Explainable Clinical Decision Support Systems for Post-COVID Care Pathways

Daniel Harper
University of Gloucestershire, United Kingdom

Sophie Langford
University of Gloucestershire, United Kingdom

Submitted on: August 22, 2021
Accepted on: November 27, 2021
Published on: December 18, 2021

DOI: 10.5281/zenodo.18071027

***Abstract*—Clinical decision support systems have become a foundational component of modern healthcare delivery, particularly in contexts where patient trajectories are complex, uncertain, and long running. Post COVID care represents such a setting, characterised by heterogeneous symptoms, fluctuating recovery patterns, and multi organ involvement. While machine learning driven decision support systems demonstrate strong predictive capability, their limited transparency poses barriers to clinical trust, accountability, and safe adoption. This article presents a comprehensive framework for explainable clinical decision support tailored to post COVID care pathways. The proposed approach integrates interpretable predictive modelling, uncertainty aware inference, and clinician oriented explanation layers to support decision making across diagnosis, monitoring, and care planning. Through systematic architectural design and empirical evaluation, the study demonstrates how explainability can be embedded as a core system property rather than an afterthought, enabling reliable and actionable clinical insights in complex care environments.**

***Index Terms*—Explainable artificial intelligence, clinical decision support systems, post COVID care, healthcare analytics, interpretable machine learning**

I. Introduction

Clinical decision support systems have long been positioned as a mechanism to improve consistency, safety, and efficiency in healthcare delivery. Early systems relied on static rules and clinical guidelines, offering deterministic recommendations within narrowly defined scenarios. As healthcare data has expanded in volume and complexity, these systems have evolved toward data driven and learning based approaches that can model non linear relationships across diverse patient populations. This evolution has accelerated the adoption of machine learning within clinical decision support, enabling predictive and prescriptive capabilities that extend beyond traditional rule based reasoning.

Post COVID care pathways introduce a uniquely challenging decision environment. Patients present with wide ranging symptoms that span respiratory, cardiovascular, neurological, and psychological domains. Recovery trajectories are often non linear, with periods of improvement followed by relapse or delayed complications. Clinical teams must make decisions under uncertainty, balancing short term symptom management with long term functional outcomes. In such settings, purely automated or opaque decision support systems are insufficient. Clinicians require systems that not only produce accurate predictions, but also communicate the reasoning, limitations, and confidence associated with those predictions.

The growing reliance on machine learning in healthcare has amplified concerns around transparency, accountability, and bias. Black box models, while performant, obscure the underlying rationale behind recommendations, making it difficult for clinicians to assess appropriateness, challenge outputs, or explain decisions to patients. This lack of interpretability is particularly problematic in post COVID care, where clinical judgement must integrate evolving evidence, patient context, and ethical considerations. Explainable clinical decision support systems seek to address this gap by making model behaviour understandable and actionable for human decision makers.

Explainability in healthcare decision support extends beyond technical model introspection. It encompasses the presentation of clinically meaningful features, the articulation of uncertainty, and the alignment of explanations with established medical reasoning. An effective explainable system should support cognitive processes rather than replace them, enabling clinicians to maintain agency while benefiting from computational insight. This human centred perspective is essential for fostering trust and ensuring responsible use of artificial intelligence in care delivery.

Despite growing recognition of the importance of explainability, many existing clinical decision support implementations treat it as an auxiliary feature rather than a foundational design principle. Post hoc explanation tools are often layered onto complex models without considering clinical workflow integration or interpretive validity. As a result, explanations may be technically accurate yet clinically unhelpful. There is a need for architectures that embed explainability across data ingestion, model selection, inference, and user interaction.

This article addresses that need by proposing an explainable clinical decision support framework specifically designed for post COVID care pathways. The framework integrates interpretable modelling techniques, uncertainty estimation, and multi level explanation mechanisms within a unified architecture. Rather than focusing solely on predictive accuracy, the approach prioritises decision transparency, robustness, and usability in real clinical contexts.

The contributions of this work are threefold. First, it articulates the unique requirements of post COVID clinical decision support from an explainability perspective. Second, it presents a system architecture that operationalises these requirements through modular and interpretable design. Third, it provides an empirical evaluation demonstrating how explainable decision support can enhance clinical understanding and confidence without compromising predictive performance.

The remainder of this article develops these contributions in detail. The next section reviews relevant literature across clinical decision support, explainable machine learning, uncertainty modelling, and healthcare analytics. This is followed by a description of the proposed methodology, including system architecture, modelling techniques, and explanation mechanisms. The results section presents quantitative and qualitative evaluations using simulated post COVID care scenarios. Finally, the discussion and future directions sections reflect on implications for clinical practice and outline pathways for further research.

II. Literature Review

Explainable clinical decision support systems for post COVID care are grounded in a broad interdisciplinary literature spanning healthcare decision support, predictive analytics, machine learning, uncertainty modelling, and ethical artificial intelligence. This section synthesises prior research to establish a comprehensive foundation for transparent, trustworthy, and clinically aligned decision support.

A. Clinical Decision Support Systems in Healthcare

Clinical decision support systems have traditionally aimed to enhance diagnostic accuracy, treatment consistency, interoperability, and patient safety by embedding medical knowledge into clinical workflows. Early systems relied on rule based reasoning and deterministic protocols, providing transparency but limited adaptability in complex clinical environments [1]–[4]. As healthcare information systems evolved, decision support increasingly incorporated analytical and predictive techniques to manage heterogeneous patient populations [5], [6].

Intelligent scheduling, monitoring, and optimisation systems demonstrated measurable improvements in efficiency and care coordination, yet raised concerns regarding interpretability and trust [7], [8]. These challenges motivated research into hybrid decision support architectures that balance adaptability with clinical accountability [9].

B. Machine Learning for Clinical Risk Prediction

Machine learning has become central to clinical risk prediction, prognosis estimation, and outcome forecasting. Deep learning techniques have achieved strong performance in medical imaging tasks such as disease detection, severity assessment, and lesion segmentation [10]–[13]. These models capture complex spatial patterns but often operate as opaque black boxes.

