



A multiscale orchestrated computational framework to reveal emergent phenomena in neuroblastoma

C. Borau^{a,*}, K.Y. Wertheim^{b,g}, S. Hervás-Raluy^a, D. Sainz-DeMena^a, D. Walker^b, R. Chisholm^b, P. Richmond^b, V. Varella^{c,d}, M. Viceconti^{c,d}, A. Montero^e, E. Gregori-Puigjané^e, J. Mestres^e, M. Kasztelnik^f, J.M. García-Aznar^a

^a Multiscale in Mechanical and Biological Engineering (M2BE), Aragon Institute of Engineering Research (I3A), Mechanical Engineering Department, University of Zaragoza, Zaragoza, Spain

^b Department of Computer Science and Insigneo Institute for In Silico Medicine, University of Sheffield, Sheffield, United Kingdom

^c Department of Industrial Engineering, Alma Mater Studiorum – University of Bologna, Bologna, Italy

^d Medical Technology Lab, IRCCS Istituto Ortopedico Rizzoli, Bologna, Italy

^e Chemotargets SL, Baldi Reixac 4, Parc Científic de Barcelona (PCB), Barcelona, Spain

^f ACC Cyfronet, AGH University of Science and Technology, Kraków, Poland

^g Centre of Excellence for Data Science, Artificial Intelligence and Modelling and School of Computer Science, University of Hull, Kingston upon Hull, United Kingdom

ARTICLE INFO

Keywords:

Neuroblastoma
Cancer
Computational modelling
Multi-scale
Framework
Cloud-based decision support system

ABSTRACT

Neuroblastoma is a complex and aggressive type of cancer that affects children. Current treatments involve a combination of surgery, chemotherapy, radiotherapy, and stem cell transplantation. However, treatment outcomes vary due to the heterogeneous nature of the disease. Computational models have been used to analyse data, simulate biological processes, and predict disease progression and treatment outcomes. While continuum cancer models capture the overall behaviour of tumours, and agent-based models represent the complex behaviour of individual cells, multiscale models represent interactions at different organisational levels, providing a more comprehensive understanding of the system. In 2018, the PRIMAGE consortium was formed to build a cloud-based decision support system for neuroblastoma, including a multi-scale model for patient-specific simulations of disease progression. In this work we have developed this multi-scale model that includes data such as patient's tumour geometry, cellularity, vascularization, genetics and type of chemotherapy treatment, and integrated it into an online platform that runs the simulations on a high-performance computation cluster using Onedata and Kubernetes technologies. This infrastructure will allow clinicians to optimise treatment regimens and reduce the number of costly and time-consuming clinical trials. This manuscript outlines the challenging framework's model architecture, data workflow, hypothesis, and resources employed in its development.

1. Introduction

Neuroblastoma is a type of cancer that affects the sympathetic nervous system, specifically the adrenal glands and nerve tissues in the neck, chest, abdomen, and pelvis. It is the most common extra-cranial solid tumour in children, with most cases occurring under the age of two [37]. In the embryo, there is a transient structure called the neural crest, which comprises the stem and progenitor cells that will go on to form the sympathetic nervous system [28]. If the MYCN oncogene gene is amplified in them and they also acquire activating mutations in the ALK oncogene, they will become neuroblastoma cells [28]. However, it

is important to note that this combination only characterises one fraction of neuroblastomas. Some neuroblastomas are characterised by TERT rearrangement, ATRX inactivation, and mutations in p53 [1]. Furthermore, even this phenotypic classification system is only a model, so the actual population structure might be more complex than that. This aggressive form of cancer is known for its ability to spread rapidly to other parts of the body and is considered one of the most challenging types of pediatric cancer to treat. Furthermore, it is an unpredictable disease as it can result in varying outcomes, from spontaneous regression without treatment in infants to fatal disease in older children even after intensive therapy. To address this, the International

* Corresponding author.

E-mail address: cborau@unizar.es (C. Borau).

<https://doi.org/10.1016/j.cmpb.2023.107742>

Received 27 April 2023; Received in revised form 19 July 2023; Accepted 31 July 2023

Available online 1 August 2023

0169-2607/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Neuroblastoma Risk Group was established to categorise patients based on their characteristics and tumour biology, in order to determine the most appropriate treatment approach [57]. The aim is to reduce therapy for low-risk patients to minimise long-term side effects, while intensifying and targeting treatments for high-risk patients to increase survival rates. The international risk classification considers factors such as the child's age, the stage of cancer, the type of tumour, the status of the MYCN gene, and the chromosomal abnormalities present in the cancer cells. Diagnosis of neuroblastoma is made through a combination of physical exams, imaging tests, and biopsy. For example, a more differentiated tumour with a higher degree of infiltration by Schwann cells is likely to be less malignant [45]. The phenotypic classification scheme is based on retrospective data about the relationship between a tumour's mutation profile and the patient's clinical outcome [1]. These were the main risk factors we considered when we set up the simulations reported in this paper. Treatment may include surgery, chemotherapy, radiotherapy, stem cell transplantation, or a combination of these. Usually, induction chemotherapy is applied to shrink the primary tumour, and it is then surgically removed. Next, treatment is consolidated by myeloablative chemotherapy with a stem cell transplant. The final stage is maintenance therapy, which involves anti-GD2 monoclonal antibodies, cytokines, and isotretinoin [54]. While there has been significant progress in treating neuroblastoma in recent years, it remains a complex and challenging disease to treat. For this reason, a great effort has been put into not only the development of novel therapies [48], but the use of computational tools that can support clinician decisions [39].

Digital twins are virtual models that can simulate and predict the behaviour of real-world objects or systems [27]. In cancer research, digital twins can be used to provide personalised representations of cancer patients using real-time data and high-performance computing. This can help to optimise treatment decisions, monitor responses and track lifestyle modifications for each patient. Digital twins for cancer care are an emerging and innovative approach that could transform cancer outcomes. As such, computational models, on which digital twins rely, are becoming increasingly important in the investigation of cancer due to their ability to analyse vast amounts of data, simulate complex biological processes (such as the growth and spread of cancer cells, the response of the immune system, and the pharmacokinetics of drugs) and ultimately make predictions about disease progression and treatment outcomes [7,26,44]. One of the major advantages of computational models is that they allow researchers to analyse data generated from various sources, including patient records, imaging studies, and molecular profiling. This information can be used to build models that reflect the underlying biology of cancer, including the genetic mutations that drive its development and progression, the interactions between cancer cells and the microenvironment, and the mechanisms by which treatments exert their effects, predicting how the disease will progress in different patients and under different conditions. This is of the utmost importance for the future of personalised medicine, where digital twins will be potentially used not only to simulate the effects of various treatments such as chemotherapy or radiotherapy, but also to help to optimise treatment regimens by predicting how changes in dosing or scheduling will impact patient outcomes based on their pharmacokinetics and the patient's specific characteristics, such as age, body mass index, and liver function. In addition, computational models can help to reduce the cost and time associated with preclinical and clinical trials by enabling researchers to test new treatments *in silico* before proceeding to animal studies and human trials. This can help identify promising treatments more quickly and reduce the number of costly and time-consuming clinical trials required to bring new treatments to the clinic.

Cancer is a complex disease that involves multiple levels of biological organisation, from molecules to cells to tissues to organs. Computational simulation is a powerful tool to study cancer dynamics and mechanisms across these levels, but it also poses a significant challenge: how to model everything over nine orders of magnitude in space and time? This

means that a comprehensive model of cancer would need to account for phenomena ranging from nanometres (such as DNA damage) to metres (such as metastasis) and from microseconds (such as molecular interactions) to years (such as tumour evolution). Such a model would require enormous computational resources, data integration, and novel mathematical methods and algorithms [10]. Continuum cancer models represent the behaviour of cells and their interactions with the microenvironment at a macro scale. These models are based on partial differential equations and represent the average behaviour of cells within a certain region [8,26,52]. For example, a continuum cancer model might represent the average growth rate of a tumour or the average concentration of a signaling molecule within the tumour. These models are well-suited for understanding the overall behaviour of a tumour, and for making predictions about the response of the tumour to various treatments. However, they do not capture individual cells' complex and heterogeneous behaviour, such as the importance of clonal evolution and other evolutionary principles [16], which can have important implications for treatment response and disease progression. Agent-based models, on the other hand, represent the behaviour of cells and their interactions with the microenvironment at a micro scale [11,24,33]. These models represent each cell as a discrete "agent" that can interact with its environment and with other agents. For example, an agent-based model might represent the behaviour of a single cell as it migrates through the tissue, responds to signalling molecules, and divides to form new cells. These models capture the complex and heterogeneous behaviour of individual cancer cells and explicitly include interactions between cancer cells and the microenvironment, such as the immune system (given the cancer community's attention on immunotherapy and the parallel push for a virtual immune system in the spirit of a digital twin) [31,66,67], in great detail but at the cost of tremendous computational power, especially when time and length scales span several orders of magnitude (i.e., biological phenomena happening up to the scale of a whole tumour over a period of months or even years). In response to stresses such as hypoxia and chemotherapeutic agents, a cancer cell exhibits sophisticated information processing capabilities, resulting in a change in phenotype. For example, different layers of regulation allow p53 to activate hundreds of genes in a context-dependent manner, potentially resulting in apoptosis, senescence, cell cycle arrest, and DNA repair, among other fates ([14]; P. [63]). Mathematical modelling, mostly based on ordinary differential equations, has been used extensively to understand the intracellular dynamics pertinent to p53 [29]. Within the context of neuroblastoma, most mathematical and computational models in the literature describe the subcellular scale [21,34]. This is important because neuroblastoma have many molecular aberrations. An ongoing effort within the neuroblastoma community is to identify targeted therapies: chemotherapeutic drugs targeting specific molecular aberrations, such as ALK inhibitors. It is, therefore, key to understanding how effects at the subcellular scale manifest at the tumour level [17].

Multiscale models alleviate computational costs by splitting the problem into multiple component models, each representing a phenomenon at a specific space-time scale. This allows the capture of interactions between system components at multiple levels of the organisation, from the molecular to the cellular and even the tissue level, therefore providing a more comprehensive representation of the system, and a better understanding of how changes at one level affect the others. Multiscale models can also provide a more accurate prediction of a system's behaviour, as they account for multiple sources of heterogeneity and nonlinear interactions. However, deciding how to split the scales and how the inter-scale information is shared is critical, and calibrating of the interconnected parameters is a challenging process.

In early 2019, 16 partners from eight European countries formed the PRIMAGE consortium to build a cloud-based decision support system, guided by imaging biomarkers and computational models, to manage neuroblastoma prognosis and diagnosis [36]. One of the objectives of

the project was to build a multi-scale model of neuroblastoma to enable patient-specific simulations of disease progression integrated into an online platform. Projects such as GoSmart [64], Computational Horizons in Cancer (CHIC) [47], EurValve [43], or Virtual Research Environment (VRE) [6] have made significant advancements in developing integrated, comprehensive frameworks, representing the state of the art for production-quality hybrid computational cloud and High Performance Computing (HPC) in silico processing of clinical cases. For example, the Go-Smart [47] provides an open-end software framework and simulation environment with the relevant physical and physiological modelling tools needed to correctly predict the result of minimally-invasive cancer treatment in terms of lesion size and shape, replicating the dedicated workflow used by interventional radiologists for patient specific planning. Similarly, the CHIC [47] project developed a series of hyper-multiscale models and repositories for four paradigmatic cancer types (neuroblastoma, non-small cell lung cancer, glioblastoma and prostate cancer) providing the community with a collaborative interface for exchanging knowledge and sharing work in an effective way. Eur-Valve introduced an integrated cloud/HPC computing solution dedicated to simulating valvular heart conditions, which is now utilized for clinical validation of the resulting decision support system. Their workflow provides data on pre-operative and predicted post-operative cardiac performance, enabling the clinicians to make the well-informed judgements. On the other hand, the VRE platform offers a versatile tool to process and analyse clinical data, including dashboards, interactive workbenches and [47] management features. However, PRIMAGE introduces a significant breakthrough by bridging the gap between cloud and HPC computing solutions, which have already proven successful in developing and validating in silico models, and their practical application in clinical settings. This considerably brings these advanced computing solutions closer to clinical use, unlocking their potential for improving patient care and outcome. At PRIMAGE's heart lies a multicellular model bridging neuroblastoma's single-cell and tumour scales, including phenomena spanning nine orders of magnitude in space and time. This document provides a detailed description of the various components involved in building the challenging framework that allows the model execution. Firstly, the nature and pre-processing of the input data required for the simulations. Secondly, we describe the model architecture, which is the blueprint for the entire system and is critical in determining the performance and effectiveness of the framework. This includes the data workflow, the particularisation and interpolation techniques required to communicate between scales, and the tools and computational resources utilised for the implementation. We then describe in detail each of the models composing the orchestrator: transport of species, agent-based cellular model, chemotherapy model and biomechanical model. We finalise by describing the calibration process and presenting an example patient-specific simulation case.

