S9. Observation, granuloma boundary and scaling

S9.1 Scope and schedule

Scope

The observation layer is every computation whose product is a recorded statistic rather than model state. It includes the granuloma constructions on which many statistics depend, and the schemes that scale granuloma counts from two to three dimensions.

S9 also specifies the single point at which an observational construct is read during agent behaviour: the test for whether a dying agent releases bacteria outside the granuloma (S9.8). The consequences of that test are confined to statistics.

Statistics incremented by agent rules, such as killing events, recruitment events and deaths by cause, are declared here with their reset semantics. The conditions under which they are incremented are specified by the documents that own the rules, chiefly S2, S3, S4, S5, S8 and M9.

The observation layer consumes no random draws. With one exception it writes no model state. The exception is the per-compartment membership flag of S9.4.

Position in the cellular timestep

Observation-layer operations occur at these points of the cellular timestep, in this order. The full schedule is in spine Element 3.

Point	Operation
Start of the timestep, immediately after the next grid is synchronised with the current grid	The code-1 reset clears every code-1 statistic (S9.2)
After recruitment, before next-state computation	The code-2 reset clears every code-2 statistic
Next-state computation and extracellular processing	Agent rules increment event statistics. The outside-granuloma test reads the region constructed at the end of the previous timestep (S9.8)
State commitment	Each agent reports its statistics immediately after its next state is committed and before dead agents are removed. Macrophages also report their intracellular containers to the intracellular histograms
After state commitment	The exposed-macrophage count is formed
Statistics block	The statistics-block sub-steps, in their fixed order

Point	Operation
After the statistics block	The T cell recruitment start test and, every T_shuffle, the agent-list shuffle (spine Element 3). Neither belongs to this layer, but the start test reads bacterial totals written in the statistics block (S9.14)
Region construction	The connected region is constructed and the region statistics written (S9.6, S9.7)
Neutrophil injection (M8.9)	Reads the region just constructed
Ring scaling	The granuloma count groups are written (S9.9, S9.10), reading the region constructed in the same timestep
MHC class II statistics	The MHC class II statistics are written (S9.15)
Last	The wall-clock duration of the timestep is recorded

An agent whose committed state is dead still reports on the timestep of its death. Whether a dead agent contributes to its type's state-indexed counts is set by that type's statistics method (S2, S3, M9).

The statistics block performs these sub-steps in fixed order:

1. extents (S9.5);
2. intracellular and total grids;
3. the grid sweep (S9.4, S9.11 to S9.15);
4. drug-concentration statistics;
5. other statistics (S9.11, S9.14, S9.15);
6. caseation statistics (S9.13).

The extents come first because later sub-steps read them. In particular, the lesion volume used for concentration conversion is derived from the cell-density extent (S9.11).

No sub-step is conditional. The intracellular-and-total-grid and drug-concentration sub-steps run every timestep over the antibiotic field list, which is empty by default, so their effect is nil. The MHC class II statistics step, after the ring scaling, is likewise unconditional: it runs every timestep and averages macrophage MHC class II pathway state, which is inert by default (S9.15).

The ring scheme reads the current timestep's region, whereas the outside-granuloma tests read the previous timestep's. The consumers of the same construct are therefore one timestep apart.

Initialisation

At initialisation, every agent created through the simulation's agent factory functions reports its statistics at creation. Initialisation ends with one call of the statistics block, which populates the grid, extent, caseation and other statistics for timestep 0.

Neither region construction nor ring scaling is performed at initialisation. Consequently:

- the region statistics and the granuloma count groups hold their cleared values at timestep 0;
- the outside-granuloma test on timestep 1 reads the default region description, which records no region

(S9.8).

During the run, an agent created by a factory function before the code-2 reset also reports at creation. Its contribution to a code-2 statistic is discarded by that reset and re-established when it reports at state commitment. A macrophage's creation report increments the per-state macrophage count at once, so recruits placed earlier in the recruitment block are counted before the code-2 reset (S8.7).

Reads and writes

Reads:

- the agent lists, and the committed state and state variables of every agent;
- grid slot contents, through the occupancy and agent-type predicates (S9.3);
- the killing count and the entrapment flag, through the caseation predicates (S5.2, S5.4);
- the TNF, IL-10, CCL2, CCL5, CXCL9 and oxygen fields, and the inert TGF-β and collagen fields;
- extracellular containers: live count, stored growth rate and per-bacterium state for the histograms;
- intracellular counts of macrophages, and of neutrophils (M9.9);
- the fibroblast-proliferation flag, false throughout by default;
- the current timestep and the parameters of S9.20.

Writes:

- the statistics object: every statistic declared in S9.2;
- the membership flag at every compartment (S9.4);
- the region description (S9.7);
- the compartment count of each extent (S9.5).

All grid reads are made on the current grid after the diffusion block, so they see the fields as advanced for this timestep.

S9.2 Statistics framework and reset semantics

Declaration

All named statistics are declared in one table, processed at compile time to generate their storage, accessors, serialisation, clearing and traversal. Each declaration is of one kind:

- a group heading, used by the traversal;
- a scalar;
- a fixed-size array, which no statistic uses directly;
- an agent-indexed array of eight entries, in the order macrophage, dendritic cell, Tgam, Tcyt, Treg, Mtb, fibroblast, neutrophil;
- a state-indexed array with one entry per state of a named agent class, including the dead state; for macrophages the entries are resting, infected, chronically infected, activated and dead.

Values are integers, double-precision reals or the platform scalar.

The derived accessors that matter for reading the specification are these. The no-argument accessor of a state-indexed statistic returns the sum over all states, the dead state included. An agent-indexed statistic has an accessor returning the sum over the eight agent types.

State commitment writes each agent's statistics before dead agents are erased. The dead entry of a state-indexed count therefore holds that timestep's deaths until the next code-2 reset.

Reset codes

Every declaration carries a reset code that sets when it is cleared to zero.

Code	Cleared	Accumulates over
0	At construction only; persists across checkpoint and restart	The whole run: a running total
1	At the start of each timestep, before the diffusion block	The whole timestep, from the diffusion block to the end of the observation block
2	After recruitment, before next-state computation	Only the part of the timestep after recruitment

Any increment to a code-2 statistic made during the diffusion block, movement or recruitment is discarded. This is why creation reports made during recruitment do not double-count.

Codes 1 and 2 are both per-timestep quantities. They differ in their accumulation window, not in persistence: a value written at the end of timestep t stays in place, unread by any rule, until it is cleared partway through timestep t + 1.

Output streams are written between timesteps (S9.18), so every stream sees the completed values of both codes. The code matters to a reader of an output file only through the window: an event counted by a code-2 statistic cannot have occurred during movement or recruitment.

Custom members

The statistics object also holds members that are not declared in the table and follow their own clearing rules.

Member	Content	Cleared
Growth-rate, biomass and reserve sample counts	Histogram bin counts set from command-line options; the reserve count is forced to 0 at construction	Never
Intracellular, extracellular and total bacterial histograms	Growth-rate and biomass distributions (S9.14)	Code-2 reset
Per-state intracellular accumulators	Summary statistics of intracellular load by macrophage state	Code-2 reset

Member	Content	Cleared
Intracellular load frequency	Count of macrophages by intracellular count (S9.14)	Code-2 reset
Two polarisation accumulators	Summary statistics of the polarisation index (S9.15)	Code-2 reset
T cell to macrophage interaction counts	Per macrophage state, for the current timestep	Code-2 reset
T cell to exposed-macrophage interaction counts	Per macrophage state	Never; cumulative
Uptake-in-caseous-surroundings histogram	Counts indexed by killing count (S9.13)	Code-1 reset
Wall-clock time per timestep	Seconds of elapsed computation	Code-1 reset, then set at the end of the timestep
The seven extents	Compartment counts (S9.5)	Recomputed at each statistics block

The wall-clock time is excluded from serialisation, so that saved states from two runs with the same executable, parameter file and seed compare exactly. It is the only recorded quantity that differs between such runs.

Of the custom members, these are written to output:

- the uptake-in-caseous-surroundings histogram and both polarisation accumulators, to the general statistics stream;
- the intracellular load frequency, the per-state intracellular accumulators and the bacterial histograms, to the internal bacteria stream;
- the wall-clock time, to the wall-time stream.

The traversal used to enumerate statistics for display does not set the content or the column order of any output stream. The stream writers do (S9.18).

Timing

`TimePerStep` is the elapsed wall-clock time of the timestep in seconds.

S9.3 Compartment predicates

S9.3 specifies the occupancy predicates and densities; spine 2.6 and S7.4 refer to it.

Slot semantics

Each compartment holds its agents in two slots (spine 2.3). An admitted agent is placed in the first empty slot in index order. An agent removed from a compartment leaves its slot empty, and the remaining occupant is not moved to a lower index.

A compartment holding a single agent may therefore hold it in either slot. A compartment whose slot-0 occupant has moved away or died, while its slot-1 occupant remains, holds that occupant in slot 1.

Every predicate examines all slots, so its value does not depend on slot order. The one observation computation that does not examine all slots is the PET-hot score (S9.12).

Predicates

Predicate	Definition	Neighbourhood
occupiedCells(p)	The compartment holds at least one immune cell	Own compartment
occupied(p)	The compartment holds an immune cell or live extracellular bacteria, or is caseated	Own compartment
$\delta(p)$, cell density	The number of compartments q in the periodic Moore neighbourhood of radius 1, including p, with occupied(q), divided by 9	Periodic, radius 1
δ_{cells}(p)	The number of compartments q in the boundary-clipped Moore neighbourhood of radius 1, including p, with occupiedCells(q), divided by 9	Clipped, radius 1
caseated(p), partiallyCaseated(p), percentCaseated(p)	As specified in S5.3 and S5.7; percentCaseated is not bounded above by 1	Own compartment
Agent-type count	The number of slots holding an agent of a given type; a variant also requires a given state	Own compartment
First macrophage-class agent	The occupant of the lowest-indexed slot holding a macrophage-class agent, or none	Own compartment

The denominator of both densities is 9 whatever the number of in-grid compartments in the neighbourhood. This affects only δ_{cells}: the periodic neighbourhood always holds nine compartments, whereas the clipped one holds six at an edge and four at a corner.

Since caseum and free bacteria count as occupied, a destroyed region, whose compartments hold only caseum or extracellular bacteria, still registers $\delta = 1$.

Dendritic cells are macrophage-class agents and report the macrophage agent type. A predicate that must distinguish them also reads the monocyte type. The agent-type count for the macrophage type therefore includes dendritic cells, and the first macrophage-class agent may be a dendritic cell. Both are absent in the default configuration. Neutrophils count as immune cells in every predicate (M9.2).

