

A systems engineering approach to validation of a pulmonary physiology simulator for clinical applications

A. Das¹, Z. Gao², P. P. Menon¹, J. G. Hardman³ and D. G. Bates^{1,*}

¹*College of Engineering, Mathematics and Physical Sciences, University of Exeter, UK*

²*Department of Electrical Engineering and Electronics, University of Liverpool, UK*

³*Division of Anaesthesia and Intensive Care, University of Nottingham, UK*

Physiological simulators which are intended for use in clinical environments face harsh expectations from medical practitioners; they must cope with significant levels of uncertainty arising from non-measurable parameters, population heterogeneity and disease heterogeneity, and their validation must provide watertight proof of their applicability and reliability in the clinical arena. This paper describes a systems engineering framework for the validation of an *in silico* simulation model of pulmonary physiology. We combine explicit modelling of uncertainty/variability with advanced global optimization methods to demonstrate that the model predictions never deviate from physiologically plausible values for realistic levels of parametric uncertainty. The simulation model considered here has been designed to represent a dynamic *in vivo* cardiopulmonary state iterating through a mass-conserving set of equations based on established physiological principles and has been developed for a direct clinical application in an intensive-care environment. The approach to uncertainty modelling is adapted from the current best practice in the field of systems and control engineering, and a range of advanced optimization methods are employed to check the robustness of the model, including sequential quadratic programming, mesh-adaptive direct search and genetic algorithms. An overview of these methods and a comparison of their reliability and computational efficiency in comparison to statistical approaches such as Monte Carlo simulation are provided. The results of our study indicate that the simulator provides robust predictions of arterial gas pressures for all realistic ranges of model parameters, and also demonstrate the general applicability of the proposed approach to model validation for physiological simulation.

Keywords: physiological simulation; systems engineering; model validation

1. INTRODUCTION

Mathematical modelling and computational simulation are becoming an increasingly important tool in the fields of medicine where *in vivo* studies are difficult, expensive or impractical [1]. However, a significant factor hampering the wider clinical exploitation of *in silico* simulation models is the current lack of rigorous procedures for model validation and verification, since doubts about model validity naturally have a strongly negative impact on the clinical applicability of any predictions arising from simulation studies. Current best practice in this area relies mainly on heuristic comparisons with previous clinical studies and other available models, together with some statistical analyses [2,3]. Indeed, many physiological models are designed and implemented using assumed values of their internal parameters without any attempt to explicitly incorporate uncertainty/variability in the model or to estimate the cumulative effect of uncertainty on model predictions. This contrasts sharply with current best

practice in the fields of systems and control engineering, where computer simulation models for safety-critical applications are routinely subjected to extensive programmes of verification and validation, using advanced analytical and computational approaches. Indeed, it could be argued that physiological simulators require special attention in this respect, since their internal parameters are often very poorly known or difficult to measure and may be subject to large variations across patient populations. As an example, consider the following equation which describes the pressure (P_{O_2})–saturation (S_{O_2}) relationship for the O_2 dissociation curve, as proposed by [4]

$$S_{O_2} = (((P_{O_2}^3 + 150 \cdot P_{O_2})^{-1} \times 23\,400) + 1)^{-1}. \quad (1.1a)$$

This relationship applies to a standard O_2 dissociation curve (at pH 7.4, temperature = 37, ignoring organophosphate affect). Appropriate correction factors (or in this case the input parameters to determine S_{O_2} as an output), such as the base excess BE, temperature T and pH correction factors also proposed by

*Author for correspondence (d.g.bates@exeter.ac.uk).

Severinghaus, are employed to get a standard O_2 tension (P_{O_2}) from the arterial O_2 tension (P'_{O_2}) to be used in the above equation (7.5005168 is the pressure conversion factor from kPa to mm Hg)

$$P_{O_2} = 7.5006168 \cdot P'_{O_2} \cdot 10^{[0.48(\text{pH}-7.4)-0.024(T-37)-0.0013 \cdot \text{BE}]}. \quad (1.1b)$$

A simple mathematical analysis reveals that, even in this very straightforward case, a level of uncertainty of only ± 5 per cent in each of these input parameters could lead to a possible ± 14 per cent variation in the resulting prediction of S_{O_2} . It should also be noted that the calculation of BE currently also relies upon determining the pH. Clearly, if equation (1.1) was used as part of a more complex nonlinear simulation model involving many other such equations, each with their own set of uncertain parameters, the possible effects on overall model predictions could be highly significant.

Interestingly, these issues also apply to computational modelling at the molecular level, and recent research in the new field of systems biology has highlighted the usefulness of the engineering concept of *robustness* for validating complex *in silico* simulation models of biochemical, genetic and signal transduction networks [5–7]. Robustness, in both engineered and natural systems, may be defined as the ability of the system to function correctly in the presence of both internal uncertainty/variability and/or external disturbances in its environment. This concept may be applied for the purposes of model validation as follows. Consider a situation where for a given set of model parameters, the outputs of a model (or of several different models) show a good match to the available experimental data. Such a model or models may be validated or invalidated by using analytical or computational approaches to evaluate the robustness of the responses of the proposed simulation model with respect to expected levels of uncertainty in the model parameters. If the model responses do not exhibit the same level of robustness as has been found for the real system *in vivo*, then this points to flaws in, or incompleteness of, the model. In addition, such analyses have the advantage of directly quantifying the cumulative effect of uncertainty in internal model parameters on the eventual predictions upon which new biological understanding or therapeutic strategies may be based. Fortunately, *robustness analysis* is now a well-established research area in the field of systems and control engineering, and many powerful approaches have been developed in recent years to quantify the robustness of complex nonlinear systems [8–11]. In this paper, we apply such approaches to compute worst-case deviations of predicted gas pressures from physiologically realistic values due to the effects of realistic levels of parametric uncertainty in a high fidelity simulation model of cardiopulmonary gas dynamics [2,3,12]. We rigorously analyse the effect of simultaneous variations in multiple uncertain parameters by formulating the problem as a multivariable nonlinear constrained optimization problem. We

also carry out a detailed comparison of the effectiveness and computational efficiency of a number of different global optimization algorithms in solving this problem and compare our results with standard statistical approaches (Monte Carlo simulation). Our aim is not to validate a particular setting of the model against data from an individual patient (the traditional approach) but to quantify its robustness (i.e. its ability to always produce realistic results for all expected levels of uncertainty in internal model parameters). Our aim is to produce a simulator that has been validated across the expected range of variation that would be found in multiple patients, since this represents a much more stringent test of the underlying model assumptions than simply checking the ability of the model to match a single set of data (it is entirely possible that a model which included significant errors might still, on occasion, produce results which matched one particular set of data).

