

Simulating biology: towards understanding what the simulation shows

Paul S. Andrews¹, Fiona Polack¹, Adam T. Sampson², Jon Timmis^{1,4}, Lisa Scott³, and Mark Coles³

¹ Department of Computer Science, University of York, UK
{psa,fiona}@cs.york.ac.uk

² Computing Laboratory, University of Kent, Canterbury, UK

³ York Centre for Immunology and Infection, University of York, UK

⁴ Department of Electronics, University of York, UK

Abstract. When building simulations of complex systems the task of validation is often overlooked. Validation helps provide confidence in the simulation by exploring the link between the models that we build and the real complex system. We investigate software engineering validation techniques from outside the area of complex systems to assess their applicability for the types of simulation we build. We then provide an example of how such techniques can be applied to a complex systems simulation of cells migrating from blood vessels into lymph nodes through the walls of the blood vessels. We suggest that explicitly stating the modelling and simulation assumptions we make is key to the process of validation. Concluding, we highlight a possible process for validating complex systems that explicitly incorporates environmental aspects.

Keywords: complex systems, simulation, modelling, validation

1 Introduction

Simulations are used to model complex systems such as biological phenomena, economies and human societies, for use in research investigation *in vivo*, *in vitro* and *in silico*. These systems are complex in the sense of having elaborate behaviour at a high level that is the consequence of

many simple behaviours at a lower level. The high-level behaviour cannot be deduced as a simple combination of low-level behaviours. Space, time and the environmental context are critical features.

In this paper, we present an initial engineering exploration in to the validity of a complex system simulation of part of the immune system, based on collaborative work with the York Centre for Immunology and Infection (CII), undertaken as part of the CoSMoS project⁵. We focus on a part of the conceptual model, and consider how, and to what extent, it is possible to validate this against the biological evidence of *in vivo* experiments. Our modelling assumptions, and validation problems, bring into consideration the importance of local environmental factors in modelling such complex systems. The analysis also highlights the limits of biological imaging based technologies that cannot provide dynamic insight in to the key features of the system; this in turn highlights areas where the biological research could focus.

Section 2 presents a review of simulation and its use in scientific research. The following section summarises biological research on the migration of lymphocytes to lymph nodes. We then explore the engineering validation of our simulation. Finally, in section 6, we propose an extension of Sargent's simulation process to express some of the complications of complex-systems modelling and simulation engineering.

2 Computer simulation: science and engineering

Computer simulation has been used to explore biological systems for many years. Traditional simulations generate output from equations (differential equations, Markovian models, etc.) that have been developed to mirror trends or behaviour in, for instance, biological populations. More recently, computer simulation has been used to model the possible effects (co-ordinated or emergent) of biological components or organisms acting in their environment. Essentially, simulations have two purposes: some are built in co-operation with research scientists in an effort to improve scientific understanding of natural systems; others are built as artificial systems to construct and explore alternative realities (either as science fiction or visions for future engineering). Here, we focus on a case study whose purpose is to contribute to scientific understanding in immunology permitting simulation of events that are difficult to experimentally validate *in vivo*. The aim is to produce a simulation in which the biologist has confidence and can help direct their wetlab experimental research.

⁵ The CoSMoS project, EPSRC grants EP/E053505/1 and EP/E049419/1, <http://www.cosmos-research.org/>, is building capacity in generic modelling tools and simulation techniques for complex systems.

Computer simulation of biological phenomena is important because static models cannot capture the dynamic features that characterise the behaviour of complex systems. For example, systems biologists are increasingly adopting conventional software engineering design diagrams to express static structures, and patterns of interaction in their models; these modelling approaches cannot express time, space or the features and consequences of large numbers of interacting instances [24]. Time and space (or at least spatial organisation) are essential to complex systems behaviour. A key aspect of space, which is also outside the remit of conventional modelling, is environmental interaction – both among components, and of components with their local environment: the behaviour of a complex system depends critically on the way that the (collective) components interact with their environment over time; failure to adequately model the environmental context naturally leads to non-realistic models of the complex system.

Thus, a scientifically-valid simulation has to extract suitable environmental aspects, at an appropriate level of abstraction; it also has to provide evidence that its environmental representation, as well as its scientific model, are adequate abstractions from the biological reality.

Despite the importance of realistic abstraction, the validity of computer simulations has not been a significant concern of simulators, except in the critical systems context. Typically, a simulation is judged by its ability to produce something like the expected results by a process that looks a bit like reality, and there is little concern for the quality or scientific relevance of the underlying simulation [11]. Simulations can be misleading, for instance because the output captures artefacts of the simulation software rather than patterns of behaviour, or where the attempt to approximate reality results in simulations whose complexity is as impenetrable as that of the observed system.

Scientific validity, like engineering validity, means that it must be possible to demonstrate, with evidence, how models express the scientific realities. Validity implies both adequate abstraction, and adequate development processes. Many computer simulations are poorly engineered: there is little attempt to record design of components or of the environmental context used in the simulation system. Assumptions and generalisations are not documented, and are thus not exposed to scientific scrutiny. An immediate result of this focus is a long-running intellectual debate about whether it is possible to do science through simulation (see [19, 22, 36, 5]). Bullock (in [36]) observes that, to assess the role and value of complex systems simulation, we need to address deep questions of comparability: we need a record of experience, of how good solutions are designed, of how to choose parameters and calibrate agents,

and, above all, how to validate a complex system simulation. To address these problems, we need principled approaches to the development of computer simulations.

For inspiration in engineering scientifically-valid simulations, we turn to two areas: non-complex critical systems, and agent modelling. Computer simulation of non-complex (but nevertheless *complicated*) systems has a long history, and the need to assure high-integrity and critical-systems models has led to a corpus of work on developing, verifying and validating simulation models. In intelligent agent modelling, there are methods for systems development that focus attention on the different aspects that need to be considered; some agent work has been taken forward for use in complex critical systems at the social, or macro, scale.

2.1 The process of simulation development

In high-integrity systems engineering, the foundations of a simulation process date from the late 1970s. For instance, Sargent (e.g. [28, 30]) presents a process (figure 1) that starts with a *problem entity*, or description of the phenomenon to be modelled. From the problem entity, a *conceptual model* is developed in a suitable representation – Sargent reviews diagrammatic models [29], and also notes mathematical or logical modelling paradigms [30]. Finally, a *computerised model* implements the conceptual model.

The simulation process is an iterative cycle, which includes an *experimentation* link between the problem entity and the computerised model. This allows the developers to trial-and-error simulation elements and settings, and to compare the results to the problem entity.

The second notable aspect of the simulation process is the explicit inclusion of verification (in the software engineering of the computerised model) and validation – both of models against reality, and of the data used to test the conceptual and computerised models. We return to this aspect in the next section.

The simulation process has much in common with conventional software engineering lifecycles – it presents a high-level summary of the necessary attributes of a development, rather than a comprehensive guide to achieving a high-quality engineered product. This area is addressed to some extent in agent-oriented modelling.

