

TUTORIAL

Agent-Based Modeling in Systems Pharmacology

J Cosgrove^{1,2,3†}, J Butler^{1,2,3†}, K Alden^{1,2}, M Read⁴, V Kumar⁵, L Cucurull-Sanchez⁶, J Timmis^{1,3,7*} and M Coles^{1,2,7*}

Modeling and simulation (M&S) techniques provide a platform for knowledge integration and hypothesis testing to gain insights into biological systems that would not be possible *a priori*. Agent-based modeling (ABM) is an M&S technique that focuses on describing individual components rather than homogenous populations. This tutorial introduces ABM to systems pharmacologists, using relevant case studies to highlight how ABM-specific strengths have yielded success in the area of preclinical mechanistic modeling.

CPT Pharmacometrics Syst. Pharmacol. (2015) 4, 615–629; doi:10.1002/psp4.12018; published online 13 November 2015.

AGENT-BASED MODELS: INTRODUCTION AND APPLICATIONS

The model-based drug discovery and development paradigm is gaining traction in the pharmaceutical industry. There has been a near constant flow of new terms introduced into the literature¹ in an attempt to capture this phenomenon: “MBDD,”² “model-facilitated” or “model-informed drug development,”³ “Quantitative and Systems Pharmacology,”⁴ and “pharmacometrics.” Large pharmaceutical companies have begun to review, quantify, and report the successes derived from the adoption of a model-based strategy, providing a thorough description of its implementation and impact.^{5–7} The US Food and Drug Administration (FDA) recently utilized mechanistic model-based methodologies to design a postmarketing clinical trial⁸; serving as a clear demonstration of the increasing confidence in and adoption of model-based techniques in pharmacology.

Many of the methodologies utilized to support the deployment of M&S strategies are recurrent across companies, forming a key component of the “learn and confirm” drug discovery and development paradigm.⁹ Such methodologies typically include, among others, pharmacokinetics and pharmacodynamics modeling (PK/PD), statistical design methods, and signaling network reconstruction methods. Such techniques can be applied to various stages of the drug development process, with the capacity to inform experimental design, “go/no-go” decisions, preclinical development, and portfolio prioritization.

Wider adoption of M & S in novel therapeutic design requires the development of techniques capable of assessing whether a putative target will yield a desired disease outcome.¹⁰ Capturing heterogeneous biological systems with phenomena occurring across distinct time and length scales in a single model is a challenging yet frequent requirement in target evaluation and selection.

ABM is used sparingly in a pharmaceutical context. However, it is equipped to address the issue of putative target evaluation, and ultimately, to provide a means of consolidat-

ing existing mechanistic understanding into a platform for hypothesis testing and decision support. The use of ABM is increasing within the basic, social, and ecological sciences, providing novel insights into complex systems and engineering challenges.¹¹

This tutorial focuses on providing an overview of the ABM methodology within a systems pharmacology context. ABM is an M&S technique characterized by emphasis on understanding how population-level behaviors emerge from the aggregate interactions of *individuals*, both with the environment and each other.¹² ABMs are composed of individual entities, known as *agents*; each agent exists in a well-defined *state*, determined by the agent’s attributes and location at a specific point in time. A transition into an alternative state is governed by a predetermined *rule-set* designed to capture the agent’s interactions with other agents or the environment. Using these rules, agents calculate how to respond to features and stimuli within their local environment. The aggregate effects of these individual decisions lead to the emergence of system-wide patterns and behaviors that are not explicitly programmed or intuitively understood from the rules alone. In designing ABMs, hypotheses can be developed as to how individual components or pathways contribute to tissue- or organism-level effects: hypotheses that may be examined further with experimental or clinical studies.¹³

In the following sections we draw on examples from both the basic sciences and the pharmaceutical development process to illustrate how an ABM approach can provide value in the context of preclinical target evaluation. Specifically, we focus on three key areas where ABMs are advantageous to other M&S techniques: the ability to serve as a platform for knowledge integration (Application 1), to incorporate heterogeneity (Application 2), and to model phenomena occurring across different temporal and spatial scales, often through incorporating other modeling techniques in a process known as *hybridization* (Application 3). The application of these strengths is further highlighted in our case study, which demonstrates how an ABM approach is used

¹York Computational Immunology Lab, University of York, York, UK; ²Centre for Immunology and Infection, University of York, York, UK; ³Department of Electronics, University of York, York, UK; ⁴Charles Perkins Centre, University of Sydney, Sydney, Australia; ⁵University of California School of Medicine, LA Jolla, California, USA; ⁶GSK Medicines Research Centre, Gunnels Wood Road, Stevenage, UK; ⁷SimOmics, York, UK. *Correspondence: M Coles (mark.coles@york.ac.uk) or J Timmis (jon.timmis@york.ac.uk); [†]Both authors contributed equally to this publication.

Received 28 January 2015; accepted 31 July 2015; published online on 13 November 2015. doi:10.1002/psp4.12018

to systematically analyze two intervention strategies in the autoimmune disease, experimental autoimmune encephalomyelitis (EAE). As the case study progresses, we identify key design and implementation issues associated with the agent-based approach, as well as engineering-inspired techniques and design principles available to tackle these issues. The supplementary information provides links to additional ABM tools (**Table S1**), resources and tutorials (**Table S2**) including hands-on examples for the interested reader.

ABM Application 1: a platform for knowledge integration, hypothesis testing and experimental design in studies of acquired immunity

As preclinical R&D programs progress, there is a need to translate knowledge between multiple researchers from various disciplines and, despite increasing knowledge, targetable mechanisms are always to some extent incompletely understood. A number of articles have highlighted the difficulties associated with embedding information in equation-based techniques and suggest that the conceptual models developed to understand a biological system may not be fully represented in the resultant equations.^{14–17}

Adopting an ABM approach permits the consolidation of existing information into a software platform capable of rapid hypothesis testing and provides a means of identifying key features of a mechanistic target.¹⁶ An ABM approach is advantageous as a knowledge-integration platform because: (i) An ABM has the ability to capture phenomena occurring across spatiotemporal scales; (ii) the highly visual output of ABMs affords an effective medium for communication within interdisciplinary teams who may not be *au fait* with advanced mathematical notation¹⁸; and (iii) an ABM permits the emergence of phenomena from low-level assumptions (discussed further in “Agent Rules”) offering a unique means of identifying knowledge gaps.¹³ When high-level properties of an ABM do not emerge as anticipated this provides evidence that the hypothesis on which the model is built is false.

The ability of ABM to provide a platform for knowledge integration, hypothesis testing, and experimental design has been used to study a key mechanistic target in vaccinology: the germinal center.¹⁹ The germinal center is a transient microenvironment where B cells enter a state of accelerated evolution to form a robust humoral immune response against a single pathogen.²⁰ Based on constraints obtained from the literature, Meyer-Hermann *et al.*²¹ exploited the emergent property of ABMs to test the veracity of different hypotheses regarding B-cell selection. In that study, the authors compared models developed for each proposed theory and were able to reject those failing to reproduce experimentally observed kinetics.²¹ This ABM was further developed to incorporate the effects of Toll-like receptor 4 (TLR4) signaling on the germinal center and inform experimental studies.²² *In silico* simulation was used to develop novel hypotheses and to identify critical timepoints and conditions to test *in vivo*. This combined experimental and theoretical approach yielded a novel mechanistic insight into the impact of TLR4 signaling on the production of high-affinity antibodies.

ABM Application 2: a tool to understand the molecular mechanisms contributing to patient variability

The use of M&S techniques to inform clinical trial design can demonstrably reduce the number of required studies and maximize the probability of success, as evidenced by several successful evaluations reviewed in the literature.³ However, patient heterogeneity can arise from genetic, molecular, and tissue levels of organization, thus making the anticipation of patient-specific responses challenging.

Top-down data-driven approaches, typically used to find patterns in existing datasets, provide value in addressing the issue of patient heterogeneity by stratifying patients on the basis of efficacy biomarkers, but often fail to provide mechanistic explanations for the disease-associated patterns that are discovered. Deriving a mechanistic understanding of how heterogeneity can arise in biological systems on a fundamental level provides a better understanding of why one patient will respond differently to another. Equation-based models such as ordinary differential equations, which typically provide averaged approximations of interactions occurring in well-mixed space, can incorporate heterogeneity to some extent through parameter perturbation. However, as system complexity increases, modeling cellular heterogeneity may necessitate the duplication of large portions of model structure for each distinct cell population in situations where parameter adjustment alone is insufficient.

The ability to describe individual molecules, cells, or patients provides a useful means of describing the underlying mechanisms, in addition to stochastic events, which contribute to heterogeneity. The ability of ABM to perform “patient-specific” *in silico* studies has been described in a hypothetical proof-of-principle study, which simulated clinical trials of anti-cytokine therapy.²³ In this study, key mechanisms relating to dynamics of the innate immune response in the context of systemic inflammatory response syndrome and multiple organ failure were described using an ABM approach. The resulting simulation of treatment regimes produced outcomes qualitatively similar to those reported in the literature, with none of the simulated therapies showing a statistically significant improvement in mortality rates. This simple model did not stratify patient subsets; however, it did provide insight into the robustness of innate immune system responses to varying therapeutic interventions, and shows how negative results yielded from quantitative systems pharmacology studies can inform key decisions in target evaluation.

ABMs have since been used to stratify patients on the basis of mechanistic understanding of spinal cord and vocal fold pathologies. An ABM approach undertaken by Solovyev *et al.*²⁴ combined data regarding blood flow, skin injury, inflammation, and ulcer formation to study the propensity of spinal cord injury (SCI) patients to undergo ulcer formation. The simulation predicted a higher propensity for ulcer formation in SCI patients, thus identifying a high-risk patient subset.²⁴ A similar approach was undertaken to develop personalized treatment strategies in the context of vocal fold injury, where a high degree of patient variability can make it difficult to predict patient-specific disease progression and treatment responses. Li *et al.*²⁵ used an ABM

approach to consolidate existing information of inflammatory mediators and used this information to determine optimized treatment strategies.

