

## RESEARCH ARTICLE

# Hybrid Automata Library: A flexible platform for hybrid modeling with real-time visualization

Rafael R. Bravo<sup>1\*</sup>, Etienne Baratchart<sup>1</sup>, Jeffrey West<sup>1</sup>, Ryan O. Schenck<sup>1,2</sup>, Anna K. Miller<sup>1</sup>, Jill Gallaher<sup>1</sup>, Chandler D. Gatenbee<sup>1</sup>, David Basanta<sup>1</sup>, Mark Robertson-Tessi<sup>1</sup>, Alexander R. A. Anderson<sup>1\*</sup>

**1** Integrated Mathematical Oncology, Moffitt Cancer Center, Tampa, Florida, United States of America, **2** Wellcome Centre for Human Genetics, University of Oxford, Oxford, United Kingdom

\* [rafael.bravo@moffitt.org](mailto:rafael.bravo@moffitt.org) (RRB); [alexander.anderson@moffitt.org](mailto:alexander.anderson@moffitt.org) (ARAA)



## OPEN ACCESS

**Citation:** Bravo RR, Baratchart E, West J, Schenck RO, Miller AK, Gallaher J, et al. (2020) Hybrid Automata Library: A flexible platform for hybrid modeling with real-time visualization. *PLoS Comput Biol* 16(3): e1007635. <https://doi.org/10.1371/journal.pcbi.1007635>

**Editor:** Roeland M.H. Merks, Universiteit Leiden, NETHERLANDS

**Received:** January 23, 2019

**Accepted:** January 6, 2020

**Published:** March 10, 2020

**Copyright:** © 2020 Bravo et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** The data found in the Results section can be generated by running the competitive release example model found in the Examples folder of HAL. See the manual for instructions on how to install HAL and run the example.

**Funding:** This work was possible through the generous support of National Institutes of Health funding, (A.R.A.A) and (M.R.T) acknowledge National Cancer Institute U54CA193489, <https://www.cancer.gov/>. (A.R.A.A) and (R.B)

## Abstract

The Hybrid Automata Library (HAL) is a Java Library developed for use in mathematical oncology modeling. It is made of simple, efficient, generic components that can be used to model complex spatial systems. HAL's components can broadly be classified into: on- and off-lattice agent containers, finite difference diffusion fields, a GUI building system, and additional tools and utilities for computation and data collection. These components are designed to operate independently and are standardized to make them easy to interface with one another. As a demonstration of how modeling can be simplified using our approach, we have included a complete example of a hybrid model (a spatial model with interacting agent-based and PDE components). HAL is a useful asset for researchers who wish to build performant 1D, 2D and 3D hybrid models in Java, while not starting entirely from scratch. It is available on GitHub at <https://github.com/MathOnco/HAL> under the MIT License. HAL requires the Java JDK version 1.8 or later to compile and run the source code.

## Author summary

In this paper we introduce the Hybrid Automata Library (HAL) with the purpose of simplifying the implementation and sharing of hybrid models for use in mathematical oncology. Hybrid modeling is used in oncology to create spatial models of tissue, typically by modeling cells using agent-based techniques, and by modeling diffusible chemicals using partial differential equations (PDEs). HAL's key components are designed to run agent-based models, PDEs, and visualization. The components are standardized and are completely decoupled, so models can be built with any combination of them. We first explore the philosophy behind HAL, then summarize the components. Lastly we demonstrate how the components work together with an example of a hybrid model, and a walk-through of the code used to construct it. HAL is open-source and will produce identical results on any machine that supports Java 8 and above, making it highly portable. We recommend HAL to modelers interested in spatial dynamics, even those outside of

acknowledge National Cancer Institute UH2CA203781 <https://www.cancer.gov/>. (D.B), (E. B), and (A.K.M) acknowledge National Cancer Institute U01CA202958-01 <https://www.cancer.gov/>. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** No authors have competing interests.

mathematical oncology, as the components are general enough to facilitate a variety of model types. A community page that provides a download link and online documentation can be found at <https://halloworld.org>.

This is a *PLOS Computational Biology* Software paper.

## Introduction

The Hybrid Automata Library (HAL) was created to assist the growing mathematical oncology community with a common framework to facilitate building and visualizing hybrid models. Hybrid models in oncology usually represent cells (both of the tumor and of the surrounding tissue) using agent-based modeling (ABMs) and the concentrations of relevant chemicals (drugs, resources and signaling molecules) as continuous partial differential equations (PDEs). These models can simulate local interactions between cells with complex internal dynamics and decision-making processes while also allowing cells to interact with the PDE concentration fields in their local environment.

Hybrid models have been widely adopted within the Mathematical Oncology community to model many aspects of cancer [1–8]. A unique strength of the hybrid modeling approach is that it allows for a mechanistic understanding of the ecological feedback between the cancer cells and their tissue environment. Cancer cell agents can be modeled as a part of a larger tissue structure and interact with the systems that normally maintain homeostasis [9–15]. Drugs may be subsequently introduced to add additional selective pressure to the model, and to observe the long-term effects on the evolving tumor [16]. A better understanding of these selection dynamics can be used to help develop more effective drug sequences to prevent cancer from resisting therapy and to develop evolutionary therapies to control cancers that cannot be cured with maximum tolerated dose [17–19]. Further realism can be incorporated by initializing spatial models with clinical or histological data [20, 21].

Whilst a number of agent-based modeling frameworks have been used for tissue modeling, HAL distinguishes itself primarily through its flexibility. Instead of a fixed modeling environment, HAL provides a lightweight set of generic tools to help modelers build their own workflows without needing to reinvent commonly used functions.

Many frameworks facilitate model building under specific assumptions. Cells can be approximated as spheres or ellipsoids, with Newtonian adhesion-repulsion dynamics used to efficiently simulate large populations. PhysiCell [22], CellSys [23], BioCellion [24], Timothy [25], Yalla [26], and Episim [27] make this assumption. Another assumption called Cellular Potts represents cells as composites of lattice positions, which can migrate over time in response to energy minimization functions. This allows for cell deformation but is typically more computationally expensive. CompuCell3D [28], Morpheus [29], and the Tissue Simulation Toolkit [30] make this assumption. HAL does not include the same depth in the domains specific to these frameworks, but uses a broader design to provide the capacity for a variety of approaches.

Some frameworks that also take a broad approach are Chaste [31], NetLogo [32], Repast [33] and MASON [34]. Chaste uses a modular system for model building, in which modular rules are composed to define behavior, and behaviors that are not currently represented can be

added as new modules. This modular approach allows for very rapid prototyping, and increases the reproducibility of results. NetLogo provides an accessible integrated model development and interaction environment, and a custom scripting language to simplify coding, making it a great choice for new modelers/coders. Repast and MASON both feature composable spatial containers and built-in visualization/interaction tools, but differ in how they handle agent scheduling. Repast uses a hierarchical nesting approach to group agents into sets that will all execute actions, while in MASON optionally repeating agent step functions are added individually per agent. A comparative list of useful features and common built-in assumptions pertinent to tissue modeling are summarized in [Table 1](#).

HAL shares many characteristics with these frameworks, but differentiates itself with a minimal, decentralized design made up of independent building blocks that are thematically similar. There is no centralized controller or scheduler, so the modeler designs the logical flow and the scheduling of interactions between model components. This removes common presuppositions or requirements made by schedulers in other frameworks (e.g. when/how models should be visualized, when their step logic should run, when model resources should be created or destroyed, etc.) and leaves these decisions up to the modeler. This cuts down on any unnecessary use of resources by the modeling system, and increases modeling flexibility. These considerations have led to a lightweight framework that is easy to use, highly flexible, and effective within the scope of hybrid modeling, agent-based modeling, and the solving of reaction-diffusion PDEs using finite differences. HAL was designed with mathematical oncology in mind, but is general enough to facilitate modeling systems from many domains (e.g. ecology [37], development, population dynamics, and network theory). While some familiarity with the Java programming language is recommended for new users of HAL, we imagine that its simplicity and explicit nature could make it a useful educational platform.

**Table 1. Comparison of HAL with other agent-based modeling frameworks commonly used in tissue modeling as of December 2019.** Features are marked according to whether there exists a built-in user accessible implementation.

Feature	HAL	Cha	Rep	Mas	Net	Phy	Cel	Bio	Tim	Yal	Epi	Com	Mor	TST
On-Lattice ABMs	✓	✓	✓	✓	✓								✓	
Off-Lattice Point ABMs	✓	✓	✓	✓	✓									
Off-Lattice Spherical ABMs	✓	✓				✓	✓	✓	✓	✓	✓			
Voronoi Tessellation ABMs		✓	✓											
Cellular Potts ABMs		✓										✓	✓	✓
Multinomial Population ABMs [35]	✓													
Diffusion PDEs	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Advection PDEs	✓	✓					✓					✓	✓	
SBML Compatible [36]						✓					✓	✓	✓	
Real-Time Visualization	✓		✓	✓	✓		✓				✓		✓	✓
Single-Model Parallelization		✓	✓			✓	✓	✓	✓	✓	✓	✓	✓	✓
Windows Compatible	✓	✓	✓	✓	✓	✓	✓		✓		✓	✓	✓	✓
Mac Compatible	✓	✓	✓	✓	✓	✓				✓		✓	✓	✓
Linux Compatible	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
User-Facing Programming Language*	J	C	R	J	N	C	C	C	C	U	I	P	M	C

Framework Abbreviations: **HAL**:Hybrid Automata Library, **Cha**:Chaste, **Rep**:Repast, **Mas**:Mason, **Net**:Netlogo, **Phy**:Physicell, **Cel**:CellSys, **Bio**:Biocellion, **Tim**:Timothy, **Yal**:Yalla, **Epi**:Episim, **Com**:CompuCell3D, **Mor**:Morpheus, **TST**:Tissue Simulation Toolkit

\* User-Facing Programming Languages: **J**: Java, **C**: C/C++, **R**: Relogo/Java/Groovy, **N**: NetLogo Programming Language, **U**: CUDA/C++, **I**: Graphical Interface, **P**: Python/XML, **M**: Morpheus model description language

<https://doi.org/10.1371/journal.pcbi.1007635.t001>

The main components of HAL consist of  $n$ -dimensional (0D,1D,2D,3D) grids that hold Agents, 1D,2D, and 3D finite difference PDE fields, 2D and 3D visualization tools, and methods for sampling distributions and data recording. In this paper we will discuss the philosophy behind these components, then look at their design and capabilities in more detail. See the manual for a complete reference on how to use these components [38].