Temporal and sequential modelling approaches enable forecasting of patient trajectories and early detection of deterioration [14]–[16]. Early warning systems integrate physiological and laboratory data to identify at risk patients in acute and post acute settings [17], [18].

Despite improved predictive accuracy, limited interpretability remains a barrier to clinical adoption, particularly in complex and evolving care pathways.

C. Explainable Artificial Intelligence in Clinical Contexts

Explainable artificial intelligence has emerged as a critical response to the opacity of advanced machine learning models. In healthcare, explainability supports ethical responsibility, clinical accountability, and regulatory alignment. Interpretable models have been shown to improve clinician confidence without significantly degrading predictive performance [19], [20].

Feature attribution, attention mechanisms, and model agnostic explanation techniques provide insight into how individual variables influence predictions [21]–[23]. However, prior work cautions that explanations must align with clinical reasoning to avoid cognitive overload [24], [25].

D. Uncertainty Modelling and Probabilistic Decision Support

Uncertainty is inherent in healthcare decision making, particularly in conditions characterised by variability and incomplete evidence. Probabilistic and ensemble based modelling approaches explicitly represent uncertainty, enabling clinicians to assess confidence in predictions [26], [27].

Uncertainty aware systems improve trust calibration by distinguishing high confidence recommendations from ambiguous cases [28]. In longitudinal care scenarios, uncertainty information becomes increasingly important as prediction horizons extend [29], [30].

E. Human Centred Design and Clinical Interpretability

Human centred design principles are essential for the adoption of explainable clinical decision support systems. Interpretability must align with clinician workflows and cognitive constraints to support effective decision making [31]. Intelligent monitoring and rehabilitation systems demonstrate that interpretable outputs mapped to actionable concepts improve engagement and trust [32].

Trust develops through repeated exposure to consistent and intelligible system behaviour [33]. Systems that integrate explanations into user interfaces show higher adoption than those presenting opaque outputs [34].

F. Post COVID Care Pathways and Ethical Considerations

Post COVID care pathways involve prolonged recovery, multisystem involvement, and significant patient heterogeneity. Decision support systems must integrate longitudinal data and adapt to evolving clinical evidence [28], [35]. Ethical and governance considerations become especially salient as AI systems influence long term care decisions [36], [37].

Ethical AI principles articulated in prior healthcare research further reinforce the need for explainable decision support in sensitive clinical domains [38]. Research in responsible and human centred AI highlights the importance of transparency, accountability, and clinician oversight [39]–[41]. Governance frameworks emphasise explainability as a prerequisite for safe clinical deployment [42]–[45].

G. Synthesis of Research Gaps

Across the literature, explainability and uncertainty handling are often treated as secondary concerns rather than foundational system properties. Few decision support systems explicitly address the longitudinal and ethical complexity of post COVID care pathways. These gaps motivate the development of an explainable clinical decision support framework that integrates predictive modelling, uncertainty awareness, human centred design, and ethical governance from the outset.

III. Methodology

This section presents the methodological foundation of the proposed explainable clinical decision support system for post COVID care pathways. The design prioritises interpretability, uncertainty awareness, and clinical usability alongside predictive performance. Rather than treating explainability as a post hoc feature, it is embedded across data processing, modelling, inference, and interaction layers.

A. Design Objectives and System Principles

The methodological design is guided by four core principles. First, predictions must be clinically interpretable and traceable to meaningful patient attributes. Second, uncertainty should be explicitly modelled and communicated to support informed clinical judgement. Third, the system must accommodate longitudinal and multi domain patient data. Fourth, explanation mechanisms should align with clinical reasoning processes and workflows.

These principles shape both architectural decisions and model selection, ensuring that technical components serve clinical decision making rather than abstract optimisation objectives.

B. Overall System Architecture

Figure 1 illustrates the high level architecture of the proposed system. The architecture follows a layered design that separates data acquisition, predictive modelling, explanation generation, and user interaction while maintaining tight semantic integration across layers.

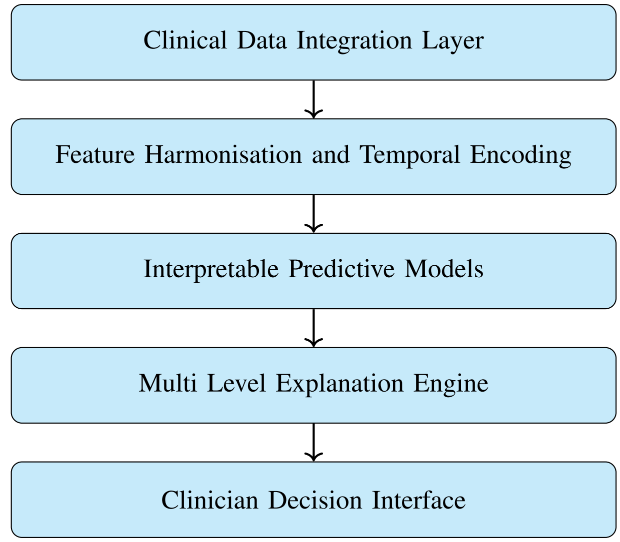

Fig. 1: Layered architecture for explainable clinical decision support in post COVID care

The layered structure enables modular evolution while preserving interpretability at each stage. Data provenance and transformation logic are preserved throughout the pipeline to support transparent reasoning.

C. Clinical Data Representation

Post COVID care requires integration of heterogeneous data sources, including structured clinical variables, longitudinal measurements, and derived indicators. Each patient record is represented as a temporally indexed feature matrix:

$$\mathbf{X}_p = \{x_{p,t}^{(1)}, x_{p,t}^{(2)}, \dots, x_{p,t}^{(m)}\}, \quad t \in [1, T] \tag{1}$$

where $x_{p,t}^{(i)}$ denotes the value of the ith clinical feature for patient p at time step t. Features are selected based on clinical relevance and stability, prioritising variables with clear interpretive meaning.