ABM Application 3: harnessing spatiotemporal resolution to probe mechanistic targets occurring across distinct space and timescales

While systems pharmacology approaches have been successfully applied to tackle key issues in oncology, for instance, with the use of an integrative systems approach to quantify anticancer drug synergy in imatinib-resistant chronic myeloid leukemia,²⁶ clinical trials in oncology still have the highest failure rate in comparison to other therapeutic areas.²⁷ In complex processes such as tumor formation, probing targetable mechanisms can be difficult owing to variability arising on multiple scales; cancerous cells adapt at genetic and molecular scales to survive in ever-changing environments, altering cellular phenotypes and therefore treatment efficacy, as documented in studies of the hypoxic environment in tumor centers.^{28–30} From a systems pharmacology perspective, there is a need to develop techniques capable of capturing mechanisms occurring across different time- and length-scales that affect disease prognosis and patient response to treatment.³¹

Ordinary differential equations, which abstract interactions occurring in physical space through the use of contact frequency terms, are suitable in many cases to model biological systems, and can be transformed into partial differential equations to incorporate space explicitly. However, when the biological system of interest is large and complex, it can be difficult to apply these techniques to describe both spatial variation and capture stochastic processes.³²

ABMs, owing to their explicit representation of space and ability to capture events occurring across multiple scales, can capture phenomena occurring on distinct spatiotemporal scales, resulting in what are typically referred to as multiscale models.³³ A common method for developing multiscale models is through the hybridization of an ABM with other techniques to simultaneously represent time and length scales at higher or lower orders of magnitude.

Athale and Deisboeck³⁴ developed an ABM to examine the effects of a molecular switch (controlled via epidermal growth factor receptor (EGFR) signaling) on tumor spatial dynamics in the brain. The model predicted that this switch could affect tumor expansion, leading to the development of novel hypotheses on the posttranslational regulation of protein expression. A particularly interesting feature of this model is the explicit modeling of feedback effects from single-cell decisions on overall tumor-growth dynamics.

A hybrid agent-based approach incorporating PK/PD modeling approaches was developed to explore the dynamics of tumor growth, infusion, and penetration of the bioreductive drug tirapazamine (TPZ).³⁵ In that study, a PK/PD methodology was applied to model factors such as hydrogen ion production, nutrient distribution, and drug concentration, while an agent-based methodology was used to model each individual cell over space and time, capturing interactions between cells and with the tumor microenvironment. Combining both types of modeling approach enabled determination that the drug was incapable of reaching the

edge of the hypoxic region of the tumor, partially due to consumption of the drug as it diffused into the tumor. This example of hybrid ABM-PK/PD model highlights the ability of ABM to integrate different modeling techniques, overcoming the respective limitations associated with using either modeling approach in isolation.

CASE STUDY: PROBING THE EFFICACY OF TWO PUTATIVE TREATMENT STRATEGIES IN AN *IN SILICO* MODEL OF AUTOIMMUNE ENCEPHALOMYELITIS

In complex multifactorial pathologies, such as multiple sclerosis, numerous mechanisms may ameliorate disease. Lead target identification and optimization requires a systematic screening of these mechanisms. In this section we describe how the ABM-specific strengths described above make this modeling approach suitable for exploring two putative intervention strategies in an experimental animal model of multiple sclerosis. We will then utilize this model as a case study for designing and developing an ABM in later sections of this tutorial.

Multiple sclerosis is a neurological disorder in which myelin, a substance that coats the neurons of the central nervous system (CNS) to ensure efficient propagation of neural signaling, is inappropriately targeted by the immune system. The resultant damage leads to a wide range of symptoms occurring across a spectrum of severity.

Much of the current understanding of multiple sclerosis pathogenesis has been obtained from animal models of the disease such as EAE.³⁶ This well-established experimental model has yielded mechanistic insights into monophasic, relapsing-remitting, or chronic forms of the disorder. Importantly, a murine EAE model has shown that disease is in part mediated through the action of encephalitogenic CD4Th1 cells.³⁷ Furthermore, the killing of encephalitogenic CD4Th1 cells in a coordinated effort from regulatory CD4+ and CD8+ T cells mediates recovery from disease.^{38,39} This led to the proposal that inhibition of T-cell receptor signaling using a monoclonal antibody against CD3 might target encephalitogenic T cells, thereby either preventing EAE or expediting its recovery. Anti-CD3 therapy was first used to treat transplant rejection (muromonab) and is currently being investigated for the treatment of other autoimmune conditions.

In addition, the importance of the spleen in recovery from EAE has been highlighted in rats: in contrast to the acute symptoms of control animals, splenectomized rats displayed symptoms of chronic disease following EAE induction.⁴⁰ To consolidate experimental findings gained from animal EAE models and to systematically investigate how splenectomy or anti-CD3 treatment could affect the progression of EAE, an agent-based model (ARTIMMUS) was developed.⁴¹ In the following sections we will use ARTIMMUS to introduce the key components of an ABM, illustrate how to develop an ABM using a principled design approach, and show how such an approach has the potential to inform key decision-making processes in the drug development process.

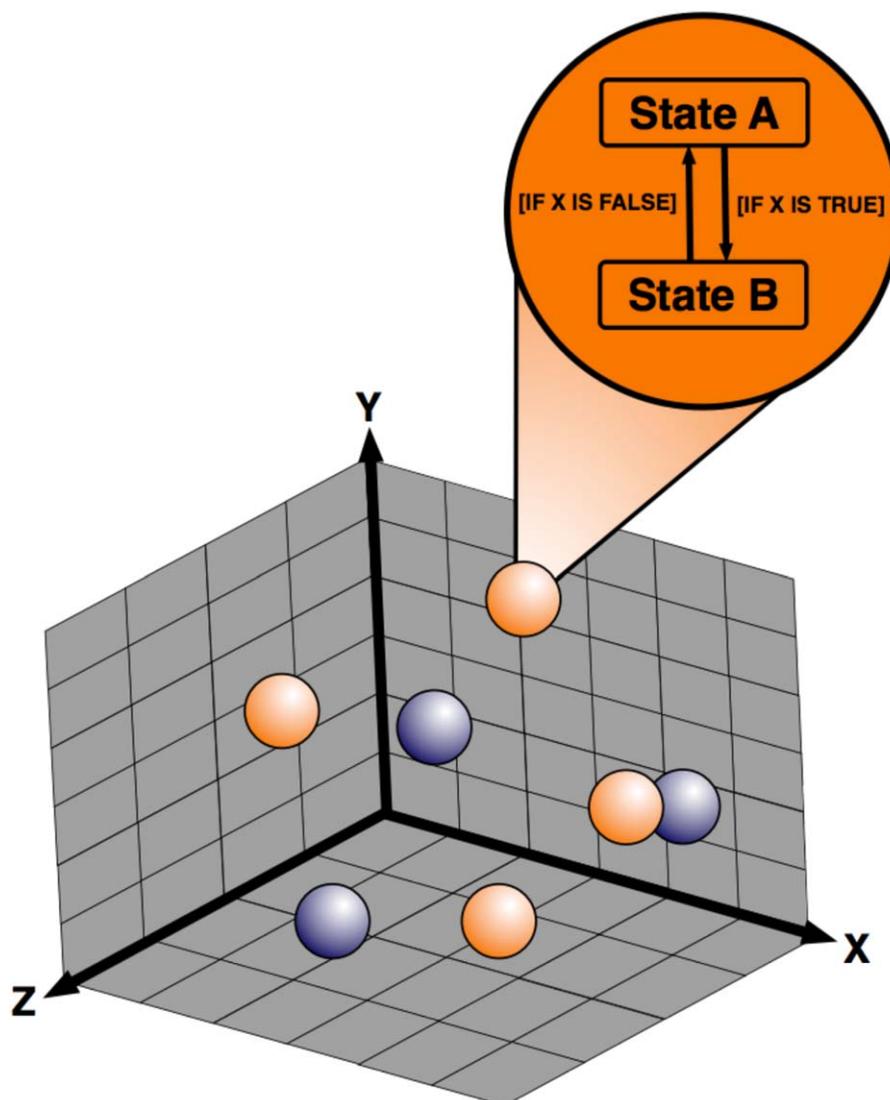


Figure 1 Structure of an ABM: Agents (shown as blue and orange spheres) are individual entities capable of maintaining their associated attributes with respect to their local environment and governing rules. The environment in which an agent exists provides a context for their interactions. The aggregate behaviors of the agents can then lead to the emergence of complex patterns and behaviors.

Key components of an ABM

An ABM contains four main components: 1) agents; 2) the environment within which the agents exist; 3) a rule-set governing relationships and means of interaction between agents and their environment; and 4) the timescale over which the simulation is executed. The system's overall behavior emerges from the collective interactions of these components over time. These key components and terminology relating to ABM are illustrated in **Figure 1** and **Table 1**.

Agents. The agents of an ABM are autonomous, self-directed entities that are designed to represent individual biological components. Each agent is associated with a set of states in which it may exist, and which define its behavior; this is referred to as a finite-state machine. Some

states are mutually exclusive, while others are orthogonal (simultaneous states), allowing finite state machines to capture complex biological behaviors. Importantly, an agent only has local information available at a given moment in time. Decisions and behaviors therefore arise due to factors in an agent's local neighborhood, or through interactions with neighboring agents.

