Appendix A. ODD-formatted model description of GranSim

Written by Christian T. Michael with assistance from U-M GPT. Reviewed by Paul Wolberg and Denise Kirschner. This version updated November 18, 2026.

This appendix describes the computational implementation of *GranSim* using the ODD protocol (Overview, Design concepts, Details) of Grimm et al. (doi: 10.18564/jasss.4259), a structured description of the formulation of agent-based models. It covers the simulation grid, the agents that act on it, the rules by which they act, the time scales on which they act and the parameters that quantify all of these.

The description, with its companion documents, gives pseudocode and equations in enough detail for the model to be reimplemented. Biological motivation, parameter sources and citations are in the main text; this appendix describes what the program does. Notation defined in Section 2.1 is used throughout.

Finally, we note that the specific parameter values we use may change between studies or as new data are available. The ranges we present here are the ones we use unless a particular manuscript indicates otherwise.

A note on the structure of this document

GranSim is a modular platform. Several components exist in interchangeable variants, for example the diffusion solver and the resolution at which cytokine signalling is represented. Several modules are optional and can be switched on or off. A study selects one option for each module slot.

The documentation is therefore a spine plus companion documents:

- This spine contains the seven ODD elements at overview level and is intended to be reusable across studies that employ GranSim.
- Submodel documents (prefix S) specify the processes summarized in Element 7.
- Module documents (prefix M) each specify one module slot: its options, how one is selected, and the algorithm of each option.
- Appendix B indexes the configuration toggles: for each optional component, the parameter or command-line option that controls it, the default value, and the source location of the gating logic.

The module-slot table (Section 7.14) lists the slots and the option selected by default. The default configuration is the set of module options and parameter values recorded in Section 7.14 and Appendix B.

1. Purpose and Patterns

1.1 Purpose

What the program is

GranSim is a two-dimensional, grid-based, hybrid agent-based model of a single Mycobacterium tuberculosis granuloma in lung tissue. Discrete agents representing macrophages (S2), three T cell types (S3) and individual mycobacteria (S4) act on a cellular timestep of ten minutes. Soluble mediators are continuous fields on the same grid (S6, S7), advanced on a faster diffusion timestep nested inside each cellular step (M1). Tissue damage is per-compartment state that accumulates with cell death and feeds back on bacterial growth and on vascular delivery (S5).

The model is mechanistic and rule-based. No rule refers to a "granuloma"; rather, a granuloma is identified after the fact by the observation layer from the spatial configuration emergent from the model rules (Section 2.6, S9). The program therefore determines what follows from a specified set of local cellular and molecular mechanisms, under a specified initial condition and over a specified interval.

General purpose

The model serves two general purposes.

The first is explanation. The model is used to determine whether a given set of local mechanisms suffices to produce granuloma-scale phenomena observed in infected tissue, and which mechanisms are responsible for which features of the outcome. Because every mechanism is a local rule and every outcome is emergent, the model supports claims that a particular combination of mechanisms does or does not suffice to produce a particular pattern.

The second is prediction. A simulation has a complete and observable history, so the model is used to generate trajectories from specified initial conditions and mechanism strengths and to compare them with measurements made at one or more times.

A third use, demonstration of ideas, is served incidentally: the architecture makes explicit which processes are represented at what resolution, and the module structure (Section 7.14) allows resolution choices to be compared directly.

Specific purpose

The specific question the program is built to answer is:

> Given a specified set of local interactions among macrophages, effector and regulatory T cells, and mycobacteria, mediated by TNF, IL-10, three chemokines, oxygen, and accumulated tissue damage, what granuloma-scale spatial organization, bacterial burden, and infection outcome follow over the course of months from the infection of a single macrophage?

The quantities the model exists to produce, in order of importance, and the submodels that produce them:

1. Total bacterial burden and its trajectory, partitioned into intracellular, replicating extracellular and non-replicating extracellular bacteria (S4, with S2 for intracellular killing and S5 for growth suppression).
2. The spatial arrangement of cell types, caseation and bacteria within the lesion (S1, S5, S8).
3. Immune cell counts by type and activation state (S2, S3, S8).

4. Lesion size (S9).

All four are recorded by the observation layer (S9).

Definitions

Three terms are defined here because they are the criteria against which model output is assessed.

Bacterial burden is the total count of live bacteria, intracellular and extracellular, summed over the grid. Under M9, bacteria held by neutrophils are included. Where compared against colony-forming unit measurements it is scaled from two to three dimensions (S9).

Caseation is the state of a grid compartment in which accumulated cell death has reached a threshold, or which is surrounded by such compartments (S5). A caseous compartment holds no immune cells, excludes further cells, suppresses local vascular delivery, and suppresses extracellular bacterial growth in its neighbourhood. The neighbourhood count that suppresses growth includes chronically infected macrophages as well as caseous compartments (S4). Caseation is the model's representation of necrotic tissue, and a contiguous set of caseous compartments its representation of a caseous core.

Infection outcome is a classification applied to a completed simulation on the basis of its bacterial burden trajectory. It is computed outside the model from recorded output and is available to no rule.

Scope

The model represents one granuloma in a fixed tissue volume over a fixed interval. It does not represent dissemination of Mtb to distant tissues, multiple concurrent lesions, or detailed/infected lymph node or blood compartments, or host-scale disease progression. The single granuloma is taken as the unit of analysis because a single uncontrolled granuloma suffices to drive progression to active disease, so granuloma-scale outcome remains informative about host-scale outcome without the host being modelled.

The default configuration represents cytokine signalling at coarse resolution (S6) and omits antibiotic treatment (M6), immunomodulation (M7), fibrosis (M4), neutrophils (M9) and dendritic cell trafficking (M2). The selections are listed in the module-slot table (Section 7.14) and the controlling settings indexed in Appendix B. The coarse resolution solves no per-agent differential equations, so no nested submodel needs its own ODD (Grimm et al. 2020; M3.1).

1.2 Patterns

The qualitative and quantitative patterns used as criteria for deciding whether the model is realistic enough for its purpose.

1. Simulated macrophage counts, T cell counts and bacterial burden fall within the range observed in longitudinal non-human primate data (S2, S3, S4, S8).
2. Granulomas in infected tissue rarely exceed about 10^6 colony-forming units except with cavitation, which the model does not represent. This model produces virtual granulomas with Mtb ranging from completely sterile to around two million CFU.
3. We observe three outcome classes: Sterilizing, controlling and uncontrolled granulomas all occur across the sampled parameter space (S2, S3, S4).
4. Like real granulomas, our virtual granulomas develop a core of infected macrophages and caseation, with a peripheral lymphocyte population (S1, S5, S8).
5. Granuloma formation begins with innate cells only, T cells arriving weeks after infection, as observed experimentally (S8; Section 5.7).
6. GranSim captures long-controlling infections. That is, a substantial fraction of granulomas keep a persistent non-zero bacterial burden over months (S2, S4, S5).

1.3 Patterns not reproduced

As recommended in Grimm et al. (Section 4.24), we describe several patterns that the model does not (and does not purport to) reproduce, so that the description does not amount to hypothesizing after results are known.

Two-dimensional geometry misrepresents cell-cell co-localization

The frequency and arrangement of contacts between cells differ between a two-dimensional lattice and three-dimensional tissue. In two dimensions a cell has at most eight neighbours rather than twenty-six, and the fraction of a neighbourhood occupied by any cell type differs correspondingly. Every direct interaction is neighbourhood-restricted (Section 4.8), so this affects killing rates, suppression rates and the coarse-grained IFN-γ action.

The two-dimensional to three-dimensional scaling (S9) corrects counts of cells and bacteria within the lesion. It corrects no contact frequency, neighbourhood composition or other geometric property. A three-dimensional build of the platform exists (M5) and is less susceptible to co-localization artefacts; we do not typically use this because of its high computational cost.

IFN-γ is non-diffusive

IFN-γ is represented as a contact- and proximity-dependent action of producing T cells on macrophages within two compartments, not as a diffusing field (S3). This cytokine is synaptic and thus this is a reasonable simplification, but this means *GranSim* may not reproduce well the impact of drugs that may specifically increase the diffusivity of, extend the half-life of, or reduce the binding rate of IFN-γ.

This is a coarse-graining (Section 4.1): the effect of IFN-γ on macrophage activation is represented, and is the sole route to activation, but the mediating quantity is not.

Spatially resolved IL-10 is unconstrained

IL-10 is a diffusing field that inhibits secretion and drives macrophage polarization (S6), but it does not act on regulatory T cell suppression, whose strength is fixed per sample (S3). No spatially resolved measurement of IL-10 is available to validate its distribution, so the IL-10 field is constrained by its consequences rather than directly.

Oxygen rates differ from their nominal values

As implemented, net oxygen consumption integrates one cellular timestep at a quarter of the rate `o2_vmaO2` implies, and supply integrates half of each cellular timestep (S7.6).

Different boundary treatments for agents and fields

Agent movement and all agent sensing treat the domain as periodic, while the diffusing fields are held at zero concentration at the boundary (Section 2.2). The two represent the tissue beyond the boundary differently. Lesions that form centrally and stay well away from the boundary are unaffected in practice, but the model should not be expected to reproduce patterns involving lesions near a tissue margin.

1.4 Code availability

The GranSim source code is not available for distribution because of intellectual-property constraints. This appendix and its companion documents are the primary reference for the algorithmic specification and are intended to allow the algorithm to be reconstructed independently.

Each parameter is reported with its implementation identifier by the document that owns its rule, and the parameter-file mechanism is described in Section 7.0, so that a reimplementation can be parameterized consistently with the simulations reported in the main text. Sampled parameters are reported as ranges, because simulations typically sample the parameter space rather than using a single calibrated parameterization; reproducing an individual simulation requires that simulation's sample file.

The documentation describes the program as it runs, including behaviour classified as defective: a defect is reproduced and identified, not corrected.

Model rules and executables are also available via the GranSim project website (http://malthus.micro.med.umich.edu/GranSim/).

Build and environment

Item	Value
Source revision	r9959
Language	C++14
Operating systems	GNU/Linux, macOS
General-purpose library	Boost v1.79
Linear algebra	GNU Scientific Library v2.7
Spectral transforms	FFTW v3.3.10
Visualization, not used for the simulations reported here	Qt v5.12.8
Build dimensionality	Two-dimensional; set at compile time

The Boost version is part of the specification as well as the build: the interval of the uniform real draw is a property of its distribution (S1.3). Dimensionality is fixed at compile time, so the two-dimensional and three-dimensional builds are separate executables (M5). Simulations were performed on national high-performance computing resources; computational requirements belong to the simulation experiments description.

2. Entities, State Variables, and Scales

This element declares the model's entities, the grid they live on, and its spatial and temporal scales. How the entities interact follows in Elements 3 and 7.

2.1 Notation and conventions

The notation below is used throughout this appendix and its companion documents.

Timescales and intervals

Symbol	Value	Meaning
Δt_a	600 s (10 min)	Cellular (agent) timestep
Δt_d	60 s	Diffusion timestep; set by the diffusion solver (M1)
Δt_m	6 s	Molecular substep; set by the diffusion solver (M1)
N_d	10	Diffusion steps per cellular timestep, $N_d = \Delta t_a / \Delta t_d$
N_m	10	Molecular substeps per diffusion step, $N_m = \Delta t_d / \Delta t_m$
T_{shuffle}	$6\,\Delta t_a = 1$ h	Agent-list shuffle interval
T_{sim}	$22176\,\Delta t_a = 154$ d	Simulation duration

There are 144 cellular timesteps per day, 6 per hour, and $N_d N_m = 100$ molecular substeps per cellular timestep.

Δt_d and Δt_m are properties of the selected diffusion solver, not independent parameters of the biological model. The conditions they must satisfy (exact division, an even N_m, a ceiling on Δt_m and an accuracy limit on Δt_d) are checked at startup and stated in M1.2; the default values satisfy all of them.

Spatial scale

Symbol	Value	Meaning
Δx	$20\ \mu\text{m} = 20 \times 10^{-4}$ cm	Compartment edge length
L	6 mm	Domain edge length
N_{grid}	300×300	Number of compartments
A_{face}	$400\ \mu\text{m}^2 = 4 \times 10^{-6}\ \text{cm}^2$	Area of one compartment face
A_{cmpt}	$1600\ \mu\text{m}^2 = 1.6 \times 10^{-5}\ \text{cm}^2$	Lateral surface of a compartment in two dimensions, four faces
V	8×10^{-12} L	Volume of one compartment
V_{mac}	4.85×10^{-12} L	Volume of a macrophage
V_{mtb}	2.12×10^{-16} L	Volume of one bacterium
N_A (`NAV`)	$6.02 \times 10^{23}\ \text{mol}^{-1}$	Avogadro's number, as represented in the code
π (`pi`)	3.141592653	π, as represented in the code

Neighbourhoods

$N_8(x)$ denotes the Moore neighbourhood of position x: x itself and its eight surrounding compartments. A Moore neighbourhood of radius r holds $(2r+1)^2$ compartments, so N_8 is the case $r = 1$; a larger radius is written explicitly, as $N_8(x, r{=}2)$ for the 25-compartment neighbourhood.

$N_4(x)$ denotes the Von Neumann neighbourhood of x: the four compartments sharing a face with x. Unlike N_8, it excludes the centre.

Moore neighbourhoods are visited in a fixed order, and some submodels name compartments by their ordinal k:

$$\begin{matrix} 6 & 7 & 8 \\ 3 & 4 & 5 \\ 0 & 1 & 2 \end{matrix}$$

with the centre at ordinal 4. The displacement for ordinal k is

$$\delta_{\text{row}} = (\lfloor k/3 \rfloor \bmod 3) - 1, \qquad \delta_{\text{col}} = (k \bmod 3) - 1,$$

where rows increase with the first coordinate and columns with the second. This is the three-dimensional formula of M5.3 restricted to $k < 9$.

A Moore shell of radius r is the outer ring of a radius-r Moore neighbourhood. Only the connected-region construction uses shells; their traversal order is given in S9.

Boundary handling in neighbourhoods

Each neighbourhood operation comes in two variants, chosen per process.

Periodic variants wrap positions outside the grid to the opposite side, so every compartment has a full set of neighbours. A coordinate c in a dimension of size n wraps as

$$\text{wrap}(n, c) = (c + n) \bmod n.$$

Regular, or boundary-clipped, variants omit positions outside the grid, so edge and corner compartments have fewer neighbours.

All agent movement and sensing use the periodic variant, with one exception: the macrophage exposure test is clipped. The other clipped uses belong to observation or to disabled modules. Section 2.2 lists every process by regime.

Distance

Two distances are used and must not be conflated. The Euclidean distance

$$d_E(p, q) = \sqrt{(p_x - q_x)^2 + (p_y - q_y)^2}$$

serves distance from the granuloma centre of mass, ring assignment in the scaling procedure and distance from the grid centre. The Moore (Chebyshev) distance

$$d_M(p, q) = \max\left(|p_x - q_x|,\ |p_y - q_y|\right)$$

is the smallest radius r whose Moore neighbourhood about p contains q. It is used wherever a submodel tells radius-1 from radius-2 neighbours. Within a periodic neighbourhood it is taken on the unwrapped offsets, so it is the shell index of the visited compartment. For $p = (0, 0)$ and $q = (3, 3)$, $d_E = 4.24$ while $d_M = 3$.

Positions

A position is a pair of integer coordinates (or a triplet for the 3D version, see module M5).

Functional forms

Michaelis–Menten-like saturating responses are written

$$H_+(C;V,K)=\frac{VC}{C+K},\qquad H_+^h(C;V,K,h)=\frac{VC^h}{C^h+K^h},$$

where C is a concentration, V the supremum and K the half-saturation value, so $H_+(K;V,K)=V/2$; h is 1 where not written. Where several species contribute additively to one recruitment or sensing term, the term is a sum of H_+ terms.

Receptor occupancy is represented by a proxy that converts a molar dissociation constant into molecules per compartment and weighs the local field value against the agent's own receptor count; its form and constants are in S6.2 and S6.3.

Threshold-and-probability tests

Many transitions require both that a signal exceed a threshold and that a Bernoulli draw succeed. The event fires if and only if

$$x>\theta \qquad \text{and} \qquad U(0,1)<1-\exp[-k\,(\min(x,x_{\text{sat}})-\theta)]\,,$$

where x is the signal, θ the threshold, k a rate constant and x_{sat} an optional ceiling; with no ceiling, $\min(x,x_{\text{sat}})=x$. The units of x, θ and k are those of each rule (S6).

Threshold comparisons truncate both operands to seven significant figures before comparing. This makes threshold outcomes insensitive to floating-point differences below the accuracy of the numerical methods, so runs with the same executable, parameter file, command line and seed give identical results.

Stochastic events

Every probabilistic event is resolved by a Bernoulli draw against its probability at the moment it is evaluated. When an agent evaluates several events in one timestep, each is drawn independently, in the order its submodel specifies. The exceptions are listed under Stochasticity (Section 4.9).

Units

Field values are in molecules per compartment unless stated otherwise. Diffusivities are in cm^2/s, degradation rate constants in s^{-1} and secretion rates in molecules per second. Oxygen is the exception: it is carried in mg/L (S7.2). The observation layer reports concentrations in pg/mL, using the lesion volume from the cell-density extent (S9).

Identifiers and terms

A parameter is named by its parameter-file identifier, of the form `Group_name` or `Group_Subgroup_name`; for example `Core_Mac_dCCL2` is the macrophage CCL2 secretion rate, declared in group `Core.Mac`. The parameter-file mechanism is described in Section 7.0, and each parameter is tabled by the document that owns its rule.

"Suppression" means the inhibition of a cell by a regulatory T cell, which the code and its variable names call deactivation.

2.2 Spatial and temporal scales

Spatial domain

The domain is a square grid of $300\times300=90000$ compartments, each of edge length $\Delta x=20$ μm, representing a $6\ \mathrm{mm}\times6\ \mathrm{mm}$ region of lung tissue. The compartment edge is about the diameter of one macrophage.

The domain size is fixed for a run and is supplied on the command line (`--dim`), not in the parameter file, because several parameters are computed from it at load time (Element 5). One integer sets both sides, so the domain is always square. The program defaults for `--dim` and `--days` differ from the default configuration's values, so a run must supply both (Appendix B, B.8).

Simulation duration

A typical run lasts 154 days: 22 weeks, or 22176 cellular timesteps.

Nested cadences

Three nested timescales operate, from slowest to fastest.