## Design and implementation

### Design philosophy

In the next section, we discuss some of the design decisions that have driven the architecture of HAL.

**Decentralization and extensibility.** HAL's components can each function independently, making HAL fundamentally decentralized. This permits any number of components to be used in a single model, with the use of spatial queries to combine components, as seen in Fig 1. Decentralization also allows modelers to choose only the components of HAL that are of interest for their project. These components can be easily mixed and matched with other software, such as using the AgentGrids with a different PDE solver, or using the GUI and Visualization components with a different modeling system. Modelers can also structure the modeling components in ways that are best suited to their workflow, for example by reusing the same model resources across multiple runs for more efficient parameter sweeping, or visualizing many models with a single visualization to compare dynamics across them, or inversely to use multiple visualizations to observe several aspects of a running model simultaneously. Decentralization also makes adding new components more manageable and easier to test without adding bulk or modifications to the existing platform.

Given the incremental nature of many scientific endeavors, we also wanted to allow models and components to be extended and modified. Java's extension architecture provides an excellent environment for layered development. As an example of the extensibility of HAL, the built-in Spherical Agent types (SphericalAgent2D, SphericalAgent3D) extend the Point Agent



**Fig 1. HAL's decentralized design helps build complex models out of simple components.** The highlighted on-lattice agent in the topmost grid searches for local overlaps with several other grids and PDEs. These overlaps can be used in a model to generate spatial interactions.

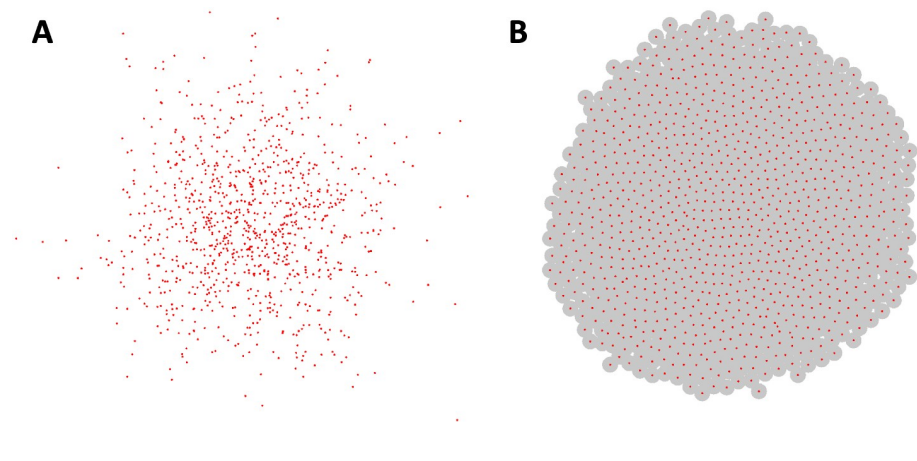
<https://doi.org/10.1371/journal.pcbi.1007635.g001>

types (AgentPT2D, AgentPT3D). By default, Point Agents have no radius and will not collide with each other. This behavior can be useful for modeling phenomena such as the Brownian motion of small particles, as visualized in Fig 2a. Spherical Agents extend Point Agents by adding an additional radius variable and velocity component variables. These properties combined with added functions for summing force vectors in response to overlap allow for a Newtonian adhesion-repulsion spherical model of spatial agents. This behavior can be useful for modeling tissue formation, as visualized in Fig 2b. A more complete description of the extension architecture of the Agent and Grid types is included in Fig D in S1 Text.

It is also possible to extend completed models using the same approach. For example, grids and agents from published models can be used as a scaffold on which to do additional studies. This allows for followup studies to focus on implementing whatever additional assumptions and functionality they need, while leaving intact the base model code with all of its published assumptions.

**Language choice.** In designing HAL we have tried to balance an adherence to performance, memory management, simplicity, stability, and extensibility. The Java language itself balances these considerations very well, making it a suitable basis for HAL. Commonly used low-level languages such as C, C++, and FORTRAN will typically run at similar speeds or faster than Java. However, these languages do not have the same safety guarantees as they permit out-of-bounds memory accesses and memory leaks. Higher level languages, such as Python, while more flexible and syntactically intuitive than Java, are typically significantly slower. The fact that Java is entirely cross-platform is also ideal. Many other languages, including the others mentioned above, don't run on a virtual machine and require maintaining multiple versions of libraries for cross-platform compatibility, and may not perform identically across platforms. Java is also one of the most commonly used and taught programming languages today, which helps facilitate the adoption of HAL by new users. Models can also be built with HAL using any languages that can interop with Java: such as Scala, Groovy, Kotlin and Clojure.

**Simplicity and stability.** An important design principle was to make HAL simple to use without sacrificing performance. Simplicity makes HAL easy to learn and forces the



**Fig 2. Off-lattice agent examples.** Each example contains 1000 agents. (A) Example of 2D Point Agents modeling Brownian motion. The Point Agents move freely and cannot collide. (B) Example of 2D Spherical Agents modeling a growing cell colony. The agents will push apart from each other to a uniform density. Agent radii are shown as gray circles around their centers. Examples Displayed using the OpenGL2DWindow object.

<https://doi.org/10.1371/journal.pcbi.1007635.g002>

components to be more generic, meaning that the same components can be applied to a greater variety of modeling problems. There is also a consistency to each framework component, such that learning to use some components is often sufficient to grasp the others, and using them in combination is intuitive.

Another key design principle is stability, which is achieved in three ways:

**Encapsulation.** By providing safe interaction functions and preventing direct interaction with component internals. For example, modelers are not permitted to directly modify the position properties of agents. Instead, they must call the provided movement functions that also update the grid position of the agents for future spatial queries.

**Defensive Programming.** By including checks in functions for invalid inputs. The program halts and throws an error message immediately when one of these problematic inputs occurs, such as inputting a diffusion rate constant that is unstable, or moving two unstackable agents to the same position. This allows the modeler to see the problem as it arises, rather than seeing its effects later down the line. While this may entail a small performance impact, it ensures that modelers are rarely confused by the behavior of components, and prevents their malfunction.

**Testing.** By testing HAL's components. HAL's algorithms are tested in a series of small test programs. These tests help ensure confidence in the algorithm implementations. Unit tests verify three main areas: PDE algorithm accuracy, ABM aggregate behavior, and mathematical utilities. We also verified L1, L2, and  $L_\infty$  convergence in time and space for the PDEGrid functions, and diffusion modeled using the PopulationGrid (see Fig E—AZ in [S1 Text](#)).

**Speed and memory management.** Much of the performance capability of HAL comes directly from its decentralized design. Having no built-in scheduler/underlying structure means that there is comparatively little work that the program does that the modeler is unaware of. This combined with the modular components and utilities allows modelers the flexibility to incorporate only the functionality that they need.

HAL also prioritizes performance in its algorithmic implementation. HAL includes efficient PDE solving algorithms, such as the ADI (alternating direction implicit) method, and uses efficient distribution sampling algorithms (see the Unit Tests section of [S1 Text](#)). The integrated visualization tools are also highly efficient, using Swing BufferedImages for lattice-based visualization, and LWJGL OpenGL for 2D and 3D polygon graphics. Whenever possible, primitives and arrays are used to store data rather than classes, which takes advantage of Java's optimization for these simpler data types. Java is also an inherently fast language, which helps efficiently execute agent behavioral logic.

There is a memory footprint consideration with most of HAL's assets. A common criticism of Java applications is that they tend to use a lot of memory and are slowed down by Java's "garbage collector" which deletes objects that are no longer being used. To sidestep these memory issues, most of the objects generated internally by HAL are recycled rather than discarded. This reuse also has a performance benefit: if a function using the same object is called many times sequentially, the object will be faster to access in the computer's memory because it was already cached from the earlier calls.

A key example of this reuse: when agents die and are removed from the model during a simulation run, the removed agents are kept internally and will be returned again for


re-initialization when a new agent is requested. Agent recycling ensures that the number of agents that the model creates over a complete model run is capped to the maximum population that exists in the model at one time.

We profiled the performance scaling of our On-Lattice and Spherical Agent Types, as well as our ADI algorithm, and they all scale approximately linearly up to around 10 million agents and diffusible lattice points, after which memory constraints become prohibitive. (see Fig A—C in [S1 Text](#)).

## Component overview

We now move from the abstract discussion of the unifying principles behind HAL to a look at its core components in more detail. Although it may seem that learning how to use these components would be a difficult task given their number and variety, all components were designed with a consistent API (Application Programming Interface), which makes changing between agent/grid types and learning their methods much easier.

**AgentGrids.** AgentGrids are used as spatial containers for agents. They come in 1D, 2D, 3D, and non-spatial types. An example usage of a 3D AgentGrid is shown in [Fig 3](#). These objects hold populations of agents that exist either bound to a lattice, or are free to move continuously. Internally, AgentGrids are composed of two data-structures: an agent list for agent iteration, and an agent lattice for spatial queries (even off-lattice agents are stored on a lattice for quick access). The agent list can be shuffled at every iteration to randomize iteration order, and the list holds onto removed agents to facilitate object recycling.



**Fig 3. An example of a 3D on-lattice hybrid model of tumor cells spreading through tissue (located in Examples/5Example3D).** The light red vertical lines represent vessels, and the blue dots represent tumor cells. The cell color goes from pink to blue depending on how much oxygen is locally available. The oxygen availability is also indicated via a heatmap on the floor of the display. With a cell radius of 10  $\mu\text{m}$ , the domain has dimensions  $1.6 \times 1.6 \times 0.4$  mm. Displayed using the OpenGL3DWindow object.

<https://doi.org/10.1371/journal.pcbi.1007635.g003>

**Agents.** There are 10 base types of agent, introduced in Table 2. The SQ and PT suffixes refer to whether the agents are imagined to exist as lattice bound squares/voxels, or as non-volumetric points in space.