We now consider modeling a key cell type from our case study, the microglia, as a state machine using the Unified Modeling Language (UML), a set of diagrammatic tools for designing software-based systems. **Figure 2** depicts the states in which a microglia may exist. Microglial cells exist in immature and mature states. While immature, they are more phagocytic than when mature. Both immature and mature microglia are able to express MHC-II molecules. Microglia do

Table 1 Table of ABM terminology

Entity	An independent element of the model, such as a cell or protein.
Finite state machine	A finite state machine consists of a set of states, which may include substates, some of which are orthogonal (simultaneous states). A finite state machine may exist in only one state for each orthogonal group at a time.
Agent	An autonomous, self-directed representation of an entity, operating as a finite-state machine.
Space	Computational representation of the physical spatial compartments within which agents are contained.
Environment	Features in space, which provide a context for agent interactions and behaviors, may contain distributions of molecular concentrations that both influence and can be influenced by agent behavior.
Neighborhood	The local environment in which an agent exists, often described as the agents adjacent to, or in contact with a specific agent.
Neighbor	An agent that exists within the neighborhood of another agent.
Model	A nonexecutable description of a system, which may be described in an abstract manner, or for a platform-specific implementation as a simulation.
Simulation	An executable implementation of a model specification.
Step	An iteration in time of a discrete-event simulation.
Hybridization	Using a combination of modeling techniques to capture aspects of the system at different scales in a tractable manner, to overcome the limitations associated with using each technique in isolation.
Multiscale	A model combining processes occurring at different orders of magnitude of time and length.
Richness	The detail contained within an agent, environment or model; comprising internal representations of properties such as cell-surface protein levels, gene expression, etc.

not exist indefinitely, and expire after some period of time (depicted as a ringed black circle in **Figure 2**), indicating that the component is destroyed.

Agent rules. An agent uses a predefined rule-set to assess its internal state in response to factors in the agent's local environment or neighborhood. Should an agent be in a situation where the requirement of a rule is met, whether due to a change in the agent's attributes or within a set location in the environment, the state of that agent is changed. An agent's rule set can range from simple Boolean statements operating over the agent's attributes, to more sophisticated mechanisms that relate agent inputs and outputs such as differential equations¹⁶ and metabolic models.¹⁷ Agent rules also offer a means of introducing stochasticity through probabilistic events, allowing for an approximation of behaviors in nondeterministic systems.

In our case study we see an example of a rules governing microglia behavior (**Figure 2**). The transition of a microglia from a highly phagocytic state to one of reduced phagocytic capacity depends on the cell undergoing maturation. Maturation occurs some time into their lifespan, represented by

λ (*maturation*), but may also be induced through perception of a sufficient concentration of type 1 cytokine. Perception of sufficient concentration of type 1 cytokine induces tumor necrosis factor alpha (TNF- α) secretion in microglia.

Agent environment. ABMs are usually spatially resolved, with agents occupying specific locations in space rather than existing in a well-mixed continuum. Space in an ABM is analogous to the physical spatial environments within which biological entities are contained and interact, enabling complex biological structures to be captured within the model. ABMs can contain multiple different spatial environments, typically referred to as compartments, and can allow agents to move between them. In ARTIMMUS a microglia may only reside in the CNS compartment while other agents such as T cells may migrate between different compartments.

Timescales and granularity. Time is typically calculated within an ABM by dividing the total time over which the biological system is being simulated into discrete time-steps. At each time-step, an agent must determine how to respond to factors in its environment, with respect to its associated rule-set and current state. In ARTIMMUS each agent must determine how to respond to external stimuli on the basis of its current state and governing rule-set every 10 minutes. Thus, the agent may show different responses to the rule-set throughout the course of the simulation. In addition, ABMs can incorporate phenomena occurring at different timescales. Population-level ordinary or partial differential equations can be used to represent soluble factors or small-scale molecules, and these can be solved on a per-agent basis within an ABM.⁵⁴ Modeling these small-scale factors through population approaches instead of as explicit agents can save considerable computational power. Furthermore, such multiscale modeling reduces the computational expense incurred if the highest resolution timescale were to be used for each phenomenon

Constructing an ABM. As with any systems pharmacology approach, ABMs are not without their limitations. ABMs are designed to directly capture attributes and states of individual components; population counts and rates of change emerge during simulation execution rather than being explicitly specified. Accordingly, ABMs can be difficult to describe succinctly, and they may require more data to develop, and significant software engineering expertise to implement. For reasons discussed in "Calibrating ABM" and "Exploring Simulation Behavior," an ABM must be executed multiple times to obtain representative outputs and to ensure robust development and confidence in results.^{42,43} As such, ABM development can require significant time and computational infrastructure. Throughout this section we describe how the key components from the previous section are developed into an ABM by dividing the process into three key stages: design, implementation, and validation. We illustrate how some ABM-associated issues can be mitigated through the application of well-established engineering practices, and detail the approaches used in the design of ARTIMMUS.

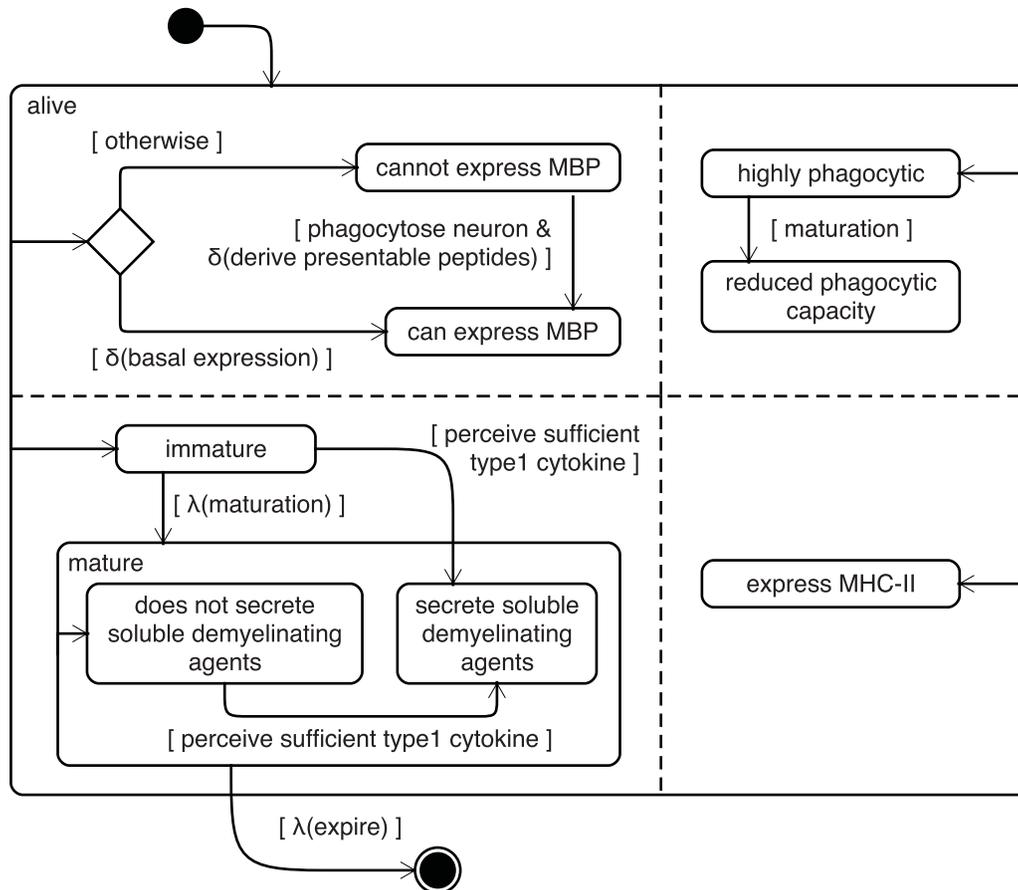


Figure 2 Microglia modeled as agents using the UML. The modeling of microglia in ARTIMMUS. Microglia exist only in the CNS. The only MHC:peptide complex that they present is MHC-II:MBP. This presentation requires the phagocytosis of a neuron and is probabilistic. A small proportion of microglia expresses MBP immediately, represented by $\lambda(\text{basal expression})$. This is to reflect the fact that the physiological turnover of neurons (which is not in itself represented in the domain model) will result in their phagocytosis by microglia and the presentation of MHC-II:MBP complexes. Microglia exist in immature and mature states. While immature they are more phagocytic than when mature. Maturation occurs some time into their lifespan, represented by $\lambda(\text{maturation})$, but may also be induced through perception of a sufficient concentration of type 1 cytokine. Perception of sufficient concentration of type 1 cytokine induces TNF- α secretion in microglia. Both immature and mature microglia are able to express MHC-II molecules. Microglia do not exist indefinitely, and expire after some period of time, represented by $\lambda(\text{expire})$. Figure adopted from Read *et al.*⁴⁰

Stage 1: designing an ABM

Principled design frameworks as standardized protocols for ABM development. The adoption of standardized software practices is a necessary step in ensuring that biomedical software is readable, scalable, and usable.^{44,45} Principled design frameworks that aim to enforce a robust model development methodology facilitate these goals. Grimm *et al.*⁴⁶ proposed a “three-block” standard protocol for describing ABMs termed ODD (overview, design concepts, and details). The three blocks are subdivided into seven stages: Purpose, State variables and scales, Process overview and scheduling, Design concepts, Initialization, Input, and Submodels. A framework for the development of complex systems models, independent of either domain or modeling techniques, termed the CoSMoS (Complex System Modeling and Simulation) process, has since evolved comprising an iterative process of model refinement and implementation.⁴⁷

Irrespective of the specific protocol, a principled design framework must provide a focus on the underlying research question and current biological understanding. This is particularly relevant when one considers that drug development relies heavily on the quality of published preclinical data²⁷; however, studies conducted by Amgen (Thousand Oaks, CA), with a view to confirming published results important to their R&D efforts, could only reproduce similar findings in 11% of cases.²⁷ If an ABM is to be used as a means of knowledge integration, then the inclusion, or omission, of specific data used to inform a model should be appropriately justified. An argumentation structure can ensure that the thought processes behind model composition and design decisions are transparent and readily available for scientific scrutiny.⁴⁸ Such an approach takes inspiration from critical systems engineering, e.g., aircraft design, where key decisions are presented as arguments over evidence.⁴⁸

Key Components of ARTIMMUS

Agents

CD4Th1 and CD4Th2 cells, CD4Treg and CD8Treg cells, microglia, dendritic cells, and neurons were identified as having critical roles in the emergence of EAE are included in ARTIMMUS as agents. **Figure 4** provides an abstract depiction of the major cell types involved in EAE autoimmunity and its associated recovery.