Sudeikat et al [33] give an insightful review of multi-agent system development methods, which, like Sargent, focuses on matching methods to the requirements of specific simulation targets. They identify as the foci of current methods: *internal architecture*, *social architecture*, *communication*, *autonomy*, *pro-activity* and *distribution*. Some of the re-

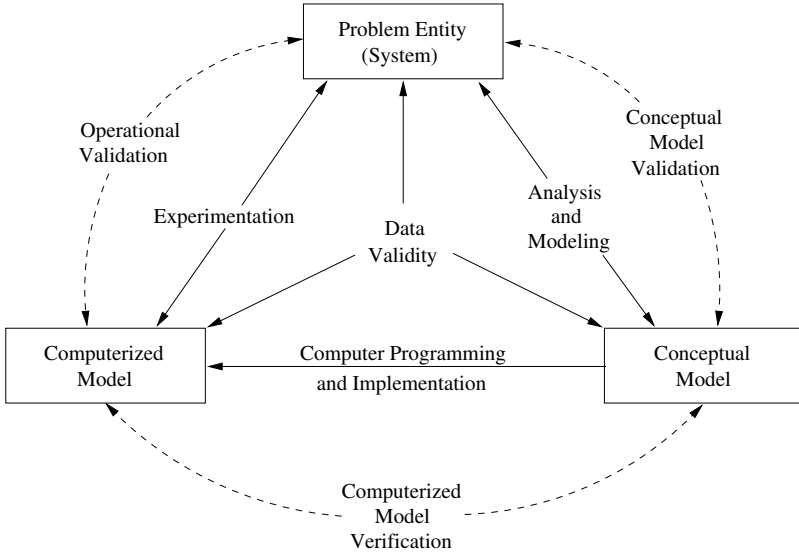


Fig. 1. Sargent’s model of the simulation development process [28]

viewed methods are sophisticated software-engineering approaches, such as Prometheus [21] (figure 2).

These methods provide an integration layer for the styles of model (usually diagrams) that are needed to express the static structures and interactions of the overall conceptual model. Implementation is often well researched, with platforms, patterns, and workbenches (see, for instance, the ACE resources, www.econ.iastate.edu/tesfatsi/ace.htm). Agent modelling methods are used in robotics, social agent systems, and similar areas, but are also entering high-integrity systems engineering: for instance, Alexander [1] uses Prometheus in the simulation of command-and-control systems. Critical systems use means that some work exists on adding validation activities to the modelling activities covered in the original methods.

Two immediate issues arise with the agent modelling methods. The first is that they are oriented to social systems – including human-scale high-integrity systems. This means that the technical emphases are on capturing the autonomy or design for learning – the BDI (beliefs-desires-intentions [12]) of agents. The second issue is that these methods do not capture the time, space, and component-quantity aspects of complex systems, or the layered abstraction aspect, noted above. The representation

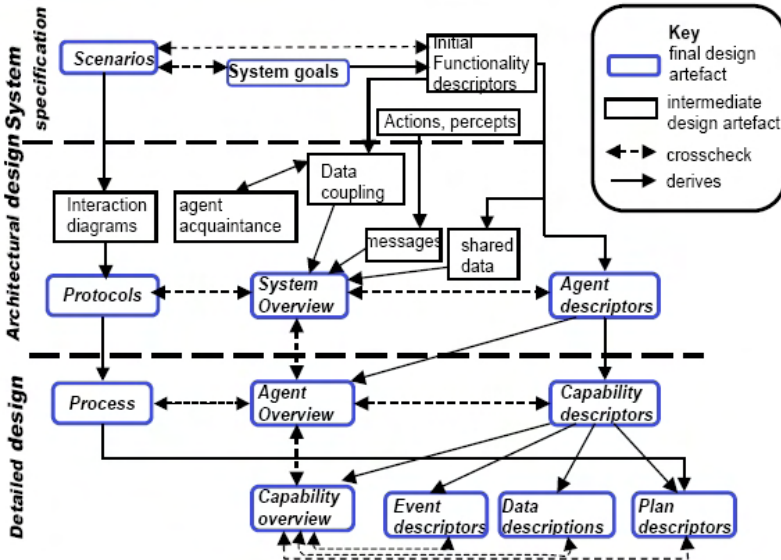


Fig. 2. The Prometheus development method [21]

of these key features are left to the inspiration of the implementer, and are thus hidden from validation scrutiny.

2.2 Verification and validation of simulations

There is a significant corpus of work from outside the area of complex systems on validating and verifying simulation models. However, although this work proffers general reminders, its direct advice is difficult to adapt to complex system simulation. Zeigler [39] presents a theory for modelling and validation of simulations; his theory is predicated on the fact of a homomorphism between conceptual models and simulations, and does not provide obvious pointers as to how the homomorphism is established – that is, verification of the development process is assumed.

Similar validation concepts come from Sargent, elaborating the validation aspects of the simulation process (figure 1, above). Sargent [30] reminds us that a *model should be developed for a specific purpose... and its validity determined with respect to that purpose*. Furthermore, Sargent notes that the level of assurance needed depends on the purpose of the simulation, and should be set independently of the development of the

Technique	Comments on Sargent's suggestions
Animation, operational graphics	Specifically, graphical visualisation, either of system behaviour or of operational parameters
Comparison to other models	Comparison to <i>valid</i> analytical models or other simulation models
Degenerate tests, extreme condition tests, parameter variability, sensitivity analysis	Typical domain-style testing of behaviour under normal and extreme input and operating conditions
Event validity	Compare the events in real and simulated systems
Face validity (ask a domain expert), traces	Appeal to logic or to domain experts to check the validity of model components or data
Historical data validation, predictive validation	Either drive a simulation with historical data and compare results to reality; or drive a simulation on current data and compare to independent predictions of future
Multi-stage validation, combining historical methods:	Three historical approaches can be combined to develop based on sound theory and assumptions, with empirical validity checks where possible.
Rationalism	The veracity of assumptions is rationally justifiable, and valid models arise from valid assumptions
Empiricism	All assumptions and outcomes are empirically validated
Positive economics	The model can predict the future, so causal relationships and mechanisms are of no concern
Internal validity	Used on stochastic models; comparison of consistency of results across runs
Turing tests	Can an expert tell that it is not the real system?

Table 1. Validation techniques for simulation development (based on descriptions from [30])

simulation – good software engineering practice. Sargent’s development process (lifecycle) for simulations explicitly incorporates verification and validation activities, and he proposes a range of approaches to validation, summarised in table 1.

Clearly, some of Sargent’s suggestions are inappropriate for complex systems work: if we knew the workings of the complex system well enough to understand event validity and traces, we would not need a computer simulation for research purposes. However, wherever such internal analyses are possible, they should be conducted. We need to be

confident that computer simulations accurately replicate contributory non-complex features. Comparison with real systems is essential, but is potentially dangerous – variants on predictive validation can lead to (accidental) construction of simulations that are self-fulfilling prophecies, whilst historical data validation tends to pick only the data that best match the simulation.

Perhaps the most useful of Sargent’s suggestions relate to analysis of assumptions – though situating this in historical theories of rationalism and empiricism tends to mask their value. A common (possibly universal) failing of research simulations is the failure to document the assumptions that they make, both about the science that underpins the models, and about the means used to create the simulation. We will return to assumptions in the case study, below.

