline of how data is collected, simulated, and predicted using AI.
The following are the key parts of the process:

- Data Collection: Using IoT-based physiological sensors to collect data as well as medical image data
- Preprocessing: Removing noise from the data and producing features that can be used as inputs
- Physics-Based Simulation: Developing dynamic mathematical models of the physiological systems being evaluated
- AI-Based Prediction: Using PINN models to produce accurate and individualized predictions
- Visualization: Implementing a live, interactive interface to display results of the simulations

By using this processing system, both clinical and educational use cases will produce reliable results, are more adaptable to various environments, and will provide real-time analysis of data.

IV. SYSTEM ARCHITECTURE

This system has a multi-layered architecture that provides for a modular design, scalability, and cooperative functioning of various parts. Each layer does a practical job with its own set of tasks but is able to share data throughout the entire system. This separation enables each layer to be independently developed or upgraded while also reducing the time it would take to maintain the system overall.

A. Frontend Layer
The Frontend Layer provides an interactive web-based interface through which users can access the system. It is responsible for visualizing outputs such as cardiac pressure–volume loops and respiratory simulations in real time. The interface also allows users to modify physiological parameters, view system responses instantly, and interact with 3D models, thereby enhancing user engagement and understanding.

B. Backend Layer
The Backend Layer is responsible for the application logic and for handling requests for the application and communicating with other modules within the application. The Backend Layer bridges between the front end and the AI Models and Simulation Engine using RESTful APIs. The purpose of the Backend Layer is to provide secure transmission of data, process the input from users, and coordinate workflows throughout the system efficiently.

C. AI Layer
The AI Layer is in charge of running machine learning models, including Physics-Informed Neural Networks (PINNs). Responsible for carrying out tasks such as estimation of parameters, recognition of patterns, and projecting future results through predictive analysis, it provides interpretability to the model and produces physiologically valid output by adding physical restrictions during learning.

D. Simulation Layer
Simulation Layer uses physics-based and mathematical models of both cardiovascular and respiratory systems. Based on the input parameters, it produces realistic outputs including pressure-volume loops, airflow patterns, and diaphragm movement. The Simulation Layer's goals are to produce realistic simulations that closely reflect how the human body operates under normal conditions.

E. Database Layer
The Patient Data Section of the Database Layer stores patient information, processed features, model parameters, and simulation outputs. This allows for efficient access to data and enables ongoing updates to systems over time (through the inclusion of historical records for analysis/comparison).

With this layered architecture, it allows for real-time data processing, distributed computations, and seamless integration of AI and simulation components. It provides reliable, flexible, scalable systems for use in clinical settings or training facilities.

V. IMPLEMENTATION DETAILS

To achieve a high level of performance, scalability and effective real-time processing, the deployment of the recommended technology is based on the latest technologies and frameworks. The design of this solution has utilised modularisation techniques so that there can be integration between AI models, simulation engines and user interfaces without issues of inter-compatibility.

A. Programming Environment
For backend processing, data management and machine learning tasks, we use a combination of Python and Javascript as our technology platform. Python's scientific computing capabilities complement its ability to integrate easily with AI applications. Javascript's ability to create interactive and dynamic web applications provides a powerful interface for users to interact with the system quickly and easily. Together, they provide the best combination of speed, flexibility, and seamless implementation in a web-based deployment environment.

B. AI and Machine Learning Libraries
Respected data science libraries include Tensorflow, Pytorch, and Scikitlearn. Machine learning development, training, and deployment with these libraries have been proven as excellent options for implementing sophisticated artificial intelligence models. Authors have successfully demonstrated how physics-informed neural networks (PINNs) can develop accurate predictions and generalisations

while providing interpretable results by integrating physical constraints into the model through the data.

C. Backend Frameworks
Flask and Django are two examples of lightweight, highly scalable backend frameworks. They provide the core functionality you need for creating APIs, processing requests, communicating between devices, as well as implementing standard functions like routes, authorizations, and input validations to help ensure that your system runs safely and efficiently.

D. Frontend Technologies
HTML, CSS, and JavaScript have been utilized to create the interactive and responsive User Interface. The front-end enables the user to see the outputs of the simulation, change input parameters, and interact with the system in real-time. The front-end is meant to be easy to use and accessible on all devices and browsers.