Agent Rules

Following *in silico* immunization, agents respond to features in their local environment based on low-level rules informed by the relevant literature as specified in **SUPPLEMENTARY FIGURE 1** of Read *et al.*⁴¹ An example is provided in **Figure 2** to illustrate this.

Environment

An agent may exist in one of five spatial compartments: the CNS, the draining cervical lymph node (CLN), the spleen, a secondary lymphoid organ (SLO), specifically a peripheral lymph node and a connecting circulatory system. This information is summarized diagrammatically in **Figure 5**.

Time

The simulation starts with immunization and typically runs for 60 days postimmunization. Time is discretized into 10-minute steps.

Although there is currently no clear consensus on a standardized protocol for developing ABMs, key features remain consistent across most methodologies: (i) designing an ABM, (ii) implementing an ABM, and (iii) validating an ABM. These processes can be further subdivided, using existing CoSMoS terminology⁴⁷ as follows: (1) defining a research context; (2) developing a nonexecutable model of the underlying biological system: the domain model; (3) developing a simulation specification: the platform model; (4) implementing a simulation: the simulation model; (5) calibration and analysis of the simulation, and (6) simulation validation. In the following sections we further describe these steps and show how they have been applied to the development of ARTIMMUS.

Defining a research context. The design and implementation decisions made when constructing an ABM are influenced by the overarching scientific objectives of the work, and simulation results are interpreted in this context. As such an ABM should not be considered a “general purpose” description of a system. By extension, an ABM should not be used outside of its original scope, or within a different research context, without considering whether the existing abstract representation of the system is still fit-for-purpose.

It is important to determine whether an ABM approach is the appropriate technique to apply for a given problem. As highlighted under “Constructing an ABM,” above, the technique is not without limitations and therefore the modeling approach that is most suited to the research question should be applied. A basic scheme for selecting a bottom-up modeling approach is presented in **Figure 3**.

Developing a domain model. The initial stage of ABM development should scope the extent of the biological sys-

tem to be captured. This entails creating a “domain model,” a nonexecutable model focusing purely on biological concerns, disregarding any consideration of how the biology will actually be implemented and simulated. The domain model should describe the states, relationships, and methods of interaction (the rule-set) for the biological entities being captured. The ARTIMMUS domain model describes the cell types to be captured, how their interactions influence one another’s behavior, and the environment in which they exist, but does not specify how they are to be implemented.

Developing a platform model. The next stage in constructing an ABM is to delineate how the domain model is to be implemented, an artifact we term the platform model. This is achieved by abstracting biological processes into mathematical constructs to move from the domain model towards a simulation. There are numerous ways in which a domain model might be implemented as a piece of software, the platform model represents the software specification that principled software engineering frameworks advocate creating, as inappropriate decisions and abstractions can render the simulation an inadequate reflection of the biology.

		SPATIALLY RESOLVED	
		NO	YES
CELLULAR HETEROGENEITY	NO	PBPK ODE	PDE
	YES	State-Based Model	ABM HYBRID-ABM

Figure 3 The capacity for various types of model to capture spatial resolution and cellular heterogeneity: When determining the appropriate modeling technique to employ it is important to consider the spatiotemporal scales relevant to the system and the heterogeneity of the entities of interest. Ordinary Differential Equations (ODEs) and Physiologically Based Pharmacokinetic (PBPK) models cannot capture systems with explicit spatial resolution (although compartmentalized systems are possible), relying on the abstract notion of well-mixed space. Partial Differential Equations (PDEs), and thereby, coupled systems of ODEs, are capable of spatial resolution, but to capture heterogeneous cellular phenotypes is often intractable. State-based modeling approaches enable heterogeneous phenotypes among cell populations but cannot in themselves capture spatial resolution (although they can model multiple, spatially disconnected compartments). ABMs incorporate state-based systems in spatial environments; as such, ABMs can capture both heterogeneous cell populations with an explicit notion of space and time.

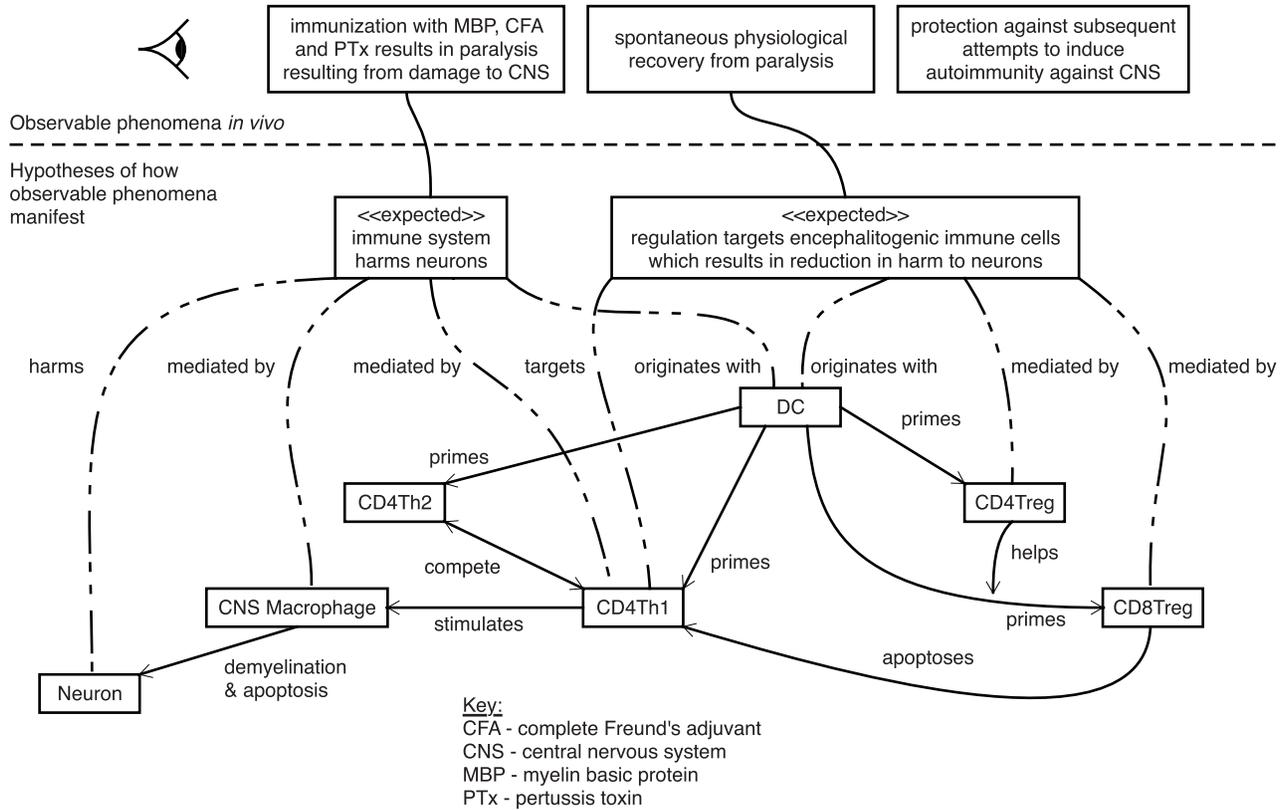


Figure 4 Capturing the emergent phenomena of EAE: An expected behaviors diagram sets the research context of the ABM. This is achieved by depicting the phenomena observed in the murine EAE model, and the behaviors manifesting from cellular interactions hypothesized to be responsible for them. Figure adopted from Read *et al.*⁴¹

The biological concepts captured in the domain model are translated into a software specification, and the architecture and organization of the simulation software are designed. It is important that any explicit reference to emer-

gent system-level properties that might be captured in the domain model, used to indicate how these properties emerge from the interactions of the biological components, are omitted from the platform model. Such properties

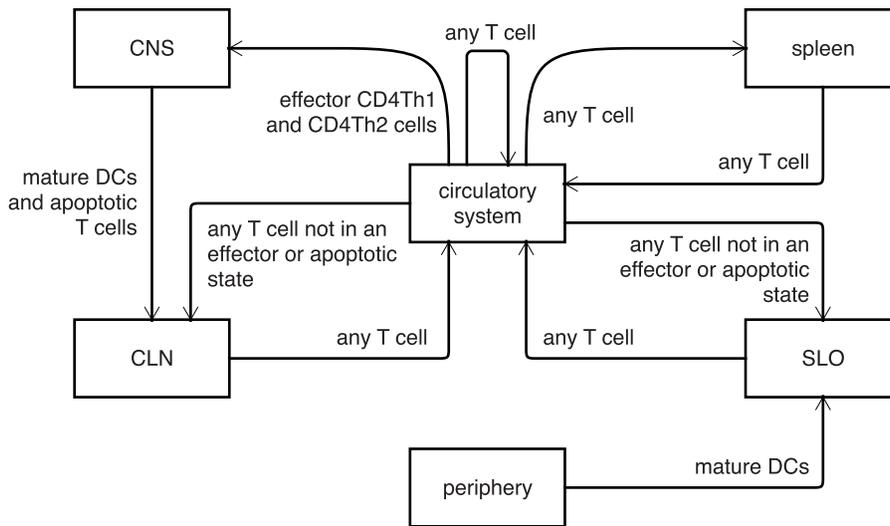


Figure 5 Spatial compartments within ARTIMMUS: The spatial compartments of the domain model, and the manner in which cells may migrate between them. Figure adopted from Read *et al.*⁴¹

should emerge in the simulation, thereby validating it as a suitable abstraction of the biology. **Figure 4** depicts how broad system-level behaviors in ARTIMMUS are hypothesized to emerge from effects that cell populations exert on one another, but was not used in implementing ARTIMMUS; **Figure 2** was used, depicting the dynamics of an individual microglia.

