Extending an established simulation: exploration of the possible effects using a case study in Experimental Autoimmune Encephalomyelitis

Richard B. Greaves[1], Mark Read[2], Jon Timmis[1,2], Paul S. Andrews[1] and Vipin Kumar[3]

[1] Department of Computer Science, Deramore Lane, University of York, UK.
{rgreaves, jtimmis, psa}@cs.york.ac.uk

[2] Department of Electronics, University of York, Heslington, York, UK.
mnr101@ohm.york.ac.uk

[3] Laboratory of Autoimmunity, Torrey Pines Institute for Molecular Studies, La Jolla, USA.

Abstract

Investigation of a biological domain through simulation can naturally lead to the desire to extend the simulation as new areas of the domain are explored. Such extension may entail the incorporation of additional cell types, molecules or entire molecular pathways. The addition of these extensions can have a profound influence on simulation behaviour, and where the biological domain is not well characterised, a structured development methodology must be employed to ensure that the extended simulation is well aligned with its predecessor. The paper presents such a methodology, relying on iterated development and sensitivity analysis, by extending an existing simulation of Experimental Autoimmune Encephalomyelitis (EAE), a disease model for Multiple Sclerosis, via inclusion of an additional regulatory pathway. We reflect on the implications of extensions which alter simulation behaviour on pre-extension results.

Keywords

Regulatory pathway; Experimental Autoimmune Encephalomyelitis; agent-based simulation; *in silico* experimentation

1. The Proposed Methodology

Computational simulation is a useful complement to wet-lab techniques, providing a means to conduct, *in silico*, experiments that might be impossible to conduct using current laboratory technology [15]. However, simulation can only ever be based upon some simplification or abstraction of the system of interest, as it is this abstraction that makes the simulation computationally tractable. Moreover, the quality of the system model employed in simulation will be dependent upon the state of knowledge of the real-world domain. Therefore it is a natural expectation that as domain knowledge increases, the working model will need to be enhanced, be that via the addition of new system components or via new modes of interaction between them. Here we discuss

the possible implications of extending existing simulation and present a methodology for the implementation of such extensions.

In implementing extensions to a simulation, an iterative development procedure which mirrors the CoSMoS methodology [1] is employed. Development commences with the design of a very simple model as there may be little in the literature to guide the development of a more sophisticated model. This developmental phase is guided by the domain expert and seeks to identify the key actors and the relevant modifications to their current behavior in the model. Additions to the basic simulation will necessitate the reparameterisation of the simulation as the previous calibration can no longer be assumed to be valid. Therefore, once the extensions to the simulation are implemented, it is then important to explore the effects of any new simulation parameters introduced to the model. One way of doing this is to systematically map simulation behavior over the full range of values of any parameters introduced to the model, referred to here as **factorial analysis** [17]. However, this will only be tractable if very few new parameters are introduced. If such an initial exploration of parameterization is not feasible, then a fuller exploration of simulation parameterization will be called for, for example using a global **sensitivity analysis** employing an efficient sampling technique, for example Latin Hyper Cube sampling [11]. This analysis will be necessary in any case following assessment of observed effects with the help of the domain expert and will serve to re-balance the effects of the extension parameters with those already included in the model.

Sensitivity analysis is a form of statistical analysis that allows the attribution of variation in system outputs to variation in inputs. Such analysis employs an efficient sampling of parameter space as a systematic mapping of parameter space would be computationally intractable. It is good practice to understand how all parameters introduced by extensions to the model change it. This information is important for reparameterisation of the simulation, which if performed properly, will entail changes to all simulation parameters, including those that are well tied to values from the literature. Abstraction entails that simulation parameters do not represent exactly the same thing as the corresponding *in vivo* values as they must compensate for other pathways and components that are omitted from the model.

Ultimately the simulation ought to be fully recalibrated to establish baseline behaviour. This need to iteratively recalibrate upon extension of a model suggests two important implications for the use of simulation in biology: the need for caution in using the quantitative results of simulation and the need for strong methodology to guide simulation development.