- The **cellular timestep** $\Delta t_a=600$ s carries movement, recruitment, all agent state transitions, bacterial growth and killing, caseation, statistics and list shuffling.
- The **diffusion timestep** $\Delta t_d=60$ s carries secretion, oxygen consumption, vascular supply and the advance of the diffusing fields. There are $N_d=10$ diffusion steps per cellular timestep.
- The **molecular substep** $\Delta t_m=6$ s carries the receptor-occupancy sinks by which cells deplete local TNF and IL-10 (S6). There are $N_m=10$ substeps per diffusion step, so 100 per cellular timestep.

By default the substep loop performs only the sink update. The per-agent ordinary differential equations that would otherwise run on this timescale belong to M3, which is disabled, so no agent ODE is solved at any point. The substep structure is kept because the sinks are applied at Δt_m resolution and because N_m enters the splitting arithmetic.

Δt_d and Δt_m are set by the selected diffusion solver, not by the biology; the alternative solvers use other values (M1.2).

Operator splitting

Within each diffusion step, the processes that act as sources and sinks for the fields are split symmetrically about the diffusion solve. Writing S for the collected source and sink operator and D for diffusion, each diffusion step evaluates

$$S(\tfrac{1}{2}\Delta t_d)\ \circ\ D(\Delta t_d)\ \circ\ S(\tfrac{1}{2}\Delta t_d).$$

This is the Strang splitting of Element 3; its rationale is given in 3.3. The N_m substeps of a diffusion step are divided equally between the two half-steps, which is why N_m must be even (M1.2). With $N_m=10$, substeps 0 to 4 run in the first half-step and 5 to 9 in the second.

The time applied per diffusion step is:

Process	Argument per call	Calls per diffusion step	Time applied per diffusion step
Secretion, every cell type	$\Delta t_d/2$	2	Δt_d
Secretion from caseum	$\Delta t_d/2$	2	Δt_d
Receptor-occupancy sinks	Δt_m per substep, $N_m/2$ substeps per call	2	$N_m \Delta t_m = \Delta t_d$
Vascular oxygen supply	$\Delta t_d/2$	2	Nominally Δt_d; see below
Oxygen consumption by macrophages	Δt_d	2	Nominally $2\Delta t_d$; see below

Per cellular timestep, therefore, each secreting agent's secretion routine is called 20 times, the smoother is entered 20 times, and the sinks are applied 100 times. (Note that we handle the units of oxygen slightly differently as detailed in S7, but no behavioural rule is governed by oxygen concentration.)

Synchronization constraints

The relationships between Δt_a, Δt_d and Δt_m are checked at startup, together with an accuracy limit on Δt_d (M1.2). The accuracy limit reads only the TNF diffusivity; no IL-10 parameter enters it, and the limit evaluated at the sampled IL-10 diffusivities is also not reached. Under the selected solver the limit is far from binding. The explicit alternative is close to its limit and also imposes a hard ceiling on Δt_d, which is checked in every configuration (Appendix B, B.4.1).

Field storage

Each field holds one value per compartment. The diffusing fields, oxygen and the M7 fields are stored with a one-compartment halo that only the diffusion solvers use; agent rules never read or write it and it appears in no output.

Double buffering of the grid

The simulation holds two complete grid objects, current and next. Once per cellular timestep, before the diffusion block, the next grid is overwritten with a copy of the current grid. Each diffusion solve reads the current grid, writes the next and exchanges the two, so after N_d solves the current grid holds the advanced fields. N_d is even, so the grid object current after the loop is the one that was current before it. The M7 fields are double-buffered with the cytokine fields.

A field that the solver neither advances nor copies, oxygen in the default configuration, therefore keeps a separate history in each grid object between refreshes. The oxygen routines write to whichever object is current when they run, and only the surviving object's history carries forward: one application of each oxygen routine per diffusion step (M1.3, "Non-transformed fields and grid buffering"; S7.1).

Two kinds of state are not double-buffered but shared by pointer between the two grid objects: the per-compartment bacterial containers and the antibiotic field objects of M6. A change to a bacterial container is visible at once through either grid.

Boundary conditions

Boundary treatment is chosen per process. Three regimes are present: periodic wrapping, boundary clipping, and the field boundary condition imposed by the diffusion solver.

Agent rules and neighbourhood sensing are periodic:

Process	Neighbourhood	Owner
Chemotactic and random movement	N_8	S1
Daughter-cell placement	N_8	S3.5
Cell density, for granuloma area and vascular gating	N_8	S9.3
Entrapment test	Radius-1 ring, within the radius-2 sweep	S5
Caseation count for extracellular growth	$N_8(r{=}2)$, centre included	S4.6, S5
Extracellular bacteria sensing	N_8	S2
Intracellular bacteria and infected-cell sensing	N_8	S2, S3
Cytokine and chemokine neighbourhood sums	N_8	S6
Coarse-grained IFN-γ action on macrophages	$N_8(r{=}2)$	S3.7
T cell and macrophage interaction counting	$N_8(r{=}2)$	S3
Target selection for killing and suppression	N_8	S3
Caseum healing sweep by activated macrophages	N_8	S2.8
Bacterial dispersal on cell death	N_8	S4.13

A few processes use boundary-clipped neighbourhoods, in which positions outside the domain are omitted, so edge and corner compartments have six or four neighbours rather than nine:

Process	Neighbourhood	Status in the default configuration
Macrophage exposure-to-bacteria test	N_8	Active rule (S2.10)

Process	Neighbourhood	Status in the default configuration
Occupied-neighbour fraction, for the connected-region granuloma	N_8	Observation (S9)
Per-agent caseation statistics	N_8	Observation (S9)
Connected-region growth and outward sweep	Moore shell	Observation (S9)
Fallback placement of the reinfecting macrophage	N_8	Disabled option (M8.3)
Local field average at a vascular source in a dense compartment	N_4	Drug and immunomodulator paths only, disabled (M6, M7)

The exposure test is thus the only active rule that uses a clipped neighbourhood. Oxygen supply at a source reads the source compartment's own value, not a neighbourhood average (S7.5).

The field boundary conditions are set by the diffusion solver. Under the selected solver the diffusing cytokine and chemokine fields are held at zero concentration (Dirichlet) at the domain edge, so these species are lost across it; the fields of disabled modules would have a no-flux (Neumann) boundary. Fields to which no spatial operator is applied, oxygen among them, have no boundary condition (M1.3, "Boundary conditions").

The two regimes differ. Agents treat the domain as a torus: an agent at the left edge senses, and may move to, the right edge. The cytokine fields instead see an absorbing edge, so a cell near it sits in a sink that interior cells do not. Near the edge TNF and IL-10 are lower, so TNF-induced apoptosis is less likely and chemotactic signals are weaker, and a cell loses part of the IL-10 it secretes. The granuloma forms at the grid centre and usually stays well away from the edge, but not necessarily in samples that produce large lesions. Section 1.3 records this as a known limitation.

Relationship to companion models

The spatial scale, a $6\ \mathrm{mm} \times 6\ \mathrm{mm}$ domain of $20\ \mu$m compartments on a 10-minute cellular timestep, is shared with the companion tumour model derived from this platform, which simulates 14 days rather than 154. The common discretization permits like-for-like comparison of lesion-forming processes across the two models.

2.3 Discrete agents

Agent classes

Entity	Biological interpretation	Grid size class	Module
Macrophage	Alveolar and recruited macrophage	Large	Always present
Tgam	IFN-γ-producing effector T cell	Small	Always present
Tcyt	Cytotoxic T cell	Small	Always present
Treg	Regulatory T cell	Small	Always present
Mtb	Individual mycobacterium	Not an occupant	Always present
Neutrophil	Neutrophil	Small	Module M9
Dendritic cell	Lung dendritic cell trafficking to the lymph node	Large	Module M2
Fibroblast	Fibroblast; the myofibroblast is one of its states	Large	Module M4

Two entities are not agents in the sense of the table.

Vascular sources are grid positions held in a list owned by the grid, not agent objects. They are immovable, have no state beyond their position, are neither created nor destroyed after initialization, and occupy no agent slot, so a compartment may hold a source and up to two agents. Whether a source supplies oxygen or recruits a cell is decided each timestep by the contents of its compartment and neighbourhood (S7, S8).

Bacteria are agent objects but not grid occupants. They are held in containers: one per compartment for its extracellular population, and one per phagocyte for that cell's intracellular population. Bacteria therefore have no position of their own, occupy no slot, are not counted by the occupancy rules and are not scheduled as actors; they neither move, secrete nor compute a next state. Their state changes are driven by their container (S4).

Common state variables

Every agent class carries the variables below. A further set is declared but inert in the default configuration (see "Declared but inert state variables").

Variable	Initial value	Meaning
`id`	Next value of a global counter	Unique agent identifier
`birthTime`	Set at creation	Reference time for age-dependent rules
`deathTime`	Set at creation	Timestep at which the agent dies of senescence
`pos`	Set at creation	Grid position
`deactivationTime`	-1	Timestep at which the agent was suppressed; -1 if not suppressed

Variable	Initial value	Meaning
`intMtbTime`	0	Cumulative timesteps during which the agent has held intracellular bacteria
`surfTNFR1`, `surfTNFR2`, `surfIL10R`	Drawn at creation	Surface receptor counts
`meanIL10R`	Type mean	Used by the IL-10 sink (S6)
`trackMolecularDynamics`	`false`	Observation flag; no effect on dynamics

Senescence is tested at the start of each agent's next-state computation (S2, S3).

Receptor heterogeneity

Each agent's surface receptor counts are drawn at creation, uniformly within a type-specific range; daughters inherit their parent's counts without a draw. The heterogeneity is active, because the cytokine submodel reads each agent's own counts, so two agents in identical surroundings can differ in their chance of TNF-induced apoptosis, pathway activation and suppression. Ranges and the draw are specified in S6.3. Bacteria built by a constructor also make these draws, which are never read; bacteria created by division make none.

Lifespans

Agent type	Senescence	Value
Macrophage, not activated	`birthTime` + `Core_Mac_maxAge`	$14400\ \Delta t_a = 100$ d
Macrophage, on activation	t + `Core_Mac_maxAgeAct`	$1440\ \Delta t_a = 10$ d
T cell, recruited	`birthTime` + `Core_Tcell_maxAge`	$432\ \Delta t_a = 3$ d
T cell, daughter	Drawn within one day of the maximum	2 to 3 d
Mtb	None	—

Activation shortens a macrophage's remaining life, and suppression restores it (S2.2). T cells live about a thirtieth as long as macrophages, so the granuloma's T cell population is sustained by recruitment and proliferation rather than by long-lived cells. Daughter lifespans are drawn so that a lineage does not die in synchrony (S3.5).

Bacteria do not senesce. They die by killing, by density-dependent death in caseum, and in principle by starvation, which is unreachable with the default values (S4.8).

Compartment occupancy

Every compartment has $N_{\text{slots}} = 2$ slots, each holding at most one agent; the slot count is fixed. An agent keeps its slot until it is removed at the end of the timestep, even if it dies earlier, and slots are not compacted; a newly admitted agent takes the first free slot.

Agents fall into two size classes. Large agents are macrophages, dendritic cells and fibroblasts; small agents are T cells and neutrophils. A compartment may hold one large agent or two small ones. In the default configuration the permitted contents are:

Contents	Permitted
Empty	Yes
One macrophage	Yes
One T cell	Yes
One macrophage and one T cell	Yes
Two T cells	Yes
Two macrophages	No
Any occupant of a caseous compartment	No

The occupant count $n(p)$ counts immune cells, not bacteria (S9.3).

Each route by which an agent enters a compartment adds conditions of its own: movement (S1.4), recruitment (S8.9), daughter placement (S3.5) and initial placement (Element 5). They are not identical. Recruitment and movement share the macrophage test; a T cell moving into a compartment that already holds one occupant is admitted only with a probability, and never beside a lone neutrophil; T cell daughters test only for a free slot. Neutrophils (M9.3) may stand in caseum.

Two consequences shape granuloma architecture. A compartment holds at most one macrophage, so the macrophage core is at most one cell per compartment dense. But a T cell may join a macrophage, so lymphocytes can interpenetrate the macrophage core rather than being confined to a surrounding ring. Caseous compartments cannot be entered, and every agent present when a compartment becomes caseous is killed at that moment (S5).

Hybrid state update

Each agent carries a current state and a pending next state, and the update is hybrid. This determines exactly when one agent's action becomes visible to another.

- **Self-transitions** are written to the pending state during the next-state pass and committed only after it. An agent that changes its own state presents its old state to every other agent for the rest of the timestep.

- **Effects on other agents**, killing above all, are written at once to both the current and the pending state of the target, so an agent killed partway through a pass is seen as dead by every agent processed after it.

Pending-state assignment is monotone toward death: once the pending state is dead, later assignments to it have no effect. An agent killed during its own rule sequence therefore stays dead, yet the rest of its sequence still runs, so behaviours late in the sequence are not under-represented merely because the agent might have been killed earlier. Some routines, among them the caseation cascade and some kill branches, assign the pending state directly instead of through this assignment (S2.3, S5).

Bacteria are outside this scheme: the container sets a bacterium's current and pending state together and starts its decay clock (S4).

Removal of dead agents

Dead agents are removed from the lists and the grid at the end of the timestep, after all states are committed (Section 3.2). Removal increments a per-compartment death counter, which no rule reads. It is distinct from the killing count that drives caseation (S5.2).

Declared but inert state variables

Every agent also carries variables of modules disabled by default: membrane-bound and receptor-bound TNF, internalized receptors, transcription and shedding rates, NF-κB and MHC class II pathway species (M3), anti-TNF complexes (M7) and cumulative bound TGF-β (M4). They are initialized and saved with the simulation state but read by no active rule; Appendix B lists them. They are mentioned here so that a reader comparing this description with a parameter file or saved state is not misled.

2.4 Continuous fields

Soluble species are scalar fields on the agents' grid, one value per compartment, in molecules per compartment unless stated otherwise. The fields fall into three categories, advanced by different mechanisms.

Category	Fields	Advanced by
Diffusing	TNF, CCL2, IL-10	The diffusion solver, once per Δt_d (M1)
Derived	CCL5, CXCL9	Set from CCL2 after each solve
Non-diffusing	O_2	Vascular supply and consumption only (S7)

Diffusing fields

Each diffusing field obeys a reaction–diffusion–decay equation,

$$\frac{\partial C}{\partial t} = D\nabla^2 C - k_{\mathrm{deg}} C + S(x,t),$$

with field-specific diffusivity D and degradation rate constant k_{deg}. The source term S collects secretion by cells in the compartment and, for the chemokines, by caseous compartments.

Field	Produced by	Role
TNF	Macrophages (NF-κB-dependent); Tgam and Tcyt designated as TNF producers	Macrophage activation pathways; apoptosis of macrophages and T cells; recruitment of all types
CCL2	Macrophages; caseous compartments	Chemotaxis of macrophages and Tgam; recruitment of macrophages and Tgam
IL-10	Infected, chronically infected and activated macrophages; Tregs	Inhibition of TNF and chemokine secretion; macrophage STAT3 pathway

The degradation rate constants belong to the biological model (S6.16); under the selected solver decay is folded into the solve (M1.3, "The propagator and incorporation of decay"), and negligible values are cut to zero after each solve (M1.3, "Normalisation and the cutoff").

Derived fields

CCL5 and CXCL9 are not diffused. After each solve they are set pointwise from CCL2,

$$C_{\mathrm{CCL5}}(p) = r_5\, C_{\mathrm{CCL2}}(p), \qquad C_{\mathrm{CXCL9}}(p) = r_9\, C_{\mathrm{CCL2}}(p),$$

with $r_5 = $ `Core_Mac_dCCL5` / `Core_Mac_dCCL2` $= 1.0$ and $r_9 = $ `Core_Mac_dCXCL9` / `Core_Mac_dCCL2` $= 2.0$. This rests on the model's assumption that the three chemokines are secreted by the same cells in fixed proportion and share one diffusivity and degradation rate (S6.14). It saves two of the six transforms per diffusion step.

The consequences are easy to miss, and a reimplementation that diffused the three chemokines independently would not reproduce this model:

- There is in effect one chemokine field read at three scalings, so the agent types' chemotactic sensitivities differ in magnitude but not in spatial pattern (S1.3).
- A direct write to CCL5 or CXCL9, such as caseum secretion, is overwritten at the next solve; only writes in the half-step after the final solve of a timestep persist into the agent phase (M1.3, "Derived assignment of CCL5 and CXCL9").
- The ratios follow the macrophage chemokine secretion parameters, which are fixed in the default configuration, so r_5 and r_9 are constant across the virtual library.

Non-diffusing field: oxygen

Oxygen is a field but is never advanced by a spatial operator: its diffusion code is unreachable in every configuration (M1.3). It changes only through supply at vascular sources, gated on local cell density and caseation, and saturable consumption by macrophages, both applied in each half-step (S7). Its spatial structure therefore comes from the placement of sources, the distribution of macrophages and the density gate. There is no lateral transport, so a depleted compartment is replenished only if it is itself a source. Oxygen is initialized uniformly rather than to zero (Section 5.5). It is carried in mg/L (S7.2); the conversion applied by the consumption term, and the resulting consumption rate,

are given in S7.6.

Fields present but typically inert

These fields stay zero throughout a run in the default configuration because the modules that write them are disabled. They are listed so that a parameter file, grid dump or saved state does not mislead.

Field	Module	In the default configuration
Shed TNFR2	M3	Allocated and zero; would diffuse and release TNF in the post-solve sweep (M1.3)
Active TGF-β	M4	Transformed on every solve, but zero
Inactive TGF-β	M4	Copied forward and degraded on every solve, but zero
Collagen	M4	Allocated and zero
Etanercept, infliximab and their TNF complexes	M7	Allocated and zero
Antibiotic fields, per drug	M6	Not created unless a drug is named

Molecular weights enter only the observation conversions (S9) and the oxygen supply term (S7.2).

2.5 Agent state variables

This section lists the variables specific to each agent type present in the default configuration, in addition to the common variables of Section 2.3. The rules that read and write them are in Element 7 and the companion documents.

Macrophage

States: resting, infected, chronically infected, activated, dead (`MAC_RESTING`, `MAC_INFECTED`, `MAC_CINFECTED`, `MAC_ACTIVE`, `MAC_DEAD`).