A further key consideration for a platform model is how users will interact with the simulation. Graphical User Interfaces are useful for assessing whether the code and individual agents perform as expected at any given moment in time, and offer an intuitive understanding of what is happening in the simulation.

Stage 2: implementing an ABM

Simulation Model. The simulation model represents a computer code implementation of the platform model of Stage 1. Code development can be time-consuming and require specific software engineering expertise; thus, to expedite matters an increasing number of toolkits and ABM-specific software libraries have become available. These offer programmers a high level of control in development while providing commonly required functionality and data-structures. A list of commonly used development tools is provided in **Table S1**.

During simulation implementation, state transitions detailed in the specification are typically expressed using control (“if – then”) statements using a programming language. Languages of the Object Oriented programming (OOP) paradigm such as JAVA⁵⁰ have proven a popular and natural means of implementing ABM.⁵¹ OOPs define *classes*, which are templates for program components and describe how they interact. A class is *instantiated* to obtain an *object* and its associated individual state. Each object is responsible for storing and maintaining its own data, while classes provide functions to strictly control how modifications are made to object data, promoting logical and well-structured code.

At each discrete time-step of the simulation (refer to “Timescales and Granularity”) factors in the agent’s environment may change, prompting an agent to respond accordingly. Care must be taken to determine the order in which agents are processed, as this could introduce bias. In many instances a randomized schedule is implemented to ensure that one agent’s decisions are not prioritized over another’s.

Calibration. A simulation will contain a number of parameters that control the mathematical constructs used to represent the agents. As these constructs may not translate directly to the biological system, it is often not possible to obtain an exact experimental value for a simulation parameter. Additionally, through the process of abstraction, a parameter controlling one mathematical construct may be accounting for additional factors, which occur in the biological system but have been omitted from the model. For example, ARTIMMUS’s “type 1 cytokine” abstracts the function of several cytokines that promote a type-1, cytotoxic adaptive immune response, such as interferon gamma (IFN- γ), interleukin 2 (IL-2), and lymphotoxin alpha (LT- α).

Designing ARTIMMUS

Employing a Principled Design Framework

The CoSMoS framework was employed to develop ARTIMMUS in a principled manner.⁴⁷

Defining a Research Context

ARTIMMUS was developed to consolidate the current mechanistic understanding of the murine EAE model and the resultant simulation was designed to systematically probe two intervention strategies: splenectomy and anti-CD3 treatment.

Choosing the Correct Technique

Similar research contexts have been addressed successfully using a differential equation approach by Fousteri *et al.*,⁴⁹ who utilized the type 1 diabetes (T1D) PhysioLab platform, developed to reproduce T1D pathophysiology in the NOD mouse, to make predictions of the impact of timing of the application of nasal insulin to prevent the onset of T1D by simulating features of untreated pathogenesis and disease outcomes for multiple interventions. In the case of ARTIMMUS, however, an ABM approach was warranted due to the heterogeneity of cell populations, the need to account for phenomena occurring on distinct time scales, and features in the cellular environments hypothesized as being important contributors to EAE.

Developing a Domain Model

As specified by CoSMoS, the purpose of a domain model is to consolidate current understanding of the biological system with no consideration to how the model might be implemented. Key phenomena and cell types, the dynamics of each cell type, and how their actions integrate to constitute onset and recovery from EAE were explicitly documented in a series of diagrams based on an extended version of the UML.^{16,43} This model was developed with extensive input from experimental collaborators to ensure that the domain model is an adequate representation of the murine EAE model.

Developing a Platform Model

With the domain model upon which the model will be based agreed, the next series of UML-based diagrams concerns their implementation as a simulation (referred to as the platform model in this tutorial). The model details the implementation of the structures, behaviors, and interactions identified in the domain model in a way that naturally translates to simulation platform technologies using the UML. ARTIMMUS employs a 2D lattice-grid for spatial representation in each compartment and these compartments are networked, allowing cells to leave one compartment and to enter another. Within these lattice grids an agent may move to any of the eight grid spaces surrounding the one in which it currently resides, or it may remain stationary. Between each grid space cytokines and soluble factors represented as concentrations that are subject to decay and diffusion.

The calibration process identifies simulation parameter values that best align the simulation’s behavior with that observed in the experimental system.

Calibration is an important step, as it establishes the baseline behavior of the simulation, to which the results of future *in silico* investigations are compared. Procedures to fit simulation results to biologically derived responses have been described in a study of the life cycle of *Mycoplasma genitalium*.⁵³ In brief, the authors fitted their whole cell computational model to datasets spanning metabolomic, transcriptomic, and proteomic

Implementing ARTIMMUS

Simulation

The simulation was implemented in the JAVA programming language and compatible MASON simulation framework.⁴⁷ Each cell type identified in Stage 1 was implemented as a Class, based on its state machine diagram (e.g., **Figure 2**).

Calibration

ARTIMMUS was calibrated against two experiments⁵⁴: the physiological recovery of mice induced into EAE, and labored recovery following abrogation of CD8Treg capacity to apoptose encephalitogenic CD4Th1 cells.

information. The resulting parameter values were then validated by determining how well independent experimental datasets were reproduced for multiple cellular functions. However, the effective calibration of complex ABM is an open research area, often requiring the use of sophisticated heuristic optimization technologies.⁵³

Stage 3: validating an ABM

Mitigating uncertainty introduced by implementation. Heterogeneity and stochasticity are key benefits of an agent-based approach, and as with biological systems, repeat experiments can lead to differing results. This variation is termed “aleatory uncertainty” (AU).⁵⁵ Multiple simulation executions must be performed to ensure that the results collected are representative of a simulation experiment, and not simply random variation. Empirical methods exist to establish the relationship between the number of simulation replicates and the effect of stochasticity on aggregated or averaged results,^{42,43,56} but these can be computationally costly to perform. They operate by contrasting samples of simulation responses obtained under the same experimental conditions, thereby estimating the influence of stochastic variation on the distribution of simulation results.

Exploring simulation behavior. Prior to using a simulation for *in silico* experimentation, it is important to appreciate how sensitive it is to perturbations in parameters. Such an analysis can be performed through application of sensitivity analysis (SA) techniques.⁵⁷ Through a systematic exploration of the parameter space, simulation inputs that have an influential effect on simulation behavior are identified and quantified, aiding the biological interpretation of simulation results; simulations that critically depend on estimated parameter values should be treated more cautiously.

SA techniques are split into two categories: local and global analysis. Local analysis techniques examine how robust the simulation is to a perturbation of a single parameter value. However, local SA techniques cannot reveal compound effects where one parameter’s influence is dependent on the value of another. Such effects may be elucidated using global analysis techniques that perturb multiple parameters simultaneously. For complex models containing a large number of parameters, tractability of the parameter search space is an important consideration. Statistical techniques need to be applied that provide an efficient exploration of the parameter space. Two global

analysis parameter sampling techniques, LHC (Latin-Hypercube) and eFAST (Extended Fourier Amplitude Sampling Test),⁴² have been successfully applied to the analysis of biological ABMs⁵⁸ to determine parameter sensitivity as other parameters are varied.

In addition to highlighting simulation sensitivity to particular parameters, SA techniques can identify key simulated biological pathways.⁴² In the context of preclinical mechanistic modeling, SA techniques can help highlight where a single or combination strategy is required. Marino *et al.*⁵⁹ developed a novel SA approach that explored simulation behaviors, examining how the number of T cells leaving the lymph node affects the bacterial load present in the lung. Uniquely, this approach demonstrates that inputs into SA need not be simulation parameters, but may be emergent properties within the simulation. Separately, SA techniques were also used by Ray *et al.*⁶⁰ to explore a key component of *Mycobacterium tuberculosis* infection: granuloma formation. Granulomas are aggregates of immune cells that can determine host response to infection. Previous experimental work highlighted how cytokine TNF- α plays a key role in the formation and maintenance of these structures, but the dynamics of this system were difficult to discern using solely experimental techniques. Through the use of ABM, Ray *et al.* isolated and examined five separate functions of TNF- α , a separation not possible experimentally. They predicted that multiple TNF- α activities and macrophage activation are key contributors to the control of infection within a granuloma.

Experimental validation. While the emergence of expected or observed biological phenomena is one indication that the ABM is an adequate representation of the biological system it captures, it is also important, where possible, that ABM undergo experimental validation.

This is illustrated by Peirce and colleagues⁶¹ in a model of therapeutic adipose-derived stromal cell (ASC) trafficking during ischemia. The model design is supported by the simulator’s ability to appropriately reproduce important aspects of ischemia and trafficking behavior. Subsequent simulations revealing the necessity for an unknown selectin-binding molecule to achieve ASC extravasation prompted further *in vitro* experimentation, which confirmed that a subpopulation of ASCs slowly rolled on immobilized P-selectin further validating the simulations predictions while gaining novel biological insights.

Using ARTIMMUS to inform intervention strategies for EAE

Calibration and sensitivity analysis data suggested that ARTIMMUS constitutes an appropriate representation of murine EAE. Consequently, two interventions were investigated: anti-CD3 treatment and splenectomy. We summarize them in this section.

