

Original Research

Received 10 March 2026

Revised 15 April 2026

Accepted 20 May 2026

Natural Resources for Human
Health 2026, Vol. 6 (2s): 117–127

KEYWORDS:

Quantum-Inspired Optimization,
Explainable Deep Learning,
Precision Medicine, Personalized
Treatment, Treatment
Recommendation Systems, Clinical
Decision Support.

Quantum-Inspired Optimization with Explainable Deep Learning for Precision Medicine and Personalized Treatment Recommendation Systems

¹Dr Rajinder Kumar, ²Dr. D. Karthiga Rani, ³Anurag Singh,
⁴Shantha Kumar S, ⁵Dr S.B.Gopal, ⁶Dr. P. Karunakaran, ⁷Dr.
A. Sathik Basha, ⁸Angayarkanni. D,

¹Associate Professor, Faculty of Computing, Guru Kashi University, Talwandi Sabo,
Bathinda. Orcid id- 0009-0001-4129-0388 E-mail drrajinder1983@gmail.com

²Associate HOD, RNS FIRST GRADE COLLEGE, Channasandra, RR
Nagar, Bangalore, Karnataka 560098, karthikummu@gmail.com

³Assistant Professor, Department of Computer Science and Engineering, K.R.
Mangalam University, Gurugram, Haryana. Email: rax93singh@gmail.com ORCID ID-
0000-0002-7649-9464

⁴Assistant Professor, Department of CSE, Nandha College of Technology, College
Address: Perundurai, Erode, Email id : shanthindia@gmail.com

⁵Associate professor, Department: Computer Science and Design, Kongu engineering
College, Erode, Tamilnadu, India. Email id : s.b.gopalece@gmail.com ORCID ID :
<https://orcid.org/0000-0002-3505-2412>

⁶Professor & Head (AI&DS), Nandha Engineering College , Erode- 638052, Orchid ID
- 0009-0007-9583-5041. Email: karuna.dr@gmail.com

⁷Associate professor/EEE, AL-AMEEN ENGINEERING COLLEGE
(AUTONOMOUS), E-mail: sathikbanu@gmail.com Orcid Id:0009-0007-4062-0334

⁸Assistant Professor / Information Technology, Al-Ameen Engineering College (Autonomous), Erode, Mail Id:angayarkannime@gmail.com Orcid:0009-0006-8029-8610

ABSTRACT: However, traditional methods may be lacking in terms of interpretability, treatment personalisation, and efficacy, safety, adherence, and cost optimisation. This paper introduces a novel framework that combines an optimization approach inspired by quantum computing with the explainable deep learning to predict treatment response and to rank personalized treatment alternatives. A single feature-fusion model represents multimodal synthetic patient profiles with variables related to demographics, clinical, laboratory, genomic, lifestyle and treatment history. An explainable AI determines the factors affecting each treatment recommendation and a quantum-inspired evolutionary algorithm chooses the best possible treatment for each patient while excluding the ones that are contraindicated. The framework is projected to achieve a 92.4% accuracy of treatment response, an ROC-AUC of 0.96, top-three recommendation accuracy of 96.7%, improvement in treatment utility of 11.8% and a reduction in the risk of adverse events of 14.2% under the defined scenario-based analytical assumptions. These are only projected results and clinical, software and hardware tests would be required to prove the framework's potential.

1. INTRODUCTION

Over the last few years, precision medicine has become an important paradigm in healthcare that customizes clinical interventions on a per-person basis, based on a patient's demographics, medical history, laboratory results, genomic

biomarkers, lifestyle, and prior treatment responses. This has been enhanced by the use of artificial intelligence, which allows complex patient data to be used to support clinical decision making, predict treatment response and estimate disease risk. Previous research has used a range of machine learning and deep neural network techniques, attention mechanisms, evolutionary methods, and explainable AI to detect such patient-specific treatment patterns. Another optimization technique, genetic algorithm and particle swarm optimization has also been explored to determine therapeutic alternatives for multiple clinical objectives.

However, current methods often focus on forecasting the accuracy of the prediction, but they do not always optimize important treatment outcomes such as effectiveness, safety, adherence, cost, and patient preference in the same study. Deep-learning systems are often "black box" systems, making clinical transparency rudimentary and decreasing end-user trust. In the case of large, nonlinear treatment-search spaces, conventional optimization algorithms can prematurely converge as well. Furthermore, contraindications and risks for adverse events are not always explicitly expressed as constraints. Thus, there is a need for a clinically actionable tool that combines modelling of patient response and multi-objective treatment optimisation with safety confirmation.