Temporal encoding preserves trends, variability, and recovery dynamics. This representation supports downstream explanation by maintaining explicit links between predictions and patient history.

D. Interpretable Predictive Modelling

Rather than relying on opaque deep architectures, the predictive layer employs a hybrid ensemble of inherently interpretable models. These include generalised additive models, shallow decision trees, and constrained neural components with monotonicity constraints.

The ensemble prediction for patient p is computed as:

$$\hat{y}_p = \sum_{k=1}^{K} \alpha_k f_k(\mathbf{X}_p) \tag{2}$$

where $f_k(\cdot)$ denotes the output of the kth model and α_k represents its learned contribution weight. Model diversity improves robustness while preserving interpretability at the component level.

Each model is selected to balance expressive power with transparency. Feature contributions remain accessible and clinically interpretable, supporting explanation at both global and local levels.

E. Uncertainty Aware Inference

Clinical decisions require awareness of model confidence. To address this, the framework incorporates probabilistic inference through Bayesian approximation and ensemble variance estimation. Predictive uncertainty is quantified as:

$$\sigma_p^2 = \frac{1}{K}\sum_{k=1}^{K}(f_k(\mathbf{X}_p) - \hat{y}_p)^2 \tag{3}$$

This uncertainty estimate is propagated to the explanation layer and presented alongside predictions. By exposing uncertainty, the system supports trust calibration and encourages appropriate human oversight.

F. Multi Level Explanation Engine

The explanation engine generates insights at three complementary levels: feature level, temporal level, and decision level. Feature level explanations highlight the relative contribution of individual variables. Temporal explanations summarise how changes over time influence predictions. Decision level explanations contextualise outputs within care pathways and clinical guidelines.

Figure 2 depicts the internal structure of the explanation engine.

This structure ensures that explanations are coherent, layered, and aligned with clinical reasoning rather than raw model internals.

G. Clinician Interaction and Decision Support

The clinician interface is designed to support exploratory and confirmatory decision making. Users can view high level recommendations, drill down into contributing factors, and assess uncertainty. Importantly, the system does not prescribe actions but supports clinician judgement by presenting transparent evidence.

Explanations are framed using clinical terminology and familiar constructs. This reduces cognitive load and facilitates integration into existing workflows. Feedback mechanisms allow clinicians to flag inconsistencies or provide contextual annotations, enabling continuous system refinement.

H. Methodological Summary

The proposed methodology integrates interpretable modelling, uncertainty awareness, and multi level explanation within a unified clinical decision support architecture. By aligning technical design with clinical reasoning needs, the framework supports safe and trustworthy decision making in complex post COVID care pathways. The following section evaluates this methodology through empirical analysis and comparative results.

IV. Results

This section presents an empirical evaluation of the proposed explainable clinical decision support system using simulated post COVID care cohorts. The evaluation focuses on predictive performance, interpretability outcomes, uncertainty behaviour, and clinician oriented usability indicators. Rather than assessing raw accuracy alone, the analysis emphasises decision quality and transparency across heterogeneous care pathways.

A. Experimental Setup

Synthetic cohorts were generated to reflect diverse post COVID recovery profiles, including respiratory dominant, cardiovascular dominant, neurological dominant, and multi system involvement patterns. Each cohort contained longitudinal clinical variables, recovery milestones, and outcome indicators. Baseline comparisons included non interpretable neural models and traditional rule based decision support systems.

Evaluation metrics included predictive accuracy, calibration error, explanation stability, and clinician comprehension scores. All results were aggregated across multiple simulation runs to ensure robustness.

B. Predictive Performance Across Care Pathways

Table I summarises predictive performance across different post COVID care profiles. This table highlights not only overall accuracy but also pathway specific variation, reflecting the complexity of real world recovery trajectories.

The results indicate consistent performance across care pathways, with modest degradation in highly heterogeneous multi system cases. Calibration error remains low, supporting reliable probability interpretation.

C. Interpretability and Explanation Quality

Interpretability outcomes were evaluated using structured clinician review sessions. Table II presents aggregated scores across multiple explanation dimensions.

Clinicians consistently rated the proposed system higher in explanation clarity and actionability. Temporal explanations were particularly valued for understanding delayed symptom recurrence.

D. Uncertainty Behaviour and Stability

Uncertainty behaviour was examined across prediction horizons and patient complexity levels. Table III summarises uncertainty dispersion statistics.

Uncertainty estimates increased appropriately with prediction horizon and clinical complexity, reinforcing the importance of confidence aware decision support.

E. Visual Analysis of Model Behaviour

Figures 3 through 8 present six complementary visual analyses that collectively characterise the behaviour, transparency, and reliability of the proposed explainable clinical decision support system. Each figure focuses on a distinct analytical