Variable	Initial value	Meaning
`state`, `nextState`	As created	Current and pending state
`intMtb`	Container with the specified count	Intracellular bacteria
`activationTime`	-1	Timestep of activation
`exposedToMtb`	`true` if created infected	Whether the cell has recently met bacteria
`exposureTime`	-1	Timestep at which exposure lapses

Three intracellular signalling pathways are represented: STAT1 and NF-κB, which are pro-inflammatory, and STAT3, which is anti-inflammatory. Each has a discrete flag with a lapse time, a graded counter and a continuous gene index (S2.4). The flags gate state transitions; a macrophage becomes activated only when STAT1 and NF-κB are both set. The counters and gene indices gate nothing and feed only a polarization index reported by the observation layer (S9). Constructor invariants on state and bacterial content are in S2.2.

T cells

All three T cell types share:

Variable	Initial value	Meaning
`cognate`	Drawn at creation	Whether the cell is Mtb-specific
`proliferationTime`	-1	Timestep at which a daughter will be produced; -1 if not stimulated
`divisionCount`	0	Divisions accrued, inherited by daughters
`macInteractionStat`, `exposedMacInteractionStat`	Zero vectors	Observation counters
`totalExposedMacInteractions`	0	Accumulator for the M8 exhaustion rule

Antigen specificity is the most consequential T cell variable. It is drawn once at creation and fixed thereafter. A non-cognate T cell moves and can die of senescence, TNF-induced apoptosis or caseation, but has no effector function: it does not secrete, proliferate, kill or suppress. The cognate fraction is small, so most T cells on the grid are inert (S3.2). Because daughters inherit the division count, a founder lineage is bounded in size (S3.5).

Tgam. States `TGAM_T0`, `TGAM_ACTIVE`, `TGAM_DOWN_REGULATED`, `TGAM_ACTIVE_DOUBLE`, `TGAM_INDUCED_REG`, `TGAM_EXH`, `TGAM_DEAD`. Three are reachable in the default configuration: active, down-regulated and dead. `TGAM_T0` is assigned only under the MHC class II component of M3, and the double-producer, induced-regulatory and exhausted states belong to M8. A Tgam also carries two producer flags, drawn at creation for cognate cells: `ifngProducer`, whether it exerts the coarse-grained IFN-γ action, and `tnfProducer`, whether it secretes TNF. Four further variables support the M8 Treg-induction option and are not updated by default. Producer status is a binary, heritable trait, not a graded rate. Because IFN-γ action is the only route by which a macrophage's STAT1 flag is set, the local density of IFN-γ producers decides macrophage activation (S3.7).

Tcyt. States `TCYT_T0`, `TCYT_ACTIVE`, `TCYT_DOWN_REGULATED`, `TCYT_EXH`, `TCYT_DEAD`; reachable by default: active, down-regulated and dead. One extra variable, `tnfProducer`.

Treg. States `TREG_T0`, `TREG_ACTIVE`, `TREG_DEAD`, with no variables beyond the shared set. Tregs do not secrete TNF, do not proliferate and cannot be suppressed; their effector action is contact suppression of neighbouring cells (S3.10).

Mycobacterium tuberculosis (Mtb)

States: live and dead (`MTB_LIVE`, `MTB_DEAD`).

Variable	Initial value	Meaning
`state`, `nextState`	`MTB_LIVE`	Live or dead
`biomass`	`Mtb_Agents_initialBiomass` $= 1$	Grows each timestep; divides above a threshold
`storedGrowthRate`	Fixed, or drawn under M8	Effective rate on the latest growth step, kept for statistics
`mtbSpecificIntraGrowthRate`, `mtbSpecificExtraGrowthRate`	`Mtb_Agents_growthRateIntMtb`, `Mtb_Agents_growthRateExtMtb`	Intrinsic per-bacterium growth rates
`generationTime`	0	Timesteps between the latest two divisions
`nrDivisions`	0	Divisions by this lineage member
`parentID`	0	Identifier of the parent
`decayClock`	0	Set at death; counts down to deletion

The intrinsic rates change only under M6 (resistance, compensation and seeding of resistant bacteria); `storedGrowthRate` records the rate after environmental modulation (S4.5). A dead bacterium stays in its container's dead list for a few timesteps before deletion and remains detectable meanwhile (S4.4, S2.10). Five further variables, among them the reserve quantities and a biomass degradation rate, are declared and read by no rule (S4.3).

Bacterial containers

Location designation	Holder	Meaning
`INTRA_CELLULAR`	A macrophage	Intracellular population
`REPLICATING_EXTRA_CELLULAR`	A compartment	Extracellular, labelled replicating
`NON_REPLICATING_EXTRA_CELLULAR`	A compartment	Extracellular, labelled non-replicating

Every compartment holds a container from the start, empty and labelled replicating. After each growth pass a compartment's container is re-labelled according to whether the extracellular growth rate there has fallen to zero (S4.4). The label is a reporting quantity: it changes no behaviour, since growth arrest follows from the environment, not from the label (S4.3). Each container keeps separate live and dead lists and exposes both a count and a biomass sum (S4).

Absent agent types

Neutrophils (M9) and dendritic cells are implemented but absent: no instance is created at any point. Fibroblasts belong to M4. Their variables are described in the module documents and indexed in Appendix B.

2.6 Compartment and environmental state

Each compartment carries non-field state: integers and scalars that do not diffuse and are stored without a halo.

Compartment state inventory

Variable	Reset	Meaning
`nKillings`	Never; halved when caseation heals	Cumulative count of cell deaths that count as tissue damage
`trappedCaseation`	Recomputed each timestep	Whether the compartment is caseous by entrapment
`healingTime`	Set and cleared by the healing rule	Timestep at which healing began; 0 if none
`nDeaths`	Never	Agents removed from this compartment; read by no rule
`nRecruitments` and the per-type recruitment counters	Never	Recruitment events here (S8, S9)
`nSecretions`	Never	Secretion events (S9)
`nCells`	Refreshed under M6	Cached cell density; read by no rule
`inGran`	Recomputed each timestep; frozen once fibrosis begins under M4	Threshold-granuloma membership; read only by the M4 pre-fibrosis counts
`ExtFIC`, `IntFIC`	Module M6	Fractional inhibitory concentrations

`nKillings` and `trappedCaseation` drive the dynamics through caseation, and `healingTime` drives its reversal; the rest are observation counters or belong to disabled modules.

The killing count is the model's record of accumulated local tissue damage. It is incremented by cell deaths that count as damage, which is not every death: senescence of a non-activated macrophage, senescence of a T cell and, under M4, the myofibroblast apoptosis proxy do not count (S5.2).

Caseation

A compartment is caseous when its killing count reaches the sampled threshold `Core_nrKillingsCaseation`, or when it is entrapped, fully ringed by caseum, whatever its own count (S5). Caseous compartments are not inert. They secrete chemokines at a reduced rate, so caseum stays a chemoattractant with no cell present; they admit and keep no agent; they suppress vascular oxygen supply; and bacteria in them are subject to density-dependent death. A partial-caseation measure is also reported; it reads the killing count alone and so is not nested within caseation (S5.7).

Occupancy, cell density and the granuloma

Several rules and the observation layer use three related occupancy predicates and the densities built on them: the immune-cell count $n(p)$; cell occupancy, at least one immune cell; and the broader occupancy that also counts live extracellular bacteria and caseum (S9.3). Cell density $\delta(p)$, the occupied fraction of the periodic N_8 in the broad sense, gates vascular supply (S7.4) and defines the threshold granuloma. Because it counts caseum and free bacteria, a region that has been entirely destroyed still registers a density of one: it measures tissue involvement, not viable cellularity.

Two non-equivalent definitions of the granuloma are recomputed each timestep, for different purposes (S9):

- The **threshold definition** counts a compartment when its cell density is at least `--area-cell-density-threshold` $= 0.5$ or it is caseous. It is local, needs no connectivity and may include isolated patches. The threshold test feeds the granuloma area, the in-granuloma cell counts, the lesion volume used for cytokine concentrations and the M6 regional statistics.
- The **connected-region definition** takes the largest contiguous region of cell-occupied compartments whose clipped neighbourhoods are sufficiently cell-occupied. It excludes caseum and free bacteria, reports one region with a centre of mass and two diameters, and drives the reported lesion diameters, the region used for three-dimensional scaling and the test of whether a dying macrophage releases bacteria outside the granuloma.

The two can disagree substantially: a lesion with a large caseous core looks smaller under the connected-region definition, which is also boundary-clipped and discards every component but the largest. The observation layer treats caseum differently again in its granuloma extents, its PET-hot proxy and its caseation statistics (S9).

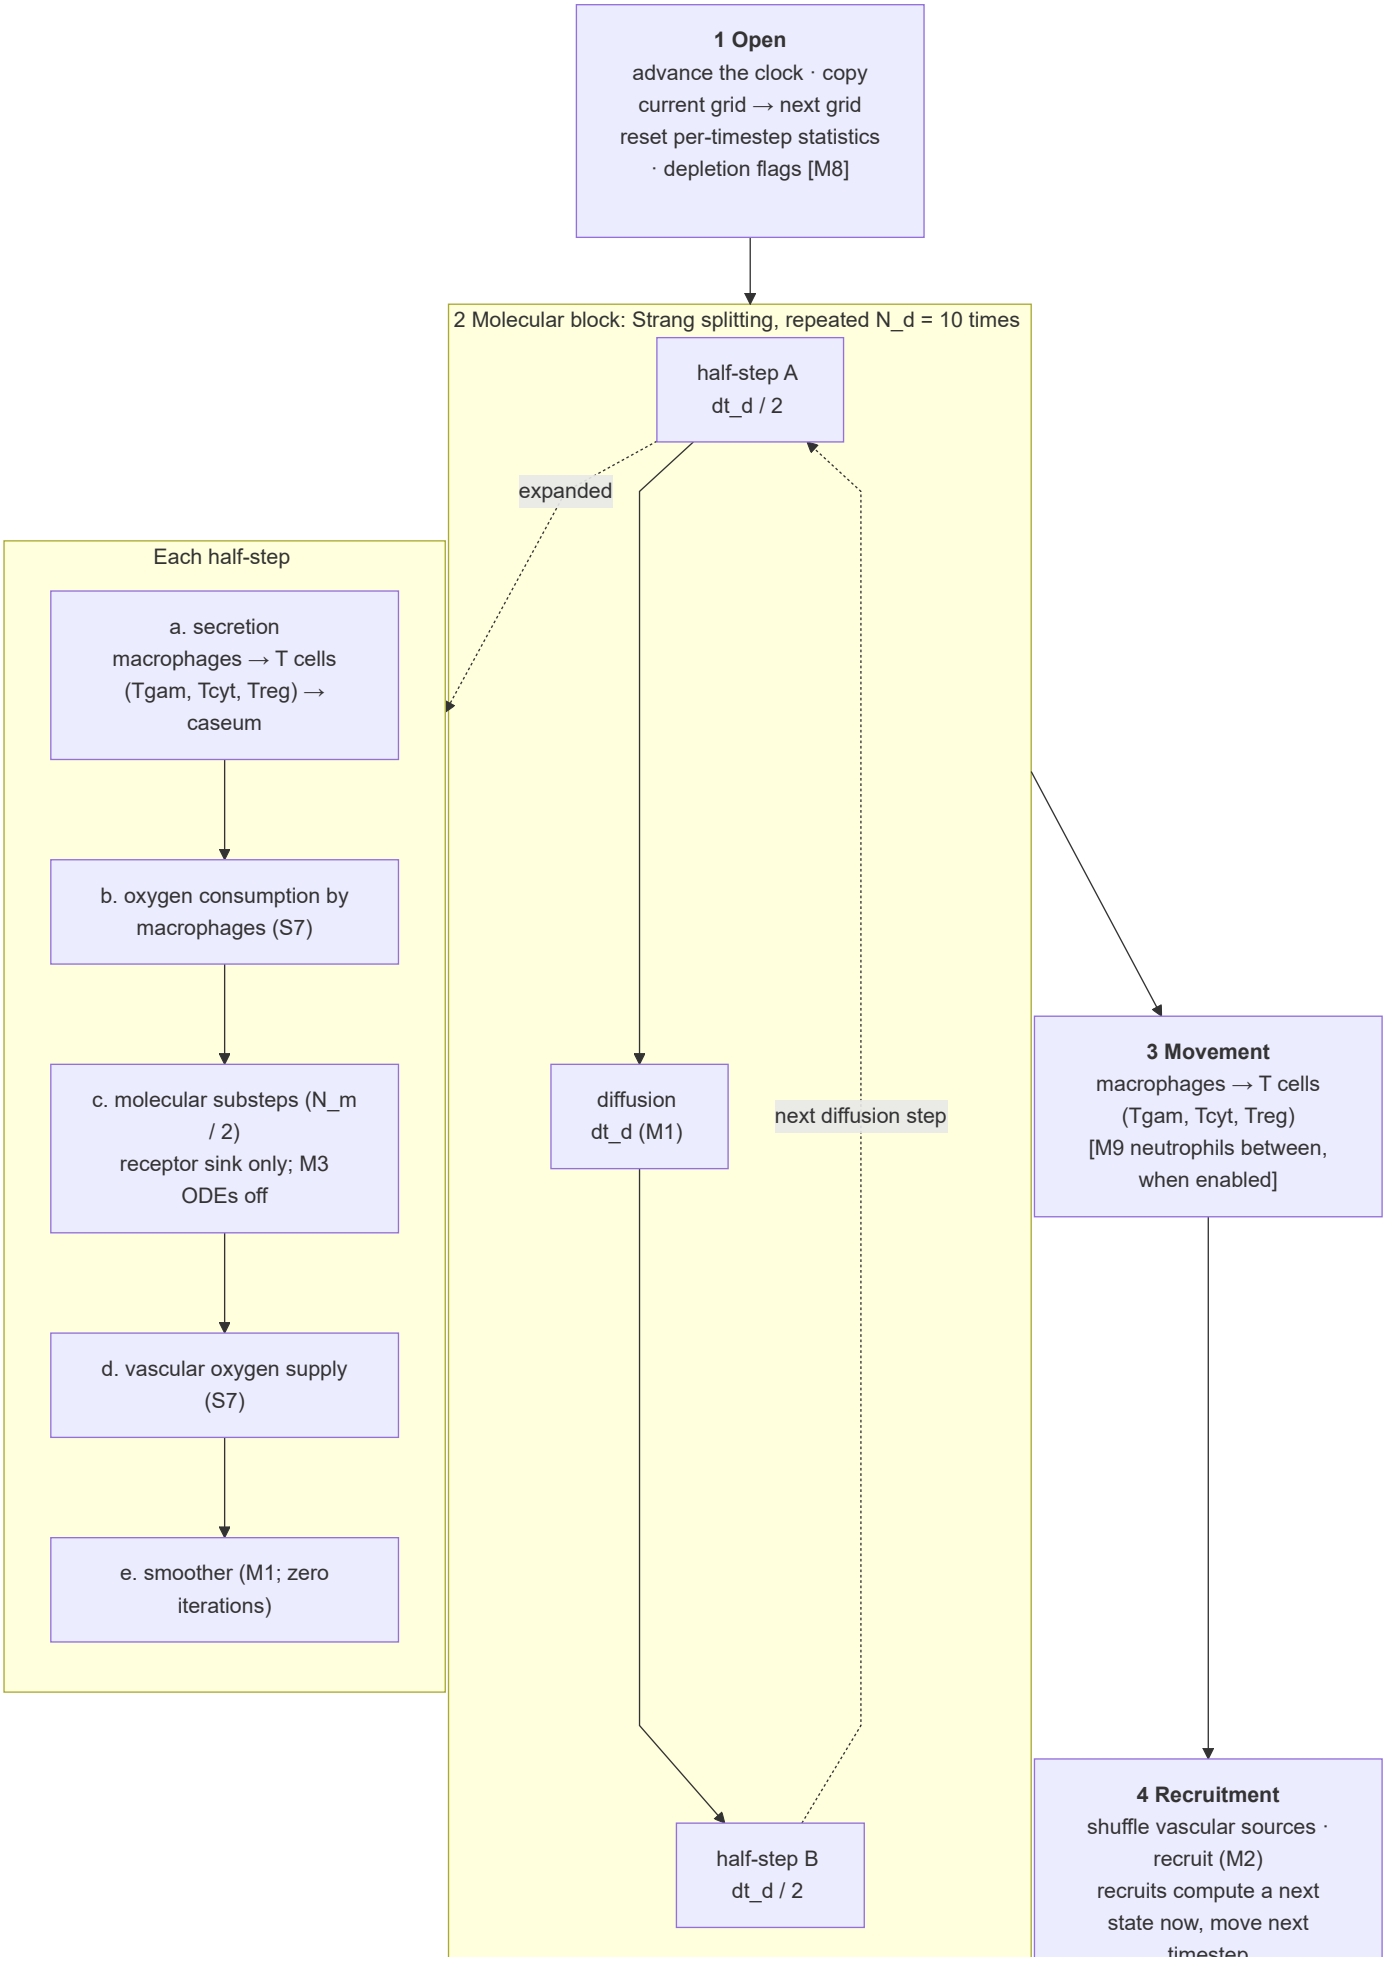
The observation layer also computes seven per-compartment granuloma extents with equivalent diameters and scaling factors, and a PET-hot score from cellular composition; both are specified in S9.