The paper is organized as follows. Section 2 provides an overview of the simulation model and describes the formulation of the model validation task as an optimization problem. Section 3 provides the results of the model validation and compares the performance of the different optimization algorithms with each other and with standard approaches. Section 4 discusses the results of the study and their implications for the validation of physiological simulation models for clinical applications. Finally, a more detailed description of the different optimization methods used to perform the analysis is provided in §5.

2. SIMULATION MODEL AND PROBLEM FORMULATION

2.1. Simulation model

The simulation model considered in this study is a new and extended MATLAB implementation of several original, physiological models originally developed as the Nottingham Physiology Simulator (NPS) [12]. The core models in the simulator are designed to represent a dynamic *in vivo* cardiopulmonary state using a mass-conserving, arithmetic set of equations based on well-established physiological principles. Designing the simulator as a set of iterating routines allows the accurate representation and observation of gradual changes in several parameters that are otherwise hard to estimate *in vivo*. The simulation of the respiratory process refines the physiological data iteratively through a sequence of established theoretical equations to reproduce complex physiological scenarios [13,14]. The lungs are modelled as a dynamical system comprising external equipment (e.g. a mechanical ventilator), anatomical and alveolar deadspaces and ventilated, perfused alveoli. In this study, 100 individual alveolar compartments have been incorporated into the model.

The generic configuration of the model is given in table 1. The inspired air consists of oxygen (19.6%), nitrogen (74%) and carbon dioxide (0.1%). The balance is made up of water vapour (6.3%) present due to the humidification of air, which is assumed to be at 37°C. The model obeys the ideal gas laws and incorporates

Table 1. Model configuration.

weight	70 kg
inspired gas	warmed and humidified
inspired flow pattern	constant flow
fraction of inhaled O ₂ , FiO ₂	0.196
tidal volume	500 ml
respiratory rate	12 bpm
inspiratory to expiratory ratio	1 : 2

the resultant effect on flow of gas with the temperature difference between inspired gas and core body temperature. In the model, it is assumed that a complete mixing of gases occurs in the alveoli. Simulating movement of oxygen, carbon dioxide, nitrogen and an anaesthetic gas causes changes in the lung volume and resulting gas concentrations. Blood flow through the lung is modelled as shunted or non-shunted blood. Blood flow across the alveoli is ‘time-sliced’ such that individual packets of blood are considered. Each packet comes to equilibrium with alveolar gases using an iterative process that includes movement of gas between alveolar and capillary compartments, resulting in a highly accurate simulation of the equilibration process. Calculation of blood oxygen and carbon dioxide content after equilibration uses standard formulae and includes the effects of BE, temperature and plasma pH in the internal blood-gas calculations. The plasma pH depends on the BE, temperature and plasma carbon dioxide content. Humidification and temperature effects on the inspired dry air concentrations are also incorporated. Barometric pressure, or the atmospheric pressure, is taken to be at 101.3 kPa. Peripheral metabolism involves simple production of carbon dioxide and extraction of oxygen, using oxygen consumption (VO₂), respiratory quotient (RQ) and cardiac output (CO).

Each alveolar compartment can be attributed a specific pulmonary vascular resistance, bronchiolar resistance and compliance in order to create a desired ventilation–perfusion distribution. Bronchiolar resistance variations as a result of expansion of the lungs are included. The bronchiolar flow can be laminar or turbulent and a common resistance to flow is included. The type of inspiratory flow during mechanical ventilation of the lungs could be preset to constant flow or constant pressure; constant flow is considered in this study. Small airway closure and the resultant recruitment pressure variation are included in the model using a simple algorithm that allows the bronchioles feeding each alveolus to collapse when the alveolus volume approaches zero (less than 1% of its volume at full respiratory capacity (FRC)); this simulates the real-world phenomenon of absorption atelectasis. Reopening is through attainment of a threshold-opening pressure, set at 10 cm H₂O above atmospheric. Note, however, that in the present investigation where a healthy lung mode is considered, the model did not suffer any significant collapse-reopening behaviour during the simulations. Figure 1 shows the ventilation–perfusion distribution created by the chosen

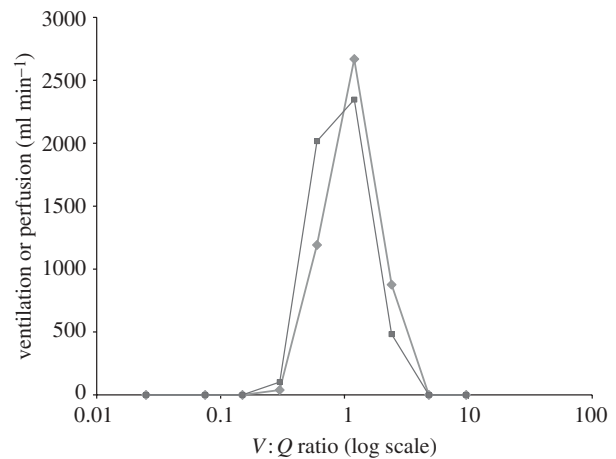


Figure 1. Ventilation–perfusion distribution in the lung model. Filled square, ventilation nominal case; filled diamond, perfusion nominal case.

lung settings in the NPS (for 100 alveolar compartments). The figure clearly shows the inhomogeneity of the lung representation considered here, which closely matches the data from [15], i.e. 95 per cent of ventilation and perfusion in a young male is confined to within the $V:Q$ ratio of 0.3 and 2.1. For further details of the simulation model and its calibration, the reader is referred to [2,3,12].

2.2. Modelling of uncertain parameters

Clearly, the number of uncertain parameters that could be considered in a complete analysis of a pulmonary system model is vast and includes both static and time-varying parameters arising from uncertainty within and between individual subjects. To illustrate our proposed approach, we have chosen a representative subset of these parameters, representing the haemoglobin (Hb) level, CO, VO₂, RQ and the core body temperature (T). These parameters are assigned a ‘nominal’ value together with an allowable range of variation/uncertainty, as shown in table 2. These values were chosen in consultation with medical specialists and are based on direct clinical experience of the levels of uncertainty that would be expected among the general patient population.

