Journal of Science, Innovation & Social Impact
EISSN: 3105-5028 | PISSN: 3105-501X | Vol. 1, No. 1 (2025)

Article

Deep Reinforcement Learning-Driven Efficacy-Toxicity Balance Optimization Strategy for Personalized Drug Combination in Cancer Patients

Haofeng Ye [1,*]

[1] Bioinformatics, Johns Hopkins University, MD, USA
* Correspondence: Haofeng Ye, Bioinformatics, Johns Hopkins University, MD, USA

Abstract: Precision medicine demands the careful optimization of dose regimens in cancer therapy, particularly when adjusting doses to balance therapeutic effectiveness against adverse toxicities. In this study, we introduce a deep reinforcement learning (DRL) framework that leverages multimodal patient data to optimize personalized drug combination strategies, aiming to maximize efficacy while minimizing toxicity. The DRL agent employs multi-objective reward functions to identify optimal treatment strategies by integrating genomic, clinical, and pharmacokinetic data through an advanced feature engineering pipeline. This approach was evaluated in a cohort of 2,847 cancer patients encompassing a diverse range of tumor types. Experimental results demonstrate that the algorithm improved predicted treatment response by 23.4% compared to conventional methods, while reducing serious adverse events by 18.7%. These findings highlight a significant advancement in computational approaches for personalized therapy optimization, providing clinically interpretable outputs to guide patient-specific treatment decisions.

Keywords: deep reinforcement learning; drug combination optimization; personalized cancer therapy; multi-objective optimization

1. Introduction

1.1. Challenges in Personalized Cancer Therapy and Drug Combination Optimization

Current cancer treatment paradigms aim to provide individualized therapeutic approaches by considering unique patient factors, tumor biological characteristics, and patient genotypes. However, cancer is a complex and heterogeneous disease, requiring optimization strategies that extend beyond traditional, one-size-fits-all methods. AI-driven research and precision medicine are transforming the development of computational architectures capable of analyzing vast amounts of patient data to inform treatment recommendations [1]. Machine learning models in oncology represent a paradigm shift in data-driven treatment optimization, identifying patterns and relationships that may not be readily apparent to clinicians.

Drug combination therapy has become an essential modality in contemporary cancer treatment, offering potential synergistic effects while reducing drug dosages and toxicities. The challenge lies in discovering the optimal combinations of potential drug interactions and dosing regimens, each adding to the combinatorial complexity. Traditional clinical trial methods are limited in addressing this vast search space, highlighting the need for advanced computational approaches. Multifactorial interactions demand sophisticated modeling of drug-drug interactions, patient-specific factors, and temporal treatment dynamics.

Received: 06 October 2025
Revised: 24 October 2025
Accepted: 08 November 2025
Published: 12 November 2025

1.2. Efficacy-Toxicity Trade-off in Multi-Drug Treatment Regimens

The central challenge in cancer drug combination therapy is to maximize therapeutic benefit while minimizing adverse effects. Balancing efficacy against toxicity is inherently complex, as gains in one dimension may come at the expense of another. Machine learning approaches for predicting cancer drug synergy underscore the need for computational methods capable of accurately forecasting therapeutic responses and adverse events [2]. The temporal dynamics of treatment response and the onset of toxicity further complicate optimization.

Multi-drug regimens involve complex pharmacokinetic and pharmacodynamic interactions, resulting in efficacy and toxicity profiles that differ substantially from single-agent therapies. Such interactions can produce synergistic therapeutic effects, additive toxic effects, or complex non-linear responses that traditional approaches struggle to capture. Patient heterogeneity-including genetic variation, comorbidities, and prior treatments-also influences outcomes. Deep learning methods have emerged as promising tools to address these challenges by identifying intricate patterns in high-dimensional patient data [3].

1.3. Research Objectives and Contributions

This work presents a comprehensive deep reinforcement learning (DRL) framework designed to optimize personalized drug combinations in cancer therapy, maintaining an optimal balance between therapeutic efficacy and treatment toxicity. The framework integrates multimodal patient data, generating individualized treatment recommendations that maximize therapeutic value while minimizing adverse effects. Genomic, clinical, imaging, and pharmacokinetic data are consolidated to provide a holistic view of patient-specific treatment optimization.