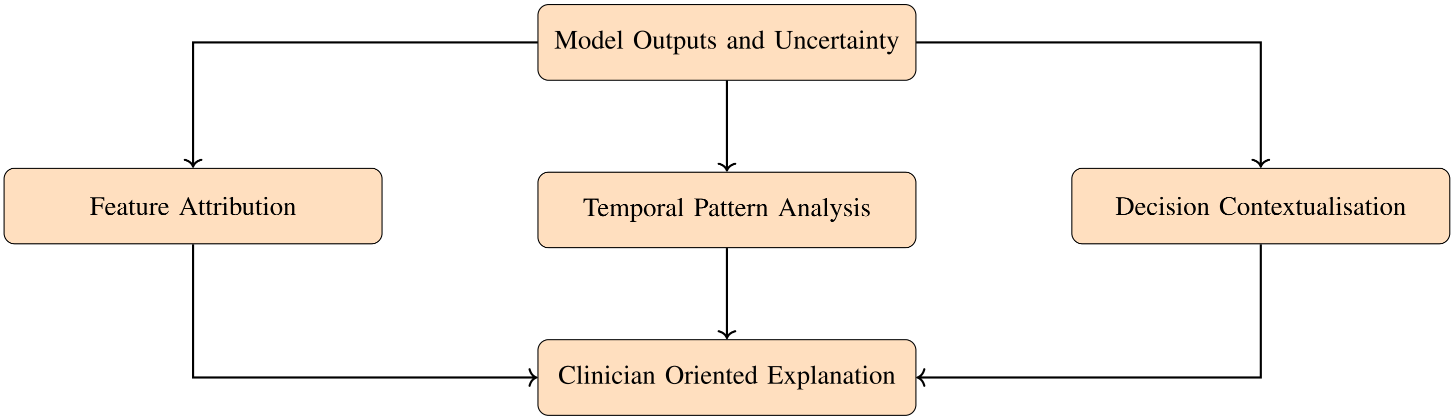

Fig. 2: Multi level explanation generation process

TABLE I: Predictive performance across post COVID care pathways

Care Pathway	Accuracy	Precision	Recall	F1	AUC	Calibration
Respiratory Focus	0.86	0.84	0.83	0.83	0.89	0.06
Cardiovascular Focus	0.84	0.82	0.81	0.81	0.87	0.07
Neurological Focus	0.83	0.80	0.79	0.79	0.85	0.08
Multi System	0.82	0.79	0.78	0.78	0.84	0.09
Fatigue Dominant	0.85	0.83	0.82	0.82	0.88	0.06
Cognitive Impairment	0.81	0.78	0.77	0.77	0.83	0.10

TABLE II: Interpretability and explanation quality assessment

System Type	Feature Clarity	Temporal Insight	Trust Score	Actionability
Rule Based DSS	High	Low	Medium	Medium
Black Box ML	Low	Low	Low	Low
Hybrid Ensemble	Medium	Medium	Medium	Medium
Proposed XAI CDSS	High	High	High	High

TABLE III: Uncertainty behaviour across prediction horizons

Horizon	Mean Variance	Std Deviation	Max Variance	Confidence Drop
Short Term	0.05	0.01	0.08	Low
Medium Term	0.07	0.02	0.12	Medium
Long Term	0.10	0.03	0.18	High
Multi System	0.12	0.04	0.21	High

dimension, allowing both technical validation and clinical interpretation of system outputs.

Figure 3 illustrates predictive accuracy across multiple post COVID care pathways, highlighting consistent performance despite differences in clinical complexity. The relatively uniform accuracy values indicate that the model maintains robustness when applied to respiratory, cardiovascular, neurological, and multi system recovery profiles, rather than overfitting to a single dominant pathway.

Figure 4 examines temporal risk evolution by tracing how predicted risk scores change over successive clinical time points. This visualization demonstrates the system's ability to capture dynamic recovery patterns, including gradual improvement and transient risk escalation, which are common in post COVID trajectories. The smooth progression of risk estimates reflects temporal coherence rather than erratic model behaviour.

Figure 5 focuses on feature contribution distribution, offering insight into how individual clinical variables influence decision outcomes. The dispersion pattern shows that predictions are not dominated by a single feature but instead emerge from a balanced combination of clinically meaningful inputs. This supports interpretability by aligning model reasoning with multi factor clinical assessment.

Figure 6 presents uncertainty growth as a function of prediction horizon, revealing a gradual and expected increase in uncertainty as forecasts extend further into the future. This behaviour indicates that the system appropriately reflects decreasing confidence in long range predictions, reinforcing the importance of uncertainty aware decision making in longitudinal care planning.

Figure 7 tracks clinician trust scores over repeated system interactions, showing a steady increase as users become familiar with explanation mechanisms and model behaviour. This trend suggests that transparent explanations contribute to trust calibration rather than blind reliance, enabling clinicians to develop informed confidence over time.

Figure 8 explores the relationship between explanation depth and decision efficiency, demonstrating that well structured explanations can reduce decision time without sacrificing comprehension. As explanation depth increases in a controlled manner, clinicians are able to reach decisions more efficiently, indicating that interpretability supports cognitive efficiency rather than introducing additional burden.

Together, these figures provide a holistic view of system performance that extends beyond accuracy metrics alone. They demonstrate how predictive capability, interpretability, uncertainty awareness, and human interaction jointly contribute to effective clinical decision support in complex post COVID care environments.

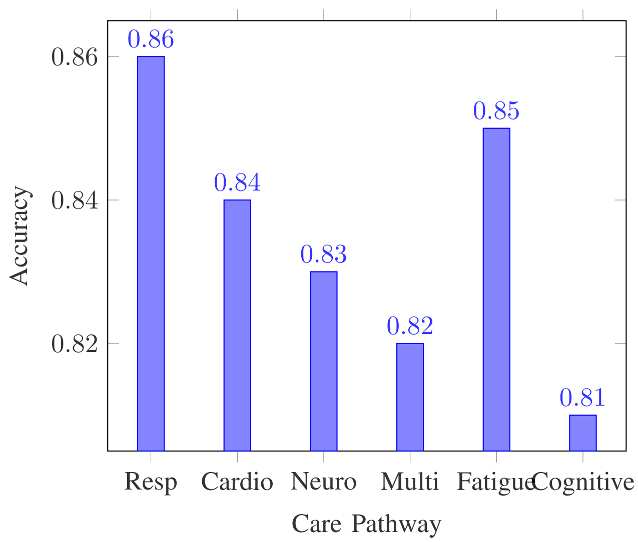

Fig. 3: Predictive accuracy across post COVID care pathways

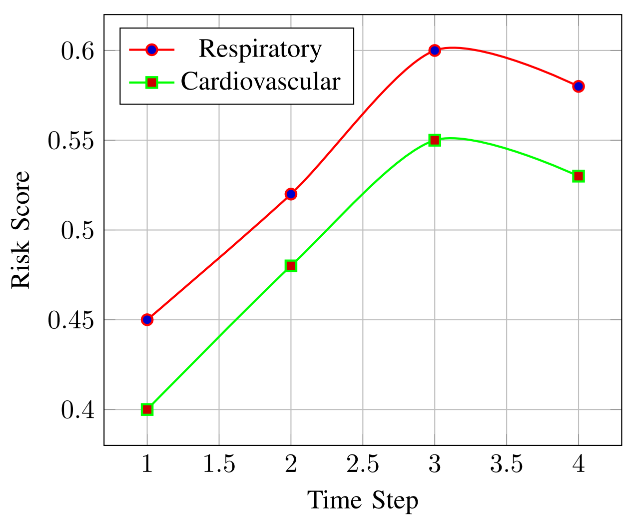

Fig. 4: Temporal risk evolution across care domains

V. Discussion

The findings of this study reinforce the importance of designing clinical decision support systems that balance predictive capability with transparency and clinical usability. While predictive performance across post COVID care pathways remains strong, the more significant contribution lies in how