All predicates are evaluated after states are committed and dead agents removed (S9.1), so every agent they examine is live.

Statistic or construction	Predicates read
Threshold definition and membership flag (S9.4)	δ, caseated, agent-type counts with monocyte type
Cell-density extent (S9.5)	δ, caseated
Other extents (S9.5)	Agent-type counts; first macrophage-class agent with monocyte type; live extracellular count
Connected-region eligibility (S9.6)	occupiedCells, δ_{cells}
Ring tallies (S9.9)	All slots directly; caseated; percentCaseated
PET-hot score (S9.12)	Slots up to the compartment's agent count
Caseation counts (S9.13)	caseated, partiallyCaseated, percentCaseated
Per-agent caseation statistic (S9.13)	The killing count against the threshold directly, not caseated

Treatment of caseation across the observation layer

The connected region (S9.6) excludes caseum, so a lesion with a large caseous core appears smaller under it. More generally, the observation layer treats caseation in four ways:

- the threshold definition (S9.4) and the cell-density extent count caseum as lesion;
- the connected region ignores it, and an enclosed caseous core lies inside its disc only through the outside-granuloma test (S9.8) and the ring tallies (S9.9);
- the PET-hot score treats it as background (S9.12);
- the caseation statistics count it by two definitions: caseation status, including entrapment, and the killing count against the threshold, excluding entrapment (S9.13).

S9.4 Threshold granuloma definition and the membership flag

The predicate

The threshold definition of the granuloma is evaluated once per compartment per timestep, in the grid sweep. A compartment satisfies it when its cell density reaches a threshold or it is caseated:

$$\mathrm{dc}(p) \iff \delta(p) \geq \delta_{\text{thr}} \ \text{ or } \ \mathrm{caseated}(p),$$

where δ is the periodic cell density of S9.3 and caseated includes caseation by entrapment.

The threshold δ_{thr} is set by the command-line option `--area-cell-density-threshold`, default 0.5, which must lie in [0, 1]. It is held by the simulation object and serialised with the state. When a saved state is loaded, the command-line value replaces the saved value, with a warning if they differ.

Because δ takes values in multiples of 1/9, the comparison is satisfied by five or more occupied compartments of nine (5/9 = 0.556) and not by four (4/9 = 0.444). It is an ordinary floating-point comparison, without the significant-figure truncation.

Every agent the sweep sees is live (S9.1). The caseation status it sees is the one established during extracellular processing earlier in the same timestep.

The membership flag

The same sweep assigns the per-compartment membership flag:

```
for each compartment p, in grid order:
    dc <- (cell density(p) >= delta_thr) or caseated(p)
    if fibroblast proliferation has not begun:
        flag(p) <- dc
    // otherwise flag(p) keeps its previous value
```

The flag is false at construction and is first assigned by the statistics call that ends initialisation.

Fibroblast proliferation never begins in the default configuration. The flag is therefore reassigned at every compartment on every timestep and equals the predicate throughout. Under M4 it freezes from the timestep of the first fibroblast proliferation, while the predicate is still evaluated afresh (M4.8).

Statistics derived from the predicate

For every compartment satisfying the predicate, the sweep examines every slot and increments one counter per occupant.

Statistic	Counted occupant
`NrOfMacsInGran`	Macrophage-class agent whose monocyte type is macrophage
`NrOfDCellsInGran`	Macrophage-class agent whose monocyte type is dendritic cell
`NrOfTcellsInGran`	Tgam, Tcyt or Treg
`NrOfFibroblastsInGran`	Fibroblast in the active state
`NrOfMyoFibroblastsInGran`	Fibroblast in the myofibroblast state
`NrOfNeutrophilsInGran`	Neutrophil

For every compartment whose membership flag is set, the sweep increments a second family of counters.

Statistic	Counted occupant
`NrOfMacsInPreFibrosisGran`	Macrophage-class agent whose monocyte type is macrophage
`NrOfFibroblastsInPreFibrosisGran`	Fibroblast in the active state

Statistic	Counted occupant
`NrOfMyoFibroblastsInPreFibrosisGran`	Fibroblast in the myofibroblast state

All these counters carry reset code 1. They count agents, not compartments, over all states of each type.

`NrOfTcellsInGran` includes Tregs, whereas the T cell extent of S9.5 excludes them, so the two T cell measures do not refer to the same population.

In the default configuration the flag equals the predicate everywhere, so the pre-fibrosis macrophage count equals the in-granuloma macrophage count on every timestep, and the fibroblast counts of both families are zero. The two families diverge only under M4, after the flag has frozen.

Under M6 the same sweep partitions each drug's total between granuloma and normal tissue using the predicate as recomputed in the sweep, not the stored flag (M6).

TNF area

The same sweep counts the compartments whose TNF value reaches a second threshold, set by the command-line option `--area-tnf-threshold`, default 0.5 molecules per compartment, which must also lie in [0, 1]:

$$\texttt{AreaTNF} = \left|\{p : C_{\mathrm{TNF}}(p) \geq 0.5\}\right| .$$

The count is taken over the whole grid, is independent of both granuloma definitions, and carries reset code 1.

The threshold is compared with a molecule count and must lie in [0, 1], so the statistic counts compartments holding at least that fraction of one TNF molecule. The two area thresholds share the [0, 1] validation.

Neither area threshold affects the state of the simulation. A change from a loaded state is warned about because it alters the reported classification, not the dynamics.

Equivalence with the cell-density extent

The cell-density extent of S9.5 evaluates the same predicate, with the same threshold read from the same simulation member, at the head of the same statistics call. The only intervening sub-step acts on antibiotic fields alone and does not alter occupancy, caseation or bacterial contents.

The number of compartments satisfying the predicate in this sweep therefore equals the cell-density extent's compartment count on every timestep. The extent's count is the one consumed by other statistics, in particular the lesion volume of S9.11. The sweep itself stores no compartment count.

S9.5 Granuloma extents

Framework

An extent is a named per-compartment predicate together with the number of grid compartments satisfying it. Seven extents are constructed with the statistics object, in a fixed order, which sets their order in the statistics traversal.

Order	Name	Predicate at compartment p
1	`AreaCellDensity`	$\delta(p) \geq \delta_{\mathrm{thr}}$ or caseated(p), as in S9.4
2	`Mtb`	The extracellular container holds at least one live bacterium, or some slot holds a macrophage-class agent or a neutrophil whose intracellular count is positive
3	`Tcell`	The compartment holds a Tgam or a Tcyt
4	`Macrophage`	The compartment holds a macrophage-class agent whose monocyte type is macrophage
5	`Neutrophil`	The compartment holds a neutrophil
6	`DCell`	The compartment holds a macrophage-class agent whose monocyte type is dendritic cell
7	`Fibroblast`	The compartment holds a fibroblast, in either state

Features of the predicates that are easily misread:

- **The Mtb predicate reads live extracellular bacteria only.** A compartment holding only dead extracellular bacteria within their decay interval does not satisfy it, nor does a caseated compartment holding no live bacteria. Its intracellular test includes dendritic cells, being on the macrophage class.
- **The T cell predicate excludes Tregs**, so the T cell extent measures the region occupied by effector T cells.
- **The macrophage predicate distinguishes macrophages from dendritic cells** by monocyte type (S9.3).

Computation

All seven counts are computed together, as the first sub-step of the statistics block:

```
for each extent e: count(e) <- 0
for each compartment p, boundary clipped:
    for each extent e, in construction order:
        if predicate_e(p):
            count(e) <- count(e) + 1
```

The sweep consumes no draws and its result is independent of iteration order. Only the seven counts are stored, and they are serialised with the statistics object.

Derived quantities

Three further quantities are derived from each count n when the statistics are enumerated for output. They are not stored. With π represented by the constant 3.141592653:

$$r = \sqrt{n/\pi}, \qquad d_{\mathrm{mm}} = 2r\,\Delta x_{\mu\mathrm{m}} \times 10^{-3} = 0.04\,r, \qquad \sigma = \max\left(\tfrac{4}{3}r,\ 1\right),$$

where r is the radius in compartments of a disc of area n compartments, d_mm the corresponding diameter in millimetres with $\Delta x_{\mu\mathrm{m}} = 20$, and σ the scaling factor. They appear in the output as `CompartmentCount`, `DiameterMM` and `ScalingFactor`, each prefixed with the extent name.

The scaling factor is the ratio of the volume of a sphere of radius r to the area of a disc of the same radius, floored at 1. The rationale is that scaling from two to three dimensions should never reduce a count. The floor binds only for $n \leq 1$. An extent with no satisfying compartment therefore reports a factor of 1 and a diameter of 0, as the neutrophil, dendritic cell and fibroblast extents do throughout by default.

A measure in square micrometres, equal to 400n, is also defined for each extent. It is not enumerated for output, and no rule reads it.

Consumers

Extent quantity	Consumer
Cell-density compartment count	Lesion volume for the concentration statistics (S9.11); the area stopping criteria (S9.20); granuloma drug statistics under M6
Macrophage, T cell, Mtb, neutrophil, dendritic cell and fibroblast scaling factors	Extent-scaled granuloma counts (S9.10)
Cell-density scaling factor	Output only; no scaled statistic uses it
All counts, diameters and factors	Output

The extents are computed after states are committed and dead agents removed, and before the connected region is constructed. The extent-scaled counts are formed later in the same timestep, so they apply this timestep's factors to counts taken from this timestep's region.

The extents are independent of the connected region and of one another. An extent may include compartments outside the region.

Under M5's three-dimensional build the radius is derived from a volume and every factor is 1 (M5).

S9.6 Connected-region construction

Eligibility

A compartment is eligible for membership in a connected region when it holds at least one immune cell and the fraction of cell-occupied compartments in its boundary-clipped Moore neighbourhood reaches `occupiedNeighborFraction`:

$$\mathrm{eligible}(p) \iff \mathrm{occupiedCells}(p) \ \text{ and } \ \delta_{\mathrm{cells}}(p) \geq \phi, \qquad \phi = 0.66,$$

with δ_{cells} as in S9.3. The comparison is an ordinary floating-point comparison.

Because the compartment itself is counted and is occupied, an interior compartment is eligible when at least five of its eight neighbours are also cell-occupied (6/9 = 0.667; 5/9 = 0.556).

The denominator is nine everywhere. An edge compartment can therefore reach at most 6/9, and is eligible only if all five of its in-grid neighbours are occupied. A corner compartment can reach at most 4/9 and is never eligible.

Caseated compartments hold no cells and are never eligible. Neither is a compartment holding only extracellular bacteria. Eligibility is evaluated on the grid state after states are committed and dead agents removed.