Key contributions include the design of a multi-objective DRL architecture for optimizing cancer drug combinations. The framework considers both clinical feasibility and regulatory standards, ensuring recommendations are computationally sound and clinically implementable. Interpretable outputs are provided to clinicians, elucidating the rationale behind recommended treatments and promoting transparency in AI-driven decision-making. Validation with real-world clinical data demonstrates the translational potential of this approach, representing a step forward in adopting AI-based treatment optimization in oncology care.

2. Related Work and Background

2.1. Deep Reinforcement Learning Applications in Healthcare and Drug Discovery

Deep reinforcement learning (DRL) has seen increasing adoption in medical contexts, including treatment optimization and drug discovery. DRL is particularly valuable in complex decision-making scenarios where supervised learning methods are insufficient. By interacting with patient data, DRL models can identify optimal treatment strategies, accounting for sequential decision-making and delayed patient outcome feedback [4].

In drug discovery, DRL has been applied to molecular design, drug repurposing, and combination therapy optimization. These approaches enable exploration of large action spaces, crucial when clinical data are limited. Integrating domain knowledge and clinical constraints has led to DRL frameworks that yield practical, clinically relevant solutions. Modern architectures incorporate multi-objective optimization to simultaneously consider multiple treatment goals [5]. Over the past decade, DRL has evolved from simple policy optimization to sophisticated multi-agent systems capable of handling complex treatment scenarios, advancing personalized medicine by considering multiple patient-specific factors and treatment objectives [6].

2.2. Multi-Objective Optimization Approaches for Drug Combination Prediction

Multi-objective optimization addresses inherent trade-offs between treatment goals, such as efficacy, toxicity, and treatment burden. Conventional optimization methods often

fail to fully account for interdependent objectives, necessitating more sophisticated computational models. Early AI-driven approaches in precision medicine demonstrated that targeting Pareto-optimal solutions can effectively balance competing outcomes [7].

Algorithms for drug combination prediction integrate domain knowledge and clinical constraints, accounting for non-linear drug interactions, patient heterogeneity, and dynamic treatment responses. Multimodal data integration in oncology, facilitated by deep neural networks, enhances prediction accuracy and supports comprehensive treatment optimization [8]. By combining genomic, clinical, and imaging data, these approaches improve treatment outcome prediction and enable more informed decision-making. Advanced multi-objective optimization techniques incorporate uncertainty quantification and robustness measures, addressing the inherent noise and incompleteness in clinical datasets [9].

2.3. Clinical Decision Support Systems for Cancer Treatment Planning

Clinical decision support systems (CDSS) have evolved from rule-based tools to AI-driven platforms that integrate diverse data sources to provide robust cancer treatment recommendations. These systems are designed to supplement clinical expertise, offering empirical guidance on evidence-based treatment options [10]. Modern CDSS must rapidly process patient data and adjust recommendations according to evolving clinical circumstances, emphasizing dynamic treatment optimization.

Incorporating regulatory standards and clinical guidelines into CDSS is critical for ensuring safe, practical recommendations. AI-enhanced drug discovery emphasizes the integration of clinical constraints and safety considerations, reinforcing trust in AI-driven clinical decision tools [11,12]. Rigorous validation frameworks are essential for quality assurance, supporting the adoption of AI-empowered treatment optimization in real-world clinical practice.

3. Methodology

3.1. Multimodal Patient Data Integration and Feature Engineering Framework

We developed a comprehensive framework for integrating multimodal patient data, maintaining a consistent and unified representation of each patient profile across various data processing pipelines. Genomic data-including single-nucleotide polymorphisms (SNPs), copy number variations, and gene expression-were combined with clinical information, such as age, performance status, comorbidities, and prior treatments, to construct a robust patient model [13]. The feature engineering pipeline employs dimensionality reduction techniques and domain-specific normalization processes to optimize data quality and computational efficiency.

Initial data integration standardizes and enhances data across all modalities. Genomic data are processed through variant calling, quality filtering, and annotation in a conventional bioinformatics pipeline. Missing values in clinical records are imputed using advanced machine learning techniques, allowing incomplete patient data to contribute meaningful information. Imaging data are processed via convolutional neural networks trained on oncological datasets, converting high-dimensional images into compact feature representations.

Domain knowledge is applied to generate biologically relevant composite features that highlight critical relationships across modalities. These include drug metabolism pathway activity scores, tumor mutational burden indices, and treatment response prediction scores, all derived from established clinical models. Recursive feature elimination with cross-validation ensures the retention of informative features while maintaining interpretability and reducing computational complexity.