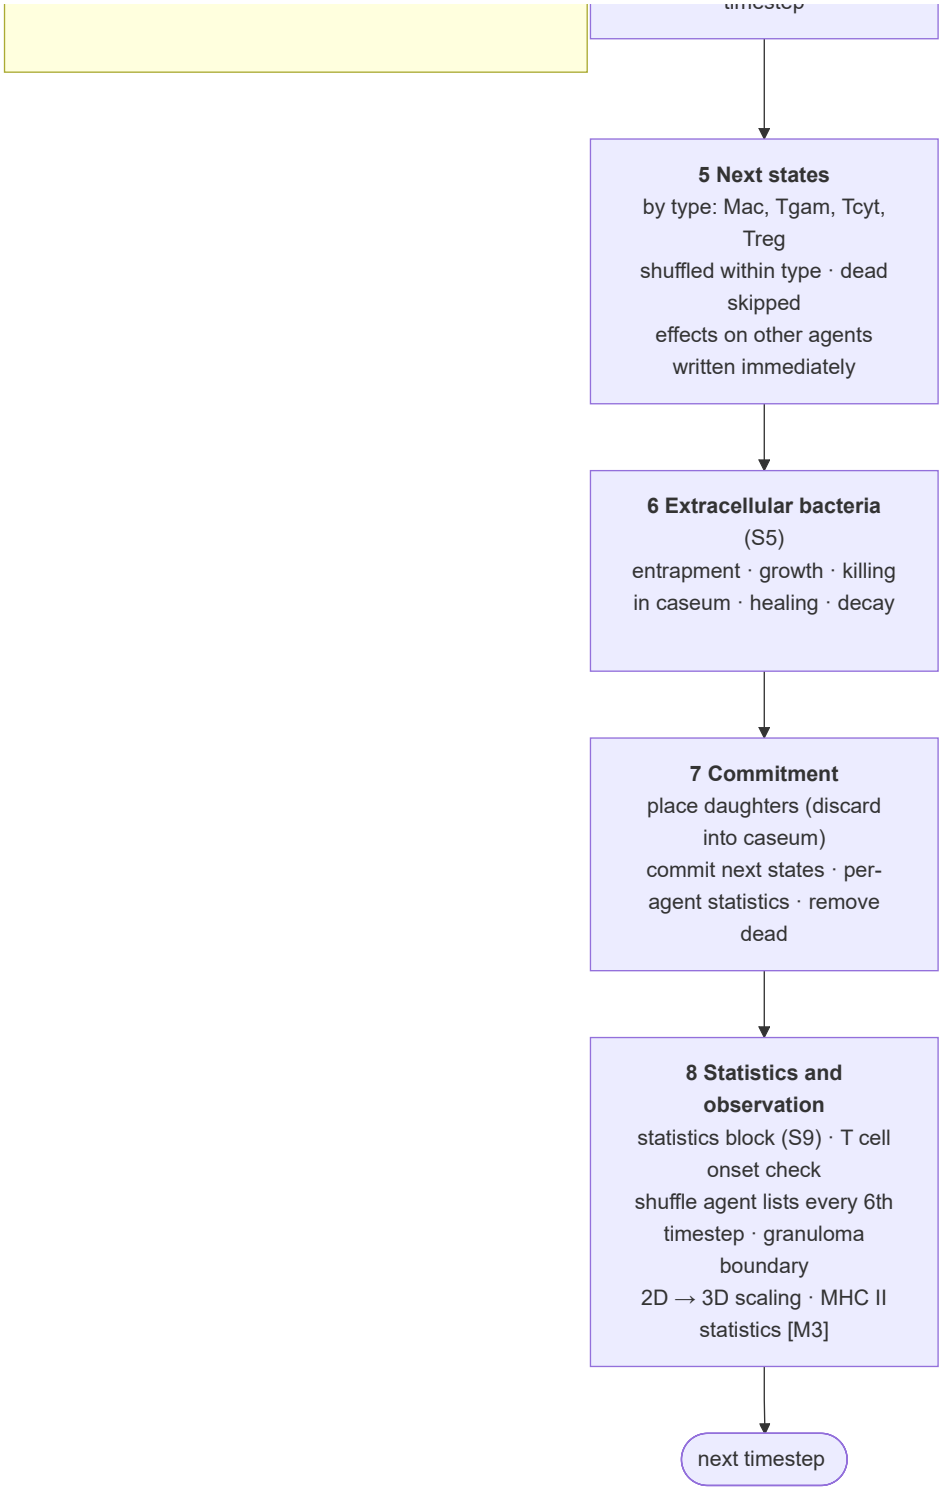
3. Process Overview and Scheduling

GranSim uses a partially asynchronous schedule. Agents are held in per-type lists and processed one list at a time in a fixed type order; the order within a list is shuffled every hour. An agent's own transitions are written to a pending next state and committed after the whole pass, whereas its effects on other agents, killing above all, take effect at once and are visible to agents processed later in the same pass (Section 2.3).

3.1 One cellular timestep: visual overview

Below is a flowchart of one cellular timestep.


1 Open
advance the clock · copy current grid → next grid
reset per-timestep statistics · depletion flags [M8]
2 Molecular block: Strang splitting, repeated N_d = 10 times
half-step A
dt_d / 2
expanded
Each half-step
a. secretion
macrophages → T cells (Tgam, Tcyt, Treg) → caseum
b. oxygen consumption by macrophages (S7)
c. molecular substeps (N_m / 2)
receptor sink only; M3 ODEs off
d. vascular oxygen supply (S7)
e. smoother (M1; zero iterations)
diffusion
dt_d (M1)
next diffusion step
half-step B
dt_d / 2
3 Movement
macrophages → T cells (Tgam, Tcyt, Treg)
[M9 neutrophils between, when enabled]
4 Recruitment
shuffle vascular sources · recruit (M2)
recruits compute a next state now, move next timestep



`[Mn]` marks a step that runs only when module *n* is enabled; all are off in the default configuration.

3.2 One cellular timestep

The timestep in outline

Opening. Each cellular timestep advances the clock first, so the first timestep processed is $t = 1$. Under M8 a reinfecting macrophage may be placed at this point, before anything else happens. The next grid is then refreshed from the current grid and the per-timestep statistics are reset.

The molecular block. The fields are brought forward by ten diffusion steps, each a diffusion solve with a source-and-sink half-step on either side (Section 2.2). In each half-step agents secrete according to their type and state (S6), caseous compartments secrete, macrophages consume oxygen, the receptor-occupancy sinks remove TNF and IL-10 on the molecular substeps, and vascular sources supply oxygen. When the block ends, the fields hold this timestep's secretion, transport and decay, including what was written in the half-step after the final solve, and these are the values every agent reads for the rest of the timestep (S1.1). The block comes first so that a signal secreted this timestep can steer neighbours this timestep. Dead agents were removed at the end of the previous timestep, so no dead macrophage secretes (S6).

Movement. Agents then move, one type at a time: all macrophages, then all T cells (Tgam, Tcyt, Treg). Each attempts a move only when its movement opportunity comes round (S1.2): T cells every timestep, macrophages at intervals that depend on their state. Moving before recruitment means incumbents settle first, and recruits enter whatever slots are left.

Recruitment. At each vascular source, in shuffled order, new cells may be recruited into free slots (S8). A recruit joins its type's list and so computes a next state on its arrival timestep, but it has not moved or secreted; it does both from the next timestep (S8.1).

Next states. The per-agent statistics are reset, and every live agent decides its next state, type by type: macrophages first, then Tgam, Tcyt and Treg. Each agent's own decision is held until commitment, but what it does to others (killing, suppression, IFN-γ stimulation, uptake of bacteria) takes effect at once. A cell due to divide reserves a compartment for its daughter, which is not yet placed.

Extracellular bacteria. One sweep over every compartment then handles the extracellular populations: the entrapment test, growth and re-designation of live bacteria, killing in caseum, healing of caseum, and decay of dead bacteria (S5.1, S4). Because this follows the next-state pass, bacteria a macrophage takes up during its own rule are grown within that rule, not by the sweep (S2).

Commitment. Promised daughters are placed; a daughter whose reserved compartment has meanwhile become caseous is discarded and counted as a killing there (S5). Every pending state is then committed, each agent reports its per-agent statistics, and dead agents are removed, freeing their slots.

Observation. The statistics block describes the committed state at the end of the timestep (S9.2). The T cell recruitment onset is latched (S8.10), the agent lists are shuffled once an hour, the connected-region granuloma is computed, and its counts are scaled to three dimensions (S9). The macrophage release test reads this region on the following timestep (S9.8).

Pseudocode for one cellular timestep

The sequence below is one cellular timestep. Branches that depend on a disabled module are kept and marked, so the control flow is complete and a configuration enabling the module can be read against the same skeleton. Lines marked [Mm] run only when that module slot is active, and are not taken in the configuration recorded in Section 7.14. Lines marked [absent] belong to agent types that never exist in this configuration.

```
for each cellular timestep:

    t <- t + 1                                  // first timestep processed is t = 1

    if t == Core_reinfectionTimestep:                               [M8]
        place the reinfecting macrophage        // M8.3

    refresh next grid from current grid         // Section 2.2

    if t == Drugs_dosageStartTime:                                  [M6]
        add randomly placed resistant bacteria

    reset per-timestep statistics
    set the cytokine depletion flags from their timesteps           [M8]
    if drugs are enabled:                                           [M6]
        cache cell density in every compartment  // read by no rule

    // ---- Molecular block, Strang split (Section 2.2) ----

    for diffusionStep = 0 .. N_d - 1:
        half-step(substeps 0 .. N_m/2 - 1)
        diffuse fields, current -> next, and exchange grids         // M1
        half-step(substeps N_m/2 .. N_m - 1)

    if TGF-beta dynamics selected:                                  [M4]
        activate TGF-beta

    // ---- Movement (S1), in movement type order ----

    move macrophages
    move dendritic cells                                            [absent]
    move neutrophils                                                [M9]
    move T cells                                // Tgam, then Tcyt, then Treg
    if TGF-beta dynamics selected:                                  [M4]
        move fibroblasts

    // ---- Recruitment (S8, M2) ----

    shuffle the vascular source list
    recruit                                     // recruits join the main lists

    reset per-agent statistics

    // ---- Next states, in next-state type order ----

    for type in (Mac, DC, Tgam, Tcyt, Treg, fibroblast, neutrophil):
        for each agent a in list[type]:         // order carried from last shuffle
            if a is dead or pending dead:
                continue
            a computes its next state           // effects on other agents
```

```
                                        // are written at once

    // ---- Extracellular bacteria (S5.1, S4) ----

    for every compartment p, in storage order:
        entrapment test at p                    // may kill occupants
        if p holds live extracellular bacteria:
            grow them; re-designate the container
            if p is caseous:
                kill by caseum density
            drug exposure, killing, efflux and counts           [M6]
        complete healing at p if due
        decay dead extracellular bacteria at p

    // ---- Commitment ----

    for each type with a proliferation list:
        for each daughter d:
            if d's compartment is caseous:
                count a killing there           // d is discarded
            else:
                append d to its main list and place it
    clear the proliferation lists and the reservation map

    for type in (Mac, DC, Tgam, Tcyt, Treg, fibroblast, neutrophil):
        for each agent a in list[type]:
            commit a's next state
            record a's per-agent statistics
        remove every dead agent, freeing its slot

    // ---- Observation (S9) ----

    mark macrophages exposed to bacteria
    statistics block                            // S9.2; includes [M6], [M3] parts
    latch the T cell recruitment onset          // S8.10
    if t is a multiple of 6:
        shuffle every agent list
    compute the connected-region granuloma
    inject neutrophils                                      [M8]
    scale counts to three dimensions
    MHC class II statistics                                 [M3]
```

The diffusion half-step

Each of the two half-steps of a diffusion step does the following, with $h = \Delta t_d/2$.

```
half-step(first substep .. last substep):

    // ---- Secretion, in type order (S6) ----
    macrophages secrete over h
    T cells secrete over h                      // Tgam, then Tcyt, then Treg
    fibroblasts secrete over h                              [M4]
    neutrophils secrete over h                              [M9]
    caseous compartments secrete over h         // S5

    // ---- Oxygen consumption (S7.6) ----
    macrophages consume oxygen over the full dt_d

    // ---- Drugs ----
    if drugs are enabled:                                   [M6]
        bacterial and macrophage drug uptake over h; partitioning

    // ---- Molecular substeps ----
    for each substep:
        integrate agent ODEs over dt_m                      [M3]
        apply receptor-occupancy sinks over dt_m // S6
        integrate fibroblast TGF-beta ODEs over dt_m        [M4]

    // ---- Vascular sources (S7) ----
    drug supply over h                                      [M6]
    immunomodulator supply over h, from its dosing start    [M7]
    oxygen supply over h

    // ---- Smoother (M1) ----
    smooth the fields                           // selected solver only
```

By default each substep applies only the sinks, which run whenever either receptor component of M3 is off. The smoother is entered in every half-step but performs zero iterations at the default setting (M1.3, "The smoother").

3.3 Why the order matters

This section explains the features of the schedule whose purpose is not evident from the pseudocode.

Operator splitting

Secretion, consumption and vascular supply act as sources and sinks for the fields. Applying them wholly before or wholly after the diffusion solve would be first-order accurate in Δt_d; splitting them symmetrically raises the source treatment to second order, at the cost of one extra evaluation per diffusion step. The substeps must divide equally between the halves, so that the sinks applied before and after the solve are of equal size; an odd N_m would break the symmetry and the second-order property with it. The diffusion scheme itself, including how decay is folded in, is specified in M1.3.

Agent-type orders

Agents are processed in lists grouped by type, in two fixed orders:

- **Movement:** macrophage, dendritic cell, neutrophil, Tgam, Tcyt, Treg, fibroblast (fibroblasts only under M4).
- **Next states and commitment:** macrophage, dendritic cell, Tgam, Tcyt, Treg, fibroblast, neutrophil.

Only macrophages and the three T cell types exist in the default configuration, so in practice both orders are macrophages, then Tgam, Tcyt and Treg.

The order matters because effects on other agents are visible at once within a pass. Macrophages decide their next states before any T cell, so a T cell acts on the macrophage configuration the macrophages have just produced, including any that became activated, chronically infected or dead. A macrophage, conversely, sees T cell actions only a timestep later: when a Tgam's IFN-γ action sets a macrophage's STAT1 flag, that macrophage has already decided its next state, so the stimulation acts from the next timestep (S3.7). Among T cells, Tregs act after the effectors, so a Treg suppressing a Tgam prevents that Tgam's actions on the next timestep, not the current one. The type orders are never shuffled.

The orders are part of the model definition. Processing the types in another order, or interleaving them, gives different trajectories.

Hybrid state update

The split between deferred self-transitions and immediate effects on others (Section 2.3) serves two purposes. Deferring an agent's own transition stops it propagating within the timestep: without it, a macrophage activated early in the pass could kill bacteria in the same timestep, and list position would decide how much each agent achieved. Applying effects on others at once stops a killed agent from being killed again, counted twice in the caseation tally, or acting further; the skip of dead and pending-dead agents at the top of the next-state pass serves the same end.

Movement rate

Movement opportunity depends on type and state (S1.2). T cells attempt to move every timestep; macrophages attempt to move at state-dependent intervals counted from their own birth time, and chronically infected macrophages never move. With the sampled intervals, T cells attempt to move four to six times as often as resting macrophages and more than a hundred times as often as infected ones. Lymphocytes therefore keep redistributing around a bacterial core that becomes spatially fixed. An opportunity is not a move: whether the attempt succeeds depends on the destination drawn and its occupancy (S1).

Deferred placement of daughter cells

A cell due to divide does not create its daughter in place. It reserves a compartment, and the daughter waits in a per-type proliferation list until commitment. This keeps a daughter from computing a next state, or dividing, on its creation timestep, which could otherwise produce a whole lineage in one timestep, and from changing what other agents do before the next timestep. A reservation map records how many slots remain in each promised compartment, so two daughters are not given the same slot (S3.5).

A cell that finds no free slot nearby does not divide and keeps its readiness, retrying on later timesteps, so proliferation is limited by space in dense regions. Daughters are placed before states are committed, so each exists on the grid and appears in the statistics from its creation timestep, though it has no behaviour then.

List shuffling

Three shuffles occur, on different cadences.

- **Agent lists**, every six timesteps (one hour). The shuffle permutes the order within each type only; it changes no position, state or occupancy. Between shuffles the order is carried forward, with recruits and committed daughters appended and dead agents removed (S1.5). It keeps a fixed order from giving some agents a lasting advantage, since an agent processed earlier may kill, suppress or claim space first.
- **The vascular source list**, every timestep, before recruitment. Sources do not compete for slots; the order matters only through the macrophage recruitment floor, which reads a count raised by each macrophage recruited, and the shuffle randomizes which sources are evaluated under it (S8.2).
- **Each bacterial container**, before its contents grow. Order within a container could matter only if the group size changed during growth; it does not, so the shuffle affects no aggregate quantity, though it decides which bacterium receives which division draw (S4.4).

Position of the statistics

Each agent reports its per-agent statistics as its state is committed, before dead agents are removed. The statistics block follows removal, so it records the completed state of the timestep. Within the block the granuloma extents come first, because other statistics depend on them (S9.2).

The connected-region granuloma is computed after the statistics block. Neutrophil injection and the three-dimensional scaling, which follow it, read the current timestep's region. The macrophage release test (S9.8) and a neutrophil routine (M9.2) read it on the following timestep, so for them it is one timestep stale.

3.4 Reproducibility

The schedule is part of the model definition, not an implementation detail. Two models with the same rules but different schedules are different models, and with immediate effects between agents the difference can be large. A conceptual replication must preserve the nesting and splitting of Section 2.2, the block order and type orders of Sections 3.2 and 3.3, and the hybrid update of Section 2.3. Above all, the molecular block precedes all agent behaviour, so agents act on fields that already reflect this timestep's secretion, diffusion and decay.

Exact reproduction of a trajectory needs more: the number and order of random draws, the significant-figure truncation of threshold comparisons, and the transform planning strategy. A reimplementation can reproduce the model's behaviour without them.

4. Design Concepts

GranSim is a mechanistic, rule-based, grid-based ABM. Granuloma-level behaviour is not prescribed: lesion formation, spatial stratification, caseation and the balance between bacterial control and escape all emerge from local interactions among agents and fields.

4.1 Basic principles

Overarching approach

Granuloma-level outcomes are not modelled directly; they arise from the interaction of individually represented cells and bacteria with one another and with a shared chemical environment. No rule refers to a granuloma as an object, and none prescribes lesion size, shape, stratification or trajectory. The granuloma is an operational construction applied afterwards by the observation layer, which picks out regions of the grid meeting a density or connectivity criterion (Section 2.6). The model's explanatory force rests on this: if realistic lesions appear, they appear because the local rules suffice to produce them.

The implementation is hybrid in Grimm et al.'s sense of an agent-based formulation coupled to another approach. Discrete entities with internal state and conditional rules are coupled to continuous fields governed by partial differential equations, the two advanced on nested timescales and exchanging information through secretion, decay and local sensing.

Principles underlying the design

These principles are embodied in the model's structure rather than in any single submodel.

1. Infection outcome is set by a balance between pro- and anti-inflammatory signalling, resolved locally. A pro-inflammatory axis acts through TNF and a coarse-grained IFN-γ action; an anti-inflammatory axis acts through IL-10 and, separately, through contact suppression by regulatory T cells. Both axes act on the same processes, chiefly macrophage activation and secretion, in opposite directions. Both are computed from local field values and local contacts, so the balance is resolved compartment by compartment, and one granuloma may contain regions on either side of it.

2. One signalling molecule can be both protective and damaging. TNF drives macrophage activation, which killing requires, and induces apoptosis of macrophages and T cells, which adds to tissue damage and so to caseation. The model gives TNF no net sign: both effects are threshold-and-probability tests on the same receptor-occupancy quantity with different thresholds and rate constants, so which predominates depends on the local concentration and the individual cell's receptors (S6.6).