2. The Proposed Case Study

The specific domain of interest for this work is Experimental Autoimmune Encephalomyelitis (EAE), a mouse model of Multiple Sclerosis. This is presented as a case study of on-going research using an Agent-based EAE simulator, ARTIMMUS [13], [14] which was developed following principled methodology and has been calibrated against *in vivo* data and which we now wish to extend to incorporate more domain detail. ARTIMMUS was developed using the CoSMoS process [1] which provides a principled framework for the development of complex system models and simulation,

promoting trust in the simulation and the results emerging from it. In accordance with the CoSMoS process explicit domain modeling was carried out, wherein specific models were created to answer specific questions in a specific domain and evaluations were guided at all times by the domain expert. The simulation has undergone an iterative calibration whereby simulation development is guided by continual improvements in the alignment of simulation predictions and *in vivo* behavior. Hence, the abstractions and assumptions made in simulation are informed by the domain expert and by empirical evidence of their appropriate capture of the domain [13], [14].

Subsequent to its development, ARTIMMUS has also been rigorously tested for simulation sensitivity to perturbations in parameter values [15]. Global **sensitivity analysis** permits the modeller to systematically examine the relative influence of each simulation parameter on simulation behavior. A robustness analysis has been employed which reveals the extent to which simulation predictions are genuinely representative of the domain rather than of underspecified parameter values arising from inconclusive domain knowledge.

An additional regulatory pathway that is believed to be influential in regulating autoimmune disease states has been identified [3]. It is known from the domain expert which cells are involved in this regulatory mechanism, but the pathway is not yet well characterised in the domain. It is therefore the intention to explore the nature of this pathway through simulation, starting by implementing a simplified model and executing exploratory simulations. The synergy of *in silico* and *in vivo* exploration and the resultant discussions with the domain expert should allow for a better characterization of the pathway via generation of experimentally testable hypotheses concerning the nature of the pathway.

The model development strategy utilized here is iterative in nature, the work presented representing an initial exploration of the pathway *in silico*. A very simplified, initial simulation has been created in which further complexity can be incrementally implemented as required since there exists very little practical guidance in the literature as to how the final simulation should behave. Since the baseline behaviour of the simulation has been characterized [14], it will be necessary to modify the parameterisation and / or the implementation of the pathway to reproduce the previously validated baseline behaviour.

The remainder of the paper addresses the disease model employed, how the additions to the simulation were implemented and the results attained. To this end, preliminary detail of the EAE model employed and the extensions to be made to it are provided in Section 3. Section 4 discusses the model of the added regulatory pathway that was implemented and Section 5 presents the results obtained from simulation based on the augmented model. Finally, Section 6 discusses the results and draws conclusions.

3. Experimental Autoimmune Encephalomyelitis

Experimental Autoimmune Encephalomyelitis (EAE) is an autoimmune disease that serves as a model for Multiple Sclerosis [10]. Multiple Sclerosis is characterised by damage to the myelin coating of nerve fibres resulting in impaired conduction of impulses along them, leading ultimately to paralysis and death. In EAE similar damage

to the myelin sheath is mediated by CD4 Th1 cells that are reactive towards various components of myelin, for example Myelin Basic Protein (MBP) [19].

This case study is focussed on the murine model of EAE [7, 18]. This model addresses the mechanisms of spontaneous recovery from EAE which is highlighted as grey dashed arrows in Figure 1.

