

This is a repository copy of *Automated multi-objective calibration of biological agent-based simulations*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/104877/>

Version: Accepted Version

Article:

Read, Mark N, Alden, Kieran, Rose, Louis M orcid.org/0000-0002-3419-2579 et al. (1 more author) (2016) Automated multi-objective calibration of biological agent-based simulations. Interface. ISSN: 1742-5662

<https://doi.org/10.1098/rsif.2016.0543>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Automated Multi-Objective Calibration of Biological Agent-Based Simulations

Mark N. Read^{1,2,*}, Kieran Alden³, Louis M. Rose⁴, Jon Timmis³

¹School of Life and Environmental Sciences, The University of Sydney, Australia.

²Charles Perkins Centre, The University of Sydney, Australia.

³Department of Electronics, The University of York, UK.

⁴Department of Computer Science, The University of York, UK.

*Corresponding author: mark.read@sydney.edu.au

Keywords: Computational Biology — Agent-Based Simulation — Calibration — Multi-Objective Optimization — Experimental Autoimmune Encephalomyelitis

Abstract

Computational agent-based simulation is increasingly used to complement laboratory techniques in advancing our understanding of biological systems. Calibration, the identification of parameter values that align simulation with biological behaviours, becomes challenging as increasingly complex biological domains are simulated. Complex domains cannot be characterised by single metrics alone, rendering simulation calibration a fundamentally multi-metric optimisation problem that typical calibration techniques cannot handle. Yet calibration is an essential activity in simulation-based science; the baseline calibration forms a control for subsequent experimentation, and hence is fundamental in the interpretation of results. Here we develop and showcase a method, built around multi-objective optimisation, for calibrating agent-based simulations against complex target behaviours requiring several metrics (termed *objectives*) to characterise. Multi-objective calibration delivers those sets of parameter values representing optimal tradeoffs in

simulation performance against each metric, in the form of a Pareto front. We use MOC to calibrate a well-understood immunological simulation against both established *a priori* and previously unestablished target behaviours. Further, we show that simulation-borne conclusions are broadly, but not entirely, robust to adopting baseline parameter values from different extremes of the Pareto front, highlighting the importance of MOC’s identification of numerous calibration solutions. We devise a method for detecting overfitting in a multi-objective context, not previously possible, used to save computational effort by terminating MOC when no improved solutions will be found. MOC can significantly impact biological simulation, adding rigour to and speeding up an otherwise time-consuming calibration process, and highlighting inappropriate biological capture by simulations that cannot be well calibrated. As such, it produces more accurate simulations that generate more informative biological predictions.

1 Introduction

Computational modelling and simulation has emerged as a tool for investigating a wide range of biological systems, spanning immunology [1][2], drug and intervention design [3][4], developmental biology [5], and ecology [6]. Biological simulation is particularly insightful when used in complement with traditional methods, such as wet-lab *in vivo* and *in vitro* work; laboratory work generates experimental data and suggests hypotheses that can be evaluated by way of their integration with simulation, which in turn can suggest further experiments or highlight areas of lacking knowledge [7][8]. Well designed, biologically-accurate simulations provide detailed spatio-temporal insight, facilitating observations and assays not possible in the real system; simulation experiments are unhampered by the ethical, practical and financial considerations inherent in biological experimentation. Research programs integrating wet-lab and simulation methods can offer a greater return on animal experimentation by generating additional insight, and hence easing the burden on experimental animals, in line with the ‘3Rs’ principles (Replacement, Reduction and Refinement).

The agent-based simulation (ABS) paradigm permits detailed and nuanced simulation of biological systems [9][3]. Simulation components are represented as explicit individual entities, *agents*, with unique states that exist within a spatial environment. Rules

specifying agent dynamics and the consequences of interaction are provided, and simulation execution allows the system-level consequences of agent-level manipulations to be observed. ABS incorporates stochastic events, and therein reflects the heterogeneity of real world natural systems. There is scope for specifying very detailed interactions using ABS, at the expense of generating large numbers of parameters: 50+ is not uncommon.

Drawing biologically meaningful conclusions from simulation requires that the mapping of the simulation to the biology is known. This can prove problematic for two reasons. First, simulations are abstract representations of their corresponding real world systems. For example, there exist at least 19 varieties of T cell, a vital component of the immune system [10]. However, rather than fully capture all their nuanced differences, a simulation is more likely to represent an abstracted subset thereof. As such, experimental measurements on a real world T cell cannot be assumed to translate directly to its simulation counterpart. Second, complex biological systems are the subject of simulation precisely because they are incompletely understood, meaning that the real world data supporting simulation design decisions and corresponding parameter values may not exist. Calibration is a critical activity in establishing the link between simulation and biology; parameter values that align simulation and real-world dynamics are identified. Furthermore, an inability to provide a good alignment points to simulation design that does not appropriately capture the biology. Calibration is used to establish a baseline simulation dynamic used as a control in subsequent experimentation, and finding appropriate values is important. Different parameter values will yield different simulation dynamics, and as such influence the conclusions drawn from experiments.

A number of approaches to calibration exist, including manual calibration [11], evolutionary algorithms [12][13], maximum likelihood estimation and various forms of regression [14]. These techniques identify parameter values by employing a single metric to align simulation dynamics with those of the real world system. However, complex biological system dynamics are not well characterised by single metrics alone. They constitute many different types of interacting component, and encompass both positive and negative feedbacks. They are highly redundant: a single component can perform many functions and any one function can be performed by several components [15][16]. As such, calibration of a complex system simulation is fundamentally a multi-metric optimisation problem; several metrics of a simulation's alignment with the biology must be simulta-

neously considered when evaluating putative parameter values. Consider, for example, cellular motility, which underlies many biological processes arising from cellular interaction. Which targets a given cell interacts with depends on both its speed and directional persistence; accurately modelling this process requires that metrics of both be considered.

In this paper, we position multi-objective optimisation-based calibration (MOC: multi-objective calibration) as an important enabling technology for simulation-based biological investigation. Given its abstractive nature, a simulation undergoing calibration will not perfectly replicate all aspects of the biology. As such, putative simulation parameter value sets will exhibit tradeoffs in their reproduction of aspects of the biology, excelling in some at the expense of others. In this context, a metric quantifying a simulation’s capture of a specific aspect of the biology is termed an *objective*. Through the use of Pareto fronts (defined in Section 3), MOC explicitly tracks the collection of simulation parameter sets exhibiting optimal tradeoffs between objectives. It is unknown if adopting baseline parameter values from different regions of the Pareto front will deliver fundamentally different conclusions from simulation-based experiments. The answer to this question is likely problem-specific, and the use of MOC allows this issue to be addressed by exposing a full range of Pareto-equivalent solutions.

Here we investigate multi-objective optimisation, specifically the NSGA-II algorithm [17], in calibrating an established immunological simulation: ARTIMMUS [18]. ARTIMMUS simulates Experimental Autoimmune Encephalomyelitis (EAE), a mouse model of multiple sclerosis [19][20]. It is a complex simulation, encompassing seven distinct cell populations that interact across five organs, and constituting 72 parameters. Its successful prior manual calibration renders it an effective test case for evaluating MOC’s applicability to simulation calibration. We demonstrate the successful calibration of ARTIMMUS using five objectives (Section 4): a range of solutions to the calibration problem, offering optimal tradeoffs against calibration objectives, are generated. Furthermore, we demonstrate that conclusions drawn from a simulation-based experiment can vary depending on exactly which calibration solution is adopted (Section 5). Hence, different calibration solution parameter values can vary downstream conclusions, highlighting MOC’s value in making these multiple solutions explicit. We show that MOC is equally applicable in generating simulation initial condition values: cellular population sizes as simulation launch. We proceed to demonstrate that MOC can identify parameter and initial condition values

that deliver previously unknown simulation dynamics, highlighting its potential beyond this well understood test case (Section 6). Lastly, we consider strategies for formulating stopping criteria for MOC, thereby preventing over-fitting and wasted computational expense when apparent improvements in simulation calibration are likely due to stochastic sampling rather than genuinely superior parameter values (Section 7). We begin by introducing ARTIMMUS (Section 2), and the MOC methodology (Section 3).

2 A testbed for calibrating biological simulations

ARTIMMUS is an agent-based simulation of an EAE protocol wherein mice induced into autoimmunity undergo a natural recovery from disease, and are thereafter resistant to disease re-induction [18][21][22]. ARTIMMUS was created, in part, to further probe the cellular interactions mediating this recovery [23][24]. It has been used to explore the mechanisms through which splenectomy, the removal of the spleen, a primary immune organ, exacerbates disease severity, and predict the outcome of T cell interaction-blocking drugs [18]. It was conceived through a collaboration of immunologists and computer scientists, and developed through a principled approach focusing on documenting how biological concepts are translated into computer code: the CoSMoS process [32]. It is written in the Java programming language.

