
A Hybrid Machine Learning Framework for Mechanistically Interpretable Latent Parameter Inference in a Spatiotemporal CAR-T Therapy Model for Solid Tumours

[Maxim Polyakov](#)*

Posted Date: 9 April 2026

doi: 10.20944/preprints202604.0650.v1

Keywords: CAR-T cell therapy; spatiotemporal modelling; reaction-diffusion equations; machine learning; model calibration; antigen escape; tumour microenvironment; computational oncology; personalised medicine



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Hybrid Machine Learning Framework for Mechanistically Interpretable Latent Parameter Inference in a Spatiotemporal CAR-T Therapy Model for Solid Tumours

Maxim Polyakov 

Department of Information Systems and Computer Modelling, Volgograd State University, Universitetsky pr., 100, Volgograd 400062, Russia; m.v.polyakov@volsu.ru

Abstract

CAR-T cell therapy remains ineffective in most solid tumours because effector cells infiltrate poorly, undergo exhaustion, and face antigen escape within an immunosuppressive microenvironment. To address this, we developed a hybrid framework that combines a mechanistic spatiotemporal model with machine learning for limited patient-specific calibration. At its core, we employed a reaction-diffusion-chemotaxis model describing functional and exhausted CAR-T cells, antigen-positive and antigen-negative tumour subpopulations, a chemoattractant, an immunosuppressive factor, and hypoxia. Gradient boosting combined with nested cross-validation was used as the primary method for parameter inference. Parameters characterising the tumour microenvironment and CAR-T cell exhaustion were recovered most robustly, whereas antigen escape and individualised initial conditions were identified substantially less accurately. As an auxiliary reference point, we also considered a direct empirical baseline for binary clinical outcomes. This baseline indicated that the observed clinical features contained a more stable signal for disease control than for objective response. A favourable response was associated with high CAR-T cell infiltration and cytotoxic potency, whereas resistance was linked to exhaustion, antigen escape, and a suppressive microenvironment. Overall, the proposed approach constitutes an interpretable proof-of-concept platform for limited patient-specific inference of latent parameters and for stratifying the mechanisms underlying response and resistance in CAR-T cell therapy for solid tumours.

Keywords: CAR-T cell therapy; spatiotemporal modelling; reaction-diffusion equations; machine learning; model calibration; antigen escape; tumour microenvironment; computational oncology; personalised medicine

1. Introduction

Chimeric antigen receptor T-cell (CAR-T) therapy has become a major advance in modern cellular immunotherapy. In B cell haematological malignancies, it yields high rates of complete remission and durable disease control in patients with refractory or relapsed disease [1–4]. These outcomes have been made possible by the rational design of CAR constructs, the optimisation of co-stimulatory domains, and the refinement of cell manufacturing protocols [4,5]. However, translating these advances to solid tumours has proved considerably more challenging, necessitating the development of new biological and computational strategies to improve therapeutic efficacy [6–11].

Solid tumours are characterised by marked spatial heterogeneity, complex tissue architecture, limited availability of tumour-specific antigens, and temporal variability in antigenic profiles [6–8,11,12]. One of the principal mechanisms underlying failure of CAR-T therapy in this group of diseases is antigen escape, whereby selective pressure leads to loss or downregulation of target expression and the emergence of resistant clones [11,13,14]. An additional limitation is the risk of on-target

off-tumour toxicity associated with the expression of potential targets in normal tissues [6,10,11,15]. Therefore, the development of CAR-T approaches for solid tumours requires the simultaneous consideration of antigen heterogeneity, safety, durability of therapeutic effect, and the potential for personalisation [12,16,17].

An equally important limitation is insufficient infiltration of CAR-T cells into tumour tissue. To exert their cytotoxic effects, effector cells must survive after administration, reach the tumour site, overcome vascular and stromal barriers, distribute throughout the tissue volume, and retain functional activity within an adverse microenvironment [6,7,10,15]. A major contributor to this problem is the mismatch between the chemokine profile produced by the tumour and the receptor repertoire of CAR-T cells, which reduces directed migration and accumulation of effector cells [18]. Accordingly, engineering strategies aimed at enhancing the recruitment and retention of immune cells within tumours are being actively developed, including the modification of CAR-T cells to secrete interleukin-7 (IL-7) and C-C motif chemokine ligand 19 (CCL19), which improves infiltration, survival, and antitumour activity in preclinical models [19]. More broadly, it is precisely these spatial constraints on CAR-T cell delivery and migration that make models capable of capturing tissue heterogeneity and local interactions between cellular populations particularly valuable [14,17]. In the wider context of spatial oncological modelling, diffusion-based approaches have already been used to describe tumour growth in heterogeneous environments, making it possible to account for the influence of tissue heterogeneity on tumour population dynamics [20].

A key determinant of therapeutic failure in solid tumours is the immunosuppressive tumour microenvironment. It encompasses hypoxia, nutrient deprivation, immunoregulatory cytokines, suppressive myeloid cells, regulatory T cells, and the expression of inhibitory ligands, which collectively impair the proliferation, cytotoxicity, and long-term persistence of CAR-T cells [6,7,15,16]. Under these conditions, CAR-T cells enter a state of functional impairment or exhaustion, characterised by reduced production of effector molecules, metabolic impairment, and sustained expression of inhibitory receptors [21–25]. In solid tumours, the hypoxic and metabolically adverse microenvironment is particularly important, as it limits the energetic capacity of T cells and promotes the development of an exhausted phenotype [26]. Thus, the spatiotemporal dynamics of CAR-T therapy are determined not only by direct cytotoxicity, but also by the balance among cell delivery, activation, exhaustion, and local mechanisms of immune suppression [11,16,17].

Despite these limitations, recent years have also brought encouraging clinical evidence supporting the fundamental applicability of CAR-T approaches to solid tumours. In recurrent high-grade gliomas, the safety and clinical activity of locoregional delivery of IL-13R α 2-targeted CAR-T cells have been demonstrated [27]. In gastrointestinal tumours, promising results have been reported for CLDN18.2-specific CAR-T cell therapy [28]. In hepatocellular carcinoma, clinical activity of GPC3-targeted CAR-T cells has been reported [29], while in neuroblastoma GD2 CAR-T cells have shown antitumour activity without pronounced on-target off-tumour toxicity [30]. These studies confirm the genuine clinical potential of CAR-T therapy for solid tumours; however, treatment response remains highly variable and depends on the biological characteristics of both the tumour and the patient, thereby necessitating a shift from population-level analytical frameworks to personalised models [10,12].

Against this backdrop, the role of quantitative and computational approaches capable of formalising CAR-T cell kinetics, tumour burden dynamics, and the influence of the microenvironment on therapeutic response is becoming increasingly important. Models of the cellular kinetics and pharmacodynamics of commercial CAR-T products have already been developed to describe cell expansion, contraction, and persistence following infusion [31–33]. To analyse interactions between tumour cells and CAR-T cells under controlled conditions, mathematical models based on real-time killing assay data have been proposed, enabling the respective contributions of CAR-T cell proliferation and exhaustion to be deconvoluted [34]. More sophisticated mechanistic approaches are also being developed that account for multiple CAR-T cell interactions with tumour cells, tumour spatial structure, antigen escape, and bystander killing effects [35–37]. In parallel, models of intracellular signal transduction

and methods for predicting the in vitro potency of cellular products are being established, and these may be used for the parameterisation and validation of broader models of immune response [38,39]. Taken together, these studies show that mathematical modelling has already become an important tool for interpreting CAR-T therapy; however, most existing models are either focused on haematological malignancies or remain insufficiently personalised to describe solid tumours.

In parallel, computational oncology has seen the emergence of a research direction focused on the development of patient-specific mechanistic models, digital twins, and hybrid frameworks that combine mechanism-oriented modelling with machine learning methods [40–43]. Such approaches are particularly valuable in settings where experimental and clinical data are fragmentary and model parameters are only partially observable. Under these conditions, machine learning methods may be used not as a “black box” for direct outcome prediction, but rather as a tool for inferring the latent parameters of a mechanistic model from clinical, molecular, and experimental features [42,43]. This strategy has already been applied to the personalisation of reaction–diffusion models of tumour dynamics using machine learning methods [44]. In the context of CAR-T therapy for solid tumours, this is particularly important because therapeutic response is determined by a combination of difficult-to-measure characteristics, including the intensity of infiltration, the rate of exhaustion, the level of immunosuppression, antigen heterogeneity, and the efficiency of local cytotoxicity [13,18,25,26].

Thus, the contemporary literature points to a need for hybrid, mechanistically interpretable models of CAR-T therapy for solid tumours that account for the spatiotemporal structure of the tumour and enable limited patient-specific personalisation. Particularly promising is the use of published experimental and clinical data for population-level calibration of the mechanistic core of the model, followed by the inference of a limited set of latent patient-specific parameters using machine learning methods. This approach should be regarded as a feasibility study and a hypothesis-generating framework that supports the stratification of mechanisms of response and resistance, rather than as a clinically validated tool for precise individual prediction [34,40,42,43].