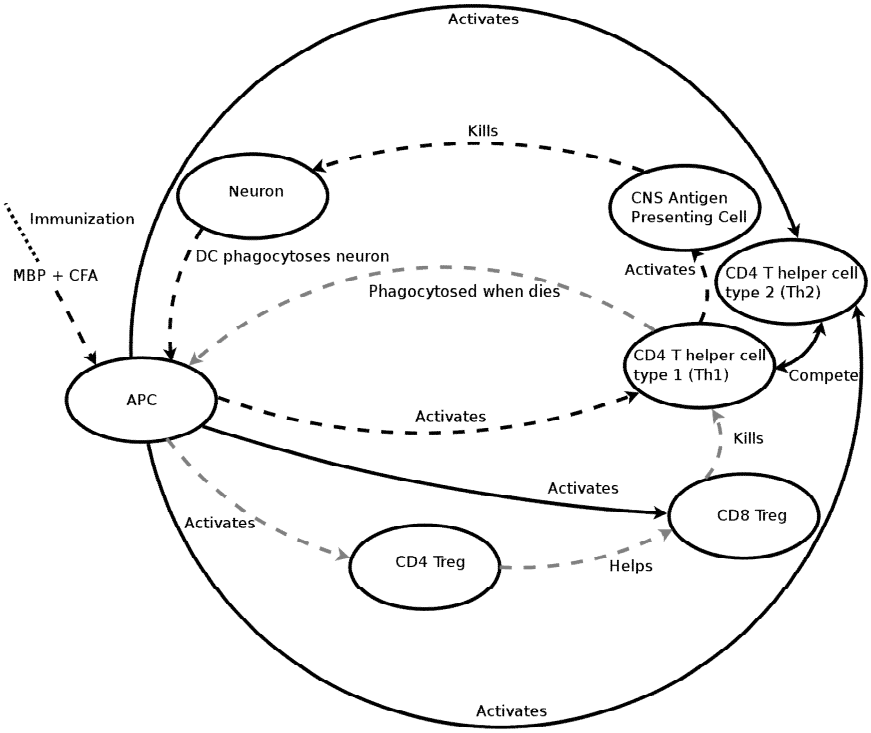

Fig. 1. An informal representation of the cells implicated in the EAE disease and spontaneous recovery cycles and their inter-relationships. The disease cycle is highlighted in black dashed lines, the recovery cycle in grey dashed lines. Dendritic cells (DCs) and macrophages are represented in the cycle as generic Antigen Presenting Cells (APCs). Figure adapted from [16].

3.1 Induction of Disease

The disease cycle is informally represented by the circuit of black dashed arrows in Figure 1. The disease can be induced in laboratory animals via inoculation with MBP, Complete Freund's Adjuvant (CFA) and Pertussis toxin [12], which provokes an immune response and stimulates DCs sufficiently that they can express costimulatory molecules. Dendritic cells (DCs) respond to MBP by phagocytosing (internalizing) and presenting it as complexes with MHC molecules. The DC then migrate to the

secondary lymphoid organs where they encounter naïve MBP-recognising CD4 T-helper cells and bind and activate them, allowing the auto-reactive CD4 Th1 to proliferate and mature into effectors.

The activated T-cells can then migrate, along with macrophages, through the blood-brain barrier into the CNS. Once within the CNS, the activated encephalitogenic T-cells can set up inflammation. The environment thus created is capable of activating microglia and macrophages, stimulating them to secrete Tumour Necrosis Factor α (TNF-α). TNF-α causes neuronal death releasing MBP which is subsequently phagocytosed and digested by macrophages leading to the presentation of MBP antigens to further naïve auto-reactive CD4 Th1 which have arisen due to CD4 Th proliferation in the CNS. Recognition of MBP presented by DCs in the cervical lymph node coupled with the inflammatory milieu fully activates further CD4 Th which can subsequently enter the CNS and perpetuate the disease cycle.

3.2 Spontaneous Recovery from Disease

Auto-reactive CD4 Th1 in the CNS will eventually reach the end of their lifespan and undergo programmed cell death (apoptosis). Apoptotic CD4 Th1 can leave the CNS and be phagocytosed by dendritic cells in the lymph nodes [12]. Phagocytosis of encephalitogenic CD4 Th1 leads to presentation of antigens derived from their T-cell receptors by DCs. Two particular TCR fragments are important in our model of EAE regulation: the Framework region 3 (Fr3) and Complementarity Determining Regions 1 and 2 (CDR1/2) [7, 18]. When presented on appropriate MHC molecules these antigens serve to activate the regulatory T-cell populations (Treg).