Seed search

Regions are discovered by examining compartments in boundary-clipped Moore shells of increasing radius about the grid centre c. The clipped shell of radius r is the perimeter of the rectangle

$$[\max(c_x - r, 0),\ \ \min(c_x + r, n_x - 1)] \times [\max(c_y - r, 0),\ \ \min(c_y + r, n_y - 1)]\,,$$

traversed counter-clockwise from its corner of least coordinates. The shell of radius 0 is c alone. The perimeter holds every in-grid compartment at Moore distance exactly r from c. Where clipping applies, it may also hold nearer compartments, which have already been examined and are skipped.

The largest radius used is

$$R_{\max} = \max\left(c_x,\ n_x - 1 - c_x,\ c_y,\ n_y - 1 - c_y\right),$$

the smallest radius by which every compartment has been reached. The centre is $c = (\lfloor n_x/2 \rfloor, \lfloor n_y/2 \rfloor)$. For the 300×300 grid, c = (150, 150) and R_max = 150.

Two sets persist across the whole search: the members of all regions constructed so far, and the discards, which are compartments examined and found ineligible. The number of unassigned compartments is

$$U = N_{\text{grid}} - |\text{members}| - |\text{discards}|\,.$$

```
regions  <- empty list
discards <- empty set
U <- N_grid
r <- 0
while r <= R_max and U > size of the largest region (0 if none):
    for each p in clipped shell(c, r), in traversal order:
        if p is a member of any region or p is in discards:
            continue
        if eligible(p):
            append flood fill(p, discards) to regions
        else:
            insert p into discards
        U <- N_grid - total members of regions - |discards|
    r <- r + 1
search radius <- min(r, R_max)

flood fill(seed, discards):
    R <- empty set
    pending <- { seed }
    while pending is not empty:
```

```
        q <- next element of pending
        if q not in R:
            if eligible(q):
                insert q into R; add q to the centre-of-mass accumulator of R
                for each q2 in clipped shell(q, 1):
                    if q2 not in R and q2 not in discards:
                        insert q2 into pending
            else:
                insert q into discards
        remove q from pending
    return R
```

The continuation test is evaluated only at the head of each shell. A shell, once begun, is completed.

Properties

- **Connectivity.** Each region is a maximal set of eligible compartments connected through Moore neighbours, without periodic wrapping. A lesion crossing the domain edge is reported as separate regions on either side, although agent sensing treats the domain as periodic.
- **Order independence.** The membership of each region depends neither on the order in which the flood fill processes pending compartments nor on where in the region the seed lies.
- **Selection.** The reported granuloma is the region with the most compartments. Among regions of equal size, the first constructed is selected: the one whose first-examined compartment lies earliest in shell order about c.
- **Safe early termination.** Every eligible compartment examined so far belongs to a constructed region, so any region not yet constructed lies wholly in unassigned compartments and has at most U members. Since selection is by strict comparison, a later region could at most tie and would not be selected. The selected region is therefore always the one a complete search would select.
- **Determinism.** The construction consumes no draws and is fully determined by the grid state.

S9.7 Region geometry and boundary statistics

Centre of mass

Each member contributes its coordinates to the region's centre-of-mass accumulator exactly once, when it is added. The accumulator keeps a running mean of each coordinate in double precision,

$$\bar{p}_k = \bar{p}_{k-1} + \frac{p_k - \bar{p}_{k-1}}{k},$$

so that after all n members have been added the centre of mass is the coordinate-wise mean

$$\bar{p} = \frac{1}{n} \sum_{p \in R} p,$$

returned as a floating-point position and not truncated to integer coordinates. The running mean may differ from the closed form in the last bits, but is deterministic given the grid state.

An accumulator to which nothing has been added returns a mean of 0 in each coordinate, so the centre of mass recorded when no region exists is (0, 0).

The centre of mass is not a compartment and need not lie within the region, for example when the region is an annulus of cells about an ineligible caseous core.

Diameters

Two diameters are computed for the selected region, in compartments:

$$d_{\max} = 2 \max_{p \in R} \lVert p - \bar{p} \rVert_2, \qquad d_{\mathrm{eq}} = 2\sqrt{n/\pi},$$

where the norm is the Euclidean distance between compartment coordinates, without wrapping, and π is represented by 3.141592653. Both are converted to millimetres by multiplying by 20×10^{-3}.

- d_max is the diameter of the smallest disc centred on $\bar{p}$ containing every member's coordinate point. Because distances are taken between compartment coordinates, it omits the half compartment beyond the outermost members.
- d_eq is the diameter of a disc whose area equals the region's compartment count, computed exactly as the extent diameter of S9.5.

For compact regions the two are close. For elongated or annular regions d_max is the larger. For small compact regions d_eq is the larger.

Statistics written

These statistics are written immediately after the construction. All carry reset code 1.

Statistic	Value
`NrRegions`	Number of regions constructed before the search stopped
`MaxMooreRadius`	The search radius of S9.6
`GranulomaCompartmentCount`	n for the selected region
`MaxGranulomaDiameterMM`	d_max in millimetres
`MinGranulomaDiameterMM`	d_eq in millimetres
`CenterOfMassX`, `CenterOfMassY`	First and second coordinates of $\bar{p}$
`CenterOfMassZ`	The constant 1
`GranulomaCenterOfMassX`, `GranulomaCenterOfMassY`	First and second coordinates of $\bar{p}$, identical to `CenterOfMassX` and `CenterOfMassY`
`GranulomaCenterOfMassZ`	The constant 0

`MinGranulomaDiameterMM` holds d_eq and `MaxGranulomaDiameterMM` holds d_max, so the "Min" value exceeds the "Max" value for small compact regions; a single compartment gives 0.0226 mm against 0 mm.

`NrRegions` equals the number of regions on the grid only when the search runs through every shell up to R_max. When it stops early, regions lying wholly in unexamined shells are not counted.

`MaxMooreRadius` is the loop counter at exit, capped at R_max. When the search stops early it is one greater than the radius of the last shell examined. When the search runs through every shell it equals R_max.

The region description

The region quantities are also held in a region description belonging to the simulation: region count, search radius, compartment count, the centre-of-mass accumulator, and both diameters in compartments and millimetres.

The description is cleared at the start of each construction and then filled. When no compartment is eligible, the region count, compartment count and diameters are zero and the accumulator is empty.

The statistics resets do not clear it. It persists until the next construction and is serialised with the simulation state, so a restart reproduces it exactly. Its readers are:

- the ring scaling of S9.9, later in the same timestep;
- neutrophil injection under M8, later in the same timestep (M8.9);
- the macrophage outside-granuloma test of S9.8, during the next timestep;
- the neutrophil outside-granuloma test at senescence, during the next timestep (M9.2). It is inert by default.

No construction is performed at initialisation. The description therefore holds its default, the no-region state, until the end of timestep 1, and the region statistics are zero in the timestep-0 output.

S9.8 The outside-granuloma test

Definition

A grid position p is classified inside the granuloma when its Euclidean distance from the region's centre of mass does not exceed half the region's maximum diameter:

$$\mathrm{inside}(p) \iff \|p - \bar{p}\|_2 \le \tfrac{1}{2} d_{\max},$$

with $\bar{p}$ and d_max as in S9.7, both in compartments. Distance is taken without wrapping, and the comparison is an ordinary floating-point comparison.

The test reads the region description and nothing else. It does not consult region membership, the threshold definition or any extent.

Relation to the region

Half of d_max is the distance from $\bar{p}$ to the farthest member, so the disc contains every member, and every member is classified inside.

The disc also contains compartments that are not members:

- ineligible compartments enclosed by the region, among them a caseous core;
- gaps in an irregular region;
- compartments between the lobes of an elongated or concave region;
- compartments of other, smaller regions that happen to lie within the disc.

The test therefore departs from region membership in one direction only: it can classify a non-member as inside, and never classifies a member as outside. For an elongated region the excess can be large, since the disc radius is set by the farthest member.

Behaviour when no region exists

When no region exists, the description holds a maximum diameter of 0 and an empty accumulator, whose mean is (0, 0) (S9.7). The test then reduces to $\mathrm{inside}(p) \iff p = (0, 0)$, so every position except the corner compartment (0, 0) is classified outside.

This state holds throughout timestep 1, because no region is constructed at initialisation, and on any timestep following one in which no compartment was eligible. Early in a simulation, before any compartment has six cell-occupied compartments in its neighbourhood, that is every timestep.

Consumer: macrophage death

The test is read by one rule in the default configuration: the update of two statistics recording bacterial release outside the granuloma by dying macrophages.

Statistic	Type	Reset code	Increment
`NrMacsDieOutsideGran`	Integer	1	One per qualifying death
`NrExtMtbOutsideGranFromMacs`	Integer	1	The number of bacteria released by that death

```
update released-bacteria statistics(mac, released):
    if not inside(mac position) and live intracellular count(mac) > 0:
        NrMacsDieOutsideGran        <- NrMacsDieOutsideGran + 1
        NrExtMtbOutsideGranFromMacs <- NrExtMtbOutsideGranFromMacs + released
```

Death route	Update called	Point of the call	released
Senescence	Yes	Before the intracellular bacteria are released to the macrophage's own compartment	The live intracellular count
Killing with dispersal by a T cell (S3.9)	Yes	After dispersal	—

Death route	Update called	Point of the call	released
Death on caseation of the compartment, by threshold crossing or by entrapment (S5)	Yes	After dispersal at fraction 1	As for killing with dispersal
TNF-induced apoptosis, Fas-mediated killing (S2.9, S3.8)	No		
Bursting of a chronically infected macrophage	No		
Clean cytotoxic kill	No		

The two statistics are incremented at senescence only. On the dispersal routes no live bacteria remain when the statistic is updated, so they record no release (S4.13).

Every route that calls the update runs before the region is rebuilt in the same timestep: senescence and T cell killing during next-state computation; caseation during next-state computation, during extracellular processing, or when a daughter cell is discarded into a newly caseated compartment. The region read is therefore always the one constructed at the end of the previous timestep.

Neither statistic is read by any rule, so the test affects no dynamics.

Neutrophils

The neutrophil counterparts, `NrNeutsDieOutsideGran` and `NrExtMtbOutsideGranFromNeuts`, have the same type and reset code and are written at neutrophil senescence (M9.2). They are identically zero by default, since no neutrophil exists. Neutrophil injection under M8 selects sites from positions classified inside, reading the current region (M8.9).

S9.9 Ring-based scaling from two to three dimensions

Purpose and timing

The ring scheme partitions a disc about the region's centre of mass into concentric annuli one compartment wide. It tallies cells, bacteria and caseation in each annulus, and scales each annulus by a radius-dependent factor. It runs once per timestep, after region construction, and reads the region description constructed earlier in the same timestep.

If the description records no region, the scheme returns without writing anything. The granuloma count groups it writes, including the extent-scaled group (S9.10), then hold the zeros set by the code-2 reset of the same timestep.