3. Tissue damage accumulates at the compartment scale and feeds back on bacterial growth. Each compartment records a cumulative count of cell deaths. When it crosses a threshold the compartment becomes caseous, kills its occupants, excludes new ones and stops local vascular supply. Caseation in turn arrests bacterial growth nearby, standing in for the hypoxic, nutrient-poor conditions of a necrotic core. The feedback is negative for bacterial replication but positive for tissue destruction, and the tension between these two consequences of one process is central to the model.

4. Bacterial growth depends on where a bacterium is, not on its genotype. Intracellular and extracellular bacteria grow under different rules, and extracellular growth is slowed and then arrested by local crowding and caseation. A bacterium carries no heritable growth phenotype in the default configuration. An arrested, non-replicating population therefore arises from the environment the model generates rather than being specified as a fraction; the replicating and non-replicating labels on containers only report it (S4.3).

5. Adaptive immunity arrives late and ramps in. T cell recruitment is off until a per-simulation onset and then rises to full rate over a further interval (Section 5.7). The model thus has an innate phase, in which bacteria grow against macrophages alone, and an adaptive phase in which effector and regulatory T cells arrive. The transition is an imposed schedule, standing in for priming and trafficking outside the simulated tissue.

6. Spatial structure is emergent. Interactions are confined to small Moore neighbourhoods, so killing, suppression and activation need proximity. Compartment occupancy is limited, so space is contested and dense regions exclude arrivals. Chemokine fields are spatially resolved and steer movement and recruitment. And vascular access is limited to fixed point sources whose oxygen supply stops where cell density is high, so the interior of a dense lesion becomes functionally avascular.

7. Cells of one type are not identical. Several properties are drawn per agent at creation and fixed thereafter: TNF and IL-10 receptor counts, antigen specificity, cytokine-producer status and initial ages. Two cells of a type in the same surroundings can therefore behave differently, which is how the model produces graded rather than switch-like population responses.

Tuneable resolution

A further principle governs construction rather than biology. Mechanisms are implemented at a selectable level of detail, and a study chooses the resolution its question needs. One biological process may be represented either by a coarse-grained algebraic relation or by ordinary differential equations for the underlying molecular species, chosen by configuration rather than by editing the model. For more on this, see our previous publication, "Tuneable Resolution" (doi: 10.1002/wsbm.1270).

The clearest instance is cytokine signalling. By default, receptor occupancy is computed in closed form from the local field value and the cell's receptor count and feeds threshold tests directly (S6). In the fine-grained alternative (M3), per-agent ODE systems track ligand, receptors, internalized complexes and transcription factors. Both target the same model-level decisions and differ in how faithfully they resolve the path from concentration to decision.

Each resolution choice is therefore a module slot; the selection is recorded in Section 7.14 and each option is specified in its M# document. A reader comparing results across studies on this platform must check the module selection before comparing mechanisms.

Scope and deliberate omissions

These are not represented in the default configuration.

Not represented	Status
Antigen presentation and MHC class II trafficking	Module M3, disabled
Lymph node priming and the blood compartment	Module M2 alternative, not selected
Dendritic cell trafficking to the draining lymph node	Absent
Neutrophils	Absent; module M9
Fibrosis, fibroblasts, collagen deposition	Module M4, disabled
Antibiotic pharmacokinetics, pharmacodynamics and resistance	Module M6, disabled
Anti-TNF therapy	Module M7, disabled
CD4 helper subsets beyond the effector and regulatory types	Not represented
$\gamma\delta$ T cells, natural killer cells, B cells	Not represented
Adaptive immune memory and clonal selection	Not represented
Vascular architecture beyond fixed point sources	Not represented
Bacterial nutrient limitation and metabolic flux	Not represented
Diffusible IFN-γ	Coarse-grained as a proximity effect of IFN-γ-producing T cells (Section 1.3, S3.7)
Three-dimensional tissue architecture	Module M5, two-dimensional selected; counts scaled afterwards (Section 1.3, S9)
Cavitation and structural damage to adjacent tissue	Not represented
Dissemination, multiple lesions, host-scale progression	Out of scope (Section 1.1)

Level at which the principles act

Principles 1, 2, 6 and 7 are implemented within submodels and are visible in their rules (S1, S2, S3, S6). Principles 3 and 4 are couplings between submodels and appear in no single rule: caseation is produced by the death rules of several cell types and consumed by the bacterial growth rule and the vascular supply rule. Principle 5 is an imposed schedule, and the scope statement (Section 1.1) a scoping decision.

None of these principles is revised by the model's output. They are design commitments, and the model is not built to test them one by one. What it tests is whether the mechanisms implementing them are jointly sufficient to reproduce the patterns of Section 1.2.

4.2 Emergence

What emerges

Every rule is local: it reads one agent's state, the contents of a small neighbourhood and field values at one or a few compartments. These outcomes therefore emerge from the rules' interaction rather than being imposed.

- **Lesion formation and persistence.** A single infected macrophage at the grid centre is contained, eliminated, or seeds a growing lesion, depending on bacterial growth, macrophage recruitment and killing, and the timing and make-up of the T cell response.
- **Spatial stratification.** No rule places macrophages centrally and lymphocytes peripherally. Stratification comes from several mechanisms together: chronically infected macrophages never move, so they accumulate where infection began; one large agent per compartment caps macrophage density, so the core saturates; T cells move far more often and are drawn inward by chemokines secreted by the core itself; and T cells may share a compartment with a macrophage, so lymphocytes are not shut out of the macrophage region (S2.7, S1.2). Whether a recognizable cuff forms, and how sharp it is, is an output.
- **Caseation and the necrotic core.** Caseation is a compartment-level threshold crossing driven by accumulated cell death. Death concentrates where infected macrophages, cytotoxic killing and TNF-induced apoptosis coincide, so caseous compartments appear first near the bacterial core and spread outward. Core geometry, the caseous fraction and whether caseum is confluent or patchy are emergent.
- **Local hypoxia.** Oxygen is supplied only at fixed sources, supply stops where cell density is high or the source is caseous, and oxygen is not transported laterally, so the interior of a dense lesion becomes depleted as a consequence of its own cellularity. Its extent is an output. By default no behavioural rule reads oxygen, so hypoxia feeds back on nothing (Section 4.7).
- **The growth-arrested bacterial population.** Extracellular growth is arrested by local crowding and caseation, so the size of the arrested population follows from the model's own spatial structure.
- **Macrophage polarization.** The distribution of each macrophage's polarization index, and its spatial arrangement, follow from local TNF and IL-10 exposure and proximity to IFN-γ-producing T cells (S2.4, S9).
- **Infection outcome.** Whether a simulation sterilizes, contains or fails to contain the bacteria is the principal emergent output, against which the patterns of Section 1.2 are assessed.

Illustrative causal chains

Emergence here is mediated by feedback loops that close through the spatial environment, not through any agent. Four are central.

Damage-containment loop, negative for bacterial growth.

$$\text{bacterial growth} \rightarrow \text{infection, recruitment and killing} \rightarrow \text{cell death} \rightarrow \text{caseation} \rightarrow \text{arrested extracellular growth nearby}$$

This loop is self-limiting and is the model's main mechanism of containment without clearance. It works through tissue damage, so containment costs necrosis. The parallel loss of oxygen supply in dense and caseous tissue does not enter it, because bacteria read the caseation count, not oxygen.

Damage loop, positive for lesion size.

$$\text{TNF} \rightarrow \text{macrophage and T cell apoptosis} \rightarrow \text{killing count} \rightarrow \text{caseation} \rightarrow \text{occupant death and exclusion} \rightarrow \text{further deaths nearby}$$

Entrapment reinforces it: a compartment ringed by caseum becomes caseous whatever its own history, so caseous regions tend to fill in rather than stay fenestrated.

Activation loop, negative for bacterial load.

$$\text{bacteria} \rightarrow \text{infected macrophages} \rightarrow \text{chemokines} \rightarrow \text{T cell recruitment and chemotaxis} \rightarrow \text{IFN-}\gamma \text{ producers within two compartments} \rightarrow \text{STAT1} \rightarrow \text{activation} \rightarrow \text{killing}$$

This is the model's mechanism of bacterial control, and it is explicitly spatial: it closes only if IFN-γ-producing T cells come within two compartments of infected macrophages. A lesion that recruits enough T cells but keeps them peripheral will not close it.

Suppression loop, positive for bacterial load.

$$\text{IL-10 from macrophages and Tregs} \rightarrow \text{reduced TNF and chemokine secretion} \rightarrow \text{reduced recruitment and activation} \rightarrow \text{bacterial growth} \rightarrow \text{more infected macrophages}$$

Treg contact suppression acts in parallel, resetting the signalling state of neighbouring macrophages and down-regulating neighbouring effector T cells.

Behaviour is decided by which loops dominate, where and when, and that cannot be read off the parameter values.

Imposed rather than emergent

Grimm et al. call distinguishing emergent from imposed outcomes the most important function of this concept. These are imposed and should not be read as results:

Imposed	Mechanism
Timing of the adaptive response	Per-simulation onset, then a linear ramp (Section 5.7)
Fraction of T cells that are antigen-specific	Per-type probability, drawn at creation
Fraction of Tgam exerting IFN-γ action	Probability, drawn at creation
Baseline resting macrophage population	Recruitment floor that bypasses every signal threshold while the macrophage count is below its initial value (S8.7)
Vascular source positions and number	Set at initialization and never changed
Caseation threshold	Sampled parameter, fixed for a run
Maximum extracellular bacterial density per compartment	Parameter in the growth-reduction term
Chemokine ratios	Fixed multiples of CCL2 (Section 2.4)
Cell lifespans	Type-specific parameters, reset on macrophage activation
Maximum T cell divisions per lineage	Parameter, inherited by daughters

The macrophage floor deserves emphasis. Recruitment is unconditional while the macrophage population is below its initial value, so the resting population does not collapse even without any chemotactic signal. Any inference from the persistence of resting macrophages is an inference about an imposed quantity.

Emergent outputs recorded

The observation layer records the quantities on which these outcomes are judged: cell and bacterial counts by type, state and location, caseation, lesion geometry, cytokine amounts, polarization and mortality and recruitment accounting. S9 gives the full inventory.

Scale of emergence

The model is specified at the scale of cells and compartments and interrogated at the scale of the lesion. That is the point of the approach: rules are stated where the biological evidence is, at what a cell does on meeting a bacterium or a cytokine, and the quantities compared with data are at the scale at which granulomas are measured. Nothing operates at an intermediate scale, and no lesion-level rule exists to be calibrated. The two-dimensional to three-dimensional scaling (S9) acts at the upper scale only and corrects counts, not mechanisms.

4.3 Adaptation

Nature of adaptation

Agents do not optimize. None evaluates alternatives against an objective; there is no fitness, utility or measure of success. Adaptation is present in Grimm et al.'s indirect form only: agents change state and behaviour in response to local conditions through conditional and probabilistic rules fixed in advance. Every adaptive response takes one of three forms:

- a threshold-and-probability test on a locally sensed quantity (Section 2.1);
- a deterministic threshold on an internal quantity, such as a bacterial count, a biomass or a timer;
- a direct consequence of contact, in which one agent's rule writes another's state, as in killing and suppression.

Macrophage responses

Response	Conditioned on	Form
STAT1 pathway	A nearby IFN-γ-producing Tgam	Contact; probabilistic at distance two
NF-κB pathway	Local TNF, neighbourhood bacteria, or the chronically infected or activated state	Threshold-and-probability, or forced by state
STAT3 pathway	Local IL-10	Threshold-and-probability
Activation	STAT1 and NF-κB both set	Deterministic conjunction
Killing or uptake of extracellular bacteria	Local bacterial count; pathway state	Probabilistic; the alternative is infection
Chronic infection; bursting	Intracellular count crossing thresholds	Deterministic
Reversion to resting	No intracellular bacteria and exposure memory lapsed	Deterministic
TNF-induced apoptosis	Local TNF	Threshold-and-probability
Secretion	State and local IL-10	State-dependent rates, IL-10-modulated
Movement	Local chemokines	Weighted random draw
Caseum healing	Activated state; a caseous neighbour	Probabilistic
Exposure memory	Recent encounter with bacteria	Timer refreshed on each encounter

The macrophage is the most responsive entity and the junction through which most signalling acts (S2). Activation needs two pathways with different stimuli, one requiring proximity to a particular T cell phenotype and the other a cytokine level or a bacterial burden; neither alone suffices.

T cell responses

Response	Conditioned on	Form
Proliferation readiness	An exposed macrophage nearby; division count below its cap	Deterministic, then a fixed delay
Proliferation	Delay elapsed; a free slot nearby	Deterministic; deferred if no space
Down-regulation; recovery	Contact with an active Treg; timer	Probabilistic; deterministic
TNF-induced apoptosis	Local TNF	Threshold-and-probability
Fas-mediated killing	Active Tgam; infected or chronically infected target	Probabilistic; target chosen at random
Cytotoxic killing	Active Tcyt; infected or chronically infected target	Probabilistic; for a chronically infected target a second draw decides the bacteriological outcome
Suppression of neighbours (Treg)	Active state	Probabilistic, per neighbour
Movement	Chemokines sensed by the type	Weighted random draw
Secretion	State; producer status	State-dependent rates

Two features limit T cell adaptability. Antigen specificity and producer status are drawn at creation and never revised, so a cell's repertoire is fixed before it meets anything. And the division count is inherited, so a lineage exhausts its proliferative capacity however favourable conditions later become (S3).

Bacterial responses

Response	Conditioned on	Form
Effective growth rate	Location; local density; local caseation	Deterministic modulation of the intrinsic rate
Division	Biomass crossing a threshold	Deterministic
Death in caseum	Local live count, caseous compartments only	Probabilistic

Bacterial adaptation is entirely environmental: a bacterium senses nothing and decides nothing, and its growth rate is computed by its container from the compartment's state (S4). Starvation is implemented but unreachable with the default values (S4.8). The replicating and non-replicating designation is a label recording whether growth has been arrested, not a bacterial response (S4.3). Under M6 bacteria would acquire heritable resistance and efflux states, a genuine within-population adaptation absent by default.

What does not adapt

Vascular sources have no state: whether a source supplies oxygen or recruits is decided each timestep by its compartment and neighbourhood. A source does not remember, deplete or remodel, and there is no angiogenesis or vascular destruction.

Lifespans are fixed at creation and revised only when a macrophage activates, which shortens its remaining life, or is suppressed, which restores it (S2.2). No agent extends its life in response to conditions.

Chemokine secretion ratios are fixed and two of the three chemokine fields are derived from the third (Section 2.4), so differential chemokine regulation cannot be represented.

Caseation reverses only through one narrow mechanism: healing initiated by an adjacent activated macrophage, followed after a delay by halving of the killing count (S5). There is no gradual resolution.

Adaptation and imposed schedule

The onset and ramp of T cell recruitment are not adaptive: under the time criterion used by default they respond to no model state (Section 5.7). The alternative bacterial criterion can delay the onset but cannot advance it, so it too does not make the onset a response to infection. The alternative recruitment option of M2 couples recruitment to a lymph node and blood model in which the response does depend on the infection, and reads no onset at all.

Summary

Adaptation is present throughout in its indirect form, condition-dependent state change by fixed rules, and absent in its direct form: no agent represents, evaluates or pursues an objective. The model's explanatory claim does not need stronger adaptation, since the behaviour of interest comes from fixed local rules interacting with a spatially structured environment.

4.4 Objectives

GranSim has no objectives in Grimm et al.'s sense. No agent has a measure of its own success, and no rule compares alternative actions against one. Every behavioural rule is a conditional or a weighted random draw against the current local environment. Chemotactic movement is the clearest case: an agent does not assess which neighbouring compartment would serve a purpose, but draws a destination with weights given by local chemokine concentrations, and the outcome of a move does not change how later draws are made.

Three features might be mistaken for objectives:

- Chemotactic weighting biases movement toward higher chemokine concentration, generally toward the infection. This is a mechanistic consequence of receptor-mediated locomotion, not a goal; an agent sensing no usable signal moves uniformly at random, with no search behaviour (S1.3).
- The classification of outcomes into sterilizing, controlling and uncontrolled is applied by the analyst afterwards; no agent has it and no rule uses it.
- The recruitment cascade (S8) divides a source's recruitment probability among the categories of cell that could be recruited there, the three T cell types counting as one category. It resembles allocation under a constraint, but nothing is maximized: the division exists to avoid enumerating combinations of cell types.

4.5 Learning

There is no machine learning in *GranSim*. No rule changes during a run, no parameter responds to outcomes, and no agent responds differently to a condition because of prior experience of it.

Several quantities accumulate over an agent's life, and accumulation is not learning: each enters a fixed rule in a fixed way, and what is kept is a count or a timer, not a revised disposition.

Accumulated quantity	Held by	Role
Pathway counters and gene indices	Macrophage	Feed the polarization index only; gate no transition (S2.4)
Exposure memory	Macrophage	Licenses nearby T cell proliferation; delays reversion to resting (S2.10)
Intracellular bacterial count	Macrophage	Crosses fixed thresholds for chronic infection and bursting
Time holding bacteria	Agent	Observation only
Division count	T cell	A budget, inherited by daughters
Interaction counts with macrophages	T cell	Feed the M8 exhaustion rule; disabled by default
Biomass	Bacterium	Crosses the division threshold
Generation time	Bacterium	Enters the M6 mutation probability; disabled by default
Killing count	Compartment	Crosses the caseation threshold; environmental state

The pathway gene indices are the likeliest to be misread, since each is a leaky integrator of recent signalling. But no rule reads them: activation is gated by the discrete flags, and the indices only feed the reported polarization index (S2.4).

Exposure memory is the one accumulated quantity with a substantive behavioural role. A macrophage is flagged as exposed after meeting bacteria, in itself or nearby, and the flag lapses a fixed interval after the last encounter (S2.10). While set, it licenses proliferation of nearby T cells and keeps an infected macrophage that has cleared its bacteria from reverting to resting. It is a decaying trace of recent contact that keeps T cell proliferation licensed briefly after bacteria are cleared, set by contact and cleared by time, with the same effect on every occasion.

Memory in the immunological sense

There is no adaptive immune memory: no memory T cell compartment, clonal expansion of a repertoire, affinity maturation or recall response. A T cell's specificity is drawn at creation from a fixed probability, so the specific repertoire does not change in composition during an infection. This follows from the scope statement (Section 1.1), since memory formation is a host- and lymph-node-scale process. The alternative recruitment option of M2 represents naive, memory and effector subsets in a lymph node and blood model, outside the tissue.