Agent objects are always bound to a grid. In their base class form, agents keep track of their position on the grid and their age. Newly created agents are not included in the same iteration loop in which they are created, to prevent infinite loops of “runaway proliferation.” The base agent classes can be extended to include additional state properties and methods as needed.

**PDEGrids.** The PDE Grids consist of either a 1D, 2D, or 3D lattice of concentrations. PDE grids contain functions that will solve reaction-diffusion equations. PDE function operations are accumulated on a separate lattice so they can be applied all at once in a simultaneous update. Currently implemented PDE solution methods include:

- Forward difference in time and 2nd order central difference in space diffusion
- ADI Diffusion [39]
- 1st order upwind finite difference advection for incompressible flows [40]
- 1st order finite volume upwind advection for compressible flows
- Modification of values at single lattice positions to facilitate reaction with agents or other sources/sinks.

Most of these methods are flexible, allowing for spatially heterogeneous diffusion rates and advection velocities as well as different boundary conditions such as periodic, Dirichlet, and zero-flux Neumann.


**PopulationGrids.** When dealing with large numbers of cells, especially at the tissue scale, it is often useful to model the spatial-temporal cellular dynamics as densities using PDEs [41–43]. This approach is effective in modeling dynamics at the population level, however it does not directly represent individuals nor rare stochastic events. To bridge between PDE modeling and agent-based stochasticity, we introduce the PopulationGrid. The PopulationGrid is similar to the AgentGrid, but stores homogeneous populations at each lattice position rather than storing individual agents. The Binomial and Multinomial functions are used to update the populations in response to probabilistic events [35]. These events may consist of migration, proliferation, death, or changes in agent state (eg. mutation). The benefit of this type of grid is that it can efficiently handle arbitrarily large populations at each lattice position (up to the

**Table 2. The 10 base agent types in HAL.** The differences between them are displayed in each column. Stackable refers to whether multiple agents can exist on the same lattice position. The complete type heirarchy used in HAL is described in Fig D in S1 Text.

Name	Spatial Dimension	Lattice Bound	Stackable
Agent0D	0	N/A	N/A
AgentSQ1D	1	yes	yes
AgentSQ1Dunstackable	1	yes	no
AgentPT1D	1	no	yes
AgentSQ2D	2	yes	yes
AgentSQ2Dunstackable	2	yes	no
AgentPT2D	2	no	yes
AgentSQ3D	3	yes	yes
AgentSQ3Dunstackable	3	yes	no
AgentPT3D	3	no	yes

<https://doi.org/10.1371/journal.pcbi.1007635.t002>





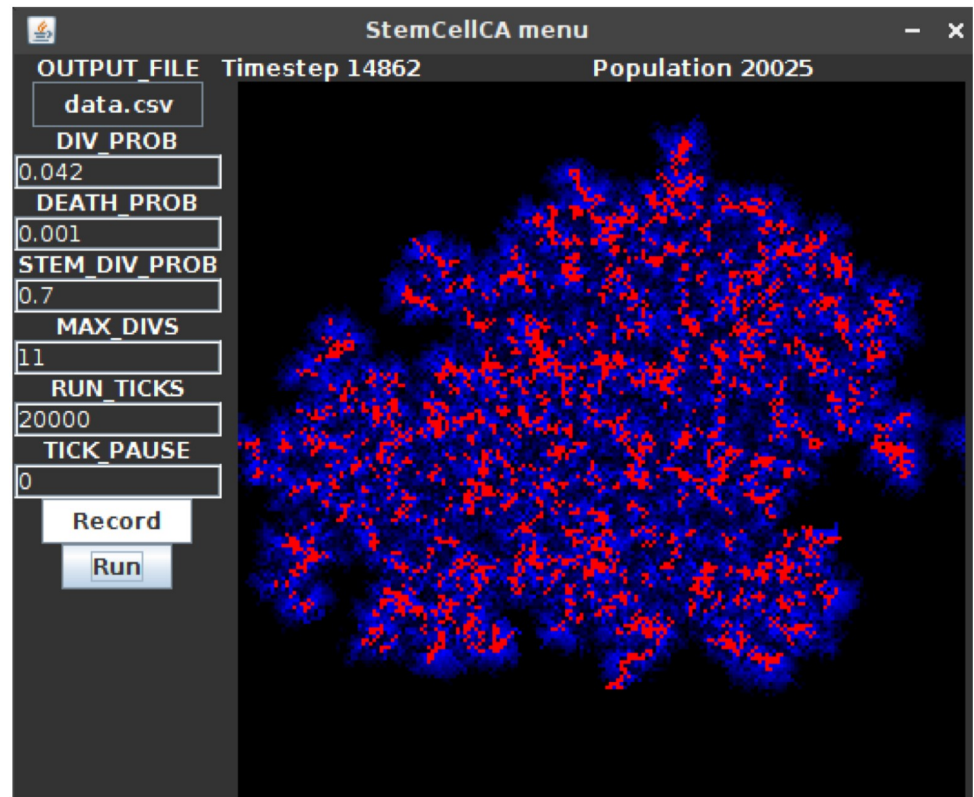
**Fig 4. An example PopulationGrid model of tumor growth and mutation in 2D (located in Examples/PopulationTumor).** At the start 1000 cells are placed in the center of the domain. Cells turn over and migrate over time. Two fitness-increasing mutations are possible, and can occur with probability  $2 \times 10^{-6}$  per cell division. The colors white, red, green, and blue indicate density of cells with no mutations, one mutation, the other mutation, and both mutations respectively. Each lattice position in the domain has a carrying capacity of 100,000 cells. If we assume that cells are growing in cubic volumes, then each position represents an overall area of  $1 \text{ mm}^3$  for a total domain size of  $150 \times 150 \times 1 \text{ mm}$ . Displayed using the GridWindow object.

<https://doi.org/10.1371/journal.pcbi.1007635.g004>

maximum size of a Java Long  $\approx 9 \times 10^{18}$ ). The downside is that all of the agents at one lattice position in a given PopulationGrid are considered identical. Changes in agent state can be implemented by shifting populations between multiple PopulationGrids or between lattice positions. Fig 4 demonstrates how this approach may be used to model stochastic occurrences of very unlikely events over large populations, such as specific driver mutations occurring in a growing tumor.

**Graphical User Interface (GUI).** The GUI building system consists of the following components:

- UIWindow: a window that displays GUI sub-components which are added in columns. the UIWindow will automatically scale to the appropriate size to fit all sub-components. The following five sub-components can be added:
  - UIGrid: a grid of pixels whose values are set individually. These are typically used to plot agent positions and diffusible concentrations, and can be easily output in GIF or PNG formats.
  - UIPlot: an extension of the UIGrid, the UIPlot is used to create real-time plots. The UIPlot will automatically resize to fit points that fall out of its bounds.
  - UILabel: a label that presents modifiable text.
  - UIButton: a button that executes a function when clicked
  - UIInputFields: various fields that accept bounded input of Integers, Doubles, Strings, Booleans, File Creation/Selection, and Combo boxes
- Window2DOpenGL/Window3DOpenGL: visualization windows that use OpenGL to efficiently render polygon graphics.
- GridWindow: A shortcut to generate a UIWindow with a single UIGrid component embedded. This simple component is used in the results section example.
- GifMaker: An object that can turn UIGrid visualization snapshots into GIFs [44].



**Fig 5. An example UIWindow GUI.** When the “Run” UIButton (bottom left) is clicked, the UIGrid component (right) displays a running model that is parameterized with the given UIInputField settings (left). In this example model based on [45], the red cells are stem cells, and the blue cells are differentiated cells. Differentiated cells have a limited number of divisions and therefore can only spread a limited distance from the stem cells. UILabels (top) show the current timestep and population size. With a cell radius of 10  $\mu\text{m}$ , the domain has dimensions  $4 \times 4$  mm.

<https://doi.org/10.1371/journal.pcbi.1007635.g005>


An example GUI that uses the UIWindow with embedded UIButtons, InputFields, UILabels, and a UIGrid is shown in Fig 5.

**Utilities.** The Util and Rand classes are used with almost every project. The Util class consists of a collection of standalone functions that solve common problems such as: Generating colors for use with the visualization tools, array manipulation, generating coordinate neighborhoods (e.g. Von Neumann, Moore, Hex, Triangular), spatial mathematical operations, multicore parallelization (typically used to run many models in parallel for parameter sweeps), functions to save and load model states, etc. The Rand class is used for generating random numbers and for sampling distributions (e.g. Gaussian, Poisson, Binomial, Multinomial—created using code adapted from the Colt and NumPy open source libraries [46, 47]) See the manual for more information [38].

**Tools.** A set of miscellaneous tool objects are included to help with specific modeling tasks, these include:


- A FileIO object that is used to read input files and output results. The object supports character and binary formats, and contains shortcuts for reading and writing delimiter separated values.
- A GenomeTracker object that internally stores phylogeny information in a searchable tree structure, and can be used to model branching processes.

- An ODESolver object that can solve ODEs numerically using Euler, Runge-Kutta 4, and Runge-Kutta-Fehlberg 4,5 integration.
- A Multi-Well Experiment object that uses multi-threading to run and display many models simultaneously. The modeler simply creates an array of initialized models, defines an update and draw step, and can then feed many models into the Multi-Well experiment object and observe divergences in dynamics. This allows modelers to intuitively seed different models or replicates of the same conditions and observe differences in their behavior over time, see Fig 6.
- An InteractiveModel object that embeds models in a graphical user interface from which the modeler can schedule modifications to parameters, such as treatment application, and interact with their model in real time. Modelers may also rewind execution to adjust settings, helping them to more quickly understand their model dynamics, and identify useful drug combinations and schedules. The InteractiveModel GUI uses a UIPlot object for the timeline, as well as several UIGrids and UIButtons for other interactive components. This tool was used as part of the development of the Cancer Crusade game [48] to test the effects of therapy on a model by Robertson-Tessi et al. [17], see Fig 7.



**Fig 6. Example of a 10x10 Multi-Well experiment where evolutionary competition of two phenotypes (red,green) shows divergent results with different random seeds.** Models are separated with red lines. With a cell radius of  $10\ \mu\text{m}$ , each well has a domain size of  $2 \times 2\ \text{mm}$ .