2.3. Formulation of the model validation problem as an optimization problem

The model validation problem considered in this study is as follows: for the realistic levels of uncertainty on key model parameters defined in table 2, guarantee that the model predictions for P'_{O_2} (current arterial O₂ pressure, kPa), P'_{CO_2} (current arterial CO₂ pressure, kPa) and pH' (current plasma pH) always stay within physiologically reasonable ranges (i.e. verify that the model’s responses exhibit the same levels of robustness to this level of uncertainty as are observed *in vivo*). Valid ranges for P'_{O_2} , P'_{CO_2} and pH' were again defined based on clinical experience and are shown in table 3.

Table 2. Nominal values, allowable uncertainty ranges and the resultant lower and upper bounds for the selected uncertain parameters in the model.

uncertain parameters	nominal value and uncertainty range	[lower bound (\underline{x}), upper bound (\bar{x})]	description, units
Hb	$140 \pm 15\%$	[119, 161]	Hb content in blood (g l^{-1})
CO	$5000 \pm 10\%$	[4500, 5500]	CO (ml min^{-1})
VO ₂	$250 \pm 20\%$	[200, 300]	VO ₂ (ml min^{-1})
RQ	$0.8 \pm 12\%$	[0.704, 0.896]	RQ (non-dimensional)
T	$37.2 \pm 0.5\%$	[37.0, 37.4]	temperature ($^{\circ}\text{C}$)

Table 3. Allowable upper and lower bounds for model predictions.

predicted parameters	[lower bound, upper bound]	description (units)
P'_{O_2}	[9, 15]	current arterial O ₂ pressure (kPa)
P'_{CO_2}	[3.5, 7.5]	current arterial CO ₂ pressure (kPa)
pH'	[7.3, 7.5]	current plasma pH

The model validation problem defined above can be formulated mathematically as an optimization problem as follows:

$$\begin{aligned} \max_x f(x) \quad \text{where } f(x) &= \|\Delta_{\text{O}_2}\|_{\infty} + \|\Delta_{\text{CO}_2}\|_{\infty} + \|\Delta_{\text{pH}}\|_{\infty} \\ \text{subject to } \underline{x} &\leq x \leq \bar{x} \end{aligned} \quad (2.1)$$

$$\begin{aligned} \Delta_{\text{O}_2} &= \left(\frac{P_{\text{O}_2} - P'_{\text{O}_2}}{P_{\text{O}_2}} \right), \Delta_{\text{CO}_2} = \left(\frac{P_{\text{CO}_2} - P'_{\text{CO}_2}}{P_{\text{CO}_2}} \right), \\ \Delta_{\text{pH}} &= \left(\frac{\text{pH} - \text{pH}'}{\text{pH}} \right). \end{aligned}$$

In the above, x is the vector of uncertain model parameters, \underline{x} and \bar{x} are vectors of lower and upper bounds, respectively, for the uncertain parameters (table 2) and $f(x)$ is the *objective function* which is maximized by the optimization algorithm. In problem (2.1), P_{O_2} , P_{CO_2} and pH refer to the nominal values of the model predictions (i.e. the values that are computed when all the uncertain parameters are set to their nominal values, over a given simulation time period). In the objective function, $\|a\|_{\infty}$ refers to the magnitude of the largest component in the vector a (i.e. $\max_{1 < i < n} |a_i|$), commonly known as the infinity norm. Thus, $\|\Delta\|_{\infty}$ measures the normalized absolute magnitude of the largest difference between the nominal model prediction and the current model prediction over the course of the simulation time window, normalized using the nominal output. In this study a simulation time window of 25 min was chosen, in order to ensure that all model outputs had settled down to steady-state values

before the objective function was evaluated. In this case, we have chosen to normalize each of the three output variables to ensure that they have an equal weighting in the objective function; however, it is clearly possible to differentially weigh the different variables in situations where deviations in one are considered more problematic or interesting than another. With the above formulation, the optimization algorithm searches for the combination of uncertain parameters within the allowable ranges that drives the predictions of the simulator as far as possible from the nominal predicted values over the course of the simulation—a schematic diagram of the model validation framework is shown in figure 2.

In order to ensure that we find the globally optimal solution to the above problem, we employ several global optimization algorithms, each of which is based on very different mathematical search principles. Once the optimal solution has been found, we check whether the corresponding ‘worst-case’ model predictions are still within the allowable ranges specified in table 3—if they are, then the model has passed that particular test. If one or more anomalous responses are found, then these need to be checked to see whether the particular combination of model parameters which produced that response is physiologically plausible, whether it corresponds to a particular condition which would be likely to produce unexpected responses etc. Whatever the results, the point of trying to ‘break’ the model is to uncover unexpected conditions or responses, which need further investigation and may (or may not) point to limitations or weaknesses in the formulation of the model. We are not trying to replace an expert analysis of the model by an automated computer program, rather we propose an analysis method which can efficiently uncover anomalous model responses for further investigation by the user. Furthermore, a comparison of the nominal versus worst-case values of the model outputs allows the impact of uncertainty on model predictions to be quantified, thus providing a valuable insight into the reliability and limitations of any new therapeutic strategies arising from the model predictions.

In this study, a local gradient-based optimization algorithm, sequential quadratic programming (SQP) [16] and two different global methods, namely genetic algorithms (GAs) [17] and mesh-adaptive direct search (MADS) [18] were used to solve the problem defined in problem (2.1). For the purposes of comparison, a standard statistical method (Monte Carlo simulation [8,19]) was also implemented. All algorithms, along with the simulation model, were coded in the MATLAB high-level modelling and computation environment [20,21].

3. RESULTS

Figures 3–5 show the results of the application of the different optimization algorithms to the problem defined in problem (2.1). In each figure, the graph on the left shows the improvement in objective function value as the algorithm proceeds (i.e. the rate of

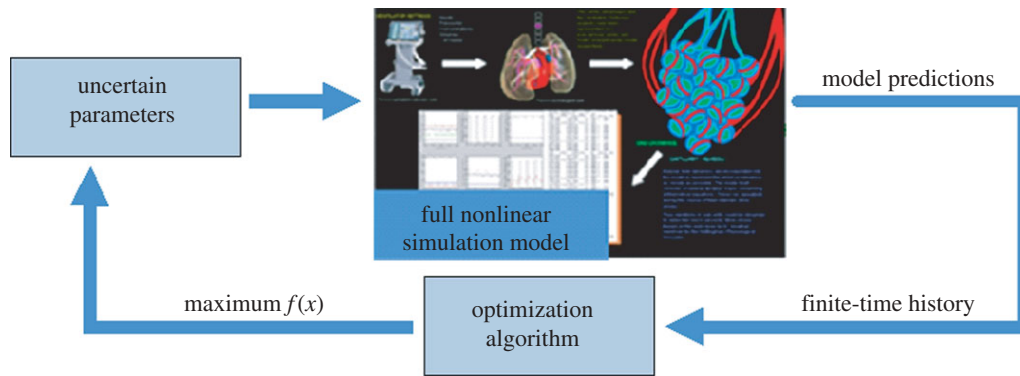


Figure 2. Optimization-based model validation framework.