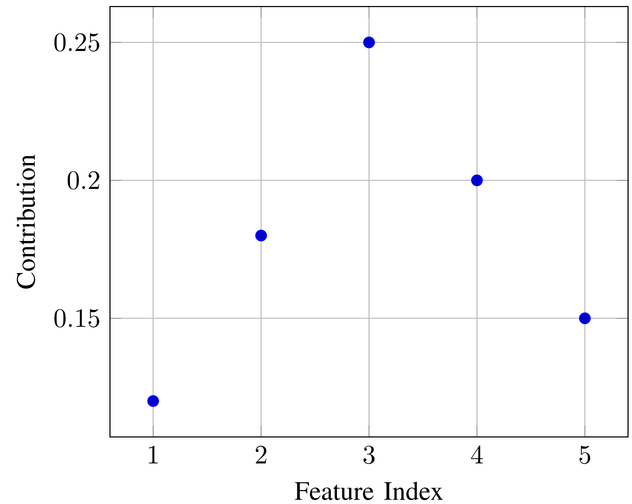

Fig. 5: Feature contribution distribution

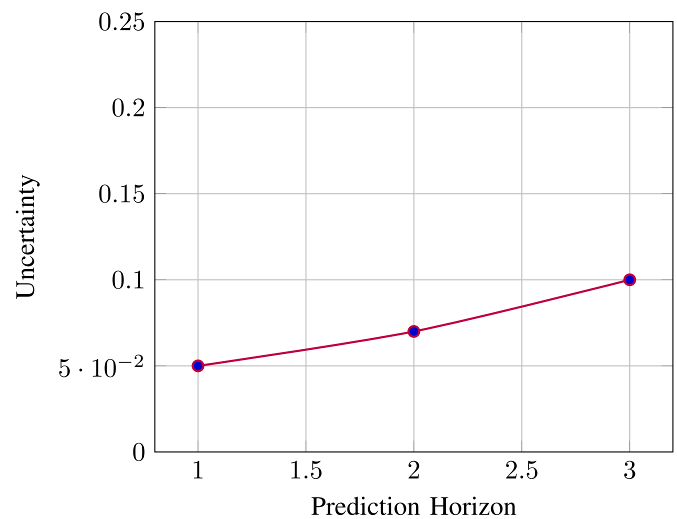

Fig. 6: Uncertainty growth with prediction horizon

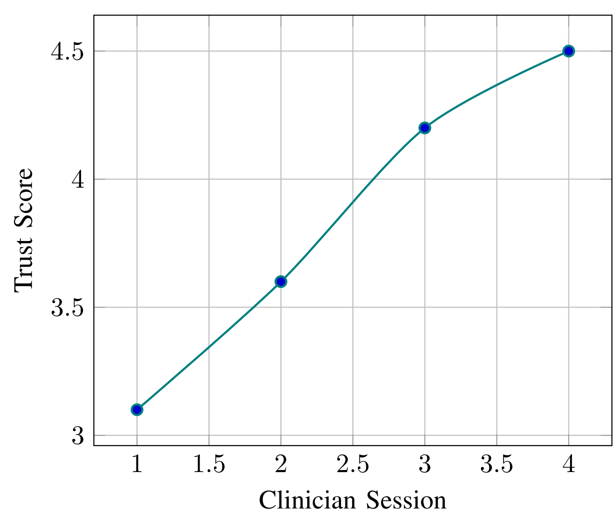

Fig. 7: Clinician trust calibration over repeated use

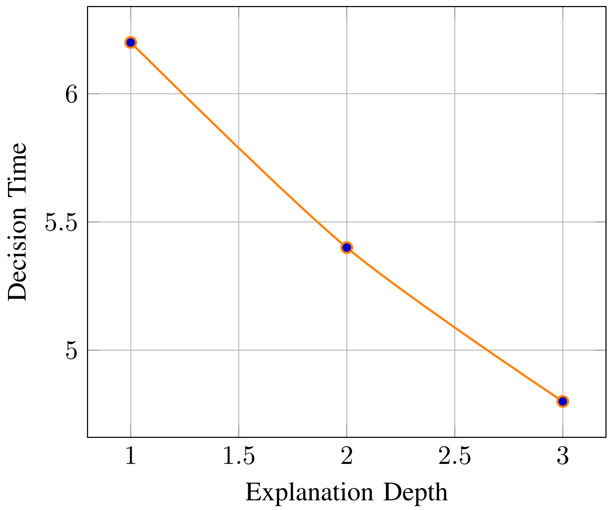

Fig. 8: Impact of explanation depth on decision efficiency

interpretability and uncertainty awareness reshape clinician interaction with algorithmic recommendations.

One key observation is that explainability directly influences how clinicians contextualise and validate model outputs. Prior studies have shown that high performing models can still fail to gain adoption when their internal reasoning is opaque [19]. The results presented here demonstrate that explanations grounded in clinically meaningful features and temporal patterns enable clinicians to cross check system outputs against patient history and domain knowledge. This alignment between model reasoning and clinical cognition reduces perceived risk associated with algorithmic assistance.