2.3 Other computer simulations for biological research

There are a number of current interdisciplinary research projects that use aspects of software engineering to produce high-quality computer simulations to support biological research.

PEPA [6, 7] is typical of several approaches that use stochastic process algebra to construct models of cell signal transduction pathways. In PEPA, complementary models are developed of a reagent view and a network view. The models are proven isomorphic with each other, and isomorphic with conventional differential equation models of transduction. Verification (that the implementation captures the conceptual model) is formal and explicit, and the ability to mimic the analytical models that the biologists create contributes to the validation of the conceptual model against the reality. This is akin to Sargent’s *comparison to other models* technique, although, in common with other differential equations, the validity of the analytical models is difficult to show.

Reactive Animation (RA) [10, 9, 14] is another robustly-engineered research simulation; it uses Rhapsody statecharts (state machines) and Live Sequence Charts (connectivity diagrams), plus data from biological experimentation, to drive powerful biological visualisations (for instance, of T cell activity in the thymus). RA reverse engineers biological systems, using a well-understood software engineering analogy [27]. RA thus exemplifies a number of Sargent’s comparatively-based validation techniques, as well as quality software engineering design and verification. In both PEPA and RA, high-quality computer engineering and attention to validation produces simulations that biologists can rely on to direct their research.

Whilst these initiatives are both scientifically and computationally successful, they are not easily generalisable. The specialised components

(such as process algebras) and proprietary tools (such as Rhapsody state charts) tie the initiatives closely to the groups that own them. This makes it hard to do a comprehensive evaluation of PEPA or RA as candidate for a general complex system development process – in other words, it is hard to generalise from these otherwise excellent initiatives in biologically-driven computer simulation.

In our work, we take inspiration from Sargent’s process and RA modelling. Our conceptual model is a very simple form of state machine, to express the possible states of a lymphocyte. Like conventional agent modelling, much of the environmental context is captured between the conceptual model and the implementation. However, we then apply Sargent’s validations, and a deviational approach to assumption generation by providing evidence for arguments (based on work by Pumfrey [25], Srivatankul [31], Allenby [2] and others applying deviational analyses to safety or security assurance work). This approaches to engineering assurance not only reveals limitations (and strengths) of our simulation, but can also be used to explore the effect of limitations of the biological knowledge of the system.

The following section provides biological background on the case study. We then review our simulations, and discuss connotations of our findings for complex systems modelling.

3 Migration of lymphocytes in the lymph node

3.1 Lymph Nodes

The mammalian immune system possesses many specialised cells that are collectively known as the leukocytes or white blood cells. The leukocytes can be divided into a number of distinct groups with different functionalities. One such group is the lymphocytes, which are vital to recognising and mounting an immune response to various harmful pathogens such as bacteria and viruses. The lymphocytes can be further classified into two distinct populations of cells: B cells and T cells. As well as specialised cells, the mammalian immune system also comprises a number of immune organs. One such organ is the lymph node, which is a small (about the size of a pea in humans) bean-shaped immune organ (figure 3) rich in lymphocytes and other leukocytes, providing a place where immune response to pathogens in the lymph may be triggered and develop. The structure of the lymph node is made up of a number of specialised areas supporting different cellular environments. There are hundreds of lymph nodes in various locations around the body.

Bodily fluid known as lymph drains into lymph node through a number of afferent lymph vessels connected to the *lymphatic system*; lymph




Fig. 3. Structure of a lymph node from [16]

leaves the lymph node through a single efferent lymph vessel. Lymph contains many different leukocytes, proteins and other particles (possibly pathogenic) that have drained from the peripheral parts of the body. The lymph node is also connected to the circulatory system via a lymphatic artery and vein. It is through the lymphatic artery that lymphocytes enter the lymph node. Lymphocytes can then migrate through specialised blood vessels (via a mechanism described below) into the functional tissue of the lymph node. Once there, lymphocytes can interact with other leukocytes that have encountered pathogens and entered the lymph node from the lymph, to initiate an appropriate immune response.

3.2 Endothelial Cells, Pericytes and High Endothelial Venules

Figure 4 summarises the main types of blood vessel present in the body: arteries are the large vessels that carry oxygenated blood from the heart; arterioles branch off the arteries carrying blood to the capillaries; capillaries are the smallest blood vessels in the body and allow the interchange of components between the blood and body tissues; venules carry the de-oxygenated blood from capillaries; and veins then carry de-oxygenated blood back to the heart.

The internal surface of all blood vessels is made up of endothelial cells. In the lymph node, a minority of the venules have plump (high)




Fig. 4. The relationship between the different types of blood vessel from http://training.seer.cancer.gov/module_anatomy/images/illu_capillary.jpg




Fig. 5. A comparison of a normal venule with a lymph node HEV from [20]. In the HEV, the ring of endothelial cells are much larger.

endothelial cells that have a significantly larger diameter than normal endothelial cells. These venules are called the high endothelial venules (HEVs) [15] (see figure 5). Areas of HEVs occur in the lymph node at various points (figure 6).

It is in HEVs that lymphocytes can migrate from the blood through the endothelial cells into the functional tissue of the lymph node. Only lymphocytes can interact with and cross HEVs, other leukocytes are excluded [20]. It is estimated that a quarter of the circulating lymphocytes migrate from the blood after entering an HEV [13].

HEVs (and other small blood vessels) are surrounded by pericytes, shown in figure 7. They are a form of vascular smooth muscle cell surrounding endothelial cells that are responsible for constriction and dilation of blood vessels. This regulates blood flow and diameter of the HEV, and thus affects the ability of lymphocytes to migrate [26]. A large influx




Fig. 6. The flow of cells in and out of a lymph node from [20]. Lymphocytes and dendritic cells enter the lymph node by two routes. Most dendritic cells enter through afferent lymph, settling near HEVs in paracortex (due both to lymph node structure and local production of chemokines). Most lymphocytes enter the lymph node across HEVs.




Fig. 7. Cross-section of a capillary from [4]. EC = endothelial cell, PC = pericyte, vSMC = vascular smooth muscle cell.

of lymphocytes into the lymph node during an immune response causes it to visibly swell. This is known as lymph node hypertrophy.

3.3 Lymphocyte Rolling and Migration in the HEV

All leukocytes are able to migrate through blood capillaries via the same mechanism. Only lymphocytes, however, can migrate through the specialised HEV in a lymph node. The process of migration is due to an adhesion cascade of various cell surface receptors and molecules. Originally identified as a three-step process of rolling, activation and arrest, the migration process has now been augmented, as shown in figure 8. Each step is initiated and regulated by specific signalling molecules and receptors [18].




Fig. 8. The leukocyte adhesion cascade from [18]. The three historical steps are shown in bold.




Fig. 9. Migration of T-cells across HEVs from [16]

To migrate, a leukocyte must pass through endothelial cells, the endothelial-cell basement membrane and pericytes. Migration through endothelial cells can be rapid (2-5 minutes), whilst penetrating basement membrane takes longer (5-15 minutes) [18].

Migration of lymphocytes from HEVs is controlled by a specific array of adhesion molecules that facilitates lymphocyte migration but bars other leukocytes (figure 9). Chemokines (chemical signalling molecules) that are produced by or adherent to HEVs are important in the control of lymphocyte migration, however, the precise mechanisms by which chemokines work *in vivo* are unclear [16]. Several chemokines are probably required for the movement of lymphocytes through HEV, but how many and in what order is unclear [20].

4 Developing a lymphocyte migration simulation

In section 2.2 we highlighted the importance of explicitly stating the purpose of a simulation as it will directly influence the design decisions and assumptions we make. Our aim is to model the migration of lymphocytes through the high endothelial venules (HEV) of the lymph nodes, and construct computer simulations of the lymphocytes as a complex system. Using these simulations we can investigate possible factors that lead to the observed lymph node hypertrophy during infection. We are interested in how the number of migrating lymphocytes changes under different conditions, thus the desired output of any simulation will be numerical data detailing lymphocyte migration rates. The simulations should enable us to test hypotheses such as *the increase in lymphocytes in the lymph node during infection is due to dilation of the high endothelial venules*.

In this section, we outline the processes that we use to develop the simulations; in the section 5 we turn to validation. Following Sargent's process, figure 1, the biological literature summarised above provides the starting context for our case study, the *problem entity*. Our *conceptual model* is first described followed by our *computerised model* or *simulations*. To aid the analysis of our modelling and simulation process, we have laid out our assumptions in table 2. In the descriptions that follow, we refer to this table where necessary.

4.1 A conceptual model

The first step in building the conceptual model is to identify the different parts of the system in which we are interested. In terms of a complex system we can consider the hypertrophy of the lymph node to be the non-linear (emergent) behaviour under investigation. The main active component of our system is a population of homogeneous lymphocytes, which interacts in an environment. This environment is simply the parts of the body with which the lymphocytes interact with respect to their migration through HEV in the lymph node.

In terms of the biology, the main cellular actors in the migration process are the lymphocytes, high-endothelial (HE) cells and pericytes. We consider HE cells and pericytes together in the form of a tube (a high-endothelial venule). Lymphocytes travel through the HEVs within the blood circulation. Outside the HEVs is the lymph node tissue, which the lymphocytes enter if they successfully migrate through the HEV. We have classified the different environments that the lymphocytes pass through as states, with transitions occurring when the lymphocyte moves

Label	Assumption
1	The detail described in section 3 is correct
2	Lymphocyte migration only takes place in the HEV areas of the lymph node
3	There is no interaction between lymphocyte agents. Lymphocytes do not collide.
4	There are no effects from external blood circulation e.g. blood flow is constant.
5	The volume of blood always there to accommodate size of HEV (i.e. enough blood to fill whatever size HEV expands to).
6	Once a suitable chemokine signal has been received by a lymphocyte, it will always migrate. Thus, subsequent stages in the adhesion cascade (see figure 8) are deterministic.
7	Lymphocytes are essentially equivalent. They express the same levels of receptors required for rolling and migration.
8	Lymphocyte will always re-enter blood circulation from the lymph node. Thus, lymphocytes can only exit the system (die) in the blood circulation state.
9	Lymphocytes are created and die at a constant rate.
10	The HEVs are homogeneous. The endothelial cells and pericytes that make up the HEV all behave the same making the HEV appear the same at all points.
11	Lymphocytes flow through the HEV at the same rate.
12	Lymphocytes can be captured and disassociate repeatedly whilst in the HEV.
13	There is no change in lymphocytes that have disassociated from an HEV wall, thus all free flowing lymphocyte are equally likely to capture.
14	Whilst passing through the HEV there is no change in the state of the lymphocyte, thus there is no distinction between new and re-circulating lymphocytes.
15	Proliferation does not occur in the lymph node.
16	The multi-stage adhesion cascade shown in figure 8 can be reduced to two probabilistic stages: capture leading to lymphocyte rolling, and migration after receiving a chemokine activation signal.
17	The number of lymphocyte chemokine receptors does not change on the time scale of the simulation.
18	Lymphocytes drain from the lymph node to the blood circulation at a constant rate.

Table 2. A list of many of our modelling and simulation assumptions. Each is given a label so that we can refer to it. This is not a complete list, but is illustrative of the kind of assumptions we make.

from one environment to the other. This is summarised by the *state diagram* in figure 10. The four key identified states are:

Blood Circulation: This state encompasses the parts of the body that the lymphocyte is in when it is not in the HEV or the lymph node tissue. It provides a place where lymphocytes can enter or leave the system. It is the state lymphocytes will be in for most of their existence. Assumptions 4, 5, 8 and 9 apply to this state.

HEV Lumen: This state describes the lymphocyte when it is flowing freely in the lumen of a HEV. Assumptions 10, 11, 13 and 14 apply to this state.

Rolling: This state represents the lymphocyte when it is rolling on the interior surface of an HEV (see figure 8). Assumptions 10 and 13 apply to this state.

Lymph Node: This state describes the lymphocyte when it is present in functional tissue of a lymph node. Assumptions 15 and 18 apply to this state.

In addition to these four states, **Start** and **Stop** states provide a means to introduce and remove a lymphocyte.

The state transitions in figure 10 map to the biology in the following ways:

Creation A newly created lymphocyte will automatically transition into the *blood circulation* state. Assumption 9 applies to this transition.

Enter HEV As a lymphocyte is transported around the body in the blood (the *blood circulation* state) it will at some point enter an area of HEV a lymph node. Assumptions 10 and 11 apply to this transition.

Exit HEV Just as a lymphocyte can enter the HEV lumen, it can also exit the HEV lumen via the blood, transiting from the *HEV lumen* state back into the *blood circulation*. Assumptions 10 and 11 apply to this transition.

Capture Whilst moving freely in the HEV lumen, a lymphocyte captures onto the endothelial wall transiting to the *rolling* state. Assumptions 7, 12, 13 and 16 apply to this transition.

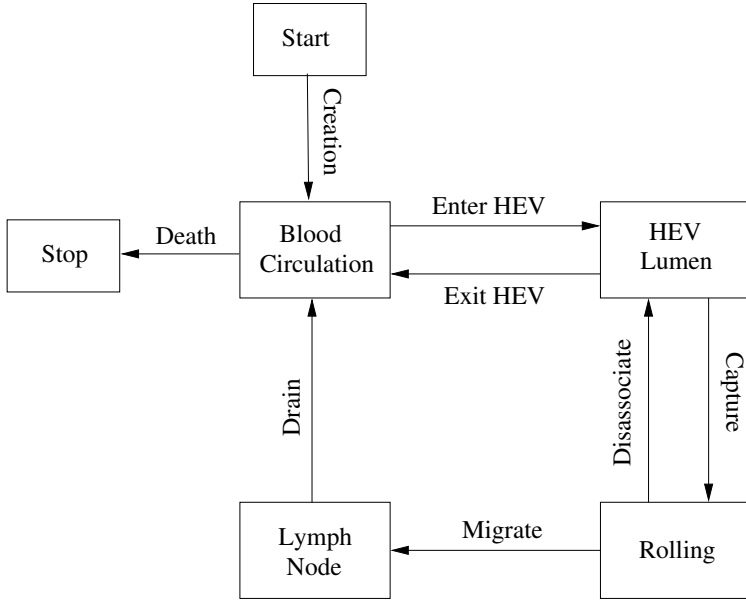


Fig. 10. A generic state transition diagram for a single lymphocyte. Boxes represent states a lymphocyte can be in and arrows represent the possible transitions between states.

Disassociate Just as the lymphocyte can transition from flowing freely in the lumen to rolling, it can also disassociate from *rolling*, moving back to the *HEV lumen*. Assumptions 7, 12, 13 and 16 apply to this transition.

Migrate During the process of rolling, a lymphocyte receives a chemokine signal from the endothelial surface of the HEV. If sufficient, this signal produces a change in the confirmation of receptors on the lymphocyte leading to a cascade that results in migration of the lymphocyte from the HEV to the functional tissues of the lymph node. Assumptions 2, 6, 7 and 16 apply to this transition.

Drain After spending time in the lymph node, a lymphocyte drains in to the lymphatic system via the efferent lymphatic vessel to rejoin the blood circulation (see figure 6). Assumption 18 applies to this transition.

Death Whilst circulating in the blood, a lymphocyte dies. Assumptions 8 and 9 apply to this transition.

Only the allowable transitions are shown, for example a lymphocyte in the *blood circulation* state can only enter the HEV via the *enter HEV* transition or cease via the *death* transition. In addition, as this is a conceptual model, the eight state transitions have a meaning, but no specific values. The values are assigned at the simulation stage, taking the form of probabilities.

The main simplification (see Assumption 16) we make with regards to lymphocyte rolling is to reduce the multi-stage adhesion cascade (represented in Fig. 8) down to two main steps. The first step, captured by the *capture* transition, models the capture of lymphocytes on to the endothelial wall. Once capture has occurred, the lymphocyte is in the *rolling* state, waiting to be activated by a chemokine signal. The second step, expressed by the *migration* transition, models the lymphocyte receiving the chemokine signal; after this it is assumed that the lymphocyte succeeds in migration. Other stages in the cascade are assumed to be either deterministic, or have such small probabilities of failing that they are insignificant.

4.2 Simulations

A simulation is best thought of as an execution of a model (such as the lymphocyte model in figure 10) over time. Typically time is implemented as atomic steps, at which each element in the simulation (each lymphocyte) updates. Within a simulation, we need to define rules to determine when a lymphocyte can transition between states. One way to achieve this is to assign each transition with a probability of occurring. These probabilities need to be extracted from the biological detail to represent what is known to happen in the biology.

We have developed two simulations of the conceptual lymphocyte migration model described above, which differ in the way spatial aspects of the environment are represented. In the first, there is no explicit coordinate system, only the four body locations in figure 10. Each of these four state spaces can contain a number of lymphocyte agents, which transit from location to location based on a set of rules. This simulation aims to capture statistically the change in lymphocyte concentrations in the lymph node as the probabilities on the capture and migration transitions change.

In the second simulation, the 3-dimensional HEV tube made up of endothelial cells, and the movement of lymphocytes through that tube, are explicitly implemented. This supports visualisation of the HEV and of the lymphocytes, with the *HEV lumen* and *rolling* states visually distinguishable via changes in colour. Again, the model is driven by the




Fig. 11. Process diagram showing how lymphocyte processes (represented as circled “L”s register with state processes (boxed)

probabilities on the transitions, but because the spatial aspect is expressed explicitly, the simulation is more obviously closer to the biology.

We refer to the two different simulations as *migration-abstract* and *migration-space* respectively. The description in this section focuses on the *migration-abstract* simulation, but draws on elements of the *migration-space* simulation when necessary. Our simulations are implemented using occam- π , a process-oriented programming (POP) language capable of massive concurrency (for details of why we choose this approach see [3]).

The migration-abstract simulation implements the state diagram shown in figure 10 for a population of lymphocyte agents over a period of time. Using the POP paradigm, each lymphocyte agent is represented as a process which is connected via a *communication channel* to one of the four body place states we are interested in (also represented as processes): blood circulation, HEV lumen, rolling on the endothelium, and in the functional tissue of the lymph node. The process network structure is shown in figure 11, and directly reflects the topology of the states in figure 10. Each of the four place state processes shown has a *shared* communication channel (shown as a bold arrow), to which any number of lymphocyte processes can be connected. Each lymphocyte process is connected to one and only one place process thus ensuring it can only be in one state at a time. According to process-oriented design rules, the lymphocyte processes act as clients to the server state processes.

Depending on the transition rules, a lymphocyte process can change the state process to which it is connected, moving from one place process to the next. Channels exist between state processes to enable this movement, which directly reflect the allowable state transitions in figure 10. For example, in figure 12 the shaded lymphocyte process moves from the blood circulation state to the HEV lumen state by disconnect-




Fig. 12. Process diagram showing how lymphocytes move between state processes

ing its communication channel from the blood circulation process and reconnecting it to the HEV lumen process.

Every lymphocyte process updates once per time step (this is achieved via the *occam- π* barrier construct). At each step, the number of lymphocyte processes associated with each place process can be recorded, to allow numerical analysis of the number lymphocytes in each state over time. A typical simulation run necessarily contains many thousands of lymphocyte processes in order to get close to the biological scale which contains millions of lymphocytes. New lymphocyte processes can be added during a simulation run and are automatically connected to the blood circulation process.

The eight lymphocyte state transitions (figure 10) in the migration-abstract simulation are each encoded a probabilistic rule. At each time step, a lymphocyte tests for its possible for transitions. For example, a lymphocyte in the rolling state can either disassociate to the HEV lumen state, migrate to the lymph node state or stay in the rolling state. To determine which of these occurs, a random number is generated for each lymphocyte and is tested against the possible transitions. If the random number falls in the range of a transition, then the lymphocyte will transit into the relevant state. If not, the lymphocyte stays in its current state.

There is a mapping between probabilities and the biological detail. To achieve an accurate simulation, we need to choose the probabilities carefully, and try to validate to be confident that they represent what they are supposed to. It is ongoing work to find good probabilities, the job of which is not trivial for a number of reasons. Often the exact biological details are poorly understood or simply not recorded. Facts can also come from many different sources, based on different experiments utilising different technologies and subjects. Many of these facts then have to be combined into a single probability. Thus when constructing these probabilities we need to document where they have come from and

how they map to the probabilities. An example of the type of detail we would use to construct a probability is that a quarter of lymphocytes entering HEVs will migrate [13]. By combining this with the numbers of lymphocytes in the body we can start to build meaning into the capture and migration probabilities.

The following descriptions examine each transition probability and suggests the type of biological detail we would need to generate their values in a simulation:

Creation: The probability for creation is distinct from the other transition probabilities as it is a function of the population size rather than the individual lymphocytes. Studies in mice show that about 1 to 2 million T cells and B cells enter the blood circulation [16]. In a homeostatic environment, the numbers of lymphocytes created should equal the number that die (see death transition below).

Enter HEV, Exit HEV: Lymphocytes continually enter areas of HEV whilst circulating in the body. The probability of a particular lymphocyte doing so needs to reflect an average amount of time not spent in HEVs. It is consequently inextricably linked to the exit HEV probability. The probability needs to encapsulate biological details such as the relative lengths of the blood circulation system and HEVs, along with the rates of blood flow.

Capture: In the description of the conceptual model, we discussed what capture means and the assumptions involved. In the simulation, the capture probability relates to the biology of the receptors of lymphocytes and endothelial cells, and the probability that the receptors on each are close enough to interact. The interaction probability can be based on the relative sizes of lymphocytes and the diameter of the HEV lumen taken from the biological data. For example, a smaller lumen diameter would increase the chances of a lymphocyte being close enough to interact. Conversely, a larger diameter may relate to a larger surface area of endothelial cell receptors to which the lymphocytes can attach, so a larger diameter would increase the probability of capture.

Disassociate: This probability is related to the capture probability in that it takes into account the strength of binding between lymphocytes and the endothelial cells. A higher concentration of endothelial cell receptors should reduce this probability.

Migrate: The migration probability is dependent on the concentration of chemokines to induce the confirmation change in the lymphocyte and/or the likelihood that the lymphocyte picks up the chemokine signal. The last probability could take into account the numbers of chemokine receptors on the lymphocyte, and the way in which endothelial cells express and present them. Assumption 17 applies.

Drain: The drain probability reflects the amount of time it takes for a lymphocyte that has migrated into the lymph node to exit to the blood circulation. This time can be dependent on whether the lymphocyte is activated, but as this frequency is usually very small, we assume the rate is constant for all lymphocytes (Assumption 18).

Death: A lymphocyte that dies will be removed from the simulation. The probability to determine this needs to take into account the average life of lymphocytes. In a homeostatic environment, death is compensated by creation. Assumption 9 applies.

All probabilities need to be scaled to fit the time step of the simulation. Each probability is also simulation specific. Our migration-space simulation implements the same conceptual model as the migration-abstract simulation, employing the same state transitions, and has a need for the same relevant probabilities. However, the actual probability values are subtly different. For example the capture probability in the migration-abstract simulation incorporates the need for a lymphocyte to be near to the endothelium for capture to occur. In the migration-space simulation, the 3-dimensional space is explicit, thus the capture probability only encapsulates the biology of receptor binding. The process of validation needs to explicitly highlight the contributions to transitions, to make the simulations transparent and open to reasoning.

5 Verification and validation of the simulations

According to Sargent's process, we need to validate the conceptual model against the problem entity (and the purpose of the simulation). Ultimately, we also have to argue the operational validity of the computerised model (i.e. determine the behaviour has sufficient accuracy for its intended purpose), but that is largely outside the scope of this paper. As we will see, validation of the conceptual model reveals the many gaps in both the science and the computerised model, these gaps inhibit exploration of any research hypothesis.

We can think of the validation process as producing an argument of validity, in the same way that critical systems developers produce arguments of safety, dependability or security. Note that just as a safety argumentation never establishes that a system is absolutely safe (no system is safe unless it is totally closed and inert), an argument of validity merely states the case for validity, exposing it to critical consideration. This is an important observation, because, in any natural complex system, we cannot expect to provide a gold-plated guarantee of equivalence between our conceptual model and the problem entity – indeed, if the model contained all the complexity needed to exactly mimic the natural system, it would be intractably large, and too complex to provide any new research insight.

An argument is expressed as a *proposition*, and is reasoned on the basis of some *premises*, to reach a *conclusion*. A variety of textual and diagrammatic techniques allow an argument to be presented with a degree of formality – exposing the premises to analysis and scrutiny (see [17, §2.6] for a succinct review). Kelly and others [17, 35] adopt the goal structuring notation (GSN), proposed in [37], to present arguments including those relating to the safety or safe design of critical systems [17] and system dependability [8], and to emphasise the role of argumentation in design [38].

In conventional safety case argumentation, basic elements are used to express an argument: goals (decomposable), strategies (to meeting the goals), justifications, assumptions, contexts, and solutions. Here, our goal is to validate the conceptual model against the problem entity. For illustration, we focus on one aspect of the conceptual model, the transition labelled *Capture*. This reveals many of the issues in validating a complex system model against a natural problem entity. Whilst we do not present a systematic analysis here, the approach could be systematised, applying a deviational approach in the way that is common in safety work (see [25]) and has more recently been used in security analysis (for example, [31]).

5.1 The Capture transition and its connotations

The *Capture* transition takes the lymphocyte from the state *HEV Lumen*, where it is moving freely within the HEV, to the state *Rolling*, where it is in the preliminary stage of the migration process. From the *Rolling* state, a lymphocyte can revert to the *HEV Lumen* state by disassociation, or continue its migration to the functional part of the lymph node. We might equate *Capture* to the second stage shown in figure 9, above.

The *Capture* transition is an abstraction from the cell biology and biochemistry – which has been well-researched (see the top panel of figure 9, above). However, validation of the *Capture* probability would require a separate, lower-level simulation of the capture bio-chemistry (or *in vivo* research). In our validation argument, the probability of capture is an area that we must expose to external review, or to further work.

However, there are other aspects of capture that we must validate, where the biological basis for our conceptual model is less well understood. We might postulate a context for capture: the chance of a lymphocyte being captured depends on (a) the density of active receptors on the endothelium; (b) the receptors on the lymphocyte; (c) the likelihood of a lymphocyte being close enough to the endothelium for capture. Thus, our simulation needs to take into account the geometry (as well as the biochemistry) of the lymphocyte and the HEV, as well as flow characteristics in the venule. Validation has a choice here. On one hand, we could attempt to simulate the three-dimensional structural biology of a HEV – capturing typical venule cell structures and their behaviour as the HEV constricts and relaxes (an animated version of figures 5 and 7) – and analysing the flow and contact characteristics that determine lymphocyte interaction with the endothelium chemicals. On the other hand, we could state our assumption that the probabilities on the transition from flowing in the *HEV Lumen* to captured in the *Rolling* state adequately captures the geometric and flow aspects of the HEV. The first option tends to improve the biological realism (and validatability) of the simulation, at the cost of complicating the conceptual basis of the model. If we opt for simplicity, we must record the assumption – which come under a general heading of *environmental factors* – so that scientists appreciate the limitations of the simulation.

In the conceptual model, the transition probabilities are constant. Each probability approximates the effect of many environmental factors, and each factor must cause fluctuations in the rate of transition in the short term. However, we assume that behaviour tend towards the norm over the timescale of the simulation (or the real biology). Again, we could postulate further research or low-level simulations to validate this assumption: for instance, we would like to understand the effect of irregular coverage of pericytes – does this give rise to an uneven longitudinal profile in the HEV, and if so, does this promote lymphocyte capture at upstream locations?

This last question identifies a potential paradox in our conceptual model: we assume constant probabilities of transition, so all lymphocyte in the *HEV Lumen* state have an equal probability of being captured; equal probability implies a homogeneous HEV environment. However,

the hypothesis seeks to relate lymphocyte volume in the lymph node with dilation of the HEV, so we are required to vary the cross-section (at least) of the HEV, and it is not obvious that this is consistent with a homogeneous HEV environment (chemically or geometrically).