In the CNS draining lymph nodes MHC-II-Fr3 complex presented by DCs is recognised and bound by CD4 Treg which become activated and in turn, owing to the inflammatory cytokines present which stimulate the DC to produce costimulatory molecules, stimulate or 'license' the DC to express the MHC molecule Qa-1 [8]. CDR1/2 derived from auto-reactive CD4 Th1 TCRs can then be presented by DCs complexed with Qa-1. CD8 Tregs are then able to recognise and bind the Qa-1-CDR1/2 complex and become activated by the DC with help from IFN-γ secreted by CD4 Treg [18]. Once activated, the CD8 Treg leave the DC and are capable of recognising and binding Qa-1-CDR1/2 presented transiently on the surface of auto-reactive CD4 Th1 [7]. CD8 Treg kill these by inducing apoptosis [2]. This cell-mediated killing serves to regulate the auto-immune response and the population of self-reactive CD4 Th1 is reduced, the spontaneous recovery from EAE and the return of cellular populations to their resting levels being essentially complete within 50 days [12]. This behaviour is reflected in the ARTIMMUS results model presented in Figure 2.

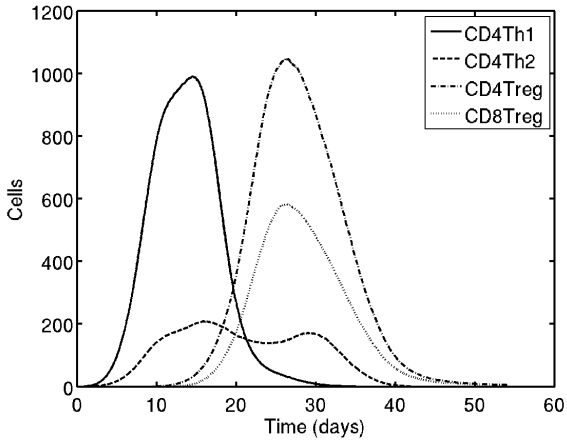

Fig. 2. The populations of T-cell effectors with time. The curves are derived from the median population levels at a given time across a set of 1000 simulator runs. The resulting curves clearly show the anticipated behaviour of the EAE system: cell populations begin at their resting levels at day 0 when immunisation occurs. The populations of the CD4 Th cells then begin to rise, peaking at around 10-15 days, with the CD4 Th1 population peaking considerably higher than CD4 Th2 (circa 1,000 cells against circa 200). By this time, the populations of Tregs will have begun to increase and these will peak around 27-30 days with CD4 Treg attaining a higher population than CD8 Treg (circa 1,000 cells against circa 600). Recovery from EAE will now be effectively complete and cell populations will fall back to their resting levels. The plot was generated by re-running the baseline experiment as described in [16].

3.3 The CD200-CD200R Immune-regulatory Pathway

The pathway added to the EAE model is thought to reduce auto-immune response by down-regulating the ability of DCs to stimulate T-cell populations. DCs are pivotal in the activation of all T-cell sub-populations and so any reduction of their capacity to bring this about could significantly scale back the size of an immune response by curtailing T-helper and Treg populations.

CD200 and CD200R are cell surface proteins, CD200 being a ligand to the CD200R receptor [5]. CD200 is expressed on a variety of cells particularly T-cells from the immune system and also on neurons [9]. The CD200 protein has also been shown to be expressed constitutively on the CD8 Treg population [4]. CD200R is similarly widely expressed with microglia and DC being able to express this receptor [9]. In the initial augmented model, only T-cell-DC interactions outside the CNS are considered in line with the iterative development methodology. When CD200 binds to CD200R, a signal is transmitted via the receptor to the DC or microglia, causing it to down-regulate production of certain proteins that are essential to the binding and activation of T-cells [5]. There is now considerable experimental evidence for the regulatory effect of CD200 on cells such as DC and macrophages [9].

4 Implementation of the Simulation

Following the iterative development methodology outlined in Section 1, a simple model of the CD200-CD200R regulatory pathway was implemented. The model allowed for the expression of CD200 by effector CD8 Treg, expression being immediate upon the Treg becoming an effector and with no requirement for local activation. DCs express CD200R immediately upon maturity. A down-regulatory signal to DCs from CD200-CD200R interaction was assumed to occur whenever CD8 Treg and DCs come into contact. The down-regulatory or 'negative' signal then probabilistically causes the DC to down-regulate expression of Qa-1 and MHC (collectively referred to as MHC) and / or of costimulatory molecules (referred to as CoStim).