<https://doi.org/10.1371/journal.pcbi.1007635.g006>



**Fig 7. An example of the InteractiveModel object, which allows the modeler to experiment with treatment strategies using a model by Robertson-Tessi et al. [17].** (A) A spatial visualization of the current model state. (B) A control panel of UIButtons allows the user to quickly restart the model, pause execution, clear all treatments, and undo previous changes. Hotkeys for these controls are in brackets. (C) A Speed Control option allows the user to easily adjust the execution speed of the model to range from evaluating as fast as possible to taking a second between timesteps, allowing for careful observation of model dynamics. (D) A timeline that will plot timestep information so that the user may observe aggregate changes over time in response to treatment. The user may also click anywhere in the timeline to backtrack to a previous time point and replay the model from there. The timeline will also automatically backtrack to recalculate any necessary frames when a treatment schedule change is made. (E) A set of sliders that allow the user to select different treatment intensities for each drug. (F) Each horizontal bar parallels the simulation timeline and displays the schedule of a different treatment. Modelers can click on regions within these bars to change regions to a new treatment intensity selected in (E). Modelers may also use the hotkeys presented in (E) to apply the selected intensities in real time as the model runs.

<https://doi.org/10.1371/journal.pcbi.1007635.g007>

## Results: Competitive release model

To demonstrate how the aforementioned principles and components of HAL are applied, we consider a simple but complete example of hybrid modeling. We implement the model of

pulsed therapy based on a recent publication by Gallaher et al. [19]. We also showcase the flexibility that the modular component approach brings by displaying three different parameterizations of the same model side by side.

## Competitive release introduction

The model in [19] describes two competing tumor-cell phenotypes: a rapidly dividing, drug-sensitive phenotype and a slower dividing, drug-resistant phenotype. There is also a diffusible drug that enters the system through the domain boundaries and is consumed by the tumor cells over time.

Every timestep (“tick”) each cell has a probability of death and a probability of division. The division probability depends on phenotype (resistant cells divide less frequently) and the availability of space (cells will divide only if there is an open space in the nearest eight grid square neighborhood or Moore neighborhood). Sensitive cells have a death rate that increases when the cells are exposed to drug, while resistant cells have a constant death rate. With a cell radius of 10  $\mu\text{m}$ , the domain size is 100 by 100 cells, or 4  $\text{mm}^2$ . The model runs for 10000 time steps, where each time step represents 2 hours, approximately 2.3 years.

The modular design of HAL allows us to test 3 different treatment conditions, each with an identical starting tumor (no drug, constant drug, and pulsed drug). An interesting outcome of the experiment is that pulsed therapy is better at managing the tumor than constant therapy. Under pulsed therapy the sensitive population is kept in check, while still competing spatially with the resistant phenotype and preventing its expansion. The rest of the section describes in detail how this abstract model is generated.

Fig 8 provides a high level look at the structure of the code. Bold font indicates where a section of the coding example is called. Table 3 provides a quick reference for the built-in HAL functions used in this example. Any functions that are used by the example but do not exist in the table are defined within the example itself and explained in detail below the code. Those fluent in Java may be able to understand the example just by reading the code and using Table 3. Built-in HAL functions and classes are highlighted in bold in the following source code to make identifying HAL’s components easier.

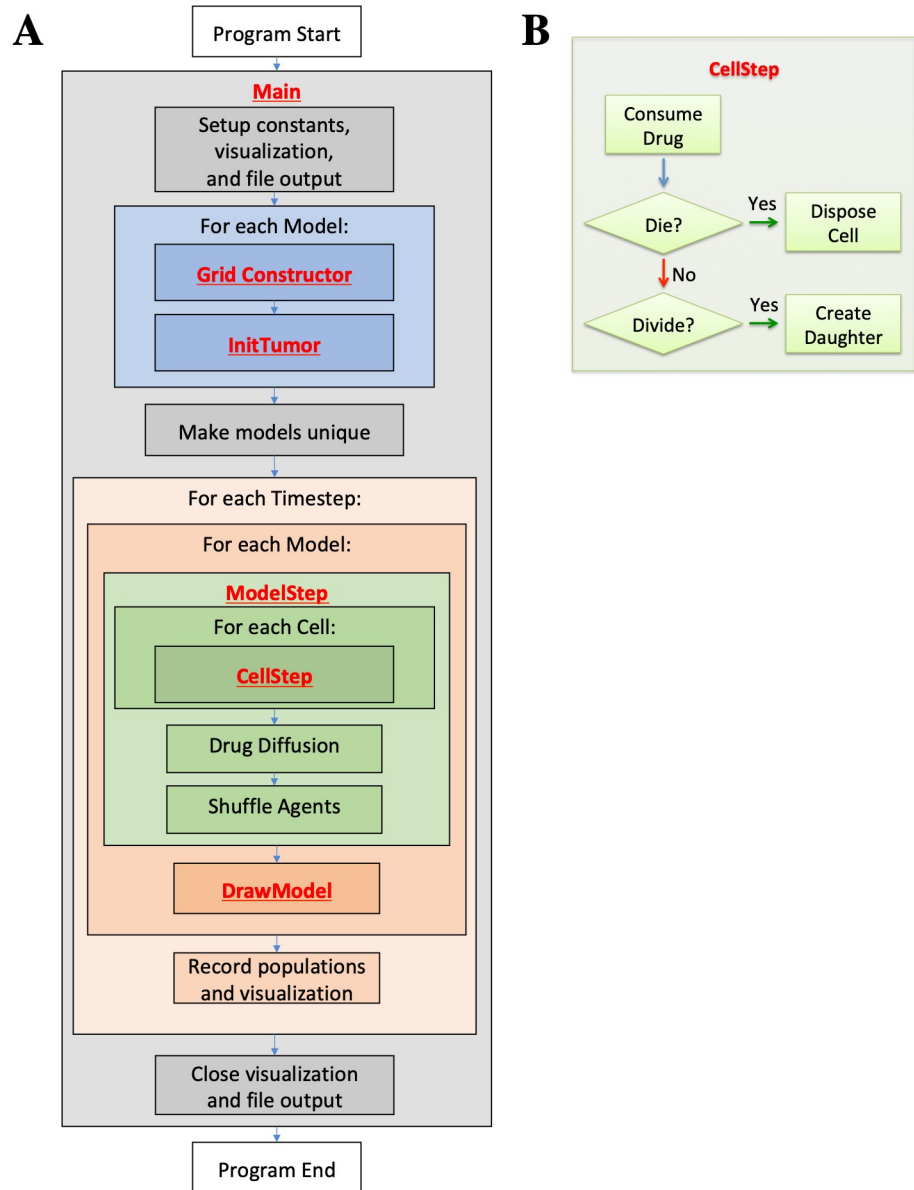
## Main function

We first examine the ‘main’ function for a bird’s-eye view of how the program is structured. Source code elements highlighted in red are built-in HAL functions and objects, and can be referenced in Table 3.

```

1 public static void main(String[] args) {
2     //setting up starting constants and data collection
3     int x = 100, y = 100, visScale = 5, tumorRad = 10, msPause = 5;
4     double resistantProb = 0.5;
5     GridWindow win = new GridWindow("Competitive Release", x * 3, y,
visScale);
6     FileIO popsOut = new FileIO("populations.csv", "w");
7     //setting up models
8     ExampleModel[] models = new ExampleModel[3];
9     for (int i = 0; i < models.length; i++) {
10        models[i] = new ExampleModel(x, y, new Rand(0));
11        models[i].InitTumor(tumorRad, resistantProb);
12    }
13    models[0].DRUG_DURATION = 0;//no drug
14    models[1].DRUG_DURATION = 200;//constant drug
15    //Main run loop
16    for (int tick = 0; tick < 10000; tick++) {

```



**Fig 8.** (A) Example program flow diagram. Bold font indicates where coding example sections are first called. (B) CellStep function flow diagram.

<https://doi.org/10.1371/journal.pcbi.1007635.g008>

```

17 win.TickPause(msPause);
18 for (int i = 0; i < models.length; i++) {
19     models[i].ModelStep(tick);
20     models[i].DrawModel(win, i);
21 }
22 //data recording
23 popsOut.Write(models[0].Pop() + "," + models[1].Pop() + "," +
    models[2].Pop() + "\n");
24 if(tick%(int)(10/models[0].TIMESTEP) == 0) {
25     win.ToPNG("ModelsDay" + (int)(tick * models[0].TIMESTEP) +
        ".png");
    
```

**Table 3. HAL functions used in the example.** Each function is a method of a particular object, meaning that when the function is called it may use properties that pertain to the object that it is called from.