2. Computational methods and system description

2.1. Input data

Anonymised patient-specific inputs for the multiscale model are collected from the PRIMAGE cloud-platform infrastructure and are provided in two formats: json (data) and DICOM (image) files. More information regarding data management and security standards can be found in Supplementary Material. Several key parameters and spatial data are automatically processed to provide the inputs needed by the computational models, namely: patient clinical data (i.e., chemotherapy treatment, histology, mutations), vascularisation (K^{Trans}) [56] and cellularity maps, and tumour geometry. Vascularisation and cellularity data are extracted from dynamic contrast enhanced (DCE) and diffusion-weighted imaging (DWI) and sequences included within the platform. The patient tumour segmentation provided by the platform is automatically implemented using a custom-made Python library,

described elsewhere [49] (code and documentation publicly available), that transforms the 2D tumour slices into a 3D finite element (FE) mesh, including the vascularity/cellularity distribution, ready for computation.

2.2. Schema of the orchestrated model

The orchestration of multiple models is typically complex and requires software framework to facilitate the task. This framework must ensure centralised management of the data flow (the handling and transformation of data between models) and the control flow (the set of instructions to convey to each model throughout the orchestration). Furthermore, the setup is a challenging process both conceptually and technically. In the following sub-sections we describe in detail how we have built the machinery that allows the patient-specific simulations.

2.2.1. Framework infrastructure

Users (clinicians) can launch simulations via a friendly user interface (UI) through the patients' profiles within the PRIMAGE web platform (see Supplementary Material). Various software and data management systems have been used to make this possible. It is worth noting that the multiscale simulation is orchestrated by a bespoke tool written using, the free and open source, Python and MySQL. The component models were also developed from scratch using different languages and software (further details described in their particular sections). The orchestrator tool is integrated with the Prometheus and Ares HPC clusters, running on a CentOS-based operating system with SLURM queuing system to manage submitted jobs, which are managed by the lmod system.¹ Importantly, the orchestrator does not submit the jobs to the HPC directly, but uses Rimrock and PLGData REST API to submit the jobs (Fig. 1).

When a simulation is initiated on the PRIMAGE web, it triggers the launch of a container on the Kubernetes platform, which is responsible for executing the entire simulation. The container generates a unique execution ID at the outset and generates an HPC job script (SLURM starting script) to prepare for the execution. The container then requests the HPC cluster to commence the simulation. The HPC analyses the starting script and allocates the requested resources based on the specified requirements. Once the required resources are in place, the simulation begins. The simulation obtains the necessary inputs from the OneData system (a high-performance data management solution that offers unified data access across globally distributed environments and multiple types of underlying storage), executes the simulation, and stages the results in OneData. At the same time, the Kubernetes container continuously monitors the status of the HPC job execution. When it is complete, it fetches the outputs from OneData and uploads the results to the PRIMAGE platform directly using the platform-specific API (Fig. 1).

It is worth noting that within the Orchestrator HPC job, the four types of modules (models) work in parallel in an interdependent way, each of them periodically polling an internal database and waiting for a specific condition to be met before starting execution. When each model completes execution, output files are written and the database is updated, triggering the following model as described in the workflow (Fig. 2). Every model has a serialiser and a de-serialiser module responsible respectively for writing the input files in the required formats and registering the outputs in the database. These outputs are read by an instance of the relation module that gathers the data and processes it. Finally, the "Master & Commander" module manages all the sub-modules, addresses all the preprocessing scripts, and handles the multiscale model results.

¹ Lmod: A New Environment Module System (<https://lmod.readthedocs.io/en/>).

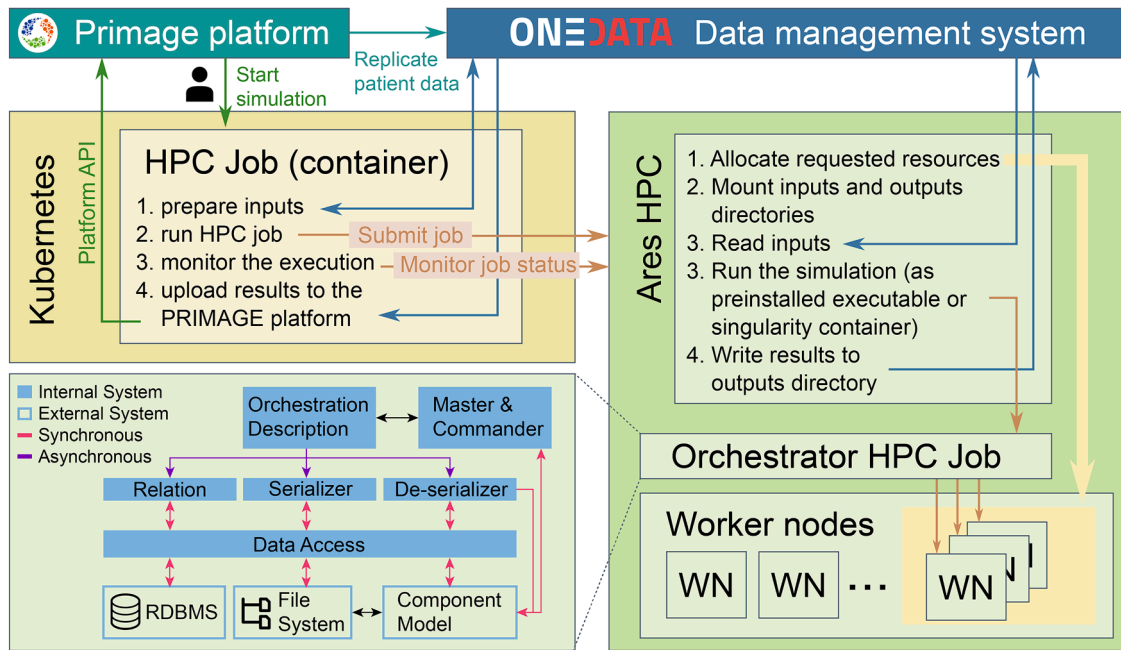


Fig. 1. PRIMAGE platform integration with the HPC cluster. When a patient's simulation is launched from the PRIMAGE web, it starts the container in the Kubernetes platform which is responsible for the whole simulation execution. The container at the beginning prepares the execution by generating a unique execution ID, generates HPC starting job script (SLURM starting script) and sends the request to the HPC cluster to start the calculation. The HPC analyses the starting script and allocates requested resources based on the requirements. Where resources are in place, the simulation starts. Simulation fetches required inputs from the OneData system, runs the actual simulation and stages the results to the OneData. Simultaneously, the Kubernetes container constantly monitors the HPC job execution status and, when it is done, fetches the outputs from the OneData and upload the results directly to the PRIMAGE platform by using the platform-specific API.

2.2.2. Data flow between models

Four interconnected component models form the core of the multiscale model presented in this work. The first, at the macro-scale, predicts via a FE scheme the diffusion-reaction of species (i.e. oxygen and nutrients) within the patient tumour depending on the vascularisation and cellularity conditions and the current geometry. The second, at the cell scale, is an agent-based model representing a few cubic millimetres of the tumour, where each cell is modelled as an autonomous agent whose behaviour depends on nutrient availability, mutations, and the presence and effects of chemotherapy. The most critical factor is, in fact, the presence and effects of chemotherapy, which are determined by a third model, at the subcellular scale, that infers them, via a machine learning scheme, based on patient clinical data and vascularity. The fourth model collates the data at the cell-level and brings it back to the macro-scale, computing the mechanical compatibility of the biomechanical deformations induced by the change in the volume occupied by the cell agents, leading to a new spatial domain and mesh for the transport model.

It is natural to assume that each finite element of the whole tumour model could be seen as a tissue unit for the agent-based cell-scale model. However, the number of elements to properly capture the details of each patient geometry could lead to unfeasible computational costs even using the most modern GPUs. For this reason, the communication between scales is done via a particularisation relation model that samples representative elements from the tumour to be simulated in detail with the agent-based model (ABM). The results are then interpolated back to all elements of the tumour geometry. A full schema of the orchestrated model including the communication between component models is depicted in Fig. 2.

2.2.3. Particularisation and interpolation

Two special types of relation modules include the processes of particularisation and interpolation, controlling the scale transition (from tumour organ to tumour cells) inside the multiscale model. These processes are responsible for reducing the computational cost of the

multiscale model by sampling representative elements to run the cell model (particularisation) and then interpolating the results obtained for all elements.

The particularisation strategy is based on binning. All elements of the mesh are separated into groups or *bins* that are defined based on one of the initial variables, which is determined before starting the simulation (e.g. oxygen level) together with the number of bins desired (degree of particularisation) (Fig. 3). The bins are built to each contain roughly the same number of elements which are allocated according to the value of a chosen variable within the corresponding bin interval. The first and last elements of each bin (the ones with the smallest/highest value) are selected as representatives for executing the ABM. Once the individual simulations finish, results are linearly interpolated to the rest of the elements within the same bin based on the min/max values and the order of elements within the group. This interpolation is done for all variables needed to run the next models and loops.

The variable determining the binning affects the quality of the interpolation and must be chosen cautiously. We found oxygen level to be the variable generating a better binning for interpolation. Ideally, the ABM response to this variable should be linear, to minimise the interpolation error (difference between interpolated results and the solution without interpolation but using all elements). Also, the number of bins impacts the estimated solution, at the cost of, mostly linearly, increasing the computational time. In our tests to find a compromise between accuracy and resources needed, the oxygen concentration presented good results with a sample size of 1%.

2.3. Component models

2.3.1. Transport of species

Transport of species at the macroscopic level is simulated via a FE model developed using the commercial software Ansys 2022 R1. The model computes the diffusion of oxygen through the tumour geometry depending on patient's K^{Trans} (a measure of vascular permeability, capillary exchange surface and blood flow) [56] as well as their