Consequences for interpretation

Apparently purposive behaviour at the lesion scale, such as containment within a bounded region or a lymphocyte cuff around an infected core, cannot be attributed to any agent pursuing it; it comes from non-purposive rules acting in a structured environment, which is the claim the model is built to support. Conversely, the model cannot address questions that need agents with objectives or learning, such as whether cells sensing the state of the whole lesion would behave differently, or whether a different allocation of effort among cell types would improve control. Because no agent has an objective, there is also nothing for an agent to anticipate (Section 4.6).

4.6 Prediction

Absence of explicit prediction

No agent anticipates a future state. Every rule is evaluated against the state of the model at that moment; no agent holds a model of its environment, extrapolates a trend or estimates a future value.

Information available at the point of decision

This is everything an agent rule may read.

Category	Available quantities
Own state	Current state and variables (Section 2.5); age
Own compartment	Field values; caseation status; killing count
Radius 1	Occupants with their types and states; field values; extracellular and intracellular bacterial counts; caseation status and killing counts
Radius 2	Occupant types and states; caseation status; used by the IFN-γ action, interaction counting and entrapment, and by bacterial containers for the growth count
Occupancy	Free slots nearby; presence of a large agent
Global	Current timestep; whether T cell recruitment has begun; the total macrophage count, read by the recruitment floor

The macrophage count is the one whole-grid quantity an agent-level rule reads. It includes dead macrophages not yet removed and the timestep's recruits so far (S8.7). The onset latch could also read the total bacterial count, but only under the bacterial criterion, which is not used by default (S8.10).

Implicit prediction

Models without explicit prediction often contain implicit prediction: a rule tacitly assumes that something observable now is a reliable proxy for something later. Each such rule is a modelling assumption that the model does not test.

- **Chemotaxis** assumes the current gradient shows where the agent should be. For a macrophage that attempts to move rarely, the lag between sensing and arrival is long relative to how fast the infection reorganizes.
- **Exposure memory** assumes recent contact with bacteria predicts continuing local infection; its interval is a parameter, not an inference.
- **Growth arrest by local caseation** assumes caseation indicates current and continuing hypoxia. Caseous compartments within radius 2 stand in for the oxygen and nutrient conditions a bacterium experiences, in place of the oxygen field. This structural proxy for a physiological measurement is the model's main implicit prediction.
- **Stopping vascular supply at high cell density** assumes a dense compartment is one whose vessels are compressed or absent; no vascular state is tracked (S7.4). Under M6 drug supply there is scaled by `vesselActivity` rather than stopped.
- **T cell proliferation readiness** assumes a nearby exposed macrophage indicates antigen. Antigen is not represented by default. Under M3 the MHC class II component adds an antigen-dependent step from the T0 state to the active state, but proliferation still needs an exposed macrophage (M3.6).
- **Division** assumes biomass indicates readiness to divide.

In each case a parameter-mediated proxy substitutes for a quantity the model does not represent or does not read. These are the points where the model's structure, not its parameter values, carries the biological assumption.

Information one timestep old

Some information read by rules is not current, which is easily confused with prediction. Under the hybrid update, an agent's self-transition is invisible to others until the end of the pass, so a neighbour's state read during the next-state pass is the state it held at the start (Section 2.3). The connected-region granuloma is computed at the end of each timestep, so the macrophage release test and an M9 routine read it one timestep late (Section 3.3). Neither is prediction: each is a choice of which snapshot a rule reads.

Prediction at the level of the study

"Prediction" has a second sense in work using this model: running a configuration forward and reporting what happens. That concerns what the modeller does, not any agent, and belongs to the simulation experiments of a study, not to this model description.

4.7 Sensing

Principles

All sensing is local and error-free. An agent reads the quantities of Section 4.6 exactly, with no noise, detection limit or delay beyond the thresholds and probabilities its rules impose. No agent senses beyond two compartments ($40\ \mu$m), and most sensing is limited to radius 1. Agent sensing uses periodic neighbourhoods, except the macrophage exposure test, which is clipped (Section 2.2).

Field sensing at the agent's own compartment

An agent reads a field at its own position, combines it with its own receptor count into an occupancy proxy and applies a threshold test (S6). Every cytokine effect acts this way.

Sensed	By	Governs
TNF	Macrophages, all T cell types	TNF-induced apoptosis
TNF	Macrophages	NF-κB pathway; IL-10 secretion of activated cells
IL-10	Macrophages	STAT3 pathway
IL-10	TNF-secreting cells	Reduction of TNF secretion
IL-10	Macrophages	Reduction of chemokine secretion
O_2	Macrophages	Their own consumption rate
CCL2, CCL5, CXCL9	Macrophages, all T cell types	Gating of chemotaxis

Because occupancy uses each agent's own receptor count, two agents of a type in one compartment sense the same value but may respond differently. Sensing is exact; the response is heterogeneous.

Field sensing across the neighbourhood

For chemotaxis an agent reads the chemokines its type senses at each compartment of its radius-1 neighbourhood, its own included, and weights the candidates accordingly. A chemokine whose value at the agent's own compartment lies outside a sampled window contributes nothing anywhere, so chemotaxis switches off rather than saturating. The types sense different chemokines, but because CCL5 and CXCL9 are fixed multiples of CCL2, every type follows the same gradient, weighted differently (S1.3).

Sensing of neighbours

An agent reads the identity, type and state of occupants of nearby compartments. Every cell–cell interaction rests on this.

Sensed	By	Governs
Macrophages and their states, radius 1	Tgam; Tcyt	Targets for Fas-mediated and cytotoxic killing
Macrophages and their states, radius 2	IFN-γ-producing Tgam	IFN-γ action on STAT1
Agents other than dead, pending-dead and non-cognate T cells, radius 1	Treg	Suppression targets
A macrophage with exposure memory, radius 1	Cognate T cells	Proliferation readiness
A macrophage holding bacteria, radius 1	Macrophages	Exposure memory
Free slots, radius 1	Proliferating cells	Daughter placement
Occupants of the destination	Moving agents	Movement admission
Macrophages, radius 2	T cells	Interaction counting, for observation and M8

Target selection does not discriminate. A Tgam or Tcyt does not pick the most heavily infected macrophage, or the nearest: it draws one macrophage-holding compartment of its neighbourhood uniformly, and attacks only if that macrophage is susceptible. A T cell beside one infected and seven uninfected macrophages has a one-in-eight chance of picking the infected one (S3.8).

Sensing of bacteria

Bacteria are not grid occupants and are sensed as counts (S2, S4):

Sensed	By	Governs
Extracellular bacteria at own compartment, live count and biomass	Macrophages	Killing versus uptake
Extracellular bacteria in the neighbourhood	Macrophages	NF-κB pathway
Live or dead extracellular bacteria, radius 1	Macrophages	Exposure memory
Live bacteria in caseous neighbours	Activated macrophages	Killing within caseum
Own intracellular count	Macrophages	Chronic infection and bursting; uptake probability

Sensed	By	Governs
Caseous compartments, radius 2	Bacterial containers	Effective growth rate
Live count in own container	Bacterial containers	Crowding and death in caseum

Dead bacteria are sensed: the exposure rule counts live and dead extracellular bacteria alike, so a compartment whose bacteria were all killed still registers as infected until they decay.

Sensing by vascular sources

A source reads its own compartment: its caseation status, which stops oxygen supply and recruitment; its cell density, which stops oxygen supply (S7.4); its free slots and occupants, which decide what may be recruited; and its TNF and chemokine values, which set recruitment eligibility and probability (S8). Oxygen supply reads the source compartment's own oxygen value (S7.5).

Sensing in the observation layer

The observation layer reads the whole grid and computes quantities no agent can access: the granuloma definitions, the connected region, the ring scaling and all statistics (S9). It affects dynamics nowhere; the one lesion-scale construct an agent rule reads, the connected region in the release test, feeds statistics only (Section 4.10).

What is not sensed in the current version of *GranSim*

- **Oxygen**, except by macrophages computing their own consumption. No behavioural rule is gated on oxygen; local hypoxia reaches bacteria only through the caseation proxy.
- **Caseation**, except in admission to a compartment, daughter placement, recruitment, the bacterial growth count, death in caseum and the healing rule.
- **The number of bacteria inside another cell.** A T cell choosing a macrophage to kill reads only its state, which distinguishes infected from chronically infected but not the burden within a state.
- **Gradients as such.** An agent reads concentrations and forms weights; it computes no derivative and has no notion of direction.
- **The granuloma**, the agent's distance from it, or any aggregate property beyond the macrophage count read by the recruitment floor.

4.8 Interaction

Modes of interaction

Four modes operate. Two are direct, one entity's rule writing another's state; two are mediated by a shared environment.

Direct contact. One agent modifies another within radius 1, except the IFN-γ action, which reaches radius 2. Effects are written at once to both the current and pending state of the target (Section 2.3).

Interaction	Actor	Target	Effect
IFN-γ action	Active IFN-γ-producing Tgam	Resting or infected macrophage, radius 2	STAT1 flag set or extended (S3.7)
Fas-mediated apoptosis	Active Tgam	Infected or chronically infected macrophage	Killed; half its bacteria, rounded down, dispersed and the rest destroyed, so a one-bacterium cell is sterilized (S3.8, S4.13)
Cytotoxic killing	Active Tcyt	Infected or chronically infected macrophage	Killed; an infected target's bacteria all destroyed; a chronically infected target's either all destroyed or half dispersed and the rest destroyed (S3.9)
Contact suppression	Active Treg	Macrophages and cognate T cells, radius 1	Signalling state reset, or down-regulation (S3.10)
Uptake of bacteria	Macrophage	Extracellular container at own compartment	Bacteria moved into the macrophage
Killing of bacteria	Macrophage	Extracellular container at own compartment	Bacteria killed
Killing in caseum	Activated macrophage	Containers of caseous neighbours	Bacteria killed, probabilistically
Intracellular killing	Activated macrophage	Own container	One bacterium per timestep
Healing of caseum	Activated macrophage	Caseous neighbour	Healing clock started

Targets of the killing interactions are chosen uniformly among neighbouring macrophages (Section 4.7). Suppression is evaluated separately for every candidate in a Treg's neighbourhood with its own draw: eighteen slots are scanned, so at most seventeen agents besides the Treg itself can be suppressed in one timestep (S3.10).

Mediated by fields. Agents secrete into the field at their compartment, the field is transported and decays, and others read it at theirs. The interaction is delayed by transport and attenuated by distance and decay, and is the only way agents influence one another beyond radius 2. Caseous compartments take part as sources with no agent present.

Mediated by recruitment. Secretion raises chemokine and TNF levels at vascular sources, which raises the chance that new agents enter there: an indirect interaction between the cells present and those that arrive, by which the lesion assembles its own infiltrate.

Mediated by space. Slot-limited occupancy (Section 2.3) makes space contested. Agents exclude one another through failed moves, blocked recruitment and deferred proliferation. A caseous compartment excludes all agents unless it heals.

Asymmetries

- A T cell may move into a compartment holding a macrophage, with a probability, but may not be recruited into one (S1.4, S8.9).
- Tregs suppress macrophages and effector T cells but are not themselves suppressed, and chronically infected macrophages are not suppressed either (S2.11, S3.10).
- The bacteriological outcome of a cytotoxic kill depends on the target's state, not the actor's: always sterilizing on an infected target, possibly releasing on a chronically infected one (S3.9).

Interaction network

The directed interactions among entity types are

$$\mathrm{Tgam} \to \mathrm{Mac}, \quad \mathrm{Tcyt} \to \mathrm{Mac}, \quad \mathrm{Treg} \to \{\mathrm{Mac}, \mathrm{Tgam}, \mathrm{Tcyt}\}, \quad \mathrm{Mac} \to \mathrm{Mtb}, \quad \mathrm{Mac} \to \mathrm{compartment}, \quad \mathrm{Mtb} \to \mathrm{Mac},$$

$$\{\mathrm{Mac}, \mathrm{Tgam}, \mathrm{Tcyt}, \mathrm{Treg}, \mathrm{caseum}\} \to \mathrm{fields} \to \{\mathrm{Mac}, \mathrm{Tgam}, \mathrm{Tcyt}, \mathrm{Treg}, \mathrm{sources}\}.$$

Bacteria act on macrophages by infection and bursting. There is no direct T cell–T cell interaction other than Treg suppression, no direct macrophage–macrophage interaction other than through exposure memory and space, and no bacterium–bacterium interaction other than the crowding term in the growth rate.

4.9 Stochasticity

Purpose

Stochasticity serves three purposes:

- **Unresolved mechanism.** A macrophage's probability of killing rather than being infected by a bacterium stands in for phagosomal events that *GranSim* does not represent. Most probabilistic rules are of this kind.
- **Heterogeneity within a type.** Receptor counts, antigen specificity, producer status and initial ages are drawn per agent, so otherwise identical cells behave differently and population responses are graded.
- **Avoiding synchrony artefacts.** List shuffling and uniform target selection keep agents early in a list, or in a fixed direction, from a systematic advantage (Section 3.3).

Inventory

Every probabilistic event is a Bernoulli or uniform draw resolved when it is evaluated (Section 2.1).

Drawn once at creation. Surface receptor counts, uniform within type ranges, for every constructed agent; daughters inherit their parent's counts without a draw (S6.3). Antigen specificity and producer status, for T cells (S3.2). Birth times of initial and recruited cells, spreading their ages (Section 5.3, S8). Daughter lifespans (S3.5). Under M8, bacterial growth rates.

Drawn at initialization. Vascular source positions, stratified (Section 5.2); resting macrophage positions (Section 5.3). The T cell onset is drawn only if its range is non-zero, which it is not by default (Section 5.7).

Drawn each timestep. Movement destinations and T cell admission into occupied compartments (S1); recruitment at each open source slot and the choice of type (S8); TNF-induced apoptosis and pathway activation (S6, S2.4); killing versus uptake and uptake by infected macrophages (S2); Fas-mediated and cytotoxic killing, their targets and, for chronically infected targets, the cytotoxic outcome (S3.8, S3.9); Treg suppression, one draw per candidate (S3.10); caseum healing and killing in caseum (S2.8); death in caseum and the asymmetric division split (S4); selection of bacteria for killing, transfer and dispersal, and of dispersal compartments (S4.13); and the agent-list, source-list and container shuffles (Section 3.3).

Departures from independent draws

Three departures from independent per-event draws affect the variance structure.

- A single draw governs caseum healing across an activated macrophage's whole neighbourhood, so healing initiations within a neighbourhood are perfectly correlated (S2.8).
- A single draw against a cumulative probability selects among mutually exclusive recruitment outcomes, so the type recruited is not decided by independent trials (S8).
- Properties drawn at creation correlate an agent's behaviour over its whole life: a macrophage with a low TNFR1 count is persistently less susceptible to TNF-induced apoptosis.

Treg suppression is not among them: each candidate gets its own draw (S3.10).

Consequences

Replicates with identical parameters diverge, and the divergence is a modelled quantity: aleatory variability across replicates is part of the model's output, and a single realization does not represent a parameter set. Because results depend on the draw sequence, exact reproduction of a trajectory needs the number and order of draws, including draws whose values are never read, such as the receptor draws of constructed bacteria (Section 3.4).

4.10 Collectives

No formal collectives

GranSim defines no collective entity: no granuloma or cluster object, and no group an agent belongs to or can read.

The granuloma is an emergent spatial collective constructed afterwards by the observation layer, in three non-equivalent ways (Section 2.6, S9):

Construction	Basis	Used for
Threshold	Per-compartment cell density or caseation	Area statistics, in-granuloma counts, lesion volume, M6 regional statistics
Connected region	Largest contiguous set of sufficiently cell-occupied compartments	Reported diameters, centre of mass, ring scaling, outside-granuloma release test
Extents	Seven per-compartment predicates	Per-cell-type equivalent diameters and scaling factors

The connected region comes closest to a collective object, with a membership set, centre of mass and diameters. It is nonetheless recomputed from scratch each timestep, has no identity across timesteps and is readable by no agent.

The one feedback from collective to agent

When a macrophage dies holding bacteria, a test against the connected region, approximated by its circumscribing circle, decides whether statistics record a release outside the lesion (S9.8). This affects statistics only, never dynamics, and uses the previous timestep's region (Section 3.3). It is the single point at which a lesion-scale construct re-enters the agent layer.

Aggregations that are not collectives

- **Agent lists** group agents by type for scheduling; they carry no shared state and no agent can read its own.
- **Bacterial containers** group bacteria by location and do carry shared state, their location label, and mediate interaction among their contents through crowding in the growth rate. They are the closest thing to a collective with causal power over members, but they belong to a compartment or host cell, not to a group the bacteria form.
- **The caseous core** is the set of compartments independently satisfying the caseation condition. Entrapment makes a compartment's status depend on its neighbours', so caseous regions fill in, but that is a local rule, not a group effect.
- **The lymphocyte cuff** is an emergent pattern with no representation anywhere.

Interpretation

Grimm et al. ask whether collectives are imposed or emergent. Here they are wholly emergent and wholly observational. The model's claim is that lesion-scale organization arises from local rules; representing the lesion as an entity with its own state would weaken that claim, and the model is built so as not to.

4.11 Observation

Observation runs at the end of each cellular timestep, after states are committed and dead agents removed, so recorded quantities describe the completed state of the timestep (Section 3.3). It does not affect dynamics; the one place an observation construct is read by a rule, the release test, feeds statistics only (Section 4.10).

The platform declares many statistics in groups, and a study chooses which output streams to write and how often. Which output streams are written, and how often, is chosen per study (S9.18); restart states are checkpoints, not observations. Some statistic groups belong to absent or disabled components and stay at zero, which is not a biological zero; S9.19 says which.

The quantities on which assessment depends are:

- **bacterial load**, total and by location, against which outcomes are classified;
- **cell counts** by type and state, over the whole grid and within the granuloma, unscaled and scaled to three dimensions;
- **lesion geometry**: the connected region's size, centre and diameters, and the extents' equivalent diameters;
- **caseation**: caseous and partially caseous compartments, healing, and cells in caseous neighbourhoods;
- **chemical environment**: field totals and maxima, and volume-normalized cytokine concentrations;
- **mortality** of cells and of bacteria, by cause;
- **recruitment**, by type and specificity, with sources blocked by caseation or crowding;
- **macrophage polarization**.

The full inventory, reset codes and definitions are in S9.

Two-dimensional to three-dimensional scaling

Because the model is two-dimensional and the data three-dimensional, granuloma cell and bacterial counts are also reported scaled, by two schemes, alongside the unscaled counts. The extent scheme applies one factor per cell type from the equivalent radius of that type's extent; the ring scheme scales counts annulus by annulus about the connected region's centre, so outer annuli weigh more and a peripherally concentrated cell type is scaled differently from a central one (S9). Both scale counts only: contact frequencies and neighbourhood composition stay those of a two-dimensional lattice (Section 1.3). In a three-dimensional build under M5 the scaled counts equal the unscaled ones (M5).

Observation at the level of the study

Quantities computed outside the simulation from its output, including the outcome classification of Section 1.1, features extracted from agent dumps and comparisons with data, belong to the simulation experiments of a study, not to the model description. The model produces the output streams; what is computed from them is a separate matter.

5. Initialization

GranSim starts from a small focus of infection in otherwise healthy lung tissue, before the onset of adaptive immunity. The initial condition is deliberately minimal: one infected macrophage carrying one intracellular bacterium, a background population of resting macrophages, and a stochastically placed set of vascular sources. No T cells and no extracellular bacteria are present at time zero. Initialization runs only for a new run; a run started from a saved state takes its initial condition from that state.

5.1 Spatial domain and order of placement

The domain is the grid of Section 2.2. Every compartment starts with no agents, a killing count and death count of zero, no caseation, no healing in progress and the granuloma membership flag false, and holds an empty bacterial container labelled replicating. All spatial structure in the initial condition therefore comes from the placement of agents and sources.

Initialization proceeds in this order. Each placement sees the compartments already filled by earlier ones, so the order matters.

1. The T cell recruitment onset is set (Section 5.7).
2. Parameter-file entries are placed, element by element: macrophages, dendritic cells, extracellular bacteria, vascular sources, fibroblasts, myofibroblasts, neutrophils (Section 5.4).
3. Random vascular sources (Section 5.2).
4. Resting macrophages (Section 5.3).
5. Initial cognate Tgam, if `Core_Tcell_Tgam_initCognateNumber` is positive; typically 0.
6. Dendritic cells, fibroblasts and neutrophils, each only if its count is positive, and fibroblasts only under M4; none by default.
7. Parameter-file field and counter values, if any (Section 5.5), and the statistics of timestep 0.

5.2 Vascular sources

Number

The number of random sources is computed at load time from a density and the number of compartments:

$$n_{\mathrm{src}} = \lfloor \texttt{Core_sourceDensity} \times N_{\mathrm{grid}} \rfloor = \lfloor 0.08762770113823116 \times 90000 \rfloor = 7886,$$

so sources occupy about 8.8 per cent of compartments. When the density is read, it takes precedence: the derived count overwrites any value given for `Core_nrSources`. Sources placed from the parameter file are additional to these.

Placement

Sources are placed by stratified sampling, so they are spread evenly rather than clustered as independent uniform draws would allow.

A border of width `Core_sourceBorderSize` $= 1$ is excluded, so no random source lies on the domain edge. The remaining 298×298 region is divided into $m = \lfloor \sqrt{n_{\mathrm{src}}} \rfloor = 88$ subgrids per dimension, $m^2 = 7744$ in all, each about 3.39×3.39 compartments. One source is placed in each subgrid, at a position drawn uniformly in floating-point coordinates, truncated to integers and shifted by the border width. Floating-point coordinates leave no compartment systematically unreachable; the price is that a truncated coordinate occasionally lands in a neighbouring subgrid. A record of compartments already holding a source, including those placed from the parameter file, is kept; if the drawn compartment already holds one, a new position is drawn within the same subgrid, up to twice the subgrid size, after which initialization fails.

The remaining $n_{\mathrm{src}} - m^2 = 142$ sources are placed by rejection sampling over the border-adjusted grid, a position being accepted if it holds no source.

If the derived count is zero, no random source is placed; this is allowed and used for testing. With no random source and no parameter-file source, there is no recruitment and no oxygen supply. Initialization fails if the total number of sources exceeds the number of compartments outside the border.

Source properties are given in Section 2.3.

5.3 Initial macrophages

Number

The number of initial resting macrophages is computed at load time:

$$n_{\mathrm{mac}} = \lfloor \texttt{Core_Mac_initDensity} \times N_{\mathrm{grid}} \rfloor .$$

The density is sampled uniformly per simulation over $[0.00479863597111917,\ 0.0102617209270614]$, at full precision, giving n_{mac} from 431 to 923 across the virtual library. As for sources, a read density overwrites any value given for `Core_Mac_initNumber`.

Placement and properties

Macrophages are placed by rejection sampling over the whole grid, with no border: a position is drawn uniformly and the macrophage placed if the macrophage admission test (Section 2.3) allows it, until the required number is placed. The placement is not stratified, so initial macrophages may cluster. Before placement, the requested number is compared with the number of compartments.

Each initial macrophage is resting, holds no bacteria, has both pro-inflammatory pathway flags clear, and has a birth time drawn uniformly from $\{-\texttt{Core_Mac_maxAge}, \ldots, -2\}$. The population therefore has a spread of ages, so its members neither senesce nor attempt to move in synchrony, since movement opportunity is tested on age (S1.2). Receptor counts are drawn as for every agent (Section 2.3).

Relation to the recruitment floor

The initial macrophage count is also read throughout the run by macrophage recruitment, which recruits unconditionally while the macrophage count is below it (S8.7). The initial number therefore sets not only the starting population but a floor maintained for the whole run, an imposed rather than emergent quantity (Section 4.2).

5.4 Initial infection

Seeding

The infection is seeded by one infected macrophage placed from the parameter file at the grid centre, $(\lfloor n_x/2 \rfloor, \lfloor n_y/2 \rfloor)$, carrying one intracellular bacterium with the initial biomass and the intracellular growth rate (Section 2.5). It is created infected, with both pro-inflammatory pathway flags clear and exposure memory set, and with a birth time drawn as for the resting macrophages, so it is of random age rather than newly arrived. The whole course of a run proceeds from this one bacterium.

Placement from the parameter file

The parameter file's Init section can place agents, bacteria and sources at given positions. A position is given as absolute coordinates, as an offset from the grid centre, or as the keyword `center`. Coordinates must lie on the grid both before and after any offset; an invalid position, or one where the entity's admission test refuses it, is fatal.

Element	Creates
`Mac`	An infected macrophage holding one live bacterium, with exposure memory set and a random birth time
`DC`	A resting dendritic cell, with a birth time drawn over the macrophage lifespan, since dendritic cells share the macrophage lifespan and states
`ExtMtb`	One live extracellular bacterium in the compartment's container
`Source`	A vascular source, in addition to the random sources; a duplicate position is fatal
`Fibroblast`, `Myofibroblast`	A fibroblast in the corresponding state, with birth time drawn from $\{-144, \ldots, 0\}$; only under M4, otherwise ignored with a warning and no draw
`Neutrophil`	An active neutrophil, in every configuration (M9.3)

The default parameter file has one entry, a `Mac` at `center`, which is the seeding above. The Init section can also set field and counter values, at one position or across the grid; this is for testing, does not update dependent statistics, and is not used in the default configuration.

Other entities placed at initialization

Dendritic cells enter only at initialization, under the selected recruitment option: their number is `Core_DCellLung_percentOfMacInitNumber` times the initial macrophage count, truncated to an integer before it is tested, 0 by default (M2.13). Dendritic cells are placed by rejection sampling under the macrophage admission predicate, after the resting macrophages and the initial cognate Tgam. Initial cognate Tgam would be created active with birth time 1, without waiting for the recruitment onset. Neutrophils placed at random use their own crowding rule (M9.3).

Consequences of the minimal start

The founding event is a single bacterium, so early extinction is possible: if the seeding macrophage kills its bacterium, or dies in a way that destroys it, the infection is sterilized at once, and the sterilizing outcome class includes trajectories decided in the first few timesteps.

The infection starts inside a cell, so the first extracellular bacteria appear only when the seeding macrophage dies, bursts or is killed. That event governs when the extracellular compartment is first populated, and so when recruitment signals begin to rise.

The grid centre is the only distinguished position in an otherwise homogeneous domain. The lesion forms there, and the connected-region analysis sweeps outward from it (S9), so the lesion and the sweep origin coincide for most of a run. Lesions that move substantially from the centre are possible but not typical.

5.5 Initial field values

All diffusing and derived fields start at zero everywhere. Of the inert fields of Section 2.4, the shed TNFR2, TGF-β, collagen and anti-TNF fields start at zero and stay there, and no antibiotic field exists unless a drug is named. The Init section can set the TGF-β and collagen fields by name, for testing; the default configuration does not.

Because the fields start empty and the only initial secretor is the seeding macrophage, the chemical environment is featureless for the first timesteps. Chemokine and TNF fields grow from a single point source, so recruitment that needs signals above threshold cannot begin until they have developed; until then only the macrophage recruitment floor recruits (Section 5.3).

Oxygen alone starts at a uniform non-zero value at every compartment, source or not:

$$C_{\mathrm{O}_2} = \texttt{O2_bloodConcentration} \times \mathrm{MW}_{\mathrm{O}_2} = 0.004 \times 32000 = 128 \ \mathrm{mg/L},$$

representing tissue in equilibrium with blood before local consumption. The field is initialized and supplied in mg/L (S7.2); the conversion applied in consumption is given in S7.6. Because oxygen is not transported (Section 2.4), this initial value is the only oxygen a compartment away from a source ever holds.

All compartment counters start at zero and no compartment is caseous or entrapped. The Init section can set a killing count directly, to create caseation at time zero for testing; the default configuration does not, so every container starts labelled replicating.

5.6 Stochastic seeding

Seed handling

Random draws come from two engines, one for real-valued and one for integer draws, seeded once at startup from the run's seed. The seed may be given or derived from the clock; in the default configuration it is given, so any run can be reproduced. The seed in use is written to the output directory. A run is thus determined by the executable, the parameter file, the command line and the seed.

What differs across replicates

Replicates sharing a parameter set differ in every quantity drawn at random (Section 4.9). At initialization:

Quantity	Varies across replicates
Vascular source positions	Yes, within the stratification
Resting macrophage positions and ages	Yes
Surface receptor counts	Yes
Seeding macrophage age	Yes
T cell recruitment onset	No in the default configuration: with a zero range it equals the sampled parameter
Numbers of sources and macrophages	No: fixed by the sampled densities
Seeding macrophage position	No: always the grid centre
Field values	No

During the run, divergence compounds through every behavioural draw. The variability is part of the model's output, not noise: a single realization does not represent a parameter set, and replicate ensembles are the unit of comparison. The sampling design and number of replicates belong to the simulation experiments, not to the model description.

5.7 Onset of adaptive immunity

T cell recruitment is off at the start and is switched on once per run, never to be switched off again. Two criteria are implemented, and exactly one is active, chosen by whether `Recruit_Prob_mtbRecEnabled` is positive (S8.10).

Under the time criterion, used by default, the onset t_{rec} is set at the top of initialization. With `Recruit_Prob_timeRecEnabledRange` $= 0$, t_{rec} equals `Recruit_Prob_timeRecEnabled`, sampled over $[3160, 4275]$ timesteps, about 22 to 30 days after infection. The effective onset is $t_{\mathrm{rec}} + 1$. The onset draw, its truncation and the latch that delays the onset by one timestep are specified in S8.10.

Under the bacterial criterion, not used by default, the flag is set when the total bacterial load first exceeds the threshold. Because the ramp counts from t_{rec} under either criterion, this criterion can delay recruitment beyond t_{rec} but not advance it (S8.10). It therefore does not make the onset a response to infection (Section 4.3).

Recruitment does not begin at full rate. Every T cell recruitment probability is multiplied by a factor that rises linearly from zero at t_{rec} to one after `Recruit_Prob_timeMaxRecDelay`, sampled over $[770, 1196]$ timesteps, about 5 to 8 days (S8.10). Full rate is therefore reached between about 27 and 38 days after infection, depending on the sample.

Interpretation

The schedule stands in for priming in the draining lymph node and trafficking to the site of infection, neither of which is simulated in the default configuration. The alternative recruitment option of M2 represents both, making the onset and size of the T cell response outputs rather than inputs.

Three consequences follow. The model has two phases by construction: before the onset the bacteria grow against macrophages alone, so the load at the onset is set by innate mechanisms, and the adaptive phase begins from whatever state they produced. Because the onset and the ramp are sampled independently of the per-type recruitment probabilities, the timing and steepness of the adaptive response vary independently of its strength. And the ramp scales recruitment probability, not the recruitment thresholds: a source whose local signals are below a type's thresholds recruits no T cell of that type whatever the ramp, so adaptive immunity needs both the schedule and enough signal from the lesion, and the second may bind early in the ramp.

Macrophage recruitment, by contrast, is active from the first timestep, subject only to its signal thresholds and the floor (Section 5.3). There is no innate counterpart to the adaptive onset.

6. Input Data

GranSim uses no external input data: no time series or other environmental driver is read during a run. Every input is static. Parameters are read once at startup from an XML parameter file, and a small number of settings are given as command-line options.

Parameters are not input data in the sense of Grimm et al. How they are declared, resolved, validated and derived is described in Section 7.0, and each parameter is tabled, with its value or sampled range, by the document that owns its rule.

7. Submodels

Element 7 opens with the parameter mechanism shared by every submodel (Section 7.0). Each of Sections 7.1 to 7.13 then explains one submodel or coupling at overview level and links the document that specifies it; the specification is in that document, not here. Section 7.14 lists the module slots.

7.0 Parameters

Declaration and addressing

Every parameter is declared once, with its type, name, group, default, units, minimum, maximum and a description. Storage, the file reader and the validity checks are generated from the declarations, so a parameter cannot be read without being declared or declared without being readable.

Parameters are addressed hierarchically. A parameter in group `Core.Mac` named `dCCL2` is written `Core_Mac_dCCL2` in this documentation and appears in the parameter file as an attribute `dCCL2` on an element `Mac` nested within `Core`. The grouping is organizational: a parameter's group does not show whether any active rule reads it.

The parameter file is XML. No other format is accepted.

The reader visits declared parameters only. An attribute or element in the file that matches no declaration is never read, and no warning is issued. Removing a declaration from `params.def` therefore turns the corresponding file entry into an inert annotation, and a misspelled name is treated like any other undeclared entry. A misspelled attribute for a defaulted parameter is not reported: the parameter takes its default, and the misspelled attribute is ignored as an undeclared entry. The leading underscore by which parameter files set an entry aside is a convention of the parameter files, not of the reader. The prefixed name matches no declaration, so the reader skips it.

Resolution

Values are resolved once, at startup:

1. If the parameter file supplies a value, it is used and the parameter is marked as read. An attribute whose value cannot be converted to the parameter's type is treated as absent.
2. Otherwise the declared default is used and the parameter is marked as not read. If no default is declared, startup fails with a missing-parameter error.
3. Computed parameters are derived from the resolved values and the grid dimensions, overwriting whatever steps 1 and 2 gave them (Section 7.0, "Computed parameters").
4. Validation repeats the range checks and then applies the cross-parameter checks.

The read flag has two uses. It decides which of a pair of mutually derivable parameters is primary, and it decides which parameters are written when a parameter file is saved from a running simulation: a parameter not read from file is omitted.

Validity ranges

A range is checked only where the declared minimum and maximum differ. A blank bound counts as 0, so a declaration with a minimum of 0 and a blank maximum is not range-checked, and a value below its minimum is accepted. Where a range is checked, both bounds are inclusive and a value outside them is fatal at startup.

Ranges are validity bounds, not plausibility bounds: a value outside one is meaningless, such as a probability outside [0, 1] or a negative count. They are unrelated to the sampled ranges of the owners' tables, which express biological plausibility and are enforced by the sampling procedure, not by the model.

Cross-parameter constraints

Some relations between parameters are checked at startup in addition to the ranges, and a violation is fatal. Each is a modelling constraint and part of the specification, so each is stated with the rule it protects: the movement bonus factors (S1), the chronic and bursting thresholds of macrophages (S2), the two caseation growth thresholds, as fractions and as computed counts (S4), and the maximum killing count that sizes the killing-count statistics (S9).

Failure at startup

Startup fails before the first timestep on any of:

- a value outside a checked range, or a violated cross-parameter constraint;
- a missing parameter with no declared default;
- a parameter file that cannot be opened or is not XML;
- a timestep combination that violates the synchronization constraints (M1.2);
- an initial entity count exceeding the available compartments, a duplicate source position, or a source that cannot be placed within its retry limit (Section 5.2).

Command-line options and the parameter file

The grid dimension, the simulated duration, the module selections, the random seed, the output streams and their cadences, and the checkpoint and batch settings are command-line options; quantities entering model equations are parameter-file entries. Settings that select among model variants or fix the domain and duration are thus kept apart from the quantities in the equations. The grid dimension is nevertheless read during parameter resolution, because computed parameters depend on it, so a parameter file is not fully meaningful apart from the grid size it was prepared for.

A run resumed from a checkpoint uses the parameter file and command line of the resuming invocation from the resumption timestep. Where a command-line option conflicts with a value stored in the checkpoint, which prevails depends on the option; Appendix B tables the rule for each.

Declared but unused parameters

The declarations cover the whole platform, including every disabled module. Every declared parameter is read, checked where its range is checked, and stored, so a parameter file must hold entries without defaults for mechanisms the configuration does not use, among them the parameters of the eleven declared drugs (M6.1) and of the lymph-node and blood compartment model (M2).

Parameters that active code reads to no effect, because a disabled factor multiplies them or a sentinel value makes the branch using them unreachable, are indexed in Appendix B, as properties of the configuration. A timestep sentinel that schedules an event beyond the last timestep disables it only for runs of that length (Appendix B).

Computed parameters

Some parameters are derived at load time, after resolution and before validation, so they are checked like supplied ones. A value supplied in the parameter file for a computed parameter is overwritten and has no effect.

Counts from densities. Four initial counts can be derived from densities: macrophages, vascular sources, neutrophils and fibroblasts. The read flags set the direction: if the density was read, the count is computed from it; if only the count was read, the density is computed from the count; if neither was read, startup fails. The forward formula and the macrophage and source values are in Sections 5.2 and 5.3; the neutrophil and fibroblast counts compute to 0 (M9, M4). The default parameter file gives `Core_Mac_initNumber` 105 and `Core_nrSources` 50; both are overwritten by the computed values (Sections 5.2 and 5.3).

