

A Machine Learning–Driven Biomathematical Model for Predicting Tumour–Immune Dynamics in Cancer Systems

Piyush Chaubey^{1*}, Suganya S², Dr Karthika Padmavathy³, Dr. Radhika Gupta⁴, Dr Shalini Kapoor⁵, Dr Abhishek Anand⁶

¹ Assistant Professor, Electrical Engineering, Government Engineering College, Bharatpur, Rajasthan, India; ² Information Technology, Easwari Engineering College, India; ³ Professor, Department of Pathology, Specialization: Pathology, Sri Lalithambigai Medical College and Hospital, Dr M.G.R Educational and Research Institute, Chennai, India; ⁴ Professor, Incharge, Sharda University, Knowledge Park, Greater Noida, India; ⁵ Professor, Department of Periodontology, SGT Dental College Hospital and Research Centre, SGT University, India; ⁶ Assistant Professor, Department of Pharmacy Practice, Teerthanker Mahaveer College of Pharmacy, TMU University, Moradabad, India

ABSTRACT

Tumour-immune interactions during immunotherapy constitute a complex, dynamic biological system characterised by pronounced inter-patient heterogeneity and sparse clinical sampling. Precision oncology is largely focused on accurately predicting the effects of therapy on immune activation, which requires integrative modelling frameworks that balance data-driven inference with mechanistic interpretability. Here, we suggest an adapted machine learning-based biomathematical model to describe the system tumour-immune response over short periods, based on pre-treatment and on-treatment melanoma tumour biopsies taken prior to and following immune checkpoint blockade. Bulk RNA-sequencing data were used to quantify changes in immune effector activity induced by therapy, measured by the cytolytic score. Regularised ElasticNet regression was used to predict changes in immune activation from paired transcriptomic differences, with patient-level cross-validation to ensure generalisability. The ElasticNet model with Yeo-Johnson target transformation achieved high out-of-fold predictive performance ($R^2 \sim 0.73$; RMSE ~ 33.4 ; MAE ~ 16.5), with strong agreement between observed and predicted immune activation changes. There was a cross-validation fold stability analysis of features that showed strong immune-based transcriptional drivers, such as genes linked to the natural killer cell and cytotoxic T-cell functions. Machine learning-derived transcriptomic immune activation scores were then superimposed onto patient-specific immune activation parameters within a simplified biomathematical tumour-immune ordinary differential equation model to provide a mechanistic understanding of therapy-induced immune stimulation without requiring long-term trajectory simulation. The derived parameters showed high inter-patient heterogeneity in immune responsiveness, as observed in biological variability. In general, the paper shows that the combination of interpretable machine learning and biomathematical modelling provides a powerful, biologically grounded approach to predicting changes in immune state and characterising the response to immunotherapy in cancer systems.

ARTICLE HISTORY

Received February 14, 2026

Accepted April 16, 2026

KEYWORDS

Tumour–Immune Dynamics,
Immunotherapy Response
Prediction,
ElasticNet Regression,
Biomathematical Modelling,
Precision Oncology

1. Introduction

1.1. Background and Context

Tumour-immune interactions are a multiscale, complex biological system in which interactions among malignant cells, immune effectors, stromal cells, and therapeutic interventions change over time. In the context of cancer immunotherapy, and immune checkpoint blockade in particular, clinical outcome is based not only on static molecular features but also on therapy-induced changes in immune activity within the tumour microenvironment. These tumour-immune interactions are, however, the key to understanding and forecasting precision oncology and enhanced patient stratification to immunotherapeutic interventions (Butner et al., 2024). Immunotherapy has shown long-term clinical benefit in subsets of patients with several different cancer types, including melanoma. However, patient responses remain highly heterogeneous and can range from deep immune stimulation and tumour regression to immune suppression or therapeutic resistance. This heterogeneity indicates inherent disparities in tumour immunogenicity, immune infiltration, and transcriptional

reprogramming (as a result of treatment). To characterise such changes in immune states, analytical models are required that can integrate molecular measurements with mechanistic models of tumour-immune interactions (Chattopadhyay et al., 2025). Mathematical and biomechanical modelling has been adopted as a principled method to study cancer systems dynamics for some time. It has enabled the formal representation of tumour growth, immune-mediated killing, and treatment effects using differential equations and similar modelling approaches. These models are interpretable and mechanistic, though they frequently rely on parameters that cannot be robustly estimated from clinical data, especially in environments with sparsely sampled longitudinal data (Dong et al., 2024). In oncological data, tumour biopsies are only sampled at several clinically significant time points, so long-term trajectories cannot be identified and instead require simplified or short-interval modelling assumptions. Recent improvements in hybrid modelling have overcome these challenges by integrating mechanistic formulations with data-driven approaches. Discrete-continuum and fractional-order tumour-immune models, for example, have shown the ability to include immunotherapy effects while also considering biological complexity (El Asraoui et al., 2025). However, such models have not been translated into patient-specific prediction because it is difficult to connect high-dimensional molecular measurements to interpretable mechanistic parameters.

In line with advances in mathematical oncology, machine learning has emerged as a potent tool in precision cancer research. Advances in regularised regression, deep learning and ensemble approaches have made it possible to predict treatment response directly from high-dimensional omics data. In immuno-oncology, machine learning has been used to predict therapeutic response and characterise immune phenotypes based on transcriptomic, radiomic and clinical features (Fountzilias et al., 2025). These approaches are good at pattern recognition and prediction, but they are often not mechanistic and thus cannot support biological interpretability or hypothesis generation (Razali, 2024).

Biomechanical and systems-level views help emphasise the importance of developing integrative approaches that relate information at the molecular level to emergent behaviour at the cancer-immune interface. Reviews of biomechanical tumour modelling point to an increasing trend of data-informed parameterisation of classical models, and bridging the gap between theory and patient-specific data (Ghahramani & Bavi, 2026). Such a convergence highlights the emergence of a new paradigm in which machine learning is not regarded as an alternative to the mechanistic modelling paradigm but as a complementary component that can guide, constrain, and personalise biomathematical model systems for cancer.