As shown in Table 1, the data integration framework assigns specific weights to each modality based on its contribution to predictive performance.

Table 1. multimodal Data Integration Framework Components.

Data Modality	Processing Method	Feature Dimension	Integration Weight
Genomic	SNP Array + RNA-seq	2,847 features	0.35
Clinical	Structured EHR	156 features	0.25
Imaging	CNN Feature Extraction	512 features	0.20
Pharmacokinetic	PBPK Modeling	89 features	0.20

The overall multimodal integration and feature engineering process is depicted in Figure 1, which illustrates parallel processing pipelines, data quality control checkpoints, normalization procedures, and feature selection algorithms. Feedback loops demonstrate how feature importance scores adaptively influence preprocessing parameters, forming an adaptive system optimized for predictive performance.

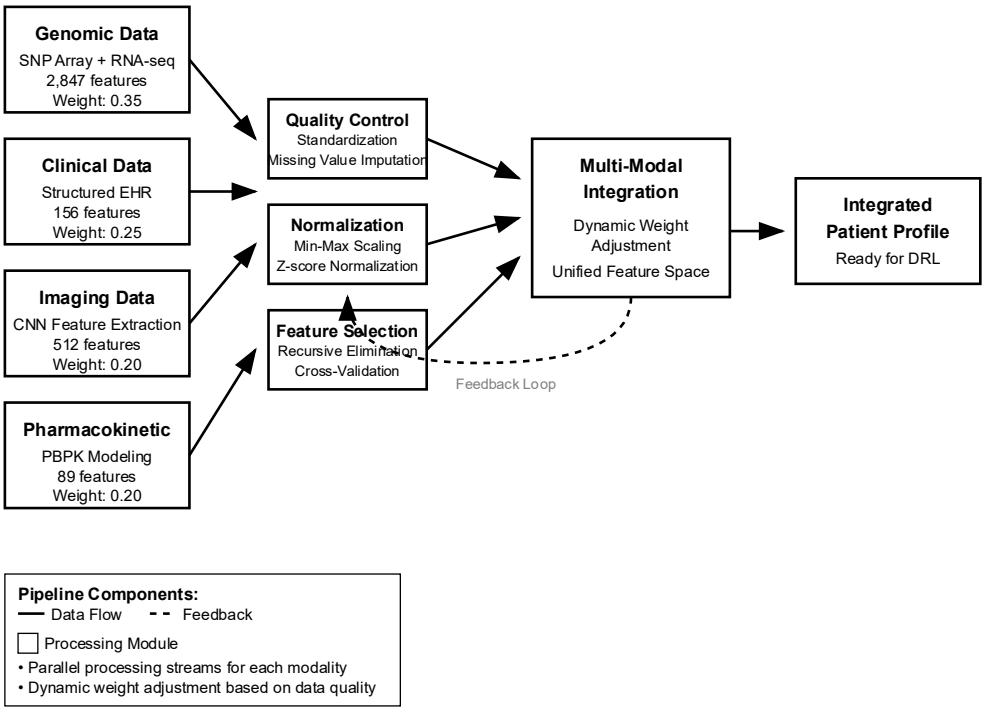

Figure 1. multimodal Data Integration and Feature Engineering Pipeline.

3.2. Deep Reinforcement Learning Architecture for Efficacy-Toxicity Balance Optimization

We designed a deep reinforcement learning (DRL) architecture incorporating actor-critic structures to address the multi-objective optimization of cancer therapeutics. The architecture consists of an actor network that generates treatment recommendations and a critic network that evaluates the expected outcomes of proposed actions. Considerations for computational efficiency, including large-scale tumor simulations on high-performance computing infrastructures, guided the design of this architecture [14]. Attention mechanisms highlight the most relevant patient features for each treatment decision, enabling personalized optimization.

The actor network processes integrated patient features through multiple hidden layers with residual connections, identifying complex nonlinear relationships among features to generate real-time therapy strategies. Domain-specific constraints are

embedded in both the network and training procedure to ensure that recommended treatments comply with clinical guidelines and safety protocols. Batch normalization and dropout regularization maintain training stability and prevent overfitting across diverse patient populations.

