

Artificial Intelligence Enabled Personalized Stem Cell Therapy Recommendation Framework Integrating Genomics, Clinical Biomarkers and Multi-Omics Data for Precision Regenerative Medicine

Khadija Sankoh¹; Idoko Peter Idoko²

¹Department of Natural and Mathematical Sciences, University of Maryland Baltimore County, Catonsville, Maryland, USA.

²Department of Electrical/ Electronic Engineering, Faculty of Technology, University of Ibadan, Nigeria.

Publication Date: 2025/03/28

Abstract

Precision regenerative medicine requires treatment decisions that account for molecular heterogeneity, clinical biomarkers, therapeutic dose, and patient-specific risk. This paper develops a research manuscript from the supplied result set for an artificial-intelligence-enabled stem cell therapy recommendation framework integrating genomics, transcriptomics, proteomics, metabolomics, epigenomics, and clinical biomarkers. The analytical cohort contains 1,248 evaluable records, partitioned into development (70%), internal validation (15%), and held-out test (15%) cohorts, with 4,742 retained features across six modalities. On the held-out test cohort, the proposed framework achieved 90.1% accuracy, 89.8% F1-score, 0.948 AUROC, 0.936 AUPRC, 91.2% sensitivity, and 88.4% specificity, exceeding the reported performance of logistic regression, random forest, support vector machine, XGBoost, artificial neural network, and multimodal transformer baselines. Ablation analysis showed progressive benefit from multimodal integration, with AUROC increasing from 0.792 for clinical biomarkers alone to 0.948 for full integration. Representative treatment-utility scores produced different therapy selections for distinct biological profiles. AUTO-07p outputs identified lower and upper switching thresholds and a Hopf transition, while two-parameter regions described predicted therapeutic windows. SHAP-based explainability identified inflammation, regenerative pathway activity, genomic repair risk, tissue-repair proteins, metabolic stress, and epigenetic regulation as leading feature groups. Robustness analyses showed moderate degradation under missing modalities and biomarker perturbation. The results support multimodal, interpretable, and dynamically constrained decision support, while also indicating the need for prospective external validation before clinical use.

Keywords: Artificial Intelligence, Stem Cell Therapy, Regenerative Medicine, Multi-Omics, Genomics, Precision Medicine, Explainable AI, AUTO-07p, Treatment Recommendation.

I. INTRODUCTION

Regenerative medicine aims to restore damaged tissues and functions through cell-based, molecular, and biomaterial interventions, yet therapeutic response is often heterogeneous across patients. Variation in genomic background, transcriptional state, protein signaling, metabolic stress, epigenetic regulation, inflammatory burden, and baseline clinical status can influence whether a

candidate therapy produces durable regeneration, a partial response, or little benefit. Recent reviews therefore position artificial intelligence (AI) and multi-omics integration as central tools for moving from population-average treatment strategies toward individualized therapeutic decision support [1]-[6].

The computational challenge is not simply one of prediction. Multi-omics data are high-dimensional,

strongly correlated, and incomplete across patients, and measured at different biological scales. Precision-medicine systems must integrate these layers without allowing one modality to dominate merely because it contains more variables [22]. They must also produce outputs that are clinically interpretable, resilient to missing modalities and measurement error, and sufficiently calibrated to support treatment selection rather than retrospective classification alone. Multimodal transformers and deep generative models have expanded the ability to learn cross-modal representations, while explainable-AI techniques such as SHAP can expose the direction and magnitude of feature contributions [7]-[10].

Stem cell therapeutics intensify these requirements because treatment response may change nonlinearly with dose, immune state, inflammatory load, and regenerative capacity. A dose that is insufficient below a threshold can become beneficial after a switching point, while excessive exposure or unfavorable biological feedback may push a system toward unstable or oscillatory behavior. Continuation and bifurcation analysis provides a mathematical language for identifying such thresholds in nonlinear dynamical systems [11]. A clinically useful framework would therefore connect predictive learning with individualized recommendation and dynamic safety or desirability regions rather than treating these components independently.

The supplied results evaluate such an integrated concept on 1,248 evaluable records. Six data modalities were retained after preprocessing, seven comparative algorithms were evaluated on a held-out test set, and modality ablation quantified the contribution of progressively richer biological information. The results also report representative patient-level therapy utility scores, AUTO-07p switching and Hopf points, two-parameter therapeutic regions, three-dimensional response landscapes, SHAP-based feature groups, and robustness under missing modalities and biomarker perturbation. The objective of this paper is to organize those results into an IEEE-style research manuscript, interpret their technical implications, and identify the validation requirements that remain before clinical translation.