Within this changing landscape, there is a growing understanding that hybrid machine learning-biomathematical frameworks are a promising direction for the future. These frameworks typically aim to preserve the interpretability and causal structure of mechanistic models and leverage machine learning to generate robust signals from complex, high-dimensional clinical data. Theoretical and practical research in the oncology field and other biomedical fields has shown that it is possible to do this integration, especially with the parameterisation or calibration of mechanistic models (not their replacement) through machine learning (Joshi et al., 2025).

1.2. Problem Statement

Nevertheless, both purely mechanistic and purely data-oriented approaches to tumour-immune interactions still have important limitations despite great progress. Classical tumour-immune differential equation models involve assumptions about growth rates, immune activation and treatment effects, which are often inferred indirectly and/or fixed a priori. Such models have a high risk of overfitting the unobservable dynamics or using poorly constrained parameters, particularly in the case of sparse clinical sampling, such as paired pre-treatment and on-treatment biopsies, and lower their predictive power (Kumar et al., 2024). On the other hand, data-driven machine learning models may operate as black boxes, even when they achieve high predictive accuracy. Although they can determine correlations between molecular features and clinical outcomes, they often provide little information about the biological processes that activate the immune response and kill tumours. This lack of interpretability creates difficulties in clinical translation, as it makes it hard to understand the mechanisms of the model, validate it, and incorporate it with biological knowledge from existing models (Lin et al., 2025). The urgent need is therefore the joint modelling frameworks. The latter are some of the features that such frameworks ought to harness: paired molecular measurements to respond to therapy-induced changes in the immune state; regularised machine learning to be robust in high-dimensional and small-sample regimes; and data-driven outputs in interpretable biomathematical models, which are likely to reflect known tumour-immune interactions. This need is critical to the further development of predictive immuno-oncology that is still scientifically interpretable and clinically relevant (Matin & Setayeshi, 2024).

1.3. Scope and Limitations

The current research is specifically dedicated to short-term changes in the immune state that can be determined from paired pre-treatment and on-treatment tumour biopsies. Rather than attempting to reconstruct continuous long-term

tumour trajectories, the framework is designed to simulate net changes in immune effector activity over the treatment interval (in keeping with the temporal resolution of the available data). Such a scope is indicative of real-world clinical limitations but allows significant inferences about the dynamics of immune responses to therapy.

Bulk RNA sequencing data is used as primary molecular input, with immune activity summarised using the cytolytic score as a functional proxy for engagement of cytotoxic effector. While this does not address cell-type-specific interactions, it is a biologically based and widely validated measure of immune activation in the tumour microenvironment. The use of bulk transcriptomics and summary measures of the immune system is a trade-off between data accessibility and biological specificity.

From a modelling point of view, the study is in a high-dimensional, small-sample regime, which requires regularised machine learning methods and strong patient-level cross-validation to avoid overfitting. The biomathematical model is deliberately sparse, neglecting detailed long-term dynamics and capturing only essential interactions between the tumour and the immune response, as well as the effects of therapy on activating the immune system. In that way, it is not supposed to project the future tumor burden trends but to characterise patient-specific immune response dynamics in an interpretable manner throughout the period of observation during treatment.

1.4. Significance of the Study

The study is part of the developing research on the area of hybrid tumor-immune modeling that illustrates a rigorously based machine learning and biomathematical hybrid based on real clinical data. Using the paired transcriptomic measurements, the proposed framework enables the capture of within-patient transitions in immune states, directly attributable to immunotherapy, and eliminates the effects of baseline inter-patient variability.

The use of regularised ElasticNet regression enables robust prediction of immune activation changes while maintaining interpretability through sparse, stable gene coefficients. Notably, these machine learning results are not treated as goals in themselves; rather, they are systematically mapped to patient-specific parameters within a mechanistic tumour-immune model. This combination offers an easy understanding of connectivity between molecular signatures and parameters of biologically significant immune activation.

Translational: In the translational view, the framework provides a roadmap towards patient-specific immunotherapy response characterisation that is predictive and mechanistically understandable. The research will provide the best oncology methods for data-driven inference and mechanistic modelling, which can complement biological knowledge, inform hypothesis formulation, and eventually provide personalised treatment choices.

1.5. Research Objectives

The primary objectives of this study are as follows:

- To capture therapy-evoked alterations to the immune state in melanoma patients using paired pre- and on-treatment bulk RNA sequencing data and cytolytic score measurements.
- To develop and validate a regularised machine learning model that can predict the change of immune activation from paired differences in transcripts in a high-dimensional small-sample setting.
- To identify stable, immune-relevant transcriptional drivers of therapy-induced immune activation using cross-validated feature stability analysis.
- To in-vitro predict tumour-immune dynamics with interpretable, patient-specific immune activation parameters using machine learning derived immune activation signatures by mapping transcriptomic scores on patient-specific immune activation parameters, which will allow predicting tumour-immune interactions with interpretable and individualized modeling.

Together, these goals comprise a coherent framework of machine learning-based biomathematical methods that will enhance predictive modelling and mechanistic understanding of tumour-immune dynamics in cancer immunotherapy.

2. Literature Review

Mortazavi and Firoozabadi (2024) have created an integrated hybrid multiscale mathematical framework to investigate the tumour immune microenvironment (TIME), with a special focus on predicting immunotherapy outcomes. Their discrete-continuum model incorporates the processes of angiogenesis, tumour growth and vascular remodelling, as well as immune cell recruitment, in which T-cell migration, proliferation, subtype conversion and interactions with tumour cells are modelled in detail. The model was able to recreate the spatial distributions of tumour-infiltrating lymphocytes by incorporating vascular dynamics and hemodynamics, and to recapitulate several phases of tumour immunoediting. Notably, their results showed that tumour mutational burden alone is insufficient to predict immune checkpoint blockade (ICB) response, and that multiple biomarkers must be combined, an idea that unambiguously drives parameterisation initiatives using data from hybrid tumour-immune systems.