Disc selection

With $\bar{p}$ and d_max as in S9.7 and $\Delta x = 20$ μm, the disc radius in micrometres is

$$R = \max\left(\tfrac{1}{2} d_{\max}\, \Delta x,\ \Delta x\right).$$

The floor gives a single-compartment region, for which d_max = 0, one annulus.

Every grid compartment is examined. A compartment p is selected when $\Delta x \|p - \bar{p}\|_2 \le R$, with distance taken without wrapping. Selection is by distance alone: a selected compartment need not be a region member, and may lie in another region or in none. The selected compartments are sorted by distance.

The selection is never empty: $\bar{p}$ is a mean of in-grid integer positions, so some compartment lies within $\sqrt{2}/2$ compartments of it, which is less than the floor.

Annuli

Annulus $k = 1, 2, \ldots, K$ has inner radius $r_1 = (k-1)\Delta x$ and outer radius $r_2 = k\Delta x$. It contains the selected compartments with

$$r_1 < \Delta x \|p - \bar{p}\|_2 \le r_2,$$

the first annulus also including distance zero. Annuli are generated while $r_2 \le R$, so $K = \lfloor R/\Delta x \rfloor$.

Selected compartments at distances in $(K\Delta x, R]$ belong to no annulus and are tallied in neither the unscaled nor the scaled totals. Whenever R is not an exact multiple of Δx, this partial outer annulus includes the farthest member of the region. For a single-compartment region, $R = 20\ \mu$m and K = 1: the member and its four face neighbours are tallied, four of them non-members. For a 3×3 block, $R = 28.3\ \mu$m and K = 1: the centre and four face neighbours are tallied, and the four corner members are dropped.

Per-compartment tallies

For each tallied compartment every slot is examined, and these quantities are added to the annulus total.

Quantity	Contribution
Macrophages by state	One per macrophage-class agent of monocyte type macrophage, at its state
Dendritic cells by state	One per macrophage-class agent of monocyte type dendritic cell
Tgam, Tcyt, Treg, neutrophils, fibroblasts by state	One per agent, at its state
Bacteria	The live extracellular count, plus the intracellular count of every macrophage-class agent and neutrophil present
Caseated compartments	One if the compartment is caseated
Fractional caseation	percentCaseated of the compartment, not bounded above by 1

Dead agents have already been removed (S9.1), so the dead-state entries are always zero.

Ring factor and totals

Each annulus total is multiplied by

$$S(r_1, r_2) = \frac{\frac{4}{3}\left(r_2^3 - r_1^3\right)}{\Delta x \left(r_2^2 - r_1^2\right)},$$

the ratio of the volume of the spherical shell between r_1 and r_2 to the volume of the annular slab of thickness Δx between the same radii, with π cancelled. With radii in micrometres the factor is dimensionless. For annulus k it reduces to

$$S_k = \frac{4}{3} \cdot \frac{3k^2 - 3k + 1}{2k - 1},$$

giving $S_1 = 1.333$, $S_2 = 3.111$, $S_3 = 5.067$ and $S_{10} = 19.02$; S_k approaches $2k$ for large k.

The ring scheme and the extent scaling of S9.10 agree for a uniformly filled disc and differ when a quantity is concentrated centrally or peripherally: peripheral cells receive larger factors.

```
if no region: return
select and sort the compartments within R of the centre of mass
unscaled <- 0; scaled <- 0
for k = 1 .. floor(R / dx):
    ring <- sum of tallies over compartments with distance in ((k-1)dx, k dx]
    unscaled <- unscaled + ring
    scaled   <- scaled + S_k * ring
write the unscaled group and the ring-scaled group
write the effective factors: scaled / unscaled, entry by entry
write the extent-scaled group (S9.10)
```

Statistics written

Group	Entries	Value
Granuloma counts, unscaled	Per-state counts for each cell type; `TotMtbUnscaled`; `TotalCaseationUnscaled`; `TotalFractionalCaseationUnscaled`	Sum of annulus totals
Granuloma counts, ring-scaled	The same entries, suffix `ScaledRing`	Sum of factor-weighted annulus totals
Effective ring factors	The same entries, suffix `ScaledRingFactor` or `RingFactor`	Scaled over unscaled, entry by entry

All entries carry reset code 2. The effective factor of an entry whose unscaled value is zero is 0, not 1, because the ratio is formed with a guarded division returning zero for a zero denominator.

Under M5's three-dimensional build the scaled entries equal the unscaled entries, and the effective factors are 1 for non-zero entries and 0 otherwise (M5).

S9.10 Extent-scaled granuloma counts

Computation

The extent-scaled group is written at the end of the ring scheme, in both builds, and only when a region exists. Each entry is the corresponding entry of the unscaled group of S9.9, the tally over the annuli of the disc, multiplied by the scaling factor σ of one extent (S9.5).

Statistic	Unscaled input	Factor taken from
`NrOfMacsScaledExtent`, by state	Macrophage counts	Macrophage extent
`NrOfTgamsScaledExtent`, by state	Tgam counts	T cell extent
`NrOfTcytsScaledExtent`, by state	Tcyt counts	T cell extent
`NrOfTregsScaledExtent`, by state	Treg counts	T cell extent
`NrOfNeutrophilsScaledExtent`, by state	Neutrophil counts	Neutrophil extent
`NrOfDCsScaledExtent`, by state	Dendritic cell counts	Dendritic cell extent
`NrOfFibrosScaledExtent`, by state	Fibroblast counts	Fibroblast extent
`TotMtbScaledExtent`	`TotMtbUnscaled`	Mtb extent

All entries carry reset code 2. The factors are those of the extents computed at the head of the statistics block of the same timestep. No extent-scaled caseation statistic exists.

Properties

Because the input is the unscaled disc tally, the extent-scaled counts inherit the disc selection and the dropped partial annulus of S9.9.

Each factor is computed from the extent's compartment count over the whole grid, not over the disc. These extents are not confined to the lesion:

- **Macrophage extent.** At most one macrophage occupies a compartment, so this extent's count is the number of macrophages on the grid, including the background resting population. The recruitment floor keeps that population near or above its initial size (S8.7), so the macrophage factor stays roughly at or above $(4/3)\sqrt{n_{\mathrm{mac}}/\pi}$ throughout the run: 15.6 at n_mac = 431 and 22.9 at n_mac = 923, the bounds of the sampled initial counts. This holds for a lesion of any size, a single infected cell included.
- **T cell extent.** It includes every compartment holding a Tgam or Tcyt anywhere on the grid, including non-cognate cells recruited at sources remote from the lesion.

Tregs are scaled by the T cell factor, although the T cell predicate excludes them.

The three absent-type extents report a factor of 1, so their extent-scaled counts equal their unscaled counts, which are zero.

S9.11 Field totals, maxima and volume-normalised concentrations

Field totals and maxima

The grid sweep accumulates these quantities over every in-grid compartment of the current grid; the halo of the padded fields is not read. Values are in the field's native units: molecules per compartment for the cytokines and chemokines (spine 2.1), mg/L for oxygen (S7.2).

Statistic	Value	Reset code
`TotTNF`, `TotIL10`, `TotCCL2`, `TotCCL5`, `TotCXCL9`	Sum of the field over the grid	1
`TotO2`	Sum of the oxygen field over the grid	1
`TotTGFBActive`, `TotTGFBInactive`, `collagen`	Sums of the three M4 fields; zero by default	1
`MaxTNFCurrentTimestep`, `MaxCXCL9CurrentTimestep`, `MaxCCL5CurrentTimestep`	Maximum of the field over the grid	1

Each maximum is the running maximum of the field and the cleared value 0. Field values are non-negative, so the result is the grid maximum. No maximum is recorded for CCL2 or IL-10.

The sweep reads the fields as they stand after the diffusion block (S9.1), that is after the second split half-step of the last diffusion step. The derived CCL5 and CXCL9 fields therefore include that half-step's direct secretion as well as their values derived from CCL2. The relation between the three chemokine fields is specified in S6.14.

The oxygen total. `TotO2` is a plain sum of mg/L values over all compartments, not a concentration, so its magnitude scales with grid size. No mechanism raises a compartment above 128 mg/L, so the total begins at its maximum of 128 N_grid (1.152×10^7 on the default grid) and declines almost monotonically, with small rises where a gate reopens or a macrophage leaves an open-gated source (S7.7). Its decline from 128 N_grid is approximately 128 times the number of exhausted compartments, perturbed by intermediate residues. No scaling to three dimensions is applied to it, so it is not comparable between the two builds.

Volume-normalised concentrations

Five concentrations in pg/mL are formed in the other-statistics sub-step, after the sweep. Let n_dc be the compartment count of the cell-density extent (S9.5), V = 8×10^{-12} L the compartment volume, and N_A = 6.02×10^{23}. The lesion volume in millilitres is

$$V_{\text{lesion}} = n_{\text{dc}}\, V \times 10^{3},$$

and for a species with grid total T molecules and molecular weight MW in g/mol,

$$c_{\text{pg/mL}} = \frac{T \cdot \text{MW} \times 10^{12}/N_A}{V_{\text{lesion}}}.$$

Statistic	Total T	MW (g/mol)
`AvgGranTNF`	`TotTNF`	17400
`AvgGranIL10`	`TotIL10`	18600
`AvgGranCCL2`	`TotCCL2`	13000
`AvgGranCCL5`	`TotCCL5`	9990
`AvgGranCXCL9`	`TotCXCL9`	11900

All five carry reset code 1. When n_dc = 0 all five are set to zero; the divisions are performed before this test and their non-finite results discarded.

The numerator is the total over the whole grid, not over the lesion, so molecules outside the cell-density extent are attributed to the lesion volume.

This conversion is the arithmetic of the chemokine-secretion rule of S6.10 at a different scale, and the two agree: for IL-10, 3.090×10^{-8} pg per molecule over 8×10^{-9} mL per compartment is 3.862 pg/mL per molecule per compartment, matching that rule's 3.862×10^{-3} ng/mL. The secretion rule normalises by a single compartment, the observation layer by the lesion.

Because the conversion multiplies by molecular weight, reported mass concentrations of CCL2, CCL5 and CXCL9 stand at 1 : 0.768 : 1.831, the molecular-weight ratios applied to the molecule ratios of S6.14 (1 : 1 : 2 at the default rates), and not at 1 : 1 : 2.

No concentration is formed for oxygen. The cytokine totals themselves are not converted; S9 forms concentrations from them.

S9.12 PET-hot score

Score

Each compartment is assigned an integer score, a proxy for uptake in positron emission tomography. The score is a sum of per-occupant weights.