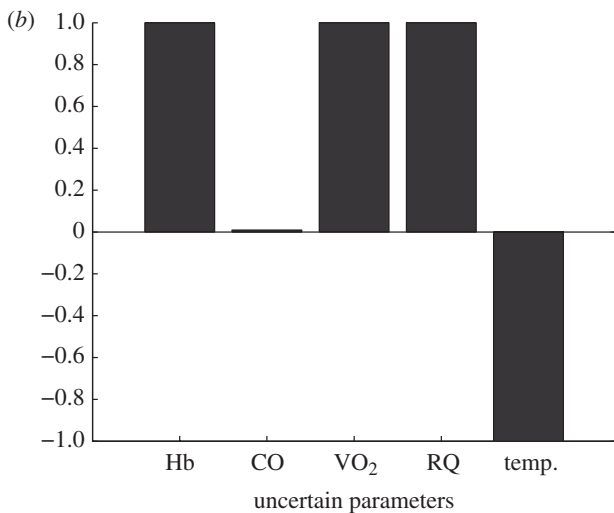
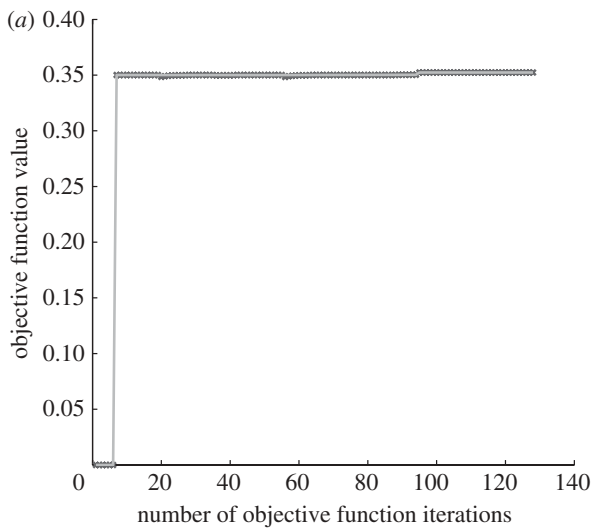


Figure 3. SQP results. (a) SQP performance and (b) optimal values of uncertain parameters.

convergence to the optimal solution). The bar graph on the right shows the uncertain parameter combination corresponding to the optimal solution found by the algorithm. Note that in the plot each uncertain parameter has been normalized between -1 and 1 (lower bound and upper bound, respectively; table 2), with a value of zero corresponding to its nominal value. The bar graphs show the optimal values of each of the five

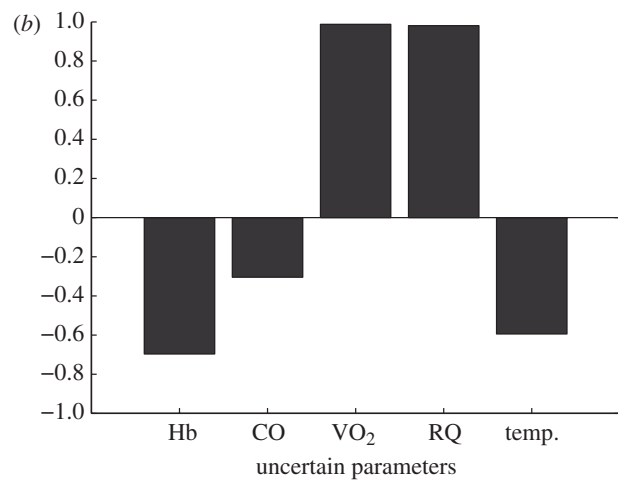
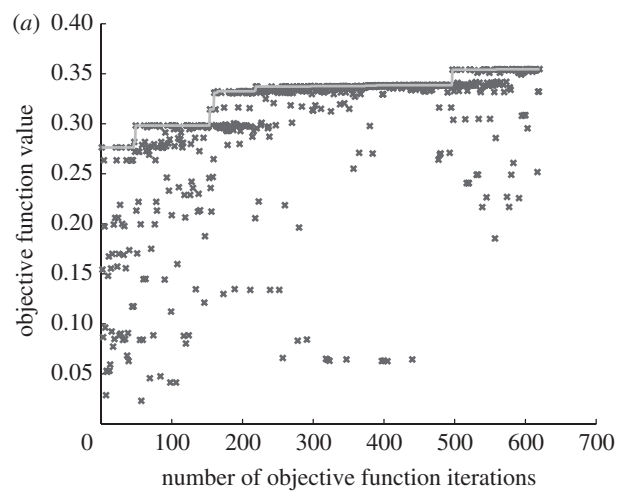


Figure 4. GA results. (a) GA performance and (b) optimal values of uncertain parameters.

uncertain parameters in the order given in table 2 (i.e. Hb, CO, VO₂, RQ and temperature). Figure 3 displays the progression and termination of the local SQP algorithm. The algorithm terminates when the magnitude of the function value improves by less than a preset tolerance, chosen as 10^{-6} . Local optimization algorithms such as SQP use gradient information to converge quickly to an optimal value in the uncertain parameter space. The surface generated by the objective function values over the multidimensional space defined by the uncertain parameters is likely to be non-convex, due

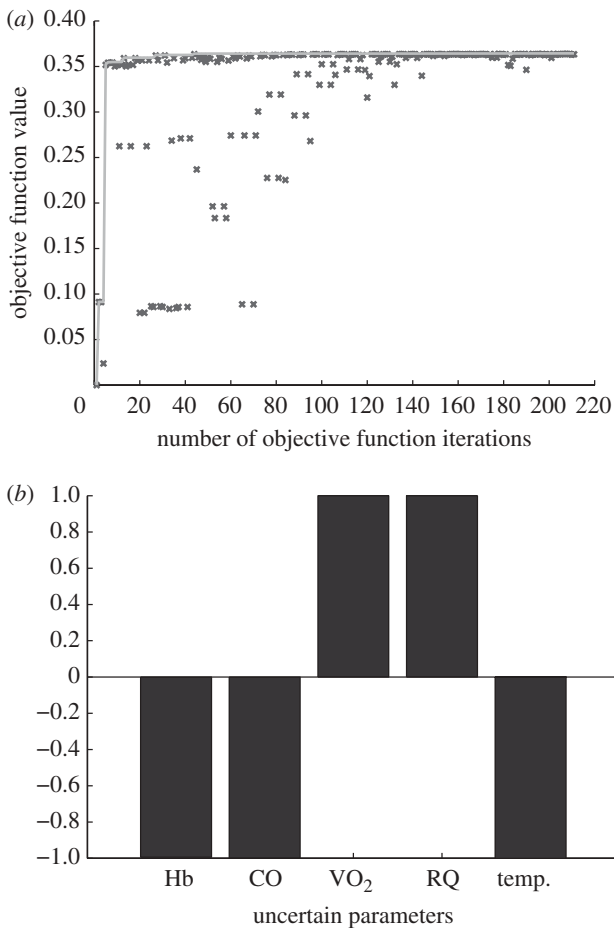


Figure 5. MADS results. (a) MADS performance and (b) optimal values of uncertain parameters.

to the highly nonlinear simulation model of the type considered here. As a result, local algorithms are likely to become trapped at local optima and will not be able to uncover the maximum possible deviation of the model responses from their nominal values. As will be shown below, this is exactly what happens in this case. Figure 3*b* gives the optimal parameter combination found by the algorithm.

Global optimization algorithms use randomization, evolutionary principles and/or heuristic search strategies to escape from local solutions and are generally accepted to have a high probability of finding globally optimal solutions, although at the cost of significantly increased computational overheads. Figure 4 shows the results of the application of the global GA optimization method to the same problem. The GA starts with an initial randomly chosen population of candidate solutions. Based on the past experience and some initial exploratory studies, we chose a population size of 20 and the algorithm terminates when the magnitude of the objective function value improves by less than the preset tolerance, again chosen as 10^{-6} . Figure 4*b* gives the final optimal parameter combination found by the algorithm. Note that the global algorithm finds a completely different worst-case parameter combination from the optimal solution found by the local algorithm, with an improved maximum value of the objective function, although with a corresponding increase in the number of simulations required.

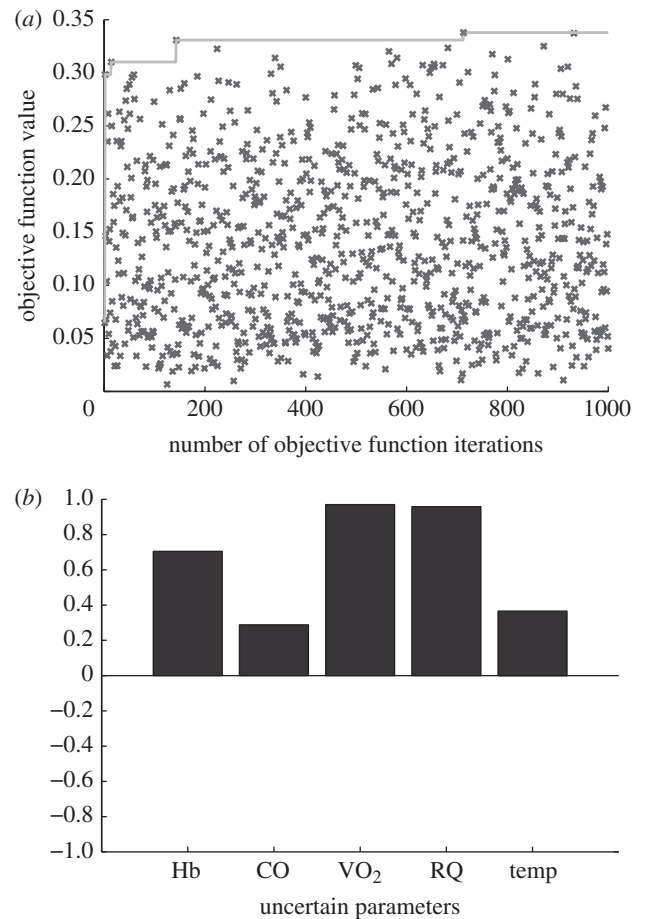


Figure 6. Monte Carlo results: x denotes the best value found over 1000 function evaluations. (a) Monte Carlo performance and (b) optimal values of uncertain parameters.

Figure 5 shows the result of the application of a deterministic global optimization algorithm known as MADS (or pattern search), which found the maximum value of the cost function among all the methods considered in this study. The function tolerance and the mesh size for the algorithm were both set to 10^{-6} . Figure 5*b* gives the final optimal parameter combination computed by the algorithm. It can be observed that this algorithm, which is based on search principles completely different from evolutionary algorithms, has varied each of the uncertain parameters in the same direction as the GA, but has taken each of the parameters to their boundaries in order to arrive at the global solution. This reveals one of the drawbacks of evolutionary algorithms, which, due to their randomized nature and lack of gradient information, often arrives in the vicinity of the global solution but then have difficulty in converging precisely to it. Interestingly, the MADS algorithm also proved to be significantly more efficient than the GA, requiring about one-third as many simulations.

For the purposes of comparison, figure 6 gives the results of a standard Monte Carlo simulation campaign using 1000 simulations. Although this number of simulations is more than three times the number required by the MADS algorithm and provides a strong statistical confidence that the true worst-case behaviour will be found (see table 6 in §4), it is apparent from figure 6

Table 4. Comparison of the results of the different optimization algorithms.

optimization algorithm	optimal parameters [Hb, CO, VO ₂ , RQ, T]	number of simulations	optimal value of the objective function
SQP	[161, 5004, 300, 0.896, 37.01]	128	0.3523
GA	[125, 4848, 299, 0.894, 37.09]	620	0.3547
MADS	[119, 4501, 300, 0.896, 37.01]	211	0.3637
MCS	[155, 5144, 299, 0.892, 37.27]	1000	0.3378

that only a local optimum for problem (2.1) has in fact been identified, with the corresponding values of the uncertain parameters being significantly different from their true worst cases.

Table 4 provides a comparison of the worst-case parameter combinations and objective function values found by the different methods considered in this study. The table also compares the computational efficiency of each approach in terms of the number of simulations required to converge to the optimal solution. Note the excellent performance of the MADS algorithm, which achieved the highest value of the cost function but required only slightly more simulations than the local SQP algorithm.