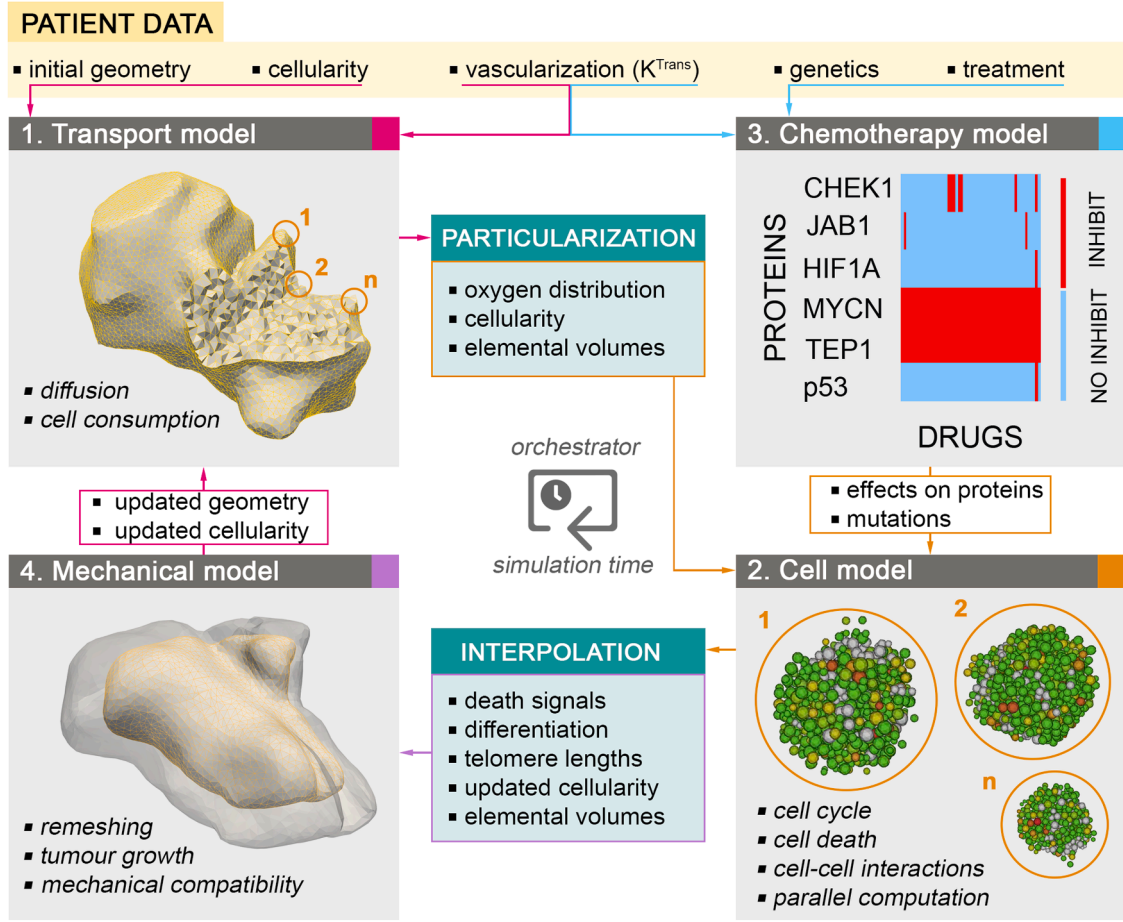


Fig. 2. Schematic representation of the multiscale model. Patient data (cellularity, vascularisation, gene profile, treatment regime) is automatically retrieved from the platform to generate a FE mesh to start the simulation. 1) A transport model at the macro scale predicts via a FE scheme the diffusion-reaction of species (i.e., oxygen) within the patient tumour depending on the vascularisation conditions and the current geometry. Results are particularised from some sampled elements to be individually simulated with 2) an agent-based model at the cell scale where each cell is modelled as an autonomous agent whose behaviour depends on nutrient availability, mutations, and the presence and effects of chemotherapy. These are determined by model 3) at the subcellular scale, which infers them via a neural network model, based on patient genetic data and vascularity. An interpolation scheme collates the data at the cell-level and returns it to the macro scale, where model 4) computes the mechanical compatibility of the biomechanical deformations induced by the cell agent-based model. Geometry and all the necessary variables are subsequently updated to start a new simulation loop.

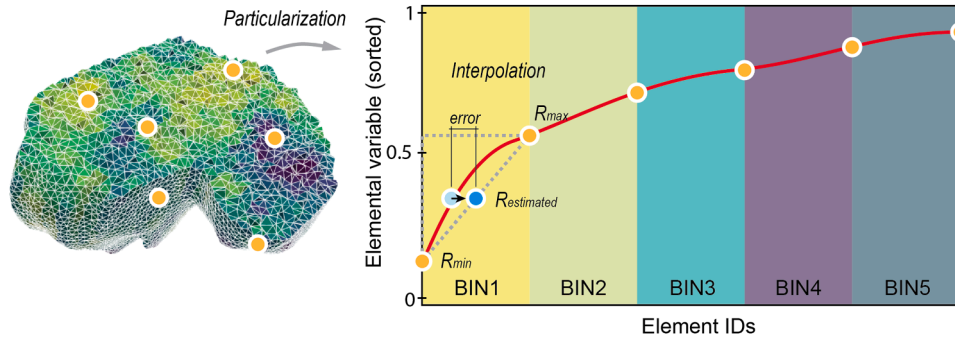


Fig. 3. Particularisation/Interpolation scheme. Elements of the mesh are grouped in different bins according to a chosen elemental variable in increasing order. The first and last elements of each bin (the ones with the smallest/highest value, represented with orange circles) are selected as representatives to perform the ABM simulations. When the individual simulations are completed, the results are interpolated linearly to the rest of the elements of each bin.

cellularity (φ^{cell}) following the equation:

$$\frac{\partial C_{ox}^{tissue}}{\partial t} = D \nabla^2 C_{ox}^{tissue} + K^{Trans} (C_{ox}^{blood} - C_{ox}^{tissue}) - \frac{A_{ox} C_{ox}^{tissue}}{k_{ox} + C_{ox}^{tissue}} \varphi^{cell} \quad (1)$$

where C_{ox}^{tissue} (unknown variable) and C_{ox}^{blood} (4124 pmol mm⁻³ [18]) are

the concentrations of oxygen at the tissue and blood respectively, D is the diffusion coefficient (0.00175 mm²/s) and A_{ox} (2200 mol m⁻³ day⁻¹) and k_{ox} (0.00464 mol m⁻³) are the oxygen uptake parameters [41]. The three terms on the right-hand side of the equation represent: i) interstitial oxygen diffusion through the tissue, ii) extravasation from the blood vessels to the tissue and iii) cell consumption.

2.3.2. Agent-based cellular model

The ABM requires a large number of parameters due to the complexity of intracellular dynamics, and a complete description of its intricacies is out of the scope of the present manuscript, where we focus on the broader vision. A full detailed description will be presented in a future publication. Nevertheless, this section provides sufficient details to allow the reader to understand the underlying phenomena and the principal mechanisms governing cell agent behaviour.

2.3.2.1. Structure of the agent-based model. The multicellular model has three components: a continuous automaton, discrete agents (representing cells), and a centre-based mechanical model. The continuous automaton describes the tumour microenvironment regarding the spatial distributions of cells and extracellular matrix. Neuroblastoma (cancerous) and Schwann (non-cancerous, matrix-producing) cells, the two main cell populations in the tumour [2], are modelled as discrete agents. While the continuous automaton mediates the effects of juxtacrine and paracrine signalling at the cellular level, a centre-based mechanical model is used to resolve cell-cell overlap. Note that in this document, within the context of the multicellular model, the words 'cell' and 'agent' are used interchangeably.

The continuous automaton uses a grid of voxels to represent the spatial domain. Each voxel is associated with a unique combination of i , j , and k , which are the voxel's spatial coordinates. Its state is defined by a vector:

$$v_{i,j,k} = (N_{i,j,k}^l, N_{i,j,k}^a, N_{i,j,k}^m, S_{i,j,k}^l, S_{i,j,k}^{lm}, S_{i,j,k}^a, S_{i,j,k}^b, M_{i,j,k}).$$

The first seven components denote the numbers of living neuroblastoma cells, apoptotic neuroblastoma cells, necrotic neuroblastoma cells, living Schwann cells, matrix-producing Schwann cells (a subset of living Schwann cells), apoptotic Schwann cells, and necrotic Schwann cells in the voxel respectively; they are non-negative integers. The last component, $M_{i,j,k}$, is the fraction of the voxel's volume occupied by extracellular matrix; $M_{i,j,k} \in [0,1]$. The voxel's 3D von Neumann neighborhood includes the voxel itself and its six orthogonally adjacent voxels.

Neuroblastoma cells are represented as discrete agents. Each agent of this type has a unique label n . Its state is described by four vectors.

- 1 The physical vector, \vec{c}_n^1 , describes the cell's spatial coordinates, the components of the net force on the cell in the same spatial dimensions, the cell's total overlap with other cells, the number of cells (including itself) within a search distance (R_{nbr}), and whether it can generate a locomotive force; $\vec{c}_n^1 = (x_n, y_n, z_n, F_n^x, F_n^y, F_n^z, o_{lp_n}, n_{br_n}, m_{ob_n})$
- 2 The cellular vector, \vec{c}_n^2 , describes the cell's position in the cell cycle, number of apoptotic signals, number of necrotic signals, critical number of necrotic signals, number of telomere units, degree of differentiation, hypoxic status, nutrient availability, if its DNA is damaged, and if its DNA is unreplicated; $\vec{c}_n^2 = (cyc_n, apop_n, nec_n, nec_n^c, tlm_n, deg_n, hyp_n, ntr_n, dd_n, du_n)$
- 3 The mutation vector, \vec{c}_n^3 , indicates whether its MYCN gene is amplified, whether its TERT gene is rearranged, whether its ATRX gene is inactivated, and the status of its ALK gene; $\vec{c}_n^3 = (MA_n, TR_n, AI_n, ALK_n)$.
- 4 The molecular vector, \vec{c}_n^4 , indicates whether the cell has sufficient ATP, whether its telomerase is active, whether its alternative lengthening of telomeres (ALT) mechanism is active, and whether its other gene products are active:

$$\vec{c}_n^4 = \begin{pmatrix} ATP_n, tel_n, ALT_n, MYCN_n, MAPK_n, JAB1_n, \\ CHK1_n, CDS1_n, CDC25C_n, ID2_n, IAP2_n, HIF_n, BNIP3_n, \\ VEGF_n, p53_n, p73_n, p21_n, p27_n, Bcl_n, BAKX_n, CAS_n \end{pmatrix}$$

Schwann cells are also represented as discrete agents. Each Schwann cell agent has a unique label s . Its state is described by two vectors and a Boolean variable. Both its physical vector and cellular vectors are analogous to those of neuroblastoma cell agents; $\vec{c}_s^1 = (x_s, y_s, z_s, F_s^x, F_s^y, F_s^z, o_{lp_s}, n_{br_s}, m_{ob_s})$ and $\vec{c}_s^2 = (cyc_s, apop_s, nec_s, nec_s^c, tlm_s, hyp_s, ntr_s, dd_s, du_s)$ respectively, with the only difference being that a Schwann cell agent does not have a variable describing its degree of differentiation. Finally, the Boolean variable ATP_s indicates if the Schwann cell has sufficient ATP.

The mechanical model is an off-lattice model allowing for continuous values of x_n, y_n, z_n, x_s, y_s and z_s . It comprises a linear force law and an equation of motion. The force law relates the overlap between two cells to the repulsive force acting between them. The overlap ($\delta_{1,2}$) between each pair of cells depends on their radii, R_1 and R_2 , which increase throughout the cell cycle, and their displacement vectors, \vec{r}_1 and \vec{r}_2 ; $\delta_{1,2} = R_1 + R_2 - ||\vec{r}_1 - \vec{r}_2||$. The magnitude of the repulsive force between this pair of cells is given by another equation: $F_{1,2} = k_1 \delta_{1,2}$ where k_1 parameterizes the force law ($2.2e^{-3}$ N/m). The force vectors connecting a cell to its peers combine to give a net force vector: (F_n^x, F_n^y, F_n^z) or (F_s^x, F_s^y, F_s^z) . If the cell is mobile ($m_{ob_n} = 1$ or $m_{ob_s} = 1$) and contact-inhibited ($n_{br_n} > n_{br_c}$ or $n_{br_s} > n_{br_c}$ with n_{br_c} being a threshold), its net force vector is multiplied by a factor of k_2 . The equations of motion describe how the cell moves in response to the net force:

$$\vec{F} = \mu(1 + M_{i,j,k}) \vec{v} \quad (2)$$

In these equations, μ denotes the tumour microenvironment's viscosity and \vec{v} the agent velocity.

Finally, the oxygen level in the microenvironment is assumed to be homogeneous in space at the scale of the ABM (the elemental value coming from the transport of species simulation) and it is represented as a dimensionless variable: $O \in [0,1]$ and the scale is O_s (72 mmHg). The O is incremented by a variable amount, P_O , in one time step. In reality, variations in P_O are caused by angiogenesis, which depends on the non-negative integral number of angiogenic signals (ang) in the system. In our simplified model, the non-negative integer number of VEGF-producing neuroblastoma cell agents in the system is N^v which depends on the number of living neuroblast agents wherein VEGF is active.

2.3.2.2. Initialisation and implementation of the agent-based model. We devised a Monte Carlo algorithm to simulate how the agent-based model's variables change over time. After the orchestrator configures the simulation and an initialisation routine, the actual simulation iterates a pre-defined sequence of operations a finite number of times.

During configuration, the simulation requires the following inputs:

- 1 The time scale is defined by its extent and grain. The simulation entails a positive integer number of time steps (steps) lasting size hours each. For example, for the designated internal step of one hour, the first time step starts when time = 0 and ends when time = 1; time tracks the simulation's progress in hours. Note that this differs from the orchestrator (macro-scale) time step (two weeks). The spatial domain is a cube in which volume V is determined by the mesh element being simulated.
- 2 The tumour's composition is defined by its histology type (hist) and grade of differentiation (grade), which are extracted from the patient's clinical profile from the PRIMAGE platform. There are seven histological categories: neuroblastoma (0), ganglioneuroblastoma (1), nodular ganglioneuroblastoma (2), intermixed ganglioneuroblastoma (3), ganglioneuroma (4), maturing ganglioneuroma (5), and mature ganglioneuroma (6) [51]. Regarding grade, the tumour can be undifferentiated (0), poorly differentiated (1), or differentiating (2) [51]. Together, hist and grade determine the cellularity, the

split between neuroblastoma and Schwann cells, and the degree of differentiation of the former. Cellularity is obtained from the patient's images.