The model implemented [6] was very simplistic owing to the lack of detailed explicit knowledge concerning the CD200-CD200R pathway. To keep the model as simple as possible CD200 or CD200R expression was not permitted within the CNS i.e on neurons or microglia and the ARTIMMUS default exclusion of CD8 Treg from the CNS [14] was maintained.

The stipulation that negative signalling probabilistically down-regulates MHC and / or CoStim expression by DCs introduces two new parameters into our model, namely the probability that a negative signal will down-regulate MHC expression and also the probability that negative signalling will down-regulate CoStim expression by DCs. We therefore seek to ascertain valid values for these new parameters, such that the added pathway will not unduly perturb simulation from its baseline behaviour. This is done with a view to re-balancing the effects of the CD200-CD200R regulatory pathway with those of the regulatory pathway already implemented in the simulation.

Since the pathway is poorly characterised it is not possible to obtain reasonable values for the probabilities from the literature. The effect of the existing simulation parameters will also be slightly perturbed by the additions to the model. These considerations therefore make it necessary to conduct a full analysis of the effects of the parameters introduced by extending the model. A factorial analysis i.e a systematic mapping was conducted of simulation behaviour across the full range of values for both of the added parameters.

5. Evaluation of the Initial Model of the CD200 Pathway

The **factorial analysis** mapped simulation behaviour for values of the probability parameters between 0% and 100% in steps of 10%, meaning a mapping of 121 separate simulations with unique pairings of probability values. Simulation behavior was described in terms of median values across 1000 simulator runs for various system properties or responses. The responses analysed were peak T-cell effector populations (for CD4 Th1, CD4 Th2, CD4 Treg and CD8 Treg), the times taken to reach these maximal populations, the CD4 Th1 population remaining at day 40 of the simulation, the mean EAE severity score and the EAE severity score at 40 days. EAE severity scores calculated from effector cell populations are intended to correspond to the clinical severity scores assigned in the laboratory on the basis of the observed extent of disease symptoms [14].

Typically, even when there was only a 10% chance that MHC expression could be reduced via negative signalling; the effect on T-cell effector populations was signifi-

cant. CD8 Treg and CD4 Treg peak populations were significantly reduced compared to the experiment with both parameters set to 0% i.e. the baseline (data presented in Figures 3 and 4 respectively). The peak population of CD4 Th1 was increased, though not significantly, presumably owing to the reduced activation by DC being balanced by the reduction in the occurrence of apoptosis by CD8 Treg.

Repetition of the factorial analysis with values for the probability parameters between 0% and 1% in steps of 0.1% still showed significant reductions in CD4 Treg and CD8 Treg populations compared to baseline values even at values of the MHC down-regulation probability as low as 0.2%.

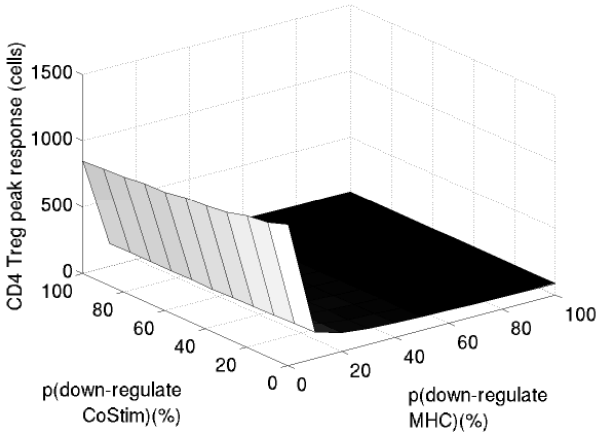

Fig. 3. Three-dimensional plot of peak CD8 Treg effector populations in a systematic mapping of simulation behaviour at values of the two probability parameters between 0% and 100% in steps of 10%. The probability parameters represent the probability that a negative signal will down-regulate MHC (i.e. MHC-II and Qa-1) expression by a DC – represented as p(down-regulate MHC) on the axes and the probability that a negative signal will down-regulate CoStim expression by DC – represented as p(down-regulate CoStim) on the axes. 121 parameter value pairings were used in generating the landscape, the baseline simulation is represented as the pairing 0%, 0% in the bottom middle of the plot.