In this study, we propose a hybrid computational approach to the personalisation of a spatiotemporal model of CAR-T therapy for solid tumours. It is based on a mechanistic description of the interactions among CAR-T cells, tumour populations, and microenvironmental factors, complemented by machine learning algorithms for inferring patient-specific parameters from published data and clinical features.

The principal aim of this study is to establish a hybrid framework in which machine learning is used to infer a reduced set of patient-specific latent parameters of a spatiotemporal model of CAR-T therapy from available clinical data, while the mechanistic core is used to provide a biological interpretation of these parameters in terms of therapeutic response and mechanisms of resistance in solid tumours.

The remainder of this paper is organised as follows. Section 2 presents the materials and methods. Section 2.1 formulates the spatiotemporal reaction-diffusion-chemotaxis model of CAR-T therapy for solid tumours. Section 2.2 describes the numerical solution of the system and the software implementation of the computational algorithm. Section 2.3 defines the problem of patient-specific model personalisation. Section 2.4 characterises the source of the clinical data and the procedure used for variable preparation. Section 2.5 considers machine learning methods for inferring the hidden patient-specific parameters of the model. Section 2.6 specifies the validation protocol and the performance metrics used. Section 3 presents the results of the study. Section 3.1 analyses the accuracy of recovering patient-specific latent parameters. Section 3.2 considers an auxiliary direct empirical baseline for binary clinical outcomes. Section 3.3 provides the main mechanistic analysis of the latent factors underlying response and resistance. Section 4 discusses the findings, the limitations of the proposed approach, and directions for its further development. Section 5 sets out the principal conclusions of the study.

2. Materials and Methods

2.1. Mechanistic Spatiotemporal Model of CAR-T Therapy for Solid Tumours

As the mechanistic core of the personalised framework, we use a previously developed spatiotemporal reaction-diffusion-chemotaxis model describing the coupled dynamics of CAR-T cells, tumour subpopulations, and microenvironmental factors [45]. In this study, machine learning methods do not replace the underlying biology of the process, but are instead used to infer hidden patient-specific parameters within an interpretable system of equations.

We consider a bounded domain $\Omega \subset \mathbb{R}^3$ corresponding to the volume of the tumour tissue and its immediate microenvironment. The model variables depend on the spatial coordinate $\mathbf{x} \in \Omega$ and time $t \geq 0$. Let $C(\mathbf{x}, t)$ denote the density of functional CAR-T cells, and let $E(\mathbf{x}, t)$ denote the density of exhausted CAR-T cells. Furthermore, let $T_A(\mathbf{x}, t)$ denote the density of tumour cells with preserved expression of the target antigen, whereas $T_B(\mathbf{x}, t)$ denotes the density of tumour cells with reduced or lost antigen expression. In addition, we introduce the following microenvironmental fields: $S(\mathbf{x}, t)$, the concentration of a chemoattractant mediating directed migration of CAR-T cells; $A(\mathbf{x}, t)$, the concentration of a generalised soluble immunosuppressive factor (such as transforming growth factor- β (TGF- β), interleukin-10 (IL-10), adenosine, among others); and $H(\mathbf{x}, t)$, the normalised degree of hypoxia.

The governing system of equations takes the following form

$$\left\{ \begin{array}{l} \frac{\partial C}{\partial t} = \nabla \cdot (D_C(\mathbf{x}) \nabla C) - \nabla \cdot (\chi_C(A) C \nabla S) + \rho_C \frac{C T_A}{K_C + C} \\ \quad - \delta_{CT} C T_A - k_e(A, H) C + \frac{p_r}{1 + A} E - d_C C, \quad (1) \\ \frac{\partial E}{\partial t} = \nabla \cdot (D_E \nabla E) + k_e(A, H) C - \frac{p_r}{1 + A} E - d_E E, \quad (2) \\ \frac{\partial T_A}{\partial t} = \nabla \cdot (D_T \nabla T_A) + r_T \left(1 - \frac{T_A + T_B}{K_T} \right) T_A - \gamma C T_A \\ \quad - \mu_{AB}(C) T_A + \mu_{BA}(C) T_B - d_{T_A} T_A, \quad (3) \\ \frac{\partial T_B}{\partial t} = \nabla \cdot (D_T \nabla T_B) + r_T \left(1 - \frac{T_A + T_B}{K_T} \right) T_B + \mu_{AB}(C) T_A \\ \quad - \mu_{BA}(C) T_B - d_{T_B} T_B, \quad (4) \\ \frac{\partial S}{\partial t} = D_S \nabla^2 S + \sigma_S T_A - \lambda_S S, \quad (5) \\ \frac{\partial A}{\partial t} = D_A \nabla^2 A + \sigma_A (T_A + T_B) + \sigma_M f(H) - \lambda_A A, \quad (6) \\ \frac{\partial H}{\partial t} = D_H \nabla^2 H + \eta (T_A + T_B) - \lambda_H H. \quad (7) \end{array} \right.$$

Here, the auxiliary relationships are defined as follows

$$\chi_C(A) = \frac{\chi_0}{1 + \alpha_A A}, \quad k_e(A, H) = k_e^0 (1 + \beta_A A + \beta_H H), \quad \mu_{AB}(C) = \mu_0 + \nu C, \quad (8)$$

$$\mu_{BA}(C) = \mu_1 \exp(-\zeta C), \quad f(H) = \frac{H}{K_H + H}. \quad (9)$$

Within the present model, infiltration is represented by a combination of diffusion and chemotaxis of CAR-T cells in heterogeneous tissue. CAR-T cell proliferation is induced by contact with antigen-positive tumour cells and is limited by saturation. Exhaustion is described as a microenvironment-regulated transition of functional cells to a dysfunctional state. Cytotoxic activity is implemented through the direct elimination of the T_A subpopulation. Antigen escape is represented as phenotypic switching between T_A and T_B , with both the direction and the intensity of switching depending on

therapeutic pressure. The immunosuppressive microenvironment and hypoxia not only alter CAR-T cell motility, but also accelerate exhaustion, thereby establishing conditions conducive to resistance.

The system of Equations (1)–(7) is supplemented by the following initial conditions

$$C(\mathbf{x}, 0) = C_0(\mathbf{x}), \quad E(\mathbf{x}, 0) = E_0(\mathbf{x}), \quad T_A(\mathbf{x}, 0) = T_{A,0}(\mathbf{x}), \quad T_B(\mathbf{x}, 0) = T_{B,0}(\mathbf{x}), \quad (10)$$

$$S(\mathbf{x}, 0) = S_0(\mathbf{x}), \quad A(\mathbf{x}, 0) = A_0(\mathbf{x}), \quad H(\mathbf{x}, 0) = H_0(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad (11)$$

where the functions $T_{A,0}$ and $T_{B,0}$ describe the initial spatial structure of the tumour and its antigenic heterogeneity, while C_0 specifies the initial distribution of the infused cellular product. No-flux boundary conditions are imposed on the boundary of the domain Ω

$$\frac{\partial T_A}{\partial n} = \frac{\partial T_B}{\partial n} = \frac{\partial E}{\partial n} = \frac{\partial S}{\partial n} = \frac{\partial A}{\partial n} = \frac{\partial H}{\partial n} = 0, \quad \mathbf{x} \in \Omega, \quad (12)$$

$$(D_C(\mathbf{x})\nabla C - \chi_C(A)C\nabla S) \cdot \mathbf{n} = 0, \quad \mathbf{x} \in \partial\Omega, \quad (13)$$

which corresponds to the absence of flux of cells and soluble factors across the outer boundary of the modelled volume.

The principal quantities used in the model and their biological interpretation are summarised in Table 1. Detailed parameter ranges are provided in [45].

Table 1. Brief notation for the principal variables and parameters of the model.

Notation	Biological meaning
C	functional CAR-T cells
E	exhausted CAR-T cells
T_A	antigen-positive tumour cells
T_B	tumour cells with reduced antigen expression
S	chemoattractant governing directed migration
A	generalised soluble immunosuppressive factor
H	normalised microenvironmental hypoxia
D_C, D_E, D_T	diffusion coefficients of the cellular populations
D_S, D_A, D_H	diffusion coefficients of the chemoattractant, the immunosuppressive factor, and the hypoxia variable
χ_0, α_A	baseline chemotactic sensitivity and its suppression by immunosuppression
ρ_C, K_C	parameters of antigen-dependent CAR-T cell proliferation
γ, δ_{CT}	intensity of cytotoxic interaction and associated functional loss of CAR-T cells
$k_e^0, \beta_A, \beta_H, p_r$	parameters of exhaustion and partial recovery of CAR-T cells
K_H	half-saturation parameter of the hypoxic effect
d_C, d_E	natural death rates of functional and exhausted CAR-T cells
r_T, K_T	tumour growth parameters
d_{T_A}, d_{T_B}	natural death rates of antigen-positive and antigen-reduced tumour cells
μ_0, μ_1, ν, ζ	parameters of antigen escape and reverse phenotypic switching
$\sigma_S, \sigma_A, \sigma_M, \eta$	production of the chemoattractant, immunosuppressive factors, and hypoxia accumulation
$\lambda_S, \lambda_A, \lambda_H$	decay of the microenvironmental fields

Thus, in the proposed approach, personalisation is based not on direct empirical prediction of clinical outcome, but on the adaptation of biologically interpretable parameters of a spatiotemporal mechanistic model. This makes it possible to preserve the causal structure of the description, support mechanistically interpretable stratification of specific resistance mechanisms, and use the model as a foundation for further analysis of individual CAR-T therapy scenarios.