- After creating the discrete neuroblastoma cell agents, their attributes must be initialised. In addition to the mutation vector \vec{c}_n^3 . It can be specified whether ALTN and p53n in the molecular vector \vec{c}_n^4 are forcibly switched on and off respectively [11].
- The tumour's microenvironment has two components. First, O needs to be initialised (inputs from the orchestrator as previously described) and $M_{i,j,k}$ (cellularity's complement). Second, the chemotherapy regimen is defined in terms of three vectors. The first two, chemostart and chemoend, have the same number of elements in agreement with the number of chemotherapeutic periods in the simulation. Each vector element indicates when (time) a chemotherapeutic period starts or ends. The remaining vector, chemoeff, describes the regimen's effects on the tumour's CHK1, JAB1, HIF, MYCN, telomerase, and p53 in different chemotherapeutic periods. Its elements are arranged in groups of six and the number of groups is the number of chemotherapeutic periods. Each element is the probability that the corresponding gene product is inhibited by chemotherapy in the corresponding period.

After configuration, each agent's randomly assigned spatial coordinates of each agent are adjusted by minimising the total overlap in the spatial domain with the centre-based mechanical model. Finally, the continuous automaton is initiated based on the agent spatial coordinates, and the oxygen supply rate (P_O) calculated according to the initial number of agents.

After initialisation, the simulation iterates the following operations steps times.

- The neuroblastoma agents are evaluated sequentially. Each agent senses the tumour microenvironment to determine whether it is hypoxic, nourished, and affected by chemotherapy, as well as the stimuli from the other agents. Subsequently updates to the status of an agent's DNA may be executed: e.g., hypoxia and having shortened telomeres damage it [42], while p53 and p73 repair damaged or unreplicated DNA [69]. Phenotypic behaviours such as differentiation, apoptosis, necrosis, as well as gene expression are updated, followed by cell cycle progression rules, with contact inhibition from the presence of neighbouring agents potentially leading to cell cycle arrest. If the agent is apoptotic or necrotic, it can be removed (e.g. by the immune system [12,55]), otherwise, it may divide to produce a daughter cell.
- The Schwann cell agents are evaluated one by one, following the same steps as in the first operation. However, as their gene products are not in the model, gene expression is not considered in this operation and the cellular phenomena are not regulated by gene products. Furthermore, these agents do not differentiate.
- After evaluating both agent groups, the total cell-cell overlap in the spatial domain is minimised using the centre-based mechanical model, thereby updating the agents' spatial coordinates. The continuous automaton is updated to reflect the spatial distributions of various cell populations, including the addition of new voxels if agents move beyond the existing outer boundary.
- The numbers of VEGF-producing neuroblastoma agents and living Schwann cell agents are used to update the progress of angiogenesis and hence the vasculature (P_O).

The oxygen level, O , is updated according to the balance between P_O and oxygen consumption. In addition, this step updates $M_{i,j,k}$ by considering the number of collagen-producing Schwann cell agents in each voxel (assuming that contact inhibition arrests collagen production).

This model, consisting of hundreds of thousands of agents per FE

simulated, was implemented using FLAMEGPU2 [13]. The FLAMEGPU2 framework takes advantage of the highly parallel architecture of GPUs to accelerate large-scale ABMs which are otherwise infeasible with traditional CPU computation.

2.3.3. Sub-cellular model of chemotherapy effects

The objective of the subcellular model is to provide predictions on the effects of chemotherapy treatment on certain proteins involved in the cell cycle, so that the agent-based cellular model can predict the emergent behaviour of the downstream cell phenotypic changes arising from these effects.

To perform this task, the subcellular model receives the treatment assigned to the patient and the vascularity of the tumour (K^{Trans}) (Fig. 2) as input data. Each chemotherapy treatment has a certain course [53] (number of cycles, number of days the cycle lasts, number of days the drugs are administered) and consists of a combination of multiple drugs. In turn, each drug has a different half-life and it exerts its action through a different set of proteins. The ability of drugs to interact with proteins is extracted from pharmacology databases [38] and complemented with computational predictions using an ensemble of six independent approaches based on chemical similarity, core substructures, quantitative structure-activity relationships, cross-pharmacology statistics and machine learning models [61] (Fig. 4).

The subcellular model collects this information and calculates intervals of presence of the drugs in the body and their effects at different points in the tumour based on K^{Trans} values. To determine the probability of protein inhibition by the drugs (ranging from 0: no inhibition to 1: total inhibition) based on K^{Trans} values, the results of the study on doxorubicin [46] have been used. In this study, it was observed that the arrival of doxorubicin to the cells began with values of $143 \cdot 10^{-6} s^{-1}$, while the necessary concentration of doxorubicin to achieve a complete response to the treatment was related to a value of K^{Trans} of $283 \cdot 10^{-6} s^{-1}$. Because of this, these values represent the lower and upper limits through which the probability values of the effects are established. Since similar studies are not known for other drugs, setting these limits on the basis of doxorubicin alone is an important assumption.

2.3.4. Macro-scale mechanical model

A continuum description for the mechanical response of the tumour is used. At the macroscopic level, the continuum-based model accounts for two compartments: the cells and the extracellular matrix (ECM). To describe the tumour kinematics, the multiplicative decomposition of the gradient deformation tensor is used. The total deformation gradient (F) is assumed to be composed by the elastic deformation (F_e) and the non-elastic one (F_{bio}).




Fig. 4. Subcellular model schema. The model receives the treatment provided to the patient and the vascularisation at different points of the tumour. If, for example, the chemotherapy treatment is CADO, this consists of 4 cycles of 21 days of duration, where the drugs are administered during the first 5 days. Furthermore, the drugs that make up this treatment are cyclophosphamide, doxorubicin and vincristine. These drugs have an average half-life in the body of 48 h and combined they inhibit the following proteins: CHEK1, JAB1, MYCN and TEP1. The model generates two output files: one with the intervals of the presence of drugs in the body and another with the probability of protein inhibition based on the K^{Trans} values, taking $143 \cdot 10^{-6} s^{-1}$ and $283 \cdot 10^{-6} s^{-1}$ as the lower and upper limits for these probabilities.

$$\mathbf{F} = \mathbf{F}_e \cdot \mathbf{F}_{bio} \quad (3)$$

\mathbf{F}_e is associated with the external loads and the stress response of the material, whereas \mathbf{F}_{bio} accounts for the change in size and shape due to volumetric growth and the new tissue produced during the incremental growth process. Non-homogeneous growth, especially when the individual elemental volume change is determined externally via the ABM, generally happens in a mechanically incompatible manner. Hence, an elastic deformation is necessary to achieve compatibility.

Assuming an isotropic growth deformation, \mathbf{F}_{bio} can be characterized by a single isotropic growth multiplier λ_{bio} :

$$\mathbf{F}_{bio} = \lambda_{bio} \mathbf{I} \quad (4)$$

so that $\lambda_{bio} = \sqrt[3]{J_{bio}}$, where J_{bio} is the net change of the volume (volume ratio) calculated by the ABM in each element. It is worth noting that this volume change is directly related to both the number and the spatial spreading of cells in the element, which was determined during the ABM by several factors including the oxygen level and the individual cell cycle, which in turn was affected by the protein alteration due to the chemotherapy treatment, whose efficacy relies on tumour vascularity. In this way, dynamics at the molecular level are linked to the macroscale as they alter both the mechanics (depending on cellularity) and the emergent behaviour of the tumour, therefore influencing its shape and evolution over time.

The internal behaviour of both cell and ECM compartments is evaluated separately. It is hypothesised that the total deformation gradient due to the ECM behaviour (\mathbf{F}_{ecm}) is equal to the deformation gradient due to cell growth (\mathbf{F}_{cell}). Therefore, in order to obtain the total stress, both cells' and ECM's stress must be accumulated:

$$\mathbf{F}_{ecm} = \mathbf{F}_{cell} \quad (5)$$

All tissues are modelled as isotropic linear elastic materials. The material properties of the different compartments are computed using a mixture rule. Therefore, the ratio of these basic types is assumed to determine the mechanical properties of the local tissue. The cell's and ECM's initial elastic moduli are set to $0.08e^{-3}$ and $0.8e^{-3}$ MPa respectively [23]. The Poisson ratio is defined as 0.38 for both compartments [23].

The continuum model is solved via FE using the commercial software Ansys® Structural Mechanical (partner of the PRIMAGE consortium), Release 19.2. The geometry is initially discretised in a linear tetrahedral mesh with an average initial element volume of 3 mm^3 . After each computation, the geometry of the tumour is updated and remeshed using the Python Gmsh lib.

2.4. Output data

When the simulations are finished, the HPC orchestrator notifies the Kubernetes platform (Fig. 1) which automatically returns a summary of the results to the platform via .xml files. These outputs are attached to the patient's analysis history, which can be consulted and visualized online within the PRIMAGE platform as well as downloaded to a local computer.

2.5. Model calibration

The calibration process in a multicomponent/multiscale model is a complex task, as there are many parameters and subtle changes at the microscopic level may have major impact at the macroscale and vice versa. Also, the models must be flexible and robust enough to accommodate a wide range of inputs, while still avoiding overfitting and poor generalisation. Ideally, data from a large number of patients before and after the treatment should be used, however, due to time and data constraints in the development of the project, we have focused the calibration on a single patient as a proof of concept of the framework.

Firstly, we pre-emptively adjusted the main parameters of each isolated model to achieve general behaviours, using also in-vitro data from literature (note that both the transport and mechanical models are deterministic and all their parameters come from literature). Secondly, in a later stage, we fine-tuned the most interscale-influential ones to capture more specific responses. In the subsequent sub-sections, we provide a detailed account of the calibration process, with particular emphasis on the ABM model, which demands a more intricate procedure due to the significant number of parameters involved.

2.5.1. Calibration of the ABM model

a) Systematic calibration with a tournament-style pipeline

In the absence of data relating model parameters to measurable emergent behaviors within PRIMAGE, we carried out an initial process of model calibration based on neuroblastoma-related data available in the literature. Initially, we calibrated parameters collectively in a tournament-style pipeline without chemotherapy. We used a Latin hypercube sampling technique to generate 3000 near-random combinations of parameters and passed them through six elimination rounds. In each round, we carried out simulations using the sets of randomly selected parameter values and attempted to reproduce a dataset found in the literature. The rounds were arranged in an increasing order of sophistication, with the first dataset—in vitro and without genomic data—being the coarsest and the sixth—clinical, patient-specific, and with known mutations—being the most precise. Specifically, these steps comprised: 1. fitting the model output to a set of in vitro data regarding neuroblastoma's growth kinetics. 2. Further assessment of the best parametric combinations from the first round by fitting the model output to in vitro data regarding neuroblastoma's hypoxic response. 3. evaluation of the remaining parametric combinations by reproducing in vitro observations of the dynamics between neuroblastoma and Schwann cells. 4. Further refinement was carried out by considering the relationship between the histological type of a tumour and the clinical outcome. 5 and 6. Finally, the last two rounds involved reproducing the clinical outcomes of patient groups with different combinations of mutations.

b) Further calibration of selected parameters accounting for PRIMAGE-relevant patient data

Following the process described above, with additional data available via PRIMAGE, we further refined selected parameters to render the agent-based model suitable for the decision support system envisioned by the PRIMAGE project. Specifically, we evaluated chemotherapy's inhibitory effects on specific gene products (and tumour shrinkage), in vivo growth kinetics, the dynamics between neuroblastoma and Schwann cells, tumour differentiation, and the pro-tumour effects of MYCN amplification. The extent of shrinkage in a real patient was used as the benchmark in parts of this step.

2.5.2. Calibration of chemotherapy effects

Regarding the chemotherapy model, its calibration entailed accurately learning and predicting the impact of drugs and their metabolites on a specific set of proteins (CHEK1, JAB1, HIF1A, MYCN, P53, TEP1) utilised to determine the cell cycle in the cellular model. The model uses the drug knowledge stored in pharmacology databases [38] as a foundation for learning. Consequently, when the model is presented with a new drug, it can forecast its impact on proteins based on similarity to the ones it was trained on. This generates a table of Boolean values that indicates whether a certain drug or metabolite (used in the Neuroblastoma chemotherapy treatments) produces the inhibition of the different proteins. The subcellular model iteratively consults this table to establish the final probability of inhibition as a function of the K^{Trans} values.