Splenectomy intervention. To simulate a splenectomy in ARTIMMUS, the entire spleen spatial compartment and its contents (including cells) are removed from the simulation, which prevents cells from entering or leaving it. This facilitated an analysis of the role that the spleen has in EAE onset and recovery. It was shown that the spleen is a major site of T

Validating ARTIMMUS

Mitigating Uncertainty

An aleatory uncertainty analysis⁴² revealed that 500 simulation replicates were needed for each experiment to reduce the influence of aleatory uncertainty.

Robustness Analyses

A form of local sensitivity analysis, which reveals how far each individual parameter may be perturbed from calibrated baseline values before scientifically significant deviations in simulation behavior occur, was performed.⁵⁴ The analysis revealed that the majority of calibrated parameters were robust to perturbations.

Global Sensitivity Analysis

The importance of various simulation components and pathways was assessed through application of a global sensitivity analysis.⁵⁵ Latin Hypercube Sampling was employed to efficiently sample parameter space, and partial rank correlation coefficients were calculated to determine how parametric variation correlated with changes in simulation behavior.⁶⁰ The analyses lead to a number of insights into the system. For example, the relatively low influence of parameters pertaining to the killing of encephalitogenic CD4Th1 cells by regulatory CD8 T cells suggested a considerable redundancy in this regulatory pathway's ability to ameliorate autoimmune behavior.

regulatory cells (Treg) priming and that splenectomy significantly reduces CD4+ and CD8+ Treg population sizes, leading to a reduced capacity to completely abrogate encephalitogenic CD4Th1 populations, which can lead to their reexpansion.⁴¹ Splenectomized groups suffered a higher mortality rate and had a greater tendency towards relapsing clinical disease.

Anti-CD3 therapeutic intervention. As a T-cell-mediated autoimmune disease, EAE is potentially treatable using anti-CD3 antibodies. Such an intervention, with varying efficacies, can be simulated through ARTIMMUS's agent-based technology. A probabilistic test is performed whenever cognate T cell and antigen binding (in the form of a peptide bound to an MHC molecule) is attempted. In a treatment simulation, an intervention is represented through higher probabilities of preventing this binding taking place. In a control simulation the lack of intervention is represented with an efficacy of 0% (i.e., with zero probability of peptide-to-MHC binding prevention).

ARTIMMUS was used to simulate different anti-CD3 antibody dosing regimens administered at the time of encephalitogenic T-cell expansion (day 4), or with the onset of clinical symptoms (day 15). *In silico* analyses of putative treatment regimens showed that anti-CD3 monoclonal antibody (mAb) treatment is not uniformly beneficial in treating EAE. Regardless of administration time, only treatments leading to 80% inhibition of CD3 signaling or greater were effective in reducing the duration of clinical episodes, and preventing clinical relapses. This finding was consistent with clinical trials of anti-CD3 intervention in new-onset type 1 diabetes, where large cumulative doses were required for maintaining beta-cell function^{62,63} and where phase III trials using low doses failed to meet their primary endpoints.⁶³

Here we show previously unpublished data from the original ARTIMMUS anti-CD3 intervention experiment, examining how neuronal death relates to varying intervention efficacies. While efficacies of 80% and less can reduce the peak rates of neuronal death, neuronal destruction persists for longer periods of time (**Figure 6b**). The cumulative number of neurons destroyed exceeds control levels for efficacies of 80% and less (**Figure 6b**). These results highlight how interventions that indiscriminately target T cells must be carefully considered, as they can interfere not only with target autoimmune T cells, but the regulatory T-cell populations that suppress them.

DISCUSSION

Agent-based modeling is a method for constructing computational models focusing on distinct components, their environment, and their interactions, as governed by a set of experimentally derived rules. ABMs can be used to model complex biological systems in which the behavior of heterogeneous individual components, stochastic events, and/or spatiotemporal considerations are important.²⁴

As with any quantitative systems pharmacology approach, the ability to integrate knowledge is a key feature of ABMs. The systematic organization of data across scales of interest and different disciplines is nontrivial. The ABM paradigm lends itself well to knowledge integration due to its intrinsically modular organization, ability to capture phenomena occurring on distinct spatiotemporal scales simultaneously, and a highly visual output. The lack of established analytical techniques used in the analyses of equation-based systems is often perceived as a shortcoming of ABM. However, it is possible to measure and analyze the visual output of ABM in the same manner as wet-lab systems,¹⁵ as demonstrated in Butler *et al.*,⁶⁴ with the development of a tool-chain for enabling existing ABMs to produce emulations of flow cytometry, immunohistochemistry, and gene expression microarrays. An existing simulation of lymphoid organogenesis was modified to produce flow cytometry data, permitting analysis with standard software tools, such that *in silico* and experimental data can be treated and presented in an identical manner. It was demonstrated that such approaches can lead to early indications of model predictions, and intuitively demonstrate these in a biological manner.⁶⁴

Systematically integrating knowledge into a mechanistically coherent output is an important driver for rational experimental design. The capacity to integrate spatial and temporal aspects can facilitate a better understanding of *in vitro* and *in vivo* systems, as demonstrated by Walker *et al.*⁶⁵ That study demonstrated that the calcium-dependent pattern of wound closure observed for an *in vitro* assay could be quantitatively reproduced *in silico* using simple rule-based dynamics. Furthermore, differences between *in silico* and *in vitro* models led to predictions for a role in wound-induced signaling events in urothelial cell cultures, showing the benefits of an iterative modeling and experimental work cycle.

Another key feature of the ABM approach is the ability to capture heterogeneity, such as that of patient populations.

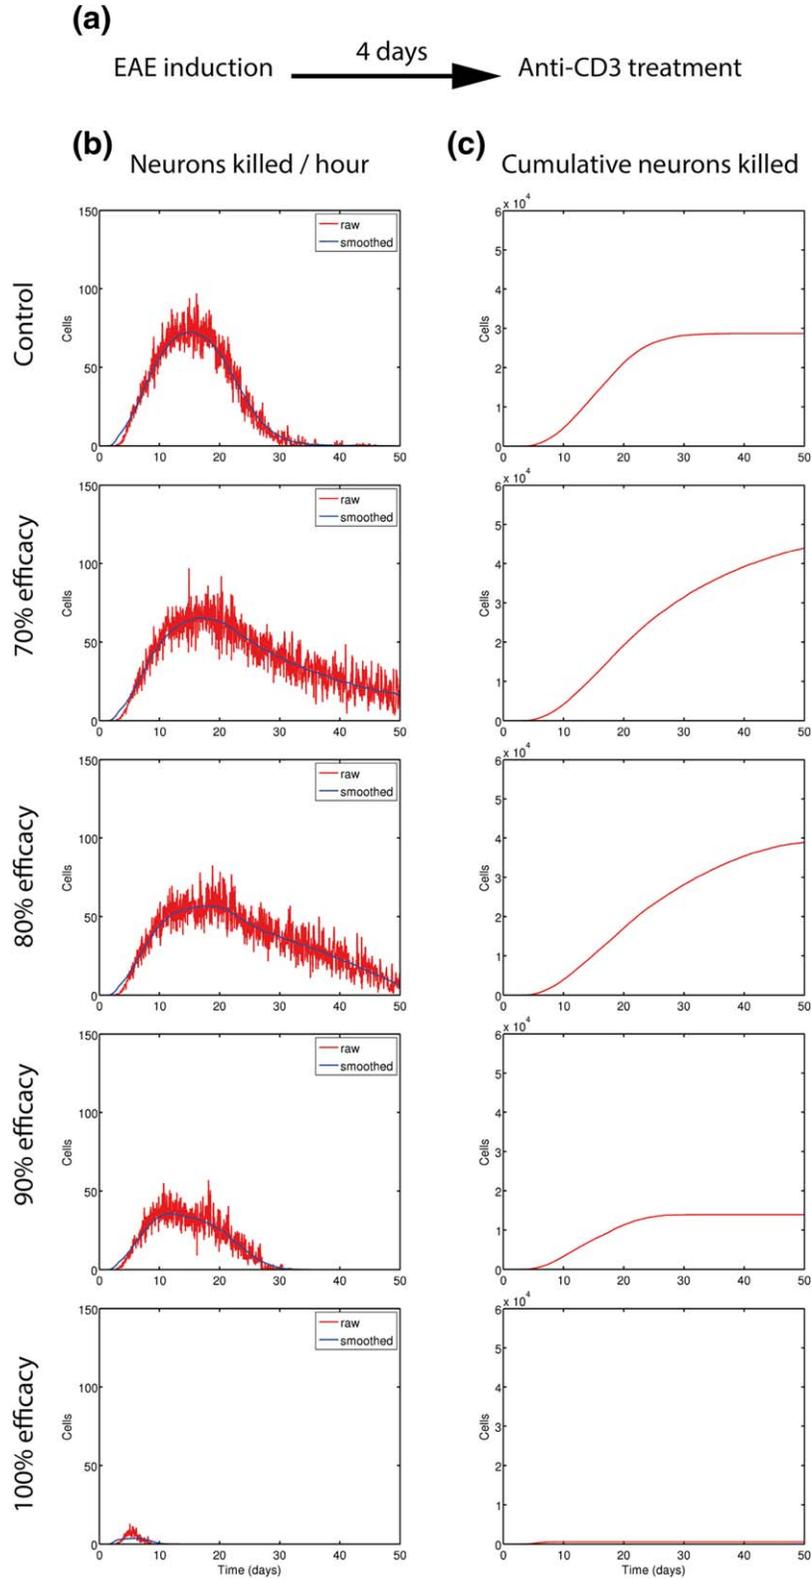


Figure 6 *In silico* anti-CD3 treatment result in a lower rate of neuronal death, but a higher number of total neurons killed for some efficacies greater than 80%: (a) EAE was induced at day 0 followed by administration of anti-CD3 at day 4. (b) The number of neurons killed per hour. (c) Cumulative count of neurons killed for varying efficacies of anti-CD3 treatment. 100% efficacy blocks all TCR:MHC-peptide bindings, 50% blocks half of all binding events, and 0% represents the control. Figure adopted from Read *et al.*⁴¹

Computational Resources Required for ARTIMMUS Development and Exploration

A single execution of ARTIMMUS takes several minutes, depending on the experiment being performed and the machine on which it is executing. Assuming 1 minute per execution, 500 simulation runs would require 8 hours of computational time. The global sensitivity analysis described above used 500 samples of parameter space, and for each sample 500 simulation executions are performed. Hence, this analysis required 173 days of computation time. Given the computational resources required, experiments with ARTIMMUS were conducted on a computational cluster, a network of machines dedicated to performing large-scale computation. The cluster could process 120 simulation executions in parallel, vastly reducing the time required to carry out experimentation, with most experiments performed within 24 hours.