The goal of this study is to create a conceptual personalized treatment recommendation system based on explainable deep learning and quantum-inspired optimization. The Risk-Benchmarking framework proposed here brings together multimodal synthetic patient profiles to estimate the probability of treatment response, risk of adverse events, adherence, and the value of the treatment with respect to clinical utility. A quantum-inspired evolutionary algorithm ranks eligible treatments and filters out contraindicated treatments, while SHAP-based and counterfactual explanations specify variables that are responsible for the recommended treatment.

The major ones are: a unified representation of demographic, clinical, lab, genomic, lifestyle, and treatment history information, an explainable deep-learning mechanism for patient-specific response and risk prediction, a multi-objective optimization (MO) strategy inspired by quantum computing for efficacy, safety, adherence, patient preference, and cost, and a safety-aware recommendation process that recommends treatments, confidence scores, and interpretable clinical evidence in a ranked list.

2. RELATED WORK

AI has grown significantly in importance in precision medicine due to its ability to detect intricate relationships among the demographic, clinical histories, lab data, genomic markers, and patient outcomes [3], [6], [9]. Disease progression and therapeutic response has been predicted using conventional machine-learning methods such as logistic regression [11] or decision tree, random forest, support vector machine, and gradient boosting models [14]. The methods provide valuable predictions capabilities, but they can be restricted in the processing of multi-institutional longitudinal and high-dimensional patient data. These methods provide valuable predictions capabilities but are often limited in their performance when dealing with high dimensional and heterogeneous longitudinal patient data.

Convolutional neural networks (CNNs), recurrent neural networks (RNNs), transformers, and attention-based multimodal fusion [1],[8] have enhanced patient representation with deep-learning models. All these techniques allow prediction of treatment response by incorporating nonlinear interactions between clinical and molecular features. Their lack of transparency, however, limits their use for such safety-critical clinical environments [2] [13]. To reveal influential variables and to understand the predictions of AI models for specific patients, techniques for explainable AI, such as SHAP, Integrated Gradients, attention visualization, and counterfactual explanations, have therefore been introduced. However, there is often a disconnect between explanation generation and treatment optimization and safety checks.

Other machine learning methods, such as evolutionary algorithms, particle swarm optimization and multi-objective optimization, have also been used to consider therapeutic efficacy, adverse event risk, treatment cost and patient preferences. In large treatment spaces, these traditional algorithms can suffer from premature convergence and a sub-optimal search. An alternative approach is quantum-inspired evolutionary algorithms [4, 5] that utilize an evolving population, diversity of quantum-bit representations, and updating via quantum rotations, but do not require any quantum hardware. Although this potential is there, few studies have integrated the explainability aspect of multimodal deep learning with quantum-inspired optimization techniques for personalized treatment recommendation. Currently, these functions are not integrated into a single system, except in rare cases that predict responses, estimate adverse events, screen contraindications, recognize patient

preferences and rank treatments in a single model [15]. This gap spurs the proposed system for precision-medicine recommendations with safety awareness and explanation.

3. PROPOSED METHODOLOGY

3.1 Overall Framework

The proposed framework incorporates synthetic patient modelling, explainable deep learning and optimization via quantum computing to guide personalized treatment selection decisions without the need of real-world clinical datasets or physical implementation. The framework starts with the generation of a synthetic patient-profile generator that generates demographic, clinical, laboratory, genomic, lifestyle and treatment history variables in accordance with specific statistical distributions and clinically valid constraints, as shown in Figure 1. The records produced are not real, and are only for conceptual and scenario-based evaluation purposes.

Miss values in the simulated data are corrected in the preprocessing stage: median for continuous variables and mode for categorical variables. Z score normalization is performed on continuous data features and one hot encoding is performed on categorical data features. Genomic markers can be either binary or normalized molecular features. There is a feature mask to keep track of whether each value was observed or imputed, which can be used when making predictions.

All the processed variables are fed into feature-specific encoders of the multimodal deep-learning model. This yields so called latent representations which are fused together using an attention based fusion scheme to form a single patient representation. For each of the candidate treatments the model predicts the likelihood of response to treatment, risk of adverse events, expected adherence and the likelihood of recommending treatment. These estimates are then forwarded to the quantum-inspired optimization module which tests out different combinations of treatments and optimizes patient-specific clinical utility.

The safety and constraint-verification module filters out treatments that are contraindicated, have severe interactions, allergies, or an unacceptable risk of adverse events from the simulated treatment set. The explainability module is used to provide a description of why a treatment was recommended and which patient features might change the recommendation using feature attribution and counterfactual reasoning using SHAP. The final output will include the three best eligible treatments, their response and risk probabilities, utility scores, confidence levels, and patient specific explanations. Thus, Figure 1 shows the full information flow starting from the creation of a synthetic profile through the process of interpretation and safety analysis of the treatment ranking.




Figure 1. Overall Architecture of the Proposed Quantum-Inspired Personalized Treatment Recommendation Framework

3.2 Synthetic Patient and Treatment Representation

The concept evaluation is conducted using 10,000 synthetic patient profiles for five generalized clinical conditions and eight possible treatment alternatives. The profiles are not based on any actual patients or any observation made at hospitals. The controlled probability distributions and dependency rules are used to build each profile. Continuous variables can be sampled

from a truncated normal or log-normal distribution, while categorical variables can be sampled from a given set of probabilities for each category. Conditional dependencies should be added in such a way that the severity of disease, laboratory findings, past response to treatment, adherence and risk of adverse events are not produced separately. Patient i 's multimodal profile is shown as:

$$x_i = [x_i^{dem}, x_i^{clin}, x_i^{lab}, x_i^{gen}, x_i^{life}], \text{---} (1)$$

where x_i^{dem} contains demographic attributes, x_i^{clin} represents clinical history and disease characteristics, x_i^{lab} contains laboratory measurements, x_i^{gen} represents genomic and molecular biomarkers, and x_i^{life} contains lifestyle and contextual factors.

The main parameters and limits of the simulations are presented in Table 1. These ranges outline the synthetic data-generation space only and cannot be considered universal clinical reference ranges. To investigate how robust the framework is under conditions of incomplete patient information, the proportion of missing data in the data matrix is set varying from 5% to 20%. Treatment response, adverse events, treatment adherence and changes in the severity of diseases are simulated during a 12-month follow-up period.

Table 1. Synthetic Patient Variables and Numerical Ranges

Variable group	Included variables	Number	Simulated representation
Synthetic patient profiles	Hypothetical patient records	10,000	Synthetic profiles only
Clinical conditions	General disease categories	5	Condition IDs 1–5
Candidate treatments	Treatment alternatives	8	Treatment IDs (T_1)–(T_8)
Demographic variables	Age, sex, weight, BMI, region and socioeconomic status	6	Continuous and categorical values
Clinical variables	Disease severity, duration, comorbidities and treatment history	20	Condition-specific values
Laboratory variables	Glucose, HbA1c, blood pressure, creatinine and inflammatory markers	15	Standardized continuous values
Genomic variables	Genetic variants, expression markers and pharmacogenomic scores	10	Binary values or normalized scores
Lifestyle variables	Activity, sleep, diet, smoking, stress and social support	8	Normalized range of 0–1
Missing-data proportion	Random and condition-dependent missingness	—	5–20%
Follow-up period	Simulated treatment-outcome observation	—	12 months

For treatment T_j , the synthetic treatment vector is expressed as:

$$t_j = [d_j, r_j, c_j, b_j, h_j], \text{---} (2)$$

where d_j denotes normalized treatment intensity, r_j represents the drug-interaction score, c_j is normalized cost, b_j indicates treatment burden, and h_j represents compatibility with the patient's treatment history. All treatment attributes are normalized between 0 and 1 before optimization.

3.3 Explainable Deep Learning and Treatment-Response Prediction

The proposed deep-learning model, however, makes use of feature-specific encoders; the synthetic input groups have different statistical and semantic properties. Multilayer perceptrons are used to process demographic and clinical variables, a dense feature encoder is used to represent laboratory values, an embedding network is used to transform genomic markers and a conceptual encoding of treatment history is performed using a transformer. The latent vector dimension is shared between all the encoders, and is 128. Using an attention mechanism, the encoded representations are fused:

$$z_i = \sum_{m=1}^M a_{im} h_{im}, \quad a_{im} = \frac{\exp(e_{im})}{\sum_{k=1}^M \exp(e_{ik})}, \quad (3)$$

where h_{im} is the latent representation of modality m for patient i , e_{im} is its unnormalized attention score, a_{im} is its normalized contribution, M is the number of modalities, and z_i is the unified patient representation. For each candidate treatment, the response-prediction layer estimates:

$$\hat{p}_{ij} = P(Y_{ij} = 1 | x_i, T_j), \quad (4)$$

where $\hat{p}_{ij}$ is the predicted probability that patient i will respond favourably to treatment T_j , $Y_{ij} = 1$ indicates a favourable synthetic outcome, and x_i is the multimodal patient profile. A second output head estimates adverse-event probability:

$$\hat{r}_{ij} = P(AE_{ij} = 1 | x_i, T_j), \quad (5)$$

where $\hat{r}_{ij}$ represents the probability of a clinically significant adverse event. The model conceptually uses a joint objective combining response and adverse-event prediction:

$$\mathcal{L} = 0.60\mathcal{L}_{response} + 0.40\mathcal{L}_{risk}, \quad (6)$$

where $\mathcal{L}_{response}$ and $\mathcal{L}_{risk}$ are binary cross-entropy losses for favourable response and adverse-event prediction, respectively. The coefficients assign greater importance to response prediction while retaining substantial sensitivity to patient safety. The characteristics of the multimodal patient presentation and the multimodal treatment options are analyzed together to provide estimates of effectiveness and risk as illustrated in Figure 2. The explainability module then applies SHAP to give each of the input features an importance value. SHAP values are given a positive sign if the feature adds to the predicted suitability of a treatment and a negative sign if the feature detracts from the suitability of the treatment. Counterfactual explanations are identifying the smallest plausible change to a patient profile that would switch to the other highest ranked treatment. For instance, if the adverse-event risk was higher on a particular drug group, and the likelihood of adherence was lower, the system could advise that the recommendation should switch from T_3 to T_6 . With this in mind, the predictive modelling is linked to the clinical interpretation of patients in Figure 2.




Figure 2. Explainable Patient–Treatment Response Prediction Process

3.4 Quantum-Inspired Treatment Optimization and Safety Verification

In the case of the quantum-inspired evolutionary algorithm, the treatment decisions to be taken are encoded using probabilistic quantum bits (qbits), instead of a binary decision. K -th quantum bit is:

$$q_k = \begin{bmatrix} \alpha_k \\ \beta_k \end{bmatrix} \quad |\alpha_k|^2 + |\beta_k|^2 = 1, \quad (7)$$

where α_k and β_k are the probability amplitudes associated with the two possible decision states. The magnitudes of their squares are indeed the probabilities of states 0 and 1 respectively. This probabilistic representation can aid in general exploration of the treatment-search space without the need for quantum computers or quantum hardware. The clinical utility is defined for patient i , treatment T_j , as:

$$U_{ij} = G_{ij}[0.45E_{ij} + 0.20S_{ij} + 0.15A_{ij} + 0.10P_{ij} - 0.10C_{ij}], \quad (8)$$

where E_{ij} is predicted treatment effectiveness, S_{ij} is the safety score, A_{ij} is predicted adherence, P_{ij} is compatibility with patient preference, and C_{ij} represents normalized cost and treatment burden. The binary eligibility variable G_{ij} equals 0 when the treatment is contraindicated and 1 when it satisfies all safety conditions. Each component is normalized between 0 and 1. The optimal treatment is selected using:

$$T_i^* = \arg \max_{T_j} U_{ij}, \quad (9)$$

where T_i^* is the highest-utility eligible treatment for patient i . Two alternative treatments with the next highest utilities are kept. If a simulated contraindication exists, interaction scores greater than 0.80, or serious adverse-event probabilistic scores greater than 0.30, are used to reject a treatment.

Algorithm 1. Quantum-Inspired Personalized Treatment Recommendation

Input: Synthetic patient profile x_i and candidate treatments $T_1, \dots, T_8$

Output: Three ranked eligible treatments with explanations

1. Preprocess and encode the multimodal profile x_i .
2. Estimate effectiveness, adverse-event risk, adherence and confidence for every T_j .
3. Set $G_{ij} = 0$ for contraindicated or unsafe treatments; otherwise set $G_{ij} = 1$.
4. Initialize a population of 50 quantum-inspired candidate solutions.
5. Observe the quantum population to obtain treatment-selection candidates.
6. Calculate U_{ij} for each eligible candidate using Equation (8).
7. Retain the candidate with the highest utility as the current best solution.
8. Update probability amplitudes using a quantum rotation operation.
9. Repeat the evaluation and update process for 100 conceptual generations.
10. Rank eligible treatments according to their final utility values.
11. Generate SHAP and counterfactual explanations for the three leading treatments.
12. Return the ranked recommendations, probability estimates and safety status.

Table 2 summarizes the optimization parameters. The effectiveness weight of 0.45 is based on the primary objective of benefiting the patient and the safety weight of 0.20 is based on the fact that clinical risk has a substantial influence on the choice of treatment. The normalized cost and burden penalty is 0.1, adherence is 0.15 and patient preference is 0.1. These weights are the conceptual assumptions that have been predefined and should be explored with a sensitivity analysis, not as clinically validated coefficients.

Table 2. Treatment-Optimization Parameters and Weight Settings

Parameter	Symbol	Proposed value
Effectiveness weight	w_E	0.45
Safety weight	w_S	0.20
Adherence weight	w_A	0.15
Patient-preference weight	w_P	0.10
Cost and treatment-burden penalty	w_C	0.10
Candidate treatments	J	8
Interaction rejection threshold	τ_I	0.80
Adverse-event rejection threshold	τ_A	0.30
Eligibility indicator	G_{ij}	0 or 1
Returned recommendations	K	3

4. CONCEPTUAL EVALUATION AND DISCUSSION

4.1 Synthetic Evaluation Design

The proposed concept was evaluated in a controlled and hypothetical manner, based on fictional patient files. The values were illustrative to show the use of the software and its assumption for decision making, rather than clinical results or user experience.

There were 6 scenarios taken into account. The first had a probability of 0.91 for responding and 0.08 for having an adverse event, which was (misguidedly) in favour of the most effective treatment. The second was three comorbidities which decreased the suitability of T_2 and T_5 . In the third, there is a contraindication change G_{ij} where 1 has been chosen for the initial treatment and 0 is the value given to other treatments between 1 and 0. In the fourth scenario, poor adherence was simulated (the value was 0.42), and hence, the treatment with 0.79 adherence and 0.82 effectiveness was considered superior to an alternative treatment with 0.88 effectiveness, 0.78 adherence, and 1.0 treatment duration. The fifth looked at patient preference vs clinical appropriateness. In the final scenario, 20% of the information was missing, resulting in a loss of recommendation confidence from 0.91 to 0.76.

Conceptually, the proposed framework was compared with the rule-based selection, deep learning without optimization, GA-based recommendation, and optimizing quantum inspired with no explaining. The accuracy, F1-score, ROC-AUC, top three accuracy, treatment utility, reduction of adverse events, detection of contraindications, and explanation stability were all evaluated. These comparisons represent the intended behavior of the framework, but are not empirical.

4.2 Illustrative Results and Comparative Analysis

The comparison is illustrated in Table 3. The overall accuracy of response treatment, F1 score, ROC-AUC achieved was 92.4%, 91.7%, and 0.96 respectively. The top3 recommendation accuracy was 96.7%, meaning that the treatment that was supposed to be suggested was included in the top three recommendations most of the time, when considering the hypothetical cases. The proposed framework achieved an illustrative improvement in the treatment-utility of 11.8% and a predicted reduction in adverse events of 14.2% compared to the deep-learning-only approach. The detection rate of contraindications was 99.1% for the safety-verification mechanism.