Occupant	Weight
Activated macrophage	6
Infected macrophage	4
Chronically infected macrophage	1
Resting macrophage	0
Tgam, Tcyt or Treg, in any state	3
Any other agent type	0

Weights for the dead macrophage state are also defined, as zero, but are never applied, because dead agents have been removed before the score is computed. Macrophage-class agents are weighted by state whatever their monocyte type. T cells in the exhausted and induced states of M8 receive the T cell weight. Fibroblasts carry no weight.

Occupants are examined by slot index, from slot 0 up to but excluding N(p), the compartment's occupied-slot count; an empty slot in that range contributes nothing. A compartment whose only occupant is in slot 1 scores 0 (S9.3).

Classification and statistics

A compartment is PET-hot when its score is at least 2, a fixed constant. A single T cell, an infected macrophage or an activated macrophage suffices; a chronically infected macrophage alone does not, but does together with any T cell.

Two statistics are formed in the grid sweep, both with reset code 1.

Statistic	Value
`TotPethot`	Sum of scores over PET-hot compartments
`nrPethotCompartments`	Number of PET-hot compartments

The score depends on cellular composition alone. It reads no field and no caseation state, and a caseated compartment, holding no cells, scores 0.

S9.13 Caseation statistics

Grid counts

The grid sweep forms three caseation counts from the predicates of S9.3 and S5.

Statistic	Value	Reset code
`NrCaseated`	Number of caseated compartments, including those caseated by entrapment	1
`NrPartialCaseated`	Number of partially caseated compartments	1
`TotFractionalCaseation`	Sum of percentCaseated over the grid	1

A compartment caseated by its killing count is also partially caseated. A compartment caseated by entrapment alone with a low killing count is not, because the partial predicate reads the killing count alone (S5.7). The two sets are therefore not nested, and the two reported counts can cross.

`TotFractionalCaseation` sums quantities not bounded above by 1. It measures accumulated killing relative to the threshold, not a number of compartments.

Agents in caseous surroundings

One count per agent type is formed in the caseation sub-step, by examining each agent in that type's list.

```
for each agent a in the type's list:
    k <- number of compartments q in the boundary-clipped Moore
         neighbourhood of radius r about a's position, including it,
         with killing count(q) >= Core_nrKillingsCaseation
    if k / (2r + 1)^2 >= Core_caseationStatFraction:
        increment the type's count
```

Here r is `Core_caseationStatRadius` = 1, truncated to an integer for the neighbourhood, so the denominator is 9. The threshold is read live from the sampled `Core_nrKillingsCaseation` (S5.9).

Statistic	Agent list	Reset code
`MacrophageCaseationRec`	Macrophages	1
`TgamCaseationRec`	Tgam	1
`TcytCaseationRec`	Tcyt	1
`TregCaseationRec`	Treg	1
`FibroblastCaseationRec`	Fibroblasts; zero by default	1
`NeutrophilCaseationRec`	Neutrophils; zero by default	1

- **Killing count, not caseation status.** The test compares the killing count with the threshold directly. A compartment caseated by entrapment alone does not count.
- **Fixed denominator.** The denominator is $(2r+1)^2$ whatever the number of in-grid compartments of the clipped neighbourhood. With the fraction 0.11, a single qualifying compartment satisfies the test (1/9 = 0.111), so the fixed denominator has no effect by default. At other settings, near the boundary the denominator exceeds the number of in-grid compartments examined.
- **Own compartment.** A live agent's own compartment never qualifies: a compartment whose killing count reaches the threshold is caseated, and its occupants are killed and removed before the statistic is computed.

Healing initiations

Statistic	Value	Reset code
`NrHealedCaseation`	Cumulative number of healing clocks started	0

The count is incremented in the activated-macrophage rule whenever a healing clock is started in a caseated compartment of the macrophage's periodic Moore neighbourhood that had none running (S2.8, S5.8). It counts initiations, not completed healings, and no statistic records completions. Initiations begun near the end of a run, too late to complete, are counted, and each healing of a compartment that caseates more than once is

counted separately.

The halving of the killing count at the end of the healing interval increments nothing. In the absence of neutrophils one halving always ends caseation, because the count cannot exceed $\theta + 1$ (S5.2); a compartment is counted again only after it re-caseates. With neutrophils the count can exceed 2θ, and a halved compartment can stay caseated (M9.8).

Bacterial uptake in caseous surroundings

A histogram counts macrophage uptake events by the killing count of the macrophage's compartment. It has `Core_maxNrKillingsCaseation` = 18 bins, indexed from 0. It is a custom member: cleared at the start of each timestep, neither serialised nor traversed (S9.2), but written to the general statistics stream as 18 columns.

An event is recorded:

- by a resting macrophage, before its kill-or-uptake decision, whenever its compartment holds live extracellular bacteria;
- by an infected macrophage, when it takes up extracellular bacteria (S2.6).

Activated macrophages record no event. With r = `Core_caseationStatMooreRadius` = 1, the event is recorded only if

$$\frac{\sum_q K(q)}{9\,\theta} \geq \texttt{Core_caseationStatMoorKillingFraction},$$

where the sum is of killing counts over the boundary-clipped Moore neighbourhood of radius r, including the macrophage's compartment, and θ is the caseation threshold. When recorded, the bin indexed by the killing count of the macrophage's own compartment is incremented. With r = 0 every event is recorded, without the neighbourhood test.

The denominator uses the radius-1 count 9 whatever the radius. At the default radius of 1 it is the full neighbourhood count; at any larger radius the sum covers more than nine compartments against a denominator of 9θ, and the quotient can exceed 1.

The index always lies within the histogram. The compartment of a living macrophage is not caseous, so its killing count is below θ and the largest index written is $\theta - 1$. The constraint `Core_maxNrKillingsCaseation` $\geq$ `Core_nrKillingsCaseation`, enforced at startup (S9.20), therefore secures the bound: a sampled threshold exceeding the declared maximum would index past the end. Neutrophil deaths can raise a killing count without limit (M9.8), but only at compartments a living macrophage cannot occupy, so they cannot overrun the histogram. No clamp is applied to the index.

Bins indexed at or above a sample's own threshold stay zero throughout: for a sample with $\theta = 12$, bins 12 to 17. Fractional caseation statistics are relative to the sampled threshold and so are not on a common scale across samples.

`NBactAct`

`NBactAct` is declared with reset code 0 in the granuloma group, beside the caseation statistics. Read by no rule.

S9.14 Bacterial statistics

Totals

Statistic	Value	Written in	Reset code
`TotExtMtb`	Sum over the grid of the live extracellular count	Grid sweep	1
`TotNonRepExtMtb`	Sum of the live extracellular count over compartments whose container's stored growth rate equals zero	Grid sweep	1
`TotIntMtb`	Agent-indexed; the macrophage entry is the sum of the intracellular count over macrophages that are not dead	State commitment and agent creation	1
`TotMtb`	`TotExtMtb` plus the sum over agent types of `TotIntMtb`	Other-statistics sub-step	1
`FixedMtbCEQ`	Cumulative division counter: set to 1 when the statistics block runs at timestep 0, then incremented by the number of daughters produced in each growth pass, intracellular and extracellular	Other-statistics sub-step; growth pass	0
`TotExtRes`, `TotIntRes`	Totals of the bacterial reserve, extracellular and intracellular	Grid sweep; state commitment	1

"Res" means reserve, not resistance. The reserve starts at zero and is only scaled at division, so both reserve totals are identically zero.

In the default configuration only the macrophage entry of `TotIntMtb` can be non-zero, so `TotMtb` is the bacterial burden of spine 1.1. With neutrophils, neutrophil-held bacteria also enter the total intracellular statistic, and therefore the burden and the bacterial criterion for the onset of T cell recruitment (M9.9). Macrophages dying on the current timestep are excluded from the intracellular sum; their bacteria have already been released, dispersed or destroyed by the death rule (S2.9).

Contributions made at creation do not double-count: recruited macrophages carry no bacteria; the seeded macrophage is created before the first reset; and a reinfecting macrophage under M8 is created before the start-of-timestep reset.

`TotNonRepExtMtb` reads the container's stored growth rate, not its location designation. Extracellular processing (S4.14) updates both only for containers holding live bacteria when they are processed. Containers that are emptied, or that receive bacteria after they have been processed in the same timestep, keep the rate and designation of their last processing. In particular, a macrophage killed by entrapment inside the extracellular pass disperses its bacteria into neighbours, some already processed; those bacteria carry their non-zero intracellular rate and are not grown again that timestep. An arrested compartment receiving such an influx then reports a non-zero mean rate and is excluded from the count, although every bacterium that was in it remains arrested.

The T cell recruitment start test reads `TotExtMtb` and `TotIntMtb` only on its bacterial-threshold branch, which is not taken by default (S8.10).

Bacterial deaths by cause

All of these carry reset code 0 and are cumulative over the run. Each is maintained by the caller, and what it accumulates depends on the routine called.

Statistic	Accumulates	Kind	Writer
`TotUptakeKillRest`	Extracellular bacteria killed by resting macrophages, including the kill-all applied when the compartment's biomass sum is below the uptake amount	Deaths	S2.5
`TotUptakeKillAct`	Extracellular bacteria killed by activated macrophages, in their own compartment and in caseous Moore neighbours	Deaths	S2.8
`TotActMacKill`	Intracellular bacteria killed by activated macrophages	Deaths	S2.8
`TotApopKill`	Bacteria destroyed in the dispersal of TNF-induced apoptosis	Destructions	S2.9
`TotFasKill`	Bacteria destroyed in the dispersal of Fas-mediated killing, about half the burden per kill	Destructions	S3.8
`TotCytoKill`	The whole burden on the clean branch; the dispersal residue on the dispersing branch	Deaths and destructions	S3.9
`TotStarveKill`	Bacteria dying below the biomass threshold; identically zero, starvation being unreachable at the default values	Deaths	S4.8
`TotCaseatKill`	Credited by density-dependent death within caseum and by dispersal on caseation or entrapment; the latter contributes nothing at fraction 1, so the statistic counts deaths within caseum	Deaths	S4.11, S5

Statistic	Accumulates	Kind	Writer
`TotDrugKill`	Bacteria killed by antibiotics; zero by default	Deaths	M6
`TotUptakeKillNeut`	Extracellular bacteria killed by neutrophils; zero by default	Deaths	M9.6

The sum of all categories is not the number of bacteria that have died, and no statistic reports that number. No statistic reports the standing dead population, the number of bacteria dispersed as distinct from destroyed, or generation times and lineage structure.

Bursting increments the burst event count (S9.15), but the bacteria it disperses are credited to no statistic, since nothing dies. An infected macrophage whose intracellular count falls below one adds the result of a kill-all to `TotDrugKill` whether or not drugs are present; intracellular counts are integers, so this adds zero.