As we explore the question of the HEV environment further, we discover that one possible conformation of the contracted HEV has the lining cells packed in tight folds; the folds relax as the HEV dilates. There must, therefore, be an intermediate point at which pockets arise between folds, which we might expect to trap lymphocytes, increasing the chance of capture. Taking an opposing view, widening the HEV means that lymphocytes are more able to move away from the venule walls, potentially causing a fall in capture. We cannot currently validate any of the assumptions about the geometry and flow of the HEV environment, but we can highlight these assumptions in relation to our simulation results.

We assume that all lymphocytes are the same size – that is, there is no differential probability across lymphocytes. This can be validated biologically: the literature gives size ranges for lymphocytes, and we can determine (by asking immunologists!) whether the range represents sizes with one lymphatic system or across a species (or what): the simulation can be driven accordingly. It is noted that much of the data taken from the literature and used for our simulation was gathered for purposes other than ours. We must, therefore, assume that the data is still applicable in our domain.

Returning to the probability of the *Capture* transition, we have used the conceptual model as the basis for two simulations. In the abstract simulation (non-spatial), the transition probability must be used to reflect the effects of all factors in HEV environment and other relevant environmental factors. However, in the spatial model, the spatial relationship between lymphocyte size and location and venule diameter and conformance is incorporated directly – the transition probabilities abstract only from the biochemistry and the surface characteristics of the HEV.

We have not identified all assumptions of our conceptual model here. For example, the model also abstracts from all biochemical factors: we assume that the probability of transition expresses any underlying variability in receptor form, binding strength, mechanisms of binding and expansion, etc. We also abstract away from the cascade details of the rolling state and the capture, disassociation and migration transitions.

5.2 Breaking with assumptions

Validating the conceptual model against the problem domain (the biology) can reveal inconsistencies in the biological detail and mismatches

with our assumptions. For example, we may have based a probability on data that has been revisited and altered by subsequent scientific research; mixed data from incompatible pieces of research; used data from fixed biological material not appropriate to the simulation. By presenting our model and simulation assumptions, and the biological detail that has influenced the design of our simulation, the validation process is made easier.

If we find a problem in an assumption that breaks the model, we need to analyse the effects and change the model accordingly. For example, assumption 15 states that proliferation (the generation of cellular clones) of lymphocytes does not occur in the lymph node. Based on experimentation with the simulator, we might decide that this has an effect. Consequently we can update the conceptual model with extra states and transitions. These changes then cascade through subsequent models to the simulator. We need to check that changing one assumption or value does not affect the others, and if it does then change these accordingly. Most incorrect assumptions will not invalidate the entire model or simulation, but just require editing. Tables of assumptions and biological details allow traceability through the modelling and simulation process. For example a domain expert might inspect them and highlight areas of inconsistencies leading to the model or simulation to be updated.

6 Discussion

Our conceptual model of the lymphocyte system represents a set of design decisions: we have chosen to abstract to certain (key) states of the lymphocyte lifecycle, and we have selected probability-based transitions as a suitable basis for experimenting with lymphocyte concentrations and HEV dilation. The consequent validation requirements are clearly dependent on these design decisions. Had we chosen to model at a different level of abstraction, or to represent lymphocyte behaviour differently, we would have different validation requirements (but a similar range of problems relating our model to the problem entity).

Disparate levels is an inherent problem of complex systems research and simulation. The validation proposals suggest that some of the conceptual model features could be validated by either lower-level simulation, or exploration of the biology (and biochemistry) at the lower level. This idea is also inherent to the CoSMoS project – the modelling and simulation platform that it seeks to develop has as one goal support for multiple simulation levels. A lesson of validation might be that we need to identify component *levels*, as well as component state-and-operations, so as to facilitate validation where biological detail is uncertain.




Fig. 13. An architecture for complex system engineering, after [23]

Elsewhere, we have proposed an architecture, and some principles, for the engineering of simulations exhibiting several layers of complex behaviour (see [34, 32, 23]). The architecture, figure 13, for complex systems proposes component specifications, higher-level system specifications, and a reconciliation (or linguistic integration) via strategic parts of the common environment. Putting together these architectural ideas with the need to identify environmental factors in determining the conceptual model (and the implementation detail of the subsequent computerised model), we propose an extension to Sargent’s process, as shown in figure 14.

There is much work still to be done in the area of validating our complex system models and simulations. This includes establishing structured ways to layout our assumptions and biological details that have influenced our designs. We also need to establish schemes for mapping between our biological details and simulator parameters. In addition we are investigating structured argumentation techniques to talk about validity of complex systems simulations. Our aim is to establish patterns of validation that are applicable to the validation of many different complex systems.

7 Summary

We have presented a selective review of engineering approaches to engineering simulations in non-complex systems and agent systems, and used




Fig. 14. Adding environmental concerns to Sargent's process (figure 1), to propose a process for complex system simulation

this to initiate a case study simulation development for part of the lymphocyte system. The simulation is based on information from biological literature and part of an ongoing research project in the CII.

Proposing the biological information as our problem entity, and exploration of proposed immune mechanisms as our goal, we have presented a simple conceptual model (from which two simulations of different abstractions have been created). Focusing on one part of the conceptual model, we discuss issues relating to validation of biological research simulations.

Drawing on work in non-complex systems simulation and agent systems modelling, we have identified possible features of a process for engineering complex systems simulations. The process is speculative, but fits our experience of simulating part of the immune system.

The paper presents a first step in an engineering approach towards scientific computer simulation; the findings are preliminary and not yet substantiated by repetition or systematic use. However, our findings are well grounded in wider work on critical systems engineering and assurance, as well as other areas of system simulation.

From the brief exploration of validation, it is clear that arguments of validity for research-oriented simulations of complex systems are going to be complicated, and often incomplete – the principle of exposing assumptions to external scrutiny is an important contribution of this paper. We expect that argumentation approaches, and deviational anal-

ysis, will contribute to the quality and visibility of validation. In short, we believe that the scepticism over use of computer simulation in complex systems research can be addressed through well-established engineering principles, just as it is being addressed in macro-scale complex systems.

8 Acknowledgements

This work is part of the CoSMoS project, funded by EPSRC grants EP/E053505/1 and EP/E049419/1.