2.6. Model testing

To evaluate the functionality of our framework implementation (producing the expected outputs for given inputs and handling various scenarios appropriately) and ensure its reliability, we have conducted extensive testing on our codebase including both manual and automated techniques. In particular, we have employed a combination of unit tests, integration tests, and system tests to thoroughly assess the robustness of our platform.

This was particularly challenging, considering that the different component models were developed by different teams, which had to ensure that each of the models was capable of working in isolation (by providing mock inputs) before connecting them to the pipeline.

For example, the ABM, the most complex component model of the orchestra built upon FLAMEGPU2 (with over two thousand unit tests), was subjected to rigorous evaluation in isolation. It is worth noting that agent based models, such as the ABM here presented, are not typically unit tested due to the desired behaviour being emergent. Furthermore, the GPU implementation within FLAMEGPU2 makes it difficult to call methods independently. Instead a process of validation was carried out to ensure the model behaved within physically realistic constraints, followed by a calibration process described in the previous section. The rest of the component models, much more simple due to their deterministic nature, were also tested by the corresponding teams before plugin a working version into the platform.

Therefore, the framework was built block by block, initially testing it locally on simplified problems. Once the architecture was considered final, it was ported to PRIMAGE HPC cluster. Both the component models and the relation modules, however, were subject to constant change due to calibration or code debugging. Unit tests of each of them were developed to be able to run them locally. These mock models can be found at https://gitlab.com/unibo_ism_open/orchestratorvph-hfv3.1/-/tree/master/Models (files with mock* prefix). The majority of the testing can be done offline, however, the functions using SLURM commands (e.g. call model, check model execution, re-call model) and the Rimrock API need a cluster to operate. Therefore, these functionalities were firstly tested isolated on the cluster and later added to the main code together with a version of each function that emulates the cluster files locally. In fact, there is a parameter in the master module (https://gitlab.com/unibo_ism_open/orchestratorvph-hfv3.1/-/blob/master/src/master.py) that allows setting the execution to either local or on Ares cluster. If set to local, a MySQL database is created in the root user and the orchestrator will call only mock models during the execution, emulating the cluster functionalities (note that the initial inputs are pre-defined for all the mock models). Once all parts were working locally, global tests were performed on the cluster now considering real inputs (patient's data). This stage relied on the support of the modelers to overcome the incompatibilities in data exchange and other potential artifacts.

3. Sample of typical system run

3.1. Description of the example of application

To demonstrate the operation of the multiscale model and the potential of the whole framework as a concept, we have picked a patient (used for calibrations) from the PRIMAGE platform as an illustrative example. This particular patient was a male diagnosed with neuroblastoma at the age of 30 months and subsequently treated with Rapid COJEC for 80 days (8 cycles of 10 days each). He was evaluated after the chemotherapy process observing good partial response (i.e. the tumour was not completely eliminated, but volume decreased significantly).

3.2. Technical parameters

The original images from the patient were retrieved from the

PRIMAGE platform and automatically processed with the *im2mesh* library, transforming 13 slices of the segmented tumour into a volumetric tetrahedral mesh of 34,483 elements (Fig. 5A). Vascularisation (K^{Trans}) and cellularity data were interpolated to the centroids of such mesh. A particularisation level of 1% was chosen to accelerate the computations (see section Orchestrator performance and computational cost).

3.3. Assumptions

To represent the behaviour of our multiscale system, we made several assumptions about the underlying mechanisms and interactions. These premises aided in managing computational complexity and the integration of diverse scales within the models. Below we provide a list of the most important assumptions, our reasoning to make them and how they affect the model behaviour:

- Uniform boundary mechanical conditions surrounding the tumour. As surrounding organ segmentation and mechanical characterization is unavailable, we consider an homogeneous elastic medium surrounding our tumour geometry for simplicity. The incorporation of different boundary conditions into the model, given the data, is immediate and would impact the shape evolution of the tumour.
- Residual mechanical stresses are not stored between orchestra steps and therefore are not taken into account. The main reasons to neglect them are: i) no surrounding tissues to interact with, ii) the ABM does not use directly the strain-stress field of the finite elements, iii) computational efficiency, iv) no data for validation.
- Only oxygen diffusion is considered. Although the transport model can indeed simulate the diffusion of any species, oxygen was chosen as a representative of nutrients in general, the main reason being the lack of clinical data. Were this kind of data available, its addition to the pipeline would be trivial and the ABM could be easily adapted to account for it.
- Constant level of oxygen at each element for the ABM simulations. We assume homogeneity within a single element (after simulating oxygen diffusion at the macroscale) because each voxel spans 30 μm in each dimension only and the diffusion coefficient of oxygen in water at room temperature is around $2 \times 10^{-5} \text{ cm}^2 \text{ s}^{-1}$ [19]. If that was not the case, and oxygen reaction-diffusion model should be solved at the microscale, within the ABM, increasing considerably the complexity of the simulation.
- Vascularization (K^{Trans}) is assumed constant throughout the simulations. The particular elemental values of (K^{Trans}) change overtime due to mesh interpolation, but the dynamics of the vascularization process itself are not modelled as we don't have data to validate with. Including it would have important implications as different treatments may affect the vascularization in diverse ways, and (K^{Trans}) plays a central role in our model as it affects both the drug and oxygen delivery which are key for the ABM simulations.
- The probability of protein inhibition by the drugs is based on K^{Trans} values, based on a study on doxorubicin [46]. Given the importance of this parameter, as stated in our previous point, new data for other drugs might have a great impact on how our model responds to treatment.

Besides, to ensure the stability of the calculations, we also assumed:

- A minimum cellularity threshold of 0.05 for patient's input data to avoid possible artifacts from the imaging process. This provides the ABM with a minimum number of agents, while having a negligible effect on the mechanical model (as the element is still basically ECM). Although there is no basis for this assumption and a tumour may contain regions without cells, it is justified because cells are likely to move from crowded regions to empty ones, so any regions devoid of cells are unlikely to be anything more than transient structures.




Fig. 5. Patient's tumour geometry (initial FE mesh). A) A volumetric stack of patient's DICOM images is automatically retrieved from the platform, obtaining the tumour mask and reconstructing the tumour volume in a 3D FE mesh needed by the models. B) Volume shaded in light blue shows the tumour shape at diagnosis. Volume shaded in orange shows the actual tumour after the chemotherapy treatment. Note that images before/after treatment are not registered, meaning that volumes are manually and arbitrarily placed for visual purposes. Gray shadow (overlapping with orange) shows the tumour geometry predicted

by the orchestrated models after the chemotherapy simulation.

- At the beginning of every loop beyond the first one, a minimum amount of SC cellularity (0.05) is assumed, mimicking the migration of SC cells from the outside of the tumour [65]. This prevents unrealistic tumour growth in elements where only neuroblastoma cells are left alive after an orchestrator loop. The threshold was chosen ad hoc. It is the parameter that deserves special attention in a future sensitivity study. The ability to control this parameter dynamically, perhaps through the use of a targeted therapy, could hypothetically give a clinician an extra therapeutic option.
- When the number of cells inside an element is extremely low (due to chemotherapy damage or other factors) the elemental volume change reported is unreliable. In order to avoid instabilities during orchestrator interpolation (Fig. 3), we assume that those elements report results (i.e. volume change, cell count) similar to the median values within the whole tumour. In other words, we transform outliers into median values to promote stability at the cost of accuracy. Although this assumption is a workaround, necessary due to the current technical limitations of our computational method, it is worth noting that this scenario occurs very rarely in our simulations so this computational fix is implemented infrequently, which allows us simulating cases of extreme tumour regression. To alleviate this, smaller time steps could be used, but notably increasing the computational cost.

3.4. Proof of concept predictions

The main goal of each simulation is to determine whether the patient's tumour grows or recedes when undergoing a specific treatment. Therefore, the overall volume reduction was the main target parameter during the calibration process as it is the outcome, measured by our model, with the most clinical relevance. For this particular patient, we were able to capture the tumour reduction ($1 - (Vol_{aftertreatment}) / (Vol_{ini})$) (98.1% real vs 95.6% predicted) as shown in Table 1 and Fig. 5B. It is worth noting that, despite the similar relative reduction, the simulated tumour is double the size of the real tumour after treatment (4,375 vs 2,054 mm³). Furthermore, although the overall relative reduction is similar, the final shapes notably differ. This cannot be controlled/adjusted from the computational side, as the global shape depends on the individual behaviour of each mesh element which in turn depends on all the models' interconnected parameters.

3.5. Orchestrator performance and computational cost

As described in the Resources section, a variety of software systems and infrastructure are needed to support the presented framework. The computational cost of each of the simulations is highly dependent on key

Table 1
Volume comparison (patient vs simulation).

	Volume	Reduction
Diagnosis	1.07e5 mm ³	–
After treatment	2054 mm ³	98.1%
Simulated	4375 mm ³	95.9%

hyper-parameters such as the level of particularisation, i.e., the percentage of representative elements of the patient's geometry mesh that will be simulated by the cellular and sub-cellular models (see Particularisation and interpolation section). To explore the impact of this parameter on performance, the typical patient case described above, was executed to reproduce a two-week progression (one time loop of the orchestrator) with a range of degrees of particularisation. The data was collected using 10 V100 GPUs on the Ares cluster.

As expected, and evident from Fig. 6, the overall wall time increases with the degree of particularisation. Two of the CPU-based component models (Chemotherapy, Mechanical) have very short times of around 1 min, whereas the remaining CPU-based component model (Transport) took around 15 min to execute. Note also that both the Transport and Mechanical models use the whole geometry of the patient's tumour, so their computational cost is unaffected by the level of particularisation. The remainder of the runtime is spent executing the agent-based model, and by the orchestrator processing inputs and outputs to each job. Evidently, when the wall-clock time using a single GPU is considered, both the time spent by the agent-based model and the orchestrator are closely linked to the number of finite elements being processed by the agent-based model, given that along with the cellularity extracted from the patient data, this is an important determinant of total number of cells/agents at the start of each simulation. As the workload for this model consists of thousands of very small jobs, it would be expected to scale linearly with the number of GPUs on which it is executed concurrently (10 in this test). Furthermore, the orchestrator is able to hide a small amount of latency, with regard to the cost of creating the agent-based model jobs. Although queue times from the HPC cluster are not included in these timings, it is worth noting that the overall queue time for all contained jobs was less than 10 min. In summary, the runtime of the orchestrator in Fig. 6 is dominated by the time spent processing inputs and outputs for the cell-model, hence why it scales with the level of particularisation. Profiling showed the bottleneck to be communication with the database: this is something which could be improved by reworking how the orchestrator stores data. In contrast, the runtime of the cell-model scales with the number of FEs it is simulating.




Fig. 6. Wall Time of Component Models and Orchestrator according to the degree of particularisation. Time is measured in hours, and the agent-based (cell) model's times represent the longest of the 10 jobs executed concurrently. Note that the contributions of the Chemotherapy and Mechanical models are negligible.

The cell model is not well suited to executing short runs (336 timesteps per loop), as the cost of initialising the model takes up as much as 10% of the runtime. However the loop duration was decided as a compromise, as less frequent communication between orchestrated models weakens the multiscale approach.

In summary, a single loop of the orchestrator, representing two weeks of treatment, can be simulated in less than two hours, and consequently a whole chemotherapy treatment course (typically 80 days) could be simulated in less than twelve hours with 10 V100 GPUs. The simulation time would likely be shorter as typically the simulated tumour would shrink reducing the computation cost, Computational time could be further accelerated by increasing the number of GPUs.

6. Discussion and conclusion

Patient-specific simulations, which use computational models to simulate patient physiology, are promising for advancing personalised medicine. However, clinicians' adoption of these simulations as part of their day-to-day routine activities is still very limited due to several technical challenges. One of the major challenges is the complexity of the models themselves, as they require detailed anatomical models that accurately represent the structure of the patient's organs and tissues. Furthermore, these models often require large amounts of input data from various sources, including medical images, laboratory results, and physiological measurements. Integrating these data sources can be complex, requiring sophisticated data management and processing techniques. Accurately parameterising these models remains a major challenge due to the technical and ethical barriers to collecting appropriate data at the range of scales required. In particular, agent-based models require data relating to cell phenotype that is simply not accessible through *in vivo* imaging (or other spatially aggregated) modalities, necessitating the use of data from *in vitro* experiments, which may have only a limited relevance to *in vivo* scenarios.