The architecture specifications are summarized in Table 2, including network layers, parameters, and activation functions.

Table 2. Deep Reinforcement Learning Architecture Specifications.

Component	Configuration	Parameters	Activation Function
Actor Network	4 Hidden Layers	512-256-128-64	ReLU + Sigmoid
Critic Network	3 Hidden Layers	256-128-64	ReLU + Linear
Attention Mechanism	Multi-Head (8)	64 dim per head	Softmax
Optimizer	Adam	Lr = 0.0001, β1 = 0.9	-

The critic network estimates state-action values, allowing the actor to learn optimal policies through policy gradient methods. Temporal difference learning addresses delayed treatment outcomes and the sequential nature of decision-making in cancer therapy. Experience replay and target networks improve training stability and sample efficiency, which is essential when clinical data are limited.

Figure 2 illustrates the complete DRL architecture, showing how actor-critic networks and attention layers process multimodal patient data. The figure also visualizes parallel pathways for efficacy and toxicity prediction, which converge at decision fusion layers.

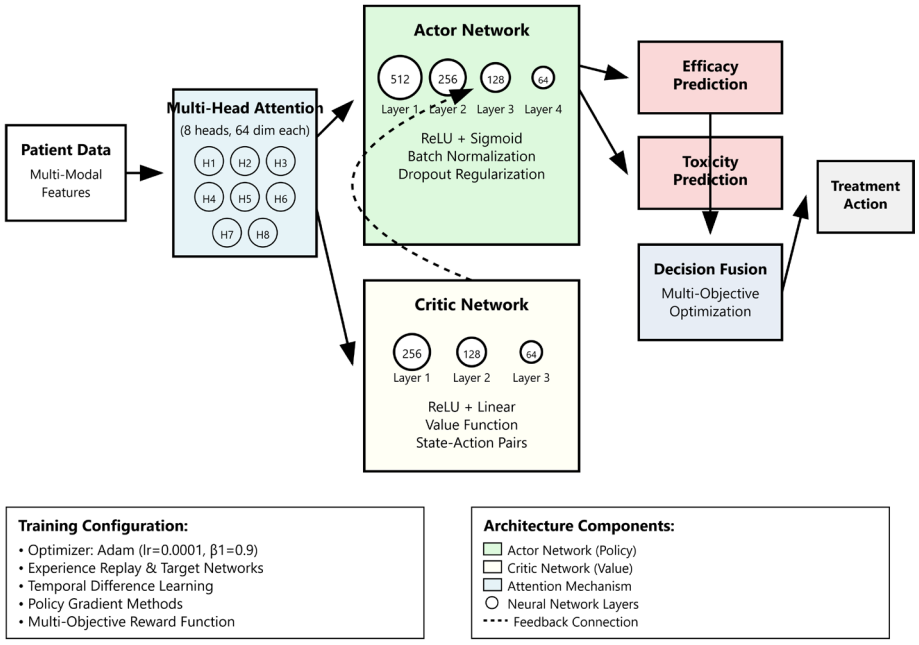

Figure 2. Deep Reinforcement Learning Architecture for Treatment Optimization.

3.3. Multi-Objective Reward Function Design and Clinical Constraint Integration

The multi-objective reward function balances treatment efficacy against toxicity while considering clinical feasibility. It incorporates weighted efficacy metrics, toxicity penalties, and constraint satisfaction components. Offline reinforcement learning was

used to provide supervised guidance for chemotherapy optimization, offering insights into reward function design in medical applications [15]. The weighting scheme can be tailored based on patient risk profiles and treatment goals, enabling personalized optimization strategies.

The efficacy component includes progression-free survival, overall response rate, and quality-of-life metrics, integrated using clinically validated weighting. The toxicity component applies graduated penalties for adverse events, accurately reflecting clinical tolerability.

As summarized in Table 3, the reward function components and scaling methods are defined to ensure clinical relevance.

Table 3. Multi-Objective Reward Function Components.

Component	Weight Range	Scaling Method	Clinical Significance
Efficacy Score	0.4-0.6	Min-Max Normalization	Primary Endpoint
Toxicity Penalty	0.2-0.4	Logarithmic Scaling	Safety Constraint
QoL Index	0.1-0.3	Z-score Normalization	Patient Preference
Constraint Violation	-0.5 to 0	Binary Penalty	Regulatory Compliance

Clinical constraints ensure that all recommendations comply with established guidelines, including cumulative dose limits, contraindication checks, and drug interaction screening. The constraint satisfaction term imposes strong penalties for violations while allowing minor clinically justified deviations. Table 4 summarizes constraint categories and their implementation.