The reverse derivation, from a count to a density, yields a density of 0 for macrophages, neutrophils and fibroblasts; no rule reads it.

Caseation growth thresholds. The two thresholds by which neighbouring caseation slows and stops extracellular growth are supplied as fractions of a neighbourhood and converted to compartment counts (S4). The conversion raises the neighbourhood width to the number of dimensions, so a parameter file is not transferable between the two- and three-dimensional builds without reconsidering the fractions (M5).

Rate rescaling. The MHC class II rate constants are declared for a five-second molecular timestep and rescaled to the one in use; thirteen further quantities are derived from them. The computation runs whether or not the fine-grained option is selected, exactly once per run and load (M3).

The neutrophil and fibroblast counts and the MHC class II quantities are computed but read by no rule in the default configuration.

Where parameter values are given

Each parameter is tabled, with its value or sampled range and its units, by the document that owns its rule (S1 to S9, M1 to M9). Observation, stopping and statistics settings are tabled in S9, the solver timesteps in M1, and disabled-module values in Appendix B.

In the default configuration 63 distinct parameters are sampled, one value per simulation held fixed for the run; most uniformly, four log-uniformly. No parameter of the bacterial or oxygen submodels is sampled, so variation in bacterial burden across the virtual library arises from the immune response and the environment it produces. The sampling scheme and the number of samples belong to the simulation experiments description.

7.1 Movement and chemotaxis

Cells move so that the lesion can assemble: recruited cells must reach infected tissue, and the arrangement of the granuloma emerges from where they go. A motile agent attempts at most one step per movement opportunity, to a compartment of its Moore neighbourhood or its own. Candidate compartments are weighted by the concentrations of the chemokines the agent's type senses, subject to a sensitivity window evaluated at the agent's own compartment, and a destination is drawn from the resulting distribution. Whether the move succeeds then depends on occupancy and caseation of the destination.

See S1; the window and the per-type sensitivities are in S1.3.

7.2 Macrophages

Macrophages are the host's first line and the bacteria's niche. They have four live states: resting, infected, chronically infected and activated. Transitions are governed by three intracellular signalling pathways at coarse resolution, by intracellular bacterial load crossing fixed thresholds, and by contact suppression. Macrophages take up and kill extracellular bacteria, become infected, and can be activated to kill intracellular bacteria. They die by one of six routes: senescence, TNF-induced apoptosis, bursting, Fas-mediated killing, cytotoxic killing, or caseation or entrapment of their compartment. Each route disposes of the bacteria the cell held differently (Section 7.13).

See S2.

7.3 T cells

T cells carry the adaptive response, which decides whether infected macrophages are cleared or the infection is contained. Three types are represented: Tgam, Tcyt and regulatory T cells. Only a fraction of recruited T cells is antigen-specific; the rest move and die without effector function. With the sampled cognate probabilities, between about one in two hundred and one in ten recruited T cells is antigen-specific, so the responding population is one to two orders of magnitude smaller than the reported T cell count. The effector functions are coarse-grained IFN-γ activation of macrophages within two compartments, Fas-mediated killing of infected macrophages, cytotoxic killing of infected macrophages with two possible outcomes for their bacteria, and contact suppression of neighbouring cells.

See S3.

7.4 Mycobacteria

Bacteria are the pathogen whose burden the model exists to predict. Each is a discrete agent carrying a biomass that grows at a per-bacterium rate and divides asymmetrically above a threshold. Bacteria are intracellular within a phagocyte, replicating extracellular or non-replicating extracellular. The extracellular growth rate is reduced by local crowding and by local caseation, the latter a proxy for hypoxia. Dead bacteria persist for a short interval, during which macrophages can still detect them. Starvation death exists in the code but cannot occur with the default values.

See S4.

7.5 Caseation and entrapment

Caseation represents the necrotic core, where tissue damage accumulates and both host cells and bacteria are constrained. Compartments accumulate a count of cell killings. When the count reaches a threshold the compartment becomes caseated and every remaining occupant is killed at once. A compartment surrounded by caseated neighbours also becomes caseated, by entrapment. Caseation is partly reversible: an activated macrophage can initiate healing of a neighbouring caseated compartment, which halves its killing count after a delay.

See S5.

7.6 Cytokines and chemokines at coarse resolution

Cytokines and chemokines carry signals between cells that are not in contact. At coarse resolution, receptor occupancy is computed in closed form from the local concentration and the agent's receptor count, and is used directly in threshold-and-probability tests for apoptosis and pathway activation. Receptor-mediated uptake is a sink term applied each molecular substep rather than explicit receptor dynamics. Secretion depends on cell state; local IL-10 inhibits macrophage chemokine secretion by a step function and TNF secretion by a logistic function.

The resolution of signalling is module slot M3. Coarse resolution keeps each agent cheap; the fine-grained option replaces the closed-form occupancy and probability tests of each signalling system by per-agent ordinary differential equations, so that receptor binding and pathway activation have their own dynamics. The default configuration uses coarse resolution for all four systems.

Default option: Coarse-grained: receptor occupancy in closed form from the local concentration and the agent's receptor count; pathway activation by threshold-and-probability tests; receptor uptake as sink terms (S6).

See S6 and M3.

7.7 Oxygen and vascular sources

Vascular sources are the lesion's supply of both oxygen and recruits. Oxygen enters at a source by a flux proportional to the difference between a fixed blood concentration and the concentration in the source's compartment, and is consumed by macrophages with saturable kinetics. A source supplies oxygen only while the cell density of its neighbourhood is below 0.8 and its compartment is not caseated. Oxygen is not transported, so its spatial structure arises solely from the geometry of sources, cell density and consumption. In each cellular timestep, supply is integrated over half a cellular timestep and consumption over one cellular timestep (Section 1.3; S7.6).

See S7; the oxygen couplings are listed in S7.8.

7.8 Recruitment

Recruitment brings cells from blood into the tissue and so sets the size and composition of the immune response. Each vascular source evaluates, once per cellular timestep and once per open slot, whether to recruit a cell. A type is eligible when at least one of its chemokine or TNF concentrations at the source exceeds a type-specific threshold, and its recruitment probability is a sum of saturating functions of the excesses. The available probability is divided among the cell categories possible at the source, the T cell types counting as one category. Macrophage recruitment also has a floor that maintains a baseline resting population.

The recruitment scheme is module slot M2. The probabilistic option reads only local signals at the source; the alternative couples recruitment to a lymph-node and blood compartment model, so that the supply of T cells depends on priming outside the lesion. The default configuration uses the probabilistic option.

Default option: Probabilistic: each source recruits from local chemokine and TNF concentrations (S8).

See S8 and M2.

7.9 Diffusion of soluble fields

Soluble mediators spread and decay between cells, which is what lets a few secreting cells shape a whole neighbourhood. Diffusion and degradation of the diffusing fields are performed by an interchangeable solver, which sets the diffusion and molecular timesteps, the boundary condition of each field, and whether a smoothing pass follows.

The solver is module slot M1. The options differ in accuracy, stability and cost: an alternating-direction explicit scheme, whose timestep is bounded by a stability ceiling (M1.4); a spectral scheme (M1.3), also available with estimate planning, which gives identical transforms in every run (`--diffusion 6`); and a Crank-Nicolson implicit scheme for the two-dimensional build only (M1.5). The default configuration uses the spectral solver (`--diffusion 5`).

Default option: Spectral: each diffusing field is advanced by transform once per diffusion step, with degradation folded into the propagator; CCL5 and CXCL9 are assigned from CCL2 (M1.3).

See M1.

7.10 Granuloma boundary definition

No rule refers to a granuloma, so the lesion must be found in the output. Two non-equivalent definitions coexist. The first is a per-compartment threshold on local cell density or caseation, used for area statistics. The second is a connected-region analysis that finds the largest contiguous region of compartments each with a sufficient fraction of occupied neighbours, and reports its centre of mass and two diameters. Neither definition gates vascular supply, which depends on the cell density of each source's neighbourhood and on caseation (Section 7.7).

See S9.

7.11 Observation and two-dimensional to three-dimensional scaling

The model is two-dimensional while the data it is compared with are three-dimensional, so cell and bacterial counts within the granuloma are scaled. Two schemes are computed: an extent scheme that applies one factor per cell type from the equivalent radius of that type's extent, and a ring scheme that scales counts annulus by annulus about the connected region's centre. Unscaled counts are also recorded.

See S9.

7.12 Stochastic event resolution

The model's behaviour depends on the sequence of random draws, not only on the probabilities. The conventions are in Section 2.1 and the inventory of draws in Section 4.9. Each submodel document specifies the draws of its rules.

7.13 Summary of submodel couplings

The submodels interact to produce granuloma-level behaviour. The tables summarize the couplings, with the mechanism and the documents that specify it. The causal loops these couplings form are described in Section 4.2.

Couplings through cells and bacteria

Coupling	Mechanism	Specified in
Bacteria to macrophage state	Uptake makes a resting macrophage infected; intracellular load crossing thresholds makes it chronically infected and then bursts it	S2, S4

Coupling	Mechanism	Specified in
Macrophage state to bacterial fate	Resting macrophages kill or are infected; activated macrophages kill intracellularly and extracellularly; the death route decides whether contained bacteria are destroyed, released or dispersed	S2, S4
Macrophage death route to bacterial distribution	Senescence releases all bacteria into the compartment; apoptosis, Fas-mediated killing and dispersing cytotoxic kills disperse $\lfloor n/2 \rfloor$ of n bacteria across the neighbourhood and destroy the rest, so a cell holding one bacterium is sterilized; clean cytotoxic kills destroy all; bursting disperses all	S2, S3, S4
Tgam to macrophage activation	Coarse-grained IFN-γ action sets the STAT1 flag within two compartments; this is the sole source of STAT1	S3, S2
Tcyt and Tgam to macrophage death	Cytotoxic and Fas-mediated killing of infected and chronically infected targets	S3, S2
Treg to macrophage and T cell function	Contact suppression resets macrophage signalling state and down-regulates effector T cells	S3, S2
Macrophage exposure to T cell proliferation	A macrophage with exposure memory in the neighbourhood licenses cognate T cell proliferation	S2, S3
Macrophage to caseum healing	An activated macrophage may initiate healing of an adjacent caseous compartment	S2, S5

Couplings through fields

Coupling	Mechanism	Specified in
Cells and caseum to chemokines	Macrophages and caseous compartments secrete CCL2; CCL5 and CXCL9 are derived from it	S6, S5
Chemokines to movement	Chemotactic weighting of the Moore neighbourhood, with type-specific sensitivities	S6, S1
Chemokines and TNF to recruitment	Per-type thresholds and saturating probability terms evaluated at each source	S6, S8
Macrophages and T cells to TNF	State-dependent secretion, inhibited logistically by local IL-10	S2, S3, S6
TNF to macrophage activation	NF-κB pathway activation by receptor occupancy	S6, S2
TNF to cell death	Apoptosis of macrophages and all T cell types	S6, S2, S3
Macrophages and Tregs to IL-10	State-dependent secretion	S2, S3, S6
IL-10 to secretion	Step-function inhibition of macrophage chemokine secretion, logistic inhibition of TNF secretion	S6
IL-10 to macrophage polarization	STAT3 pathway activation	S6, S2
Cells to oxygen	Consumption by macrophages	S7 (S7.8)

Couplings through compartment state

Coupling	Mechanism	Specified in
Cell death to caseation	Accumulated killing count reaching a threshold	S5
Caseation to cell survival	All occupants killed on reaching the threshold; subsequent entry excluded	S5
Caseation to caseation	A compartment ringed by caseum becomes caseous by entrapment	S5
Caseation and chronic infection to bacterial growth	The combined count within radius two reduces or stops extracellular growth	S5, S4
Caseation to bacterial death	Density-dependent death within caseous compartments	S5, S4
Caseation to chemokine secretion	Caseous compartments secrete at a reduced rate with no cell present	S5, S6
Caseation to vascular delivery	No oxygen supply at caseous sources	S5, S7 (S7.8)
Caseation to recruitment	Caseous sources recruit nothing	S5, S8
Cell density to vascular delivery	No oxygen supply where the density of the source's neighbourhood reaches 0.8	S7 (S7.8)
Occupancy to movement, recruitment and proliferation	Slot limits and the size-class rule constrain all three	S1, S8, S2, S3

Couplings through the schedule

Coupling	Mechanism	Specified in
Adaptive onset to T cell presence	Recruitment gated until a sampled timestep, then ramped	Section 5.7, S8
Macrophage population floor to recruitment	Unconditional recruitment while the count is below its initial value	Section 5.3, S8
Granuloma boundary to bacterial release statistics	The previous timestep's connected region decides whether a release is recorded as outside the lesion	S9

What is not coupled

Five absences close feedbacks that might be expected.

Oxygen gates no behavioural rule. Macrophages consume it and it is reported, but no cell or bacterium reads it. Hypoxia acts on bacteria through the caseation proxy, so the oxygen field and the bacterial growth rule are not coupled, although both represent the same physiology (S7.8).

IL-10 does not reach regulatory T cell suppression. Its only effects are the STAT3 test and the two secretion inhibitions; the suppression probability is fixed per sample (S3, S6).

Bacterial burden influences recruitment only through secretion. There is no direct bacterial signal, so a lesion with many bacteria but few secreting cells recruits poorly.

Caseation does not influence chemotaxis directly. A caseous compartment secretes and so appears as a chemokine source, but no rule reads caseation when computing movement weights; an agent is kept out only by the admission test.

The two granuloma definitions of Section 2.6 are coupled neither to one another nor to the extents. All three are computed independently from the same grid state and may disagree.

7.14 Index of module slots and optional components

GranSim's modules let one platform serve studies that need different resolutions, geometries or treatments. Each slot is specified by its M# document, and Appendix B gives, for each, the controlling parameter or command-line option, the value used and the location of the gating logic. Slots M1 to M3 are explained with the submodels they serve (Sections 7.9, 7.8 and 7.6).

Slot	Component	Options	Default configuration
M1	Diffusion solver	Alternating-direction explicit; spectral; spectral with estimate planning; Crank-Nicolson ADI (two-dimensional only)	Spectral (`--diffusion 5`)
M2	Recruitment scheme	Probabilistic from local signals; lymph-node and blood compartment model	Probabilistic
M3	Molecular resolution	Coarse-grained or fine-grained, chosen independently for TNF/TNFR, IL-10/IL-10R, NF-κB and MHC class II	Coarse-grained for all four
M4	TGF-β and fibrosis	None; coarse-grained; receptor ODEs. Fibroblasts exist under both non-none options (M4.1)	None
M5	Dimensionality	Two-dimensional; three-dimensional	Two-dimensional
M6	Antibiotics and pharmacology	Off; on, with a drug-interaction sub-slot	Off
M7	Immunomodulation	Off; systemic dosing of TNF, IL-10, Etanercept or Infliximab	Off
M8	Optional perturbations	Reinfection; cytokine depletion; T cell exhaustion; Treg induction; growth-rate randomization; post-adaptive intracellular growth factor; neutrophil injection	All off
M9	Neutrophils	Absent; present through any of four creation paths	Absent

M4 TGF-β and fibrosis

TGF-β and fibrosis matter for lesions in which scarring walls off bacteria, but they add an agent type and several fields. With a non-none option, bound TGF-β reduces macrophage killing and TNF secretion, gates T cell proliferation and down-regulates cytotoxic T cells, and fibroblasts are created; when they become collagen-secreting myofibroblasts is specified in M4. The default configuration selects none.

Default option: None: the module's executed expressions leave the trajectory unchanged, and no fibroblast exists (M4.1).

M5 Dimensionality

Dimensionality is chosen when the executable is built. It changes no process, order or timestep, only the geometry every block reads: grid shape, neighbourhood sizes and traversal orders, normalizing constants, the diffusion solver and the observation formulae. The three-dimensional build is less prone to the co-localization artefacts of Section 1.3 but is much more costly. The default configuration uses the two-dimensional build.

Default option: Two-dimensional: a square lattice of n^2 compartments, as the spine and the S# documents specify (M5.1).

M6 Antibiotics and pharmacology

The antibiotic module serves treatment studies. With drugs named, antibiotics are dosed into a whole-body plasma model, delivered at vascular sources, diffuse and bind in tissue and enter macrophages, and kill bacteria with a probability set by the local concentration, modified by efflux and acquired resistance. The order in which drugs are named is part of the configuration. The default configuration names no drug.

Default option: Off: no drug is named, so no antibiotic field or drug state exists (M6.1).

M7 Immunomodulation

The immunomodulation module serves studies of host-directed therapy. With an agent dosed, exogenous TNF or IL-10 is delivered into the existing fields, and the anti-TNF antibodies Etanercept and Infliximab diffuse, decay and bind soluble and membrane TNF. The default configuration doses nothing.

Default option: Off: no agent is dosed; the antibody fields are allocated and stay zero (M7.1).

M8 Optional perturbations

M8 is a set of independent perturbations of the baseline, each with its own switch, used to probe how a granuloma responds to a change in conditions. Any combination may be enabled. All are off in the default configuration.

Default option: All off: each perturbation has a nil effect and consumes no random draw (M8.1).

M9 Neutrophils

Neutrophils are a small agent type that moves by its own chemotactic rule, kills or takes up extracellular bacteria, secretes TNF and may stand in caseum. No single switch selects them: they are held absent by closing each of four creation paths. No neutrophil exists in the default configuration.

Default option: Absent: no neutrophil is created, but under the probabilistic recruitment option the neutrophil recruitment test still runs and consumes one draw at every open source slot (M9.1).

Components absent from the default configuration

Component	Status
Dendritic cells and lymphatic trafficking	No dendritic cell is created (Section 5.4; M2)
Fibroblasts, myofibroblasts, collagen	Absent with M4 none
Neutrophils	Absent (M9)
Antibiotic fields and bacterial drug response	Absent with M6 off
Anti-TNF agents and their complexes	Fields allocated and zero with M7 off

Companion documents

Document	Content
S1	Movement and chemotaxis
S2	Macrophages
S3	T cells
S4	Mycobacteria
S5	Caseation and entrapment
S6	Cytokines and chemokines at coarse resolution
S7	Oxygen and vascular sources
S8	Recruitment
S9	Observation, granuloma boundary, and scaling
M1	Diffusion solver module
M2	Recruitment scheme module
M3	Molecular resolution module
M4	TGF-β and fibrosis module
M5	Dimensionality module
M6	Antibiotics and pharmacology module
M7	Immunomodulation module

Document	Content
M8	Optional perturbations module
M9	Neutrophil module
Appendix B	Index of configuration toggles