2.2. Numerical Solution of the Model and Software Implementation

The system of Equations (1)–(7) is solved numerically using the finite difference method on a uniform three-dimensional rectangular grid with operator splitting by physical processes. This construction makes it possible to separate the diffusive component from the nonlinear reaction-chemotaxis terms, thereby providing an acceptable compromise among stability, accuracy, and computational cost.

The computational domain is defined as the rectangular parallelepiped

$$\Omega = [0, L_x] \times [0, L_y] \times [0, L_z], \quad (14)$$

which is discretised by a uniform grid

$$x_i = i \Delta x, i = 0, \dots, N_x - 1, y_j = j \Delta y, j = 0, \dots, N_y - 1, z_k = k \Delta z, k = 0, \dots, N_z - 1, \quad (15)$$

where

$$\Delta x = \frac{L_x}{N_x - 1}, \quad \Delta y = \frac{L_y}{N_y - 1}, \quad \Delta z = \frac{L_z}{N_z - 1}. \quad (16)$$

In the implementation, the domain dimensions and grid resolution are specified by the parameters L_x , L_y , L_z and n_x , n_y , n_z .

For the diffusion terms, a standard second-order central finite difference approximation in space is used. For any variable $u \in \{C, E, T_A, T_B, S, A, H\}$, the Laplacian operator is approximated by a seven-point stencil

$$\begin{aligned} (\nabla_h^2 u)_{i,j,k} \approx & \frac{u_{i+1,j,k} - 2u_{i,j,k} + u_{i-1,j,k}}{\Delta x^2} + \frac{u_{i,j+1,k} - 2u_{i,j,k} + u_{i,j-1,k}}{\Delta y^2} \\ & + \frac{u_{i,j,k+1} - 2u_{i,j,k} + u_{i,j,k-1}}{\Delta z^2}. \end{aligned} \quad (17)$$

The chemotactic term in the equation for functional CAR-T cells is discretised in conservative flux form, thereby reducing numerical artefacts and remaining consistent with the physical interpretation of directed migration.

Time integration is based on Strang splitting. Denoting by

$$U = \begin{pmatrix} C \\ E \\ T_A \\ T_B \\ S \\ A \\ H \end{pmatrix} \quad (18)$$

the vector of state fields, we represent the system in the form

$$\frac{dU}{dt} = \mathcal{R}(U) + \mathcal{D}(U), \quad (19)$$

where \mathcal{R} includes the reaction and chemotactic terms, whereas \mathcal{D} includes diffusion. One time step of size Δt is then computed as

$$U^{n+1} = \Phi_{\mathcal{R}}\left(\frac{\Delta t}{2}\right) \Phi_{\mathcal{D}}(\Delta t) \Phi_{\mathcal{R}}\left(\frac{\Delta t}{2}\right) U^n. \quad (20)$$

The reaction-chemotaxis sub-system is integrated using an explicit two-stage second-order Runge-Kutta scheme

$$K_1 = \mathcal{R}(U^m), \quad K_2 = \mathcal{R}(U^m + hK_1), \quad (21)$$

$$U^{m+1} = U^m + \frac{h}{2}(K_1 + K_2). \quad (22)$$

To enhance stability for stiff parameter sets, the implementation allows the explicit step to be subdivided into substeps using the parameter `exp_substeps`. After each substep, a projection onto the physically admissible domain is performed to prevent the occurrence of negative concentration and density values.

The diffusion sub-system for each variable is solved implicitly. For an arbitrary field u , the following relation is used

$$\frac{u^* - u^{\text{rhs}}}{\tau} = D_u \nabla_h^2 u^*, \quad \text{or} \quad \left(I - \tau D_u \nabla_h^2 \right) u^* = u^{\text{rhs}}, \quad (23)$$

where u^{rhs} is the known value of the corresponding component of the solution, obtained after the preceding reaction-chemotaxis substep and used as the right-hand side in the implicit treatment of diffusion. The resulting linear system is solved using the Jacobi iterative method. The number of iterations and the stopping criterion are specified by the parameters `diff_iters` and `diff_tol`.

Reflecting boundary conditions are imposed for all variables. In the software implementation, this is achieved by referencing the nearest admissible grid node whenever an index falls outside the domain, which is equivalent to a zero normal derivative and, consequently, to zero flux across the boundary. The initial conditions are generated by a separate procedure. In the test configuration, a spherical tumour nodule located at the centre of the domain and a layer of CAR-T cells positioned near one of the boundaries are considered.

The numerical solver is implemented in C++, using the `Params`, `Grid3D`, `Field`, `State`, and `Deriv` structures to store parameters, grid geometry, and state fields. To accelerate computation, loops over grid nodes are parallelised using `OpenMP`. The results are saved in CSV format, with each record containing the time, node indices, coordinates, and the values of all seven variables ($t, i, j, k, x, y, z, C, E, T_A, T_B, S, A, H$). This output organisation is convenient for subsequent visualisation, statistical analysis, and integration with external computational modules.

Figure 1 schematically shows the structure of the computational solver. At the highest level, the implementation comprises parameter and grid initialisation, generation of the initial conditions, the time-stepping loop, and result storage. Within each time step, Strang splitting is applied: the reaction-chemotaxis sub-system is integrated using an explicit RK2 scheme, whereas the diffusion sub-system is solved implicitly for each component of the state vector using the Jacobi iterative method.

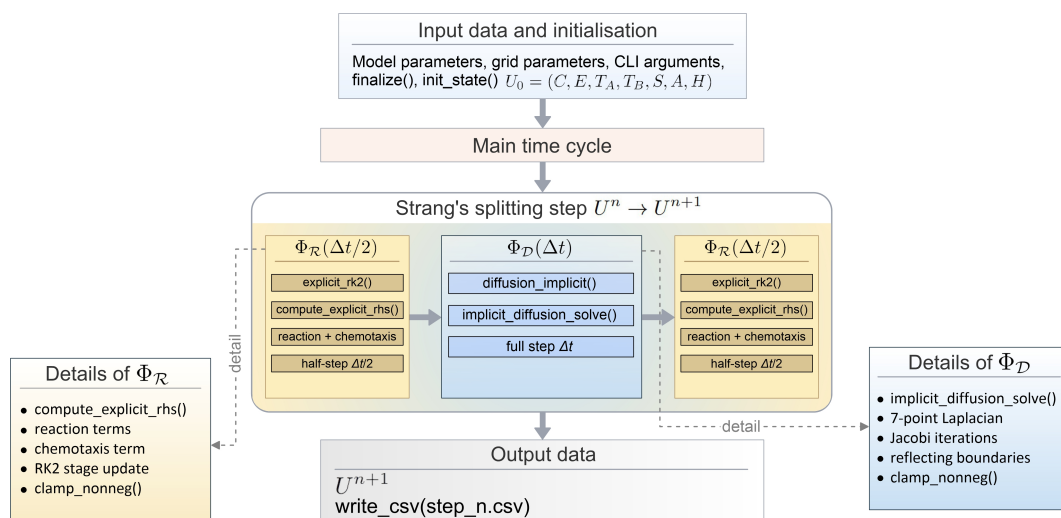


Figure 1. Schematic of the numerical solver architecture and the structure of a single time step based on Strang splitting.

To ensure reproducibility of the results, the full source code of the numerical solver has been made available in an open GitHub repository: <https://github.com/maximvpolyakov/CAR-T-model> (accessed on 6 April 2026).

2.3. Formulation of the Limited Inverse Problem of Patient-Specific Model Personalisation

In this work, personalisation is understood as a limited parametric adaptation of a spatiotemporal mechanistic model of CAR-T therapy to an individual patient. The objective is not to predict clinical outcome directly from observable features, but to estimate a set of latent parameters that enables the mechanistic description to be aligned with the available observations.

The full set of model parameters is represented as

$$\Theta = \Theta^{\text{POP}} \cup \Theta^{\text{pat}} \cup \Theta^{\text{tr}}, \quad (24)$$

where Θ^{POP} denotes the population-level parameters, pre-calibrated using literature and experimental data, Θ^{pat} denotes the patient-specific parameters, and Θ^{tr} denotes the treatment-regimen parameters. Here, Θ^{tr} includes the observed characteristics of CT041 administration and the associated initial conditions, whereas from the set Θ^{pat} we identify a limited subset for personalisation

$$\tilde{\Theta}^{\text{pat}} \subseteq \Theta^{\text{pat}}. \quad (25)$$

The reduced vector of patient-specific parameters is defined as

$$\tilde{\theta}^{\text{pat}} = (\theta_{\text{inf}}, \theta_{\text{exh}}, \theta_{\text{cyt}}, \theta_{\text{esc}}, \theta_{\text{env}}, \theta_{\text{init}}), \quad (26)$$

where θ_{inf} characterises the effective infiltration of CAR-T cells into the tumour, θ_{exh} denotes the propensity for functional exhaustion, θ_{cyt} denotes the integral antitumour potency, θ_{esc} denotes the intensity of antigen escape, θ_{env} denotes the severity of the immunosuppressive microenvironment, and θ_{init} denotes the individualised initial conditions associated with tumour volume and baseline antigenic heterogeneity. The exact composition of this vector may subsequently be refined in accordance with data availability and the results of identifiability analysis.