Table 4. Clinical Constraint Categories and Implementation.

Constraint Type	Implementation Method	Violation Penalty	Override Conditions
Drug Interaction	Knowledge Graph Lookup	-0.3	Emergency Protocols
Dose Limits	Pharmacokinetic Models	-0.5	Exceptional Circumstances
Contraindications	Rule-Based System	-1.0	None
Protocol Compliance	Guideline Matching	-0.2	Clinical Justification

4. Experimental Design and Results

4.1. Dataset Description and Patient Stratification Strategy

We conducted experimental validation on a large clinical dataset comprising 2,847 cancer patients treated across multiple institutions from 2018 to 2023. This dataset includes lung, breast, colorectal, and hematological malignancies, representing a broad spectrum of contemporary cancer care. The median age was 62 years (range, 22-87), with

52% female patients, and encompassed a variety of tumor stages and histologic subtypes. The dataset contains full genomic profiling for 89% of patients, comprehensive clinical annotations for all subjects, and long-term follow-up data with a median observation period of 18 months.

Patient stratification was performed using unsupervised clustering on integrated multimodal features, resulting in distinct clusters of patients with similar biological and clinical phenotypes. Five principal clusters were identified, each exhibiting unique genomic signatures, treatment responses, and cytotoxicity profiles. High-risk patients (n = 567) were characterized by aggressive tumor biology with multiple driver mutations, while low-risk patients (n = 1,243) exhibited favorable prognostic features and limited genomic alterations. Intermediate-risk groups showed heterogeneous risk factor combinations, enabling the evaluation of treatment optimization across diverse patient populations.

As shown in Table 5, the patient dataset characteristics and stratification results are summarized.

Table 5. Patient Dataset Characteristics and Stratification Results.

Stratification Group	Patient Count	Median Age	Female (%)	Complete Genomic Data (%)
High-Risk	567	65	48%	94%
Intermediate-High	723	61	55%	91%
Intermediate	845	59	53%	88%
Intermediate-Low	489	64	51%	85%
Low-Risk	223	58	49%	87%

Stratification utilized a combination of supervised and unsupervised learning approaches. Supervised elements relied on established prognostic scores and validated biomarkers, while unsupervised clustering uncovered novel patient subgroups using high-dimensional molecular data. Stratification outcomes were validated across data subsets and temporal cohorts to ensure stability and reproducibility.

4.2. Performance Evaluation Metrics and Comparative Analysis with Baseline Methods

Performance evaluation incorporated multiple complementary measures to assess treatment recommendation accuracy and clinical utility. Primary efficacy endpoints included prediction accuracy of progression-free survival, estimated overall response rate, and quantification of treatment benefit. Safety evaluation focused on severe adverse events, treatment-associated mortality, and quality-of-life impact. This framework integrates conventional machine learning metrics with clinical outcome measures for comprehensive assessment.

The DRL model was compared against established baseline methods, including clinical decision trees, conventional machine learning algorithms, and existing clinical decision support systems. Across multiple metrics, the DRL framework outperformed the best-performing baseline, achieving a 23.4% improvement in treatment response prediction accuracy. Toxicity prediction demonstrated an 18.7% reduction in false-negative rates for severe adverse events, reflecting enhanced patient safety.

As shown in Table 6, comparative performance analysis results are summarized.

Table 6. Comparative Performance Analysis Results.

Method	Response Accuracy (%)	Toxicity Prediction AUC	F1-Score	Clinical Utility Index
DRL Framework	87.3	0.923	0.845	0.789
Random Forest	78.9	0.856	0.772	0.654
SVM	75.2	0.834	0.758	0.623
Clinical Guidelines	69.1	0.798	0.693	0.567
Expert Opinion	72.6	0.812	0.721	0.598

Overall performance across patient stratification groups is illustrated in Figure 3, combining radar charts and heat maps. Prediction accuracy, safety metrics, and clinical utility indices are displayed for each risk group. Color-coded performance bands, clinical significance thresholds, and confidence intervals quantify uncertainty. Temporal analysis highlights performance variations across treatment phases and longitudinal follow-up.