ARTIMMUS has previously undergone a by-hand, manual calibration [11], and was shown to reflect the dynamics of the real world disease [18]. The process demanded close collaboration between the simulation developer and an immunologist who informed the work, helping bridge biological data and concepts to simulation constructs and output. This manual calibration took two weeks, and entailed an iterative process through which simulation code and parameter value changes that might explain perceived discrepancies between simulation and biological system dynamics were identified and explored in turn. Those best aligning simulation with biological dynamics were adopted before repeating the process. This calibration approach is akin to a non-population, manual, greedy local search wherein the best immediate improvement is always adopted.

Despite delivering a well-calibrated result for ARTIMMUS, this calibration search strategy presents several potential pitfalls. It is entirely plausible that the manual search does not find the global optimum parameter set that best aligns simulation dynamics with

those of the biological system. As a greedy search strategy, its result is highly dependent on the search's starting position, and complex landscapes where one parameter's influence on simulation dynamics critically depends on the values held by others are particularly challenging. The existence of multiple solutions to the calibration problem can go entirely undetected. Lastly, manual calibration is time consuming, and agent-based simulation's stochastic nature furthers compound these challenges. It is these issues that collectively motivated the present automated MOC approach.

Here we provide a brief summary of EAE and ARTIMMUS to aid understanding of the sections that follow; a comprehensive description may be found in the supplementary materials of [18]. Figure 1A provides an abstract overview of the major cell types involved in EAE, and their relationships to one another. EAE is induced through injection of neuronal fragments which are internalized by dendritic cells (DCs) which then direct the growth of a T cell population (CD4Th1, abbreviated to Th1) targeting these fragments. These Th1 cells enter the central nervous system (CNS), where they stimulate CNS-resident macrophages into secreting $\text{TNF-}\alpha$, which in turn damages neurons. The resultant neuronal fragments are internalized by further populations of DCs, which direct further Th1 activities, perpetuating the autoimmune cycle. Recovery from autoimmunity is through the actions of two populations of regulatory T cell, CD4Treg and CD8Treg cells, so named as they regulate the activities of other T cells. The natural life-cycle of a Th1 cell results in its eventual death and internalization by DCs, which derive fragments therefrom and direct the growth of CD4Treg and CD8Treg cells targeting the Th1 cell population. CD4Tregs play an essential role in facilitating the development of CD8Treg cells. CD8Treg cells can directly kill Th1 cells, interrupting their natural life-cycle and preventing the perpetuation of autoimmunity. Th2 cells directly compete with Th1 cells, as both arise from a common progenitor and they each perform downstream activities that promote their own development. The reduced severity of the autoimmune environment arising from the action of CD8Treg cells favours the growth of Th2 cells over Th1 cells, which do not directly harm neurons and hence do not contribute to this autoimmune process. Figure 1B shows a time-series graph of T cell population sizes in ARTIMMUS.

3 Multi-objective calibration (MOC) methodology

We present here an overview of the multi-objective calibration (MOC) concept, detailing how we employ multi-objective optimisation technology to calibrate simulation parameters and initial conditions. A graphical overview is supplied in Figure 2.

Firstly, we define the desired (*target*) ARTIMMUS dynamics, Figure 2A. In this manuscript targets are expressed as peak cell population sizes, the times at which those peaks occur, or the cell population sizes at a given time. Target dynamics might represent known biological results to be reproduced, or hypothetical outcomes of interest. In this study we adopt the dynamics of a previous manual calibration of ARTIMMUS, so as to evaluate MOC on a well-understood problem; thereafter we employ MOC to obtain hypothetical dynamics not known possible *a priori*. We note that many other aspects of simulation performance can constitute target dynamics, depending on the context and simulation being calibrated. The expression of targets as distributions reflects the stochastic nature of biological systems and agent-based simulations, wherein repeat experiments can yield slightly different results.

MOC seeks to identify parameter values that best align simulation with target dynamics. As such, we define metrics, termed *objectives*, that quantify the alignment between the two. As illustrated in Figure 2B (left), we employ the the non-parametric Kolmogorov-Smirnov statistic in our objectives, which quantifies the difference in target and simulation dynamics for a given set of simulation parameter values. Rather than contrasting the medians of two distributions, as many statistics do, the KS statistic quantifies the biggest distance between two distributions' cumulative distribution functions. As such, its use here facilitates the calibration of a distribution's shape, not simply its median or mean. We consider this a strength of our approach; as may be seen in the sections that follow, MOC is capable of reproducing distributions of behaviour, not simply averages. Each set of simulation parameter values is termed a 'candidate solution', and its corresponding simulation performance is evaluated against each objective individually. By evaluating many candidate solutions we identify regions of parameter space providing close alignment with target dynamics (Figure 2B, right). Importantly, the regions that satisfy each objective differ. In practice, it is computationally intractable to fully explore parameter space as suggested by the heatmaps in this Figure, particularly when many parameters are investigated. Instead, a heuristic (guided) search strategy is employed that samples

parameter space, evaluates performance, and decides from where to extract the next candidate solutions based on the results. In this study we employ NSGA-II as our guided search engine [17], but we believe other multi-objective optimisation technologies could be successfully substituted. NSGA-II maintains a population of candidate solutions, and employs (heavily abstracted) principles of genetic recombination, mutation and natural selection to generate and evaluate successive generations of superior candidate solutions. Hence, NSGA-II is an iterative algorithm. We refer the readers to [17] for more detail on NSGA-II. Here we have employed the ‘inspyred’ python module NSGA-II implementation.

We identify those candidate solutions that constitute optimal tradeoffs in performance against each objective, referred to simply as *solutions*, Figure 2C. The set of solutions is termed the *Pareto front*. These solutions are *Pareto-equivalent*: no solution has been found that offers an improvement in one objective without a worsening in another. Pareto-equivalent solutions may reside in disparate regions of parameter space, and the ability to recognize this is a key strength of MOC. Though these regions of parameter space may be Pareto-equivalent for the given target simulation behaviour, they could yield very different behaviours when subjected to further downstream experimentation, and as such lead to different simulation-borne conclusions. In this study we investigate this phenomenon for a given experiment in ARTIMMUS.

We note that it is possible to derive a great many targets and objectives for complex system simulations. Increasing the number of objectives increases the difficulty of the calibration problem, and the computational resource required to address it; in the field of optimisation this is known as the ‘curse of dimensionality’. Hence, employing fewer, uncorrelated objectives is considered good practice: it encourages the identification of good quality solutions whilst minimising the resources required to do so.

3.1 Selecting candidates from the Pareto front

Upon completion MOC delivers a Pareto front of Pareto-equivalent solutions, representing optimal tradeoffs between the calibration objectives. Deciding which solution adopt as the baseline simulation parameter values is an application-specific problem. For the present study we have developed a function, $\Lambda(c)$, which assesses candidate solution c against the criteria below. We select the candidate with the lowest Λ value when presenting the results of calibration below. Λ is calculated as follows.

Let Ω represent the set of calibration objectives, and $KS_o(c)$, $o \in \Omega$ as the corresponding Kolmogorov-Smirnov score for candidate c on objective o . $\overline{KS}(c)$ represents the mean objective score for candidate c . The Λ score is calculated as:

$$\Lambda(c) = \alpha \cdot \overline{KS}(c)^2 + \sum_{o \in \Omega} \left(KS_o(c) - \overline{KS}(c) \right)^2 \quad (1)$$

Low Λ scores are achieved through low mean objective KS scores, and balanced KS scores across all objectives. α specifies the relative importance of these two components. When $\alpha = 1$, both measures contribute equally to Λ . Lower mean KS scores are prioritised with $\alpha > 1$, and *vice versa*. We employ $\alpha = 1$ throughout. We note that Λ is unit-less, and as such is not explicitly reported here; it is used only to extract one candidate solution from a Pareto front, presented as the chief result of calibration in the results that follow.

4 Successful re-calibration of ARTIMMUS

We demonstrate MOC by re-calibrating ARTIMMUS, taking as target dynamics those of the previous manually-calibrated simulation dynamics [18]. As these dynamics are known to be obtainable, and at least one set of parameter values that produce them are known, we are able to evaluate MOC's performance.

With 5 objectives MOC successfully reproduced the manually-calibrated ARTIMMUS dynamics, as demonstrated in Figure 3. The objectives used were:

- the peak Th1 cell population size (Figure 3B)
- the time at which the peak occurred (Figure 3C)
- the Th2 population size at 30 days (Figure 3D)
- the peak population sizes of both CD4Treg and CD8Treg cells (Figures 3E and F).

The corresponding target distributions of values are also shown in Figure 3.