For patient p , two groups of observable quantities are distinguished. The first forms the feature vector

$$\mathbf{x}^{(p)} \in \mathcal{X}, \quad (27)$$

which may include tumour type, pretreatment characteristics, the level of CLDN18.2 expression, the CT041 dose, and other variables available before or at the early stages of therapy. The second group forms the vector of observed responses

$$\mathbf{y}_{\text{obs}}^{(p)} \in \mathcal{Y}, \quad (28)$$

which comprises quantities comparable to the outputs of the mechanistic model, including the response category according to RECIST 1.1, characteristics of peripheral CAR-T cell kinetics, the grade of cytokine release syndrome (CRS), peak inflammatory markers, and time to event measures represented in a quantitatively interpretable form.

The personalisation problem is formulated as a limited inverse problem. The aim is to determine a vector $\tilde{\theta}_{\text{pat}}^{(p)}$ that, for fixed θ^{POP} and known $\theta_{\text{tr}}^{(p)}$, provides the best possible agreement between the model and the patient's observed data. To this end, the following Cauchy problem is solved:

$$U_t = \mathcal{M}(U; \theta^{\text{POP}}, \tilde{\theta}_{\text{pat}}^{(p)}, \theta_{\text{tr}}^{(p)}), \quad U(\mathbf{x}, 0) = U_0(\mathbf{x}; \tilde{\theta}_{\text{pat}}^{(p)}, \theta_{\text{tr}}^{(p)}), \quad (29)$$

where $U = (C, E, T_A, T_B, S, A, H)^T$ is the vector of state fields, and the operator \mathcal{M} is defined by the system in Section 2.1. On the basis of the numerical solution, the observable functionals

$$\hat{y}^{(p)} = \mathcal{G}(U^{(p)}), \quad (30)$$

are computed and compared with $y_{\text{obs}}^{(p)}$.

In general form, the parameter-estimation problem is written as

$$\tilde{\theta}_{\text{pat}}^{(p)} = \arg \min_{\theta \in \tilde{\Theta}^{\text{pat}}} \left[\mathcal{L}(\mathcal{G}(U(\theta)), y_{\text{obs}}^{(p)}) + \mathcal{R}(\theta) \right], \quad (31)$$

where \mathcal{L} is the mismatch function between the model-derived and observed quantities, whereas \mathcal{R} is a regularisation term that constrains biologically implausible parameter values.

Because the clinical data available from the literature are incomplete, personalisation in the present study is necessarily partial. Not all components of $\tilde{\theta}^{\text{pat}}$ can be reliably inferred for every patient. Therefore, the primary focus is placed on those effective parameters that are both influential for the model and supported by the available variables described in Section 2.4.

In this formulation, machine learning methods do not replace the mechanistic model, but instead approximate the mapping

$$\mathcal{F}_{\phi} : \mathcal{X} \rightarrow \tilde{\Theta}^{\text{pat}}, \quad (32)$$

which maps the patient's observable features to the reduced set of latent parameters. Accordingly, the overall study pipeline takes the form

$$x^{(p)} \xrightarrow{\mathcal{F}_{\phi}} \tilde{\theta}_{\text{pat}}^{(p)} \xrightarrow{\mathcal{M}} U^{(p)} \xrightarrow{\mathcal{G}} \hat{y}^{(p)}.$$

At the first stage, the latent parameters are estimated; at the second, the individualised dynamics of the system are computed; and at the third, clinically interpretable output characteristics are extracted. The overall logic of the proposed personalisation framework is schematically shown in Figure 2.

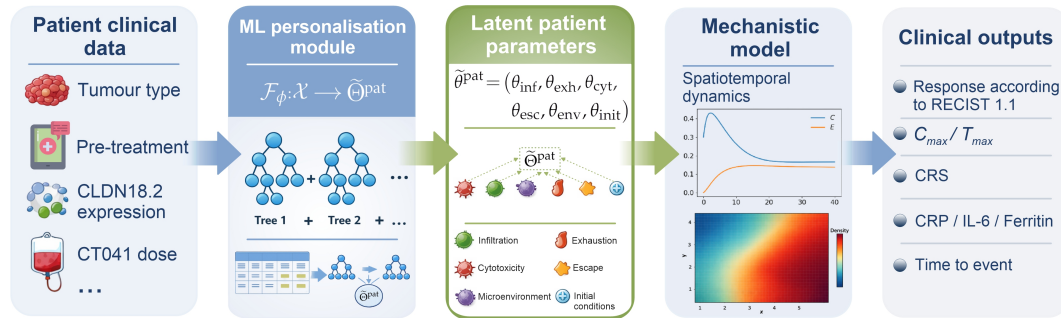


Figure 2. General scheme of the hybrid personalisation pipeline for the spatiotemporal model of CAR-T therapy based on patient clinical data.

2.4. Clinical Data Source and Variable Preparation

As the primary literature source, we used a clinically coherent study [46], which reported interim results of a study of CT041, a CLDN18.2-specific CAR-T therapy, in patients with previously treated tumours of the digestive system. This choice was driven not by an intention to confine the analysis to a single study, but by the need to minimise inter-study heterogeneity in reporting formats while at the same time preserving a minimally sufficient set of observations for the partial constraint of latent parameters in the mechanistic model and their subsequent clinical interpretation.

Only those variables that could be unambiguously matched to model variables or used as personalisation features were included in the working dataset. These comprised tumour type, the level of CLDN18.2 expression, information on prior therapy, the CT041 dose, response category according to RECIST 1.1, measures of objective response rate (ORR), disease control rate (DCR), progression-free

survival (PFS), and overall survival (OS), the grade of cytokine release syndrome (CRS), characteristics of peripheral CAR-T cell expansion and persistence, as well as peak post-infusion inflammatory markers. Variables lacking a clear quantitative interpretation, temporal reference, or specification of measurement units were excluded from the analysis.

Information extraction was performed from the main text of the article, the tables, and the supplementary materials. Priority was given to data presented in tabular form or as explicit numerical values. Graph digitisation was used only when no tabulated representation was available. For each variable, the measurement units, temporal reference point, and level of data aggregation were recorded. After extraction, temporal characteristics were normalised relative to the first CT041 infusion, response categories were coded according to RECIST 1.1, survival endpoints were expressed in months, and the drug dose was retained as the absolute number of administered CAR-T cells.

A separate harmonisation procedure was applied to variables related to antigen expression and cellular kinetics. Because CLDN18.2 positivity in the original study was defined immunohistochemically as a staining intensity of at least 2+ in no fewer than 40% of tumour cells, this variable was encoded as an ordinal categorical variable. For the quantitative analysis of CAR-T cell kinetics, the primary metrics used were C_{max} , T_{max} , and persistence duration. Inflammatory markers, including CRP, IL-6, and ferritin, were considered in terms of their peak post-infusion values.

The use of a single clinically coherent data source at this stage reduced variability associated with inter-study differences in design, endpoint composition, and completeness of reporting, thereby providing a more homogeneous basis for the initial feasibility assessment of the hybrid pipeline. Missing values were handled conservatively. When exact numerical values were unavailable, neither model-based nor multiple imputation was applied; instead, the corresponding variables were excluded from quantitative calibration and used only for qualitative interpretation.

As a result, a harmonised set of variables was assembled, including tumour type, pretreatment characteristics, CT041 dose, the level of CLDN18.2 expression, response category, ORR/DCR, PFS and OS, CRS grade, peak inflammatory markers, and characteristics of peripheral CAR-T cell kinetics. This dataset provides a reproducible foundation for population-level calibration of the mechanistic core, limited inference of latent patient-specific parameters, and their subsequent mechanistic interpretation.

After variable selection, harmonisation, and rescaling to a comparable range, the resulting patient-oriented feature set was further analysed by means of a heat map (Figure 3). This visualisation provides an intuitive overview of inter-patient variability and the structure of the harmonised set of clinical features across the cohort as a whole, and introduces the subsequent machine-learning stage, in which this same feature set is used to infer the reduced patient-specific parameters of the model.




Figure 3. Heat map of the harmonised set of patient clinical features used for personalisation of the spatiotemporal model of CAR-T therapy. Colour indicates normalised feature values.

2.5. Machine Learning for Inference of Patient-Specific Model Parameters

In this work, the machine learning module is used to approximate the mapping

$$\mathcal{F}_\phi : \mathcal{X} \rightarrow \tilde{\Theta}_{\text{pat}},$$

which reconstructs the reduced vector of mechanistic model parameters from the patient's clinical features. The target values are the reference estimates

$$\tilde{\theta}_{\text{ref}}^{(p)} = \begin{pmatrix} \theta_{\text{inf}}^{(p)} \\ \theta_{\text{exh}}^{(p)} \\ \theta_{\text{cyt}}^{(p)} \\ \theta_{\text{esc}}^{(p)} \\ \theta_{\text{env}}^{(p)} \\ \theta_{\text{init}}^{(p)} \end{pmatrix},$$

obtained by solving the limited inverse problem for each patient.