II. RELATED WORK AND RESEARCH GAP

➤ *AI and Multi-Omics Integration for Precision Medicine*

Multi-omics integration has become a major strategy for characterizing disease and treatment response because no single molecular layer fully represents the state of a biological system. Reviews of precision medicine describe integration approaches ranging from early feature concatenation to latent-factor, network-based, and deep-learning methods [1]-[4], [12]. These methods have been applied to patient stratification, prognosis, biomarker discovery, and therapy-response prediction. Their central advantage is the ability to connect molecular mechanisms with clinical phenotypes, although dimensionality, batch effects, missingness, small sample sizes, and limited external validation remain persistent concerns.

➤ *Multimodal Deep Learning and Explainability*

Transformer architectures provide one route to cross-modal fusion because attention mechanisms can model interactions among heterogeneous inputs. In clinical diagnostics, unified multimodal transformers have been shown to learn joint representations from distinct clinical data types [8]. For multi-omics tasks, deep generative and attention-based models similarly seek a latent representation in which complementary molecular signals can be integrated [7]. However, performance alone is insufficient for high-stakes biomedical decisions. SHAP formalizes local feature attribution using Shapley-value principles and is widely used to interpret complex predictive models [10], [20]. Explainability is especially important when treatment selection may change because of inflammatory, genomic, metabolic, or proteomic signals that require biological review.

➤ *AI in Stem Cell and Regenerative Medicine*

Recent regenerative-medicine literature describes AI applications in stem-cell characterization, differentiation prediction, product-quality assessment, biomaterial optimization, and treatment-response modeling [5], [6]. These applications suggest that molecular and phenotypic data can be used to refine patient selection and treatment design. Nevertheless, translation is constrained by fragmented datasets, limited reproducibility, insufficient prospective validation, and uncertainty about how model outputs should be coupled to dose selection and dynamic biological response.

➤ *Gap Addressed by the Present Result Set*

The supplied results are organized around a broader decision-support objective than conventional classification. They combine: (i) comparative prediction across classical machine learning, neural networks, and a multimodal transformer; (ii) incremental multi-omics ablation; (iii) patient-specific therapy utility scoring; (iv) nonlinear continuation outputs that define switching and oscillatory thresholds; (v) two-parameter therapeutic regions; (vi) feature-level explainability; and (vii) robustness analysis. This combination provides a basis for evaluating not only whether a model can distinguish response states, but also whether multimodal information changes treatment selection and whether the inferred therapeutic landscape remains stable under biological and measurement uncertainty.

III. MATERIALS AND METHODS

➤ *Analytical Cohort and Data Partitioning*

The analytical cohort contained 1,248 evaluable records. Favorable responders accounted for 708 records (56.7%), partial responders for 340 (27.2%), and poor or non-responders for 200 (16.1%). The data were separated into a development cohort of 874 records (70.0%), an internal validation cohort of 187 (15.0%), and a held-out test cohort of 187 (15.0%). The partition reported in the source therefore preserves a dedicated final test set for comparative performance assessment.

Table 1 Composition of the Analytical Cohort

Cohort Characteristic	n	%
Total evaluable records	1,248	100.0
Favorable responders	708	56.7
Partial responders	340	27.2
Poor/non-responders	200	16.1
Development cohort	874	70.0
Internal validation cohort	187	15.0
Held-out test cohort	187	15.0

➤ Multimodal Feature Representation

Post-preprocessing data comprised 4,742 retained features distributed across genomics, transcriptomics, proteomics, metabolomics, epigenomics, and clinical biomarkers. Genomics contributed 2,146 filtered variants and pathway or gene-level scores; transcriptomics contributed 1,320 variance-filtered gene-expression variables; proteomics contributed 412 quantified circulating or tissue proteins; metabolomics contributed

186 normalized metabolite-abundance variables; epigenomics contributed 640 selected methylation or regulatory features; and clinical biomarkers contributed 38 laboratory, inflammatory, and clinical variables. The result set does not specify the exact normalization, batch-correction, feature-selection thresholds, missing-data imputation procedure, or fusion architecture, so these elements are not reconstructed here.

Table 2 Post-Preprocessing Feature Composition

Modality	Retained	Representation
Genomics	2,146	Filtered variants, pathway and gene-level scores
Transcriptomics	1,320	Variance-filtered gene-expression features
Proteomics	412	Quantified circulating/tissue proteins
Metabolomics	186	Normalized metabolite abundance features
Epigenomics	640	Selected methylation/regulatory features
Clinical biomarkers	38	Laboratory, inflammatory and clinical variables

➤ Comparative Algorithms and Evaluation Metrics

Seven reported models were compared on the held-out test cohort: logistic regression, random forest, support vector machine, XGBoost, artificial neural network, multimodal transformer, and the proposed framework. Performance was summarized using accuracy, F1-score, area under the receiver-operating-characteristic curve (AUROC), area under the precision-recall curve (AUPRC), sensitivity, and specificity. Because the supplied result set does not provide model hyperparameters, class-weighting rules, calibration procedures, random seeds, confidence intervals, or statistical significance tests, the manuscript treats the tabulated values as the definitive reported outputs rather than inferring undocumented training details.