Growth-rate extrema and histograms

The six growth-rate extrema, `MaxExtMtbGrowthRate`, `MinExtMtbGrowthRate` and the intracellular and total pairs, carry reset code 0.

Three histograms of bacterial growth rate and biomass, for intracellular, extracellular and all bacteria, are custom members cleared at the code-2 reset, so each is a per-timestep snapshot.

- The intracellular histogram is fed by every macrophage at state commitment, including macrophages dying on that timestep, and by every macrophage at creation.
- The extracellular histogram is fed by every compartment's container in the grid sweep.
- Each has ten growth-rate and ten biomass bins, set by `--rand-growth-samples` and `--biomass-samples` (S9.20). Each is advanced per live bacterium only while its sample count and the container's live count are positive.
- A reserve histogram is filled for each, but never written, because the output writers emit reserve columns only when the reserve sample count is positive, and it is forced to zero (S9.2).

Biomass. Each bacterium's biomass b is normalised as $(b - 0.5)/1.5$ and binned with the top bin as a ceiling. Biomass lies in [0.86, 2) (S4.8), so only bins 2 to 9 are ever populated.

Growth rate. The quantity binned is the bacterium's stored, fully modulated growth rate (S4.5), normalised against the object's own bounds and clamped into bins 0 to 9. The bounds are those of the four randomisation parameters, read as histogram limits whether or not randomisation is on:

- intracellular: `Mtb_Agents_growthRateIntMtbMin` and `Mtb_Agents_growthRateIntMtbMax`, [0.001, 0.005];
- extracellular: `Mtb_Agents_growthRateExtMtbMin` and `Mtb_Agents_growthRateExtMtbMax`, [0.0005, 0.00325];
- total: the lesser minimum and greater maximum of the two, [0.0005, 0.005].

`Mtb_minGrowthBin` and `Mtb_maxGrowthBin` are not the bounds (S9.20).

At the default values the intracellular histogram is a point mass in bin 8, the intracellular rate being unmodulated by default (S4.7). An arrested extracellular bacterium, with a stored rate of zero, is clamped into bin 0, as is any bacterium growing below the lower bound, so extracellular bin 0 combines arrested and slow-growing bacteria and cannot recover the non-replicating count. Under M8 growth-rate randomisation the histograms and extrema report the compounded rates (M8.7).

Intracellular load frequency

A custom member counts infected and chronically infected macrophages by intracellular count. It is fed at state commitment and creation, and cleared at the code-2 reset. It has $\lfloor$ `Core_Mac_nrIntMtbBurstCInf` $\rfloor + 1$ entries, indexed by count, so the number of entries varies with the sampled burst threshold, from 25 to 42. It is serialised but not traversed, and is written only to the internal bacteria stream, which is not selected in the default configuration (S9.18).

A macrophage's intracellular count can exceed the burst threshold before the burst test is evaluated, when many bacteria divide in one timestep. When the count is at least the number of entries, the highest load is recorded outside the histogram's reported range.

A per-state set of summary accumulators of intracellular count is likewise fed at state commitment. It is neither serialised nor traversed.

Extracellular density exceedance

One further count is formed during extracellular processing rather than in the statistics block. `NrCompartmentsWhereLiveExtMtbExceedsMax`, reset code 1, counts compartments whose live extracellular count exceeds `Mtb_growthExtMtbBound` (strict comparison). It is tested for each compartment holding at least one live extracellular bacterium, immediately before that compartment's extracellular growth step and before the caseum death rule (S4). It therefore records the count left by agent behaviour in the same timestep, before growth.

S9.15 Agent-state and macrophage statistics

Agent state counts

Each agent reports to the statistics after its next state is committed and before dead agents are removed (S9.1), and also at creation.

For macrophages the report increments `NrOfAgents` for the macrophage type and `NrOfMacs` at the entry for the committed state. Both carry reset code 2, so creation reports made during recruitment are discarded. A macrophage dying on the current timestep is counted once, in the dead entry; the live macrophage count is the sum over the four live states. The corresponding T cell and neutrophil reports are specified in S3 and M9.9.

The cumulative counts `TotalMac` and the T cell analogues carry reset code 0. For macrophages, `TotalMac` is incremented:

- on each transition into the infected, chronically infected and activated states;
- at initialisation, once per initial resting macrophage and once for the seeded infected macrophage.

It counts entries into a state, not distinct cells: a macrophage that becomes infected, reverts to resting and is infected again is counted twice. Reversion by suppression increments nothing. A reinfecting macrophage under M8 is not counted. Increments for recruited agents are specified in S9.17.

Macrophage function

Statistic	Value	Reset code
`MacNFkB`, `MacStat1`, `MacDeact`	By state, the number of macrophages with the NF-κB flag, the STAT1 flag, or an active suppression	2
`MacM1Phenotype`, `MacM2Phenotype`	Number of macrophages with polarisation index above and below 1	2
`MacApoptosisTNF`	By prior state, TNF-induced apoptoses	0
`MacNFkBTNF`	By state, TNF-induced NF-κB stimulations; non-zero only for resting and infected macrophages	0
`MacActivation`, `MacDeactivation`	By prior state, activations and suppressions	0
`NormMacApopTNF`, `NormMacNFkBTNF`	By state, the cumulative event count divided by the cumulative entries into that state	1
`MacExposedToMtb`	Number of macrophages carrying exposure memory after dead agents are removed	2
`MacrophageWithIntMtb`	Number of macrophages holding live intracellular bacteria at the end of their own next-state computation	1
`BurstMci`	Burst events	0
`TotalIntMtbTime`	Sum of per-macrophage time holding bacteria, added whenever that time is reset	0
`AvgIntMtbTime`	`TotalIntMtbTime` divided by the cumulative infected entries, or 0	1

The pathway, polarisation and suppression counts include macrophages dying on the current timestep.

The polarisation index. A macrophage's index is (G_stat1 + G_nfkb)/(2 G_stat3), with each gene index floored at 1 (S2.4). A macrophage is classed M1 when the index exceeds 1 and M2 when it is below 1. An unstimulated macrophage has all three gene indices equal to 1 and an index of exactly 1, so it is classed as neither.

The NF-κB counter is TNF-attributed. It records stimulations attributable to TNF and excludes those arising from the neighbourhood bacterial count or from the state (S6.5). It is incremented only from the resting and infected handlers.

The normalised quantities. The denominator of each is the cumulative entries into that state for macrophages plus those for dendritic cells, and a guarded division returns 0 for a non-positive denominator. Event counts and denominators are both cumulative, so each is a ratio of cumulative counts: the fraction of cells ever in a state that have died by, or been stimulated by, TNF. It is not a per-timestep rate.

Exposure. `MacExposedToMtb` is also incremented once at initialisation for the seeded macrophage.

`MacrophageWithIntMtb` is counted during next-state computation, so it includes macrophages killed later in the same pass.

Residence time. `AvgIntMtbTime` divides the cumulative residence total by cumulative infected-state entries, including dendritic-cell entries when those cells are present. The denominator counts entries rather than distinct cells, and can include repeated infection of the same cell. The infected handler increments a macrophage's residence counter unconditionally, whereas the activated handler increments it only while the live count is positive. A macrophage that clears its bacteria stays infected for at least five days, and on each of those timesteps banks one timestep into the total, about 720 per cell. The total is banked on every timestep the count is zero, so it receives many small contributions from one cell rather than one equal to its residence.

Polarisation accumulators. Two accumulators are fed with every macrophage's polarisation index at state commitment: the first takes all macrophages, the second those for which any of the three pathway counters is non-zero. Neither is serialised. Both are written to the general statistics stream as count, maximum, minimum, mean and standard deviation, the second under the label `MacNotResting`.

Movement blocking

The macrophage entry of `NrTimesWhereMoveBlockedByBigCell`, reset code 1, is incremented on each movement attempt whose chosen compartment has a free slot and contains a large agent. The test is made on the chosen compartment before the admission predicate (S1.4). When the chosen compartment is the macrophage's own, the macrophage itself is a large agent, so a macrophage that stays in place while its compartment has a free slot is counted as blocked.

A solitary macrophage therefore increments the statistic on about one movement attempt in nine where no usable chemokine signal exists, and with the weighted probability of its own ordinal where one does, which is typically largest near a chemokine source. The statistic therefore includes these self-draws as well as blocking by other macrophages, and the self-draw contribution varies in space. Dynamics are unaffected.

The statistic is declared in the neutrophil group. Its neutrophil entry counts every blocked neutrophil attempt (M9.4).

MHC class II inputs

Three averages are formed in the grid sweep over compartments holding an agent of macrophage type, with reset code 1. Each is a sum divided by the number of such compartments. They are accumulated through the same function that supplies the M3 MHC class II dynamics.

Statistic	Summand per compartment
`AvgTgamMN1`	Number of compartments within Moore distance 1 holding at least one Tgam of any state

Statistic	Summand per compartment
`AvgTgamMN2`	Number of compartments at Moore distance exactly 2 holding at least one Tgam
`AvgExtMtbMN1`	Live plus dead extracellular bacteria within Moore distance 1

All neighbourhoods are periodic, with distance taken before wrapping. The averages are computed every timestep although only M3 reads them. With no macrophage-type agent on the grid the divisor is zero and the recorded values are non-finite. The recruitment floor (S8.7) makes this unlikely by default but does not exclude it.

`AverageMhcExPepOut` and `NrMacCanActivateTcell` are likewise computed every timestep, from MHC class II pathway state that is inert by default. The distance statistics of the same group and `NrOfTcellActivations` are written only when MHC class II dynamics are enabled, and are zero by default (M3.6).

S9.16 T cell statistics

The T cell statistics are written by the T cell rules of S3. Only cognate cells carry producer status, so only cognate cells contribute to the producer counts. The per-state counts distinguish cognate from non-cognate cells for all three types, but the in-granuloma counts (S9.4) and the T cell extent (S9.5) do not, so the T cell extent includes non-cognate cells, which are the majority of T cells (S3.2).

Producer counts (reset code 1). The TNF-producing Tgam count is incremented for a producing Tgam that is not down-regulated. The TNF-producing Tcyt count is incremented for a producing Tcyt only while its suppression timer is set. Both are incremented before the death branches, so a cell dying on the timestep is still counted.

The IFN-γ-producing Tgam count is incremented once per timestep per producing Tgam that reaches its active handler, whether or not a target was found; it counts active and double-producer cells only. A producing Tgam dying of senescence or apoptosis on a timestep is not counted on it, so this count and the TNF-producer counts differ in whether a dying cell is counted. No statistic records the IFN-γ stimulations delivered or received.