Each candidate solution generated by NSGA-II was assessed through 200 replicate simulation executions. The target distributions against which candidates are contrasted are derived from 500 replicates generated with the previous manual-calibration parameter values. The manual-calibration's replicates need be executed once only and stored, they do not change. In contrast, assessment of candidates is computationally costly because

so many are generated; a figure of 200 replicates per candidate was selected to strike a balance between experimental sensitivity and computational cost. A previous analysis of parametric perturbation in ARTIMMUS established that contrasting distributions comprising 200 replicate executions was sufficient to detect ‘small’ changes in $\frac{2}{3}$ of simulation behaviour metrics, and ‘medium’ in the remainder [11]. Hence, we consider 200 replicates to offer sufficient sensitivity in differentiating candidate performances. These effect size categories arise from the analysis’s use of the Vargha-Delaney *A test* [25], which provides interpretation guidelines. For reference, the *A test* is a non-parametric effect magnitude test representing the probability that a randomly selected member of one distribution is larger than a randomly selected member of the other. An *A test* score of 0.5 indicates the two distributions are indistinguishable (using this test). Values of 1 and 0 indicate no overlap in the two distributions. A single calibration exercise required around 5 days on a dedicated computational cluster able to execute 120 simulations simultaneously; each single simulation replicate takes around 2-10 minutes to execute, depending on the parameter values used.

We have successfully applied MOC to both ARTIMMUS parameter values and initial conditions, but focus here on the former. Initial condition calibration results are reported in the supplementary materials. Calibration was performed over 8 ARTIMMUS parameters which all pertain to presentation of substances to T cells, particularly Th1 and Th2 cells, and their resultant development. The biology captured in these parameters is outlined in supplementary Figure S1, and we note that a thorough understanding of this biology is not required to appreciate our results. These parameters were selected for the reasons that ascertaining their values experimentally would be challenging and they all relate to a critical aspect of the biology: the perpetuation of autoimmunity, and (for some) it’s amelioration (as Treg cell development is also directed by DCs). Hence, by successfully calibrating parameter values that are highly influential on simulation dynamics we demonstrate MOC’s potential. Parameters were given a constrained range of values that the MOC process could assign, being zero to twice their manually-calibrated range, as shown in Table 1. In exploring the space of putative parameter values, NSGA-II maintained a population of 64 candidate solutions which were subject to genetic recombination and mutation (see [17]) over 32 generations of natural selection, wherein only the best 64 solutions (i.e. those on or near the Pareto front) were retained in the successive

generation.

This calibration exercise was repeated three times for both parameters and initial conditions. Figure 3 shows the solution with the lowest Λ score from one such parameter calibration. The remaining two are shown in supplementary Figures S2 and S3. The calibrated simulation dynamics closely resemble the target distributions in all cases. The three parameter calibration exercises generated, respectively, Pareto fronts constituting 82, 87 and 112 Pareto-equivalent solutions. The ranges of parameter values represented across the Pareto fronts' solutions in each independent calibration exercise are shown in Figure 4, as are the baseline manually-calibrated values. In all but one case the baseline parameter value sat within the range of non-outlier MOC-derived values, the exception being *Th1_diff80* in exercise 3. Hence, we conclude that MOC is an effective means of calibration: it has repeatedly reproduced ARTIMMUS dynamics that were known possible, and has identified similar solutions, in the form of parameter values, that do so.

Next we investigated how the space of ARTIMMUS parameter values relates to the space of successful target dynamic reproductions, i.e., tradeoffs in objective values. We find statistically significant ($p < 0.01$) differences between calibration exercises' distributions of calibrated parameter values for 7 of 8 parameters, Figure 4. This corresponds to 19 of 24 (79%) of pairwise comparisons. Further, 75% (18/24) pairwise comparisons register a KS value ≥ 0.3 . For context, a KS value of 1.0 indicates no overlap between 2 distributions. In contrast, this degree of variation is not observed in Pareto fronts' objective values, depicted in Figure 5. Here we instead find statistically significant differences in only 27% (4/15) pair-wise calibration comparisons, and only 27% (4/15) of comparisons register $KS \geq 0.3$. We find no evidence of objectives that are harder to calibrate than others; the smallest objective values are < 0.05 in all cases, and the median objective values all lie under 0.17.

Together, these data suggest a redundancy in the ability for parameter values to deliver particular objective scores. This corresponds to a landscape wherein parameter values mapped to objective values is relatively flat, as a wide range of ARTIMMUS parameter values deliver relatively similar objective scores. The results of using MOC to calibrate ARTIMMUS initial conditions are reported in supplementary Section S1, and supplementary Figures S4, S5 and S6. They are qualitatively identical to our findings in calibrating parameters, and support the conclusions drawn here.

An obvious question is, why does MOC not deliver any perfectly calibrated solutions, wherein all objective scores are 0.0? The best solutions, determined by their minimal Λ values, in each calibration exercise are shown in Table 2. Objective KS values ranged from 0.05 to 0.14 (and 0.03 to 0.12 for initial conditions). We attribute the inability to deliver a perfect calibration to the stochastic nature of ARTIMMUS, wherein 200 replicate executions for a given candidate yields sufficient variation so as to deliver objective KS scores of ≥ 0.05 . There is a risk that improvements in objective KS values that are already so small cannot be confidently attributed to an actual improved simulation calibration, as opposed to stochastic variation between simulation replicates. Section 7, below, explores a method for terminating the MOC process on the premise that further effort will not deliver better quality solutions.

These data collectively highlight the challenges in exactly calibrating (i.e. $KS=0.0$) simulations to several objectives simultaneously. As such, we consider in the next Section the implications on experimental results of adopting baseline simulation values from different extremes of the Pareto front.

5 Scientific significance of imperfect calibration

As demonstrated above, MOC delivers a host of solutions to a given calibration problem, each representing an optimal tradeoff in calibration criteria (see Figure 2). It falls on the simulation developer to decide which to adopt baseline parameter values in subsequent experimentation. There is a risk that whilst calibration solutions lying in different regions of parameter space give rise to Pareto-equivalent solutions, they do not behave in a consistent manner when further experiments are performed. In such a case, a simulation-based experiment would lead to different conclusions depending on which calibration result was adopted as the baseline. In this section we investigate the extent to which this phenomenon holds.

The manually-calibrated ARTIMMUS simulation was previously used to elucidate the effect of removing a central immune organ, the spleen (a *splenectomy*), in EAE-induced animals [18]. Previous experiments had demonstrated that splenectomy in rats prior to the induction of EAE increased the mortality rate and hampered recovery [26]. Simulating splenectomy in ARTIMMUS revealed the spleen as a primary site for the generation of

autoimmunity-combating CD4Treg and CD8Treg cells. The reduced Treg populations resulting from the spleen’s removal prior to EAE-induction were unable to completely abrogate the autoimmunity-inducing Th1 populations, allowing for their re-expansion, and thus facilitating increased disease severity and relapses.

Here we explore whether the results of splenectomy in ARTIMMUS differ when baseline parameter values are adopted from disparate extremes of the Pareto front. The experimental procedure is highlighted in Figure 6. First, Pareto front solutions representing the extreme values, both low and high, of objective KS measures are identified. These solutions represent extremes in the range of simulation dynamics encapsulated within the Pareto front. For each solution 200 simulation replicates are performed for both control and splenectomy groups. Key performance indicators (KPI) are extracted from the resultant distributions of 200 simulation executions in each group. The performance indicators used are identical to those of the original ARTIMMUS splenectomy experiment [18]: the peak population sizes for each T cell population in the simulation, the times at which these peaks are reached, and the number of Th1 cells remaining at day 40 (giving a total of 9). For each KPI, the distributions of values obtained for control and splenectomy groups are contrasted using the Vargha-Delaney *A test* [25], as per the original experiment [18]. This procedure is repeated for each of the three calibration exercises reported in Section 4. The resultant *A test* scores are shown in Figure 6’s tables. Also shown, for context, are the *A test* scores of the original ARTIMMUS experiment [18].

Broadly speaking, the splenectomy results generated by Pareto-equivalent solutions are consistent with one another, and with the original experiment. There exceptions, however, wherein differences in *A test* scores reported for solution and the original experiment differed substantially: g23c60 in exercise 1, and g6c35 and g30c58 in exercise 2. These differences occurred for ‘Th1 at 40d’, ‘Th2 peak’ and ‘Th2 Time’ KPIs. Of interest, three of these solutions were obtained from the region of the Pareto front where alignment with target Th2 peak population size was poorest. In the case of g23c60 and g6c35, exercises 1 and 2 respectively, the parameter values were sufficient to return Th1 population size at 40 days to control group levels, despite the splenectomy ($A=0.58$ and 0.56 ; 0.5 indicates no difference). This is significant, as the principle conclusion of the original experiment was that splenectomy reduces Treg population sizes to levels unable to suppress Th1 cell populations and abrogate autoimmunity. The time series T cell population dynamics of

both these solutions under control and splenectomy are shown in supplementary Figure S9. In both cases the peak Th1 population sizes are smaller than in the original experiment (see Figure 6)), and the Th2 population sizes are substantially larger. Based on this we hypothesize that despite reduced Treg population sizes resulting from splenectomy, the altered balance between Th1 and Th2 populations which compete with one another is sufficient to abrogate the Th1 population at day 30 in these solutions.