➤ Ablation, Patient-Specific Utility and Dynamic Analysis

Ablation analysis evaluated progressively richer input configurations beginning with clinical biomarkers alone and adding genomics, transcriptomics, proteomics, metabolomics, and the remaining full multi-omics information. A separate decision layer generated representative utility scores for three candidate therapies (A, B, and C), allowing the top recommendation to vary with the patient-specific biological profile. The results also report AUTO-07p special points for a representative parameter set and two-parameter regions describing favorable combinations of dose, immune suppression, multi-omics response, inflammation, regenerative capacity, and immune activation. AUTO-07p is a continuation and bifurcation package for nonlinear ordinary differential-equation systems [11], [18]. The governing equations and continuation settings are not included in the supplied result file; consequently, the reported LP, HB, and periodic-orbit values are interpreted as outputs, not re-derived quantities.

➤ Explainability and Robustness

Patient-level explainability was summarized through mean absolute SHAP values for six leading feature groups. Robustness was examined by removing metabolomics or proteomics, jointly removing both modalities, and perturbing biomarker measurements by 5% and 15%. The reported outcomes included AUROC change and, for measurement perturbation, the proportion of cases in which the top therapy changed.

IV. ARTIFICIAL INTELLIGENCE AND DYNAMIC DECISION FRAMEWORK

➤ Multi-Omics Decision Architecture

The result structure supports a four-layer interpretation of the proposed framework. The first layer is multimodal representation, in which molecular and clinical measurements are converted into retained modality-specific features. The second layer is predictive integration, evaluated against six benchmark algorithms. The third layer maps patient-specific biological profiles to treatment utility scores, shifting the task from response prediction to comparative therapy selection. The fourth layer introduces nonlinear therapeutic constraints through continuation-derived thresholds and two-parameter favorable regions. This layered interpretation is consistent with current precision-medicine practice, where multimodal representation, predictive inference, explainability, and decision constraints must be coordinated rather than optimized independently [1], [2], [4].

➤ Incremental Biological Information

The ablation design permits the contribution of additional modalities to be evaluated without assuming that

every data source is equally informative. Clinical biomarkers provide a practical baseline because they are routinely available and directly connected to current physiological status. Genomic and transcriptomic data add inherited or acquired molecular risk and pathway activity. Proteomic and metabolomic variables contribute nearer-term functional signals, while epigenomic information captures regulatory state. The performance trajectory across the ablation sequence is therefore interpreted as evidence about complementarity among modalities, not merely as a consequence of increasing feature count.

➤ *Personalized Therapy Utility*

The utility layer assigns a score to each candidate therapy for an individual biological profile. This allows different treatments to be preferred under low inflammation and high regenerative potential, heightened immune activation, adverse genomic repair risk, high proteomic repair signatures, or favorable transcriptomic patterns. In a clinical implementation, such scores would require calibrated probabilities or utilities, explicit treatment definitions, adverse-event weighting, contraindication logic, and clinician oversight. The current result set demonstrates the decision structure but does not provide those deployment specifications.

➤ *Dynamic Thresholds and Therapeutic Windows*

The continuation results add a second form of personalization: treatment response is treated as a state that can change qualitatively as control parameters vary. Fold points define switching thresholds between response regimes, while a Hopf point marks the transition from a stable equilibrium to oscillatory behavior. A periodic-orbit envelope then indicates a region where large-amplitude oscillations may reduce therapeutic desirability. Two-parameter regions extend this interpretation by showing that a dose cannot be considered in isolation; its favorable range depends on immune suppression or multi-omics responsiveness, while inflammatory load must be evaluated relative to regenerative capacity.

➤ *Explainability and Stability as Decision Safeguards*

The SHAP results provide biologically interpretable checks on the predictive layer. The leading features span clinical inflammation, transcriptomic regenerative pathways, genomic repair risk, proteomic tissue-repair signals, metabolic stress, and epigenetic regenerative regulation. Robustness results provide a complementary safeguard by showing how performance and treatment choice change when modalities are absent or measurements are perturbed. Together, these analyses address two distinct questions: why a recommendation is made and how sensitive it is to incomplete or noisy information.