Pinhasi and Yizhak (2025) used machine learning and single-cell RNA sequencing to identify the gene and cellular signatures associated with immune checkpoint therapy response in melanoma. They came up with a predictive gene signature of treatment response, based on XGBoost and Boruta feature selection, across all cancer types, comprising 11 interpretable genes. Their application of SHAP values revealed complex, non-linear interactions among genes at the single-cell scale, underscoring the significance of interpretability in machine learning models in immuno-oncology. This work shows how stable molecular signatures derived using regularised or feature-selection-driven ML strategies can be used to gain insights into immune activation modelling.

Qin et al. (2024) have reviewed predictive biomarkers and computational models of immunotherapy efficacy, including mechanistic systems biology models and machine-learning-based data-driven models. Their analysis included PK/PD models, PDE-based tumour-immune models, quantitative systems pharmacology and agent-based models, as well as regression, ensemble learning and deep learning techniques. Importantly, the authors highlighted that hybrid frameworks that merge mechanistic understanding with machine learning have the greatest potential for translational immunotherapy prediction, supporting the arguments for integrated ML-biomathematical frameworks.

Ricker et al. (2024) provided a historical and a prospective view on the use of computational science in immuno-oncology, with a focus on the increased information on tumour-immune interactions provided by high-throughput sequencing and multi-omics data. Their assessment highlighted the shortcomings of descriptive computational studies and emphasised the use of predictive models to relate molecular immune states to clinical outcomes. One of the crucial future solutions to the issues of data integration, interpretability, and clinical translation identified by the authors is hybrid modelling approaches. One suggestion was a radiomics-based machine learning pipeline to predict immunotherapy response using CT imaging, proposed by Rundo et al. (2025). In their dopamine-inspired, neuro-modulated architecture, high predictive performance was demonstrated for metastatic cancer groups, showing the performance of deep learning in response classification. However, the model was mostly phenomenological and had low biological interpretability. This underscores an important limitation of purely data-driven approaches and the need for hybrid strategies that do not lose mechanistic relevance.

Salman et al. (2024) have conducted a systematic review of approaches related to the deep learning and mathematical modelling of glioblastoma for its detection and progression. Their comparison showed that deep learning models outperform mathematical models in classification performance, whereas mathematical models are more interpretable and have a biological foundation. The authors concluded that future cancer modelling efforts should focus on hybrid models that balance predictive power and mechanistic understanding and are very compatible with the design philosophy of biomathematical models that employ ML.

Shutaywi et al. (2025) proposed a mathematical model of tumour-immune interactions, which is based on a fractional-order dynamics of immune responses that incorporates nonlocal and memory responses in immune dynamics. Their work showed complex oscillatory and chaotic behaviour that is mediated by immune tumour coupling and the sensitivity of the results to important parameters. The model was mathematically rich but not directly integrated with molecular patient data, a characteristic of a gap that can be closed by parameterising it with ML.

Uthamacumaran et al. (2025) developed an AI-powered hybrid ecological model by combining Lotka-Volterra equations with evolutionary optimisation and reinforcement learning for predicting the dynamics of oncolytic viral therapy. Their framework had good predictive accuracy and identified biologically validated immune and inflammatory biomarkers. This study is a good example of how machine learning can be applied to infer biologically meaningful parameters into mechanistic models. It therefore provides a conceptual parallel with ML-driven immune activation parameterisation.

Zanib et al. (2025) proposed a hybrid compartmental model with artificial neural networks for the study of the dynamics of COVID-19. Although they applied this work outside the field of oncology, their work showed how machine learning can improve parameter estimation and prediction accuracy in mechanistic systems of limited observability. The methodological concepts of hybridisation and short-interval state modelling are directly applicable to the tumour-immune system when clinical sampling is sparse.

Zhang et al. (2025) proposed a deep learning architecture named TILScout, which can predict the density of tumour-infiltrating lymphocytes based on 28 different types of cancer with the use of histopathology images. Their results established strong correlations between the inferred immune infiltration, cancer stage and immune checkpoint gene expression. This study strengthens the key role of immune activation and infiltration metrics as functional biomarkers that can serve as target variables or mechanistic state proxies in tumour-immune modelling.

3. Methodology

3.1 Dataset Description and Study Design

This study employs a retrospective computational research design that combines machine learning and

biomathematical modelling to investigate tumour-immune dynamics during immunotherapy. The analysis is based on publicly available data from the Gene Expression Omnibus (GEO) database (GSE91061), titled "Molecular portraits of tumour mutational and micro-environmental sculpting by immune checkpoint blockade therapy", released on January 23, 2018. The dataset contains tumour biopsy samples from patients with melanoma treated by immune checkpoint blockage and includes gene expression profiling by high-throughput RNA sequencing (RNA-seq) of *Homo sapiens*.

The total number in the cohort is 109 RNA-seq samples from 65 patients, comprised of 51 pre-treatment and 58 on-treatment tumour biopsies. To model the dynamics of the immune response, this paper will concentrate on the group of patients who have both pre- and on-treatment samples and can assess the within-patient longitudinal dynamics of immune state changes caused by treatment. This paired design is especially appropriate for dynamic modelling, because it minimises the inter-individual biological variability and increases sensitivity to immune changes associated with treatment. The data present bulk RNA-seq data from tumour tissue, as well as curated immune phenotyping data, including a cytolytic score, a functional proxy for immune effector T cells in the tumour microenvironment. Clinical and sample metadata, such as patient identifiers and timepoint annotations, can be used to precisely match transcriptomic and immune measurements to treatment stages.

Overall, the study design is aimed at the short-interval dynamics of the immune system by leveraging discrete measurements taken before and after treatment to approximate changes in the underlying continuous tumour-immune system. This method captures the real clinical data limitations, in that dense longitudinal sampling is often not possible, but is also highly compatible with mechanistic modelling schemes.