Supporting the notion that solutions' results are relatively consistent, the direction of change in solutions' KPIs resulting from splenectomy differs from the original experiment in only a minority of cases. Further, this occurs only in KPIs for which the original experiment reports a comparatively small change between splenectomy and control, the largest being in exercise 2 when the original experiment reports a change of $A=0.66$, which was not interpreted as significant.

We have conducted the same investigation on Pareto-equivalent solutions generated under the three independent initial condition calibration exercises (supplementary Section S1). Detailed analysis is reported in supplementary Section S2 and Figure S10; briefly, divergences between initial condition solution and original experiments were smaller than reported here for parameters. We take this to indicate that the initial parameters investigated were less influential on simulation behaviour than the parameters investigated here.

In summary, the conclusions that would be drawn from adopting baseline parameters values from disparate Pareto-equivalent solutions are mostly, but not completely, consistent with one another and with the original splenectomy experiment. There were two notable exceptions, and they underscore the importance of considering the range of simulation performances that satisfy a calibration exercise. Making these explicit through Pareto fronts is a strength of the MOC approach. It remains important to, where possible, further evaluate Pareto-equivalent solutions in the context of domain knowledge and expertise, which might have ruled out the two exceptions noted above, as the Th2 population size is abnormally large compared to the Th1 population. Where this is not possible, where no grounds to discard some Pareto-equivalent solutions exist, we advise that experiments are performed in replicate adopting a wide range of calibration solutions and that conclusions are drawn after taking stock of the full range of results generated. This is particularly important if quantitative, rather than qualitative, results are sought; our

present data show more divergence between calibration solutions and original experiment in the quantitative case.

6 Multi-objective calibration delivers previously unseen disease phenotypes

In Section 4, above, MOC successfully reproduced simulation dynamics known to exist by virtue of a prior manual calibration. To further demonstrate MOC’s generality and utility, we now derive simulation dynamics not known to exist *a priori*.

ARTIMMUS’s baseline behaviour constitutes a period of autoimmunity followed by recovery, reflecting typical biological disease [21][22]. However, disease susceptibility and severity vary considerably between mouse strains and between mice within a given strain [27][28]. Furthermore, depletion or incapacitation of CD4Treg and CD8Treg cells leads to exacerbated disease symptoms [29][30]. Here we investigate the capacity for ARTIMMUS to reproduce persisting disease symptoms of varying severity. To reflect potential genetic differences between mouse strains, we calibrate over initial conditions specifying cell population sizes, and a parameter controlling the efficiency of Th1 killing by CD8Treg cells; together comprising 9 variables. In this experiment we are implicitly investigating whether variation in these basal population sizes and the efficiency of the CD8Treg-Th1 killing pathways could explain the differences in autoimmune phenotypes observed between mouse strains and individuals therein.

Three persisting disease severities are investigated, ranging from mild to severe. These are captured by defining the distribution of Th1 cells remaining at 60 days as a target for calibration, captured as a Gaussian distribution. Mild, moderate and severe disease are represented with mean (μ) and standard deviation (σ) values of $\mu=50$ & $\sigma=10$, $\mu=200$ & $\sigma=100$, and $\mu=500$ & $\sigma=200$ respectively. To ensure an aggressive onset of autoimmunity, consistent with animal models, a second calibration target distribution of $\mu=1000$ & $\sigma=200$ Th1 cells at 15 days is employed.

Each persisting autoimmunity severity is independently calibrated three times, representatives of which are shown in Figure 7 (the remainder are shown in supplementary Figures S11, S12 and S13). Automated calibration successfully delivers the required median number of cells in most cases, with $KS \leq 0.2$ in 6 of the 9 calibrations. However, the

spread of the ‘Th1 cells at 60 days’ distribution for mild persisting disease is notably less well calibrated, with all three calibrations delivering $KS > 0.3$.

Together, these data support the general applicability of MOC to problems where a simulation’s ability to deliver a desired dynamic is not known *a priori*. These data also suggest that the heterogeneity in disease severities observed in experimental animals could be attributed to differences in basal population sizes and regulatory pathway efficiency.

7 When to stop MOC

A key consideration in any optimisation task is the stopping criteria. For MOC, underpinned by the NSGA-II optimisation algorithm, this equates to determining when to stop calibration.

Overfitting describes the case where the simulation being calibrated starts to capture the noise in the target distributions, rather than the trends those distributions represent. This is a particular issue when target distributions do not contain many samples, as might be the case if they represent biological experiments (Figure 8A). For example, studies involving experimental animals can require their sacrifice to collect data. As such, it is considered unethical (and is practically cumbersome) to collect hundreds of samples, and 5 to 10 are more typical. These smaller sample sizes are unlikely to perfectly capture the underlying distribution that would emerge if thousands of samples were available. Overfitting is said to have occurred when the calibrated simulation better reflects these 5-10 samples than their underlying distribution, as illustrated in Figure 8B.

A common strategy in single-objective (not MOC, which is multi-objective) problems for determining when to terminate an optimisation process is to segregate the available data into two parts, termed ‘training’ and ‘validation’ datasets. The training dataset is used as normal to search for improved solutions, akin to MOC’s target data. The validation dataset is used as an independent check for overfitting of solutions to the training data set. Such a case of overfitting is depicted in Figure 8C. Both the training and validation data roughly reflect the underlying distribution, from which they were sampled. The candidate solution more closely resembles the training dataset than either the underlying distribution or the validation dataset, hence, it is overfitted. As illustrated in Figure 8D, in the earlier stages of optimisation successive candidate solutions that

better capture the training dataset will also better capture the validation data. It is only when overfitting starts to occur that performance against the validation data worsens whilst performance against training data continues to improve. It is at this point that the optimisation process is best terminated.

MOC is, however, a multi-objective optimisation problem, and it is unclear in the literature how this overfitting detection strategy ought be applied. We propose here a novel strategy for detecting overfitting in multi-objective problems based on co-membership of solutions to both training and validation dataset Pareto fronts (P_t and P_v), maintained throughout the calibration process (Figure 8E). The overfittedness at a given point in the optimisation process is reflected in the proportion of P_t members that are not members of P_v . The following algorithm performs the calculation:

```

     $m \leftarrow 0$ 
    for all  $i \in P_t$  do
        if  $i \in P_v$  then
             $m \leftarrow m + 1$ 
        end if
    end for
    return  $1 - (m/\text{size}(P_t))$ 

```

A proportion of 0 indicates that all training dataset Pareto solutions are also members of the validation Pareto front. At the other extreme, a value of 1 indicates that the training dataset Pareto front has been completely over-fitted, as none of its members are Pareto optimal with respect to the validation dataset. A threshold level of over-fitting at which the optimisation process (i.e., MOC) is to be terminated can be selected by the simulation experimenter.

We investigated different overfitting thresholds for MOC termination in the three ARTIMMUS parameter recalibration exercises reported in Section 4 above. An additional 214 simulation replicates using manually-calibrated parameter values were acquired to use as a validation dataset, constituting a 70-30 (500-215) training-validation data split. The validation dataset Pareto front for each iteration of the MOC algorithm (generation) was determined, and the overfittedness calculated. Figure 9A shows how, as MOC progresses, the proportion of overfitted candidate solutions on the training (target) dataset increases for each of the three calibration exercises. Figure 9B shows the point at which MOC

calibration would have been terminated should a given overfittedness threshold have been selected. Had we employed a overfittedness termination threshold of 0.5, wherein half of the training dataset Pareto front is overfitted, calibration would have terminated at generation 14, 15 or 23 (for exercises 1, 2 and 3 respectively) instead of 32. Given that each of these calibration exercises required around 7 days to complete on a dedicated computing cluster, this speed-up is substantial. We note that these combined training and validation datasets constitute 714 data points, considerably exceeding what might be obtained from real biological experiments. We anticipate that with fewer data points overfitting will occur sooner in the MOC process.