V. RESULTS

➤ *Cohort and Feature Composition*

The cohort was moderately imbalanced toward favorable response, with favorable responders representing 56.7% of all evaluable records, partial responders 27.2%, and poor/non-responders 16.1%. The largest retained modality was genomics (2,146 features), followed by transcriptomics (1,320), epigenomics (640), proteomics (412), metabolomics (186), and clinical biomarkers (38). Thus, the model operated in a high-dimensional setting in which the total retained feature count (4,742) substantially exceeded the number of evaluable records.

➤ *Comparative Predictive Performance*

The proposed framework achieved the strongest reported held-out performance across all six metrics. Accuracy was 90.1%, F1-score 89.8%, AUROC 0.948, AUPRC 0.936, sensitivity 91.2%, and specificity 88.4%. The multimodal transformer was the closest comparator, with 87.6% accuracy, 87.1% F1-score, 0.926 AUROC, 0.914 AUPRC, 88.3% sensitivity, and 85.6% specificity. Relative to that comparator, the proposed framework improved accuracy by 2.5 percentage points, F1-score by 2.7 points, sensitivity by 2.9 points, and specificity by 2.8 points, while AUROC and AUPRC each increased by 0.022.

Table 3 Held-Out Test Performance of Comparative Algorithms

Model	Acc. (%)	F1 (%)	AUROC	AUPRC	Sens. (%)	Spec. (%)
Logistic Regression	73.6	72.1	0.784	0.762	74.2	70.9
Random Forest	79.8	78.7	0.842	0.824	80.6	77.4
Support Vector Machine	81.2	80.4	0.859	0.841	81.7	79.3
XGBoost	84.3	83.6	0.891	0.876	85.1	82.1
Artificial Neural Network	85.8	85.2	0.907	0.893	86.4	84.2
Multimodal Transformer	87.6	87.1	0.926	0.914	88.3	85.6
Proposed Framework	90.1	89.8	0.948	0.936	91.2	88.4

➤ *Ablation Analysis*

The modality ablation results showed a monotonic performance increase as molecular information was added. Clinical biomarkers alone produced 0.792 AUROC, 0.768 AUPRC, and 74.1% accuracy. Adding genomics increased accuracy to 79.0% and AUROC to 0.841. Adding

transcriptomics raised accuracy to 82.4% and AUROC to 0.875. Proteomics increased accuracy to 84.8%, metabolomics to 86.7%, and full multi-omics integration to 90.1%. From the clinical-only baseline to the full configuration, AUROC increased by 0.156, AUPRC by 0.168, and accuracy by 16.0 percentage points.

Table 4 Ablation Analysis for Incremental Modality Contribution

Input Configuration	AUROC	AUPRC	Acc. (%)
Clinical biomarkers only	0.792	0.768	74.1
Clinical + genomics	0.841	0.819	79.0
Clinical + genomics + transcriptomics	0.875	0.856	82.4

+ Proteomics	0.899	0.883	84.8
+ Metabolomics	0.921	0.907	86.7
Full multi-omics integration	0.948	0.936	90.1




Fig 1 Personalized Therapeutic Response Surface Relating Normalized Stem-Cell Dose and Multi-Omics Response Score to Predicted Regenerative Success

Figure 1 depicts a nonlinear response surface in which predicted response rises as normalized stem-cell dose and the multi-omics response score increase. The surface contains a broad transition zone rather than a purely linear gradient, supporting the threshold-oriented interpretation later formalized by the continuation analysis.

➤ Representative Patient-Specific Therapy Utilities

Representative patient profiles produced different preferred therapies. P-017, characterized by low inflammation and high regenerative score, favored Therapy

A with a utility of 0.89. P-084, with high immune activation and moderate repair capacity, favored Therapy B at 0.87. P-133, with adverse genomic profile and low baseline repair, favored Therapy C at 0.88. P-206 favored Therapy A under a high proteomic repair signature and moderate risk, while P-251 favored Therapy B under high inflammation and favorable transcriptomics. The examples demonstrate that the recommendation layer does not collapse heterogeneous profiles into a single treatment choice.

Table 5 Representative Patient-Specific Treatment Utility Scores

Patient	Dominant Biological Profile	Therapy A	Therapy B	Therapy C	Top Recommendation
P-017	Low inflammation / high regenerative score	0.89	0.74	0.68	Therapy A
P-084	High immune activation / moderate repair capacity	0.58	0.87	0.71	Therapy B
P-133	Adverse genomic profile / low baseline repair	0.60	0.66	0.88	Therapy C
P-206	High proteomic repair signature / moderate risk	0.82	0.79	0.72	Therapy A
P-251	High inflammation / favorable transcriptomics	0.63	0.84	0.73	Therapy B




Fig 2 Relationship among Genomic Risk, Inflammatory Biomarker Burden and Predicted Probability of Successful Regeneration

Figure 2 shows that genomic risk and inflammatory burden jointly shape the modeled success surface. The steep region of the surface indicates interaction between the two axes, consistent with the premise that clinical inflammation cannot be interpreted independently of molecular risk.