3.2 Data Preprocessing and Feature Construction

Raw RNA-seq expression data were obtained as transcripts per million (TPM) values to ensure comparability across samples. TPM normalisation accounts for sequencing depth and gene length, making it suitable for analysis across samples. To stabilise variance and reduce the effect of extreme expression values, all TPM values were transformed using a $\log_2(1 + \text{TPM})$ transformation, defined as $\log(1 + \text{TPM})$. The identification of genes was standardised to official gene names using Ensembl or GeneID identifiers, based on a standardised human genome annotation reference. Genes without valid symbols were removed from downstream analysis to ensure biological interpretability. In instances where more than one row was mapped to the same gene symbol, expression values were deterministically aggregated to a single value using their mean, resulting in one representative expression profile for each gene. To take advantage of the paired biopsy design, differences in transcriptomic expression were calculated for both patients by subtracting baseline (pre-treatment) expression values from on-treatment expression values. This resulted in a paired-difference matrix of therapy-induced transcriptional changes, controlled for baseline inter-patient heterogeneity. Gene removal before modelling. Because of the high dimensionality of the transcriptomic space, genes with near-zero variation across samples were filtered before modelling, thereby reducing noise and computational cost. This filtering step, based on variance, is used to remove uninformative features and does not affect genes that exhibit significant biological diversity in their roles in immune activation.

3.3 Definition of Immune State Variables

Immune effector activity was measured using the cytolytic score, a composite measure based on the expression of cytotoxic effector genes such as GZMB and PRF1. This score has been extensively validated as a good surrogate for the cytotoxicity of T cells and natural killer cells within tumour tissue. It is relevant in immunotherapy response studies.

For each patient, the baseline immune state was defined as the cytolytic score of the pre-treatment biopsy and the on-treatment immune state was defined analogously using the on-treatment biopsy. The primary immune outcome of interest was the change in immune state, defined as the difference between on-treatment and pre-treatment cytolytic scores. This amount captures therapy-related immune activation/suppression and provides a discrete estimate of the immune system's dynamics during treatment.

The use of cytolytic score as a state variable is biologically justified because immune-mediated tumour control activities under checkpoint blockade are, in great part, mediated by cytotoxic effector mechanisms. Each bulk RNA-seq result, even though it cannot resolve cell-type-specific dynamics, gives a functionally significant summary of immune effector interactions in the tumour microenvironment, as the cytolytic score.

3.4 Machine Learning Model for Immune State Prediction

To model the relationship between transcriptomic changes and immune state transitions, we used ElasticNet regression, a regularised form of linear modelling well-suited to high-dimensional, small-sample biological data. ElasticNet is a combination of L1 (lasso) and L2 (ridge) penalties, which allow for both feature selection and coefficient shrinkage, resulting in better generalisation in the presence of correlated predictors. The model input was paired differences in transcriptomic data, and the target variable was Oliver's change in immune state, calculated from the cytolytic scores. All

model training and evaluation were performed using patient-level K-fold cross-validation, ensuring that each patient's data was used in only one-fold and that no information leaked from the training set to the testing set.

Because immune state changes showed a skewed, heavy-tailed distribution, a Yeo-Johnson target transformation was used in the modelling pipeline to stabilise variance while retaining the ability to model negative values. Nested cross-validation was used to optimise model hyperparameters such as regularisation strength and the L1/L2 mixing ratio, with training folds entirely held out.

Model performance was evaluated using several complementary measures, including root mean squared error (RMSE), mean absolute error (MAE), coefficient of determination (R²), and Pearson and Spearman correlation coefficients, which provided a comprehensive assessment of both accuracy and rank-order agreement.

3.5 Model Interpretability and Feature Stability

To obtain a biologically interpretable result, the ElasticNet coefficients were extracted from each cross-validation fold and aggregated into consensus effects at the gene level. For each gene, the average coefficient across folds measured its average contribution in one direction via immune activation, and the average absolute coefficient indicated its importance. Feature stability was determined by calculating the selection frequency, defined as the percentage of cross-validation folds in which a gene was selected (i.e., assigned a non-zero coefficient). Genes with high selection frequency were interpreted as robust molecular drivers rather than fold-specific artefacts. This approach allowed us to identify immune-relevant transcriptional programs, such as cytotoxic T cell markers, interferon signalling genes and natural killer cell-associated transcripts, which provide mechanistic insight into immune activation dynamics under immunotherapy.

3.6 Biomathematical Tumour–Immune Model Formulation

To incorporate machine learning results into a mechanistic model of tumour-immune interaction, we developed a biomathematical model of tumour-immune interaction using a system of ordinary differential equations (ODEs). The model describes the interaction between the tumour burden and immune effector activity, which can be identified with sparse longitudinal sampling.

Let $T_i(t)$ denote the latent tumour burden and $E_i(t)$ denote the immune effector activity for the patient i at time t . The dynamics of tumour growth are modelled using logistic dynamics, which are modulated by immune-mediated killing, while immune effector dynamics include recruitment, therapy-induced activation, decay, or exhaustion. The model is given by:

$$\begin{aligned}\frac{dT_i(t)}{dt} &= rT_i(t) \left(1 - \frac{T_i(t)}{K}\right) - \kappa E_i(t)T_i(t) \\ \frac{dE_i(t)}{dt} &= \rho \frac{T_i(t)}{h + T_i(t)} + \alpha_i u(t) - \delta E_i(t)\end{aligned}$$

Here, r is the intrinsic tumour growth rate, K is the tumour carrying capacity, and κ is a measure of the immune-mediated tumour killing efficiency. Immune effector recruitment is mediated by tumour antigenicity, with a saturated function described by a maximal rate ρ and a half-saturation constant h . The parameter δ represents immune effector decay, which is caused by natural turnover or exhaustion.