Gradient boosting on decision trees was used as the primary method. This approach is well suited to small datasets, accommodates mixed tabular data, and is capable of capturing non-linear relationships among features. In addition, tree-based ensembles are typically more robust and interpretable than more complex neural network architectures under conditions of limited data availability. The model hyperparameters were tuned in the inner loop of nested cross-validation using `GridSearchCV`, with the mean absolute error as the optimisation criterion. The following hyperparameter grid was considered: number of trees $n_{\text{estimators}} \in \{100, 200\}$, maximum tree depth $d_{\text{max}} \in \{2, 3\}$, learning rate $\eta \in \{0.05, 0.1\}$, and the minimum number of samples per leaf $\in \{1, 2\}$, with the subsampling parameter fixed at `subsample = 0.8`. Restricting tree depth and learning rate in this way was intended to reduce the risk of overfitting in the setting of a small sample size and to preserve the stability of out-of-fold estimates.

Separate regression models were trained for each component of the reduced vector. A distinct model was constructed for each component of the reduced vector

$$f_j(\cdot; \phi_j) : \mathcal{X} \rightarrow \mathbb{R}, \quad j = 1, \dots, d,$$

so that

$$\mathcal{F}_\phi(x) = \begin{pmatrix} f_1(x; \phi_1) \\ \vdots \\ f_d(x; \phi_d) \end{pmatrix}.$$

This construction makes it possible to tune the models in accordance with differences in the identifiability and statistical properties of individual biological parameters.

To ensure compliance with biological constraints, each parameter θ_j is transformed into a continuous variable before model training. Let the admissible range of the parameter be given by the interval $[a_j, b_j]$. The following transformation is then used

$$z_j^{(p)} = \log \frac{\theta_{\text{ref},j}^{(p)} - a_j + \varepsilon}{b_j - \theta_{\text{ref},j}^{(p)} + \varepsilon},$$

where $\varepsilon > 0$ is a small constant. After prediction, the inverse transformation is applied:

$$\hat{\theta}_j^{(p)} = a_j + (b_j - a_j) \sigma(\hat{z}_j^{(p)}),$$

where $\sigma(\cdot)$ is the logistic function. This guarantees that the predicted values remain within the admissible biological domain.

To improve robustness to noise in the pseudo-labels, training is performed using the robust Huber loss function

$$\mathcal{J}_j(\phi_j) = \frac{1}{|P_j|} \sum_{p \in P_j} w_p \ell_\delta(f_j(x^{(p)}; \phi_j) - z_j^{(p)}) + \lambda \Omega(\phi_j),$$

where P_j denotes the set of patients for whom an estimate of the parameter θ_j is available, w_p denotes the observation weights, $\Omega(\phi_j)$ is the regularisation term, and ℓ_δ is the Huber loss. This formulation reduces the influence of errors associated with incomplete clinical data and the inaccuracy of the inverse problem solution.

Continuous features are log-transformed, whereas categorical and ordinal variables are retained in their original form. Hyperparameter tuning is performed exclusively on the training subsets within nested cross-validation.

After training, the predicted parameter vector

$$\hat{\theta}_{\text{pat}} = \mathcal{F}_\phi(x)$$

is passed to the mechanistic model in order to compute the system dynamics and the resulting prediction of therapeutic response. Thus, machine learning serves as a tool for personalising the latent parameters of the mechanistic model.

2.6. Validation Protocol and Performance Metrics for Latent Parameter Recovery and Baseline Prediction of Clinical Outcomes

Validation of the proposed approach in the present study is structured around two tasks, both directly reflected in Section 3. The first task is to assess the quality of out-of-fold recovery of the reduced vector of patient-specific latent parameters,

$$\tilde{\Theta}^{\text{pat}} = (\theta_{\text{inf}}, \theta_{\text{exh}}, \theta_{\text{cyt}}, \theta_{\text{esc}}, \theta_{\text{env}}, \theta_{\text{init}})$$

from the available patient clinical features. The second task is an auxiliary evaluation of a direct empirical baseline for binary clinical outcomes, aimed at determining what prognostic signal is contained in the observed features themselves, without the stage of inferring the latent parameters of the mechanistic model. The main mechanistic analysis of the latent factors underlying response and resistance, presented later in the Results section, is regarded as an interpretative stage and is not accompanied by a separate formal performance metric.

Because of the small sample size, evaluation is performed at the patient level using nested cross-validation. In the outer loop, K -fold cross-validation with patient-specific splitting is applied. All data from a given patient are assigned entirely either to the training set or to the test set, thereby preventing information leakage. Within the training portion of each outer split, all stages of model development are carried out, including feature transformation, handling of missing values, hyperparameter tuning, and training. Any preprocessing statistics are estimated exclusively on the training subset of the corresponding fold and are then applied to the test patients without further adjustment.

For the outer split m , the mapping

$$\mathcal{F}_\phi^{(m)} : \mathcal{X} \rightarrow \tilde{\Theta}^{\text{pat}} \quad (33)$$

is trained. For a test patient $p \in \mathcal{D}_{\text{test}}^{(m)}$, an out-of-fold estimate of the latent parameters is computed from the feature vector $x^{(p)}$ as

$$\hat{\theta}_{\text{pat}}^{(p)} = \Pi_{\mathcal{A}}(\mathcal{F}_\phi^{(m)}(x^{(p)})), \quad (34)$$

where $\Pi_{\mathcal{A}}$ denotes the projection onto the set of biologically admissible values. The predicted values are compared with the reference estimates $\tilde{\theta}_{\text{ref}}^{(p)}$, obtained independently by solving the limited inverse problem using the patient's available observable data. These reference values are interpreted not as true biological parameters, but as pseudo-labels consistent with the mechanistic model and the available measurements.

For component-wise evaluation of the recovery quality of each latent component θ_j , four metrics are used. The first metric is the mean absolute error

$$\text{MAE}_{\theta_j} = \frac{1}{N_{\text{test}}} \sum_p \left| \hat{\theta}_j^{(p)} - \theta_{\text{ref},j}^{(p)} \right|. \quad (35)$$

The second metric is the root mean square error

$$\text{RMSE}_{\theta_j} = \sqrt{\frac{1}{N_{\text{test}}} \sum_p \left(\hat{\theta}_j^{(p)} - \theta_{\text{ref},j}^{(p)} \right)^2}. \quad (36)$$

The third metric is the Spearman rank correlation coefficient ρ_{θ_j} , which characterises how well the inter-patient ranking of the corresponding latent component is preserved. The fourth metric is the coefficient of determination

$$R_{\theta_j}^2 = 1 - \frac{\sum_p \left(\hat{\theta}_j^{(p)} - \theta_{\text{ref},j}^{(p)} \right)^2}{\sum_p \left(\theta_{\text{ref},j}^{(p)} - \bar{\theta}_{\text{ref},j} \right)^2}, \quad (37)$$

where $\bar{\theta}_{\text{ref},j}$ denotes the mean value of the reference component across the test patients. The joint use of ρ , MAE, RMSE, and R^2 makes it possible to assess not only the absolute error, but also the preservation of rank structure and the proportion of explained variability.

In the second validation block, a direct empirical baseline is considered for two binary endpoints: achievement of an objective response and disease control. Its role is limited to that of an external reference point and does not replace the main mechanism-oriented study pipeline. The quality of binary prediction is evaluated on the basis of out-of-fold predictions using balanced accuracy, the F_1 score, AUROC, and AUPRC. Balanced accuracy is used as a more robust measure under conditions of potential class imbalance; the F_1 score characterises the trade-off between recall and precision for the positive class; AUROC assesses ranking quality across the full range of thresholds; and AUPRC is particularly informative for rare positive outcomes.

In addition, for each binary endpoint, agreement between the observed and predicted cohort-level frequency of the positive outcome is analysed. For this purpose, the observed proportion of patients with a positive outcome is compared with the corresponding proportion reconstructed from the out-of-fold predictions of the baseline model. This assessment does not replace individual classification metrics, but allows evaluation of how well the model reproduces the aggregate event frequency at the level of the entire cohort.

All final estimates are calculated at the patient level from out-of-fold predictions. For the metrics of the direct empirical baseline, 95% confidence intervals obtained by patient-level bootstrap estimation are additionally reported. Under conditions of a small sample size, the main emphasis is placed on comparing metric magnitudes across tasks and latent components, rather than on formal testing of differences among a large number of competing models.

3. Results

3.1. Inference of Patient-Specific Latent Model Parameters

At the first stage, the quality of the patient-specific latent representation, expressed as a reduced vector, was evaluated $\tilde{\Theta}_{\text{pat}} = (\theta_{\text{inf}}, \theta_{\text{exh}}, \theta_{\text{cyt}}, \theta_{\text{esc}}, \theta_{\text{env}}, \theta_{\text{init}})$.