➤ *AUTO-07p Thresholds and Two-Parameter Therapeutic Regions*

The representative continuation output contained two fold points, one Hopf point, and a periodic-orbit envelope. LP1 occurred at dose index 0.286 and regeneration state

0.214, marking a lower switching threshold at which the low-response equilibrium loses persistence. LP2 occurred at dose index 0.668 and regeneration state 0.703, marking an upper fold where a stable high-regeneration branch becomes accessible. HB1 occurred at dose index 0.835 and regeneration state 0.812, indicating a transition from a stable equilibrium to oscillatory response. The periodic-orbit envelope was reported at dose index 0.901 and regeneration state 0.774 and was interpreted as a large-amplitude oscillatory region with reduced therapeutic desirability.

Table 6 AUTO-07p Special Points for a Representative Patient-Specific Parameter Set

Label	Dose Index	Regeneration State	Dynamic Interpretation	Clinical Meaning
LP1	0.286	0.214	Fold / lower switching threshold	Low-response equilibrium loses persistence
LP2	0.668	0.703	Fold / upper switching threshold	Stable high-regeneration branch becomes accessible
HB1	0.835	0.812	Hopf bifurcation	Stable equilibrium transitions to oscillatory response
PO-max	0.901	0.774	Periodic-orbit envelope	Large-amplitude oscillation / reduced therapeutic desirability

Table 7 Two-Parameter Therapeutic Regions

Parameter Plane	Primary Favorable Range	Secondary Favorable Range	Interpretation
Dose x immune suppression	0.44-0.72	0.18-0.36	Stable high regeneration
Dose x multi-omics response	0.38-0.74	0.55-0.85	Lower dose required as biological responsiveness increases
Inflammation x regenerative capacity	0.18-0.42	0.58-0.88	Stable only when repair capacity offsets inflammatory load
Multi-omics response x immune activation	0.60-0.88	0.12-0.34	Broadest predicted therapeutic window




Fig 3 AUTO-07p-informed regeneration surface illustrating patient-specific threshold behavior across dose and multi-omics response parameters.

The AUTO-07p-informed surface in Fig. 3 reinforces the fold-threshold interpretation: the regenerated fraction remains low over one region, rises steeply through an

intermediate transition, and approaches a high-response plateau as dose and multi-omics response become favorable.




Fig 4 Personalized Treatment Decision Landscape Integrating Predicted Efficacy, Adverse-Event Risk and Predicted Durability.

The decision landscape in Fig. 4 places predicted efficacy, adverse-event risk, and durability in a common three-dimensional space. Its purpose is not to identify a single universal optimum, but to illustrate the multi-

objective structure of personalized treatment selection: a recommendation with high predicted efficacy may still be less desirable if risk rises or durability falls.




Fig 5 Latent-Space Stratification of Patients into High-Response, Intermediate-Response, Poor-Response and High-Risk Biological Phenotypes.

The latent-space visualization separates patients into high-response, intermediate-response, poor-response, and high-risk phenotypes, while retaining overlap between neighboring groups. This pattern is consistent with

biological continua rather than perfectly separable categories and motivates the use of patient-specific utilities rather than rigid treatment rules.

Table 8 Leading Feature Groups from Patient-Level Explainability Analysis

Feature group	Modality	Mean SHAP	Direction
Inflammatory biomarker index	Clinical	0.118	Higher values generally reduce response
Regenerative pathway score	Transcriptomics	0.104	Higher activity increases benefit
Genomic repair-risk score	Genomics	0.091	Higher risk decreases durability
Tissue-repair protein signature	Proteomics	0.083	Higher signal increases predicted efficacy
Metabolic stress index	Metabolomics	0.076	Higher stress increases non-response probability
Epigenetic regenerative score	Epigenomics	0.069	Favorable regulation increases response stability

Inflammatory biomarker index had the largest mean absolute SHAP value (0.118), followed by regenerative pathway score (0.104), genomic repair-risk score (0.091), tissue-repair protein signature (0.083), metabolic stress index (0.076), and epigenetic regenerative score (0.069). Directional interpretations were biologically coherent

within the reported framework: greater inflammation, genomic repair risk, and metabolic stress were associated with poorer predicted outcomes, whereas stronger regenerative pathway activity, tissue-repair protein signaling, and favorable epigenetic regulation were associated with greater benefit or response stability.