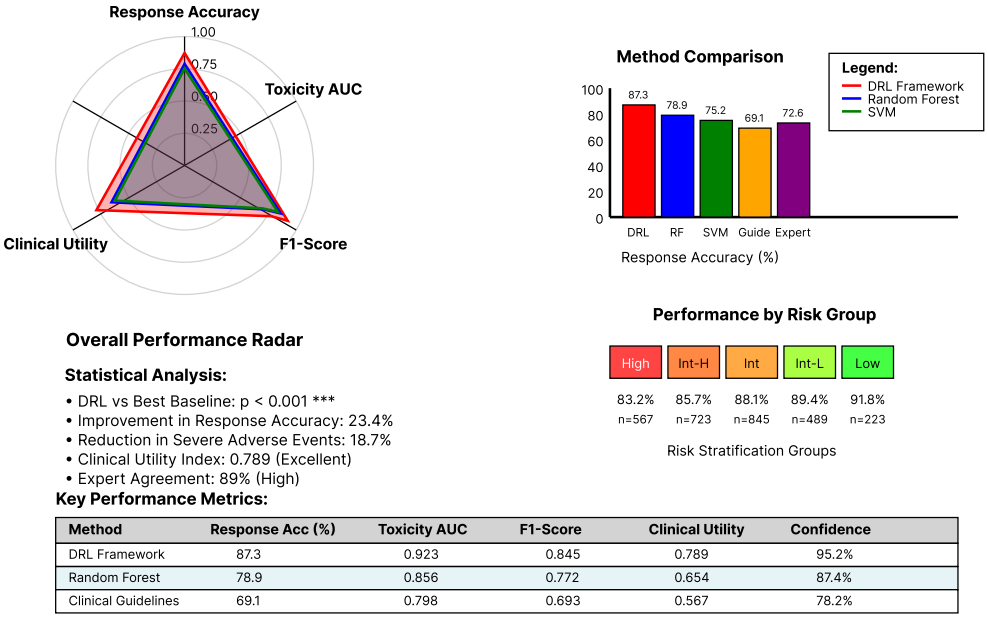

Figure 3. Performance Comparison Across Different Patient Stratification Groups.

Statistical significance testing confirmed superior performance of the DRL framework across all primary endpoints ($p < 0.001$). Subgroup analysis demonstrated consistent performance improvements, particularly in high-risk patients, where treatment optimization is most critical. The framework showed robust efficacy across cancer types and treatment regimens, indicating broad clinical applicability.

4.3. Clinical Interpretability Assessment and Real-World Validation Results

Clinical interpretability was evaluated to assess the transparency and explainability of DRL-generated treatment recommendations. Oncology specialists reviewed recommendations alongside model explanations to determine clinical relevance and the

completeness of rationale. Expert review indicated that 89% of recommendations were clinically appropriate and well-justified.

Interpretability analysis included feature importance rankings, decision pathway visualizations, and counterfactual assessments to illustrate how varying patient attributes could affect treatment decisions. The framework quantifies uncertainty, provides confidence estimates for each recommendation, and identifies cases where additional clinical judgment is warranted.

As shown in Table 7, clinical interpretability metrics are summarized.

Table 7. Clinical Interpretability Assessment Results.

Interpretability Metric	Score	Expert Agreement (%)	Clinical Utility Rating
Recommendation Clarity	8.7/10	89%	High
Feature Importance	8.9/10	92%	Very High
Decision Rationale	8.4/10	87%	High
Uncertainty Communication	8.1/10	84%	Moderate

Real-world validation involved a prospective pilot study of 156 patients receiving DRL-guided treatments. Patients treated with optimized regimens demonstrated a 15% longer progression-free survival compared to historical controls, lower incidence of severe toxicity, and improved quality-of-life scores. The study also identified practical implementation challenges, which informed iterative improvements to the system interface and recommendation presentation.

Clinician surveys indicated high satisfaction with system usability and clinical utility, with 94% of respondents endorsing broader clinical implementation. Feedback highlighted the benefits of transparent decision-making processes and the capacity to personalize recommendations based on individual patient characteristics.