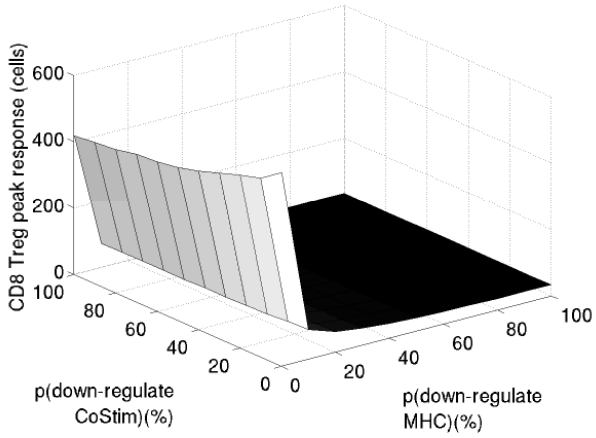

Fig. 4. Three-dimensional plot of peak CD4 Treg effector populations in a systematic mapping of simulation behaviour at values of the two probability parameters between 0% and 100% in steps of 10%. The probability parameters represent the probability that a negative signal will down-regulate MHC (i.e. MHC-II and Qa-1) expression by a DC – represented as p(down-regulate MHC) on the axes and the probability that a negative signal will down-regulate CoStim expression by DC – represented as p(down-regulate CoStim) on the axes. 121 parameter value pairings were used in generating the landscape, the baseline simulation is represented as the pairing 0%, 0% in the bottom middle of the plot.

The **factorial analysis** highlights a need to further develop our model of the CD200-CD200R regulatory pathway, the existing model exerting too drastic an effect on T-cell population priming and almost certainly not being a fair representation of the domain. In practice, the model leads to a severe reduction in T-cell effector populations, which is believed to be excessive. This arises from the impact that CD200 expressed by Treg has on all T-cell populations (illustrated informally in Figure 5). It is difficult to establish exactly how one should go about implementing such a sparsely characterised pathway in the model, because of the potentially far reaching consequences of reducing the capacity of DCs to prime T-cells in this particular instance. Via the expression of CD200, CD8 Treg can impact not only the priming of all other T-cell sub-populations, but also their own priming by DCs.

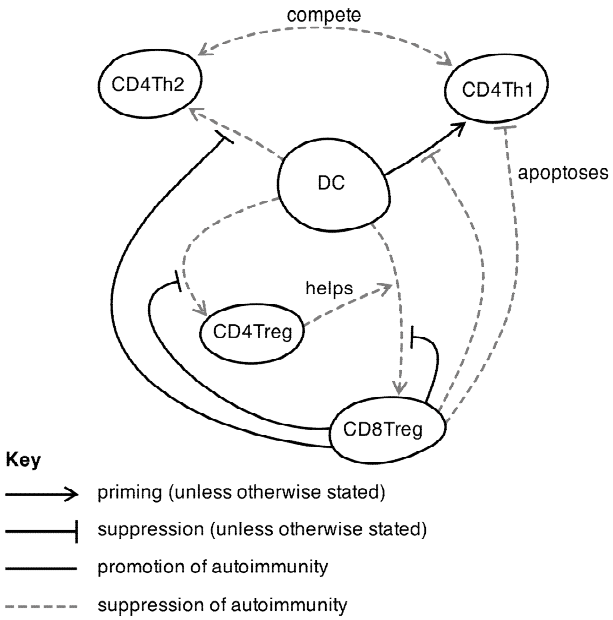

Fig. 5. An informal description of the interactions of the different T-cell populations of our model with each other and with dendritic cells (DCs). Via CD200 expression, CD8 Treg can modulate DC ability to prime all T-cell populations, including CD8 Treg and can this significantly impact T-cell effector populations. Interactions marked as solid black lines tend to promote autoimmunity, whereas those indicated as grey broken lines serve to reduce it. Figure reproduced from [14].