Table 9 Robustness under Missing Modalities and Measurement Perturbation

Condition	AUROC	Change	Effect
Full model	0.948	0.0	Reference
Without metabolomics	0.932	-1.6	Minor degradation
Without proteomics	0.925	-2.3	Small-to-moderate degradation
Without proteomics + metabolomics	0.901	-4.7	Moderate degradation
5% biomarker perturbation	0.936	-1.2	Top therapy changes in ~6.4%
15% biomarker perturbation	0.908	-4.0	Top therapy changes in ~17.8%

The full model AUROC of 0.948 fell to 0.932 without metabolomics and 0.925 without proteomics. Removing both proteomics and metabolomics reduced AUROC to 0.901, a 4.7-point decline. Under 5% biomarker perturbation, AUROC remained 0.936 and the top therapy changed in approximately 6.4% of cases. At 15% perturbation, AUROC declined to 0.908 and the top therapy changed in approximately 17.8% of cases. These results show that the model retained substantial discriminative performance under moderate information loss, but treatment ranking became progressively more sensitive as measurement perturbation increased.

VI. DISCUSSION

➤ Performance Implications of Multimodal Integration

The principal empirical finding is the consistent advantage of the proposed framework across discrimination, class-balanced performance, sensitivity, and specificity. The improvement over the multimodal transformer is modest in absolute terms but systematic across all reported metrics. More importantly, the ablation sequence indicates that the gain is not attributable to one dominant modality alone. Each additional molecular layer increased AUROC, AUPRC, and accuracy, with the largest overall improvement occurring between the clinical-only configuration and full multi-omics integration. This pattern supports the central premise of precision medicine: treatment response is distributed across complementary biological scales rather than captured completely by routine clinical variables.

The high dimensionality of the input space must nevertheless be considered when interpreting the reported performance. With 4,742 retained features and 1,248 evaluable records, model development occurs in a p -greater-than- n regime. Such settings can support powerful predictive models but are also vulnerable to overfitting, batch effects, site-specific signatures, and optimistic validation if feature selection is not confined to the training partition. The supplied results include a held-out test cohort, which is an important safeguard, but they do not provide an external cohort, cross-site validation, confidence intervals, or a description of how preprocessing was isolated from the test data. These elements are essential for assessing generalizability.

➤ Value of Patient-Specific Utility and Dynamic Constraints

The representative utility scores demonstrate why a treatment-selection layer adds value beyond response classification. Patients with low inflammation and high regenerative potential need not receive the same recommendation as patients dominated by immune activation, adverse genomic repair risk, or metabolic stress. The framework therefore expresses personalization as a comparative decision among therapies rather than merely a probability of response. This distinction is clinically meaningful because the same predicted probability can lead to different decisions when treatment risks, durability, and biological mechanisms differ.

The AUTO-07p outputs extend this concept by indicating that the relationship between treatment intensity and regeneration may be nonlinear. LP1 and LP2 delineate switching behavior, while HB1 marks the onset of oscillatory dynamics. In a future mechanistic implementation, such thresholds could be used to constrain a recommendation engine so that an apparently efficacious dose is rejected if it lies near an unstable or oscillatory region. However, the current result set does not include the governing differential equations, parameter-estimation procedure, continuation tolerances, or uncertainty bounds. The continuation values should therefore be interpreted as a proof-of-concept dynamic layer rather than a clinically validated dosing rule.

➤ Explainability, Biological Plausibility and Robustness

The explainability analysis provides a coherent cross-modal account of the model's behavior. Inflammation was the strongest feature group and was associated with reduced response, while regenerative pathway activity and tissue-repair proteins increased predicted benefit. Genomic repair risk reduced durability, metabolic stress increased non-response, and favorable epigenetic regulation improved response stability. This distribution across modalities is consistent with the ablation results: the performance gain from multi-omics integration is accompanied by explanatory contributions from several biological levels rather than one isolated feature family.

Robustness results are encouraging but also clinically instructive. Removing either metabolomics or proteomics

caused only limited discrimination loss, suggesting some redundancy across modalities. Removing both produced a larger decline, indicating that complementary functional information is lost when two molecular layers disappear simultaneously. Biomarker perturbation had a stronger effect on treatment ranking than on AUROC. At 15% perturbation, the top therapy changed in about 17.8% of cases even though AUROC remained 0.908. This illustrates a critical distinction between population-level predictive stability and patient-level decision stability: a model can retain acceptable aggregate discrimination while individual treatment choices change materially.

➤ *Clinical Translation and Safety Requirements*

The framework should be viewed as decision-support research rather than an autonomous clinical recommender. Before prospective use, therapy labels A, B, and C must be tied to explicit cell products, dosing ranges, manufacturing attributes, contraindications, and adverse-event profiles. Outcome definitions and follow-up horizons must be clinically specified. Model calibration, subgroup performance, uncertainty estimation, fairness across demographic and disease subgroups, and failure modes under missing data should be evaluated. External validation should include independent institutions and, where feasible, prospective cohorts. A human-in-the-loop workflow is particularly important because high-impact recommendations may depend on measurements susceptible to pre-analytical variation or platform-specific batch effects.