The immunotherapy input, $u(t)$ is modelled as a step function which is activated at the start of treatment. The parameter α_i represents the therapy-induced immune activation strength of the patient, deduced from machine learning outputs, and is responsible for the magnitude of immune activation under treatment.

Given that the GSE91061 dataset only includes paired pre-treatment and on-treatment biopsies, the model is interpreted over a short time interval $[t_0, t_1]$, where t_0 corresponds to the pre-treatment biopsy and t_1 to the on-treatment biopsy. As a result, the model is concerned with net state transitions rather than with detailed, continuous trajectories. This short-formulated model does not overfit the unobservable dynamics, without sacrificing biologically meaningful mechanisms of interaction.

3.7 Integration of Machine Learning Outputs into the Biomathematical Model

Machine learning outputs were incorporated into the biomathematical model via a parameter-mapping approach to link changes in transcriptomics to the dynamics of immune activation. From the ElasticNet regression model, a consensus vector of coefficients was found $\hat{\beta}$ which represented stable gene-level contributions of immune state change across the cross-validation folds.

For each patient i , paired transcriptomic differences were computed as:

$$X_{i,\Delta} = X_{i,\text{on}} - X_{i,\text{pre}},$$

where $X_{i,\text{pre}}$ and $X_{i,\text{on}}$ denote log-transformed transcriptomic profiles at baseline and on-treatment timepoints, respectively. A transcriptomic immune activation score was then defined as:

$$s_i = \sum_{j=1}^p w_j X_{i,\Delta,j},$$

where w_j denotes stability-weighted ElasticNet coefficients that account for both effect size and selection frequency across folds. The resulting score s_i was standardised across patients to obtain a normalised immune activation index z_i .

This standardised score was mapped to the patient-specific immune activation parameter α_i using a monotonic link function:

$$\alpha_i = g(\alpha_0 + \alpha_1 z_i),$$

where α_0 and α_1 are scaling parameters and $g(\cdot)$ is chosen to ensure biologically valid parameter values (e.g., enforcing positivity of activation rates). This formulation ensures that higher transcriptomic immune activation signatures correspond to stronger therapy-induced immune stimulation within the mechanistic model.

The scaling parameters α_0 and α_1 were calibrated by minimising the discrepancy between the observed immune state change, defined as the difference in cytolytic scores between pre-treatment and on-treatment biopsies, and the mechanistic immune state change predicted by the ODE model:

$$\min_{\alpha_0, \alpha_1} \sum_{i=1}^N (E_i(t_1) - E_i(t_0) - \Delta E_i^{\text{obs}})^2.$$

This endpoint-matching calibration ensures consistency between data-driven predictions and mechanistic immune dynamics while preserving interpretability. Through this integration, the proposed framework yields a machine learning–driven biomathematical model capable of capturing patient-specific tumor–immune dynamics under immunotherapy.

3.8 Ethical Considerations

All analyses were performed on de-identified, publicly available data, and no new patient data were collected. As such, the study did not need institutional review board approval. The study conforms to ethical guidelines for analysing secondary data, but is primarily concerned with transparency, reproducibility, and responsible interpretation of findings.

4. Results

4.1 Dataset Characteristics and Paired Sample Summary

After quality control and alignment of transcriptomic and immune phenotyping data, a subset of 43 melanoma patients from the GSE91061 cohort with paired pre-treatment and on-treatment tumour biopsies was retained for analysis, resulting in 86 tumour samples used for paired longitudinal modelling (Riaz et al., 2017). The paired design ensured that the downstream analyses of immune state transitions focused on within-patient variability rather than cross-sectional variability.

Immune activation was measured by the cytolytic score, and the change in the level of immune activation was defined as the difference between on-treatment and pre-treatment cytolytic scores, as originally described in the source cohort (Riaz et al., 2017). The distribution of immune activation showed substantial variation across patients, with values ranging from a negative change (immune suppression) to large positive shifts, indicating strong immune activation. Following gene annotation, symbol harmonisation, and variance-based filtering, the final molecular feature space contained on the order of tens of thousands of genes, reflecting the high dimensionality of bulk transcriptomic data. This high-dimensional environment, relative to the scale of the patient population, explains why regularised modelling solutions and cross-validation should be used with caution.

Overall, the data set represents a biologically realistic immunotherapy cohort that is characterised by sparse longitudinal sampling, high dimensionality of the molecular data and high heterogeneity in immune response across patients (Riaz et al., 2017).

4.2 Performance of Machine Learning Models

To predict changes in immune activation based on paired transcriptomic differences, multiple regression models were tested using patient-level K-fold cross-validation, with training and test patients strictly separated. The application of ElasticNet regression to the raw immune activation change showed good predictive performance with a cross-validated coefficient of determination (R^2) of about 0.71 and high Pearson and Spearman correlations between the predicted and observed immune changes. However, performance was sensitive to extreme responders because of the heavy-tailedness

of the target variable. Ridge regression, though stable, consistently underperformed ElasticNet. The standardised shrinkage introduced by Ridge reduced sensitivity to biologically significant signals, resulting in lower explained variance and higher prediction error. To correct target skewness, a Yeo-Johnson transformation was applied to the immune activation change in the ElasticNet framework. This modification brought a systematic improvement across all metrics, with a higher model explained variance ($R^2 \sim 0.73$), lower mean absolute error, and greater robustness across folds. Notably, the predictions were inversely transformed to the original scale of immune activation (which maintained interpretability).

Table 1: Cross-Validated Performance Comparison of Machine Learning Models

Model	Target Handling	RMSE	MAE	R^2	Pearson r	Spearman ρ
ElasticNet	Raw ΔE	~ 34.5	~ 19.3	~ 0.71	~ 0.89	~ 0.84
Ridge Regression	Raw ΔE	~ 45.1	~ 30.3	~ 0.51	~ 0.72	~ 0.64
ElasticNet	Yeo-Johnson ΔE	~ 33.4	~ 16.5	~ 0.73	~ 0.86	~ 0.83

Table 1 summarises the out-of-fold predictive performance of each evaluated model. ElasticNet with target transformation strikes the best balance among accuracy, robustness, and interpretability and is therefore chosen as the main machine learning model.