Function	Object	Action
NewAgentSQ(INDEX)	AgentGrid2D	Returns a new agent, placed at the center of the square at the provided position INDEX.
ShuffleAgents(RNG)	AgentGrid2D	Usually called after every timestep to shuffle the order of agent iteration.
GetTick()	AgentGrid2D	Returns the current grid timestep.
ItoX(INDEX), ItoY(INDEX)	AgentGrid2D	Converts from a grid position INDEX into the x and y components that point to the same grid position.
MapHood(NEIGHBORHOOD,X,Y)	AgentGrid2D	Finds all position indices in the provided neighborhood, centered around X,Y that don't fall out of bounds of the AgentGrid2D. Writes these indices into the NEIGHBORHOOD argument, and returns the number that were found.
MapEmptyHood(NEIGHBORHOOD)	AgentSQ2D	Finds all position indices in the provided neighborhood, centered around the agent, that do not have an agent occupying them. Writes these indices into the NEIGHBORHOOD argument, and returns the number that were found.
G	AgentSQ2D	Gets the grid that the agent occupies.
Isq()	AgentSQ2D	Gets the position index of the grid square that the agent occupies.
Dispose()	AgentSQ2D	Removes the agent from the grid and from iteration.
Get(INDEX)	PDEGrid2D	Returns the concentration of the PDE field at the given index.
Mul(INDEX, VALUE)	PDEGrid2D	Multiplies the concentration at the given INDEX by VALUE and adds the result to the current concentration when Update() is called
DiffusionADI(RATE)	PDEGrid2D	Applies diffusion using the ADI method with the rate constant provided. A reflective boundary is assumed. The result is applied when Update() is called.
DiffusionADI(RATE, BOUNDARY_COND)	PDEGrid2D	Applies diffusion using the ADI method with the RATE constant provided. The BOUNDARY_COND value diffuses from the grid borders. The result is applied when Update() is called.
Update()	PDEGrid2D	Applies all state changes simultaneously to the PDEGrid
SetPix(INDEX, COLOR)	GridWindow	Sets the color of a pixel.
TickPause(MILLISECONDS)	GridWindow	Pauses the program between calls to TickPause. The function automatically subtracts the time between calls from MILLISECONDS to ensure a consistent timestep rate for visualization.
ToPNG(FILENAME)	GridWindow	Writes out the current state of the UIWindow to a PNG image file.
Close()	GridWindow	Closes the GridWindow.
ProbScale(PROB, DURATION)	Util	Scales the probability that an event will occur in 1 unit of time to the probability that the event will occur at least once in the DURATION.
RGB(RED, GREEN, BLUE)	Util	Returns an integer with the requested color in RGB format. This value can be used for visualization.
HeatMapRGB(VALUE)	Util	Maps the VALUE argument (assumed to be between 0 and 1) to a color in the heat colormap.
CircleHood(INCLUDE_ORIGIN, RADIUS)	Util	Returns a set of coordinate pairs that define the neighborhood of all squares whose centers are within the RADIUS distance of the center (0, 0) origin square. The INCLUDE_ORIGIN argument specifies whether to include the origin in this set of coordinates.
MooreHood(INCLUDE_ORIGIN)	Util	Returns a set of coordinate pairs that define a Moore neighborhood around the (0, 0) origin square. The INCLUDE_ORIGIN Boolean specifies whether we intend to include the origin in this set of coordinates.
Write(String)	FileIO	Writes the STRING to the output file.
Close()	FileIO	Closes the output file.
Double()	Rand	Generates a random double value in the range [0 - 1)

<https://doi.org/10.1371/journal.pcbi.1007635.t003>

```

26     }
27   }
28   //closing data collection
29   popsOut.Close();
30   win.Close();
31 }

```

Lines 3-4. Defines all of the constants that will be needed to setup the model and display.

- Creates a GridWindow of RGB pixels for visualization and for generating timestep PNG images.  $x*3$ ,  $y$  define the dimensions of the pixel grid. the  $x$  variable is multiplied

by 3 so that 3 models can be visualized side by side in the same window. The last argument is a scaling factor that specifies that each pixel on the grid will be viewed as a 5x5 square of pixels on the screen.

6. Creates a file output object that will write to a file called populations.csv.
8. Creates an array with 3 entries that will be populated with models.
- 9-12. Fills the model list with models that are initialized identically, with identical random number generators. Each model will hold and update its own cells and diffusible drug. See the Grid Definition and Constructor section and the InitTumor Function section for more details.
- 13-14. Setting the DRUG\_DURATION constant creates the only difference in the 3 models being compared. In models[0] no drug will be administered. In models[1] drug administration will be constant (DRUG\_DURATION is set equal to DRUG\_CYCLE). In models[2] drug will be administered periodically (the default value of DRUG\_DURATION is 40). See the ExampleModel Constructor and Properties section for the default model initialization.
16. Executes the main loop for 10000 timesteps.
17. Requires every iteration of the loop to take a minimum number of milliseconds. This slows down the execution and display of the model and makes it easier for the viewer to follow.
18. Loops over all models to update them.
19. Advances the state of the agents and diffusibles in each model by one timestep. See the Model Step Function for more details.
20. Draws the current state of each model to the window. See the Draw Model Function for more details.
23. Writes the population sizes of each model every timestep to allow the models to be compared.
- 24-25. Every 10 days, writes the state of the model as captured by the GridWindow to a PNG file.
- 29-30. After the main for loop has finished, the FileIO object and the visualization window are closed, and the program ends.

### ExampleModel constructor and properties

This section explains how the grid is defined and instantiated.

```

1 public class ExampleModel extends AgentGrid2D<ExampleCell> {
2   //model constants
3   public final static int RESISTANT = RGB (0, 1, 0), SENSITIVE = RGB
      (0, 0, 1);
4   public double TIMESTEP = 2.0/24; //2 hours per timestep
5   public double SPACE_STEP = 20; //um
6   public double DIV_PROB_SEN = ProbScale (0.5, TIMESTEP);
7   public double DIV_PROB_RES = ProbScale (0.2, TIMESTEP);
8   public double DEATH_PROB = ProbScale (0.02, TIMESTEP);
9   public double DRUG_DEATH = ProbScale (0.8, TIMESTEP);

```



```

10 public double DRUG_START = 20/TIMESTEP;
11 public double DRUG_PERIOD = 15/TIMESTEP;
12 public double DRUG_DURATION = 2/TIMESTEP;
13 public double DRUG_UPTAKE = -0.03 *TIMESTEP;
14 public double DRUG_DIFF_RATE = 0.02*60*60*24*(TIMESTEP/
    (SPACE_STEP*SPACE_STEP));
15 public double DRUG_BOUNDARY_VAL = 1.0;
16 //internal model objects
17 public PDEGrid2D drug;
18 public Rand rng;
19 public int [] divHood = MooreHood (false);
20
21 public ExampleModel (int xDim, int yDim, Rand rng) {
22     super (xDim, yDim, ExampleCell.class);
23     this.rng = rng;
24     drug = new PDEGrid2D (xDim, yDim);
25 }

```

1. The ExampleModel class, which is user defined and specific to this example, is built by extending the generic AgentGrid2D class. The extended grid class requires an agent type parameter, which specifies the type of agent that will live on the grid. To meet this requirement, the <ExampleCell> type parameter is added to the declaration.
3. Defines RESISTANT and SENSITIVE constants, which are created by the Util RGB function. These constants serve as both colors for drawing and as labels for the different cell types.
- 4-5. All grids in HAL assume unit spatial dimensions, and operations such as diffusion assume unit timesteps. This means that all rate/probability constants should be scaled such that when applied, they operate at a consistent resolution chosen for the grids.
- 6-9. Probabilistic constants are scaled from their original probabilities (per day) into probabilities per 2 hours, such that the expected value per day is the same.
- 10-12. Defines treatment application constants. These values can be reassigned after model creation to test different treatment schedules. In the main function, the DRUG\_DURATION variable is modified for the Constant-Drug, and Pulsed Therapy experiment cases.
- 13-15. Defines the PDE constants, the DRUG\_DIFF\_RATE is scaled from  $\text{um}^2/\text{second}$  to  $\text{um}^2/\text{day}$ , and is then scaled by the timestep and space-step.
17. Declares that the model will contain a PDEGrid2D, which will hold the drug concentrations. The PDEGrid2D can only be initialized when the x and y dimensions of the model are known, which is why we do not create the PDEGrid2D object until the constructor function is called.
18. Declares that the Grid will contain a Random number generator (the Rand object), but takes it in as a constructor argument to allow the modeler to seed the generator if desired for consistent output.
19. Creates a neighborhood using the MooreHood function. The MooreHood function generates a set of coordinates that define the Moore Neighborhood (the 8 closest coordinates to a central origin), centered around (0, 0). The false argument declares that we do not want to include the origin in the neighborhood, just the 8 coordinates around that position. The neighborhood is stored in the format  $[0_10_2, \dots, 0_m, x_1, y_1, x_2,$

$y_2, \dots, x_n, y_n]$ . The leading zeros are written to when MapHood is called, and will store the position indices that the neighborhood maps to. See the CellStep function for more information, and the InitTumor Function Line 3 for another example of the use of neighborhoods

21. Defines the model constructor, which takes as arguments the x and y dimensions of the model and a random number generator (a Rand object).
22. Calls the AgentGrid2D constructor with super, passing it the x and y dimensions of the model, and the ExampleCell Class. This Class is used by the Grid to generate a new cell when the NewAgentSQ function is called.
- 23-24. The Rand argument is assigned and the drug PDEGrid2D is defined with matching dimensions.

### InitTumor function

The next segment of code is a function from the ExampleModel class that defines how the tumor is first seeded after the ExampleModel is created.

```

1 public void InitTumor(double radius, double resistantProb) {
2   //get a list of indices that fill a circle at the center of the
   grid
3   int[] tumorNeighborhood = CircleHood(true, radius);
4   int hoodSize = MapHood(tumorNeighborhood, xDim / 2, yDim / 2);
5   for (int i = 0; i < hoodSize; i++) {
6     if (rng.Double () < resistantProb) {
7       NewAgentSQ(tumorNeighborhood[i]).type = RESISTANT;
8     } else {
9       NewAgentSQ(tumorNeighborhood[i]).type = SENSITIVE;
10    }
11  }
12 }
```

1. The arguments passed to the InitTumor function are the approximate radius of the circular tumor being created and the probability that each created cell will be of the resistant phenotype.
3. Sets the tumorNeighborhood array using the CircleHood function, which stores coordinates in the form  $[0_1, 0_2, \dots, 0_n, x_1, y_1, x_2, y_2, \dots, x_n, y_n]$ . The x,y coordinate pairs define a neighborhood of all squares whose centers are within the radius distance of the center (0, 0) origin square. The leading 0s are used by the MapHood function to store the mapped indices. The Boolean argument specifies that the origin will be included in this set of squares, thus making a completely filled circle of squares.
4. Uses the MapHood function to map the neighborhood defined above to be centered around xDim/2,yDim/2 (the dimensions of the AgentGrid). The results of the mapping are written as position indices to the beginning of the tumorNeighborhood array. MapHood returns the number of valid indices found, and this will be the size of the starting population.
5. Loops from 0 to hoodSize, allowing access to each mapped position index in the tumorNeighborhood.

6. Samples a random number in the range  $[0 - 1)$  and compares to the `resistantProb` argument to set whether the cell should have the resistant phenotype or the sensitive phenotype.
- 7-9. Uses the `NewAgentSQ` function to place a new cell at each `tumorNeighborhood` position. In the same line we also specify that the phenotype should be either resistant or sensitive, depending on the result of the `rng.Double()` call.