8 Discussion

Simulation represents a powerful tool to advance the investigation of biological systems, particularly when used in tandem with traditional approaches. As more complex biological systems become the subject of simulation a challenge in their calibration emerges: complex biological systems cannot be characterised by single metrics alone. There exist technologies capable of identifying parameter values that align simulation dynamics with some desired target, but these operate on single metrics. Even in cases where parameter values can be ascertained experimentally, seemingly avoiding the need for calibration, the abstract nature of simulation can complicate their direct adoption. Here we have demonstrated how biological agent-based simulation parameter values can be derived using multi-objective optimisation, an approach we have termed Multi-Objective Calibration (MOC). Multi-objective optimisation algorithms find solutions to problems simultaneously described by more than one metric. In MOC the desired characteristics of the simulation, which can represent either established biological data to be reproduced or some desired hypothetical simulation outcome, are expressed as distributions. Importantly, several such characteristics can be expressed, and MOC identifies those sets of parameter values that deliver optimal tradeoffs against each.

We evaluated MOC on a well understood simulation, using it to reproduce a previous manual calibration effort and therein delivering a solution that was known to be possible. The ARTIMMUS simulation was used, which simulates a mouse multiple sclerosis disease model [18]. MOC delivered around 90 unique parameter value combinations,

each of which provided an optimal tradeoff in performance against the 5 target ARTIMMUS characteristics specified. This range of possible calibration solutions was unknown *a priori*; the previous manual calibration of ARTIMMUS having delivered only one such solution [11]. It would ordinarily fall on the simulation user to select one solution (set of parameter values) to adopt as a baseline for subsequent simulation experimentation. We investigated the significance of selecting solutions representing different extremes of tradeoffs in delivering target simulation characteristics. A previous experiment with ARTIMMUS determined that removing the spleen, an important immune system organ, resulted in exacerbated autoimmune symptoms. The results of re-performing this experiment with different MOC solutions adopted as baseline parameter values were broadly, but not absolutely, similar. Hence, adopting different calibration solutions can lead to different experimental conclusions. It is a strength of MOC that this range of solutions is made explicit. Where possible, we recommend that MOC solutions be evaluated against biological data to discard those that represent biologically unrealistic parameter values or behaviours. Where this is not possible, we advocate performing experiments in replicate using multiple MOC solutions such that the full range of possible results be established before conclusions are drawn.

We demonstrated MOC in deriving simulation behaviours that were not known possible *a priori*: varying degrees of persisting autoimmunity in ARTIMMUS. MOC can be applied to both parameters and initial conditions, at the same time, as demonstrated in these calibration exercises. We do not consider simulation parameter values and initial conditions as independent; a poor selection of initial condition values coupled with appropriate parameter values can still fail to deliver the desired simulation dynamic. MOC's successful delivery of these previously unknown simulation dynamics presents an interesting use case for MOC. It could be used to identify which parameters, and hence components and pathways, need be manipulated to resolve a simulated disease state, therein highlighting candidate therapeutic targets. Furthermore, for disease simulations that incorporate potential interventions, MOC can be used to determine optimal intervention strategies that exploit synergies between several treatment options.

We surmise that MOC can support model selection and development. Accurately simulating a biological system requires both an appropriate model of the biology, and appropriate parameter values for that model. There typically exist several options for how

to represent a biological concept in simulation, the most suitable of which is often unclear. Models must strike a balance between including sufficient complexity to accurately reflect the biology’s dynamics, whilst remaining sufficiently simplistic to offer insight. The unsuccessful calibration of a given model of the biology can lead to two conclusions; first, that the calibration process was simply unsuccessful in finding a solution that does exist, a risk we argue is greatly lessened through MOC; or second, that the model is incapable of replicating the biological dynamics in question. In this latter case, MOC can inform simulation design, where a succession of putative models can be evaluated until calibration is successful. The possibility of directly applying MOC to the space of biological abstractions, rather than parameter values, is intriguing, though extremely challenging technically. Here, MOC would search for which cells were represented, and how. This would encompass their interactions with one another, opting to ignore some found to be irrelevant to the biological phenomenon of interest, or *vice versa*. The level of detail through which molecular secretions and expressions were represented could also be determined; is variable expression level necessary, or does simply ‘present’ vs ‘not’ suffice? The challenge herein lies in building an agent-based simulation infrastructure capable of capturing all these possibilities, and allowing the automated optimisation process to manipulate them. The aforementioned point still applies, for each possible model, the space of parameter values must also be investigated, as an accurate reflection of biology requires both an appropriate model and corresponding parameter values. Hence, MOC would be applied in a nested fashion, firstly over the space of biological representations, and therein over the space of parameter values for each model.

Although our present investigation has employed an agent-based simulation, MOC is applicable to other simulation paradigms also, such as ordinary differential equations (ODE). Application to non-stochastic simulations, such as ODEs, requires significantly less computational power, as there is no need to obtain simulation replicates in assessing a candidate solution’s fitness. We note that, from our experience in building them, not all biological simulations are as computationally costly to execute and calibrate as ARTIMMUS. Each MOC calibration exercise has taken up to a week of time on a dedicated computational facility. In this regard, terminating the MOC process when a threshold level of overfitting is detected is pertinent (see Figure 8). Overfitting was detected in all three of our ARTIMMUS parameter recalibration exercises, and selecting a threshold

of 0.5, wherein half of the MOC solutions at a given point no longer represent optimal performance tradeoffs in an independent test, could as much as halve the computational effort required.

The ability to detect overfitting in a multi-objective context is a novel contribution of this work. Though a common strategy for stopping a single-objective optimisation process, it was previously unclear how to deploy this strategy in a multi-objective context [31]. There is another condition under which we feel it pertinent to terminate the MOC process. The goal of MOC is to find parameter values yielding simulation dynamics that closely resemble some target. As this alignment increases, and differences in solutions' simulation performances reduce, it is possible that seemingly better alignments in fact represent sampling artefacts arising from the stochastic simulation, rather than genuinely superior parameter values. We note that detecting this in a statistically robust manner is challenging, and as such we highlight it as potential further work.

This work fits within the context of a wider framework for supporting complex system simulation, the CoSMoS framework [32]. CoSMoS advocates explicitly recording, typically through graphical modelling [33], how real world concepts are translated into computer code, and the implicit assumptions therein. In this context, MOC can help in relating simulation results to biological data. The case where a distribution of results emerges from a given biological experiment, even to the point where replicates or individuals within an experiment exhibit completely different outcomes, can be handled in MOC by defining bi-modal (or multi-modal) target distributions. A scenario wherein MOC unexpectedly delivers several distinct and unconnected simulation phenotypes, rather than a continuum of points on the Pareto front, is interesting. This can either suggest the existence of additional phenotypes to look for in the biology, or if this can be ruled out, suggests instead that the model being calibrated fails to accurately capture the biology. This latter case is an example of how MOC could drive simulation design and development, as covered above. Related work on supporting the link of simulation to biology proposes the construction of an argument wherein a claim such as 'this simulation is an adequate representation of the biology' is supported by explicitly cited evidence [34]. In this context, application of MOC can raise confidence that appropriate parameter and initial condition values have been identified. The range of possible values can be contrasted against biological literature and data, excluding those deemed implausible. Subsequent

simulation experiments can be performed in replicate with those that remain, therein highlighting the full range of results that are plausible in absence of better reason to rule out particular parameter values. We argue that drawing conclusions from this nature of simulation experimentation, and making explicit the full range of parameter values that satisfy the calibration problem, leads to more robust conclusions.

In summary, our novel application of multi-objective optimisation in MOC presents the multi-objective optimisation community with a new field of application, and one we feel has considerable scope for growth. Importantly, it provides fundamental support for a critical aspect of simulation-based biological experimentation: identifying parameter values and initial conditions that align simulations with a complex target behaviour.

Competing interests

We have no competing interests.

Authors' contributions

MR conceived and designed the study, carried out all data generation and analysis reported herein, and drafted the manuscript. KA conceived the study and drafted the manuscript. LR conceived the study. JT conceived the study and drafted the manuscript.

Acknowledgment

We thank Dr. Fabian Held, Prof. Susan Stepney and Dr. Simon Poulding for engaging discussions on statistics.

Funding

MR is supported by the David & Judith Coffey Life Lab. JT is part-funded by The Royal Society.