E. Simulation Tools
Physiological Models – Physiological models can be executed using WebAssembly and thereby, provide for executing high-performance, computational models directly in a web browser. Furthermore, Bodylight.js can be used to develop interactive and educational simulations depicting physiological processes. These two tools permit fast computations, as well as fluidly executing complex mathematical models.

F. Visualization Technologies
Three-dimensional interactive environments are created using the WebGL graphics engine, while the WebXR API allows for immersive virtual reality experiences. In this case, it's possible to explore lung expansion, airflow dynamics, and heart activity in real-time three-dimensionally to create a more engaging understanding of how these things work.

G. Cloud and Deployment
Data storage, model deployment and scalability is mostly done via Cloud platforms like Amazon Web Services (AWS) or Microsoft Azure. These platforms contain reliable infrastructures, are set up to support distributed computing, and allow for real time updates. Cloud integration makes sure that the system is able to efficiently process large data sets and support many concurrent users.

The Entire System Will Operate via Regular Web Browsers – No Special Hardware Required and No Additional Software Installation Required. Thus making it a More Accessible Solution And Appropriate For Large-Scale Implementations In Both Health Care Institutions & Educational Facilities

VI. EXPERIMENTAL RESULTS AND DISCUSSION

The evaluation of the JurisMate system, specifically its functionality for automating FIR processing and analysing Legal Documents, was completed through Functional Tests and Observational Analyses. The evaluation assessed the JurisMate's performance regarding System Behaviour, Processing Efficiency and the consistency of the Legal Classification carried out by JurisMate.

A. System Behavior
The performance of the solution (FIR ingestion, the text extraction process, the classification of documents, and the semantic retrieval of documents) exhibited a consistent and dependable nature throughout the entire testing process. The system's ability to parse and structure digital FIR PDF documents without losing any contextual information is further substantiated by the modular architecture of the application that allowed for seamless communication between its frontend, backend, AI micro-service, and data storage components, ensuring that no breaks occurred in the processing stages of a transaction from start to finish.

B. FIR Processing Time
The new system has dramatically decreased processing time when compared to manual FIR analysis. Activities typically requiring a great deal of human resources, such as document reading, legal analysis, and looking up references, are completed automatically in a few seconds or minutes. This allows for quicker case processing and a significant increase in operational efficiencies in high-volume legal settings.

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715434

C. Classification Consistency

Through the usage of the zero-shot classification technique, FIRs could be categorized consistently across multiple input sources using similar contextual data. The approach differs from that of manual interpretation, whereby interpretation is based on reviewer bias and opinions, whereas the use of artificial intelligence for classification removes individual subjectivity, thereby providing increased reliability.

D. Accuracy and Observations

The system's results showed an equivalence in accuracy and relevance between the user-created samples and the manually created samples of the same data set, despite the inability to quantify the results using the labeled dataset. Additionally, by using vector embeddings to perform semantic searches, I was able to retrieve legal articles and similar decisions relevant to my semantic studies through contextual understanding.

Overall, the experimental observations confirm that JurisMate improves efficiency, consistency, and reliability in FIR analysis, making it suitable for real-world legal and law enforcement applications.