The reference estimates of the latent parameters demonstrated pronounced inter-patient heterogeneity (Figure 4). The latent profiles differed substantially across patients and could not be reduced to a single averaged pattern, indicating the heterogeneity of the resulting latent space.




Figure 4. Reference estimates of the patient-specific latent model parameters. Rows correspond to individual patients, columns correspond to the components of the reduced vector $\hat{\Theta}_{pat}$, and colour indicates the normalised value of the corresponding parameter.

The quality of out-of-fold recovery of the latent parameters from the observed clinical features proved to be heterogeneous across the components of the latent vector (Figure 5 and Table 2). The best results were obtained for the parameter θ_{env} ($\rho = 0.858$, MAE = 0.021, RMSE = 0.031, $R^2 = 0.831$). Good agreement between the reference and predicted values was also observed for θ_{exh} ($\rho = 0.585$, $R^2 = 0.388$). For θ_{inf} and θ_{cyt} , recovery quality was moderate. Despite positive rank correlations ($\rho = 0.569$ and 0.563 , respectively), the coefficient of determination remained limited ($R^2 = 0.182$ and 0.230). For θ_{esc} , predictive performance remained low ($R^2 = -0.027$), whereas θ_{init} did not show reliable agreement between the reference and predicted estimates ($\rho = -0.023$, $R^2 = -0.357$).

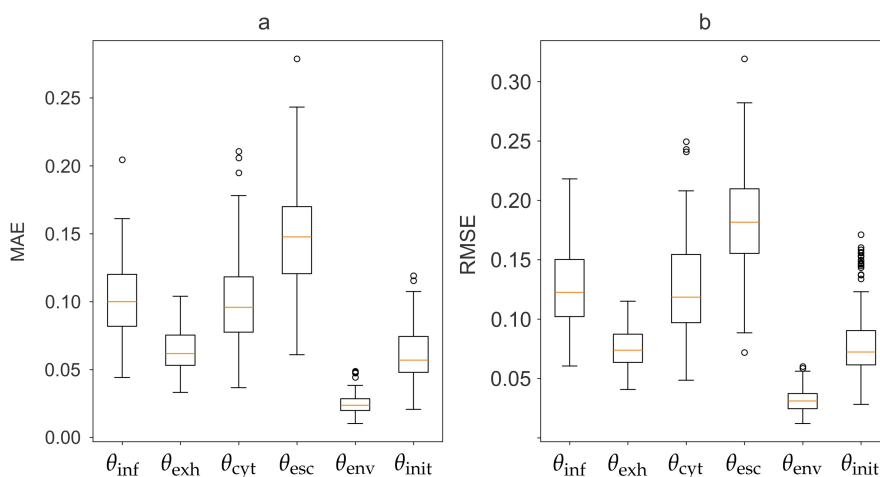


Figure 5. Distributions of out-of-fold recovery errors for the latent model parameters: (a) box plots of MAE values for the individual components of the reduced vector; (b) box plots of RMSE values for the same parameters.

The error distributions shown in Figure 5 further confirm differences in the reproducibility of individual latent components. The lowest MAE and RMSE values were observed for θ_{env} , whereas θ_{esc} exhibited the highest median errors and the widest spread. For θ_{exh} , the error distributions remained comparatively narrow, while for θ_{inf} and θ_{cyt} they occupied an intermediate position.

Table 2. Summary performance metrics for out-of-fold recovery of the latent model parameters.

Parameter	ρ	MAE	RMSE	R^2
θ_{inf}	0.569	0.100	0.130	0.182
θ_{exh}	0.585	0.062	0.076	0.388
θ_{cyt}	0.563	0.095	0.126	0.230
θ_{esc}	0.300	0.154	0.192	-0.027
θ_{env}	0.858	0.021	0.031	0.831
θ_{init}	-0.023	0.068	0.100	-0.357

The results show that recovery accuracy differs substantially across latent components. In the current configuration, the most reproducible parameters are those associated with the microenvironment and exhaustion, whereas reliable recovery was not achieved for the parameters related to antigen escape and, in particular, the individualised initial conditions.

3.2. Auxiliary Empirical Baseline for Binary Clinical Outcomes

We next consider an auxiliary empirical baseline that makes it possible to assess the extent to which the available observable features themselves contain signal for distinguishing binary clinical outcomes at the patient level. The analysis was performed for two endpoints – achievement of an objective response and disease control – and served as an external reference point for the subsequent interpretation of the contribution of mechanistic latent-parameter inference. Baseline performance was evaluated in an out-of-fold setting using balanced accuracy, F_1 score, AUROC, and AUPRC (Figure 6).

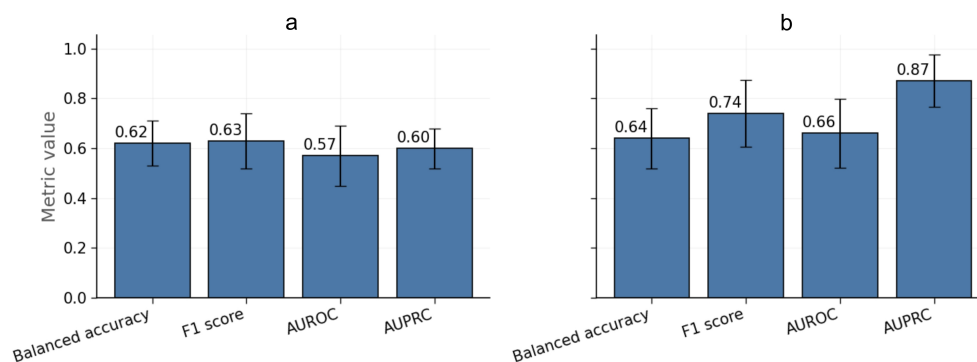


Figure 6. Out-of-fold predictive performance for binary clinical outcomes using the direct empirical approach: (a) for the endpoint of achieving an objective response; (b) for the endpoint of achieving disease control. The bars correspond to mean metric values, and the error bars represent 95% confidence intervals obtained by bootstrap estimation.

For the endpoint related to achievement of an objective response, classification performance remained limited: balanced accuracy was 0.62, the F_1 score was 0.63, AUROC was 0.57, and AUPRC was 0.60. For the disease-control endpoint, the metrics were higher: balanced accuracy was 0.64, the F_1 score was 0.74, AUROC was 0.66, and AUPRC was 0.87. Thus, disease control was predicted from the observed features substantially better than objective response.

Agreement between the observed and predicted cohort-level frequencies was additionally evaluated (Figure 7). For the endpoint of achieving an objective response, the observed proportion of patients with a positive outcome was 0.49, whereas the predicted frequency was 0.51. For the endpoint of achieving disease control, the corresponding values were 0.73 and 0.76, respectively.




Figure 7. Agreement between observed and predicted cohort-level frequencies for binary clinical outcomes: (a) for the endpoint of achieving an objective response; (b) for the endpoint of achieving disease control. Shown are the observed proportion of patients with a positive outcome and the corresponding proportion reconstructed from the out-of-fold predictions.

Overall, the direct empirical approach showed that the observed clinical features carry only limited signal for individual prediction of the more stringent endpoint of objective response, whereas for disease control this signal is somewhat more robust. Thus, the auxiliary baseline should not be regarded as a standalone final result of the study, but rather as an external reference point against which the principal value of the proposed approach lies in translating clinical heterogeneity into a space of interpretable latent mechanisms.

3.3. Mechanistic Interpretation of Latent Factors Underlying Response and Resistance

After evaluating the quality of recovery of the patient-specific latent parameters and considering the auxiliary empirical baseline, the latent space was then interpreted in terms of the mechanisms underlying response and resistance. To this end, the values of the components of the reduced vector $\tilde{\Theta}_{pat} = (\theta_{inf}, \theta_{exh}, \theta_{cyt}, \theta_{esc}, \theta_{env}, \theta_{init})$ were compared across patients with positive and negative clinical outcomes. Effect size was assessed using Cliff's delta separately for objective response and for disease control (Figure 8).




Figure 8. Latent mechanisms associated with clinical response and resistance. For each component of the reduced vector $\tilde{\Theta}_{pat}$, the effect size according to Cliff's delta is shown for the comparison between patients with positive and negative outcomes. Positive values correspond to higher parameter values in the group with a favourable outcome, whereas negative values correspond to higher parameter values in the group without response or without disease control. The error bars represent 95% confidence intervals.

For objective response, positive effects were observed primarily for the parameters related to CAR-T cell cytotoxic potency and infiltration, whereas negative effects were found for the parameters associated with exhaustion, antigen escape, microenvironmental suppressiveness, and initial conditions. Accordingly, the probability of response increases when preserved effector function is combined with effective tumour infiltration by CAR-T cells, whereas resistance is associated with intensified immunosuppression, loss of the antigenic target, and dysfunction of the effector pool.

For disease control, the direction of the effects was preserved, although the contrast between the groups was less pronounced. The parameters most consistently associated with lack of disease control were those related to antigen escape and unfavourable initial conditions, whereas the contributions of infiltration and cytotoxic potency remained positive. This is consistent with the notion that achieving disease control requires a less stringent combination of favourable mechanisms than does the development of an objective response.