The combination of bottom-up modeling approaches with top-down data-driven approaches, such as bioinformatics analyses of high-throughput genomic data, could present an interesting approach for developing personalized medicine strategies. For example, molecular signatures in the form of biomarker panels could be used as model inputs for subsequent *in silico* simulation of stratified patient populations.⁶⁶

Throughout the tutorial we have discussed the differences between ABM and other modeling techniques. It is important to note, however, that modeling paradigms are not mutually exclusive, and the combination of ABM with other approaches can facilitate the integration of data across different temporal and/or spatial scales. This allows a single model to examine how processes such as molecular diffusion, occurring over a timescale of milliseconds, can affect cellular motility occurring over minutes and hours. ABMs enable phenomena to be captured at the appropriate scale through independent discretization of each scale of interest.

In this tutorial we have highlighted the challenges associated with an agent-based approach, but reason that such issues can be largely mitigated through the application of well-established software engineering practices. A number of principled approaches to simulation development and applications are emerging with the aim of ensuring that simulation results are appropriately interpreted.^{28,36,38,56} Achieving a general consensus on a standardized development framework can facilitate model communication, and therefore simulation reproducibility, peer-review, and repurposing; this therefore should be a fundamental goal in the development of ABM techniques in a systems pharmacology context. CoSMoS, the principled design framework used in the development of ARTIMMUS, for example, shares many key features with the methods for physiological model qualification in drug discovery proposed by *Friedrich and Savic*.⁶⁷ employing concepts from engineering, statistics, complex systems modeling and related fields to address questions of relevance, dealing with uncertainty, dealing with variability, and matching test data in a principled manner.⁵⁸

A key challenge facing the wide acceptance of principled design frameworks is the generation of appropriate model documentation that is clear to an interdisciplinary team, yet concise enough to be practicable.⁶⁸ However, by following

existing software engineering principles models can be expressed in discipline-independent structures and language using visual notations. Visual notations can ensure that the design process is (i) easy for nonspecialists to interpret (provided they are familiar with the syntax), (ii) explicit (i.e., can be interpreted objectively, *not* subjectively), and (iii) accessible. These notations can also be used to provide an argument that the ABM and resultant simulator are a fit-for-purpose representation of the underlying biological domain of interest.⁶⁹ Aside from formal mathematical notations,⁷⁰ a number of well-established notations exist and are addressed in **Table S1**.

The case study in modeling EAE has highlighted how principled simulation development techniques can give rise to well-engineered simulations, and establishes the link between simulation and the original biological system. We have demonstrated the range of interventions that can be simulated through ABM, from removal of an entire spatial compartment to fine-grained manipulations of cell-signaling pathways. Furthermore, the effects of interventions across a heterogeneous and stochastic biological system can be probed: the EAE simulation was used to examine how the range of disease progressions experienced varied with intervention efficacy, be it relatively mild or terminal.⁴¹

The agent-based approach, when applied appropriately and in a principled manner, offers unique advantages in the study of highly complex biological systems. Taken together, the studies referred to in this tutorial highlight how ABM has been used to gain novel insights and give an indication into how the approach will be useful for future studies in systems pharmacology.

Acknowledgments. The first two authors contributed equally to this work. Partly funded by the Wellcome Trust (ref: 097829) through the Centre for Chronic Diseases and Disorders (C2D2). J.C. is supported by a studentship from the Wellcome Trust 4-year PhD programme (WT095024MA): Combating Infectious Disease: Computational Approaches in Translation Science. J.B. is supported by a studentship from the Engineering and Physical Sciences Research Council (EPSRC). M.R. thanks the Judith and David Coffey Life Lab for their support. M.C. is funded in part by the Medical Research Council grants (G0601156 and MR/K021125/1), NC3Rs (NC/K999527/1) and the Human Frontiers Science Program (RGP0006/2009-C). J.T. is partly funded by the Royal Society and the Royal Academy of Engineering and EPSRC grant EP/K040820/1. V.K. is supported by R01 grants from the National Institutes of Health, Bethesda, MD. The authors would like to thank Michael Meyer Hermann and Andreas Deutsch for permitting use of their scheme which is adapted in figure 3.

Conflict of Interest. J.T. and M.C. are Directors of SimO-mics Ltd. L.C.-S. is currently a permanent employee of GSK.

Author Contributions. M.C., J.C., J.B., K.A. M.R., L.C.-S., and J.T. wrote the article; M.C., K.A., M.R., V.K., and J.T. designed the research; J.C., J.B., K.A., and M.R. performed the research; M.C., J.C., J.B., K.A., M.R., V.K., and J.T. analyzed the data.