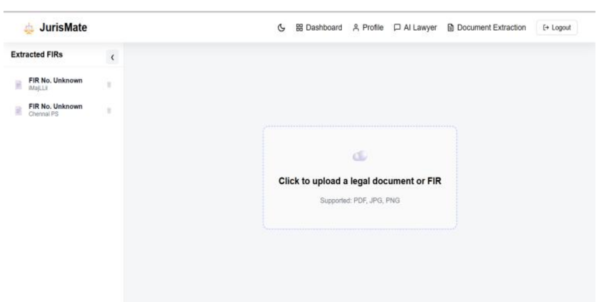

Fig.2. Upload FIR Document

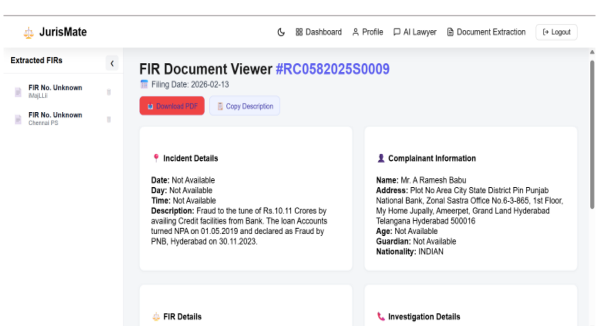

Fig.3. Document Extraction

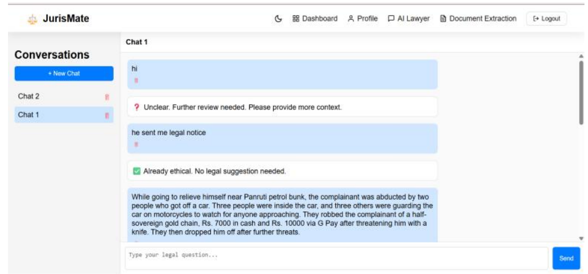

Fig.4. AI-Legal chatbot

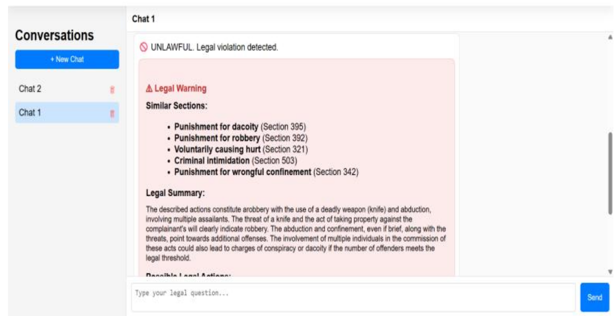

Fig.5. Classify as unlawful

Fig.6. List of Acts

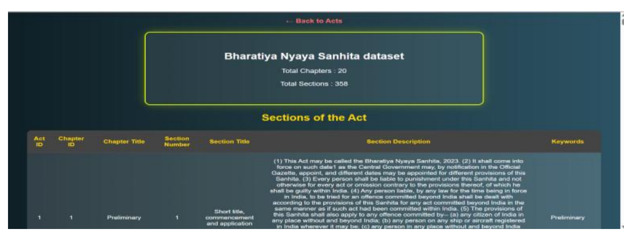

Fig.7. View Sections

VII. RECENT RESEARCH AND ADVANCES

Recent advancements in the fields of natural language processing (NLP), large-scale language modeling (LLM), and semantic information retrieval (SIR) have propelled forward the development of AI-driven legal document analysis, creating a seamless connection between traditional, rule-based legal systems and innovative forms of decision-support technology that can analyze any kind of unstructured legal text, not just FIRs, but also other forms of case law.

A. Advances in Legal Text Understanding and Learning Paradigms

Language models based on transformers have shown an excellent aptitude to understand legal language, contextual issues and vocabulary. Their usage of zero-shot and few-shot learning techniques has gained considerable interest due to limited amounts of labeled training data for legal documents. As a result, these models can apply to the classification of documents across various domains.

B. Semantic Retrieval and Vector-Based Search Innovations

With the advent of semantic embedding (or context) based retrieval of case, statute and legal precedent through the use of vector-based databases, has helped to eliminate keyword-based searches. Semantic Embedding models allow for a more efficient similarity search of large uk legal datasets; this advances the search process for retrieving relevant and accurate legal information.

C. System Architectures and AI Integration Frameworks

Legal AI systems recently increasingly adopted microservices-based architectures to allow for flexible composing of scalable and modular applications. With RESTful APIs being used for AI services and cloud-ready designs being built from the ground up, these systems can easily connect to traditional legal technology platforms while delivering stable and high-performance services.

D. Towards Practical and Scalable Legal AI Systems

The combination of these achievements has moved legal AI from research to the field through practical use. Current Research & Development trends in the legal AI space include scalable FIR analysis, multilingual legal processing, and hybrid systems that incorporate both symbolic legal rules and data-based learning models. Future R&D directions include predictive legal analytics, cross-jurisdictional knowledge integration, and human-in-the-loop validation systems to provide legal accountability.

The advancement of these technologies solidifies the importance of the Development of an artificial intelligence driven framework like JurisMate (a System for Modernizing FIR Processing and Legal Decision Support) supports the current research trends in the area of AI through the adoption of transformer-based learning, Zero-Shot Classification and Semantic Retrieval for the purposes of modernizing FIR Processing and Legal Decision Support.