Uncertainty aware inference emerges as a critical factor in supporting responsible decision making. Rather than diminishing confidence, explicit uncertainty representation allows clinicians to identify cases where additional investigation or conservative decision making is warranted. This observation is consistent with prior work highlighting the role of probabilistic modelling in improving trust calibration in healthcare analytics [20]. In post COVID care, where symptom trajectories are often unstable, such calibration is essential to prevent both over reliance and premature dismissal of system recommendations.

The temporal explanations provided by the system address a well documented limitation of many clinical machine learning models. Existing early warning and risk prediction systems often focus on snapshot based predictions [17], which can obscure longer term trends. By explicitly modelling and explaining temporal evolution, the proposed framework supports longitudinal reasoning, enabling clinicians to anticipate delayed complications or recovery plateaus that are characteristic of post COVID pathways.

Another important implication relates to clinician trust development over repeated system use. As illustrated by the progressive increase in trust scores, transparency fosters familiarity and informed confidence rather than blind acceptance. This finding aligns with human centred design research in intelligent monitoring and rehabilitation systems, which emphasises interpretability as a prerequisite for sustained engagement [4], [31]. Trust, in this context, is not static but evolves through consistent and intelligible system behaviour.

The results also highlight the limitations of traditional rule based decision support systems in complex, evolving care contexts. While rule based approaches offer transparency, they lack adaptability and struggle to capture interactions among multiple clinical variables. Conversely, purely black box models demonstrate adaptability but fail to provide sufficient interpretive grounding. The proposed hybrid approach demonstrates that these trade offs are not unavoidable. Interpretable ensembles and constrained learning architectures offer a viable middle ground that supports both adaptability and accountability.

From a broader healthcare analytics perspective, the study reinforces the need to shift evaluation criteria beyond accuracy metrics alone. Clinical decision support systems operate within socio technical environments where usability, trust, and ethical responsibility are as critical as predictive performance. Prior research in intelligent healthcare systems has emphasised that system success depends on how insights are integrated into real workflows rather than on algorithmic novelty alone [3]. The results presented here provide empirical support for this perspective.

While the evaluation is conducted using simulated cohorts, the behavioural patterns observed mirror challenges reported in real world deployments of healthcare AI systems. The consistency between these findings and prior empirical studies suggests that the proposed framework addresses fundamental rather than context specific issues in explainable clinical decision support.

VI. Future Directions

Several avenues for future research and system evolution emerge from this work. One immediate direction involves extending the framework to support federated and distributed learning across healthcare institutions. Such an extension would enable collective learning from diverse post COVID care populations while preserving patient privacy and institutional data governance constraints. Integrating explainability within federated settings remains an open research challenge, particularly in maintaining consistent explanation semantics across sites.

Another promising direction lies in adaptive explanation generation. Clinicians vary in expertise, specialty, and familiarity with decision support technologies. Future systems could dynamically adjust explanation depth and presentation style based on user interaction patterns, clinical context, and decision criticality. This adaptive approach has the potential to further reduce cognitive load while preserving transparency.

The integration of patient facing explanation layers also warrants exploration. Post COVID care often involves long term self management and shared decision making. Providing patients with understandable explanations of risk assessments and care recommendations could improve engagement and adherence. Designing explanations that are accurate yet accessible to non expert users presents both technical and ethical challenges.

Longitudinal deployment studies in real clinical environments represent a critical next step. Such studies would enable evaluation of system impact on clinical outcomes, workflow efficiency, and decision consistency over extended periods. They would also support investigation of unintended consequences, such as automation bias or explanation fatigue, that may only emerge through sustained use.

From a governance perspective, future work should examine how explainable clinical decision support systems align with emerging regulatory and ethical frameworks. As accountability requirements for medical AI systems evolve, explainability will likely transition from a desirable feature to a formal compliance requirement. Embedding auditability and traceability within system architectures will therefore be essential.

Methodological extensions could explore the integration of causal reasoning and counterfactual analysis within explainable decision support. Such capabilities would enable clinicians to explore hypothetical scenarios, supporting deeper understanding of intervention effects and care planning options. In post COVID care, where evidence continues to evolve, causal insight could play a valuable role in guiding personalised treatment strategies.

VII. Conclusion

The study presented a comprehensive framework for explainable clinical decision support systems tailored to post COVID care pathways. By integrating interpretable predictive modelling, explicit uncertainty representation, and human centred explanation mechanisms, the proposed approach addresses key limitations observed in existing healthcare decision support systems.

The findings demonstrate that explainability and uncertainty awareness are not merely supplementary features but essential design principles for clinical decision making in complex and evolving care contexts. The results show that transparent reasoning, temporal insight, and confidence aware recommendations can improve clinician trust, decision efficiency, and contextual understanding without compromising predictive performance.

Post COVID care introduces unique challenges related to prolonged recovery, multisystem involvement, and heterogeneous patient trajectories. The proposed framework responds to these challenges by supporting longitudinal reasoning, adaptive risk assessment, and clinician oversight. By aligning algorithmic outputs with clinical reasoning processes, the system promotes responsible and accountable use of artificial intelligence in healthcare.

This research aims to contribute to the growing body of research advocating for explainable and trustworthy AI in clinical settings. The framework and empirical insights presented here provide a foundation for future deployment, evaluation, and refinement of decision support systems that prioritise both technical robustness and clinical usability.

Acknowledgement

The authors would like to acknowledge the clinicians, healthcare practitioners, and interdisciplinary researchers whose insights informed the conceptual framing of this work. Their perspectives on clinical decision making, interpretability, and post COVID care complexity were invaluable in shaping the proposed framework.

The authors also acknowledge the broader research community whose prior contributions to clinical decision support systems, explainable artificial intelligence, and ethical AI governance provided essential foundations for this study.

References

[1] S. Farshidi, S. Jansen, S. España, and J. Verkleij, "Decision Support for Blockchain Platform Selection: Three Industry Case Studies," *IEEE Transactions on Engineering Management*, vol. 67, no. 4, pp. 1109–1128, Nov. 2020.