A number of feasible alternative models have been proposed which we shall not concern ourselves with too deeply here, however, one possible line of investigation concerns a stepped approach to the down-regulation of MHC and CoStim by DCs. In this alternative model each negative signal would cause a partial probabilistic down-regulation of MHC or CoStim rather than the total probabilistic down-regulation currently implemented. This stepped down-regulation would require several signals to completely down-regulate MHC or CoStim expression by DC, each signal compounding the effect of previous signals.

Of more pressing import are the implications for the flexibility of simulation as a research tool if extensions to well characterised (calibrated and validated) simulation result in serious perturbations of simulation behaviour as our extensions appear to do. Our current model is simplistic and the pathway of interest is incompletely understood. We believe the work has validity as it is of an exploratory nature and its strength lies in its ability to investigate hypotheses proposed by immunologists and feed back to them the results of simulating this system. In this way, we aim to help shape immunological thinking and thus guide wet-lab experimentation.

6. Discussion and Conclusions

The paper presents a methodology which is applicable to the extension of established models in order to explore new pathways and components. In this instance, EAE has been employed as a case study, entailing the extension of the ARTIMMUS simulation via addition of a model of the CD200 pathway which considers only the DC-CD8 Treg interactions. The resulting simulation behavior leads us to consider the wider implications of extending simulations in computational immunology.

Simulation represents a means of integrating biological data. A parameterized simulation, once properly calibrated, can serve as a tool for formulating and testing hypotheses relating to the domain. In the natural course of research it is possible that new pathways or components will be identified as being influential in the system, and simulation provides a means whereby a preliminary exploration of these may be conducted. However, as often little is known of the domain, a structured and principled approach to this exploration is required. An appropriate methodology for such exploration has been presented here.

The methodology has been applied to ARTIMMUS, which has been rigorously developed and calibrated against the real-world system [15]. We are therefore confident that ARTIMMUS is representative of the domain. The CD200 regulatory pathway has been identified by the domain expert as having an influential role in the regulation of dendritic cell ability to activate T-cell populations during autoimmunity. In order to conduct preliminary simulation of the pathway, it was necessary to incorporate an abstraction of the pathway into our current simulation. However, the domain is not well characterised and therefore a rigorous approach is required for implementing models of the pathway. An initial, simple model was developed in line with the described methodology. **Factorial analysis** reveals the mechanism to be too simplistic and so further models, such has that described briefly in Section 5, will be explored. However, the severe impact of the implemented pathway on simulation T-cell populations leads us to consider certain important philosophical issues concerning the use of simulation in exploring immunology. Extensions to simulation can rebalance the influence of cells and pathways within a simulation and as such quantitative measures of their influence may change, and this can have implications for previous quantitative results.

Firstly, we must be cautious in extracting quantitative results from simulation and qualitative results will be, at best, acceptable in terms of the current domain knowledge. Secondly, although the advantages that simulation has to offer as a complement to wet-laboratory experimentation are formidable, it is, however, clear that strong methodologies are required to guide simulation development to appropriate levels of abstraction. These methodologies will also serve to ensure that simulation adequately captures the domain, that modellers continually evaluate simulation performance against real domain observations and that simulation is capable of indicating when something influential, and possibly as yet unidentified *in vivo*, is missing from the simulation. A failure of simulation to fully represent real-world dynamics should lead to further *in vivo* exploration or development of the simulation.

Simulation has the potential to offer a powerful complement to *in vivo* study. However, for that potential to be realized, the simulation community needs to employ

methodologies that instill confidence that simulation results are accurate and can clearly motivate further development (of the simulation) when this is not the case.

Acknowledgement

We wish to acknowledge Dr Carl Ritson for allowing us access to the CoSMoS computing cluster at the University of Kent in order to carry out the simulation reported here.