1. Allerheiligen, S.R.B. Impact of modeling and simulation: myth or fact? *Clin. Pharmacol. Ther.* **96**, 413–415 (2014).
2. Lalonde, R.L. et al. Model-based drug development. *Clin. Pharmacol. Ther.* **82**, 21–32 (2007).
3. Bonate, P.L. Clinical trial simulation in drug development. *Pharm. Res.* **17**, 252–256 (2000).
4. Allerheiligen, S.R.B. Next-generation model-based drug discovery and development: quantitative and systems pharmacology. *Clin. Pharmacol. Ther.* **88**, 135–137 (2010).
5. Milligan, P.A. et al. Model-based drug development: a rational approach to efficiently accelerate drug development. *Clin. Pharmacol. Ther.* **93**, 502–514 (2013).
6. Visser, S.a.G., de Alwis, D.P., Kerbusch, T., Stone, J.A. & Allerheiligen, S.R.B. Implementation of quantitative and systems pharmacology in large pharma. *CPT Pharmacomet. Syst. Pharmacol.* **3**, 1–10 (2014).
7. Visser, S., Manolis, E., Danhof, M. & Kerbusch, T. Modeling and simulation at the interface of nonclinical and early clinical drug development. *CPT Pharmacomet. Syst. Pharmacol.* **2**, 1–3 (2013).
8. Peterson, M. & Riggs, M. FDA advisory meeting clinical pharmacology review utilizes a quantitative systems pharmacology (QSP) model: a watershed moment? *CPT Pharmacomet. Syst. Pharmacol.* **4**(3), 189–192 (2015).
9. Sheiner, L.B. Learning versus confirming in clinical drug development. *Clin. Pharmacol. Ther.* **61**, 275–291 (1997).
10. Schmidt, B.J., Papin, J.A. & Musante, C.J. Mechanistic systems modeling to guide drug discovery and development. *Drug Discov. Today*. **18**, 116–127. (2013).
11. Helbing, D. *Social Self-Organization* (Springer, Berlin, 2012).
12. Bauer, A.L., Beauchemin, C.A.A. & Perelson, A.S. Agent-based modeling of host-pathogen systems: the successes and challenges. *Inf. Sci.* **179**, 1379–1389 (2009).
13. Chavali, A.K. et al. Characterizing emergent properties of immunological systems with multi-cellular rule-based computational modeling. *Trends Immunol.* **29**, 589–599 (2008).
14. An, G.C. Translational systems biology using an agent-based approach for dynamic knowledge representation: an evolutionary paradigm for biomedical research. *Wound Repair Regen.* **18**, 8–12 (2010).
15. An, G., Mi, Q., Dutta-Moscato, J. & Vodovotz, Y. Agent-based models in translational systems biology. *Wiley Interdiscip. Rev. Syst. Biol. Med.* **1**, 159–171 (2009).
16. Hunt, C.A., Kennedy, R.C., Kim, S.H.J. & Ropella, G.E.P. Agent-based modeling: a systematic assessment of use cases and requirements for enhancing pharmaceutical research and development productivity. *Wiley Interdiscip. Rev. Syst. Biol. Med.* **5**, 461–480 (2013).
17. Hunt, C.A. & Ropella, G.E.P. Moving beyond in silico tools to in silico science in support of drug development research. *Drug Dev. Res.* **72**, 153–161 (2011).
18. Brynne, L., Bresell, A. & Sjögren, N. Effective visualization of integrated knowledge and data to enable informed decisions in drug development and translational medicine. *J. Transl. Med.* **11**, 250 (2013).
19. Meyer-Hermann, M., Figge, M.T. & Toellner, K.-M. Germinal centres seen through the mathematical eye: B-cell models on the catwalk. *Trends Immunol.* **30**, 157–164 (2009).
20. Victora, G.D. & Mesin, L. Clonal and cellular dynamics in germinal centers. *Curr Opin Immunol.* **28**, 90–96 (2014).
21. Meyer-Hermann, M.E., Maini, P.K. & Iber, D. An analysis of B cell selection mechanisms in germinal centers. *Math. Med. Biol. J.* **23**, 255–277 (2006).
22. Garin, A. et al. Toll-like receptor 4 signaling by follicular dendritic cells is pivotal for germinal center onset and affinity maturation. *Immunity* **33**, 84–95 (2010).
23. An, G. In silico experiments of existing and hypothetical cytokine-directed clinical trials using agent-based modeling. *Crit. Care Med.* **32**, 2050–2060 (2004).
24. Solovyev, A., Mi, Q., Tzen, Y.-T., Brienza, D. & Vodovotz, Y. Hybrid equation/agent-based model of ischemia-induced hyperemia and pressure ulcer formation predicts greater propensity to ulcerate in subjects with spinal cord injury. *PLoS Comput. Biol.* **9**, e1003070 (2013).
25. Li, N.Y.K. et al. A patient-specific in silico model of inflammation and healing tested in acute vocal fold injury. *PLoS ONE* **3**, e2789 (2008).
26. Winter, G.E. et al. Systems-pharmacology dissection of a drug synergy in imatinib-resistant CML. *Nat. Chem. Biol.* **8**, 905–912 (2012).
27. Begley, C.G. & Ellis, L.M. Drug development: raise standards for preclinical cancer research. *Nature* **483**, 531–533 (2012).
28. Mao, Q. et al. A tumor hypoxic niche protects human colon cancer stem cells from chemotherapy. *J. Cancer Res. Clin. Oncol.* **139**, 211–222 (2013).
29. Mimeault, M. & Batra, S.K. Hypoxia-inducing factors as master regulators of stemness properties and altered metabolism of cancer- and metastasis-initiating cells. *J. Cell. Mol. Med.* **17**, 30–54 (2013).
30. Marte, B. Tumour heterogeneity. *Nature* **501**, 327–327 (2013).
31. Wolkenhauer, O. et al. Enabling multiscale modeling in systems medicine. *Genome Med.* **6**, 21 (2014).
32. Mirsky, H.P., Miller, M.J., Linderman, J.J. & Kirschner, D.E. Systems biology approaches for understanding cellular mechanisms of immunity in lymph nodes during infection. *J. Theor. Biol.* **287**, 160–170 (2011).
33. Dada, J.O. & Mendes, P. Multi-scale modelling and simulation in systems biology. *Integr. Biol.* **3**, 86–96 (2011).
34. Athale, C.A. & Deisboeck, T.S. The effects of EGF-receptor density on multiscale tumor growth patterns. *J. Theor. Biol.* **238**, 771–779 (2006).
35. Kazmi, N., Hossain, M.A. & Phillips, R.M. A hybrid cellular automaton model of solid tumor growth and bioreductive drug transport. *IEEEACM Trans. Comput. Biol. Bioinform. IEEE ACM.* **9**, 1595–1606 (2012).
36. Baxter, A.G. The origin and application of experimental autoimmune encephalomyelitis. *Nat. Rev. Immunol.* **7**, 904–912 (2007).
37. Governanm J. Autoimmune T cell responses in the central nervous system. *Nat. Rev. Immunol.* **9**, 393–407 (2009).
38. Kumar, V. & Sercarz, E. An integrative model of regulation centered on recognition of TCR peptide/MHC complexes. *Immunol. Rev.* **182**, 113–121 (2001).
39. Kumar, V. Homeostatic control of immunity by TCR peptide-specific Tregs. *J. Clin. Invest.* **114**, 1222–1226 (2004).
40. Ben-Nun, A., Ron, Y. & Cohen, I.R. Spontaneous remission of autoimmune encephalomyelitis is inhibited by splenectomy, thymectomy or ageing. *Nature* **288**, 389–390 (1980).
41. Read, M. et al. Determining disease intervention strategies using spatially resolved simulations. *PLoS ONE* **8**, e80506 (2013).
42. Read, M., Andrews, P.S., Timmis, J. & Kumar, V. Techniques for grounding agent-based simulations in the real domain: a case study in experimental autoimmune encephalomyelitis. *Math. Comput. Model. Dyn. Syst.* **18**, 67–86 (2012).
43. Alden, K. et al. Spartan: a comprehensive tool for understanding uncertainty in simulations of biological systems. *PLoS Comput. Biol.* **9**, e1002916 (2013).
44. Bersini, H., Klatzmann, D., Six, A. & Thomas-Vaslin, V. State-transition diagrams for biologists. *PLoS ONE* **7**, e41165 (2012).
45. Ince, D.C., Hatton, L. & Graham-Cumming, J. The case for open computer programs. *Nature* **482**, 485–488 (2012).
46. Grimm, V. et al. A standard protocol for describing individual-based and agent-based models. *Ecol. Model.* **198**, 115–126 (2006).
47. Andrews, P. et al. The CoSMoS process, version 0.1: a process for the modelling and simulation of complex systems. Technical Report. YCS-2010-453, University of York (2010).
48. Alden, K. et al. Using argument notation to engineer biological simulations with increased confidence. *J. R. Soc. Interface* **12**, 20141059 (2015).
49. Foustier, G. et al. Virtual optimization of nasal insulin therapy predicts immunization frequency to be crucial for diabetes protection. *Diabetes* **59**, 3148–3158 (2010).
50. Eckel, B. *Thinking in Java: The Definitive Introduction to Object-Oriented Programming in the Language of the World-Wide Web* (Prentice Hall, Upper Saddle River, NJ, 2002).
51. Johnson, C.G., Goldman, J.P. & Gullick, W.J. Simulating complex intracellular processes using object-oriented computational modelling. *Prog. Biophys. Mol. Biol.* **86**, 379–406 (2004).
52. Karr, J.R. et al. A whole-cell computational model predicts phenotype from genotype. *Cell* **150**, 389–401 (2012).
53. Read, M., Tripp, M., Leonova, H., Rose, L. & Timmis, J. Automated calibration of agent-based immunological simulations. MIT Press. 847–845 (2013).
54. Read, M. *Statistical and Modelling Techniques to Build Confidence in the Investigation of Immunology Through Agent-Based Simulation* (University of York, UK, 2011).
55. Helton, J.C. *Uncertainty and Sensitivity Analysis for Models of Complex Systems. Computational Methods in Transport: Verification and Validation* (Springer, Berlin, 2008).
56. Vargha, A. & Delaney, H.D. A critique and improvement of the “CL” common language effect size statistics of McGraw and Wong. *J. Educ. Behav. Stat.* **25**, 31–32 (2000).
57. Saltelli, A., Chan, K. & Scott, E.M. *Sensitivity Analysis* (Wiley, New York, 2008).
58. Alden, K., Timmis, J., Andrews, P.S., Veiga-Fernandes, H. & Coles, M.C. Pairing experimentation and computational modeling to understand the role of tissue inducer cells in the development of lymphoid organs. *Front. Immunol.* **3**, 172 (2012).
59. Marino, S., Hogue, I.B., Ray, C.J. & Kirschner, D.E. A methodology for performing global uncertainty and sensitivity analysis in systems biology. *J. Theor. Biol.* **254**, 178–196 (2008).
60. Ray, J.C.J., Flynn, J.L. & Kirschner, D.E. Synergy between individual TNF-dependent functions determines granuloma performance for controlling Mycobacterium tuberculosis infection. *J. Immunol.* **182**, 3706–3717 (2009).
61. Bailey, A.M., Lawrence, M.B., Shang, H., Katz, A.J. & Peirce, S.M. Agent-based model of therapeutic adipose-derived stromal cell trafficking during ischemia predicts ability to roll on P-selectin. *PLoS Comput. Biol.* **5**, e1000294 (2009).
62. Keymeulen, B. et al. Insulin needs after CD3-antibody therapy in new-onset type 1 diabetes. *N. Engl. J. Med.* **352**, 2598–2608 (2005).
63. Sherry, N. et al. Teplizumab for treatment of type 1 diabetes (Protégé study): 1-year results from a randomised, placebo-controlled trial. *Lancet* **378**, 487–497 (2011).
64. Butler, J., Timmis, J. & Coles, M. Novel approaches to the visualization and quantification of biological simulations by emulating experimental techniques. ALIFE 14 Fourteenth Conf. Synth. Simul. Living Syst. **14** (2014).
65. Walker, D.C., Hill, G., Wood, S.M., Smallwood, R.H. & Southgate, J. Agent-based computational modeling of wounded epithelial cell monolayers. *IEEE Trans. NanoBioscience* **3**, 153–163 (2004).
66. Zhang, L., Strouthos, C.G., Wang, Z. & Deisboeck, T.S. Simulating brain tumor heterogeneity with a multiscale agent-based model: linking molecular signatures, phenotypes and expansion rate. *Math. Comput. Model.* **49**, 307–319 (2009).
67. Friedrich, C.M., & Savic, R.M. Facilitating drug discovery and development with mechanistic physiological models that are “fit for purpose”: introducing a model

- qualification method. <<http://www.rosaandco.com/pressReleases/rosaMQMWhitePaper2012.pdf>>.
68. Thorne, B.C., Bailey, A.M. & Peirce, S.M. Combining experiments with multi-cell agent-based modeling to study biological tissue patterning. *Brief. Bioinform.* **8**, 245–257 (2007).
 69. Polack, F.A.C. Choosing and adapting design notations in the principled development of complex systems simulations for research. *Proceedings of the Modelling of the Physical World Workshop.* **6**, 1–6 (2012).
 70. Franziska Hinkelmann, D.M. A mathematical framework for agent based models of complex biological networks. *Bull. Math. Biol.* **73**, 1583–1602 (2010).

© 2015 The Authors. *CPT: Pharmacometrics & Systems Pharmacology* published by Wiley Periodicals, Inc. on behalf of American Society for Clinical Pharmacology and Therapeutics. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Supplementary information accompanies this paper on the *CPT: Pharmacometrics & Systems Pharmacology* website (<http://www.wileyonlinelibrary/psp4>)