## ModelStep function

This section looks at the model's step function which is executed once per timestep by each Model.

```

1 public void ModelStep(int tick) {
2   ShuffleAgents(rng);
3   for (ExampleCell cell: this) {
4     cell.CellStep();
5   }
6   double periodTick = (tick - DRUG_START) % DRUG_PERIOD;
7   if (periodTick > 0 && periodTick < DRUG_DURATION) {
8     //drug will enter through boundaries
9     drug.DiffusionADI(DRUG_DIFF_RATE, DRUG_BOUNDARY_VAL);
10  } else {
11    //drug will not enter through boundaries
12    drug.DiffusionADI(DRUG_DIFF_RATE);
13  }
14  drug.Update();
15 }
```

2. The `ShuffleAgents` function randomizes the order of iteration so that the agents are always looped through in random order.
- 3-4. Iterates over every cell on the grid, and calls the `CellStep` function on every cell.
- 6-7. The `periodTick` variable stores at what point in the drug delivery cycle the tick is, and the `If` statement checks whether the tick is in the right part of the drug cycle to apply drug. (See the `Grid Definition and Constructor` section for the values of the constants involved, the `DRUG_DURATION` variable is set differently for each model in the `Main Function`)
9. If it is time to add drug to the model, the `DiffusionADI` function is called. `DiffusionADI` uses the `ADI` method which is more stable than 2D Euler and allows us to take larger steps. The additional argument to the `DiffusionADI` function specifies the boundary condition value `DRUG_BOUNDARY_VAL`. This causes the drug to diffuse into the `PDEGrid2D` from the boundary. Here we assume that drug is only delivered from the boundaries of the domain
12. Without the second argument the `DiffusionADI` function assumes a zero-flux boundary, meaning that drug concentration cannot escape or enter through the sides of the model. Therefore the only way for the drug concentration to decrease is via uptake by the Cells. See the `CellStep` function section, line 6, for more information.
14. `Update` is called to apply the reaction and diffusion changes to the `PDEGrid`.

## CellStep function and cell properties

We next look at how the ExampleCell Agent is defined and at the CellStep function that runs once per Cell per timestep. The G property that is referenced many times in this section is a built-in agent property that gives access to the ExampleGrid object that the cell lives on.

```

1 class ExampleCell extends AgentSQ2Dunstackable<ExampleModel> {
2   public int type;
3
4   public void CellStep() {
5     //uptake of Drug
6     G.drug.Mul(Isq(), G.DRUG_UPTAKE);
7     double deathProb, divProb;
8     //Chance of Death, depends on resistance and drug concentration
9     if (this.type == RESISTANT) {
10      deathProb = G.DEATH_PROB;
11    } else {
12      deathProb = G.DEATH_PROB + G.drug.Get(Isq()) * G.DRUG_DEATH;
13    }
14    if (G.rng.Double() < deathProb) {
15      Dispose();
16      return;
17    }
18    //Chance of Division, depends on resistance
19    if (this.type == RESISTANT) {
20      divProb = G.DIV_PROB_RES;
21    } else {
22      divProb = G.DIV_PROB_SEN;
23    }
24    if (G.rng.Double() < divProb) {
25      int options = MapEmptyHood(G.divHood);
26      if (options > 0) {
27        G.NewAgentSQ(G.divHood[G.rng.Int(options)].type = this.
28          type;
29      }
30    }
31 }

```

1. The ExampleCell class is built by extending the generic AgentSQ2Dunstackable class. The extended Agent class requires the ExampleModel class as a type argument, which is the type of Grid that the Agent will live on. To meet this requirement, we add the <ExampleModel> type parameter to the extension.
2. Defines a cell property called “type”. Each Cell holds a value for this field. If the value is RESISTANT, the Cell is of the resistant phenotype, if the value is SENSITIVE, the cell is of the sensitive phenotype. The RESISTANT and SENSITIVE values are defined in the ExampleGrid as constants (See the ExampleModel Constructor and Properties, line 3).
6. The G property is used to access the ExampleGrid object that the Cell lives on. G is used often with agent functions as the AgentGrid is expected to contain any information that is not local to the individual agent. Here it is used to get the drug PDE-Grid2D. The drug concentration at the index that the Cell is currently occupying (Isq ()) is then multiplied by the drug uptake constant, thus modeling local drug uptake by the Cell.

7. Defines deathProb and divProb variables, these will be assigned different values depending on whether the ExampleCell is RESISTANT or SENSITIVE.
- 9-12. If the cell is resistant, the deathProb variable is set to the DEATH\_PROB value alone, if the cell is sensitive, an additional term is added to account for the probability of the cell dying from drug exposure, using the concentration of drug at the cell's position (Isq())
- 14-16. Samples a random number in the range [0 – 1) and compares to deathProb to determine whether the cell will die. If so, the built-in agent Dispose() function is called, which removes the agent from the grid, and then return is called so that the dead cell will not divide.
- 19-22. Sets the divProb variable to either DIV\_PROB\_RES for resistant cells, or DIV\_PROB\_SEN for sensitive cells.
24. Samples a random number in the range [0 – 1) and compares to divProb to determine whether the cell will divide.
25. If the cell divides, the MapEmptyHood function is used, which checks the positions in the divHood (the Moore neighborhood as defined in the ExampleModel Constructor and Properties section, line 11) around the Cell, and writes the position indices that do not contain any agents into the divHood. MapEmptyHood returns the number of valid empty positions found.
- 26-27. If there are one or more valid division options, a new daughter cell is created using the NewAgentSQ function and its starting location is chosen by randomly sampling the divHood array to pull out one of its valid locations. The other daughter is assumed to occupy the same location as the mother cell. Finally with the.type = this.type statement, the phenotype of the newly placed daughter cell is inherited from the mother cell.

## DrawModel function

We next look at the DrawModel Function, which is used to display a summary of the model state on a GridWindow object. In this program, DrawModel is called once for each model per timestep; see the main function section for more information.

```

1 public void DrawModel(GridWindow vis, int iModel) {
2   for (int x = 0; x < xDim; x++) {
3     for (int y = 0; y < yDim; y++) {
4       ExampleCell drawMe = GetAgent(x, y);
5       if (drawMe != null) {
6         vis.SetPix(x + iModel * xDim, y, drawMe.type);
7       } else {
8         vis.SetPix(x + iModel * xDim, y, HeatMapRGB(drug.Get(x, y)));
9       }
10    }
11  }
12 }

```

- 2-3. Loops over every lattice position of the grid being drawn, xDim and yDim refer to the dimensions of the model.
4. Uses the GetAgent function to get the Cell that is at the x,y position.

- 5-6. If a cell exists at the requested position, the corresponding pixel on the GridWindow is set to the cell's phenotype color. To draw the models side by side, the pixel being drawn is displaced to the right by the model index.
- 7-8. If there is no cell to draw, then the pixel color is set based on the drug concentration at the same index, using the built-in heat colormap.

## Imports

The final code snippet looks at the imports that are needed. Any modern Java IDE should generate import statements automatically.

```
1 package Examples._6CompetitiveRelease;
2 import HAL.GridsAndAgents.AgentGrid2D;
3 import HAL.GridsAndAgents.PDEGrid2D;
4 import HAL.Gui.GridWindow;
5 import HAL.GridsAndAgents.AgentSQ2Dunstackable;
6 import HAL.Tools.FileIO;
7 import HAL.Rand;
8 import static Examples._6CompetitiveRelease.ExampleModel.*;
9 import static HAL.Util.*;
```

1. The package statement specifies where the file exists in the larger project structure
- 2-7. Imports all of the classes that we will need for the program.
8. Imports the static fields of the model so that we can use the type names defined there in the Agent class.
9. Imports the static functions of the Util file, which adds all of the Util functions to the current namespace, so we can natively call them. Statically importing Util is recommended for every project.

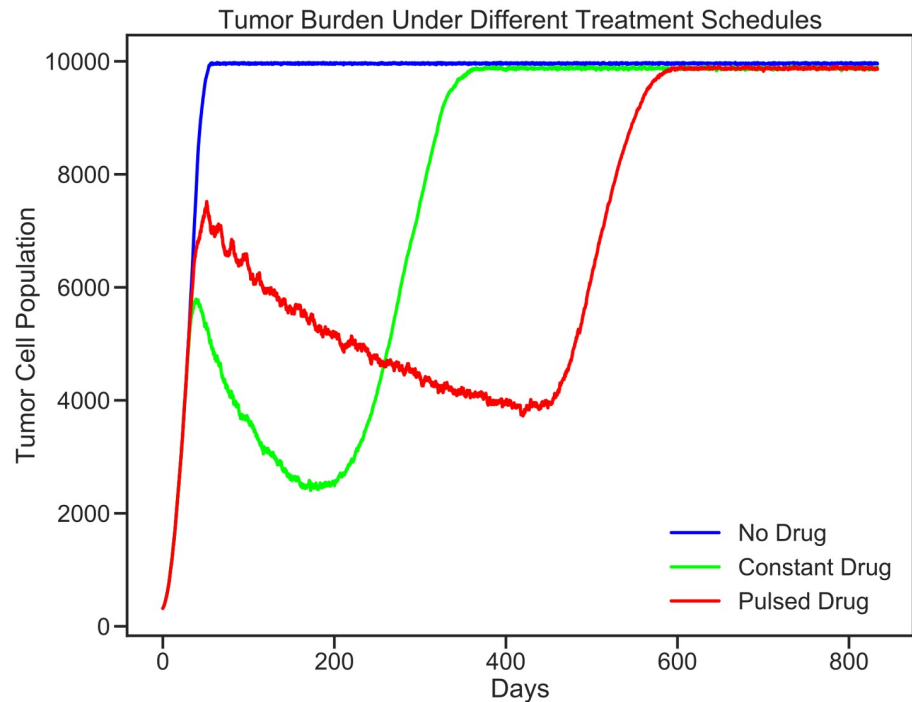
## Model results

[Fig 10](#) displays the model visualization at Day 0, Day 20, Day 200, Day 500, and Day 650 recorded from the GridWindow ToPNG function. [Fig 9](#) displays the population sizes as recorded by the FileIO Write function at the end of every timestep.