References

- [1] Bauer AL, Beauchemin CA, Perelson AS. Agent-Based Modeling of Host-Pathogen Systems: The Successes and Challenges. *Inf Sci (Ny)*. 2009;179(10):1379–1389.
- [2] Harris TH, Banigan EJ, Christian Da, Konradt C, Tait Wojno ED, Norose K, et al. Generalized Lévy Walks and the Role of Chemokines in Migration of Effector CD8+ T Cells. *Nature*. 2012;486(7404):545–548.
- [3] An G, Mi Q, Dutta-Moscato J, Vodovotz Y. Agent-Based Models in Translational Systems Biology. *Wiley Interdiscip Rev Syst Biol Med*. 2009;1(2):159171.
- [4] Pienaar E, Dartois V, Linderman JJ, Kirschner DE. In Silico Evaluation and Exploration of Antibiotic Tuberculosis Treatment Regimens. *BMC Systems Biology*. 2015;9(1):79.
- [5] Thorne BC, Bailey AM, DeSimone DW, Peirce SM. Agent-Based Modeling of Multicell Morphogenic Processes during Development. *Birth Defects Research Part C, Embryo Today: Reviews*. 2007;81(4):344–353.
- [6] Senior AM, Charleston MA, Lihoreau M, Buhl J, Raubenheimer D, Simpson SJ. Evolving Nutritional Strategies in the Presence of Competition: a Geometric Agent-Based Model. *PLoS Computational Biology*. 2015;11(3):e1004111.
- [7] Kitano H. Systems Biology: a Brief Overview. *Science*. 2002;295(5560):1662–4.
- [8] Kitano H. Computational Systems Biology. *Nature*. 2002;420:206–210.
- [9] Cosgrove J, Butler J, Alden K, Read M, Kumar V, Cucurull-Sanchz L, et al. Agent-Based Modelling in Systems Pharmacology. *CPT: Pharmacometrics & Systems Pharmacology*. 2015;4(11):615629.
- [10] Dong C, Martinez GJ. T cells: the Usual Subsets; 2010. *Nature Reviews Immunology* poster: www.nature.com/nri/posters/tcellsubsets.
- [11] Read M, Andrews PS, Timmis J, Kumar V. Techniques for Grounding Agent-Based Simulations in the Real Domain: a case study in Experimental Autoimmune Encephalomyelitis. *Mathematical and Computer Modelling of Dynamical Systems*. 2012;18(1):67–86.

- [12] Calvez B, Hutzler G. Automatic Tuning of Agent-Based Models Using Genetic Algorithms. In: Sichman JS, Antunes L, editors. Multi-Agent-Based Simulation VI. Springer Berlin Heidelberg; 2006. p. 41–57.
- [13] Fabretti A. On the Problem of Calibrating an Agent Based Model for Financial Markets. *Journal of Economic Interaction and Coordination*. 2012;8(2):277–293.
- [14] Burnham KP, Anderson DR. Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach (2nd ed). vol. 172; 2002.
- [15] Cohen IR. Tending Adam’s Garden : Evolving the Cognitive Immune Self. Elsevier Academic Press; 2004.
- [16] Kitano H. Biological Robustness. *Nature Reviews Genetics*. 2004;5(11):826–837.
- [17] Deb K, Pratap A, Agarwal S, Meyarivan T. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*. 2002;6(2):182–197.
- [18] Read M, Andrews PS, Timmis J, Williams RA, Greaves RB, Sheng H, et al. Determining Disease Intervention Strategies using Spatially Resolved Simulations. *PLOS ONE*. 2013;8(11):e80506.
- [19] Baxter AG. The Origin and Application of Experimental Autoimmune Encephalomyelitis. *Nature Reviews Immunology*. 2007;7(11):904–912.
- [20] Pachner AR. Experimental Models of Multiple Sclerosis. *Current Opinion in Neurology*. 2011;24(3):291–299.
- [21] Kumar V, Sercarz E. An Integrative Model of Regulation Centered on Recognition of TCR Peptide/MHC Complexes. *Immunol Rev*. 2001;182:113–121.
- [22] Kumar V. Homeostatic Control of Immunity by TCR Peptide-Specific Tregs. *J Clin Invest*. 2004;114(9):1222–1226.
- [23] Greaves RB, Read M, Timmis J, Andrews PS, Butler JA, Gerckens BO, et al. In Silico Investigation of Novel Biological Pathways: the Role of CD200 in Regulation of T Cell Priming in Experimental Autoimmune Encephalomyelitis. *Biosystems*. 2013;112(2):107–121.

- [24] Williams RA, Greaves R, Read M, Timmis J, Andrews PS, Kumar V. In Silico Investigation into Dendritic Cell Regulation of CD8Treg Mediated Killing of Th1 Cells in Murine Experimental Autoimmune Encephalomyelitis. BMC Bioinformatics. 2013;14 Suppl 6:S9.
- [25] Vargha A, Delaney HD. A Critique and Improvement of the CL Common Language Effect Size Statistics of McGraw and Wong. Journal of Educational and Behavioral Statistics. 2000;25(2):101–132.
- [26] Ben-Nun A, Ron Y, Cohen IR. Spontaneous Remission of Autoimmune Encephalomyelitis is Inhibited by Splenectomy, Thymectomy or Ageing. Nature. 1980;288(5789):389–390.
- [27] Levine S, Sowinski R. Experimental Allergic Encephalomyelitis in Inbred and Outbred Mice. Journal of Immunology. 1973;110(1):139–43.
- [28] Gold R. Understanding Pathogenesis and Therapy of Multiple Sclerosis via Animal Models: 70 Years of Merits and Culprits in Experimental Autoimmune Encephalomyelitis Research. Brain. 2006;129(8):1953–1971.
- [29] Beeston T, Smith TR, Maricic I, Tang X, Kumar V. Involvement of IFN- γ and Perforin, but not Fas/FasL Interactions in Regulatory T Cell-Mediated Suppression of Experimental Autoimmune Encephalomyelitis. Journal of Neuroimmunology. 2010;15(229):91–97.
- [30] Kumar V, Stellrecht K, Sercarz E. Inactivation of T Cell Receptor Peptide-Specific CD4 Regulatory T Cells Induces Chronic Experimental Autoimmune Encephalomyelitis (EAE). J Exp Med. 1996;184(5):1609–1617.
- [31] Dos Santos EM, Sabourin R, Maupin P. Overfitting Cautious Selection of Classifier Ensembles with Genetic Algorithms. Information Fusion. 2009;10(2):150–162.
- [32] Bown J, Andrews PS, Deeni Y, Goltsov A, Idowu M, Polack FAC, et al. Engineering Simulations for Cancer Systems Biology. Current Drug Targets. 2012;13(12):1560–1574.

[33] Read M, Andrews PS, Timmis J, Kumar V. Modelling Biological Behaviours with the Unified Modelling Language: an Immunological Case Study and Critique. Journal of the Royal Society Interface. 2014;11(9):20140704.

[34] Alden K, Andrews PS, Polack FAC, Veiga-fernandes H, Coles MC, Timmis J, et al. Using Argument Notation to Engineer Biological Simulations with Increased Confidence. Journal of the Royal Society Interface. 2015;12:20141059.

Figure and table captions

Figure 1. The ARTIMMUS simulation, used as a testcase for evaluating MOC.

A, the the major cell types represented in ARTIMMUS, and their key influences on one another. Red and green arrows respectively indicate activities that perpetuate autoimmunity or mediate recovery. Figure adapted from [11]. **B**, the baseline dynamic of ARTIMMUS, depicting four T cell population sizes over time. The simulation behaviour depicted here forms a calibration target for MOC in Section 4. Lines correspond to like-coloured cells in Figure A; these colours are maintained throughout the manuscript. Error bars capture 90% of the data derived from 500 simulation executions, timeseries lines indicate median population sizes at each time point.

Figure 2. Overview of the Multi-Objective Calibration (MOC) concept. A,

The desired (target) simulation dynamics are defined as distributions (only 2 shown): the desired distributions of peak cell number and the times at which these occur. Distributions are depicted as histograms, or the corresponding cumulative distribution functions describing the proportion of samples in the distribution (y axis) that hold a given value or less (x axis). **B**, the capacity for putative simulation parameter (only 2 shown) values, termed *candidate solutions*, to reproduce target dynamics is evaluated. The Kolmogorov-Smirnov (KS) statistic quantifies the difference between target and a given candidate solution's simulation performance (left); this metric is termed an *objective*. By sampling and evaluating regions of parameter space we identify those that provide good alignment with a given objective, illustrated through greyscale heatmaps (right). No single region of parameter space maximizes performance against all objectives (only 2 shown), there

exist inherent tradeoffs. A heuristic (guided) search strategy, NSGA-II, is employed to strategically sample parameter space. **C**, *solutions* representing optimal tradeoffs in performance against each objective are identified, collectively termed the *Pareto front* (left). These solutions are *Pareto-equivalent* (pink): no solution has been found that represents an improvement in one objective without a worsening in another. Sub-optimal candidate solutions are discarded (blue). Pareto-equivalent solutions may reside in disparate regions of parameter space(right).