Figure 7*a–d* shows the main result of our model validation study—for the worst-case deviations from the nominal model predictions due to parametric uncertainty in the model the responses of the model always stay within the physiologically realistic bounds defined in table 3. The shaded cubical region in figure 7*a* represents the allowable ranges of each output as given in table 3. It is clear that the worst-case model predictions never violate the boundaries of the region over the course of the simulation. Figure 7*b* and *c* displays the worst-case result as given by each of the different optimization algorithms.

To obtain an understanding of the relative effects of the different uncertain parameters on the objective function, a sensitivity analysis was performed around the nominal point (figure 8*a*) as well as at the worst-case point (figure 8*b*) in the uncertain parameter space. As shown in the figures, VO₂ has the strongest effect on the objective function relative to the other parameters. Reassuringly, this is to be expected from physiological considerations as, of all the parameters, VO₂ will be the one which has the most direct effect on the quantity of oxygen and carbon dioxide in the blood. VO₂ is the oxygen taken away from the blood (reduction in P_{O_2}), which would result in higher carbon dioxide production at the peripheral tissues (higher P_{CO_2}). Blood–acid relationships suggest that this increase would also result in a significant decrease in the levels of pH in the blood. All these factors indicate that variations in this parameter would result in significant deviations from the nominal P_{O_2} , P_{CO_2} and pH values, i.e. it will have a large influence on the value of the objective function.

Figure 9 shows the ventilation–perfusion (VQ) distribution of the lung corresponding to the worst-case model configuration in comparison to the nominal distribution in figure 1. It can be observed that the differences in the values of the uncertain parameters

such as CO (CO has decreased from 5000 ml min⁻¹ in the nominal configuration to 4504 ml min⁻¹ in the worst-case configuration) have resulted in a change in the VQ distribution, as shown in the figure. However, this difference is relatively small, and the resulting distribution is still representative of a typical VQ distribution for a normal lung as suggested by Wagner *et al.* [15]. This provides further confidence in the ability of the model to accurately represent the dynamics of a healthy lung over the range of uncertain parameters considered in the study.

The optimal uncertain parameter combination uncovered by our analysis suggests that the worst-case physiological response of the model corresponds to anaemic conditions (low Hb and low CO) alongside high VO₂, which could be considered a ‘reasonable’ prediction for the worst case of a pulmonary simulator based on physiological intuition. However, the respiratory system has an intrinsic response to low oxygen levels in blood which is to restrict the blood flow in the pulmonary blood vessels, known as hypoxic pulmonary vasoconstriction (HPV). To investigate the effect of HPV, a simple function, resembling the stimulus response curve suggested by Marshall *et al.* [22], was incorporated into the simulator to gradually constrict the blood vessels as a response to low alveolar oxygen tension. Using the best performing optimization technique from the previous section (MADS), the new worst case was then found to be as given in table 5 and figures 10 and 11. As table 5 shows, once the HPV function is introduced, a contradictory worst-case uncertain parameter combination is computed which is different from the ‘intuitive’ predictions found in the previous analysis. Figure 11*a* shows that including HPV causes a lower deviation from nominal for O₂ in the worst case and thus explains the reduced value of the objective function found by the optimization algorithm (where the worst case would be the maximum deviation from nominal). It is interesting that the major difference in the worst-case parameters seem to be the increased values of Hb and CO, parameters which typically when increased would be expected to show a relatively small but direct incline in the arterial O₂, causing in this case the deviation from the nominal to be reduced. However, as shown in figure 11*b*, the HPV added to the model attempts to correct for this hypoxemia in blood, resulting instead in an increase in arterial P_{O_2} for the particular parameter configuration which would otherwise have reduced the arterial P_{O_2} . This clearly demonstrates the benefits of the proposed validation framework in uncovering extreme physiological simulator responses earlier in the developing stage.