4.3 Patient-Specific Immune Activation Predictions

Out-of-fold predictions from the main ElasticNet model showed high consistency across patients, with predicted changes in immune activation following the observed values. Scatter analysis showed an excellent linear relationship between predicted and observed immune activation, indicating that the model captured both the directionality and the magnitude of changes in immune responses. Patients with large positive immune activation changes (hyper-responders) were overall predicted with a little attenuated magnitude, reflecting a regularisation effect of the model. Conversely, patients with minimal or negative immune activation were consistently identified, demonstrating the model's ability to discriminate non-responders and immune-suppressed cases.

Notably, there was no systematic bias in prediction errors across cross-validation folds, suggesting that the learned transcriptomic-immune relationship can be generalised.




Figure 1: Observed vs. Predicted Immune Activation Change (ΔE)

Each point in Figure 1 represents an individual patient. Predictions were made only by models that did not include the patient in the training data. The tight agreement with the identity line shows excellent agreement between predicted and observed changes in immune activation.

4.4 Identification of Molecular Drivers of Immune Activation

To identify transcriptional programs underlying the dynamics of immune activation, the ElasticNet coefficients were aggregated over the cross-validation folds to obtain a consensus gene effect. Genes were ranked according to the mean absolute magnitude of the coefficients and the frequency with which they were chosen in different folds.

The most highly ranked genes were enriched for immune effector and cytotoxic pathways, including CD8+ T-cell markers and natural killer and T-cell receptor signalling. Genes involved in interferon signalling and immune activation all

had positive coefficients, while genes associated with immune suppression or exhaustion had negative coefficients. Feature stability analysis showed that many top genes were selected across all cross-validation folds, suggesting stable, non-artifactual contributions.

Table 2: Top Molecular Drivers of Immune Activation Dynamics

Gene Symbol	Mean Coefficient (β)	Mean Absolute Coefficient	Non-zero Frequency
CD8B	Positive	High	1.00
GZMB	Positive	High	1.00
PRF1	Positive	High	1.00
NKG7	Positive	High	1.00
IFNG	Positive	Moderate	0.95

Table 2 lists the representative genes consistently associated with immune activation. Positive coefficients are associated with greater immune effector activity, and high selection frequency is associated with robustness across cross-validation folds.

4.5 ML-Driven Parameterisation of the Biomathematical Model

Machine learning-based transcriptomic immune activation scores, derived from stability-weighted ElasticNet coefficients, were linked to patient-specific immune activation parameters within the biomathematical tumour-immune model. The deduced activation parameters were widely dispersed, with substantial heterogeneity in immune stimulation by therapy among patients. Inferred activation parameters were higher in patients predicted to have strong immune activation and lower in immune-suppressed and non-responsive patients. The inferred parameters showed positive correlations with the observed and predicted changes in immune activation, indicating internal consistency between the data-driven and mechanistic elements of the framework. This parameterisation allowed patients to be stratified into distinct immune activation regimes, providing a mechanistic explanation for transcriptomic heterogeneity. Although continuous tumour dynamics were not modelled due to limited time information, the estimated parameters reflect patient-specific immunodynamics over the course of treatment.




Figure 2: Distribution of Inferred Immune Activation Parameters

The histogram and density estimate show patient-specific, therapy-induced immune activation strengths (α_i), estimated from transcriptomic immune activation scores derived using the machine learning framework. Figure 2 shows substantial inter-patient variability, with clear low- and high-activation regimes, indicating heterogeneity in immune responsiveness to immunotherapy.

4.6. Summary of Key Findings

- The GSE91061 cohort shows high heterogeneity in immune activation in immunotherapy.
- ElasticNet regression with target transformation for robust and interpretable immune state change prediction.
- Stable immune-related transcriptional programs are involved in immune activation dynamics.
- Machine learning outputs can be translated into mechanistic model parameters successfully, allowing tumour-

immune dynamics to be interpreted successfully on a patient-specific basis.

5. Discussion

The current paper presents a machine learning-based biomathematical model for analysing short-term tumour-immune interactions during immune checkpoint blockade, based on paired transcriptomic measurements. By combining regularised machine learning and a parsimonious mechanistic tumour-immune model, this work shows how to quantitatively infer, predict, and mechanistically interpret therapy-induced immune activation at the patient level. The results yield several important insights into immune response heterogeneity, model performance, and the value of hybrid modelling strategies in immuno-oncology.

5.1. Interpretation of Key Results

A key finding of this study is the high degree of inter-patient heterogeneity in therapy-induced immune activation, as measured by changes in cytolytic score between pre-treatment and on-treatment biopsies. The observed distribution includes immunosuppression, moderate activation, and hyperactivation of immunotherapy, which are well documented in melanoma patients. This heterogeneity has a biological significance and is consistent with the existing diversity of tumour immune microenvironment and immune responsiveness. The ElasticNet regression model showed high predictive power ($R^2 \sim 0.73$ after transforming the target), suggesting that therapy-induced immune activation is largely encoded in paired changes in the transcriptome. Importantly, by using regularisation, the feature selection was stable in the high-dimensional, small-sample setting, reducing overfitting while ensuring interpretability. Feature stability analysis showed that immune effector genes such as CD8B, GZMB, PRF1, NKG7, and IFNG consistently contributed to predictions, strengthening their biological relevance as drivers of immune activation. These results provide a basis for testing the premise that organised transcriptional programs, rather than individual genes, are responsible for an effective immune response to checkpoint blockade. The mapping of machine-learning-derived immune activation signatures to patient-specific parameters within a biomathematical tumour-immune model is an important conceptual contribution. The values for the inferred immune activation parameters showed a wide distribution, similar to the heterogeneity in changes in cytolytic scores. This correspondence demonstrates internal consistency between the predictions from the data and the mechanistic explanation, providing a biologically oriented explanation of the transcriptomic variability. Rather than focusing on machine learning outputs as black-box predictions, the framework integrates the machine learning output into an interpretable dynamical system, allowing patient-specific mechanistic characterisation of the immune response.