Figure 3. Multi-objective calibration (MOC) successfully re-calibrates ARTIMMUS parameters against 5 objectives. The best solution's, that with lowest A score, target simulation dynamics are shown. The solution dataset comprises 200 simulation replicates, the target comprises 500. **A**, T cell population sizes over time, for both target (dotted line) and solution (solid line). The median values from each dataset at the given point in the time series are plotted. **B-F**, cumulative distribution functions showing alignment of solution and target distributions of values for each objective, with titles giving KS values. These graphs show the distribution of calibration target values obtained in each dataset: the y-axis indicates the proportion of items in the distribution holding a value less than or equal to the corresponding x-axis value. Objectives are: B, peak CD4Th1 population size cell; C, time at which this peak occurs; D, CD4Th2 population size at 30 days; E, peak CD4Treg population size; F, peak CD8Treg population size. These data represent the first of three independent recalibration experiments.

Figure 4. Automated re-calibration of ARTIMMUS parameters delivers solutions approximating the original manually-calibrated parameter values. Box plots are shown for each of three independent calibration exercises. The horizontal green line represents the manually-calibrated parameter values. Calibration was performed over 5 objectives: the peak population sizes of Th1, CD4Treg, CD8Treg cells, the time at which the Th1 population peaks, and the number of Th2 cells at 30 days. Parameters subject to calibration are listed in Table 1, see Figure S1 for an explanation of their operation in ARTIMMUS. Values shown above each plot are the Kolmogorov-Smirnov scores between

distributions, shown to one significant figure; the associated p-values are: *, $p < 0.01$ and **, $p < 0.001$. Outliers in boxplots are defined as lying beyond the first or third quartiles by 1.5 times the interquartile range.

Figure 5. The range of objective values that constitute the Pareto front derived through MOC re-calibration of ARTIMMUS. Box plots are shown for each of three independent calibration exercises. These objective values correspond to the Pareto front and associated ARTIMMUS parameter values of Figure 4. Calibration was performed against five objectives: **A**, the peak population size of Th1 cells; **B**, the time at which this occurred; **C**, the number of Th2 cells at 30 days; **D**, the peak population size of CD4Treg cells; **E**, the peak population size of CD8Treg cells. Statistical and boxplot formatting are as in Figure 4.

Figure 6. Do different regions of MOC's Pareto front of solutions give rise to different results in subsequent experimentation? Top, an overview of the experimental procedure. **1**, Pareto front members representing objective value extremes are identified (only two objectives shown in example). **2**, The simulation parameters represented by such members are adopted in performing a control and splenectomy experiment, with 200 replicate simulations in each group. **3**, Key performance indicators are extracted from the resultant distributions of simulation dynamics, indicated here is the peak CD4Treg population size within each individual simulation. **4**, Performance indicators are statistically contrasted for splenectomy and control experiments. These statistics are examined across different Pareto front members, thereby gauging the extent to which experimental results critically depend on which Pareto-equivalent parameter values are adopted in simulation. **Tables**, columns represent extreme Pareto front solutions, defined as having either the highest or lowest KS value for each of the five objectives used in calibration (see Section 4). The objective KS value scores are shown in parentheses. Only the first occurrence of each solution is shown, with subsequent entries indicated by '-'. Rows indicate the difference between control and splenectomy simulations based on each solution according to key indicators of simulation behaviour, as measured by the Vargha-

Delaney *A test* [25]. The original *A test* scores for the manually-calibrated simulation are shown ('orig'), as is the biggest difference in *A test* score observed between manually- and automatically-calibrated simulations ('diff'). Values highlighted in red represent four differences in candidate and original *A test* scores that are notably larger than differences observed elsewhere. 'D.c.' indicates 'direction change', where there exists at least one candidate with for which the *A test* score lay on the other side of 0.5 from the original, indicating that the distribution of values under splenectomy increased in the original experiment but decreased for the candidate (or *vice versa*).

Figure 7. Employing MOC to discover parameter and initial condition values delivering simulation dynamics not known to exist *a priori*: persisting autoimmune states of varying severity. Three severities are explored, represented as columns. They are, a mean of 50, 200 or 500 Th1 cells at 60 days (with standard deviations of 10, 100 and 200 respectively). A second objective is employed in all cases, 1000 Th1 cells at 15 days, which drives the establishment of autoimmunity. Each severity is calibrated in three independent experiments, and shown here are the solutions exhibiting lowest Λ values from a representative calibration of each experiment. The first row of graphs depicts the median T cell time-series. The second row shows the candidate's performance against an objective of 1000 Th1 at 60 days (standard deviation = 200). The last row depicts the second objective, the (respective) number of T cells at 60 days.

Figure 8. Terminating MOC when overfitting occurs. Overfitting describes the case when solutions generated by an optimisation process, e.g. MOC, better resemble the target data than the underlying distribution from which it was drawn. **A**, in many contexts, such as animal experiments, only limited samples of a phenomenon can be obtained. The samples will broadly, but not exactly, reflect the underlying distribution. **B**, an overfitted candidate solution more closely resembles the target data than the underlying distribution from which the target data was drawn. Detecting this is difficult because the true underlying distribution cannot be absolutely known. **C**, a common strategy in single-objective optimisation problems is to divide the available data into two, a training

dataset and a validation dataset. The training dataset is used as the target in obtaining successively better quality solutions. The validation dataset is used as an independent check. Overfitting is detected when solutions more closely resemble the training dataset than validation dataset. This is illustrated in **D**, where early solutions generally offer improved performance against both datasets. It is only in later stages that solutions so closely reflect the target dataset that they diverge from the validation dataset. This is when the process should be stopped. **E**. Overfitting can be detected in multi-objective optimisation, such as MOC, by maintaining Pareto fronts of optimal solutions against both training and validation data independently. The degree of overfitting is reflected in the proportion of training data Pareto front solutions that are not members of the validation data Pareto front.

Figure 9. Empirical results for detecting overfitting in MOC, and when to terminate the process accordingly. We generated a validation dataset using ARTIMMUS’s previous manually-calibrated parameter values, and retrospectively analysed how overfitted MOC solutions would have been on the three MOC calibration exercises reported in Section 4. **A**, The overfittedness, defined as the proportion of MOC Pareto front solutions that are not also members of a similar Pareto front maintained for the validation data, at each MOC generation. **B**, the generation at which MOC would have been terminated for a given overfittedness threshold value.

Table 1. The ARTIMMUS parameters (top) and initial conditions (bottom) subject to calibration, their baseline (manually-calibrated) values, and the lower and upper bounds of values they may be assigned during MOC.

Table 2. The best solution, being that with the lowest Λ value, arising from each of three independent calibration exercises. Shown are each of the five objective KS values. We independently investigated the calibration of both ARTIMMUS parameters (top) and initial conditions (bottom). High quality calibrations, as indicated by low KS values, were

910 obtained in all cases.

911 **Figures and tables**



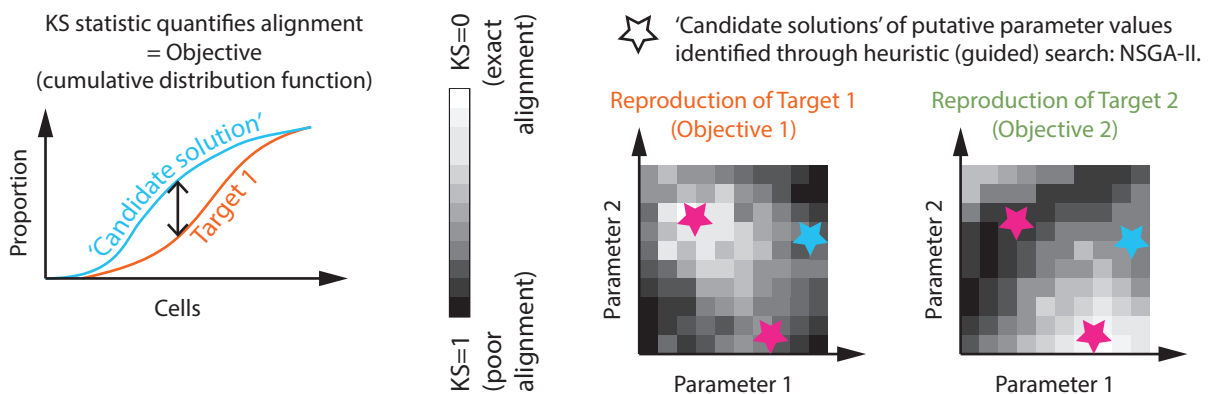


Figure 1:

A. Define target simulation dynamics



B. Evaluate putative simulation parameters' reproduction of target dynamic



C. Identify Pareto-optimal solutions (the Pareto front)



Figure 2:

Calibrate ARTIMMUS parameters, exercise 1




Figure 3:




Figure 4:


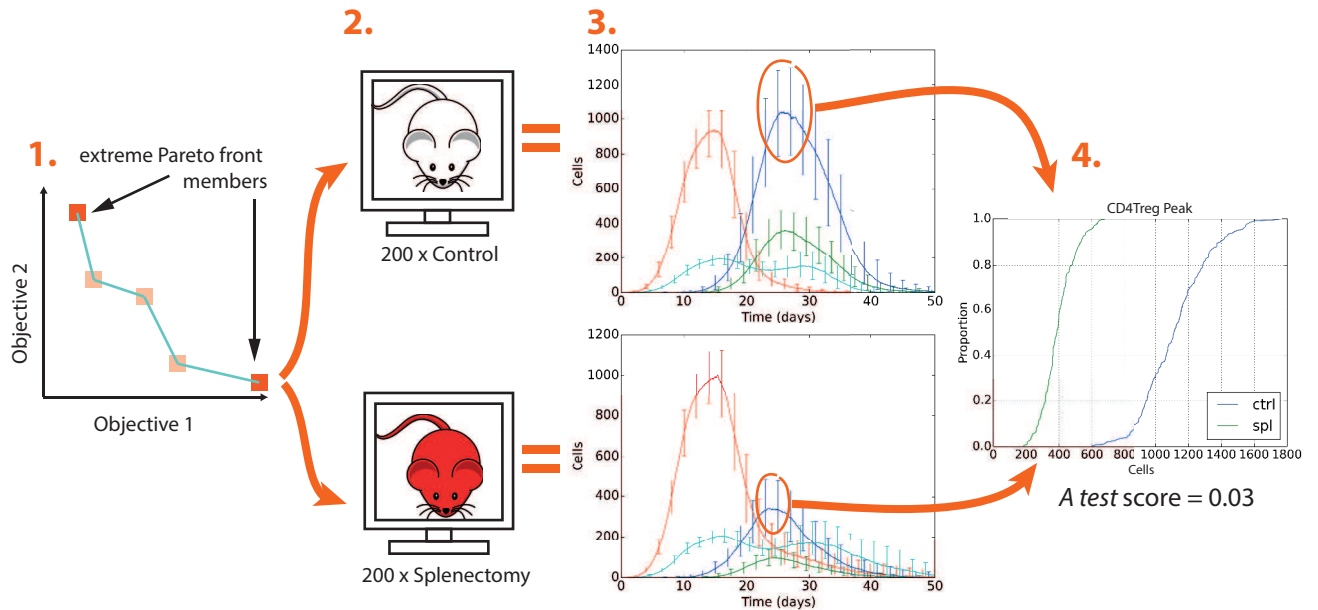


Figure 5:



Calibration Exercise 1, Vargha-Delaney A Test Scores													
Response	Th1Peak		Th1Time		Th2at30d		CD4TregPeak		CD8TregPeak		orig	diff	d.c.
	g31c32 (KS=0.04)	g11c33 (0.81)	g29c37 (0.03)	g10c13 (0.60)	g31c6 (0.03)	g23c60 (0.73)	g30c0 (0.04)	g23c22 (0.65)	g10c13 (0.04)	g23c22 (0.51)			
Th1 Peak	0.70	0.73	0.65	0.67	0.67	0.66	0.71	0.67	-	-	0.62	0.11	
Th1 Time	0.53	0.51	0.55	0.42	0.53	0.52	0.52	0.52	-	-	0.47	0.08	Y
Th2 Peak	0.60	0.74	0.66	0.66	0.67	0.67	0.61	0.62	-	-	0.66	0.08	
Th2 Time	0.47	0.53	0.68	0.46	0.51	0.56	0.53	0.47	-	-	0.58	0.11	Y
CD4Treg Peak	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-	-	0.00	0	
CD4Treg Time	0.23	0.24	0.26	0.26	0.20	0.28	0.20	0.24	-	-	0.21	0.07	
CD8Treg Peak	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-	-	0.00	0	
CD8Treg Time	0.24	0.24	0.29	0.27	0.23	0.25	0.25	0.26	-	-	0.23	0.06	
Th1 at 40d	0.98	0.96	0.92	0.89	0.96	0.58	0.96	0.94	-	-	0.95	0.37	

Calibration Exercise 2, Vargha-Delaney A Test Scores													
Response	g30c25 (KS=0.03)	g6c35 (0.94)	g30c34 (0.03)	g9c63 (0.33)	g15c14 (0.04)	g30c58 (0.98)	g17c61 (0.04)	g14c54 (0.77)	g9c56 (0.04)	g14c54 (0.67)	orig	diff	d.c.
Th1 Peak	0.63	0.69	0.66	0.66	0.71	0.65	0.65	0.66	0.68	-	0.62	0.09	
Th1 Time	0.53	0.49	0.55	0.47	0.50	0.53	0.54	0.51	0.48	-	0.47	0.08	Y
Th2 Peak	0.66	0.67	0.65	0.70	0.68	0.39	0.65	0.61	0.70	-	0.66	0.27	Y
Th2 Time	0.56	0.51	0.63	0.63	0.50	0.31	0.54	0.48	0.54	-	0.58	0.27	Y
CD4Treg Peak	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	-	0.00	0.01	
CD4Treg Time	0.27	0.30	0.23	0.19	0.20	0.21	0.17	0.22	0.29	-	0.21	0.09	
CD8Treg Peak	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	-	0.01	0.01	
CD8Treg Time	0.27	0.26	0.23	0.17	0.21	0.20	0.22	0.23	0.28	-	0.23	0.06	
Th1 at 40d	0.96	0.56	0.94	0.91	0.94	1.00	0.97	0.97	0.82	-	0.95	0.39	

Calibration Exercise 3, Vargha-Delaney A Test Scores													
Response	g5c39 (KS=0.03)	g23c10 (0.99)	g28c47 (0.04)	g13c21 (0.48)	g11c6 (0.04)	g5c39 (0.82)	g24c62 (0.03)	g23c10 (0.98)	g23c40 (0.04)	g23c10 (0.94)	orig	diff	d.c.
Th1 Peak	0.65	0.65	0.63	0.70	0.68	-	0.65	-	0.70	-	0.62	0.08	
Th1 Time	0.43	0.49	0.49	0.42	0.47	-	0.47	-	0.52	-	0.47	0.05	Y
Th2 Peak	0.67	0.64	0.70	0.68	0.67	-	0.67	-	0.67	-	0.66	0.04	
Th2 Time	0.59	0.51	0.52	0.50	0.47	-	0.62	-	0.57	-	0.58	0.11	Y
CD4Treg Peak	0.07	0.02	0.00	0.00	0.00	-	0.00	-	0.00	-	0.00	0.07	
CD4Treg Time	0.30	0.35	0.21	0.24	0.26	-	0.20	-	0.19	-	0.21	0.14	
CD8Treg Peak	0.04	0.01	0.00	0.00	0.00	-	0.00	-	0.00	-	0.00	0.04	
CD8Treg Time	0.27	0.40	0.17	0.26	0.27	-	0.20	-	0.21	-	0.23	0.17	
Th1 at 40d	0.93	0.83	0.97	0.90	0.92	-	0.97	-	0.95	-	0.95	0.12	

Figure 6:




Figure 7:




Figure 8:




Figure 9:

Parameters calibrated			
Parameter	Baseline value	Lower bound	Upper bound
<i>APC_immatureDuration</i>	48	0	96
<i>APC_matureDuration</i>	110	0	220
<i>APC_phagocytosisToPeptide</i>	0.02	0	0.04
<i>CNSM_MBPExpressionProbability</i>	0.2	0	0.4
<i>DCT1_cytokineSecretionRate</i>	10	0	20
<i>DC_T2CytokineRatio</i>	0.17	0	0.34
<i>Th1_diff00</i>	0.05	0	0.1
<i>Th1_diff80</i>	0.85	0	1.0
Initial conditions calibrated			
Initial condition	Baseline value	Lower bound	Upper bound
<i>numTh</i>	40	0	80
<i>numCD4Treg</i>	30	0	60
<i>numCD8Treg</i>	30	0	60
<i>numCNS</i>	500	0	1000
<i>numCNSMacrophage</i>	75	0	150
<i>numDC</i>	10	0	20
<i>numDCCNS</i>	40	0	80
<i>numDCSpleen</i>	100	0	200

Table 1:

Calibration on parameters					
Calibration exercise	Objective KS value				
	Th1Peak	Th1Time	Th2at30d	CD4TregPeak	CD8TregPeak
1	0.06	0.10	0.08	0.06	0.07
2	0.08	0.06	0.06	0.05	0.07
3	0.05	0.08	0.14	0.08	0.05

Calibration on initial conditions					
Calibration exercise	Objective KS value				
	Th1Peak	Th1Time	Th2at30d	CD4TregPeak	CD8TregPeak
1	0.06	0.08	0.04	0.03	0.06
2	0.04	0.08	0.10	0.11	0.12
3	0.06	0.06	0.06	0.07	0.05

Table 2: