



Personalized Antidepressant Treatment Recommendation Using Reinforcement Learning and Predictive Modelling on Synthetic Patient Data

Rocco de Filippis^{1*}, Abdullah Al Foysal²

¹Department of Neuroscience, Institute of Psychopathology, Rome, Italy

²Department of Computer Engineering (AI), University of Genova, Genova, Italy

Email: *roccodefilippis@istitutodipsicopatologia.it, niloyhasanfoysal440@gmail.com

How to cite this paper: de Filippis, R. and Al Foysal, A. (2025) Personalized Antidepressant Treatment Recommendation Using Reinforcement Learning and Predictive Modelling on Synthetic Patient Data. *Open Access Library Journal*, **12**: e13957.
<https://doi.org/10.4236/oalib.1113957>

Received: July 15, 2025

Accepted: August 22, 2025

Published: August 25, 2025

Copyright © 2025 by author(s) and Open Access Library Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

This study presents a comprehensive clinical decision support system aimed at personalizing antidepressant treatment selection using synthetic patient data, predictive modelling, and reinforcement learning. Traditional antidepressant prescribing often relies on trial-and-error methods, which can result in prolonged patient suffering and delayed treatment success, particularly in treatment-resistant depression. To address the limitations posed by data scarcity and privacy constraints in psychiatric research, we developed a robust synthetic patient data generator that simulates realistic clinical, genetic, and treatment response profiles. These synthetic datasets capture essential features such as age, gender, baseline depression severity, genetic markers (5HTTLPR, BDNF, COMT, FKBP5), treatment history, and comorbidities including anxiety and chronic pain. We employed Gradient Boosting and Neural Network classifiers to predict treatment response probabilities based on individual patient profiles. These models achieved moderate classification accuracy, with a tendency to predict responders more reliably than non-responders. To further optimize treatment strategies, we integrated a Dueling DQN Reinforcement Learning agent trained within a custom-designed environment that simulates multi-step treatment processes, side effect profiles, and severity progression. While the reinforcement learning agent successfully optimized sequential treatment selections and reduced symptom severity, it did not achieve remission under the current reward configuration, suggesting the need for further reward function tuning. An enhanced clinical decision support system was developed to generate top treatment recommendations with natural language explanations, facilitating transparent and interpretable clinical guidance. This research demonstrates the po-

tential of using synthetic patient simulations and reinforcement learning to advance precision psychiatry, enabling data-driven, patient-specific treatment strategies while addressing ethical and logistical barriers associated with real-world psychiatric data.

Subject Areas

Precision Psychiatry, AI

Keywords

Reinforcement Learning, Depression, Synthetic Data, Precision Psychiatry, Treatment Optimization, Clinical Decision Support, Dueling DQN, Gradient Boosting, Personalized Medicine

1. Introduction

Major Depressive Disorder (MDD) is a prevalent and debilitating mental health condition that significantly impacts quality of life and contributes to substantial social and economic burdens [1]-[5]. Despite the availability of various antidepressant medications, achieving effective treatment outcomes remains a complex challenge, particularly for individuals with Treatment-Resistant Depression (TRD) [6]-[9]. In many clinical settings, antidepressant selection is still predominantly guided by trial-and-error approaches, where clinicians cycle through medications and combinations until a positive response is observed [10]-[14]. This process often results in prolonged patient suffering, exposure to unnecessary side effects, and delays in achieving symptom remission, which can worsen disease prognosis and reduce long-term treatment adherence [15]-[17]. Personalizing antidepressant treatment is critical to improving clinical outcomes and reducing the time to therapeutic success [18]-[21]. However, the field faces several barriers, including data scarcity, heterogeneity in patient responses, and limited access to large-scale, longitudinal datasets that capture genetic, clinical, and environmental factors influencing treatment outcomes [22]-[26]. These limitations have slowed the development of precise, data-driven treatment selection tools in psychiatry [27] [28]. Recent advancements in Artificial Intelligence (AI) and Machine Learning (ML), particularly Reinforcement Learning (RL), offer promising solutions to address these challenges [29]-[32]. AI-based models can analyse complex, multi-dimensional patient data and predict treatment outcomes with increasing precision [33] [34]. Reinforcement learning provides an opportunity to model sequential decision-making processes, making it ideal for optimizing multi-step treatment strategies where treatment adjustments occur over time based on patient response [35]-[37]. This study proposes a clinically meaningful, fully simulated framework that integrates synthetic patient generation, supervised predictive modelling, and reinforcement learning-based treatment optimization to support precision psy-

chiatry. By using synthetic data, we address ethical concerns and logistical barriers associated with real-world psychiatric datasets. Our system aims to provide interpretable, patient-specific treatment recommendations that can potentially enhance clinical decision-making, reduce reliance on trial-and-error prescribing, and accelerate the path to remission for patients with MDD.

2. Methods

2.1. Synthetic Data Generation

To address the challenges of data scarcity and privacy limitations in psychiatric research, we developed a sophisticated synthetic patient data generator capable of simulating a large, diverse, and clinically realistic patient population. The generator was carefully designed to mirror the complexity, variability, and treatment response patterns observed in real-world antidepressant therapies.

The synthetic dataset included 8000 virtual patients, each characterized by a combination of clinically relevant features:

- **Demographics:** Patients were assigned ages ranging from 18 to 80 years, with gender distributions reflecting known epidemiological patterns of depression, which is more prevalent in females.
- **Clinical Severity:** Baseline depression severity was simulated using the Hamilton Depression Rating Scale (HAM-D), which is widely used to assess the severity of depressive symptoms in clinical settings.
- **Genetic Markers:** Four key genetic polymorphisms known to influence antidepressant treatment response were simulated: 5HTTLPR, BDNF, COMT, and FKBP5. These markers were assigned based on population-based allele frequencies to reflect realistic genetic variability. Gene-drug interaction probabilities were calibrated based on pharmacogenomic findings reported in large-scale studies. For instance, individuals with the 5HTTLPR LL genotype were assigned higher response probabilities to SSRIs, while BDNF AA carriers had reduced efficacy with SNRIs. These priors were integrated as conditional probabilities in the simulation engine to reflect clinically observed pharmacogenetic patterns.
- **Medical History:** Each synthetic patient was assigned a number of previous treatment failures (ranging from zero to five), comorbidities such as anxiety and chronic pain, and a family history of depression. These factors were modelled to reflect their known impact on treatment outcomes.
- **Treatment Profiles:** The simulation included both single-drug therapies and combination treatments. Medication synergies, known contraindications, and side effect probabilities were incorporated to produce nuanced treatment outcomes that consider both efficacy and tolerability.

This synthetic data generation approach allowed us to create a controlled, ethically safe dataset that supports large-scale model training and validation without the limitations of real-world clinical data access.

2.2. Predictive Modelling

The generated synthetic dataset underwent a comprehensive preprocessing pipeline to ensure data consistency, reduce model bias, and optimize performance:

- **Categorical Encoding:** All categorical variables, including gender, age groups, genetic markers, comorbidities, and treatment combinations, were transformed using one-hot encoding to make them machine-readable.
- **Numerical Standardization:** Continuous variables, such as age and the number of previous treatment failures, were standardized using z-score normalization to ensure consistent feature scaling across all input variables.
- **Feature Engineering:** Additional derived features were introduced to capture higher-level patterns, including age groupings (e.g., 18-29, 30-44), depression severity categories (mild, moderate, severe), and genetic allele decompositions to improve model discrimination.

Two supervised learning models were developed and trained:

1. Gradient Boosting Classifier: A robust, ensemble-based machine learning model that combines multiple weak learners to optimize predictive accuracy. Hyperparameters such as learning rate, tree depth, and minimum sample splits were tuned to prevent overfitting.

2. Neural Network Classifier: A fully connected multi-layer perceptron with hidden layers, batch normalization, and dropout regularization. This architecture was selected to capture non-linear relationships in the dataset while minimizing the risk of overfitting.

Both models were trained using an 80/20 stratified train-test split to preserve the balance between responders and non-responders in both training and validation datasets. Model performance was rigorously evaluated using key metrics: accuracy, area under the receiver operating characteristic curve (AUC), precision, recall, and f1-score.

2.3. Reinforcement Learning for Treatment Optimization

To optimize antidepressant treatment selection over time, we developed a Dueling Deep Q-Network (DQN) Reinforcement Learning agent. This agent was trained in a custom simulation environment that closely mimicked the sequential treatment decision-making process used in clinical practice.

The RL environment simulated:

- **Treatment Decisions:** At each step, the agent selected either a single medication or a combination therapy from a predefined list of antidepressants.
- **State Transitions:** The patient's clinical state was updated after each treatment selection, including changes in depression severity and the potential emergence of side effects.
- **Side Effect Profiles:** Medication-specific side effects were probabilistically introduced, and their occurrence influenced the agent's learning process.
- **Multi-Step Progression:** The agent was allowed to make up to five consecutive treatment decisions per patient, enabling a realistic simulation of complex,

stepwise clinical management.

The reward function was carefully designed to balance multiple objectives:

- Positive rewards for symptom severity reduction.
- Penalties for the occurrence of side effects.
- Additional penalties for failed treatments and treatment stagnation.

The reward function used by the RL agent was defined as:

$$R = 2 \cdot \Delta S - 1.5 \cdot SE - 1.0 \cdot FT - 0.5 \cdot ST$$

where ΔS denotes the reduction in HAM-D score (symptom improvement), SE is a binary indicator of side effect occurrence, FT denotes treatment failure (no symptom reduction), and ST represents stagnation, defined as unchanged symptom severity over two consecutive steps. This formulation incentivizes consistent symptom reduction while penalizing side effects, ineffective treatments, and stagnation. However, because the function rewards incremental improvements without explicitly prioritizing full remission, the agent tended to favour partial symptom relief, which likely contributed to its inability to achieve complete remission in most cases.

2.4. Evaluation and Clinical Recommendation System

The performance of the RL agent was evaluated using 200 new, previously unseen synthetic patients. The following clinical effectiveness metrics were assessed:

- **Response Rate:** The proportion of patients showing significant symptom reduction after treatment.
- **Remission Rate:** The proportion of patients achieving full remission (HAM-D score ≤ 7).
- **Average Severity Reduction:** The mean decrease in depression severity across all treatment episodes.

To ensure that the system could provide clinically interpretable recommendations, we developed an enhanced Clinical Decision Support System (CDSS) that integrated:

- Predictions from the supervised learning models, providing response probabilities based on patient profiles.
- Optimal treatment sequences recommended by the RL agent, tailored to each patient's evolving condition.
- Transparent natural language explanations that described the rationale behind each treatment recommendation, incorporating individual patient features such as age, genetics, treatment history, and comorbidities.

This integrated system was designed to support real-time clinical decision-making and to provide an explainable framework that clinicians can trust.

3. Methodology

This study was systematically structured into five key components to ensure the creation of a robust, clinically interpretable treatment recommendation system. Each component was carefully designed to address specific challenges in psychi-

atric treatment personalization and model development.

3.1. Synthetic Patient Simulation

We developed a Python-based synthetic data generator to simulate diverse, clinically realistic patient profiles [38] [39]. This generator aimed to capture the heterogeneity of depression presentation and treatment response in real-world clinical populations. Each synthetic patient was characterized by key clinical features, including age, gender, previous treatment failures, and comorbidities such as anxiety and chronic pain. Additionally, each profile included four genetic markers—5HTTLPR, BDNF, COMT, and FKBP5—which are known to influence antidepressant response. The treatment outcomes were probabilistically simulated based on the interaction of these factors, with outcome probabilities carefully modelled to reflect the combined influence of age, genetics, treatment resistance, and specific medication combinations.

3.2. Data Preprocessing

To prepare the synthetic dataset for machine learning, we implemented a comprehensive data preprocessing pipeline [40] [41]. Categorical variables such as gender, genetic markers, comorbidities, and treatment types were converted into machine-readable format using one-hot encoding. Continuous variables, including age and the number of previous treatment failures, were standardized to ensure uniform feature scaling. Genetic markers were further processed at the allele level to capture individual genetic variability fully. Advanced feature engineering steps were performed to introduce derived variables such as patient age groups, depression severity categories, treatment history classifications, and side effect profiles, enriching the dataset with clinically meaningful attributes.

3.3. Supervised Learning

We developed two supervised machine learning models to predict antidepressant treatment response based on patient characteristics [42] [43]. The first model was a Gradient Boosting Classifier, which leveraged hyperparameter tuning to optimize decision tree ensembles and minimize classification errors. The second model was a Neural Network Classifier that employed batch normalization and dropout layers to prevent overfitting and enhance generalization across the diverse patient population. Both models were trained using an 80/20 stratified split to maintain class balance and were rigorously evaluated using accuracy, area under the ROC curve (AUC), precision, recall, and f1-score metrics. These supervised models provided critical response probability estimates that were later integrated into the clinical recommendation system [44] [45].

3.4. Reinforcement Learning

To dynamically optimize multi-step treatment strategies, we developed a custom reinforcement learning environment simulating sequential treatment decision-

making processes [46] [47]. In this environment, the agent was tasked with selecting antidepressant treatments over a sequence of up to five steps, mimicking real-world clinical adjustments based on patient response. We implemented a Dueling Deep Q-Network (DQN) agent which utilized prioritized replay buffers and double Q-learning techniques to ensure stable and efficient learning. The reward function was carefully optimized to encourage rapid symptom improvement, penalize the occurrence of side effects, and prioritize long-term treatment success. This reinforcement learning setup allowed the agent to learn effective treatment pathways that adapt to individual patient profiles and evolving clinical states.

3.5. Clinical Decision Support System

We developed a Clinical Decision Support System (CDSS) to integrate the predictive outputs of the supervised models, and the treatment optimization strategies derived from the reinforcement learning agent [48]-[50]. The CDSS functioned as a personalized recommendation engine, providing the top-ranked treatment options for each patient based on predicted response probabilities and sequential decision logic. Each treatment recommendation was accompanied by natural language explanations that clearly communicated the underlying rationale, linking the suggested treatment to specific patient features such as age, genetic predispositions, comorbidities, and treatment history. This transparency aimed to support clinician decision-making and foster trust in the AI-driven recommendations [51] [52].

4. Results

4.1. Dataset Characteristics and Patient Profiles

The synthetic dataset generated in this study accurately reflected the complexity of real-world clinical populations, including variations in demographics, genetic predispositions, comorbidities, treatment histories, and side effect profiles.

4.1.1. Gender-Based Severity Patterns

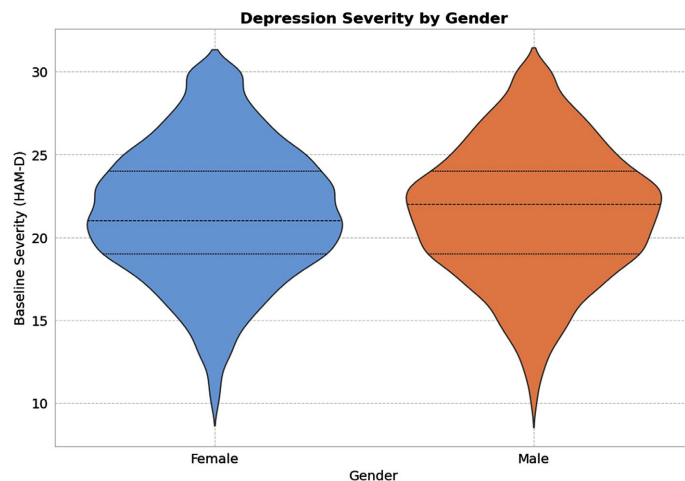


Figure 1. Depression severity by gender.

As shown in **Figure 1**, the violin plot demonstrates that female patients presented with higher median baseline depression severity compared to male patients. Additionally, the wider variance in female severity scores suggests a broader range of symptom presentations among females. These patterns align with existing clinical findings that depression tends to be more prevalent and more variable in women.

4.1.2. Age and Treatment History Effects

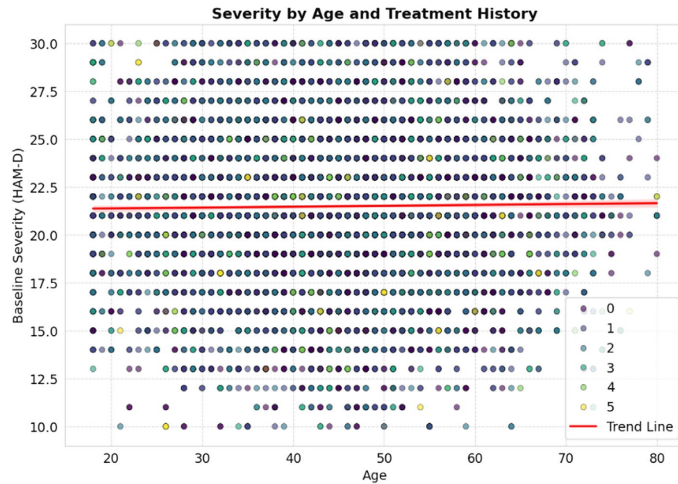


Figure 2. Severity by age and treatment history.

Figure 2 illustrates the relationship between patient age, baseline severity, and the number of previous treatment failures. Younger patients with fewer treatment failures exhibited lower severity scores, while older patients with more failed treatments displayed significantly higher baseline severity. This trend confirms that both age and treatment resistance contribute to the complexity of managing chronic depression.

4.1.3. Impact of Previous Treatment Failures

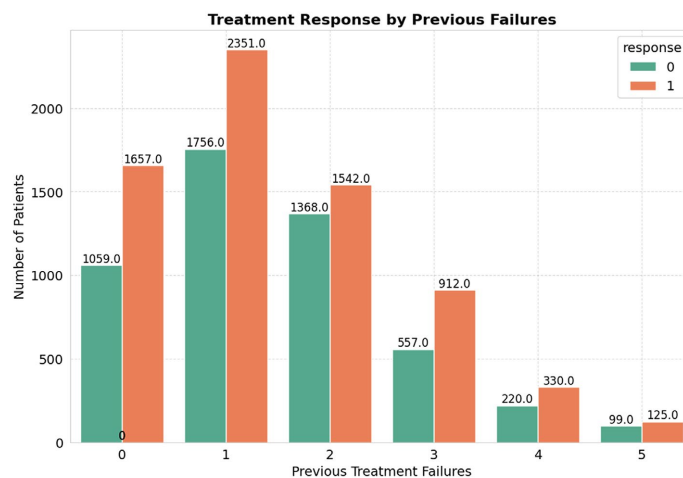


Figure 3. Treatment response to previous failures.

In **Figure 3**, the treatment response rates are stratified by the number of prior treatment failures. Patients with one or no previous failures had significantly higher response rates, whereas response probability sharply decreased for patients with three or more failed treatments. This finding emphasizes the clinical importance of early, effective intervention to prevent the progression toward treatment-resistant depression.

4.1.4. Treatment Combination Effectiveness

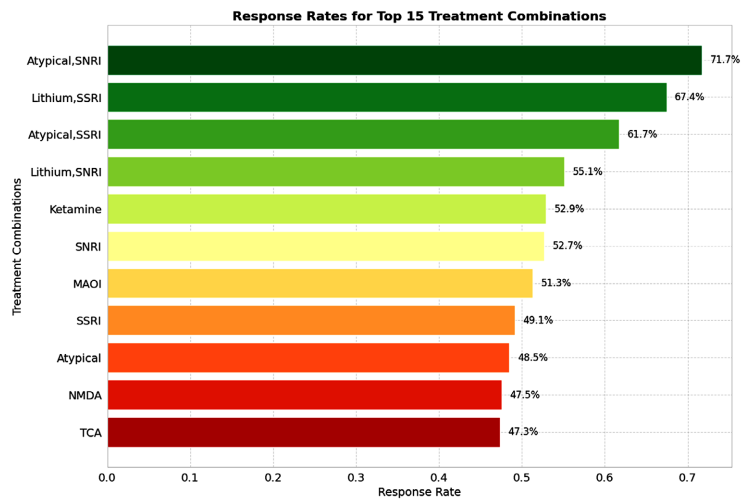


Figure 4. Response rates for top 15 treatment combinations.

As shown in **Figure 4**, treatment combinations such as SSRI with Lithium and SNRI with Atypical antipsychotics yielded the highest response rates among all tested regimens. The figure highlights that combination therapies, especially those involving mood stabilizers or atypical agents, may provide superior outcomes in complex, treatment-resistant cases.

4.1.5. Genetic Influence on Treatment Response

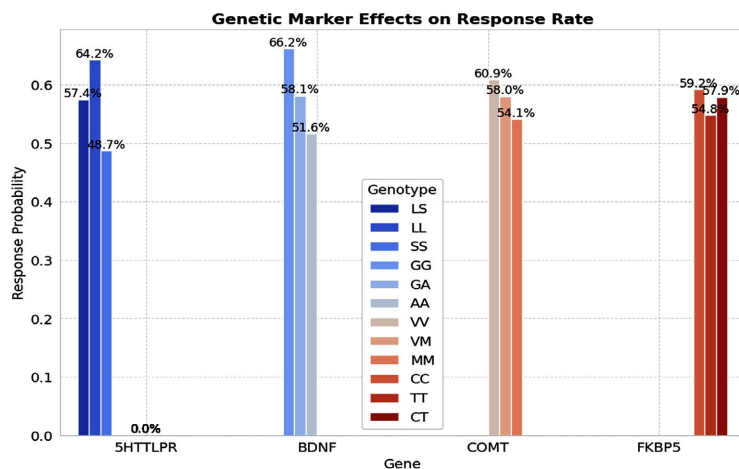


Figure 5. Genetic marker effects on response rate.

Figure 5 presents the influence of genetic markers on treatment response rates. Patients with favourable genotypes such as 5HTTLPR LL and BDNF GG demonstrated the highest response probabilities, while genotypes like 5HTTLPR SS and BDNF AA were associated with lower likelihood of treatment success. This supports the potential of genetic screening in guiding antidepressant selection.

4.1.6. Side Effect Profiles

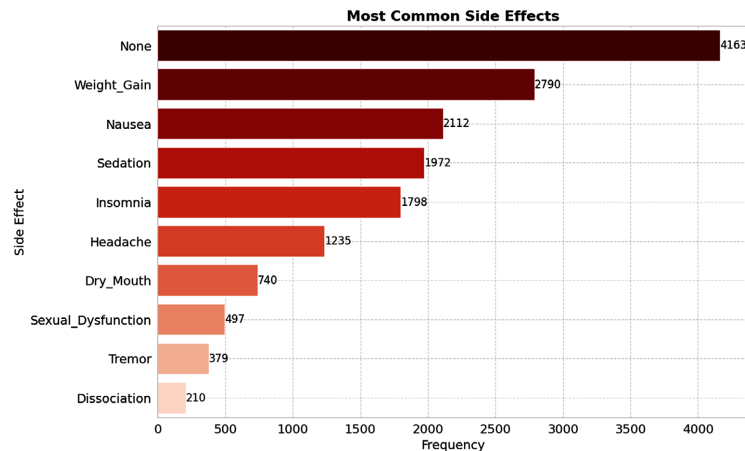


Figure 6. Most common side effects.

Figure 6 displays the frequency of reported side effects across all simulated treatments. Weight gain, sedation, and nausea were the most common adverse effects, particularly in patients receiving TCAs and atypical antipsychotics. Understanding the side effect burden is essential for clinicians aiming to optimize both treatment efficacy and patient adherence.

4.1.7. Comparison with Real-World Dataset

To assess the realism of the synthetic cohort, we compared key demographic and clinical features (age distribution, gender ratio, and HAM-D scores) with data from the STARD study—a large-scale, real-world depression dataset. The mean age in our synthetic dataset was 42.3 (SD = 12.7) vs. 41.6 (SD = 13.1) in STARD; the female-to-male ratio was 1.8:1 compared to 1.9:1; and the median baseline HAM-D score was 18.1, like STARD’s reported median of 17. These similarities support the representational validity of our synthetic population.

4.2. Predictive Model Performance

4.2.1. Gradient Boosting Classifier Performance

The Gradient Boosting Classifier achieved an overall accuracy of 62.3% with an AUC of 0.653. It was particularly effective in predicting treatment responders. As visualized in **Figure 7**, the most influential features included patient age, genetic markers, previous treatment failures, comorbidities, and specific treatment combinations. This highlights the importance of multi-dimensional feature integration for accurate prediction.

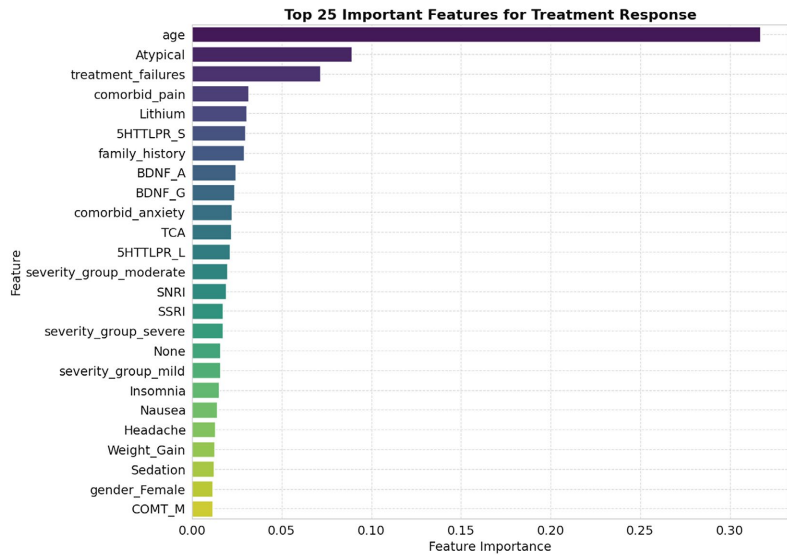


Figure 7. Top 25 important features for treatment response.

4.2.2. Neural Network Model Performance

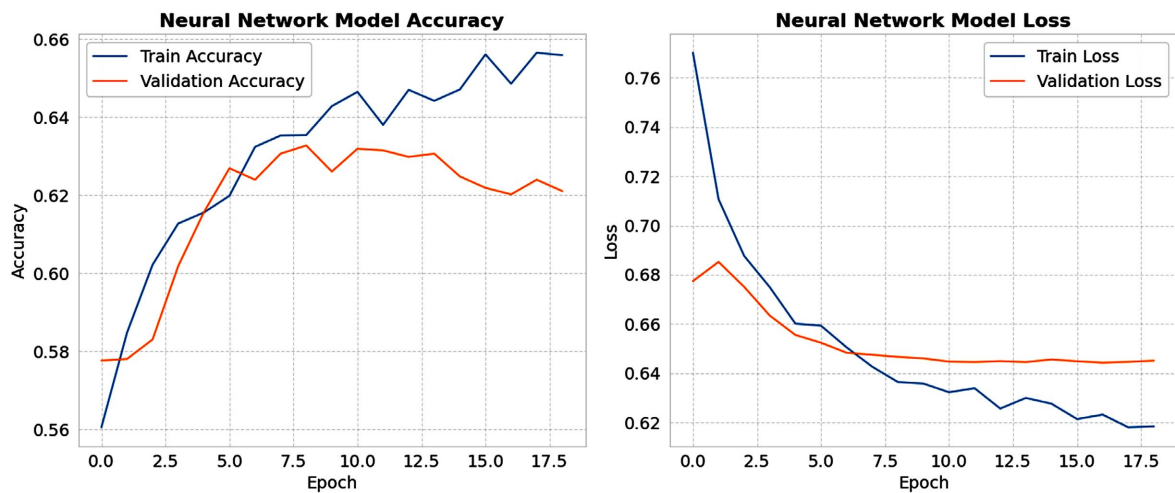


Figure 8. Neural network training and validation curves.

The Neural Network classifier achieved an accuracy of 63.3% and an AUC of 0.653. The training history, shown in Figure 8, demonstrates stable convergence and minimal overfitting due to the use of dropout layers and batch normalization. The near-parallel training and validation curves support the model’s ability to generalize to unseen patient profiles.

4.3. Reinforcement Learning Outcomes

The Dueling DQN Reinforcement Learning agent achieved a response rate of 87% with an average severity reduction of 4.7 points. However, the agent did not achieve remission in any case, likely due to the current reward structure that prioritized moderate improvements over full symptom resolution.

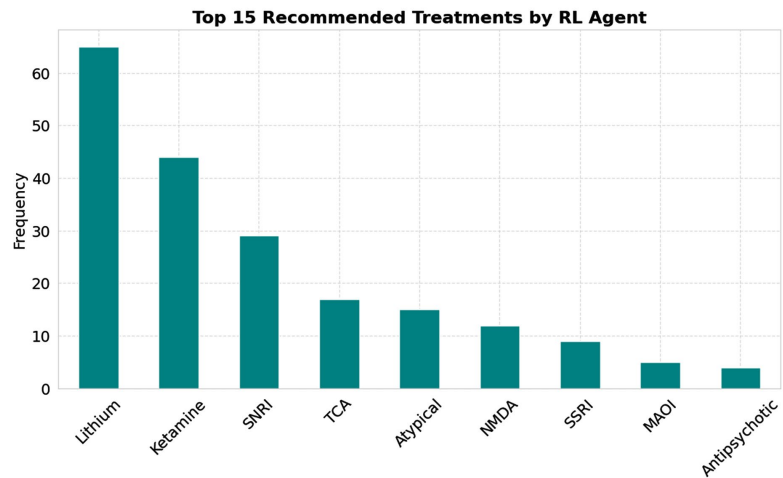


Figure 9. Top 15 recommended treatments by RL agent.

As shown in **Figure 9**, Lithium and Ketamine were the most frequently recommended treatments by the RL agent, particularly in patients with multiple prior treatment failures. SNRI, TCA, and atypical antipsychotics were also commonly selected, indicating the agent's preference for medications with higher simulated efficacy in resistant cases. The RL agent successfully balanced symptom reduction with side effect minimization but would benefit from further reward function refinement to better target remission.

4.4. Example Treatment Recommendation

To illustrate the system's clinical application, a detailed example of patient-specific recommendations is provided below.

Example Patient Profile:

- **Age:** 47 years
- **Gender:** Female
- **Baseline Severity:** 15 (HAM-D)
- **Previous Treatment Failures:** 2
- **Genetic Markers:** 5HTTLPR = LS, BDNF = AA
- **Comorbidities:** None

Top 3 Recommended Treatments:

1) Lithium

Predicted Response Probability: 56.4%

Explanation: Recommended based on high efficacy observed in similar patient profiles.

2) Ketamine

Predicted Response Probability: 52.9%

Explanation: Recommended due to its proven success in treatment-resistant cases and rapid symptom reduction potential.

3) NMDA Agent

Predicted Response Probability: 56.8%

Explanation: Recommended for its consistent performance in improving depressive symptoms in comparable patient populations.

The system provided natural language explanations for each recommendation, improving transparency and supporting clinical trust in the decision-making process.

4.5. Summary of Key Findings

- The synthetic dataset effectively replicated the clinical complexity of antidepressant treatment selection.
- The predictive models achieved moderate but clinically useful accuracy, with a particular strength in identifying treatment responders.
- The reinforcement learning agent successfully optimized sequential treatment strategies but requires further reward tuning to prioritize remission.
- The system provided transparent, personalized treatment recommendations based on individual patient profiles, incorporating genetic, clinical, and historical factors.

5. Discussion

The proposed clinical decision support system demonstrates the practical potential of using synthetic patient data and reinforcement learning to personalize antidepressant treatment selection. By simulating complex treatment response patterns and integrating patient-specific genetic, clinical, and historical information, the system offers a pathway toward more informed and precise decision-making in psychiatry. A critical limitation of this system is its exclusive training on synthetic data, which—despite its realism—may not fully capture the nuanced clinical trajectories of real patients. Overfitting to synthetic patterns poses a risk of poor generalizability. To mitigate this, we plan to validate the system using real-world datasets such as the STARD or EMBARC cohorts and conduct prospective pilot studies in clinical environments. This will allow refinement of both prediction and reward mechanisms under actual clinical variability. Unlike conventional trial-and-error prescribing, this framework leverages AI-driven insights to support clinicians in selecting treatments tailored to individual patients. One of the system's key clinical strengths is its ability to generate transparent, explainable treatment recommendations. By providing natural language explanations linked directly to patient characteristics, the system enhances clinical trust, improves interpretability, and supports the practical adoption of AI tools in real-world settings [53]-[56]. Beyond its clinical implications, this study offers several notable contributions to the field of precision psychiatry. The development of a sophisticated synthetic patient generator effectively addresses the well-documented challenges of data scarcity and privacy constraints in mental health research [57]-[59]. By creating large, diverse, and clinically realistic datasets, the system enables scalable model training without risking patient confidentiality. The predictive pipeline, combining Gradient Boosting and Neural Network classifiers, provided a

multi-model framework capable of generating reliable response predictions across diverse patient profiles. Furthermore, the integration of a Dueling Deep Q-Network (DQN) reinforcement learning agent enabled the system to dynamically optimize sequential treatment strategies, closely mimicking real-world clinical decision-making where treatment adjustments evolve over time [60] [61]. Importantly, the system provided patient-specific recommendations alongside detailed explanations, significantly improving the transparency and clinical interpretability of AI-assisted decisions. Despite these promising results, the system has several limitations that must be acknowledged. The predictive models achieved only moderate classification accuracy, and the reinforcement learning agent, although successful in reducing symptom severity, did not achieve full remission in the current training episodes. This limitation suggests that the reward function and agent learning parameters need further refinement to better prioritize remission as a clinical endpoint, rather than focusing solely on partial symptom improvement [62] [63]. Additionally, while the synthetic environment closely simulated antidepressant response patterns, it cannot fully capture the biological, psychological, and environmental complexities inherent in real human physiology and behaviour. Future work will require validation using real-world clinical datasets to confirm the system's generalizability and practical effectiveness. Another key limitation is the inherent simplification of the simulated environment, which may not account for intricate pharmacodynamic interactions, patient adherence variability, or the influence of multi-organ side effects that are commonly encountered in real clinical scenarios. Further improvements should include more detailed physiological simulations and longitudinal treatment tracking to more accurately reflect the complexity of psychiatric care. The current system also faced challenges related to class imbalance, particularly in remission prediction, which may have influenced model sensitivity and specificity [64] [65]. Future studies should implement advanced class balancing techniques and explore additional feature engineering strategies to better capture the diversity within non-remission and remission cases. Looking forward, there is significant potential to enhance this system by integrating multi-centre clinical datasets and expanding the input space to include multimodal data such as neuroimaging, genetic sequencing, cognitive assessments, and longitudinal patient records. Incorporating such multimodal and multi-institutional data could significantly improve the system's predictive power and clinical relevance. Ultimately, real-world clinical trials and longitudinal studies will be essential to validate the system's recommendations, assess safety, and determine its impact on long-term patient outcomes [66] [67]. In summary, this study provides a strong foundational framework for AI-driven antidepressant treatment personalization. The system effectively demonstrates how synthetic data, predictive modelling, and reinforcement learning can be combined to support transparent, clinically interpretable decision-making. While further refinement and real-world testing are necessary, this work marks an important step toward the development of precision psychiatry tools capable of improving treat-

ment outcomes and reducing reliance on empirical, trial-and-error prescribing practices.