5.2. Comparison with Existing Literature

The results of this study are in line with and expand previous hybrid modelling work in immuno-oncology. Hybrid mechanistic-machine learning approaches have been used in the past to enhance predictive accuracy for immunotherapy situations in which clinical data is combined with mechanistic model outputs (Butner et al., 2024). However, many such studies have been done to predict long-term survival or model continuous tumour burden. In contrast to this, the current work explicitly considers sparse longitudinal sampling, with short-interval immune state transitions, in a way that is characteristic of the actual availability of biopsies. Multiscale and fractional tumour-immune models have given us rich theoretical knowledge on immune-tumour interactions and treatment effects (El Asraoui et al., 2025). Although they are rich dynamic models, they tend to rely on parameters that are difficult to estimate from patient-specific molecular data. The current framework is complementary to these approaches, as it demonstrates how transcriptomic signatures can inform mechanistic parameterisation in a data-limited environment. Purely data-driven machine learning analysis, such as deep learning or ensemble approaches, has been used to obtain impressive results when classifying immunotherapy responses (Fountzilias et al., 2025; Lin et al., 2025). However, such models are frequently not directly mechanistic. The present research overcomes this limitation by applying a linear, regularised model whose results can be directly mapped to biological activities in a tumour-immune ODE system, thereby bypassing the trade-off between predictive accuracy and mechanistic understanding. Individual cell and bulk-level results are also consistent with findings of feature-level studies that have associated cytotoxic immune programs with immunotherapy responsiveness (Dong et al., 2024). While approaches that work with single cells offer higher resolution, they are not always feasible in a clinical setting. The current study demonstrates that bulk RNA-seq, even when analysed in a paired, mechanistically informed fashion, can continue to provide biologically significant insights into immune dynamics.

5.3. Implications for Precision Oncology

Translationally, this paper highlights the importance of immune dynamics rather than describing immune states. By

focusing on within-patient changes in immune activation, the framework minimises confounding by baseline between-patient variability and directly assesses the effects of the therapy. This approach is consistent with the clinical reality of immunotherapy, in which dynamic immune engagement, rather than baseline immune infiltration, is critical to determining therapeutic benefit. The incorporation of machine learning outputs into a mechanistic model enables patient-specific interpretation of immune responsiveness, which may be useful for treatment stratification or hypothesis generation in future studies. Although the framework itself is not meant to make clinical decisions at this point, it provides a methodological basis for interpretable, patient-specific immune modelling in precision oncology.

5.4. Limitations

There are several limitations to be acknowledged. First, the analysis is based on bulk RNA-seq data and on the cytolytic score as a proxy for immune effector activity. Although it is biologically based, this method fails to address interactions between cell types as well as spatial heterogeneity in the tumour microenvironment. Second, the small sample size of paired biopsy studies requires regularisation and limits the complexity of the machine-learning and mechanistic components. Third, the biomathematical model is deliberately simplified and deals with short-interval state transitions rather than continuous long-term tumour trajectories. As such, it cannot be used to predict future tumor burden and long-term clinical outcomes.

5.5. Future Directions

Future work could expand this framework in several ways. Biological resolution can be improved by including additional immune proxies, e.g., immune cell deconvolution estimates or spatial transcriptomic features, without loss of interpretability. Adding the ability to consider multi-timepoint datasets would also allow reconstruction of parts of dynamic immune processes and validation of the inferred parameters over longer time horizons. Lastly, interlinkage with other omics modalities or imaging-based immune capabilities might also increase predictive capability and mechanistic depth and facilitate the creation of genuinely multi-scale patient-specific tumour-immune models. Overall, this paper shows that parameterisation of biomathematical tumour-immune models using machine learning is not only possible but also informative, subject to feasible constraints from real clinical data. The presented framework brings predictive precision and mechanistic interpretability together, adding to the growing list of hybrid modelling methods expected to advance the field of precision immuno-oncology.

6. Conclusion

This paper introduces a machine learning-based biomathematical model of predicting and interpreting short-period tumour immune dynamics in cancer models during immunotherapy. Using paired pre-treatment and on-treatment bulk RNA-seq data from melanoma patients, we showed that therapy-induced immune activation, quantified by changes in cytolytic score, can be robustly predicted from transcriptomic differences using regularised machine learning. ElasticNet regression achieved strong performance in cross-validation and enabled stable identification of immune-relevant transcriptional drivers. We used the outputs of these data-driven models to convert transcriptomic signatures into patient-specific parameters of immune activation in a parsimonious tumour-immune ordinary differential equation model, thereby providing a mechanistic understanding of immune response heterogeneity. These results have significant consequences for accurate oncology and computational immunology. First, they point out the importance of modelling changes in immune states, rather than just static immune measurements, to provide a more biologically meaningful representation of immunotherapy response. Second, the proposed hybrid framework shows that machine learning and mechanistic modelling are not competing paradigms but rather complementary ones. Machine learning helps us infer from high-dimensional molecular data, and biomathematical modelling provides interpretability and biological grounding, together supporting patient-specific immune characterisation. Based on these results, we recommend the novel use of hybrid modelling strategies for immunotherapy research, especially in settings where longitudinal sampling is not feasible. The structure included here can serve as a roadmap for incorporating transcriptomic data into mechanistic models without overfitting the unobservable dynamics, thereby enhancing interpretability and translational significance. Future studies should build on this research to extend it to larger and more diverse patient cohorts, incorporate additional immune and spatial biomarkers, and examine multi-timepoint datasets to enable reconstruction of continuous tumour-immune trajectories. Further refinement of the mechanistic component and validation against clinical outcomes will be the key to moving towards clinical decision support. On the whole, this study adds a philosophical and explainable modeling approach in the study and prediction of tumour-immune interactions in cancer systems.