For a two-dimensional interpretation of the latent space, the landscape of the predicted response probability was examined in the $(\theta_{inf}, \theta_{env})$ plane, reflecting the balance between effective CAR-T cell infiltration and microenvironmental suppressiveness (Figure 9).




Figure 9. Landscape of the predicted probability of response in a two-dimensional projection of the latent space. Colour indicates the predicted probability of response in the $(\theta_{inf}, \theta_{env})$ plane. Triangles correspond to patients with an objective response, whereas circles correspond to patients without an objective response. The light contour lines indicate levels of equal predicted probability.

The probability of response is determined not by any single parameter, but by the patient's position within the latent space (Figure 9). More favourable regions correspond to the combination of high infiltration and a less pronounced suppressive microenvironment, whereas displacement towards the opposite region is accompanied by a lower probability of response. The separation between responders and non-responders is continuous rather than discrete, indicating the existence of intermediate states between sensitivity and resistance.

Thus, the latent space of the model reflects mechanistically interpretable differences among patients. Response is associated with higher CAR-T cell infiltration and cytotoxic potency, whereas resistance is associated with exhaustion, antigen escape, and a suppressive microenvironment.

4. Discussion

The results obtained indicate that hybrid personalisation of a spatiotemporal model of CAR-T therapy in solid tumours is feasible, although the informativeness of clinical features is not uniform across different latent mechanisms. The parameters most robustly recovered were those associated with the immunosuppressive microenvironment and CAR-T cell exhaustion, whereas antigen escape and individualised initial conditions remained substantially less identifiable. The integral effects of the microenvironment and dysfunction are reflected across several observable channels, including CAR-T cell kinetics, inflammatory markers, and toxicity, whereas spatial antigen heterogeneity and the initial intratumoural architecture manifest themselves only indirectly in aggregated clinical data [6,7,11,21,26,47]. Accordingly, differences in recovery accuracy reflect not only algorithmic limitations, but also the fundamental incompleteness of the observations.

The latent space in the proposed model does not impose a rigid division of patients into responders and non-responders, but rather defines a continuous landscape of sensitivity and resistance. This indicates that response to CAR-T therapy in solid tumours is determined by the combined effects of infiltration, preservation of effector function, antigenic stability, and the level of local suppression, rather than by a single dominant mechanism [10,12,17]. Within this framework, a higher probability of response corresponds to the combination of effective infiltration and a less pronounced suppressive microenvironment, whereas resistance is associated with intensified exhaustion, antigen escape, and unfavourable initial conditions. This interpretation is consistent with the literature, in which inadequate CAR-T cell trafficking, chemokine mismatch, hypoxia, and inhibitory signalling are regarded as key barriers to therapy in solid tumours [16,18,19,47].

The lower predictability of objective response than of disease control is unsurprising. Disease control may be achieved under conditions of partial suppression of tumour growth, whereas objective response requires a deeper degree of cytoreduction and, consequently, more precise information on the spatial distribution of the target, the local density of infiltration, and the selection of antigen-negative clones. This is consistent both with the clinical data on CLDN18.2-specific CAR-T therapy [28,46] and with current understanding of the role of antigen escape and functional T-cell dysregulation in the development of resistance [13,14,23,25,48]. From the perspective of mathematical modelling, the results extend the line of work in which interpretable parameters are used to decompose the contributions of proliferation, exhaustion, spatial constraints, and antigen escape, while advancing it to the level of limited patient-specific adaptation [31,32,34–37].

From a theoretical perspective, the study supports the rationale for hybrid frameworks in which machine learning does not replace the mechanistic model, but is instead used to infer latent parameters from incomplete clinical data [40,42,43]. From a practical perspective, this makes it possible to regard latent parameters as a language for stratifying mechanisms of resistance: insufficient infiltration points to the potential value of strategies that enhance CAR-T cell trafficking and retention, dominance of microenvironmental suppression and exhaustion indicates the need for armoured constructs or combined immunomodulation, and a high risk of antigen escape highlights the relevance of multi-antigen platforms [16,17,19,47,48]. In a broader context, this approach is consistent with the development of patient-specific digital twins in oncology and digital medicine [49–51].

The present study should be interpreted in the context of an early stage in the development of a hybrid patient-specific model of CAR-T therapy for solid tumours. The cohort size and the predominant use of a single clinically coherent data source limit the statistical generalisability of the results; however, at this stage, such a design made it possible to minimise inter-study heterogeneity and to assess the internal consistency of the proposed personalisation pipeline. Importantly, the latent patient-specific parameters in the problem considered here are not directly observable; therefore, the study used surrogate mechanistic target variables obtained within the framework of a limited inverse problem, and these should be regarded as an operational approximation rather than as independent ground-truth labels. In addition, spatial intratumoural processes are reflected only indirectly in the available clinical data. Consequently, the results presented here primarily support the feasibility and

mechanistic interpretability of the hybrid approach, rather than its readiness for clinical application. In this sense, the study constitutes a hypothesis-generating proof-of-concept framework that provides a foundation for subsequent external validation in independent cohorts and for expansion of the set of multimodal observations [27,28,49–51].

5. Conclusions

The aim of the present study was to develop and evaluate a hybrid framework for patient-specific calibration of a spatiotemporal model of CAR-T therapy in solid tumours, in which a reaction-diffusion-chemotaxis mechanistic core is complemented by machine learning to infer latent parameters from clinical data.

The results show that the proposed hybrid framework makes it possible to partially infer biologically interpretable latent parameters of a spatiotemporal model of CAR-T therapy from a limited set of clinical observations and to use them for the mechanistic interpretation of response and resistance scenarios. The parameters most reliably identified were those associated with the immunosuppressive microenvironment and CAR-T cell exhaustion, whereas antigen escape and individualised initial conditions remained substantially less observable. The auxiliary direct baseline model based on clinical features showed that, in the absence of mechanistic structure, the more stringent clinical endpoint exhibits only limited predictability. Thus, the current version of the approach demonstrates the feasibility and interpretability of limited patient-specific personalisation, but requires further independent external validation, expansion of the spectrum of tumour targets, and integration of more spatially informative data sources.

Funding: The study was funded by the Russian Science Foundation grant No. 25-71-00041, <https://rscf.ru/project/25-71-00041/> (accessed on 6 April 2026).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the author.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Maude, S.L.; Frey, N.; Shaw, P.A.; Aplenc, R.; Barrett, D.M.; Bunin, N.J.; Chew, A.; Gonzalez, V.E.; Zheng, Z.; Lacey, S.F.; et al. Chimeric antigen receptor T cells for sustained remissions in leukemia. *N. Engl. J. Med.* **2014**, *371*, 1507–1517.
2. Maude, S.L.; Laetsch, T.W.; Buechner, J. et al. Tisagenlecleucel in children and young adults with B-cell lymphoblastic leukemia. *N. Engl. J. Med.* **2018**, *378*(5), 439–448.
3. Neelapu, S.S.; Locke, F.L.; Bartlett, N.L. et al. Axicabtagene ciloleucel CAR T-cell therapy in refractory large B-cell lymphoma. *N. Engl. J. Med.* **2017**, *377*(26), 2531–2544.
4. June, C.H.; O'Connor, R.S.; Kawalekar, O.U. et al. CAR T cell immunotherapy for human cancer. *Science* **2018**, *359*(6382), 1361–1365.
5. Sadelain, M.; Brentjens, R.; Rivière, I. The basic principles of chimeric antigen receptor design. *Cancer Discov.* **2013**, *3*(4), 388–398.
6. Newick, K.; O'Brien, S.; Moon, E.; Albelda, S.M. CAR T cell therapy for solid tumors. *Annu. Rev. Med.* **2017**, *68*, 139–152.
7. Wagner, J.; Wickman, E.; DeRenzo, C.; Gottschalk, S. CAR T cell therapy for solid tumors: Bright future or dark reality? *Mol. Ther.* **2020**, *28*(11), 2320–2339.
8. Li, J.; Li, W.; Huang, K. et al. Chimeric antigen receptor T cell (CAR-T) immunotherapy for solid tumors: Lessons learned and strategies for moving forward. *J. Hematol. Oncol.* **2018**, *11*(1), 22.
9. Filley, A.C.; Henriquez, M.; Dey, M. CART immunotherapy: Development, success, and translation to malignant gliomas and other solid tumors. *Front. Oncol.* **2018**, *8*, 453.
10. Ramakrishna, S.; Highfill, S.L.; Walsh, Z. et al. Prospects and challenges for use of CAR T cell therapies in solid tumors. *Expert Opin. Biol. Ther.* **2020**, *20*(5), 503–516.
11. Rafiq, S.; Hackett, C.S.; Brentjens, R.J. Engineering strategies to overcome the current roadblocks in CAR T cell therapy. *Nat. Rev. Clin. Oncol.* **2020**, *17*(3), 147–167.