Developing and validating such models is a time-consuming and technically challenging process, requiring expertise in computational modeling and clinical medicine. Another significant challenge is the need for high-performance computing resources to run the simulations, as they will likely be computationally intensive and will require substantial computational power to run in a reasonable amount of time. Clinicians often do not have access to these resources, or even if they do, they may not have the technical expertise to operate and maintain them. Finally, the outputs of patient-specific simulations can be difficult to interpret and integrate into clinical decision-making. The results are often presented in complex numerical data or 3D visualisations that may be challenging for clinicians to understand and use effectively.

The PRIMAGE [36],[47], aims to address many of these challenges, in an ambitious effort to provide a state-of-the-art decision support system for neuroblastoma paediatric cancer that offers a comprehensive platform for storing patient data, including medical images and clinical data [3,15,50,59,60]. This platform has been developed with the input and guidance of leading clinicians and healthcare professionals, who will be the primary users of the system. The user interface, developed by the partner company Quibim, has been designed to be intuitive and user-friendly, ensuring that clinicians can quickly and easily access the data they need to make informed decisions about the best course of treatment for their patients. The system is designed to connect to a powerful computational infrastructure that allows for the evolution of neuroblastoma tumour growth through advanced computational simulations. Examples of the user interface can be found in Supplementary Material.

There exists diverse literature involving neuroblastoma modelling focused both at the microscale, describing in detail intracellular mechanisms, and the macroscale or whole-tumour level [5,9,20,22,35], which are based on differential equations but either lacking generality of mechanistic details. For example, the model described in [20] uses ordinary differential equations to describe the pharmacokinetics of

bevacizumab and its effects on neuroblastoma growth and vascularisation: a drug-specific problem. Conversely, [5] develop a more general model that reduces neuroblastoma progression to two phenomena, growth and dissemination, only. More details are needed to reveal non-linear phenomena. A platform to conduct *in silico* trials of targeted therapies and immunotherapies requires non-linear phenomena to be integrated across multiple different scales in space and time. On the other hand, the mechanistic details to be integrated must be broad enough to include drug targets at the genome scale.

In this work, we present what we believe is the first multiscale orchestrated framework that integrates the multiple scales involved in neuroblastoma and the automatic execution of computational simulations on high-performance equipment. We underline the theoretical and technical challenges associated with connecting models at different scales, highlighting the vast amount of resources, software, and technical knowledge that is required to build a platform capable of handling such complex simulations. Indeed, developing such a platform requires not only the latest software and hardware technology but also the input of highly skilled professionals from a range of disciplines, which have been provided by the different partners of the PRIMAGE consortium.

However, apart from the assumptions previously discussed, there are some limitations related to our approach. For example, at the subcellular scale, we have a machine learning model to predict the effects of drugs on the proteins in the ABM. Still, we do not have a dynamic model of the signalling pathways and gene regulatory networks that link these proteins together. In addition, the mutational profile is set at the beginning of each simulation, where the ABM clones compete for space, expanding, or shrinking. However, they cannot interconvert, and, as such, our model at its current state is incapable of modelling therapy resistance. It could certainly be adapted as drug resistance is a critical issue in positive patient outcomes, and we have already started carrying out parallel studies relating to clonal competition and resistance using a mathematical modelling approach [25].

Besides, further calibration would be needed for a broader application of the model. Both a higher number of patients and more combinations of hyper-parameters (at the orchestrator level) should be explored to test the actual robustness of the orchestrator. At the macroscale, we are currently assuming uniform boundary mechanical conditions surrounding the tumour mass, which leads, together with many other factors, to the final shape of the tumour. With the available data, we can calibrate model parameters to match a patient's tumour volume after treatment, but we cannot accurately predict the final tumour shape. To overcome this, we would need image segmentations of the surrounding organs, as well as their mechanical characterization. This, per se, is feasible and easy to implement into the current model. The reality, however, is that only the tumour is segmented in the clinical practice for the vast majority of the patients, although this could change in the near future with the increasing use of AI tools. In any case, it is still not fully clear how tumour volume and shape relate to prognosis. For instance, [4] found that none of the methods of primary tumour response assessment, including measurement in one dimension as per Response Evaluation Criteria in Solid Tumours (RECIST) or measurement in three dimensions as per INRC, was predictive of outcome. In fact, the authors concluded that primary tumour response in children with high-risk neuroblastoma should be evaluated in accordance with RECIST criteria, using the single longest dimension. More recent studies have been more successful in finding correlations. For example [32] found that PET-based intratumour heterogeneity was a strong independent prognostic factor in neuroblastoma. In the clinical high-risk group, patients with high metabolic heterogeneity showed significantly poorer outcomes compared to those with relatively homogeneous tumours. This is something that our model could explore if given the appropriate data. Others such as (J. X. [62]) aimed to elucidate the prognostic significance of tumour size, but in surgery performed on neuroblastoma patients, finding an optimal cutoff of 4 cm for overall survival. Interestingly, they claimed that this cutoff value can only

identify unfavourable neuroblastoma patients with diagnosis at distant-stage disease or differentiated grade tumour but not with regional and local or undifferentiated tumour. Further research in this direction would be very valuable to link our predictions with potential patient outcome.

Despite these limitations and uncertainties, from a research point of view, our platform can still be useful in assisting drug discovery, for example, by exploiting evolutionary principles to optimise chemotherapy schedules. Furthermore, the governing dynamics are not the same before chemotherapy, when there are many cancer cells, and afterwards, when the population is small, fragmented, and genetically less diverse [16]. Interactions between tumour cells, non-cancerous cells and the local microenvironment are highly complex and not fully understood. Our current cellular ABM incorporates some aspects of this behaviour, but there remains scope to increase this functionality to include additional relevant behaviours e.g. recruitment of Schwann cells into the tumour [65]. More broadly, the framework can meet the need to accelerate drug development by identifying high-priority targets and drugs, especially the combination of modern targeted therapies and immunotherapies with chemotherapy [40]. Examples are the combination of ALK inhibitors with chemotherapy and the combination of anti-GD2 monoclonal antibodies with chemotherapy [58]. It could also be useful for basic neuroblastoma research. For example, MYCN is known to have a non-linear relationship with the ARF/MDM2/p53 axis and hence the clinical outcome [68]. Different cellular stresses are encoded in different p53 dynamics, which are decoded in a context-dependent manner to trigger the correct cellular responses. This set of information processing mechanisms has many layers of regulation (P. [63]). The mechanisms have been modelled extensively at the sub-cellular level [29]. Our platform can link them to cell-cell interactions and tumour-level population dynamics.

In summary, we have successfully implemented a complex framework that enables the simulation of patient-specific tumour growth across multiple scales on a computational cluster. All the necessary code to achieve this framework has been extensively documented to be easily adapted and extended by new developers in future projects. The development of this framework opens up new opportunities to better understand and treat cancer, making it a significant achievement in the field of computational oncology, hoping to inspire further research and development in this area in the years to come. We believe that this work represents a major step forward in terms of its technical, theoretical, and practical implications, as it closes the gap between the predictive capabilities of computational models and their practical use by clinicians. Despite the complexity of the underlying algorithms and computational processes involved, the platform has been designed to be highly user-friendly, with clinicians able to launch simulations with just a few clicks, without needing to have any technical expertise in the underlying processes. The potential benefits of this decision support system for cancer are significant, providing healthcare professionals with a powerful tool to improve patient outcomes and ultimately save lives. By providing a comprehensive platform for storing and analysing patient data, and enabling clinicians to make more informed decisions based on simulations, this system has the potential to transform the way we approach cancer treatment and care.

7. Hardware and software specifications

The multiscale model requires significant CPU and GPU resources, plus the PRIMAGE platform needs to be able to run multiple multiscale simulations at once. These reasons lead us to choose HPC infrastructure as our target execution platform. In fact, Cyfronet AGH, partner of the

project, has provided PRIMAGE with access to three high-performance computing clusters, each holding a position within the TOP500 HPC systems internationally. These have been used for developing, testing and evaluating the complete orchestrated model. The preliminary tests of the multiscale simulation were performed on the Prometheus HPC.² It uses the Slurm workload manager system³ connected with the computation grant negotiated by the PRIMAGE [47] with the PLGrid infrastructure [30]. During tests and model initial calibrations we used 106, 100 CPU hours and 18,110 GPU hours. As of now, the orchestrator and all models have been moved to a newer and more powerful HPC cluster - Ares⁴ also hosted by Cyfronet, on which the models previously calibrated on the PRIMAGE patient cohort will be run. We are also planning additional speed improvement to our models by using a dedicated GPU HPC cluster (based on NVIDIA A100 GPU cards) - Cyfronet Athena⁵ in future developments of the project.

Code availability

The source code of the orchestrator, including the component model codes, is available under the terms of a MIT license and can be downloaded from <https://gitlab.com/primageproject/orchestrator>.

Additionally, the source code of the agent based model that can be run standalone is available under the terms of a MIT license and can be downloaded from <https://github.com/primagesheffield/flamegpu2-neuroblastoma>.

Credit authorship contribution statement

Borau C.: writing, review and editing of the manuscript, coordination of orchestra model communication and calibration, formal analysis, co-author of the Python library used to generate the input data (geometry, mesh, cellularity and vascularisation). **Wertheim K.Y.:** writing, review and editing of the manuscript, construction and calibration of the multicellular model. **Hervas-Raluy S.:** development of the mechanical model. **Sainz-DeMena D.:** development of the species transport model; co-author of the Python library used to generate the input data (geometry, mesh, cellularity and vascularisation). **Walker D.:** supervision of the Sheffield team, manuscript revision. **Chisholm R.:** implementation of the Cell model in high performance CUDA. **Viceconti, M.:** supervision of the Bologna team, manuscript revision. **Varella, V.:** implementation of the orchestrator framework. **Richmond P.:** creator of FLAMEGPU, supervision of the Sheffield team and high performance implementation. **Kasztelnik M.:** supervision of the Cyfronet team, support for running simulations on the HPC clusters. **Montero A.:** implementation of the sub-cellular model of chemotherapy effects. **Gregori-Puigjané, E.:** implementation of the protein inhibition models for the sub-cellular model of chemotherapy effects. **Mestres J.:** supervision of the Chemotargets team, manuscript revision. **García-Aznar J. M.:** supervision of the Zaragoza team, manuscript revision.

Declaration of Competing Interest

The authors declare that there is no conflict of interests.

² HP Enterprise Apollo 8000 Gen9, 2403 TFlops, based Intel Xeon E5-2680v3 CPUs, NVIDIA Tesla K40 and V100 GPUs, <https://www.cyfronet.pl/en/computers/15226,artykul,prometheus.html>.

³ <https://slurm.schedmd.com/>.

⁴ Huawei CH121L V5 Liquid-Cooled, 3.5 PFlops for the CPU (Intel Xeon Platinum) parts, and over 500 TFlops for the GPU (NVIDIA Tesla V100) parts, 290th on the Top500 list, https://www.cyfronet.pl/en/computers/18827,artykul,ares_supercomputer.html.

⁵ FormatServer THOR ERG21, 7.7 PFlops for the GPU (NVIDIA Tesla A100), 105th on the Top500 list, <https://www.cyfronet.pl/en/19073,artykul,athena.html>.

Acknowledgements

This work has received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 826494 (PRIMAGE). We gratefully acknowledge Poland's high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2022/015627. SHR was supported by the Government of Aragon (2019–23). We also declare that there are no competing interests to report in relation to this research project

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.cmpb.2023.107742](https://doi.org/10.1016/j.cmpb.2023.107742).