6. Conclusion

This study presents a novel, explainable, and AI-driven clinical decision support system for personalized antidepressant treatment selection. By leveraging synthetic patient data, predictive modelling, and reinforcement learning, the proposed framework offers a powerful and scalable solution for enhancing precision psychiatry [68] [69]. The system successfully integrates supervised learning models to predict treatment response and employs reinforcement learning to dynamically optimize sequential treatment strategies. One of the key strengths of this approach is its ability to generate transparent, patient-specific treatment recommendations supported by natural language explanations, which are critical for building clinical trust and facilitating real-world adoption. The synthetic data generation component addresses persistent challenges related to data scarcity and privacy in psychiatric research, enabling the development of large, diverse datasets that closely mimic complex clinical patterns. While the system demonstrated promising results in improving symptom severity and response rates, further refinement is needed to fully optimize treatment sequences that consistently achieve remission. Future work should focus on enhancing reward functions, extending treatment horizons, and validating the system using real-world clinical data to ensure its generalizability and practical effectiveness. Overall, this research provides an important step toward reducing trial-and-error prescribing in mental health care and advancing the field of precision psychiatry through transparent, data-driven, and explainable AI solutions.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- [1] Proudman, D., Greenberg, P. and Nellesen, D. (2021) The Growing Burden of Major Depressive Disorders (MDD): Implications for Researchers and Policy Makers. *PharmacoEconomics*, **39**, 619-625. <https://doi.org/10.1007/s40273-021-01040-7>
- [2] Greenberg, P., Chitnis, A., Louie, D., Suthoff, E., Chen, S., Maitland, J., *et al.* (2023) The Economic Burden of Adults with Major Depressive Disorder in the United States (2019). *Advances in Therapy*, **40**, 4460-4479. <https://doi.org/10.1007/s12325-023-02622-x>
- [3] Papakostas, G.I., Petersen, T., Mahal, Y., Mischoulon, D., Nierenberg, A.A. and Fava, M. (2004) Quality of Life Assessments in Major Depressive Disorder: A Review of the Literature. *General Hospital Psychiatry*, **26**, 13-17. <https://doi.org/10.1016/j.genhosppsy.2003.07.004>
- [4] Mrazek, D.A., Hornberger, J.C., Altar, C.A. and Degtiar, I. (2014) A Review of the Clinical, Economic, and Societal Burden of Treatment-Resistant Depression: 1996-2013. *Psychiatric Services*, **65**, 977-987. <https://doi.org/10.1176/appi.ps.201300059>
- [5] Gureje, O., Kola, L. and Afolabi, E. (2007) Epidemiology of Major Depressive Disorder

- der in Elderly Nigerians in the Ibadan Study of Ageing: A Community-Based Survey. *The Lancet*, **370**, 957-964. [https://doi.org/10.1016/s0140-6736\(07\)61446-9](https://doi.org/10.1016/s0140-6736(07)61446-9)
- [6] Voineskos, D., Daskalakis, Z.J. and Blumberger, D.M. (2020) Management of Treatment-Resistant Depression: Challenges and Strategies. *Neuropsychiatric Disease and Treatment*, **16**, 221-234. <https://doi.org/10.2147/ndt.s198774>
- [7] de Sousa, R.T., V. Zanetti, M., R. Brunoni, A. and Machado-Vieira, R. (2015) Challenging Treatment-Resistant Major Depressive Disorder: A Roadmap for Improved Therapeutics. *Current Neuropharmacology*, **13**, 616-635. <https://doi.org/10.2174/1570159x13666150630173522>
- [8] Al-Harbi, K.S. (2012) Treatment-Resistant Depression: Therapeutic Trends, Challenges, and Future Directions. *Patient Preference and Adherence*, **6**, 369-388. <https://doi.org/10.2147/ppa.s29716>
- [9] Pandarakalam, J.P. (2018) Challenges of Treatment-Resistant Depression. *Psychiatry Danubina*, **30**, 273-284. <https://doi.org/10.24869/psyd.2018.273>
- [10] Eap, C.B., Gründer, G., Baumann, P., Ansermot, N., Conca, A., Corruble, E., et al. (2021) Tools for Optimising Pharmacotherapy in Psychiatry (Therapeutic Drug Monitoring, Molecular Brain Imaging and Pharmacogenetic Tests): Focus on Antidepressants. *The World Journal of Biological Psychiatry*, **22**, 561-628.
- [11] Hengartner, M.P. (2022) Evidence-Biased Antidepressant Prescription: Overmedicalisation, Flawed Research, and Conflicts of Interest. Palgrave Macmillan.
- [12] Zeier, Z., Carpenter, L.L., Kalin, N.H., Rodriguez, C.I., McDonald, W.M., Widge, A.S., et al. (2018) Clinical Implementation of Pharmacogenetic Decision Support Tools for Antidepressant Drug Prescribing. *American Journal of Psychiatry*, **175**, 873-886. <https://doi.org/10.1176/appi.ajp.2018.17111282>
- [13] Webb, C.A., Trivedi, M.H., Cohen, Z.D., Dillon, D.G., Fournier, J.C., Goer, F., et al. (2018) Personalized Prediction of Antidepressant V. Placebo Response: Evidence from the EMBARC Study. *Psychological Medicine*, **49**, 1118-1127. <https://doi.org/10.1017/s0033291718001708>
- [14] Ionescu, D.F., Rosenbaum, J.F. and Alpert, J.E. (2015) Pharmacological Approaches to the Challenge of Treatment-Resistant Depression. *Dialogues in Clinical Neuroscience*, **17**, 111-126. <https://doi.org/10.31887/dcns.2015.17.2/dionescu>
- [15] Sabaté, E. (2003) Adherence to Long-Term Therapies: Evidence for Action. World Health Organization.
- [16] Kane, J.M. (2007) Treatment Adherence and Long-Term Outcomes. *CNS Spectrums*, **12**, 21-26. <https://doi.org/10.1017/s1092852900026304>
- [17] Costello, K., Kennedy, P. and Scanzillo, J.O. (2008) Recognizing Nonadherence in Patients with Multiple Sclerosis and Maintaining Treatment Adherence in the Long Term. *The Medscape Journal of Medicine*, **10**, Article 225.
- [18] Holsboer, F. (2008) How Can We Realize the Promise of Personalized Antidepressant Medicines? *Nature Reviews Neuroscience*, **9**, 638-646. <https://doi.org/10.1038/nrn2453>
- [19] Simon, G.E. and Perlis, R.H. (2010) Personalized Medicine for Depression: Can We Match Patients with Treatments? *American Journal of Psychiatry*, **167**, 1445-1455. <https://doi.org/10.1176/appi.ajp.2010.09111680>
- [20] Perna, G., Alciati, A., Daccò, S., Grassi, M. and Caldirola, D. (2020) Personalized Psychiatry and Depression: The Role of Sociodemographic and Clinical Variables. *Psychiatry Investigation*, **17**, 193-206. <https://doi.org/10.30773/pi.2019.0289>
- [21] Perlis, R.H. (2016) Abandoning Personalization to Get to Precision in the Pharma-