12. Kirtane, K.; Elmariah, H.; Chung, C.H.; Abate-Daga, D. Adoptive cellular therapy in solid tumor malignancies: Review of the literature and challenges ahead. *J. Immunother. Cancer* **2021**, *9*(7), e002723.
13. Majzner, R.G.; Mackall, C.L. Tumor antigen escape from CAR T-cell therapy. *Cancer Discov.* **2018**, *8*(10), 1219–1226.
14. Labanieh, L.; Majzner, R.G.; Mackall, C.L. Programming CAR-T cells to kill cancer. *Nat. Biomed. Eng.* **2018**, *2*(6), 377–391.
15. Petty, A.J.; Heyman, B.; Yang, Y. Chimeric antigen receptor cell therapy: Overcoming obstacles to battle cancer. *Cancers* **2020**, *12*(4), 842.
16. Tian, Y.; Li, Y.; Shao, Y.; Zhang, Y. Gene modification strategies for next-generation CAR T cells against solid cancers. *J. Hematol. Oncol.* **2020**, *13*(1), 54.
17. Andrea, A.E.; Chiron, A.; Mallah, S. et al. Advances in CAR-T cell genetic engineering strategies to overcome hurdles in solid tumors treatment. *Front. Immunol.* **2022**, *13*, 830292.
18. Zhang, P.-F.; Wang, C.; Zhang, L.; Li, Q. Reversing chemokine/chemokine receptor mismatch to enhance the antitumor efficacy of CAR-T cells. *Immunotherapy* **2022**, *14*(6), 459–473.
19. Adachi, K.; Kano, Y.; Nagai, T. et al. IL-7 and CCL19 expression in CAR-T cells improves immune cell infiltration and CAR-T cell survival in the tumor. *Nat. Biotechnol.* **2018**, *36*(4), 346–351.
20. Polyakov, M.V.; Ten, V.V. Simulation tumor growth in heterogeneous medium based on diffusion equation. *Int. J. Mod. Phys. C* **2024**, *35*(1), 2450010.
21. Wherry, E.J.; Kurachi, M. Molecular and cellular insights into T cell exhaustion. *Nat. Rev. Immunol.* **2015**, *15*(8), 486–499.
22. Schietinger, A.; Greenberg, P.D. Tolerance and exhaustion: Defining mechanisms of T cell dysfunction. *Trends Immunol.* **2014**, *35*(2), 51–60.
23. Thommen, D.S.; Schumacher, T.N. T cell dysfunction in cancer. *Cancer Cell* **2018**, *33*(4), 547–562.
24. van der Leun, A.M.; Thommen, D.S.; Schumacher, T.N. CD8⁺ T cell states in human cancer: Insights from single-cell analysis. *Nat. Rev. Cancer* **2020**, *20*(4), 218–232.
25. Dolina, J.S.; Van Braeckel-Budimir, N.; Thomas, G.D.; Salek-Ardakani, S. CD8⁺ T cell exhaustion in cancer. *Front. Immunol.* **2021**, *12*, 715234.
26. Schurich, A.; Magalhaes, I.; Mattsson, J. Metabolic regulation of CAR T cell function by the hypoxic microenvironment in solid tumors. *Immunotherapy* **2019**, *11*(4), 335–345.
27. Brown, C.E.; Hibbard, J.C.; Alizadeh, D. et al. Locoregional delivery of IL-13R α 2-targeting CAR-T cells in recurrent high-grade glioma: A phase 1 trial. *Nat. Med.* **2024**, *30*(4), 1001–1012.
28. Qi, C.; Liu, C.; Gong, J. et al. Claudin18.2-specific CAR T cells in gastrointestinal cancers: Phase 1 trial final results. *Nat. Med.* **2024**, *30*, 2224–2234.
29. Shi, D.; Shi, Y.; Kaseb, A.O. et al. Chimeric antigen receptor-glypican-3 T-cell therapy for advanced hepatocellular carcinoma: Results of phase I trials. *Clin. Cancer Res.* **2020**, *26*(15), 3979–3989.
30. Straathof, K.; Flutter, B.; Wallace, R. et al. Antitumor activity without on-target off-tumor toxicity of GD2-chimeric antigen receptor T cells in patients with neuroblastoma. *Sci. Transl. Med.* **2020**, *12*(571), eabd6169.
31. Mueller, K.T.; Waldron, E.; Grupp, S.A. et al. Clinical pharmacology of tisagenlecleucel in B-cell acute lymphoblastic leukemia. *Clin. Cancer Res.* **2018**, *24*(24), 6175–6184.
32. Stein, A.M.; Grupp, S.A.; Levine, J.E. et al. Tisagenlecleucel model-based cellular kinetic analysis of chimeric antigen receptor-T cells. *CPT Pharmacometrics Syst. Pharmacol.* **2019**, *8*(5), 285–295.
33. Chaudhury, A.; Zhu, X.; Chu, L. et al. Chimeric antigen receptor T cell therapies: A review of cellular kinetic-pharmacodynamic modeling approaches. *J. Clin. Pharmacol.* **2020**, *60*(S1), S147–S159.
34. Sahoo, P.; Yang, X.; Abler, D. et al. Mathematical deconvolution of CAR T-cell proliferation and exhaustion from real-time killing assay data. *J. R. Soc. Interface* **2020**, *17*(162), 20190734.
35. Li, R.; Sahoo, P.; Wang, D. et al. Modeling interaction of glioma cells and CAR T-cells considering multiple CAR T-cells bindings. *Immunoinformatics* **2023**, *9*, 100022.
36. Kara, E.; Jackson, T.L.; Jones, C. et al. Mathematical modeling insights into improving CAR T cell therapy for solid tumors with bystander effects. *NPJ Syst. Biol. Appl.* **2024**, *10*, 105.
37. Polyakov, M.V.; Tuchina, E.I. Reaction-Diffusion Model of CAR-T Cell Therapy in Solid Tumours with Antigen Escape. *Computation* **2026**, *14*(1), 3.
38. Rohrs, J.A.; Siegler, E.L.; Wang, P.; Finley, S.D. ERK activation in CAR T cells is amplified by CD28-mediated increase in CD3 ζ phosphorylation. *iScience* **2020**, *23*(4), 101023.
39. Logun, M.; Lyon, S.M.; Chow, V.A. et al. Label-free in vitro assays predict the potency of anti-disialoganglioside chimeric antigen receptor T-cell products. *Cytotherapy* **2023**, *25*(6), 670–682.

40. Butner, J.D.; et al. Mathematical modeling of cancer immunotherapy for personalized clinical translation. *Nat. Comput. Sci.* **2022**, *2*(12), 785–796.
41. Wu, C.; Lorenzo, G.; Hormuth, D.A. et al. Integrating mechanism-based modeling with biomedical imaging to build practical digital twins for clinical oncology. *Biophys. Rev.* **2022**, *3*(2), 021304.
42. Procopio, A.; Cesarelli, G.; Donisi, L. et al. Combined mechanistic modeling and machine-learning approaches in systems biology – a systematic literature review. *Comput. Methods Programs Biomed.* **2023**, *240*, 107681.
43. Lorenzo, G.; Ahmed, S.R.; Hormuth, D.A. et al. Patient-specific, mechanistic models of tumor growth incorporating artificial intelligence and big data. *Annu. Rev. Biomed. Eng.* **2024**, *26*(1), 529–560.
44. Polyakov, M. Integration of Machine Learning to Personalise a Mathematical Model of Tumour Dynamics Based on Reaction-Diffusion Equations. In *2025 7th International Conference on Control Systems, Mathematical Modeling, Automation and Energy Efficiency (SUMMA)*; Russian Federation: Lipetsk, Russia, **2025**; pp. 776–779.
45. Polyakov, M. Spatiotemporal Modelling of CAR-T Cell Therapy in Solid Tumours: Mechanisms of Antigen Escape and Immunosuppression. *Computation* **2026**, *14*(4), 87.
46. Qi, C., Gong, J., Li, J. et al. Claudin18.2-specific CAR T cells in gastrointestinal cancers: Phase 1 trial interim results. *Nat. Med.* **2022**, *28*, 1189–1198.
47. Taylor, C.A.; Glover, M.; Maher, J. CAR-T cell technologies that interact with the tumour microenvironment in solid tumours. *Expert Rev. Clin. Immunol.* **2024**, *20*(8), 849–871.
48. Lin, H.; Yang, X.; Ye, S.; Huang, L.; Mu, W. Antigen escape in CAR-T cell therapy: Mechanisms and overcoming strategies. *Biomed. Pharmacother.* **2024**, *178*, 117252.
49. Katsoulakis, E.; Wang, Q.; Wu, H.; Shahriyari, L.; Fletcher, R.; Liu, J.; Achenie, L.; Liu, H.; Jackson, P.; Xiao, Y.; et al. Digital twins for health: A scoping review. *npj Digit. Med.* **2024**, *7*, 77.
50. De Domenico, M.; Allegri, L.; Caldarelli, G.; d’Andrea, V.; Di Camillo, B.; Rocha, L.M.; Rozum, J.; Sbarbati, R.; Zambelli, F. Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *npj Digit. Med.* **2025**, *8*, 37.
51. Aghamiri, S.S.; Amin, R. The Potential Use of Digital Twin Technology for Advancing CAR-T Cell Therapy. *Curr. Issues Mol. Biol.* **2025**, *47*(5), 321.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.