7. References

- Butner, J. D., Dogra, P., Chung, C., Koay, E. J., Welsh, J. W., Hong, D. S., Cristini, V., & Wang, Z. (2024). Hybridizing mechanistic modeling and deep learning for personalized survival prediction after immune checkpoint inhibitor immunotherapy. *npj Systems Biology and Applications*, 10(1), 88.
- Chattopadhyay, A. K., Unkundiye, A. P. N., & Pearce, G. (2025). Hybrid Machine Learning and Mathematical Modeling for Tumor Dynamics Prediction: Comparing SPIONs against mNP-FDG. *arXiv preprint arXiv:2505.21094*.
- Dong, Y., Chen, Z., Yang, F., Wei, J., Huang, J., & Long, X. (2024). Prediction of immunotherapy responsiveness in melanoma through single-cell sequencing-based characterization of the tumor immune microenvironment. *Translational Oncology*, 43, 101910.
- El Asraoui, H., Hilal, K., & El Hajaji, A. (2025). Hybrid discrete-continuum modeling of tumor-immune interactions: Fractional time and space analysis with immunotherapy. *Mathematical Methods in the Applied Sciences*, 48(4), 5126-5153.
- Fountzilias, E., Pearce, T., Baysal, M. A., Chakraborty, A., & Tsimberidou, A. M. (2025). Convergence of evolving artificial intelligence and machine learning techniques in precision oncology. *NPJ digital medicine*, 8(1), 75.
- Ghahramani, M., & Bavi, O. (2026). Biomechanical modeling of glioblastoma progression: a comprehensive review from classic mathematical frameworks to data-driven strategies: M. Ghahramani and O. Bavi. *Biomechanics and Modeling in Mechanobiology*, 25(1), 1.
- Joshi, S., Nadaf, H., Jabade, M., & Pimpalekar, S. (2025). Oncology in the AI Era: Transforming Cancer Care Through Intelligent Diagnosis and Treatment. *Artificial Intelligence in Oncology: Cancer Diagnosis and Treatment, Medical Imaging, and Personalized Medicine*, 595-608.
- Kumar, R. P., Chauhan, V., & Chauhan, N. (2024). Oncogenic Genes Identification Using NLP and Ensemble Learning Approach. 2024 International Conference on Artificial Intelligence and Emerging Technology (Global AI Summit),
- Lin, R.-B., Zhou, L., Lin, Y.-C., Yu, Y., Yang, H.-C., & Yu, C.-W. (2025). A machine learning framework for cancer prognostics: integrating temporal and immune gene dynamics via ARIMA-CNN. *Biomedicine*, 13(11), 2751.
- Matin, H. N., & Setayeshi, S. (2024). A computational tumor growth model experience based on molecular dynamics point of view using deep cellular automata. *Artificial Intelligence in Medicine*, 148, 102752.
- Mortazavi, S. M. A., & Firoozabadi, B. (2024). Towards a framework for predicting immunotherapy outcome: a hybrid multiscale mathematical model of immune response to vascular tumor growth. *Biomechanics and Modeling in Mechanobiology*, 23(6), 2243-2264.
- Pinhasi, A., & Yizhak, K. (2025). Uncovering gene and cellular signatures of immune checkpoint response via machine learning and single-cell RNA-seq. *NPJ Precision Oncology*, 9(1), 95.
- Qin, Y., Huo, M., Liu, X., & Li, S. C. (2024). Biomarkers and computational models for predicting efficacy to tumor ICI immunotherapy. *Frontiers in Immunology*, 15, 1368749.
- Razali, A. B. (2024). Critique of Genetic Mutation Rates in Evolving Populations: A Biomathematical Approach. *Letters in biomathematics*, 11(1), 11-20.
- Riaz, N., Havel, J. J., Makarov, V., Desrichard, A., Urba, W. J., Sims, J. S., Hodi, F. S., Martín-Algarra, S., Mandal, R., & Sharfman, W. H. (2017). Tumor and microenvironment evolution during immunotherapy with nivolumab. *Cell*, 171(4), 934-949. e916.
- Ricker, C. A., Meli, K., & Van Allen, E. M. (2024). Historical perspective and future directions: computational science in immuno-oncology. *Journal for immunotherapy of cancer*, 12(1), e008306.
- Rundo, F., Spata, M. O., Pino, C., & Battiato, S. (2025). Dopamine-Inspired Neuro-Modulated Radiomic Pipeline for Computed-Tomography-Driven Prediction of Immunotherapy Response in Metastatic Cancer Treatment. Proceedings of the IEEE/CVF International Conference on Computer Vision,
- Salman, M., Das, P. K., & Mohanty, S. K. (2024). A systematic review on recent advancements in deep learning and mathematical modeling for efficient detection of glioblastoma. *IEEE Transactions on Instrumentation and Measurement*, 73, 1-34.

- Shutaywi, M., Shah, Z., & Jan, R. (2025). A robust study of the dynamics of tumor–immune interaction for public health via fractional framework. *The European Physical Journal Special Topics*, 234(8), 1971-1990.
- Uthamacumaran, A., Kiyokawa, J., & Wakimoto, H. (2025). AI-Driven Hybrid Ecological Model for Predicting Oncolytic Viral Therapy Dynamics. *arXiv preprint arXiv:2501.10620*.
- Zanib, S. A., Tariq, R., Abbas, N., Nazir, A., & Shatanawi, W. (2025). A hybrid mathematical and machine learning approach for COVID-19 dynamics: SA Zanib et al. *International Journal of Dynamics and Control*, 13(7), 274.
- Zhang, H., Chen, L., Li, L., Liu, Y., Das, B., Zhai, S., Tan, J., Jiang, Y., Turco, S., & Yao, Y. (2025). Prediction and analysis of tumor infiltrating lymphocytes across 28 cancers by TILScout using deep learning. *NPJ Precision Oncology*, 9(1), 76.