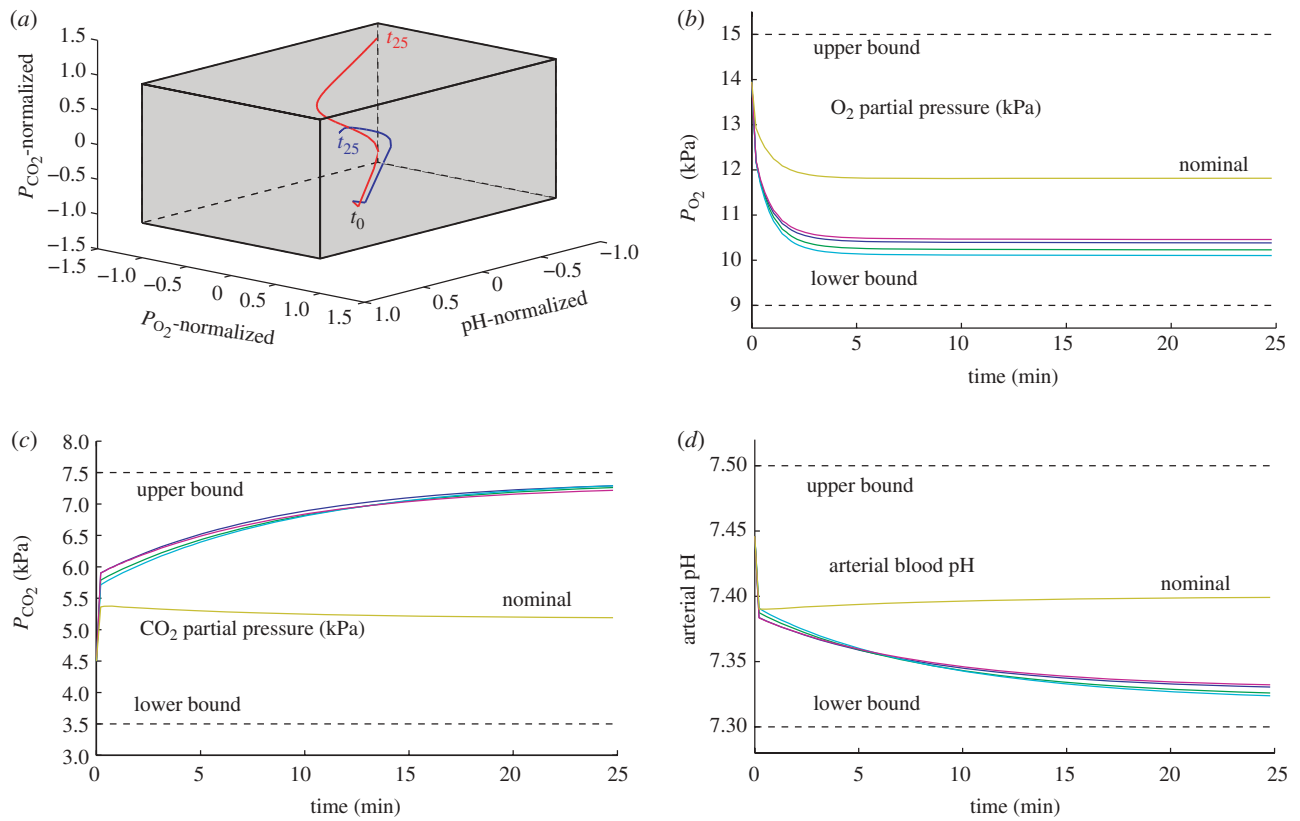


Figure 7. (a) Nominal (blue) and worst-case (red) model predictions for normalized P_{O_2} , P_{CO_2} and pH. t_0 , simulation starting time; t_{25} , simulation end time, *box*, allowable values. (b) Nominal and worst-case model predictions for P_{O_2} . (c) Nominal and worst-case model predictions for P_{CO_2} . (d) Nominal and worst-case model predictions for pH. (b–d) Dark blue, SQP; green, GA, light blue, PS; pink, MCS.

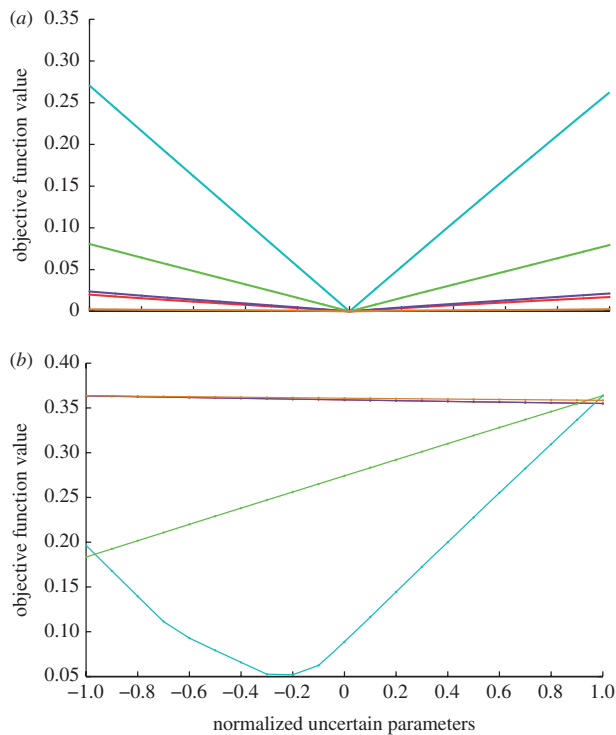


Figure 8. Sensitivity analyses at the (a) nominal case and (b) worst case. Red, Hb; violet, CO; light blue, VO_2 ; light green, RQ; orange, temp.

4. DISCUSSION AND CONCLUSIONS

Monte Carlo simulation is generally considered the ‘gold standard’ in many different areas of science and technology for estimating the effects of parameter variations on the outputs of complex dynamical systems. However, Monte Carlo simulation suffers from an exponential increase in computation times as the required confidence levels for the statistical analysis are increased. As shown in table 6 (taken from [23]), to be within a probability range of 0.75–0.95 of estimating the worst-case behaviour of a system to within ± 20 per cent, only 25 simulations are required. If, on the other hand, we need to be within a probability range of 0.94–0.999 of estimating the worst-case behaviour of a system to within ± 1 per cent, then 40 000 simulations are required. Since this number of simulations is usually totally impractical for complex physiological simulation models of the type considered here, much lower confidence levels are often accepted in practice. As shown by the example presented in this paper, Monte Carlo analysis can fail to find global solutions even when very large numbers of simulations are employed, thus seriously compromising the integrity of the overall model validation strategy. Interestingly, this issue has recently been widely recognized throughout the aerospace industry, which spends vast sums each year on

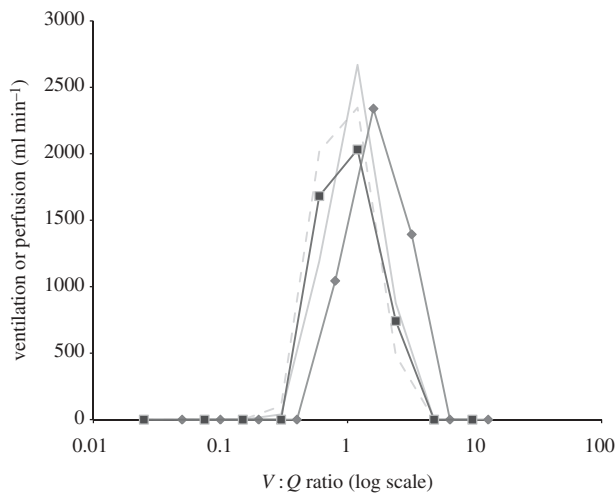


Figure 9. Ventilation (V)–perfusion (Q) distribution of the lung for the worst case. Solid grey line, ventilation nominal case; dashed line, perfusion nominal case; line with diamond, ventilation worst case; line with square, perfusion worst case.

Table 5. Comparison of the effect of HPV on the model.

	worst-case parameters	optimal value of the objective function
without HPV	[119, 4501, 300, 0.896, 37.01]	0.3637
with HPV	[161, 5500, 300, 0.896, 37.03]	0.3159

Monte Carlo simulation campaigns to validate new safety-critical flight systems. This has led to a drive to investigate alternative validation strategies based on the use of global optimization algorithms to intelligently search for worst-case behaviour, and promising results in this direction have recently appeared in the research literature [8–11]. Similar initiatives are also underway in the field of molecular systems biology, where robustness analysis of complex biochemical network models using global optimization methods has resulted in many new insights and improvements in the modelling of these complex systems [7].