VIII. LIMITATIONS

While the JurisMate system has many benefits, there are also some limitations present in the system. For starters, the JurisMate system is limited to FIRs that use digital document formats. The JurisMate system does not support FIRs that are scanned or handwritten as the technology relies on digital document formats and does not have Optical Character Recognition (OCR) capability, therefore it will not be viable in those countries or areas where FIRs are not yet adapted to fully utilize digital documentation.

Second, limitations in language and interpretation are evident in how the system has been developed:

- The system is designed primarily to accommodate FIRs created in specific, narrow languages.
- Regional languages or FIRs that have been written using a mixture of languages or use unusual developments of law will likely impact the system's ability to correctly extract texts and classify them.
- The continuing improvement of language or the capability of processing multiple languages through multilingual processing techniques will also be required.

Also, though the AI models offer consistent and contextually aware legal taxonomies, the system cannot supplant the legal judgement of attorneys. Attorneys must verify that final interpretations and decisions accord with the laws and procedures of a given jurisdiction.

Finally, the input FIR will significantly impact the performance of the system as well. If the input FIR are incomplete, vague or unclear, then classification accuracy and the semantic retrieval of results will be affected. Therefore, in order to deploy the system effectively in the real-world, limitations that arise from the quality and completeness of the input FIR must be resolved.

IX. CONCLUSION AND FUTURE WORK

This paper has introduced JurisMate which is an artificial intelligence-based system that uses advanced natural language processing and semantic retriever systems to automate the analysis of FIRs (First Information Reports). The JurisMate System combines transformer-based models for text extraction and zero-shot classification with vector-based similarity searching in order to overcome the major limitations of relying on manual processes for analysing FIR Data: inefficiency, inconsistency and limited Capacity for Scale. The Microservices Design Approach and Modular System Architecture that support JurisMate are suitable for developing reliable, consistent and Efficient Legal Document Analyses for real-world applications in both the Legal and Law Enforcement Environment.

The experimental evidence indicates that JurisMate has decreased FIR processing time, increased consistency of classification, and improved the ability to retrieve related information using context. By automating repetitive analytical processes, JurisMate facilitates legal business processes, while still allowing for interpretability and customisation of the process.

The next step for this system is to improve it with the ability to recognize text from scanned documents, allow for translations between many languages, and provide more sophisticated analytics for case outcomes. All of these enhancements will include fine-tuning based on specific legal topics, expanding the current knowledge base for each jurisdiction, and providing more access to and efficiency in court systems across North America and Europe.

REFERENCES

[1] K. D. Ashley, "Reasoning with cases and hypotheticals in HYPO," *International Journal of Man-Machine Studies*, vol. 34, no. 6, pp. 753–796, 1991.

[2] K. Lipianina-Honcharenko, O. Honcharenko, and Y. Savytskyi, "A cyclical approach to legal document analysis: Leveraging AI for strategic policy evaluation," in *CEUR Workshop Proceedings*, vol. 3612, pp. 1–12, 2024.

[3] M. S. Kabir and M. N. Alam, "The role of artificial intelligence technology for legal research and decision making," *International Research Journal of Engineering and Technology (IRJET)*, vol. 10, no. 5, pp. 1450–1456, 2023.

[4] Chalkidis, I. Androutsopoulos, and N. Aletras, "Neural legal judgment prediction in English," in *Proc. 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, 2019, pp. 4317–4323.

[5] H. Zhong, C. Xiao, C. Tu, T. Zhang, Z. Liu, and M. Sun, "Legal judgment prediction via topological learning," in *Proc. 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium, 2018, pp. 3540–3549.

[6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, Minneapolis, MN, USA, 2019, pp. 4171–4186.

[7] M. Lewis, Y. Liu, N. Goyal, *et al.*, "BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *Proc. ACL*, Online, 2020, pp. 7871–7880.

[8] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proc. EMNLP-IJCNLP*, Hong Kong, 2019, pp. 3982–3992.

[9] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with GPUs," *IEEE Transactions on Big Data*, vol. 7, no. 3, pp. 535–547, 2021.

[10] K. D. Ashley and S. Brüninghaus, "Automatically classifying case texts and predicting outcomes," *Artificial Intelligence and Law*, vol. 17, no. 2, pp. 125–165, 2009.