References

- [1] R. Alexander. *Using Simulation for Systems of Systems Hazard Analysis*. PhD thesis, Department of Computer Science, University of York, 2007.
- [2] K. Allenby and T. P. Kelly. Deriving safety requirements using scenarios. In *5th IEEE International Symposium on Requirements Engineering (RE'01)*. IEEE Computer Society Press, 2001.
- [3] Paul S. Andrews, Adam T. Sampson, John Markus Bjorndalen, Susan Stepney, Jon Timmis, Douglas N. Warren, and Peter H. Welch. Investigating patterns for the process-oriented modelling and simulation of space in complex systems. In *To appear: Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*. MIT Press, 2008.
- [4] G. Bergers and S. Song. The role of pericytes in blood-vessel formation and maintenance. *Neuro-Oncology*, 7(4):452–464, 2005.
- [5] J. Bryden and J. Noble. Computational modelling, explicit mathematical treatments, and scientific explanation. In *Artificial Life X*, pages 520–526. MIT Press, 2006.
- [6] M. Calder, S. Gilmore, and J. Hillston. Modelling the influence of RKIP on the ERK signalling pathway using the stochastic process algebra PEPA. *Transactions on Computational Systems Biology VII*, 4230:1–23, 2006.
- [7] M. Calder, S. Gilmore, J. Hillston, and V. Vyshemirsky. Formal methods for biochemical signalling pathways. In *Formal Methods: State of the Art and New Directions*. Springer, 2008.
- [8] G Despotou and T Kelly. Design and development of dependability case architecture during system development. In *25th International System Safety Conference*. System Safety Society, 2007.
- [9] S. Efroni, D. Harel, and I. R. Cohen. Reactive animation: realistic modeling of complex dynamic systems. *IEEE Computer*, 38(1):38–47, 2005.
- [10] S. Efroni, D. Harel, and I. R. Cohen. Emergent dynamics of thymocyte development and lineage determination. *PLoS Computational Biology*, 3(1):0127–0135, 2007.
- [11] J. M. Epstein. Agent-based computational models and generative social science. *Complexity*, 4(5):41–60, 1999.

- [12] M. Georgeff, B. Pell, M. Pollack, M. Tambe, and M. Wooldridge. The belief-desire-intention model of agency. In *ATAL'98*, volume 1555 of *LNCS*, pages 1–10. Springer, 2000.
- [13] J. Girard and T. Springer. High endothelial venules (HEVs): specialized endothelium for lymphocyte migration. *Immunology Today*, 15:449–457, 1995.
- [14] D. Harel, Y. Setty, S. Efroni, N. Swerdlin, and I. R. Cohen. Concurrency in biological modeling: Behavior, execution and visualization. *FBTC 2007: Electronic Notes in Theoretical Computer Science*, 194(3):119–131, 2008.
- [15] P. G. Herman, I. Yamamoto, and H. Z. Mellins. Blood microcirculation in the lymph node during the primary immune response. *The Journal of Experimental Medicine*, 136:697–713, 1972.
- [16] C. A. Janeway, P. Travers, M. Walport, and M. J. Shlomchik. *Immunobiology: The Immune System in Health and Disease (6th Edition)*. Garland Science Publishing, 2005.
- [17] T. P. Kelly. *Arguing safety – a systematic approach to managing safety cases*. PhD thesis, Department of Computer Science, University of York, 1999. YCST 99/05.
- [18] K. Ley, C. Laudanna, M. I. Cybulsky, and S. Nourshargh. Getting to the site of inflammation: the leukocyte adhesion cascade updated. *Nature Reviews Immunology*, 7(9):678–689, 2007.
- [19] G. F. Miller. Artificial life as theoretical biology: How to do real science with computer simulation. Technical Report Cognitive Science Research Paper 378, School of Cognitive and Computing Sciences, University of Sussex, 1995.
- [20] M. Miyasaka and T. Tanaka. Lymphocyte trafficking across high endothelial venules: dogmas and enigmas. *Nature Reviews Immunology*, 4(5):360–370, 2004.
- [21] L Padgham and M Winikoff. Prometheus: A methodology for developing intelligent agents. In *AOSE III*, volume 2585 of *LNCS*, pages 174–185. Springer, 2003.
- [22] E. Di Paolo, J. Noble, and S. Bullock. Simulation models as opaque thought experiments. In *Artificial Life VII*, pages 497–506. MIT Press, 2000.
- [23] F. Polack, S. Stepney, H. Turner, P. Welch, and F. Barnes. An architecture for modelling emergence in CA-like systems. In *ECAL*, volume 3630 of *LNAI*, pages 433–442. Springer, 2005.
- [24] F. A. C. Polack, T. Hoverd, A. T. Sampson, S. Stepney, and J. Timmis. Complex systems models: Engineering simulations. In *ALife XI*. MIT press, 2008. to appear.
- [25] D. J. Pumfrey. *The Principled Design of Computer System Safety Analyses*. PhD thesis, Department of Computer Science, University of York, 2000.
- [26] H. K. Rucker, H. J. Wynder, and W. E. Thomas. Cellular mechanisms of CNS pericytes. *Brain Research Bulletin*, 51(5):363–369, 2000.

- [27] A. Sadot, J. Fisher, D. Barak, Y. Admanit, M. J. Stern, E. J. A. Hubbard, and D. Harel. Towards verified biological models. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2007.
- [28] R. G. Sargent. An exposition on verification and validation of simulation models. In *17th Winter Simulation Conference*, pages 15–22. ACM, 1985.
- [29] R. G. Sargent. The use of graphical models in model validation. In *18th Winter Simulation Conference*, pages 237–241. ACM, 1986.
- [30] R. G. Sargent. Verification and validation of simulation models. In *37th Winter Simulation Conference*, pages 130–143. ACM, 2005.
- [31] T. Srivatanakul. *Security Analysis with Deviational Techniques*. PhD thesis, Department of Computer Science, University of York, UK, 2005. <http://www.cs.york.ac.uk/ftpdir/reports/YCST-2005-12.pdf>.
- [32] S. Stepney, F. Polack, and H. Turner. Engineering emergence. In *ICECCS'06*, pages 89–97. IEEE Computer Society, 2006.
- [33] J. Sudeikat, L. Braubach, A. Pokahr, and W. Lamersdorf. Evaluation of agent-oriented software methodologies – examination of the gap between modeling and platform. In *AOSE 2004*, volume 3382 of *LNCS*, pages 126–141. Springer, 2004.
- [34] H. Turner, S. Stepney, and F. Polack. Rule migration: Exploring a design framework for emergence. *Int. J. Unconventional Computing*, 3(1):49–66, 2007.
- [35] R. A. Weaver. *The Safety of Software – Constructing and Assuring Arguments*. PhD thesis, Department of Computer Science, University of York, 2003. YCST-2004-01.
- [36] M. Wheeler, S. Bullock, E. Di Paolo, J. Noble, M. Bedau, P. Husbands, S. Kirby, and A. Seth. The view from elsewhere: Perspectives on alife modelling. *Artificial Life*, 8(1):87–100, 2002.
- [37] S. Wilson, J. McDermid, P. Fenelon, and P. Kirkham. No more spineless safety cases: A structured method and comprehensive tool support for the production of safety cases. In *2nd International Conference on Control and Instrumentation in Nuclear Installations (INEC'95)*, 1995.
- [38] W. Wu and T. Kelly. Towards evidence-based architectural design for safety-critical software applications. In *Architecting Dependable Systems*, volume 4615 of *LNCS*. Springer, 2007.
- [39] B. P. Zeigler. A theory-based conceptual terminology for m&s vv&a. Technical Report 99S-SIW-064, Arizona Center for Integrative Modeling and Simulation, 1999. <http://www.acims.arizona.edu/PUBLICATIONS/publications.shtml>.