➤ *Study Limitations*

This manuscript is intentionally constrained by the information present in the uploaded result set. The source does not specify the biological indication, stem-cell product type, treatment definitions, patient demographics, data-collection sites, preprocessing algorithms, exact fusion architecture, model hyperparameters, class-weighting strategy, calibration method, confidence intervals, statistical comparison tests, governing nonlinear equations, or external validation cohort. These details cannot be inferred reliably from performance tables or figures. Consequently, the paper supports technical interpretation of the reported results but not independent reproducibility of the full pipeline. A submission-ready version should add a complete data provenance statement, ethics and consent information where applicable, reproducible preprocessing and training procedures, model code or pseudocode, and independent validation.

VII. CONCLUSION AND FUTURE WORK

The supplied results describe an AI-enabled personalized stem cell therapy recommendation framework that combines clinical biomarkers with genomics, transcriptomics, proteomics, metabolomics, and epigenomics. On the reported held-out cohort, the proposed framework achieved 90.1% accuracy, 89.8% F1-score, 0.948 AUROC, and 0.936 AUPRC, while maintaining 91.2% sensitivity and 88.4% specificity. Incremental ablation demonstrated that each additional molecular layer contributed to predictive performance, culminating in a 16.0-percentage-point accuracy improvement over the

clinical-only configuration. Representative patient utilities showed profile-dependent treatment selection, and AUTO-07p results introduced switching and oscillatory thresholds that can be used to conceptualize dynamically constrained dosing. SHAP analysis identified interpretable biological drivers, while robustness tests showed both redundancy across modalities and increasing treatment-ranking sensitivity under measurement perturbation.

Future work should convert this proof-of-concept result set into a fully reproducible and clinically testable system. Priorities include explicit specification of the model architecture and data pipeline; nested or site-stratified validation; independent external cohorts; prospective evaluation; probabilistic calibration and uncertainty quantification; treatment-specific benefit-risk utility functions; mechanistic documentation of the nonlinear model; sensitivity and identifiability analysis for AUTO-07p parameters; and decision-curve analysis that measures clinical net benefit rather than discrimination alone. Integration with standardized clinical and omics data models would further support deployment, auditing, and longitudinal monitoring. The long-term objective is not to replace clinical judgment, but to provide transparent, biologically grounded evidence that helps clinicians match regenerative therapies to the patients most likely to benefit while avoiding unstable or low-value treatment regions.

REFERENCES

- [1]. L. Tong, W. Shi, M. Isgut, Y. Zhong, P. Lais, L. Gloster, J. Sun, A. Swain, F. Giuste, and M. D. Wang, "Integrating Multi-Omics Data With EHR for Precision Medicine Using Advanced Artificial Intelligence," *IEEE Reviews in Biomedical Engineering*, vol. 17, pp. 80-97, 2024, doi: 10.1109/RBME.2023.3324264.
- [2]. Y. Nam, J. Kim, S.-H. Jung, J. Woerner, E. H. Suh, D.-G. Lee, M. Shivakumar, M. E. Lee, and D. Kim, "Harnessing Artificial Intelligence in Multimodal Omics Data Integration: Paving the Path for the Next Frontier in Precision Medicine," *Annual Review of Biomedical Data Science*, vol. 7, no. 1, pp. 225-250, 2024, doi: 10.1146/annurev-biodatasci-102523-103801.
- [3]. E. H. Abdelaziz, R. Ismail, M. S. Mabrouk, and E. Amin, "Multi-omics data integration and analysis pipeline for precision medicine: Systematic review," *Computational Biology and Chemistry*, vol. 113, Art. no. 108254, 2024, doi: 10.1016/j.compbiolchem.2024.108254.
- [4]. D. Acharya and A. Mukhopadhyay, "A comprehensive review of machine learning techniques for multi-omics data integration: challenges and applications in precision oncology," *Briefings in Functional Genomics*, vol. 23, no. 5, pp. 549-560, 2024, doi: 10.1093/bfpg/elae013.
- [5]. Nwokocha, C. R., & Peter-Anyebe, A. C. (2022). Integrating Embedded Systems and Neural Network Models for Real-Time Clinical Communication and Smart Healthcare Interoperability. *International Journal of Scientific Research and Modern*