References

- [1] S. Ackermann, M. Cartolano, B. Hero, A. Welte, Y. Kahlert, A. Roderwieser, C. Bartenhagen, E. Walter, J. Gecht, L. Kerschke, R. Volland, R. Menon, J. M. Heuckmann, M. Gartlgruber, S. Hartlieb, K.O. Heinrich, K. Okonechnikov, J. Altmüller, P. Nürnberg, M. Fischer, A mechanistic classification of clinical phenotypes in neuroblastoma, *Science* 362 (6419) (2018) 1165–1170, <https://doi.org/10.1126/SCIENCE.AAT6768>.
- [2] I.M. Ambros, P.F. Ambros, Schwann cells in neuroblastoma, *Eur. J. Cancer* 31 (4) (1995) 429–434, [https://doi.org/10.1016/0959-8049\(95\)00051-J](https://doi.org/10.1016/0959-8049(95)00051-J).
- [3] C. Baeza-Delgado, L. Cerdá Alberich, J.M. Carot-Sierra, D. Veiga-Canuto, B. Martínez de las Heras, B. Raza, L. Martí-Bonmatí, A practical solution to estimate the sample size required for clinical prediction models generated from observational research on data, *Eur. Radiol. Exp.* 6 (1) (2022) 1–10, <https://doi.org/10.1186/S41747-022-00276-Y/TABLES/5>.
- [4] R. Bagatell, K. McHugh, A. Naranjo, C. Van Ryn, C. Kirby, P. Brock, K.A. Lyons, L. J. States, Y. Rojas, A. Miller, S.L. Volchenboum, T. Simon, B. Krug, S. Sarnacki, D. Valteau-Couanet, D. Von Schweinitz, B. Kammer, C. Granata, L. Pio, J. Nuchtern, Assessment of Primary Site Response in Children With High-Risk Neuroblastoma: an International Multicenter Study, *J. Clin. Oncol.* 34 (7) (2016) 740, <https://doi.org/10.1200/JCO.2015.63.2042>.
- [5] S. Benzekry, C. Sentis, C. Coze, L. Tessonnier, N. André, Development and Validation of a Prediction Model of Overall Survival in High-Risk Neuroblastoma Using Mechanistic Modeling of Metastasis, *JCO Clin. Cancer Inform.* 5 (2021) 81–90, <https://doi.org/10.1200/cci.20.00092>.
- [6] Berlin Institute of Health. (2020). Virtual Research Environment Architecture - BIH At Charité. <https://www.bihhealth.org/de/translation/netzwerk/digitale-medizin/bihcharite-virtual-research-environment/virtual-research-environment-architecture>.
- [7] H. Chen, J. Li, Y. Wang, P.K.S. Ng, Y.H. Tsang, K.R. Shaw, G.B. Mills, H. Liang, Comprehensive assessment of computational algorithms in predicting cancer driver mutations, *Genome Biol.* 21 (1) (2020), <https://doi.org/10.1186/S13059-020-01954-Z>.
- [8] C. Colson, H.M. Byrne, P.K. Maini, Combining Mechanisms of Growth Arrest in Solid Tumours: a Mathematical Investigation, *Bull. Math. Biol.* 84 (8) (2022) 1–24, <https://doi.org/10.1007/S11538-022-01034-2/FIGURES/8>.
- [9] M. Dahlhaus, A. Burkovski, F. Hertwig, C. Mussel, R. Volland, M. Fischer, K. M. Debatin, H.A. Kestler, C. Beltinger, Boolean modeling identifies Greatwall/MASTL as an important regulator in the AURKA network of neuroblastoma, *Cancer Lett.* 371 (1) (2016) 79–89, <https://doi.org/10.1016/J.CANLET.2015.11.025>.
- [10] B. de Melo Quintela, S. Hervás-Raluy, J.M. García-Aznar, D. Walker, K. Y. Wertheim, M. Viceconti, A theoretical analysis of the scale separation in a model to predict solid tumour growth, *J. Theor. Biol.* 547 (2022), 111173, <https://doi.org/10.1016/J.JTBI.2022.111173>.
- [11] A. Deutsch, J.M. Nava-Sedeño, S. Syga, H. Hatzikirou, BIO-LGCA: a cellular automaton modelling class for analysing collective cell migration, *PLoS Comput. Biol.* 17 (6) (2021), e1009066, <https://doi.org/10.1371/JOURNAL.PCBI.1009066>.
- [12] J.L. Dunster, H.M. Byrne, J.R. King, The resolution of inflammation: a mathematical model of neutrophil and macrophage interactions, *Bull. Math. Biol.* 76 (8) (2014) 1953–1980.
- [13] FLAME GPU. (n.d.). Retrieved March 10, 2023, from <https://flamegpu.com/>.
- [14] L. Friedel, A. Loewer, The guardian's choice: how p53 enables context-specific decision-making in individual cells, *FEBS J.* 289 (1) (2022) 40–52, <https://doi.org/10.1111/FEBS.15767>.
- [15] M. Gabelloni, L. Faggioni, R. Borgheresi, G. Restante, J. Shortrede, L. Tumminello, C. Scapicchio, F. Coppola, D. Cioni, I. Gómez-Rico, L. Martí-Bonmatí, E. Neri, Bridging gaps between images and data: a systematic update on imaging biobanks, *Eur. Radiol.* 32 (5) (2022) 3173–3186, <https://doi.org/10.1007/S00330-021-08431-6/FIGURES/4>.
- [16] R.A. Gatenby, J.S. Brown, Integrating evolutionary dynamics into cancer therapy, *Nat. Rev. Clin. Oncol.* 17 (11) (2020) 675–686, <https://doi.org/10.1038/s41571-020-0411-1>, 2020 17:11.
- [17] E.G. Greengard, Molecularly Targeted Therapy for Neuroblastoma, *Children* 5 (10) (2018), <https://doi.org/10.3390/CHILDREN5100142>.
- [18] J. Grote, R. Stisskind, P. Vaupel, Oxygen diffusivity in tumor tissue (DS-Carcinosarcoma) under temperature conditions within the range of 20–40°C, *Plügers Arch.* 372 (1) (1977) 37–42, <https://doi.org/10.1007/BF00582204>, 1977 372:1.
- [19] P. Han, D.M. Bartels, Temperature Dependence of Oxygen Diffusion in H₂O and D₂O, *J. Phys. Chem.* 100 (13) (1996) 5597–5602, <https://doi.org/10.1021/JP952903Y>.
- [20] Y. He, A. Kodali, D.I. Wallace, Predictive Modeling of Neuroblastoma Growth Dynamics in Xenograft Model After Bevacizumab Anti-VEGF Therapy, *Bull. Math. Biol.* 80 (8) (2018) 2026–2048, <https://doi.org/10.1007/S11538-018-0441-3>.
- [21] M.R. Hidalgo, A. Amadoz, C. Çubuk, J. Carbonell-Caballero, J. Dopazo, Models of cell signaling uncover molecular mechanisms of high-risk neuroblastoma and predict disease outcome, *Biol. Direct* 13 (1) (2018) 1–12, <https://doi.org/10.1186/S13062-018-0219-4/FIGURES/4>.
- [22] M.R. Hidalgo, A. Amadoz, C. Çubuk, J. Carbonell-Caballero, J. Dopazo, Models of cell signaling uncover molecular mechanisms of high-risk neuroblastoma and predict disease outcome, *Biol. Direct* (1) (2018) 13, <https://doi.org/10.1186/S13062-018-0219-4>.
- [23] M.T. Islam, S. Tang, C. Liverani, S. Saha, E. Tasciotti, R. Righetti, Non-invasive imaging of Young's modulus and Poisson's ratio in cancers in vivo, *Sci. Rep.* 10 (1) (2020) 1–12, <https://doi.org/10.1038/s41598-020-64162-6>, 2020 10:1.
- [24] M. Italia, F. Dercole, R. Lucchetti, Optimal chemotherapy counteracts cancer adaptive resistance in a cell-based, spatially-extended, evolutionary model, *Phys. Biol.* 19 (2) (2022), <https://doi.org/10.1088/1478-3975/AC509C>.
- [25] M. Italia, K.Y. Wertheim, S. Taschner-Mandl, D. Walker, F. Dercole, Mathematical Model of Clonal Evolution Proposes a Personalised Multi-Modal Therapy for High-Risk Neuroblastoma, *Cancers* 15 (7) (2023) 1986, <https://doi.org/10.3390/CANCERS15071986/S1>.
- [26] E.J. Jordan, K. Patil, K. Suresh, J.H. Park, Y.P. Mosse, M.A. Lemmon, R. Radhakrishnan, Computational Algorithms for In Silico Profiling of Activating Mutations in Cancer, *Cell. Mol. Life Sci.* 76 (14) (2019) 2663, <https://doi.org/10.1007/S00018-019-03097-2>.
- [27] M.N. Kamel Boulos, P. Zhang, Digital Twins: from Personalised Medicine to Precision Public Health, *J. Pers. Med.* (8) (2021) 11, <https://doi.org/10.3390/JPM11080745>.
- [28] P. Kamenova, A.v. Artemov, M.E. Kastriti, L. Faure, T.K. Olsen, J. Otte, A. Erickson, B. Semsch, E.R. Andersson, M. Ratz, J. Frisén, A.S. Tischler, R.R. de Krijger, T. Boudierlique, N. Akkuratova, M. Vorontsova, O. Gusev, K. Fried, E. Sundström, I. Adameyko, Single-cell transcriptomics of human embryos identifies multiple sympathoblast lineages with potential implications for neuroblastoma origin, *Nat. Genet.* 53 (5) (2021) 694–706, <https://doi.org/10.1038/s41588-021-00818-x>, 2021 53:5.
- [29] E. Kim, J.Y. Kim, J.Y. Lee, Mathematical Modeling of p53 Pathways, *Int. J. Mol. Sci.* (20) (2019) 20, <https://doi.org/10.3390/IJMS20205179>.
- [30] Kitowski, J., Turala, M., Wiatr, K., Dutka, L., Bubak, M., Szeplieniec, T., Radecki, M., Sterzel, M., Mosurska, Z., Pajak, R., Slota, R., Kurowski, K., Palak, B., Balcerak, B., Bala, P., Filocha, M., & Tylman, R. (2012). Polish computational research space for international scientific collaborations. Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 7203 LNCS(PART 1), 317–326. [10.1007/978-3-642-31464-3_32/COVER](https://doi.org/10.1007/978-3-642-31464-3_32/COVER).
- [31] R. Laubenbacher, A. Niarakis, T. Helikar, G. An, B. Shapiro, R.S. Malik-Sheriff, T. J. Segó, A. Knapp, P. Macklin, J.A. Glazier, Building digital twins of the human immune system: toward a roadmap, *NPJ Digital Med.* 5 (1) (2022) 1–5, <https://doi.org/10.1038/s41746-022-00610-z>, 2022 5:1.
- [32] C. Li, S. Wang, C. Li, Y. Yin, F. Feng, H. Fu, H. Wang, S. Chen, Improved risk stratification by PET-based intratumor heterogeneity in children with high-risk neuroblastoma, *Front. Oncol.* 12 (2022), 896593, <https://doi.org/10.3389/FONC.2022.896593/BIBTEX>.
- [33] E.A.B.F. Lima, D. Faghihi, R. Phillely, J. Yang, J. Virostko, C.M. Phillips, T. E. Yankeelov, Bayesian calibration of a stochastic, multiscale agent-based model for predicting in vitro tumor growth, *PLoS Comput. Biol.* 17 (11) (2021), e1008845, <https://doi.org/10.1371/JOURNAL.PCBI.1008845>.
- [34] S.D. Lombardo, M. Presti, K. Mangano, M.C. Petralia, M.S. Basile, M. Libra, S. Candido, P. Fagone, E. Mazzon, F. Nicoletti, A. Bramanti, Prediction of PD-L1 Expression in Neuroblastoma via Computational Modeling, *Brain Sci.* 9 (9) (2019) 221, <https://doi.org/10.3390/BRAINSCI9090221>, 2019, Vol. 9, Page 221.
- [35] S.D. Lombardo, M. Presti, K. Mangano, M.C. Petralia, M.S. Basile, M. Libra, S. Candido, P. Fagone, E. Mazzon, F. Nicoletti, A. Bramanti, Prediction of PD-L1 Expression in Neuroblastoma via Computational Modeling, *Brain Sci.* 9 (9) (2019), <https://doi.org/10.3390/BRAINSCI9090221>.
- [36] L. Martí-Bonmatí, Á. Alberich-Bayarri, R. Ladenstein, I. Blanquer, J.D. Segrelles, L. Cerdá-Alberich, P. Gkontra, B. Hero, J.M. García-Aznar, D. Keim, W. Jentner, K. Seymour, A. Jiménez-Pastor, I. González-Valverde, B. Martínez de las Heras, S. Essiaf, D. Walker, M. Rochette, M. Bubak, E. Neri, PRIMAGE project: predictive in silico multiscale analytics to support childhood cancer personalised evaluation empowered by imaging biomarkers, *Eur. Radiol. Exp.* 4 (1) (2020) 1–11, <https://doi.org/10.1186/S41747-020-00150-9/TABLES/4>.
- [37] K.K. Matthay, J.M. Maris, G. Schleiermacher, A. Nakagawara, C.L. Mackall, L. Diller, W.A. Weiss, Neuroblastoma, *Nat. Rev. Dis. Primers* 2 (1) (2016) 16078, <https://doi.org/10.1038/nrdp.2016.78>.
- [38] D. Mendez, A. Gaulton, A.P. Bento, J. Chambers, M. de Veij, E. Félix, M. P. Magariños, J.F. Mosquera, P. Mutowo, M. Nowotka, M. Gordillo-Marañón, F. Hunter, L. Junco, G. Mugumbate, M. Rodríguez-Lopez, F. Atkinson, N. Bosc, C. J. Radoux, A. Segura-Cabrera, A.R. Leach, ChEMBL: towards direct deposition of