Killing (reset code 0). Fas-mediated apoptosis events and cytotoxic kill events are incremented once per successful kill. `TotFasKill` and `TotCytoKill` are bacterial categories (S9.14), not counts of macrophages killed. Each kill also adds the target's banked residence counter to `TotalIntMtbTime`.

T cell apoptosis by TNF (reset code 0). One unindexed counter is incremented by all three T cell types, so output cannot distinguish the type that died, nor cognate from non-cognate losses.

Suppression (reset code 0). The per-state suppression counters are incremented by the targets, one set each for macrophages, Tgam and Tcyt, broken down by the target's state at the time. No statistic records suppression from the Treg's side.

Proliferation. Divisions, deferred proliferations and divisions reaching the maximum are each counted per agent type, both per timestep (reset code 2) and cumulatively (reset code 0). A division is counted once; a deferral once per failed placement attempt; and a division reaching the maximum once per division whose resulting count equals the cap. The deferral count is a count of blocked attempts, not of blocked cells. A division reaching the cap produces two cells at the cap but increments its statistic once, so the number of cells that have exhausted their capacity is twice the reported figure. The per-state stimulated-cell counts count cells whose timer is set, both those waiting for the delay and those blocked for space. Tregs inherit the proliferation variables, which stay at their initial values, so their stimulated-cell counts are always zero.

Interaction counters. The macrophage interaction sweep of S3.6 writes `NrTcellsInteractWithMacs` and `NrTcellsInteractWithExposedMacs`, so both are active by default. Each is advanced at most once per T cell per timestep, so they count distinct cells and are bounded by the T cell population. The per-state interaction counts (S9.2) count interactions, and can exceed the population by a factor of up to 25. T cells in the exhausted and induced states of M8 are reported under their own states in the per-state counts.

S9.17 Recruitment and per-compartment statistics

Per-compartment counters

Per-compartment counters are cumulative over the run, are never reset and are read by no rule. They appear only in grid dumps and the per-compartment state.

- **Death counter.** Counts agents removed from the compartment. No rule reads it.
- **Secretion counter.** Incremented at the end of each macrophage state branch that secretes, on both the NF-κB and coarse paths; never by T cells or caseum. It therefore counts macrophage secretion events, up to 20 per secreting macrophage per cellular timestep. The resting branch increments it only in its NF-κB-positive sub-branch, and the infected branch in both.
- **Recruitment counters.** A general counter, incremented once per recruitment event at that compartment, and one per type, for macrophage, Tgam, Tcyt, Treg and neutrophil. Under option B of M2 only the general counter is incremented, and only for monocytes recruited above the floor (M2.9). The dendritic-cell counter is declared and never incremented.

Recruitment statistics

Recruitment statistics are written by S8 and, under option B, by M2. Their reset codes decide whether a column is a per-timestep count or a running total.

Statistic	Granularity	Reset code
Recruited count, per type, with cognate and non-cognate variants	Per timestep	1
Eligible-source count, per type	Per timestep	1
Source-active count, per type	Per timestep	1
Blocked-source count and blocked-by-large-agent count, macrophage and neutrophil only	Per timestep	1
Caseated-source count	Per timestep	1
Sources where extracellular bacteria exceed the neighbourhood maximum, neutrophil only	Per timestep	1
Recruitment rate, per type, with cognate and non-cognate variants	Recomputed each timestep	1

Statistic	Granularity	Reset code
Cumulative recruited count, per type, with cognate and non-cognate variants	Whole run	0
Cumulative per-state totals, per type	Whole run	0

The cumulative per-state totals are incremented by each construction routine and are the denominators of the normalised macrophage statistics and of `AvgIntMtbTime`.

Eligible against source-active. The eligible-source count is incremented inside a type's possibility predicate, at slot 0, when the type's eligibility test holds. It is written before any admission test and before the caseation gate, so it measures signal availability, not recruitment opportunity. For T cells it is non-zero from the first timestep at which the local signal crosses threshold, well before T cell recruitment can begin. The source-active count is incremented by the cascade, once per slot attempt, when the type is possible: eligibility, admission, the probability draw and, for T cells, the onset flag; a macrophage under the floor (S8.7) needs admission only. A source at which a type is possible on both slots contributes two. Neither count bounds the other. Under option B the source-active count is incremented whenever the source is not full (M2.9).

Blocking. The macrophage blocked-source and blocked-by-large-agent counts are incremented under identical conditions in the default configuration, since within a slot attempt the source is neither caseated nor full, so the two are equal by default. Each counts sources holding a macrophage, at most once per source per timestep. The caseated-source count and the macrophage blocked-source count are disjoint. No blocked-source statistic exists for T cells. The neutrophil blocking counts do not coincide (M9.9).

Rates. After the source loop, one routine sets each recruitment rate to the cumulative recruited count divided by the current timestep, once per timestep, with all, cognate and non-cognate variants, for each of six agent types: eighteen rates. Extra-granuloma recruits are excluded. Each rate is therefore the mean number of recruitment events per timestep over the run to date. The rates are maintained but not written to any stream (S9.18); the per-timestep recruited counts, which are written, give the instantaneous figure.

The rates are computed for every type, including those never recruited. Cognate and non-cognate rates are computed for macrophages, neutrophils and dendritic cells although their cognate counters are never incremented, so eight of the eighteen are identically zero by default: those six, and the neutrophil and dendritic-cell totals. Option B's rate and counter behaviour, including its Treg rates, is specified in M2.9.

S9.18 Output streams

Streams

Each output stream is enabled by command-line options and written to the output directory.

Stream	Enabled by	Row content	In the default configuration
General statistics	`--stats-timespec`	One row per output timestep	Daily

Stream	Enabled by	Row content	In the default configuration
Agent dump	`--agent-dump-file`, `--agent-dump-timespec`	One row per agent	Weekly
Grid dump	`--grid-dump-timespec`, `--grids-to-dump`	One line per dump, one file per grid	Weekly; TNF and killing count
Internal bacteria	`--moi-timespec`	One row per output timestep	Not selected
Wall time	`--walltime`	One row per timestep	Not selected
Molecular tracking, bacterial tracking, T cell to macrophage interaction	Their timespec and file options	One row per tracked entity	Not selected

Saved states and restart checkpoints are not observation streams.

Every statistic is computed on every timestep whatever output is requested; the output options choose only what is written and when.

When output occurs

Output is written after each timestep has completed, and once at timestep 0 after initialisation when a run starts from timestep 0. Every recorded value therefore describes the completed state of its timestep (S9.2). The timestep-0 row reflects the initialisation statistics call, so its region and granuloma count statistics are zero (S9.1).

A stream writes on each timestep in its timespec, and also on the final timestep of the run: the configured last timestep, or the timestep at which a stopping criterion is met. Every stream's last row is therefore the end of the run. No stopping criterion is enabled by default (S9.20).

General statistics content

The general statistics stream carries these blocks.

- Agent counts by type and state, with the T cell producer counts and the cognate, cognate-stimulated, non-cognate and non-cognate-stimulated counts by type and state.
- The unscaled, ring-scaled, effective-ring-factor and extent-scaled granuloma count groups (S9.9, S9.10). Each type's leading effective-factor column is its scaled total divided by its unscaled total, not a sum over states.
- The uptake-in-caseous-surroundings histogram, 18 columns (S9.13).
- Bacteria: neutrophil and macrophage intracellular totals; total intracellular; extracellular; replicating and

non-replicating extracellular; total; `FixedMtbCEQ`; the reserve totals; and the granuloma bacteria groups. The total-bacteria column is computed at writing, as the intracellular plus the extracellular totals.

- Agents in caseous surroundings, deaths by cause and the M6 drug-interaction averages (S9.13, S9.14).
- Field totals; region statistics (S9.7); `AreaTNF` and the extents (S9.4, S9.5).
- Lymph node and blood; recruitment and sources; caseation; cell events; proliferation; polarisation; cumulative state entries and `AvgIntMtbTime`; MHC class II; normalised macrophage rates; anti-TNF; PET-hot; threshold granuloma; outside-granuloma release; extra-granuloma recruitment; T cell exhaustion.

Of the recruitment statistics, the neutrophil `NrSourcesWhereLiveExtMtbExceedsMax` is written for neutrophils only, and the per-timestep recruited count for every type. `TotalNeutrophil`, the neutrophil source-active count, and the cumulative recruited counts and recruitment rates of every type are maintained but not written.

These are not written to any stream selected in the default configuration:

- the five `AvgGran` concentrations;
- the three field maxima;
- the six growth-rate extrema;
- the declared `TotMtb` statistic, whose value the total-bacteria column reproduces.

Agent dump

The agent dump writes, for every agent in the agent lists, its timestep, identifier, agent type, state and position. Bacteria are not dumped. It follows the removal of dead agents, so no dead agent appears; every neutrophil row records the active state, and no intracellular count is written (M9.9).

Grid dump

A grid dump writes the timestep, then every compartment value in row-major order, without the halo.

S9.19 Statistics of absent and disabled components

Several statistic groups are declared, allocated, serialised and written, but stay identically zero in the default configuration, because the components that write them are absent or belong to modules not selected. A column of zeros in an output file is then a property of the configuration, not of the simulated biology. More generally, some outputs are not generated, or take a different form, when the module or component that produces them is not enabled: for example the neutrophil statistics (M9.9) and the three-dimensional extent and ring scaling (M5.12).

Statistics	Component	Reason zero
Neutrophil group; neutrophil, dendritic-cell and fibroblast entries of every granuloma count group; the corresponding in-granuloma counts	Neutrophils (M9) and dendritic cells absent; fibroblasts belong to M4	No such agent is created
Dendritic-cell group	Dendritic cells absent	No dendritic cell is created

Statistics	Component	Reason zero
`TotTGFBActive`, `TotTGFBInactive`, `collagen`, the fibroblast pre-fibrosis counts and fibroblast statistics	M4 (M4.8)	The fields are zero and no fibroblast exists
Anti-TNF group	M7 (M7.10)	Accumulated only when an anti-TNF agent is dosed
Lymph node and blood group; crowded-source count; dendritic-cell stimulation and trafficking events; lymphatic-exit statistics	Option B of M2 (M2.9, M2.10)	Written only by that option
Extra-granuloma recruitment statistics	Option B of M2	Written only by that option
Drug-interaction averages and pre-drug counts; per-drug antibiotic statistics	M6	Accumulated only during drug killing; no antibiotic object is created
MHC class II distance statistics and T cell activation counts	M3 (M3.6)	Written only with MHC class II dynamics enabled

The neutrophil recruitment statistics are the exception within the neutrophil group: the eligible-source, large-agent and exceedance counts are written by the neutrophil predicate at every open slot and are generally non-zero, although no neutrophil exists (M9.9).