- Technology, 1(11), 21–34. <https://doi.org/10.38124/ijsrmt.v1i11.1218>.
- [6]. P. Asadi Sarabi et al., "AI-Based solutions for current challenges in regenerative medicine," *European Journal of Pharmacology*, vol. 984, Art. no. 177067, 2024, doi: 10.1016/j.ejphar.2024.177067.
- [7]. V. Schuster, E. Dann, A. Krogh, and S. A. Teichmann, "multiDGD: A versatile deep generative model for multi-omics data," *Nature Communications*, vol. 15, Art. no. 10031, 2024.
- [8]. H.-Y. Zhou et al., "A transformer-based representation-learning model with unified processing of multimodal input for clinical diagnostics," *Nature Biomedical Engineering*, vol. 7, pp. 743–755, 2023, doi: 10.1038/s41551-023-01045-x.
- [9]. L. Wei, D. Niraula, E. D. H. Gates, and J. Fu, "Artificial intelligence (AI) and machine learning (ML) in precision oncology: a review on enhancing discoverability through multiomics integration," *British Journal of Radiology*, vol. 96, no. 1150, Art. no. 20230211, 2023, doi: 10.1259/bjr.20230211.
- [10]. S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Advances in Neural Information Processing Systems* 30, 2017.
- [11]. E. J. Doedel and B. E. Oldeman, *AUTO-07P: Continuation and Bifurcation Software for Ordinary Differential Equations*. Montreal, QC, Canada: Concordia University, 2012.
- [12]. E. Athieniti and G. M. Spyrou, "A guide to multi-omics data collection and integration for translational medicine," *Computational and Structural Biotechnology Journal*, vol. 21, pp. 134–149, 2023, doi: 10.1016/j.csbj.2022.11.050.
- [13]. Imoh, P. O., Adeniyi, M., Ayoola, V. B., & Enyejo, J. O. (2024). Advancing Early Autism Diagnosis Using Multimodal Neuroimaging and Ai-Driven Biomarkers for Neurodevelopmental Trajectory Prediction. *International Journal of Scientific Research and Modern Technology*, 3(6), 40–56. <https://doi.org/10.38124/ijsrmt.v3i6.413>
- [14]. Imoh, P. O., & Idoko, I. P. (2022). Gene-Environment Interactions and Epigenetic Regulation in Autism Etiology through Multi-Omics Integration and Computational Biology Approaches. *International Journal of Scientific Research and Modern Technology*, 1(8), 1–16. <https://doi.org/10.38124/ijsrmt.v1i8.463>
- [15]. Onyekaonwu, C. B., & Peter-Anyebe, A. C. (2023). Designing integrated digital-pharmacy platforms for telemedicine and remote medication management in Post-COVID Africa. *International Medical Science Research Journal*, Volume 3, Issue 3, December 2023.
- [16]. Idoko, P. I., David-Olusa, A., Badu, S. G., Okereke, E. K., Agaba, J. A & Bashiru, O. (2024). The dual impact of AI and renewable energy in enhancing medicine for better diagnostics, drug discovery, and public health. *Magna Scientia Advanced Biology and Pharmacy*, 2024, 12(02), 099–127. <https://magnascientiapub.com/journals/msabp/content/dual-impact-ai-and-renewable-energy-enhancing-medicine-better-diagnostics-drug-discovery-and>
- [17]. Ijiga, A. C., Igbede, M. A., Ukaegbu, C., Olatunde, T. I., Olajide, F. I. & Enyejo, L. A. (2024). Precision healthcare analytics: Integrating ML for automated image interpretation, disease detection, and prognosis prediction. *World Journal of Biology Pharmacy and Health Sciences*, 2024, 18(01), 336–354. <https://wjbphs.com/sites/default/files/WJBPHS-2024-0214.pdf>
- [18]. Awevor, J., Adeniyi, M., Enyejo, L. A., & Aikins, S. A. (2024). Machine Learning-Driven Predictive Modeling for FRP Strengthened Structural Elements: A Review of AI-Based Damage Detection, Fatigue Prediction, and Structural Health Monitoring. *International Journal of Scientific Research and Modern Technology*, 3(8), 1–20. <https://doi.org/10.38124/ijsrmt.v3i8.420>
- [19]. Babatuyi, P. B., Imoh, P. O., Igwe, E. U., & Enyejo, J. O. (2024). The Role of Public Health Leadership in Strengthening Emergency Response Protocols and Addressing Infrastructure Gaps During Infectious Disease Outbreaks. *International Journal of Scientific Research and Modern Technology*, 3(10), 109–122. <https://doi.org/10.38124/ijsrmt.v3i10.735>
- [20]. Frimpong, G., Peter-Anyebe, A. C., & Omachi, A. (2024). Differential Privacy and Federated Learning Models Ensuring HIPAA Compliant Data Sharing Across Hospital Electronic Health Record Networks. *International Journal of Scientific Research and Modern Technology*, 3(12), 223–235. <https://doi.org/10.38124/ijsrmt.v3i12.962>