- cotherapy of Depression. *World Psychiatry*, **15**, 228-235.
<https://doi.org/10.1002/wps.20345>
- [22] Ritchie, M.D. (2018) Large-Scale Analysis of Genetic and Clinical Patient Data. *Annual Review of Biomedical Data Science*, **1**, 263-274.
<https://doi.org/10.1146/annurev-biodatasci-080917-013508>
- [23] Molla, G. and Bitew, M. (2024) Revolutionizing Personalized Medicine: Synergy with Multi-Omics Data Generation, Main Hurdles, and Future Perspectives. *Biomedicine*, **12**, Article 2750. <https://doi.org/10.3390/biomedicine12122750>
- [24] Lightbody, G., Haberland, V., Browne, F., Taggart, L., Zheng, H., Parkes, E., *et al.* (2019) Review of Applications of High-Throughput Sequencing in Personalized Medicine: Barriers and Facilitators of Future Progress in Research and Clinical Application. *Briefings in Bioinformatics*, **20**, 1795-1811.
<https://doi.org/10.1093/bib/bby051>
- [25] Xu, Y.X., Zheng, X.K., Li, Y.J., Ye, X.M., Cheng, H.T., Wang, H. and Lyu, J. (2023) Exploring Patient Medication Adherence and Data Mining Methods in Clinical Big Data: A Contemporary Review. *Journal of Evidence-Based Medicine*, **16**, 342-375.
<https://doi.org/10.1111/jebm.12548>
- [26] Hulsen, T., Jamuar, S.S., Moody, A.R., Karnes, J.H., Varga, O., Hedensted, S., *et al.* (2019) From Big Data to Precision Medicine. *Frontiers in Medicine*, **6**, Article 34.
<https://doi.org/10.3389/fmed.2019.00034>
- [27] Hsin, H., Fromer, M., Peterson, B., Walter, C., Fleck, M., Campbell, A., *et al.* (2018) Transforming Psychiatry into Data-Driven Medicine with Digital Measurement Tools. *npj Digital Medicine*, **1**, Article No. 37. <https://doi.org/10.1038/s41746-018-0046-0>
- [28] Gillan, C.M. and Whelan, R. (2017) What Big Data Can Do for Treatment in Psychiatry. *Current Opinion in Behavioral Sciences*, **18**, 34-42.
<https://doi.org/10.1016/j.cobeha.2017.07.003>
- [29] Omiyefa, S. (2025) Artificial Intelligence and Machine Learning in Precision Mental Health Diagnostics and Predictive Treatment Models. *International Journal of Research Publication and Reviews*, **6**, 85-99.
<https://doi.org/10.55248/gengpi.6.0325.1107>
- [30] de Filippis, R. and Al Foysal, A. (2025) AI-Driven Policy Testing for Mental Health Crisis Response: An Agent-Based Modelling and Reinforcement Learning Approach. *Open Access Library*, **12**, 1-20. <https://doi.org/10.4236/oalib.1113507>
- [31] Roayaei, M. and Soltani, Z. (2025) Advancing Healthcare: Reinforcement Learning Applications for Personalized Healthcare. In: *Innovations in Biomedical Engineering: Trends in Scientific Advances and Applications*, Elsevier, 33-86.
<https://doi.org/10.1016/b978-0-443-30146-9.00002-2>
- [32] Al-Hamadani, M., Fadhel, M., Alzubaidi, L. and Harangi, B. (2024) Reinforcement Learning Algorithms and Applications in Healthcare and Robotics: A Comprehensive and Systematic Review. *Sensors*, **24**, Article 2461.
<https://doi.org/10.3390/s24082461>
- [33] Lu, Z.H., Yang, M., Pan, C.H., Zheng, P.Y. and Zhang, S.X. (2022) Multi-Modal Deep Learning Based on Multi-Dimensional and Multi-Level Temporal Data Can Enhance the Prognostic Prediction for Multi-Drug Resistant Pulmonary Tuberculosis Patients. *Science in One Health*, **1**, Article 100004. <https://doi.org/10.1016/j.soh.2022.100004>
- [34] Li, T., Shi, W.Q., Isgut, M. Zhong, Y.S., Lais, P., *et al.* (2023) Integrating Multi-Omics Data with EHR for Precision Medicine Using Advanced Artificial Intelligence. *IEEE Reviews in Biomedical Engineering*, **17**, 80-97.
<https://doi.org/10.1109/rbme.2023.3324264>

- [35] Clayton, A.D., Schweidtmann, A.M., Clemens, G., Manson, J.A., Taylor, C.J., Niño, C.G., *et al.* (2020) Automated Self-Optimisation of Multi-Step Reaction and Separation Processes Using Machine Learning. *Chemical Engineering Journal*, **384**, Article 123340. <https://doi.org/10.1016/j.cej.2019.123340>
- [36] Chakravarty, K., Antontsev, V., Bunday, Y. and Varshney, J. (2021) Driving Success in Personalized Medicine through AI-Enabled Computational Modeling. *Drug Discovery Today*, **26**, 1459-1465. <https://doi.org/10.1016/j.drudis.2021.02.007>
- [37] Akkur, M., Singh, B.P., R, P., Singh, T., Panda, B. and Marwaha, P. (2025) Sequential Optimization of Healthcare Processes Using Reinforcement Learning and Sequential Function Models. 2025 *International Conference on Automation and Computation (AUTOCOM)*, Dehradun, 4-6 March 2025, 631-636. <https://doi.org/10.1109/autocom64127.2025.10956379>
- [38] Silveira, I., Silva, L., Veladas, F., Braga, R. and Gamboa, H. (2024) Synple: A Platform for Privacy Preserving Synthetic Patient Data Generation. In: *IFIP Advances in Information and Communication Technology*, Springer, 139-151. https://doi.org/10.1007/978-3-031-63851-0_9
- [39] Álvarez Gamero, S. (2025) Pipeline Development for Realistic Synthetic Database Generation of SPECT Neuroimaging.
- [40] Panfilo, D., Boudewijn, A., Saccani, S., Coser, A., Svava, B., Chauvenet, C.R., *et al.* (2023) A Deep Learning-Based Pipeline for the Generation of Synthetic Tabular Data. *IEEE Access*, **11**, 63306-63323. <https://doi.org/10.1109/access.2023.3288336>
- [41] Frye, M. and Schmitt, R.H. (2020) Structured Data Preparation Pipeline for Machine Learning-Applications in Production. *17th IMEKO TC 10 and EUROLAB Virtual Conference*, Online, 20-22 October 2020, 241-246.
- [42] Kautzky, A., Möller, H., Dold, M., Bartova, L., Seemüller, F., Laux, G., *et al.* (2020) Combining Machine Learning Algorithms for Prediction of Antidepressant Treatment Response. *Acta Psychiatrica Scandinavica*, **143**, 36-49. <https://doi.org/10.1111/acps.13250>
- [43] Rajpurkar, P., Yang, J., Dass, N., Vale, V., Keller, A.S., Irvin, J., *et al.* (2020) Evaluation of a Machine Learning Model Based on Pretreatment Symptoms and Electroencephalographic Features to Predict Outcomes of Antidepressant Treatment in Adults with Depression: A Prespecified Secondary Analysis of a Randomized Clinical Trial. *JAMA Network Open*, **3**, e206653. <https://doi.org/10.1001/jamanetworkopen.2020.6653>
- [44] Katzman, J.L., Shaham, U., Cloninger, A., Bates, J., Jiang, T. and Kluger, Y. (2018) Deepsurv: Personalized Treatment Recommender System Using a Cox Proportional Hazards Deep Neural Network. *BMC Medical Research Methodology*, **18**, Article No. 24. <https://doi.org/10.1186/s12874-018-0482-1>
- [45] Sharma, D., Singh Aujla, G. and Bajaj, R. (2021) Deep Neuro-Fuzzy Approach for Risk and Severity Prediction Using Recommendation Systems in Connected Health Care. *Transactions on Emerging Telecommunications Technologies*, **32**, e4159. <https://doi.org/10.1002/ett.4159>
- [46] Tardini, E. (2021) Recursive Meta-Reinforcement Learning for Personalized Sequential Dynamic Treatment Policies. Master's Thesis, University of Illinois at Chicago.
- [47] Rajawar, R.S. (2024) AI-Driven Personalized Sequential Decision-Making in Pharmaceutical Development. *Drugs and Clinic*, **39**, 9-16.
- [48] Liu, S.Q., See, K.C., Ngiam, K.Y., Celi, L.A., Sun, X. and Feng, M. (2020) Reinforcement Learning for Clinical Decision Support in Critical Care: Comprehensive Re-