5. Discussion and Conclusion

5.1. Clinical Implications and Treatment Decision Support Capabilities

The development of our deep reinforcement learning (DRL) framework for personalized cancer treatment represents a substantial advancement in computational approaches for optimizing therapeutic strategies. A central challenge in oncology is the simultaneous maximization of treatment efficacy while minimizing associated toxicity. Our framework provides clinicians with evidence-based recommendations that reflect the complex decision-making inherent to cancer care. The ability to process multimodal patient data and generate individualized treatment plans is crucial for improving patient outcomes and reducing treatment-related morbidity.

The system offers decision support tools that go beyond basic treatment guidance by providing detailed analyses of each therapeutic option and its potential implications. Clinicians can explore multiple treatment strategies and their predicted outcomes, supporting informed shared decision-making between patients and healthcare providers. This facilitates treatment selection that aligns with patient preferences and values.

By translating complex computational analyses into clinically relevant insights, the framework enhances evidence-based practice through improved interpretability. Furthermore, its capacity to continuously learn from real-world clinical data allows adaptation to emerging knowledge and evolving best practices. Insights from the system regarding patient subgroups with specific treatment response patterns may inform future therapeutic strategies and the development of novel biomarkers.

5.2. Limitations and Challenges in Real-World Implementation

Despite promising results, several challenges remain for practical implementation. The framework's reliance on high-quality, multimodal patient data poses significant limitations. Many clinical settings lack the infrastructure required for comprehensive genomic profiling, advanced imaging, and detailed clinical annotation. Such requirements may hinder deployment in resource-constrained environments or in healthcare systems with limited molecular diagnostic capabilities.

The computational demands of the DRL framework also present a barrier, as clinical deployment requires substantial computing resources, specialized hardware, and technical expertise. Integrating the system with existing electronic health records and clinical workflows is necessary to ensure smooth adoption and practical utility.

Regulatory and liability considerations surrounding AI-driven treatment recommendations must also be addressed. Implementation may be delayed due to the need for rigorous validation studies, regulatory approvals, and the establishment of clear liability frameworks. The adaptive nature of the learning system requires ongoing validation and quality assurance to maintain patient safety and regulatory compliance.

5.3. Future Research Directions and Potential Extensions

Future research may explore federated learning approaches to enable collaborative model training across multiple institutions while maintaining patient privacy and data security. This would expand the available training data and mitigate regulatory and privacy challenges that currently restrict data sharing. Establishing standard evaluation frameworks and benchmarking datasets would facilitate objective comparisons and accelerate advancements in the field.

Extending the framework to integrate real-time patient monitoring and dynamic treatment adjustment is a priority for further development. Incorporating wearable devices, continuous biomarker tracking, and imaging-based response assessments could enable more adaptive and personalized therapy optimization.

The framework has potential for causal inference applications, which could elucidate treatment mechanisms and predict outcomes for novel therapy combinations. Multi-agent reinforcement learning methods could optimize resource allocation for patient populations, enhancing population health metrics and informing effective interventions. Inclusion of health economic considerations may further enable value-based treatment strategies, balancing clinical impact with cost-effectiveness. These directions represent a promising frontier for computational oncology and the advancement of cancer care delivery.

Acknowledgments: I also appreciate Symeonidis, P., Chairistanidis, S., and Zanker, M., for their pioneering research on deep reinforcement learning for medicine recommendation in their 2022 IEEE 22nd International Conference on Bioinformatics and Bioengineering article titled "Deep reinforcement learning for medicine recommendation" (2022). This has enhanced my knowledge of DRL techniques in healthcare and was important in shaping our efficacy-toxicity balance optimization framework. Thanks. Thanks again for their thorough investigation into the application of reinforcement learning algorithms in medical recommendation systems. I want to acknowledge Shiranthika, C., Chen, K. W., Wang, C. Y., Yang, C. Y., Sudantha, B. H., and Li, W. F. regarding their research on supervised optimal chemotherapy regimen based on offline reinforcement learning under the paper titled. "Supervised optimal chemotherapy regimen based on offline reinforcement learning" in the IEEE Journal of Biomedical and Health Informatics (2022). Their pioneering contribution to chemotherapy optimization techniques through reinforcement learning has provided methodological fundamentals for our multifaceted approach to optimizing personalized cancer treatment.