Extra-granuloma recruitment statistics. Six statistics are declared, three cumulative (code 0) and three per timestep (code 1). The general statistics stream carries nine columns, the three per-timestep statistics for each T cell type; the cumulative ones are not written. All are zero under the selected recruitment option.

Anti-TNF statistics. All have reset code 1: they are cleared at the start of each timestep and none is cumulative.

Quantities not recorded. No statistic records:

- any receptor occupancy, any secretion or sink amount, or either IL-10 modulation factor (S6), so the degree of IL-10 inhibition in force cannot be recovered from output;
- the IL-10 or CCL2 field maximum;
- the per-source attribution of any field, between macrophages, T cells and caseum;
- the secretion and death counters of S9.17, except through grid dumps;
- healing completions, or the compartments currently healing;
- entrapped compartments as distinct from threshold-caseous ones, so the contribution of entrapment to the reported caseous extent can be recovered only from a grid dump of the killing count;
- caseous compartments by killing count;
- occupants killed by the cascade or by entrapment, as distinct from other deaths of their types.

S9.20 Parameters

S9.20 lists the parameter-file entries, command-line options and fixed quantities whose role is observational. None of them affects model dynamics through this layer.

Parameter-file entries

Identifier	Value or range	Units	Meaning
`occupiedNeighborFraction`	0.66	fraction	Minimum δ_{cells} for connected-region eligibility (S9.6)
`Core_maxNrKillingsCaseation`	18[1]	integer, bins	Bins of the uptake-in-caseous-surroundings histogram (S9.13)
`Core_caseationStatRadius`	1	compartments	Radius of the per-agent caseation statistic, converted to an integer (S9.13)
`Core_caseationStatFraction`	0.11	fraction	Minimum fraction of qualifying compartments for the per-agent statistic (S9.13)
`Core_caseationStatMooreRadius`	1	compartments	Radius of the uptake test; 0 removes the neighbourhood test (S9.13)
`Core_caseationStatMoorKillingFraction`	0.65	fraction	Minimum neighbourhood killing fraction for the uptake test (S9.13)

[1] `Core_maxNrKillingsCaseation` $\geq$ `Core_nrKillingsCaseation` is enforced at startup. It keeps the uptake histogram's index in bounds (S9.13), not only its width.

`occupiedNeighborFraction` is ungrouped, and is range-checked on [0, 1].

Parameters owned elsewhere and read here: `Core_nrKillingsCaseation`, sampled, and `Core_partialCaseationFraction` (S5.9); `Mtb_growthExtMtbBound`, the four growth-rate bounds, `Mtb_Agents_divisionThreshold` and `Mtb_Agents_deathThreshold` (S4.16), which also set the internal-bacteria histogram limits and labels; `Core_Mac_nrIntMtbBurstCInf`, sampled (S2.12), which sets the width of the load frequency.

Not histogram bounds: `Mtb_minGrowthBin` and `Mtb_maxGrowthBin` are read by nothing. The growth-rate histograms are bounded by the four randomisation parameters `Mtb_Agents_growthRateIntMtbMin`, `Mtb_Agents_growthRateIntMtbMax`, `Mtb_Agents_growthRateExtMtbMin` and `Mtb_Agents_growthRateExtMtbMax` (S9.14).

Read but without effect by default: `--reserve-samples` and `Mtb_Agents_maxReserve`. The reserve sample count is zeroed on the statistics object, and the output writers emit reserve columns only when it is positive.

Command-line options

Option	Default	Value in the default configuration	Use
`--area-cell-density-threshold`	0.5	0.5	δ_{thr} of the threshold definition and the cell-density extent (S9.4, S9.5)
`--area-tnf-threshold`	0.5	0.5	Threshold for `AreaTNF` (S9.4)
`--rand-growth-samples`	10	10	Growth-rate bins of the bacterial histograms; unrelated to growth-rate randomisation
`--biomass-samples`	10	10	Biomass bins of the bacterial histograms
`--reserve-samples`	10	10 requested	None; replaced by 0
`--stats-timespec`	Empty	Every 144 timesteps	General statistics stream
`--agent-dump-file`, `--agent-dump-timespec`	Empty	Every 1008 timesteps	Agent dump
`--grid-dump-timespec`, `--grids-to-dump`	Empty	Every 1008 timesteps; the TNF and killing-count fields	Grid dump
`--moi-timespec`, `--walltime`, and the tracking and interaction options	Not set	Not set	Streams not selected
`--stop-on-clearance`	Not set	Not set	Stop when the total bacterial load is zero

Both area thresholds must lie in [0, 1], or startup fails. Both are stored with the simulation state; when a saved state is loaded, the command-line values replace the stored ones, with a warning if they differ. The grid size, set by `--dim`, fixes the search origin and R_max of S9.6 (spine 2.2).

Stopping criteria

The stopping criteria and the clearance option read observation statistics. They are evaluated after the output step of each timestep, so the stopping timestep's row is written to every stream before the run ends (S9.18). Let M be the extracellular total plus the sum over agent types of the intracellular total (S9.14), and n_dc the compartment count of the cell-density extent (S9.5).

Parameter pair	Value	Condition for stopping at timestep t
`mtbStopppingTimeStep`, `mtbStoppingThreshold`	0, 50	t equals the timestep, the timestep is positive, and M < 50
`mtbStopppingTimeStep2`, `mtbStoppingThreshold2`	0, 10^9	t equals the timestep, the timestep is positive, and M > 10^9
`areaCellDensityStopppingTimeStep`, `areaCellDensityStoppingThreshold`	0, 1.8×10^7	t equals the timestep, the timestep is non-zero, and n_dc < 1.8×10^7
`areaCellDensityStopppingTimeStep2`, `areaCellDensityStoppingThreshold2`	0, 3×10^8	t equals the timestep, the timestep is non-zero, and n_dc > 3×10^8
`--stop-on-clearance`	Not set	M = 0

The identifiers are ungrouped, and spelled as shown. Each criterion is tested at a single timestep only. A timestep of zero, or one beyond the final timestep, disables a criterion; all four are disabled by default. The bacterial criteria require a positive timestep and the area criteria a non-zero one, but a negative timestep never equals t, so the difference has no effect.

The area criteria compare a compartment count, at most N_grid = 90000 on the default grid, with thresholds of 1.8×10^7 and 3×10^8. At these thresholds an enabled first criterion is met at its test timestep in every run, and the second is never met. Under M5's three-dimensional build the comparison differs (M5.13).

Fixed structural constants

The compartment edge (20 μm), face area A_{face} (400 μm^2) and volume ($V = 8 \times 10^{-12}$ L), and the two slots per compartment, are spine 2.1 and 2.3 constants. The constants specific to this layer are:

- micrometre-to-millimetre factor 10^{-3}; π represented as 3.141592653;
- extent scaling factor max((4/3)r, 1) (S9.5) and ring factor $S(r_1, r_2)$ (S9.9);
- Moore neighbourhood size 9 in both densities and in the uptake-test denominator;
- Avogadro's number 6.02×10^{23}; molecular weights TNF 17400, IL-10 18600, CCL2 13000, CCL5 9990 and CXCL9 11900 g/mol; unit factors 10^{12} pg/g and 10^3 mL/L (S9.11);
- PET-hot weights 6, 4, 1 and 0 for activated, infected, chronically infected and resting macrophages and 3 for any T cell, and the PET-hot threshold 2 (S9.12);
- MHC class II input radius 2 (S9.15);
- reserve sample count 0, replacing the command-line value;
- guarded division returning 0 for a zero denominator: effective ring factors, normalised rates and averages.

S9.21 Interfaces

The observation layer contains no module-specific algorithm. Each module affects it only through the state it reads or the statistics it populates.

With	Provides or receives	See
Spine 1.1	Provides the bacterial burden, `TotMtb`	S9.14
Spine 2.1, 2.3	Receives the compartment constants, the ordinal convention, the slot count and first-free-slot placement	S9.3, S9.20
Spine 2.2	Receives the grid size and run length (`--dim`, `--days` or `--timesteps`)	S9.6, S9.18
Spine 2.6	Links S9.3 for the occupancy predicates and densities	S9.3
Spine Element 3	Receives the schedule; provides the observation points and the statistics-block order	S9.1
Spine 4.11	Links S9 for the observation design concept and the streams	S9.18
S1	Receives the movement routines whose attempts feed the blocking statistic	S9.15
S2	Receives the macrophage rules that write the macrophage statistics, the residence counter and the polarisation index	S9.8, S9.14, S9.15
S3	Receives the T cell rules that write the T cell statistics	S9.16
S4	Receives the container state, the stored growth rates, the division counter and the bacterial death routes	S9.14
S5	Receives the caseation predicates, the killing count, healing and `Core_nrKillingsCaseation`	S9.3, S9.13
S6	Receives the fields, the derived chemokine relation (S6.14) and the TNF-attributed event flags	S9.11, S9.15, S9.16
S7	Receives the oxygen field; S7.4 links S9.3 for the cell density	S9.3, S9.11
S8	Receives the recruitment events and counters; provides the bacterial totals read by the onset test	S9.14, S9.17
M1	Advances the fields read by the totals; no further effect	S9.11
M2	Option B writes the lymph node and blood, extra-granuloma, crowded-source and lymphatic-exit statistics, and changes the source-active count and the per-compartment counters (M2.9)	S9.17, S9.19
M3	Gives the MHC class II averages and pathway statistics live values; the distance statistics and activation counts use the centre of mass of the current region (M3.6)	S9.15, S9.19

With	Provides or receives	See
M4	Freezes the membership flag from the first fibroblast proliferation (M4.8); makes the fibroblast and TGF-β statistics non-zero (M4.8)	S9.4, S9.19
M5	Three-dimensional build: extent radius from a volume, every extent factor 1, ring-scaled counts equal to unscaled, 27-compartment neighbourhoods and meaningful third coordinates (M5.12); comparability of lesion concentrations between builds (M5.5)	S9.5, S9.9, S9.10
M6	Drug partitioning by the recomputed threshold predicate, drug columns and sub-steps; under fluoroquinolone dosing death within caseum is suspended, so `TotCaseatKill` does not grow (M6.14)	S9.4, S9.14
M7	Writes the anti-TNF and membrane-bound TNF statistics (M7.10)	S9.19
M8	Injection reads the current region (M8.9); exhaustion and Treg induction fill further T cell states; reinfection adds an infected macrophage that `TotalMac` does not count; randomisation compounds the binned growth rates	S9.7, S9.12, S9.14, S9.15, S9.16
M9	Neutrophil statistics are written by the neutrophil rules (M9.9); the neutrophil senescence test reads the previous timestep's region (M9.2)	S9.7, S9.8, S9.19
Appendix B	Indexes the observation options, the stopping criteria and the inert reserve settings	S9.18, S9.20