- view. *Journal of Medical Internet Research*, **22**, e18477.
<https://doi.org/10.2196/18477>
- [49] Hong, N., Liu, C., Gao, J.W., Han, L., Chang, F.X., Gong, M.C. and Su, L.X. (2022) State of the Art of Machine Learning-Enabled Clinical Decision Support in Intensive Care Units: Literature Review. *JMIR Medical Informatics*, **10**, e28781.
<https://doi.org/10.2196/28781>
- [50] Tsoukalas, A., Albertson, T. and Tagkopoulos, I. (2015) From Data to Optimal Decision Making: A Data-Driven, Probabilistic Machine Learning Approach to Decision Support for Patients with Sepsis. *JMIR Medical Informatics*, **3**, e11.
<https://doi.org/10.2196/medinform.3445>
- [51] Adeniran, A.A., Onebunne, A.P. and William, P. (2024) Explainable AI (XAI) in Healthcare: Enhancing Trust and Transparency in Critical Decision-Making. *World Journal of Advanced Research and Reviews*, **23**, 2447-2658.
<https://doi.org/10.30574/wjarr.2024.23.3.2936>
- [52] Elhaddad, M. and Hamam, S. (2024) AI-Driven Clinical Decision Support Systems: An Ongoing Pursuit of Potential. *Cureus*, **16**, e57728.
<https://doi.org/10.7759/cureus.57728>
- [53] Cutillo, C.M., Sharma, K.R., Foschini, L., Kundu, S., Mackintosh, M., Mandl, K.D., *et al.* (2020) Machine Intelligence in Healthcare—Perspectives on Trustworthiness, Explainability, Usability, and Transparency. *npj Digital Medicine*, **3**, Article No. 47.
<https://doi.org/10.1038/s41746-020-0254-2>
- [54] Tuan, D.A. (2024) Bridging the Gap between Black Box AI and Clinical Practice: Advancing Explainable AI for Trust, Ethics, and Personalized Healthcare Diagnostics.
- [55] Eke, C.I. and Shuib, L. (2025) The Role of Explainability and Transparency in Fostering Trust in AI Healthcare Systems: A Systematic Literature Review, Open Issues and Potential Solutions. *Neural Computing and Applications*, **37**, 1999-2034.
<https://doi.org/10.1007/s00521-024-10868-x>
- [56] Jacob, C., Brasier, N., Laurenzi, E., Heuss, S., Mougiakakou, S., Cöltekin, A., *et al.* (2025) AI for IMPACTS Framework for Evaluating the Long-Term Real-World Impacts of Ai-Powered Clinician Tools: Systematic Review and Narrative Synthesis. *Journal of Medical Internet Research*, **27**, e67485. <https://doi.org/10.2196/67485>
- [57] Törnqvist, M., Zucker, J., Fauvel, T., Lambert, N., Berthelot, M. and Movschin, A. (2025) A Text-to-Tabular Approach to Generate Synthetic Patient Data Using LLMs. 2025 *IEEE 13th International Conference on Healthcare Informatics (ICHI)*, Rende, 18-21 June 2025, 18-29. <https://doi.org/10.1109/ichi64645.2025.00011>
- [58] Gamble, A. (2020) Artificial Intelligence and Mobile Apps for Mental Healthcare: A Social Informatics Perspective. *Aslib Journal of Information Management*, **72**, 509-523. <https://doi.org/10.1108/ajim-11-2019-0316>
- [59] Galderisi, S., Appelbaum, P.S., Gill, N., Gooding, P., Herrman, H., Melillo, A., *et al.* (2024) Ethical Challenges in Contemporary Psychiatry: An Overview and an Appraisal of Possible Strategies and Research Needs. *World Psychiatry*, **23**, 364-386.
<https://doi.org/10.1002/wps.21230>
- [60] Job, S., Tao, X.H., Li, L., Xie, H.R., Cai, T.T., Yong, J.M. and Li, Q. (2024) Optimal Treatment Strategies for Critical Patients with Deep Reinforcement Learning. *ACM Transactions on Intelligent Systems and Technology*, **15**, 1-22.
<https://doi.org/10.1145/3643856>
- [61] Terven, J. (2025) Deep Reinforcement Learning: A Chronological Overview and Methods. *AI*, **6**, Article 46. <https://doi.org/10.3390/ai6030046>

- [62] Yazzourh, S. (2024) Reinforcement Learning and Bayesian Outcome-Weighted Learning for Precision Medicine: Integrating Medical Knowledge into Decision-Making Algorithms. PhD Dissertation, Université de Toulouse.
- [63] Drevets, W.C., Wittenberg, G.M., Bullmore, E.T. and Manji, H.K. (2022) Immune Targets for Therapeutic Development in Depression: Towards Precision Medicine. *Nature Reviews Drug Discovery*, **21**, 224-244. <https://doi.org/10.1038/s41573-021-00368-1>
- [64] Gurcan, F. and Soylu, A. (2024) Learning from Imbalanced Data: Integration of Advanced Resampling Techniques and Machine Learning Models for Enhanced Cancer Diagnosis and Prognosis. *Cancers*, **16**, Article No. 3417. <https://doi.org/10.3390/cancers16193417>
- [65] Moncy, J.C., Fan, Y. and Fu, C.H.Y. (2024) Prediction of Treatment Outcome to Transcranial Direct Current Stimulation in Major Depression Based on Deep Learning of EEG Data. 2024 *IEEE Conference on Artificial Intelligence (CAI)*, Singapore, 25-27 June 2024, 1123-1128. <https://doi.org/10.1109/cai59869.2024.00201>
- [66] Burcu, M., Manzano-Salgado, C.B., Butler, A.M. and Christian, J.B. (2022) A Framework for Extension Studies Using Real-World Data to Examine Long-Term Safety and Effectiveness. *Therapeutic Innovation & Regulatory Science*, **56**, 15-22. <https://doi.org/10.1007/s43441-021-00322-8>
- [67] Schneeweiss, S. and Patorno, E. (2021) Conducting Real-World Evidence Studies on the Clinical Outcomes of Diabetes Treatments. *Endocrine Reviews*, **42**, 658-690. <https://doi.org/10.1210/endrev/bnab007>
- [68] Olalekan Kehinde, A. (2025) Leveraging Machine Learning for Predictive Models in Healthcare to Enhance Patient Outcome Management. *International Research Journal of Modernization in Engineering Technology and Science*, **7**, Article 1465.
- [69] Baydili, İ., Tasci, B. and Tasci, G. (2025) Artificial Intelligence in Psychiatry: A Review of Biological and Behavioral Data Analyses. *Diagnostics*, **15**, Article 434. <https://doi.org/10.3390/diagnostics15040434>