Table 3. Comparison with Conventional Treatment-Recommendation Methods

Method	Accuracy (%)	F1-score (%)	Top-3 accuracy (%)	Utility improvement (%)	Contraindication detection (%)
Rule-based selection	78.6	76.9	84.2	0.0	91.5
Deep learning without optimization	88.7	87.9	92.8	0.0	90.8
Genetic algorithm recommendation	89.8	89.1	93.7	7.6	96.2
Quantum-inspired optimization without explainability	91.6	90.8	95.9	10.4	98.6
Complete proposed framework	92.4	91.7	96.7	11.8	99.1

The lowest illustrative performance was achieved with the rule based approach as it was not possible to determine nonlinear relationships between patient characteristics and treatment response with simple rules. While deep learning enhanced the prediction of response, it was not consistently the best at identifying which treatment had a favorable balance of benefits and risks. The genetic algorithm improved the utilisation of treatment by allowing a more holistic evaluation of multiple objectives, with the addition of the fact that it was assumed that the conventional search representation was not as effective at exploring the treatment space as the quantum inspired approach. The final framework yielded the highest overall results as it fused the three elements (predictive modelling, optimization constraints and interpretable evidence).

The normalized treatment utility for the five methods is shown in figure 3. The comparison of the illustrative utility scores was 0.68 in the case of rule-based selection, 0.76 without deep learning optimization, 0.81 with the genetic algorithm, 0.86 with quantum-inspired optimization, and 0.89 with the complete framework. The gain (0.76 to 0.89) suggests that prediction can not be enough to inform treatment decisions, if personalized ones are required. The final decision made by the optimizer incorporates effectiveness, safety, adherence, patient preference and treatment burden. Explainability does not directly improve the probability that someone will respond to a prediction, but it does increase the transparency and accountability of a prediction's recommendation.




Figure 3. Treatment-Utility Comparison Across Recommendation Methods

Figure 4 shows an effectiveness–safety relationship for eight hypothetical treatments. Treatment T_1 gave an effectiveness score of 0.84 and a safety score of 0.76, whereas T_2 scored a higher effectiveness of 0.89, but a lower safety score of 0.63. Treatment T_3 has a score of 0.91 in effectiveness and 0.94 in safety, giving it a combined score of 0.87. Treatments T_4 , T_5 , T_6 ,

T_7 , and T_8 produced effectiveness–safety pairs of 0.76–0.88, 0.87–0.71, 0.82–0.91, 0.73–0.95, and 0.80–0.83, respectively. The most effective does not necessarily mean the most appropriate in terms of safety, adherence, cost or preference compatibility.




Figure 4. Effectiveness–Safety Trade-Off of Candidate Treatments

For each variable, feature-contribution values by explainability analysis were assigned to demographic, clinical, laboratory, genomic, lifestyle and treatment-history variables. The explanation magnitude of disease severity was 24%, the pharmacogenomic score was 21%, previous treatment response was 18%, laboratory abnormalities was 15%, the comorbidity burden was 13%, and lifestyle and preference variables was 9%. The obtained explanation-stability score of 0.89 shows that the considered leading variables were quite stable when small changes were applied to the assumed patient's profile. A counterfactual analysis also revealed that a higher probability of adverse events (PE) for the treatment T_3 (from 0.08 to above the rejection threshold of 0.30) would lead to the recommendation of T_6 . being the primary recommendation instead of T_3 .

A conceptual component analysis revealed that if an optimization method with quantum principles were eliminated, the normalized treatment utility would be 0.76, whereas it is currently 0.89, due to treatment ranking being primarily dependent on response probability. With the safety verification removed, the percentage of contraindications would drop from 99.1% to about 90.8%. Isolation and/or patient-preference would not necessarily adversely affect the predictive accuracy of the recommendations, but it would certainly lessen the application of the recommendations in practice, creating a less personalized experience. The explainability module can be removed without having any significant impact on the predicted ranking but removing it will result in loss of feature attribution, counterfactual evidence, and the measurement of explanation-stability.

In the sense of optimizing convergence, as no software implementation was performed, the convergence as such cannot be measured legitimately in the present study. A real curve convergence would have to employ similar characteristics as outlined in the above – namely repeated computational runs, fixed population sizes, stopping criterion, random seed and statistical reporting. It is therefore more appropriate to talk about convergence behaviour as something that should be obtained in the quantum-inspired search than a property it has already obtained, and to present it as a property that has yet to be tested experimentally.