- bioassay data, *Nucleic Acids Res.* 47 (D1) (2019) D930–D940, <https://doi.org/10.1093/NAR/GKY1075>.
- [39] J. Metzcar, Y. Wang, R. Heiland, P. Macklin, A Review of Cell-Based Computational Modeling in Cancer Biology, *JCO Clin. Cancer Inform.* 3 (3) (2019) 1–13, <https://doi.org/10.1200/CCI.18.00069>.
- [40] L. Moreno, G. Barone, S.G. DuBois, J. Molenaar, M. Fischer, J. Schulte, A. Eggert, G. Schleiermacher, F. Speleman, L. Chesler, B. Georger, M.D. Hogarty, M.S. Irwin, N. Bird, G.B. Blanchard, S. Buckland, H. Caron, S. Davis, B. de Wilde, A.D. J. Pearson, Accelerating drug development for neuroblastoma: summary of the Second Neuroblastoma Drug Development Strategy forum from Innovative Therapies for Children with Cancer and International Society of Paediatric Oncology Europe Neuroblastoma, *Eur. J. Cancer (Oxford : 1990)* 136 (2020) 52–68, <https://doi.org/10.1016/J.EJCA.2020.05.010>.
- [41] F. Mpekris, S. Angeli, A.P. Pirentis, T. Stylianopoulos, Stress-mediated progression of solid tumors: effect of mechanical stress on tissue oxygenation, cancer cell proliferation, and drug delivery, *Biomech. Model. Mechanobiol.* 14 (6) (2015) 1391–1402, <https://doi.org/10.1007/S10237-015-0682-0>.
- [42] B. Muz, P. de la Puente, F. Azab, A.K. Azab, The role of hypoxia in cancer progression, angiogenesis, metastasis, and resistance to therapy, *Hypoxia* 3 (2015) 83.
- [43] P. Nowakowski, M. Bubak, T. Bartysiński, T. Gubała, D. Haręźlak, M. Kasztelnik, M. Malawski, J. Meizner, Cloud computing infrastructure for the VPH community, *J. Comput. Sci.* 24 (2018) 169–179, <https://doi.org/10.1016/J.JOCS.2017.06.012>.
- [44] R. Nussinov, H. Jang, C.J. Tsai, F. Cheng, Review: precision medicine and driver mutations: computational methods, functional assays and conformational principles for interpreting cancer drivers, *PLoS Comput. Biol.* (3) (2019) 15, <https://doi.org/10.1371/JOURNAL.PCBI.1006658>.
- [45] C. Okamatsu, W.B. London, A. Naranjo, M.D. Hogarty, J.M. Gastier-Foster, A. T. Look, M. LaQuaglia, J.M. Maris, S.L. Cohn, K.K. Matthay, Clinicopathological characteristics of ganglioneuroma and ganglioneuroblastoma: a report from the CCG and COG, *Pediatr. Blood. Cancer* 53 (4) (2009) 563–569.
- [46] J. Park, Y. Zhang, N. Vykhodtseva, F.A. Jolesz, N.J. McDannold, The kinetics of blood brain barrier permeability and targeted doxorubicin delivery into brain induced by focused ultrasound, *J. Control. Release* 162 (1) (2012) 134, <https://doi.org/10.1016/J.JCONREL.2012.06.012>.
- [47] Project | CHIC. (2017). <https://chic-vph.eu/project/>.
- [48] B. Qiu, K.K. Matthay, Advancing therapy for neuroblastoma, *Nat. Rev. Clin. Oncol.* 19 (8) (2022) 515–533, <https://doi.org/10.1038/s41571-022-00643-z>, 2022 19:8.
- [49] D. Sainz-Demena, J. Manuel García-Aznar, M. Ángeles Pérez, C. Borau, Im2mesh: a Python Library to Reconstruct 3D Meshes from Scattered Data and 2D Segmentations, Application to Patient-Specific Neuroblastoma Tumour Image Sequences, *Appl. Sci.* 12 (22) (2022) 11557, <https://doi.org/10.3390/AP122211557>, 2022, Vol. 12, Page 11557.
- [50] D. Sainz-DeMena, W. Ye, M.Á. Pérez, J.M. García-Aznar, A finite element based optimization algorithm to include diffusion into the analysis of DCE-MRI, *Eng. Comput.* 38 (5) (2022) 3849–3865, <https://doi.org/10.1007/S00366-022-01667-W/FIGURES/8>.
- [51] H. Shimada, I.M. Ambros, L.P. Dehner, J. Hata, V.V. Joshi, B. Roald, Terminology and morphologic criteria of neuroblastic tumors: recommendations by the International Neuroblastoma Pathology Committee, *Cancer Interdisciplin. Int. J. Am. Cancer Soc.* 86 (2) (1999) 349–363.
- [52] R.J. Shipley, P.W. Sweeney, S.J. Chapman, T. Roose, A four-compartment multiscale model of fluid and drug distribution in vascular tumours, *Int. J. Numer. Method Biomed. Eng.* 36 (3) (2020), e3315, <https://doi.org/10.1002/CNM.3315>.
- [53] T. Simon, B. Hero, J.H. Schulte, H. Deubzer, P. Hundsdoerfer, D. von Schweinitz, J. Fuchs, M. Schmidt, V. Prasad, B. Krug, B. Timmermann, I. Leuschner, M. Fischer, T. Langer, K. Astrahantseff, F. Berthold, H. Lode, A. Eggert, 2017 GPOH Guidelines for Diagnosis and Treatment of Patients with Neuroblastic Tumors, *Klin. Padiatr.* 229 (3) (2017) 147–167, <https://doi.org/10.1055/S-0043-103086>.
- [54] V. Smith, J. Foster, High-Risk Neuroblastoma Treatment Review, *Children* 5 (9) (2018), <https://doi.org/10.3390/CHILDREN5090114>.
- [55] P. Song, J. An, M.H. Zou, Immune clearance of senescent cells to combat ageing and chronic diseases, *Cells* 9 (3) (2020) 671.
- [56] P.S. Tofts, A.G. Kermode, Measurement of the blood-brain barrier permeability and leakage space using dynamic MR imaging. 1. Fundamental concepts, *Magn. Reson. Med.* 17 (2) (1991) 357–367, <https://doi.org/10.1002/MRM.1910170208>.
- [57] V.P. Tolbert, K.K. Matthay, Neuroblastoma: clinical and Biological Approach to Risk Stratification and Treatment, *Cell Tissue Res.* 372 (2) (2018) 195, <https://doi.org/10.1007/S00441-018-2821-2>.
- [58] E.R. Tucker, I. Jiménez, L. Chen, A. Bellini, C. Gorrini, E. Calton, Q. Gao, H. Che, E. Poon, Y. Jamin, B. Martins Da Costa, K. Barker, S. Shrestha, J.C. Hutchinson, S. Dhariwal, A. Goodman, E. del Nery, P. Gestraud, J. Bhalshankar, G. Schleiermacher, Combination Therapies Targeting Alk-Aberrant Neuroblastoma in Preclinical Models, *Clin. Cancer Res.* (2023), <https://doi.org/10.1158/1078-0432.CCR-22-2274>.
- [59] D. Veiga-Canuto, L. Cerdà-Alberich, C. Sangüesa Nebot, B. Martínez de las Heras, U. Pötschger, M. Gabelloni, J.M. Carot Sierra, S. Taschner-Mandl, V. Düster, A. Canete, R. Ladenstein, E. Neri, I. Martí-Bonmatí, Comparative Multicentric Evaluation of Inter-Observer Variability in Manual and Automatic Segmentation of Neuroblastic Tumors in Magnetic Resonance Images, *Cancers* 14 (15) (2022) 3648, <https://doi.org/10.3390/CANCERS14153648>, 2022, Vol. 14, Page 3648.
- [60] M. Viceconti, M.A. Juárez, C. Curreli, M. Pennisi, G. Russo, F. Pappalardo, Credibility of in Silico Trial Technologies-A Theoretical Framing, *IEEE J. Biomed. Health Inf.* 24 (1) (2020) 4–13, <https://doi.org/10.1109/JBHI.2019.2949888>.
- [61] D. Vidal, R. García-Serna, J. Mestres, Ligand-based approaches to in silico pharmacology, *Methods Mol. Biol.* 672 (2011) 489–502, https://doi.org/10.1007/978-1-60761-839-3_19.
- [62] J.X. Wang, Z.Y. Cao, C.X. Wang, H.Y. Zhang, F.L. Fan, J. Zhang, X.Y. He, N.J. Liu, J. Bin Liu, L. Zou, Prognostic impact of tumor size on patients with neuroblastoma in a SEER-based study, *Cancer Med.* 11 (14) (2022) 2779–2789, <https://doi.org/10.1002/CAM4.4653>.
- [63] P. Wang, H.Y. Wang, X.J. Gao, H.X. Zhu, X.P. Zhang, F. Liu, W. Wang, Encoding and Decoding of p53 Dynamics in Cellular Response to Stresses, *Cells* 12 (3) (2023), <https://doi.org/10.3390/CELLS12030490>.
- [64] Weir, P., Ellerweg, R., Payne, S., Reuter, D., Alhonnoro, T., Voglreiter, P., Mariappan, P., Pollari, M., Park, C.S., Voigt, P., van Oostenbrugge, T., Fischer, S., Kalmar, P., Futterer, J., Stiegler, P., Zangos, S., Flanagan, R., Moche, M., & Kolesnik, M. (2016). Go-Smart: open-Ended, Web-Based Modelling of Minimally Invasive Cancer Treatments via a Clinical Domain Approach. https://www.researchgate.net/publication/324081147_Go-Smart_Open-Ended_Web-Based_Modelling_of_Minimally_Invasive_Cancer_Treatments_via_a_Clinical_Domain_Approach?channel=doi&linkId=5abcc6f30f7e9bfc04561b81&showFulltext=true.
- [65] T. Weiss, S. Taschner-Mandl, L. Janker, A. Bileck, F. Rifatbegovic, F. Kromp, H. Sorger, M.O. Kauer, C. Frech, R. Windhager, C. Gerner, P.F. Ambros, I. M. Ambros, Schwann cell plasticity regulates neuroblastic tumor cell differentiation via epidermal growth factor-like protein 8, *Nat. Commun.* 12 (1) (2021) 1–19, <https://doi.org/10.1038/s41467-021-21859-0>, 2021 12:1.
- [66] K.Y. Wertheim, B.L. Puniy, I. Alyssa La Fleur, A.R. Shah, M. Barberis, T. Helikar, A multi-approach and multi-scale platform to model CD4+ T cells responding to infections, *PLoS Comput. Biol.* 17 (8) (2021), e1009209, <https://doi.org/10.1371/JOURNAL.PCBI.1009209>.
- [67] J. Wienke, M.P. Dierselhuys, G.A.M. Tytgat, A. Künkele, S. Nierkens, J.J. Molenaar, The immune landscape of neuroblastoma: challenges and opportunities for novel therapeutic strategies in pediatric oncology, *Eur. J. Cancer* 144 (2021) 123–150, <https://doi.org/10.1016/J.EJCA.2020.11.014>.
- [68] A. Zafar, W. Wang, G. Liu, W. Xian, F. McKeon, J. Zhou, R. Zhang, Targeting the p53-MDM2 Pathway for Neuroblastoma Therapy: rays of Hope, *Cancer Lett.* 496 (2021) 16, <https://doi.org/10.1016/J.CANLET.2020.09.023>.
- [69] E. Zaika, J. Wei, D. Yin, C. Andl, U. Moll, W. El-Rifai, A.I. Zaika, p73 protein regulates DNA damage repair, *FASEB J.* 25 (12) (2011) 4406–4414.