This example illustrates the power of HAL's approach to model building. Writing relatively little complex code, we setup a three model experiment with nontrivial dynamics along with methods to collect data and visualize the models. We now briefly review the model results.

As can be seen in [Fig 10](#), at Day 0 and Day 20 (right before drug application starts), all 3 models are identical. At Day 200 the differences in treatment application show different effects: when no drug is applied, the rapidly dividing sensitive cells quickly fill the domain, while when drug is applied constantly, the resistant cells start to overtake the sensitive population. Pulsed drug kills some sensitive cells, but leaves enough alive to prevent growth of the resistant cells. At Day 500, the resistant cells have begun to emerge from the center of the pulsed drug model. At Day 650, all domains are filled.

As can be seen in [Fig 9](#), the pulsed therapy is the most effective at preventing tumor growth long-term, however the resistant cells ultimately succeed in breaking out of the tumor center and out-competing the sensitive cells on the fringes of the tumor. It may be possible to contain a population of sensitive and resistant cells for longer by using a different pulsing schedule or



**Fig 9. FileIO population output.** This plot summarizes the changes in tumor burden over time for each model. This plot was constructed using data accumulated in the program output populations.csv file. Displayed using the Java XChart package.

<https://doi.org/10.1371/journal.pcbi.1007635.g009>

by modifying the treatment schedule in response to the tumor growth (adaptive therapy). As the presented model is primarily an example, we do not explore how to improve treatment further. For a more detailed exploration of the potential of adaptive therapy for prolonging competitive release, see [19].


## Availability and future directions

### How to download and contribute

HAL is publicly available on GitHub, at <https://github.com/MathOnco/HAL>. A manual is included that walks the user through installation and serves as a coding reference. For those interested in using HAL, downloading and setting up the project is a good first start. From there running and examining the included examples is recommended, as they do a good job of summarizing HAL's capabilities. Modelers can contribute tools that they develop by making pull requests to the repository.

### Future directions

**Cross model validation.** Having many different paradigms to choose from adds several complications to modeling: It can take significant effort to build a model from scratch under one paradigm, and then significant additional effort to migrate the model to a different paradigm. By adding more modeling approaches with a consistent interface, HAL will lower the model migration barrier and allow modelers to test the merits of many paradigms in their investigation, and to validate their results by seeing whether they hold true across paradigms [49]. Note that our goal is not to recreate all of the functionality of the pre-existing frameworks



**Fig 10. Selected model visualizations.** Blue cells are drug sensitive, Green cells are drug resistant, background heatmap colors show drug concentration. The model domain size is 100 by 100 cells, or 4 mm<sup>2</sup>

<https://doi.org/10.1371/journal.pcbi.1007635.g010>

that support these paradigms, it is to provide their core algorithms in an accessible and consistent format so that users can easily choose from and compare them.

To take HAL's modeling flexibility further we hope to incorporate additional modeling paradigms that are commonly used in agent-based modeling of cells. A potential addition is a



Voronoi Agent type, which would use Voronoi tessellation [50] to find the cell's nearest neighbors and determine cell volume. We are also considering including modeling paradigms that construct cells out of smaller subunits, such as Deformable Ellipsoid Cell Modeling [51], as this would allow us to model the mechanics of tissue formation and cell migration in more detail.

**Bridging spatial scales.** We also hope to explore the possibility of adjusting spatial scales for both our PDEs and Agents. For PDEs, this is a readily understood problem, and we intend to add scalable PDEGrids to HAL soon. However, for agent-based modeling the process of changing scales while preserving dynamics is not so well defined, though we imagine that it may be possible under certain assumptions. This would be useful for helping us bridge the divide between cell level and tissue/organ/tumor level dynamics, as the number of cells involved at these scales are orders of magnitude greater than what desktop machines can tractably model. The PopulationGrid modeling system is an example of this.

**Assumption modules.** A common modeling task is exploring how combinations of different assumptions influence model behavior. The included ModuleSetManager object helps design models specifically with this in mind. The design entails providing code “hooks” so that code can be added to influence specific agent decisions and model events, (e.g. whether an agent will reproduce). Modelers can then write assumption modules that will influence these events (e.g. by altering the probability of reproduction based on an environmental factor that would otherwise be ignored).

This approach allows modelers to combine and remove assumption modules without having to worry about breaking the model. This facilitates easy exploration of the space of assumptions until ones suitable for understanding biological phenomena are found. We are very excited about the potential of this approach for collaborative projects and for building increasingly complex models by encapsulating the complexity into manageable parts, and hope to improve on the tools for this paradigm as we explore its potential.

**Advanced scheduling.** Taking inspiration from Repast, SWARM, and MASON, another expected extension is the inclusion of optional schedulers to facilitate more complex methods of iterating through agents than simply looping over each grid. This is not intended to replace the simple grid iteration approach, but instead should augment it with optional complex methods. An AgentList object is currently included to begin to address this. It allows modelers to make selective lists of agents for more flexible iteration.

**Building a community.** HAL has already seen adoption within the labs at the Integrated Mathematical Oncology department of Moffitt Cancer Center and beyond. We certainly hope that more outside users will be interested in its potential. As the user-base for HAL grows, we plan to extend the base of resources around the platform. The current set of resources that exist for new users to get started are the manual [38], a website with an online version of the manual [52] and a playlist of YouTube videos [53]. We intend to increase HAL's online presence by including a website with a code repository to make sharing models and tools easier.

## Conclusion

Cancer is a complex and heterogeneous disease whose mathematical study is still being developed. To make better progress in this endeavor, it is helpful to have a set of highly generic tools that encapsulate the key components of spatial modeling so that researchers can produce efficient models quickly without being constrained in their approach, nor in the complexity of the systems that they can produce. HAL is our attempt to achieve this.

HAL was made easily extensible so that researchers can build models and more specific tools on top of HAL's generic base. The hope is that by this process HAL will grow into a

powerful tool set that will help standardize and coordinate hybrid modeling in mathematical oncology.

We recommend HAL to anyone building spatial models for oncology, as the tools presented are primarily geared toward this end. However, given the amount of overlap and cross talk between the approaches used in different modeling applications, we hope that modelers outside of mathematical oncology will also take interest and contribute, to our mutual benefit.

## Supporting information

**S1 Text. Supplementary materials.**  
(PDF)

## Author Contributions

**Conceptualization:** Rafael R. Bravo, Alexander R. A. Anderson.

**Funding acquisition:** Mark Robertson-Tessi, Alexander R. A. Anderson.

**Methodology:** Rafael R. Bravo, Etienne Baratchart, Anna K. Miller, Jill Gallaher, Chandler D. Gatenbee, Alexander R. A. Anderson.

**Software:** Rafael R. Bravo, Etienne Baratchart, Chandler D. Gatenbee.

**Supervision:** Alexander R. A. Anderson.

**Validation:** Rafael R. Bravo, Etienne Baratchart, Ryan O. Schenck, Anna K. Miller.

**Visualization:** Rafael R. Bravo, Jeffrey West, Ryan O. Schenck, David Basanta.

**Writing – original draft:** Rafael R. Bravo, Mark Robertson-Tessi, Alexander R. A. Anderson.

**Writing – review & editing:** Rafael R. Bravo, Etienne Baratchart, Jeffrey West, Anna K. Miller, Jill Gallaher, Mark Robertson-Tessi, Alexander R. A. Anderson.

## References

1. Anderson AR. A hybrid mathematical model of solid tumour invasion: the importance of cell adhesion. *Mathematical medicine and biology: a journal of the IMA*. 2005; 22(2):163–186. <https://doi.org/10.1093/imammb/dqi005> PMID: 15781426
2. Rejniak KA, Anderson AR. Hybrid models of tumor growth. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine*. 2011; 3(1):115–125. <https://doi.org/10.1002/wsbm.102> PMID: 21064037
3. Anderson AR, Chaplain M. Continuous and discrete mathematical models of tumor-induced angiogenesis. *Bulletin of mathematical biology*. 1998; 60(5):857–899. <https://doi.org/10.1006/bulm.1998.0042> PMID: 9739618
4. Gerlee P, Anderson AR. An evolutionary hybrid cellular automaton model of solid tumour growth. *Journal of theoretical biology*. 2007; 246(4):583–603. <https://doi.org/10.1016/j.jtbi.2007.01.027> PMID: 17374383
5. Norton KA, Gong C, Jamalian S, Popel AS. Multiscale agent-based and hybrid modeling of the tumor immune microenvironment. *Processes*. 2019; 7(1):37. <https://doi.org/10.3390/pr7010037> PMID: 30701168
6. Dormann S, Deutsch A. Modeling of self-organized avascular tumor growth with a hybrid cellular automaton. *In silico biology*. 2002; 2(3):393–406. PMID: 12542422
7. Kim Y, Stolarska MA, Othmer HG. A hybrid model for tumor spheroid growth in vitro I: theoretical development and early results. *Mathematical Models and Methods in Applied Sciences*. 2007; 17(supp01):1773–1798. <https://doi.org/10.1142/S0218202507002479>
8. Sanga S, Frieboes HB, Zheng X, Gatenby R, Bearer EL, Cristini V. Predictive oncology: a review of multidisciplinary, multiscale in silico modeling linking phenotype, morphology and growth. *Neuroimage*. 2007; 37:S120–S134. <https://doi.org/10.1016/j.neuroimage.2007.05.043> PMID: 17629503

9. Basanta D, Anderson A. Homeostasis Back and Forth: An Ecoevolutionary Perspective of Cancer. *Cold Spring Harbor perspectives in medicine*. 2017; 7(9). <https://doi.org/10.1101/cshperspect.a028332> PMID: 28289244
10. Basanta D, Strand DW, Lukner RB, Franco OE, Cliffl DE, Ayala GE, et al. The Role of Transforming Growth Factor- $\beta$ -Mediated Tumor-Stroma Interactions in Prostate Cancer Progression: An Integrative Approach. *Cancer research*. 2009; 69(17):7111–7120. <https://doi.org/10.1158/0008-5472.CAN-08-3957> PMID: 19706777
11. Kim E, Rebecca V, Fedorenko IV, Messina JL, Mathew R, Maria-Engler SS, et al. Senescent fibroblasts in melanoma initiation and progression: an integrated theoretical, experimental, and clinical approach. *Cancer research*. 2013;. <https://doi.org/10.1158/0008-5472.CAN-13-1720>
12. Anderson AR, Weaver AM, Cummings PT, Quaranta V. Tumor morphology and phenotypic evolution driven by selective pressure from the microenvironment. *Cell*. 2006; 127(5):905–915. <https://doi.org/10.1016/j.cell.2006.09.042> PMID: 17129778
13. Araujo A, Cook LM, Lynch CC, Basanta D. An integrated computational model of the bone microenvironment in bone-metastatic prostate cancer. *Cancer research*. 2014; 74(9):2391–2401. <https://doi.org/10.1158/0008-5472.CAN-13-2652> PMID: 24788098
14. Scianna M, Bell C, Preziosi L. A review of mathematical models for the formation of vascular networks. *Journal of theoretical biology*. 2013; 333:174–209. <https://doi.org/10.1016/j.jtbi.2013.04.037> PMID: 23684907
15. Metzcar J, Wang Y, Heiland R, Macklin P. A review of cell-based computational modeling in cancer biology. *JCO clinical cancer informatics*. 2019; 2:1–13. <https://doi.org/10.1200/CCI.18.00069>
16. Chamseddine IM, Rejniak KA. Hybrid modeling frameworks of tumor development and treatment. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine*. 2019; p. e1461. <https://doi.org/10.1002/wsbm.1461> PMID: 31313504
17. Robertson-Tessi M, Gillies RJ, Gatenby RA, Anderson AR. Impact of metabolic heterogeneity on tumor growth, invasion, and treatment outcomes. *Cancer research*. 2015; 75(8):1567–1579. <https://doi.org/10.1158/0008-5472.CAN-14-1428> PMID: 25878146
18. Zhang J, Cunningham JJ, Brown JS, Gatenby RA. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature communications*. 2017; 8(1):1816. <https://doi.org/10.1038/s41467-017-01968-5> PMID: 29180633
19. Gallaher JA, Enriquez-Navas PM, Luddy KA, Gatenby RA, Anderson AR. Adaptive vs continuous cancer therapy: Exploiting space and trade-offs in drug scheduling. *bioRxiv*. 2017;.
20. Karolak A, Rejniak KA. Micropharmacology: an in silico approach for assessing drug efficacy within a tumor tissue. *Bulletin of mathematical biology*. 2018; p. 1–19.
21. Alfonso J, Talkenberger K, Seifert M, Klink B, Hawkins-Daarud A, Swanson K, et al. The biology and mathematical modelling of glioma invasion: a review. *Journal of the Royal Society Interface*. 2017; 14(136):20170490. <https://doi.org/10.1098/rsif.2017.0490>
22. Ghaffarizadeh A, Heiland R, Friedman SH, Mumenthaler SM, Macklin P. PhysiCell: an open source physics-based cell simulator for 3-D multicellular systems. *PLoS Comput Biol*. 2018; 14(2):e1005991. <https://doi.org/10.1371/journal.pcbi.1005991> PMID: 29474446
23. Hoehme S, Drasdo D. A cell-based simulation software for multi-cellular systems. *Bioinformatics*. 2010; 26(20):2641–2642. <https://doi.org/10.1093/bioinformatics/btq437> PMID: 20709692
24. Kang S, Kahan S, McDermott J, Flann N, Shmulevich I. Biocellion: accelerating computer simulation of multicellular biological system models. *Bioinformatics*. 2014; 30(21):3101–3108. <https://doi.org/10.1093/bioinformatics/btu498> PMID: 25064572
25. Cytowski M, Szymańska Z, Umiński P, Andrejczuk G, Raszkowski K. Implementation of an agent-based parallel tissue modelling framework for the Intel MIC architecture. *Scientific Programming*. 2017; 2017. <https://doi.org/10.1155/2017/8721612>
26. Germann P, Marin-Riera M, Sharpe J. ya|| a: GPU-powered Spheroid Models for Mesenchyme and Epithelium. *Cell systems*. 2019;. <https://doi.org/10.1016/j.cels.2019.02.007> PMID: 30904379
27. Sütterlin T, Kolb C, Dickhaus H, Jäger D, Grabe N. Bridging the scales: semantic integration of quantitative SBML in graphical multi-cellular models and simulations with EPISIM and COPASI. *Bioinformatics*. 2012; 29(2):223–229. <https://doi.org/10.1093/bioinformatics/bts659> PMID: 23162085
28. Swat MH, Thomas GL, Belmonte JM, Shirinifard A, Hmeljak D, Glazier JA. Multi-scale modeling of tissues using CompuCell3D. *Methods in cell biology*. 2012; 110:325. <https://doi.org/10.1016/B978-0-12-388403-9.00013-8> PMID: 22482955
29. Starruß J, de Back W, Bruschi L, Deutsch A. Morpheus: a user-friendly modeling environment for multi-scale and multicellular systems biology. *Bioinformatics*. 2014; 30(9):1331–1332. <https://doi.org/10.1093/bioinformatics/btt772> PMID: 24443380

30. Graner F, Glazier JA. Simulation of biological cell sorting using a two-dimensional extended Potts model. *Physical review letters*. 1992; 69(13):2013. <https://doi.org/10.1103/PhysRevLett.69.2013> PMID: 10046374
31. Mirams GR, Arthurs CJ, Bernabeu MO, Bordas R, Cooper J, Corrias A, et al. Chaste: an open source C++ library for computational physiology and biology. *PLoS computational biology*. 2013; 9(3):e1002970. <https://doi.org/10.1371/journal.pcbi.1002970> PMID: 23516352
32. Tisue S, Wilensky U. Netlogo: A simple environment for modeling complexity. In: *International conference on complex systems*. vol. 21. Boston, MA; 2004. p. 16–21. Available from: <https://ccl.northwestern.edu/papers/netlogo-iccs2004.pdf>.
33. Collier N. Repast: An extensible framework for agent simulation. *The University of Chicagos Social Science Research*. 2003; 36:2003.
34. Luke S, Cioffi-Revilla C, Panait L, Sullivan K, Mason A. A new multi-agent simulation toolkit. In: *Proceedings of the 2004 swarmfest workshop*. vol. 8. Department of Computer Science and Center for Social Complexity, George Mason University Fairfax, VA; 2004. p. 316–327. Available from: <http://cobweb.cs.uga.edu/~maria/pads/papers/mason-SwarmFest04.pdf>.
35. Lampoudi S, Gillespie DT, Petzold LR. The multinomial simulation algorithm for discrete stochastic simulation of reaction-diffusion systems. *The Journal of chemical physics*. 2009; 130(9):094104. <https://doi.org/10.1063/1.3074302> PMID: 19275393
36. Hucka M, Finney A, Sauro HM, Bolouri H, Doyle JC, Kitano H, et al. The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models. *Bioinformatics*. 2003; 19(4):524–531. <https://doi.org/10.1093/bioinformatics/btg015> PMID: 12611808
37. Anderson A, Sleeman B, Young I, Griffiths B. Nematode movement along a chemical gradient in a structurally heterogeneous environment: 2. Theory. *Fundamental and applied nematology*. 1997; 20(2):165–172.
38. Bravo R. HAL Manual; 2018. Available from: <https://github.com/MathOnco/HAL/blob/master/manual.pdf>.
39. Peaceman DW, Rachford HH Jr. The numerical solution of parabolic and elliptic differential equations. *Journal of the Society for industrial and Applied Mathematics*. 1955; 3(1):28–41. <https://doi.org/10.1137/0103003>
40. Courant R, Isaacson E, Rees M. On the solution of nonlinear hyperbolic differential equations by finite differences. *Communications on Pure and Applied Mathematics*. 1952; 5(3):243–255. <https://doi.org/10.1002/cpa.3160050303>
41. Swanson KR, Alvord EC Jr, Murray J. A quantitative model for differential motility of gliomas in grey and white matter. *Cell proliferation*. 2000; 33(5):317–329. <https://doi.org/10.1046/j.1365-2184.2000.00177.x> PMID: 11063134
42. Ambrosi D, Preziosi L. On the closure of mass balance models for tumor growth. *Mathematical Models and Methods in Applied Sciences*. 2002; 12(05):737–754. <https://doi.org/10.1142/S0218202502001878>
43. Byrne H, Preziosi L. Modelling solid tumour growth using the theory of mixtures. *Mathematical medicine and biology: a journal of the IMA*. 2003; 20(4):341–366. <https://doi.org/10.1093/imammb/20.4.341> PMID: 14969384
44. Meister P. gifAnimation processing library; 2015. Available from: <https://github.com/extrapixel/gif-animation>.
45. Poleszczuk J, Macklin P, Enderling H. Agent-based modeling of cancer stem cell driven solid tumor growth. In: *Stem Cell Heterogeneity*. Springer; 2016. p. 335–346.
46. Oliphant TE. *A guide to NumPy*. vol. 1. Trelgol Publishing USA; 2006.
47. CERN. Colt; 2004. Available from: <http://dst.lbl.gov/ACSSoftware/colt/>.
48. Bravo R, Robertson-Tessi M, West J, Anderson AR. *Cancer Crusade*; 2018. Available from: <https://cancercrusadegame.com>.
49. Kursawe J, Baker RE, Fletcher AG. Impact of implementation choices on quantitative predictions of cell-based computational models. *Journal of Computational Physics*. 2017; 345:752–767. <https://doi.org/10.1016/j.jcp.2017.05.048>
50. Bock M, Tyagi AK, Kreft JU, Alt W. Generalized voronoi tessellation as a model of two-dimensional cell tissue dynamics. *Bulletin of mathematical biology*. 2010; 72(7):1696–1731. <https://doi.org/10.1007/s11538-009-9498-3> PMID: 20082148
51. Anderson Alexander R K, Chaplain Mark A J. *Single-Cell-Based Models in Biology and Medicine*. illustrated ed. Springer Science & Business Media; 2007.
52. Jeffrey West RB. *Hybrid Automata Library*; 2018. Available from: <https://halloworld.org>.
53. Bravo R. HAL Tutorial 1: Setup; 2018. Available from: <https://www.youtube.com/watch?v=yjTmH3qORFQ&t=43s>.