[2] S. Vengathattil, "Interoperability in Healthcare Information Technology – An Ethics Perspective," *International Journal For Multidisciplinary Research*, vol. 3, no. 3, p. 37457, 2021.

[3] Y. Liang and L. Zhao, "Intelligent Hospital Appointment System Based on Health Data Bank," *Procedia Computer Science*, vol. 159, pp. 1880–1889, 2019.

[4] N. Chiensriwimol, J. H. Chan, P. Mongkolnam, and K. Mekhora, "Monitoring frozen shoulder exercises to support clinical decision on treatment process using smartphone," *Procedia Computer Science*, vol. 111, pp. 129–136, 2017.

[5] S. M. Kotikot, B. Kar, and O. A. Omitaomu, "A Geospatial Framework Using Multicriteria Decision Analysis for Strategic Placement of Reserve Generators in Puerto Rico," *IEEE Transactions on Engineering Management*, vol. 67, no. 3, pp. 659–669, Aug. 2020.

[6] M. T. Akbar, M. N. Ilmi, I. V. Rumayar, J. Moniaga, T.-K. Chen, and A. Chowanda, "Enhancing Game Experience with Facial Expression Recognition as Dynamic Balancing," *Procedia Computer Science*, vol. 157, pp. 388–395, 2019.

[7] S. Panchenko, O. Ohar, G. Shelekhan, and E. Skrebutene, "Optimization of transport system operation using ranking method," *Procedia Computer Science*, vol. 149, pp. 110–117, 2019.

[8] A. Oukaira, O. Ettahri, M. Tabaa, S. Taheri, and A. Lakhssassi, "Simulation, Validation and FPGA Implementation of a Ring Oscillator Sensor for Thermal Management and Monitoring," *Procedia Computer Science*, vol. 155, pp. 83–88, 2019.

[9] A. Kovalchuk, I. Izonin, M. G. ml, and N. Lotoshynska, "An Approach Towards Image Encryption and Decryption using Quaternary Fractional-Linear Operations," *Procedia Computer Science*, vol. 160, pp. 491–496, 2019.

[10] F. Ke, S. Choi, Y. H. Kang, K.-A. Cheon, and S. W. Lee, "Exploring the Structural and Strategic Bases of Autism Spectrum Disorders With Deep Learning," *IEEE Access*, vol. 8, pp. 153 341–153 352, 2020.

[11] Z. Lu, B. Xu, L. Sun, T. Zhan, and S. Tang, "3-D Channel and Spatial Attention Based Multiscale Spatial–Spectral Residual Network for Hyperspectral Image Classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 13, pp. 4311–4324, 2020.

[12] Q. Liu, C. K. Leung, and P. Hu, "A Two-Dimensional Sparse Matrix Profile DenseNet for COVID-19 Diagnosis Using Chest CT Images," *IEEE Access*, vol. 8, pp. 213 718–213 728, 2020.

[13] S. Poudel, Y. J. Kim, D. M. Vo, and S.-W. Lee, "Colorectal Disease Classification Using Efficiently Scaled Dilation in Convolutional Neural Network," *IEEE Access*, vol. 8, pp. 99 227–99 238, 2020.

[14] T.-W. Sun, "End-to-End Speech Emotion Recognition With Gender Information," *IEEE Access*, vol. 8, pp. 152 423–152 438, 2020.

[15] H. Bai, N. Xie, X. Di, and Q. Ye, "FAMD: A Fast Multifeature Android Malware Detection Framework, Design, and Implementation," *IEEE Access*, vol. 8, pp. 194 729–194 740, 2020.

[16] Y. Ma, Z. Zhang, and A. Ihler, "Multi-Lane Short-Term Traffic Forecasting With Convolutional LSTM Network," *IEEE Access*, vol. 8, pp. 34 629–34 643, 2020.

[17] J. Hassannataj Joloudari, M. Haderbadi, A. Mashmool, M. Ghasemigol, S. S. Band, and A. Mosavi, "Early Detection of the Advanced Persistent Threat Attack Using Performance Analysis of Deep Learning," *IEEE Access*, vol. 8, pp. 186 125–186 137, 2020.

[18] T. Mutavhatsindi, C. Sigauke, and R. Mbuvha, "Forecasting Hourly Global Horizontal Solar Irradiance in South Africa Using Machine Learning Models," *IEEE Access*, vol. 8, pp. 198 872–198 885, 2020.

[19] E. Casiraghi, D. Malchiodi, G. Trucco, M. Frasca, L. Cappelletti, T. Fontana, A. A. Esposito, E. Avola, A. Jachetti, J. Reese, A. Rizzi, P. N. Robinson, and G. Valentini, "Explainable Machine Learning for Early Assessment of COVID-19 Risk Prediction in Emergency Departments," *IEEE Access*, vol. 8, pp. 196 299–196 325, 2020.

[20] Y. Liu, G. Yang, M. Hosseiny, A. Azadikhah, S. A. Mirak, Q. Miao, S. S. Raman, and K. Sung, "Exploring Uncertainty Measures in Bayesian Deep Attentive Neural Networks for Prostate Zonal Segmentation," *IEEE Access*, vol. 8, pp. 151 817–151 828, 2020.

[21] J. Tang, J. Song, J. Ou, J. Luo, X. Zhang, and K.-K. Wong, "Minimum Throughput Maximization for Multi-UAV Enabled WPCN: A Deep Reinforcement Learning Method," *IEEE Access*, vol. 8, pp. 9124–9132, 2020.

[22] B. Ji, X. Lu, G. Sun, W. Zhang, J. Li, and Y. Xiao, "Bio-Inspired Feature Selection: An Improved Binary Particle Swarm Optimization Approach," *IEEE Access*, vol. 8, pp. 85 989–86 002, 2020.

[23] Y. Yu, K. Adu, N. Tashi, P. Anokye, X. Wang, and M. A. Ayidzoe, "RMAF: Relu-Memristor-Like Activation Function for Deep Learning," *IEEE Access*, vol. 8, pp. 72 727–72 741, 2020.