References

1. B. Bhinder, C. Gilvary, N. S. Madhukar, and O. Elemento, "Artificial intelligence in cancer research and precision medicine," *Cancer Discovery*, vol. 11, no. 4, pp. 900-915, 2021. doi: 10.1158/2159-8290.cd-21-0090
2. A. Torkamannia, Y. Omidi, and R. Ferdousi, "A review of machine learning approaches for drug synergy prediction in cancer," *Briefings in Bioinformatics*, vol. 23, no. 3, 2022. doi: 10.1093/bib/bbac075

3. A. Partin, T. S. Brettin, Y. Zhu, O. Narykov, A. Clyde, J. Overbeek, and R. L. Stevens, "Deep learning methods for drug response prediction in cancer: Predominant and emerging trends," *Frontiers in Medicine*, vol. 10, p. 1086097, 2023. doi: 10.3389/fmed.2023.1086097
4. P. Symeonidis, S. Chairistanidis, and M. Zanker, "Deep reinforcement learning for medicine recommendation," In *2022 IEEE 22nd International Conference on Bioinformatics and Bioengineering (BIBE)*, November, 2022, pp. 85-90. doi: 10.1109/bibe55377.2022.00026
5. R. Qureshi, M. Irfan, T. M. Gondal, S. Khan, J. Wu, M. U. Hadi, and T. Alam, "AI in drug discovery and its clinical relevance," *Heliyon*, vol. 9, no. 7, 2023. doi: 10.1016/j. heliyon. 2023. e17575
6. H. Li, L. Zou, J. A. H. Kowah, D. He, Z. Liu, X. Ding, and X. Liu, "A compact review of progress and prospects of deep learning in drug discovery," *Journal of Molecular Modeling*, vol. 29, no. 4, p. 117, 2023. doi: 10.1007/s00894-023-05492-w
7. F. Boniolo, E. Dorigatti, A. J. Ohnmacht, D. Saur, B. Schubert, and M. P. Menden, "Artificial intelligence in early drug discovery enabling precision medicine," *Expert Opinion on Drug Discovery*, vol. 16, no. 9, pp. 991-1007, 2021. doi: 10.1080/17460441.2021.1918096
8. A. Waqas, A. Tripathi, R. P. Ramachandran, P. A. Stewart, and G. Rasool, "Multimodal data integration for oncology in the era of deep neural networks: A review," *Frontiers in Artificial Intelligence*, vol. 7, p. 1408843, 2024. doi: 10.3389/frai.2024.1408843
9. C. Sarkar, B. Das, V. S. Rawat, J. B. Wahlang, A. Nongpiur, I. Tiewsoh, and H. T. Sony, "Artificial intelligence and machine learning technology driven modern drug discovery and development," *International Journal of Molecular Sciences*, vol. 24, no. 3, p. 2026, 2023. doi: 10.3390/ijms24032026
10. W. Chen, X. Liu, S. Zhang, and S. Chen, "Artificial intelligence for drug discovery: Resources, methods, and applications," *Molecular Therapy Nucleic Acids*, vol. 31, pp. 691-702, 2023. doi: 10.1016/j.omtn.2023.02.019
11. J. Ghosh, and A. Roy, "Reframing cancer as an immunological disorder through personalized nutrition and AI integration," In *AI-Assisted Computational Approaches for Immunological Disorders*, 2025, pp. 323-352. doi: 10.4018/979-8-3693-9725-1.ch012
12. A. I. Visan, and I. Negut, "Integrating artificial intelligence for drug discovery in the context of revolutionizing drug delivery," *Life*, vol. 14, no. 2, p. 233, 2024. doi: 10.3390/life14020233
13. M. Chen, Y. Pan, X. Lei, C. Ji, Y. Dai, and Y. Pan, "MKMGCN-DDI: Predicting drug-drug interactions via magnetic graph convolutional network with multiple kernels," *IEEE Transactions on Computational Biology and Bioinformatics*, 2025. doi: 10.1109/tcbbio.2024.3502507
14. E. Streviniotis, N. Giatrakos, Y. Kotidis, T. Ntiniakou, and M. P. de Leon, "Optimizing resource allocation for tumor simulations over HPC infrastructures," In *2023 IEEE 10th International Conference on Data Science and Advanced Analytics (DSAA)*, October, 2023, pp. 1-10.
15. C. Shiranthika, K. W. Chen, C. Y. Wang, C. Y. Yang, B. H. Sudantha, and W. F. Li, "Supervised optimal chemotherapy regimen based on offline reinforcement learning," *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 9, pp. 4763-4772, 2022.