References

[1] Andrews, P.S., Polack, F.A.C., *et al.*: CoSMoS Process: A Process for the Modelling and Simulation of Complex Systems, University of York, Technical Report YCS-2010-453 (2010)

[2] Beeston, T., Smith, T.R.F., *et al.*.: Involvement of IFN-g and perforin, but not Fas/FasL interactions in regulatory T cell-mediated suppression of experimental autoimmune encephalomyelitis, Journal of Neuroimmunology, vol. 229, pp. 91-97 (2010)

[3] Feuer, R.: Tickling the CD200 Receptor. A Remedy for those Irritating Macrophages, The American Journal of Pathology, vol. 171, pp. 396-398, (2007)

[4] Fanchiang, S.S., Cojocaru, R., *et al.*: Global expression profiling of peripheral Qa-1 restricted CD8aa+TCRab+ regulatory T cells reveals innate-like features: implications for immune-regulatory repertoire, Human Immunology, In Press (2011)

[5] Gorczynski, R.M., Chen, Z.,Kai, Y., Lee, L., Wong, S., Marsden, P.A.: CD200 is a Ligand for All Members of the CD200R Family of Immunoregulatory Molecules, Journal of Immunology, vol. 172, pp. 7744-7749 (2004)

[6] Greaves, R.: Computational Modelling of Treg Networks in Experimental Autoimmune Encephalomyelitis, MSc by Research Thesis, Department of Computer Science, University of York (2011)

[7] Kumar, V, Sercarz, E.: An Integrative Model of Regulation Centred on Recognition of TCR / MHC Complexes, Immunological Reviews, vol. 182, pp. 113-121 (2001)

[8] Kumar, V., Stellrecht, K., *et al.*: Inactivation of T Cell Receptor Peptide-specific CD4 Regulatory T Cells Induces Chronic Experimental Autoimmune Encephalomyelitis (EAE), Journal of Experimental Medicine, vol. 184, pp. 1609-1617 (1996)

[9] Liu, Y., Bando, Y., *et al.*: CD200R1 Agonist Attenuates Mechanisms of Chronic Disease in a Murine Model of Multiple Sclerosis, Journal of Neuroscience, vol. 30, pp. 2025-2038 (2010)

[10] Lublin, F.D.: Relapsing Experimental Allergic Encephalomyelitis an Autoimmune Model of Multiple Sclerosis, Springer Seminars in Immunopathology, vol. 8, pp. 197-208 (1985)

[11] McKay, M.D., Beckmann, R.J., et al.: A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. Technometrics, vol. 21, pp. 239-245 (1979)

[12] Pender, M.P., Experimental Autoimmune Encephalomyelitis. In Autoimmune Neurological Disease, Pender, M.P. and McCombe, P.A. (Eds.), Cambridge: Cambridge University Press, pp. 26-88 (1995)

[13] Read, M, Andrews et al., Nature Biotechnology (in preparation) (2011)

[14] Read, M.: Statistical and Modelling Techniques to Investigate Immunology Through Agent-based Simulation, PhD Thesis, Department of Computer Science, University of York (2011)

[15] Read, M., Andrews, P., *et al.*: Techniques for Grounding Agent-based Simulations in the Real Domain: a Case-study in Experimental Autoimmune Encephalomyelitis, Mathematical and Computer Modelling of Dynamical Systems, DOI:10.1080/13873954.2011.601419 (2011).

[16] Read, M., Timmis, J, *et al.*: A domain model of experimental autoimmune encephalomyeltis. In: Stepney, S., Welch, P.H., Andrews P.S., Timmis, J. (Eds.) CoSMoS, pp 9-44, Luniver Press (2009)

[17] Saltelli,A., Chan, K., *et al.*: Sensitivity Analysis. Wiley series in probability and statistics. Wiley (2000)

[18] Tang, X., Smith, T.R.F., *et al.*: Specific Control of Immunity by Regulatory CD8 T Cells, Cellular and Molecular Immunology, vol. 2, pp. 11-19 (2005)

[19] Van den Bark, A.A., Gill, T, *et al.*: A Myelin Basic Protein-specific T Lymphocyte Line that Mediates Experimental Autoimmune Encephalomyelitis, Journal of Immunology, vol. 135, pp. 223-228 (1985)