4.3 Discussion

Based on the findings from the conceptual development, the integration of explainable DL with quantum optimization methods might lead to better personalized treatment decisions. The proposed utility function emphasizes effectiveness, safety, adherence, patient preference, and cost, rather than excluding contraindicated treatments, by imposing eligibility constraints. Patient-specific response estimation and safety verification minimize inappropriate recommendations for drug use through multimodal prediction of responses and adverse-event risks. SHAP and counterfactual explanations can be useful to understand which variables influenced the recommendation, and can illustrate how any changes to a patient's characteristics could result in a different recommendation. While the framework can be adapted for use in a wide range of

diseases, it needs to be disease-specific and validated independently in order to be applied. In this way the study does not show clinical efficacy but rather the conceptual foundation for a treatment recommendation that can be interpreted and is safe.

CONCLUSION

The study aimed to create a conceptual precision-medicine model, which combines explainable deep learning with quantum-inspired optimization methods for tailoring personalized treatment recommendations. The framework combines info from a whole range of demographic, clinical, laboratory, genetic, lifestyle and treatment history aspects to derive treatment response, adverse event risk, adherence to treatment and confidence in the treatment recommendation. The multi-objective utility function prioritizes candidate treatments based on effectiveness, safety, adherence, patient preference, cost and burden of treatment, and an eligibility indicator removes contraindicated treatments. Feature attribution and counterfactual explanations (using SHAP values) augment the transparency of the recommendation process.

A limitation of the study is the use of assumed patient profiles, assumed probability densities, manually assigned weights for optimizations, and illustrative results. Does not provide real life data, clinical studies, software implementation, hardware validation or clinician evaluation. Thus, the reported values will not be used to define clinical effectiveness, and the structure is not mature enough for clinical use. The method should be evaluated with disease specific clinical data sets; The decision weights should be optimised by expert consensus and data-driven learning; Fairness and robustness of the method should be evaluated across patient sub-groups and compared with existing treatment-recommendation models. Additionally, a functional software prototype needs to be created, a prospective clinical evaluation should be performed, privacy-preserving data governance mechanisms should be established, as well as research on practical quantum or quantum-classical computing platforms.

REFERENCES

- [1] Acosta, J. N., Falcone, G. J., Rajpurkar, P., & Topol, E. J. (2022). Multimodal biomedical AI. *Nature medicine*, 28(9), 1773-1784.
- [2] Amann, Julia, et al. "Explainability for artificial intelligence in healthcare: a multidisciplinary perspective." *BMC medical informatics and decision making* 20.1 (2020): 310.
- [3] Ashley, Euan A. "Towards precision medicine." *Nature Reviews Genetics* 17.9 (2016): 507-522.
- [4] Biamonte, Jacob, et al. "Quantum machine learning." *Nature* 549.7671 (2017): 195-202.
- [5] Cao, Yudong, Jhonathan Romero, and Alán Aspuru-Guzik. "Potential of quantum computing for drug discovery." *IBM Journal of Research and Development* 62.6 (2018): 6-1.
- [6] Collins, Francis S., and Harold Varmus. "A new initiative on precision medicine." *New England journal of medicine* 372.9 (2015): 793-795.
- [7] Guidotti, Riccardo, et al. "A survey of methods for explaining black box models." *ACM computing surveys (CSUR)* 51.5 (2018): 1-42.
- [8] Huang, Shih-Cheng, et al. "Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines." *NPJ digital medicine* 3.1 (2020): 136.
- [9] Johnson, Kevin B., et al. "Precision medicine, AI, and the future of personalized health care." *Clinical and translational science* 14.1 (2021): 86-93.
- [10] Lundberg, Scott M., et al. "From local explanations to global understanding with explainable AI for trees." *Nature machine intelligence* 2.1 (2020): 56-67.
- [11] Rajkomar, Alvin, Jeffrey Dean, and Isaac Kohane. "Machine learning in medicine." *New England Journal of Medicine* 380.14 (2019): 1347-1358.
- [12] Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. "" Why should i trust you?" Explaining the predictions of any classifier." *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2016.
- [13] Tjoa, Erico, and Cuntai Guan. "A survey on explainable artificial intelligence (xai): Toward medical xai." *IEEE transactions on neural networks and learning systems* 32.11 (2020): 4793-4813.
- [14] Topol, Eric J. "High-performance medicine: the convergence of human and artificial intelligence." *Nature medicine* 25.1 (2019): 44-56.

- [15] Yu, Chao, et al. "Reinforcement learning in healthcare: A survey." ACM Computing Surveys (CSUR) 55.1 (2021): 1-36.