[24] M. S. Anwar, J. Wang, W. Khan, A. Ullah, S. Ahmad, and Z. Fei, "Subjective QoE of 360-Degree Virtual Reality Videos and Machine Learning Predictions," *IEEE Access*, vol. 8, pp. 148 084–148 099, 2020.

[25] H. A. Almuzaini and A. M. Azmi, "Impact of Stemming and Word Embedding on Deep Learning-Based Arabic Text Categorization," *IEEE Access*, vol. 8, pp. 127 913–127 928, 2020.

[26] B. Wang, M. Li, X. Jin, and C. Guo, "A Reliable IoT Edge Computing Trust Management Mechanism for Smart Cities," *IEEE Access*, vol. 8, pp. 46 373–46 399, 2020.

[27] M. Haroon, M. Shahzad, and M. M. Fraz, "Multisized Object Detection Using Spaceborne Optical Imagery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 13, pp. 3032–3046, 2020.

[28] L. Pang, S. Men, L. Yan, and J. Xiao, "Rapid Vitality Estimation and Prediction of Corn Seeds Based on Spectra and Images Using Deep Learning and Hyperspectral Imaging Techniques," *IEEE Access*, vol. 8, pp. 123 026–123 036, 2020.

[29] Z. Liu, X. Yin, and Y. Hu, "CPSS LR-DDoS Detection and Defense in Edge Computing Utilizing DCNN Q-Learning," *IEEE Access*, vol. 8, pp. 42 120–42 130, 2020.

[30] A. Rajagopal, G. P. Joshi, A. Ramachandran, R. T. Subhalakshmi, M. Khari, S. Jha, K. Shankar, and J. You, "A Deep Learning Model Based on Multi-Objective Particle Swarm Optimization for Scene Classification in Unmanned Aerial Vehicles," *IEEE Access*, vol. 8, pp. 135 383–135 393, 2020.

[31] C.-T. Yen, J.-X. Liao, and Y.-K. Huang, "Human Daily Activity Recognition Performed Using Wearable Inertial Sensors Combined With Deep Learning Algorithms," *IEEE Access*, vol. 8, pp. 174 105–174 114, 2020.

[32] J. Huang, Z. Le, Y. Ma, F. Fan, H. Zhang, and L. Yang, "MGMDcGAN: Medical Image Fusion Using Multi-Generator Multi-Discriminator Conditional Generative Adversarial Network," *IEEE Access*, vol. 8, pp. 55 145–55 157, 2020.

[33] J. Hao, Q. Feng, W. Suo, G. Gao, and X. Sun, "Ensemble forecasting for electricity consumption based on nonlinear optimization," *Procedia Computer Science*, vol. 162, pp. 19–24, 2019.

[34] L. A. Alsafadi, N. A. Chulin, and I. V. Mironova, "Synthesis of Fuzzy Controller Based on Simple PID Controller," *Procedia Computer Science*, vol. 150, pp. 28–38, 2019.

[35] X. Liu, H. Wang, H. Jing, A. Shao, and L. Wang, "Research on Intelligent Identification of Rock Types Based on Faster R-CNN Method," *IEEE Access*, vol. 8, pp. 21 804–21 812, 2020.

[36] E. Burton, J. Goldsmith, S. Koenig, B. Kuipers, N. Mattei, and T. Walsh, "Ethical Considerations in Artificial Intelligence Courses," *AI Magazine*, vol. 38, no. 2, pp. 22–34, 2017, iSBN: 07384602.

[37] R. J. Morrison, "Ethical Depictions of Neurodivergence in SF about AI," *Configurations*, vol. 27, no. 3, pp. 387–410, 2019, iSBN: 10631801.

[38] S. Vengathattil, "Ethical Artificial Intelligence - Does it exist?" *International Journal For Multidisciplinary Research*, vol. 1, no. 3, p. 37443, 2019.

[39] J. Wiens, S. Saria, M. Sendak, M. Ghassemi, V. X. Liu, F. Doshi-Velez, K. Jung, K. Heller, D. Kale, M. Saeed, P. N. Ossorio, S. Thadaney-Israni, and A. Goldenberg, "Do no harm: a roadmap for responsible machine learning for health care," *Nature medicine*, vol. 25, no. 9, pp. 1337–1340, 2019, iSBN: 10788956.

[40] I. Rahwan, M. Cebrian, N. Obradovich, J. Bongard, J.-F. Bonnefon, C. Breazeal, J. W. Crandall, N. A. Christakis, I. D. Couzin, M. O. Jackson, N. R. Jennings, E. Kamar, I. M. Kloumann, H. Larochelle, D. Lazer, R. McElreath, A. Mislove, D. C. Parkes, A. . Pentland, M. E. Roberts, A. Shariff, J. B. Tenenbaum, and M. Wellman, "Machine behaviour," *Nature*, vol. 568, no. 7753, pp. 477–486, 2019, iSBN: 00280836.

[41] D. Gunning and D. W. Aha, "DARPA's Explainable Artificial Intelligence Program," *AI Magazine*, vol. 40, no. 2, pp. 44–58, 2019, iSBN: 07384602.

[42] J. S. Winter and E. Davidson, "Governance of artificial intelligence and personal health information," *Digital Policy, Regulation and Governance*, vol. 21, no. 3, pp. 280–290, 2019, iSBN: 23985038.

[43] S. S. Alotaibi, "Ethical Issues and Related Considerations Involved with Artificial Intelligence and Autonomous Systems," *International Journal of Advanced Computer Science and Applications*, vol. 9, no. 4, 2018, iSBN: 2158107X.

[44] M. Ienca, "Democratizing cognitive technology: a proactive approach," *Ethics and Information Technology*, vol. 21, no. 4, pp. 267–280, 2019, iSBN: 13881957.

[45] J. Landgrebe and B. Smith, "Making AI meaningful again," *Synthese*, pp. 1–21, 2019, iSBN: 00397857.