Simulation models which are intended for direct application in healthcare research and clinical practice face extremely stringent requirements; they must cope with high levels of uncertainty due to non-measurable parameters, population heterogeneity and disease heterogeneity, and simultaneously, their validation procedures must offer watertight proof of their applicability and reliability in the clinical arena. Clearly, such expectations have not to date been met by the vast majority of simulators developed in academic studies. Typically, adequate statistical results when model credibility is tested on a single idealized subject have been substituted for proper population-based model validation. Modern techniques from systems and control engineering such as those considered here offer new hope that simulators can be developed which acceptably represent true levels of population and disease heterogeneity, and hence provide robust and

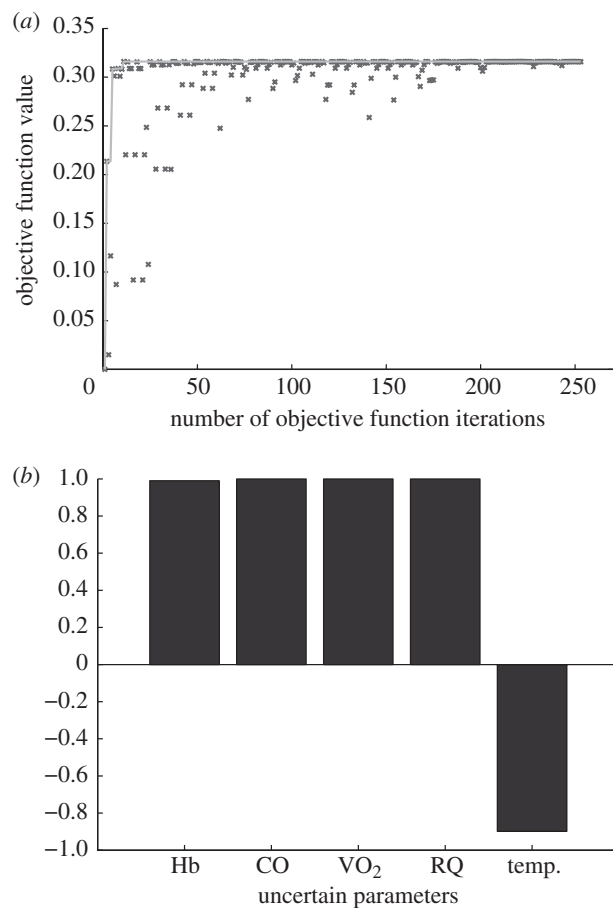


Figure 10. Worst-case optimization result with the inclusion of HPV (a) MADS performance and (b) worst-parameter combination.

reliable predictions upon which novel therapeutic strategies may be based. As demonstrated by the results presented in this paper, the time is now ripe for such optimization-based approaches to play an important role in transforming physiological simulators into powerful biomedical engineering tools for a direct application in clinical practice.

5. OPTIMIZATION METHODS

In this section, we give a brief description of the different optimization algorithms that were used to generate the results presented in §3.

5.1. Sequential quadratic programming

Local gradient-based optimization techniques provide fast convergence rates to a local optimum. When used on a problem with many local optima, the quality of the resulting solution then completely depends on the proximity of the initial ‘guess’ for the values of uncertain parameters to the global optimum. Since, for this problem, no *a priori* information is available on which combination of parameter values is most likely to generate worst-case behaviour in the simulation model, the uncertain parameters are simply set to their nominal values and the optimization algorithm proceeds from this point. The particular algorithm used in this study

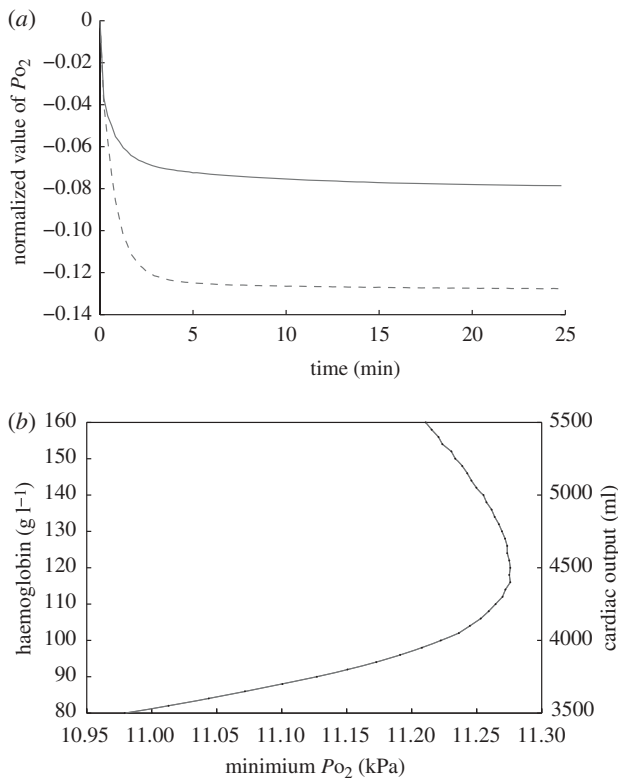


Figure 11. (a) Worst-case P_{O_2} for model with HPV (solid line) and without HPV (dashed line). (b) Arterial P_{O_2} characteristic with decreasing Hb and CO with HPV.

implements a method based on SQP [16]. A medium-scale optimization scheme is chosen, where the algorithm solves a quadratic programming subproblem at each iteration—see [20] for full details.

5.2. Genetic algorithms

GAs are general purpose stochastic search and optimization procedures based on evolutionary principles [24]. Evolutionary algorithms aim to generate a population of fittest candidates by implementing mathematically a range of evolutionary concepts such as selection, mutation, recombination etc. Owing to their ease of application in problems with large and small parameter search spaces, GAs have become a popular search and optimization technique in many fields of science and engineering [9,17]. Owing to their stochastic nature, GAs are generally capable of converging to the global optimum even in highly non-convex parameter spaces—convergence can be slow, however, and global algorithms such as GAs typically require much longer computation times than local gradient-based methods.

In GAs, a randomly selected population of candidates (first generation) undergoes a repetitive process of reproduction, where selection is based on the value of the objective function (also known as the fitness function). Every generation, the best candidates from the previous generation and candidates obtained through mutation, crossover and recombination form a new population. The algorithm can terminate if (i) the allowed number of iterations are exceeded, (ii) the objective function improves by less than a preset

Table 6. Number of Monte Carlo trials required to achieve a desired estimation uncertainty with known probability [23].

uncertainty probability range	number of Monte Carlo trials based on per cent of estimation uncertainty (%)				
	20	15	10	5	1
0–0.683	7	12	25	100	2500
0.750–0.954	25	45	100	400	10 000
0.890–0.997	57	100	225	900	22 500
0.940–0.999	100	178	400	1600	40 000

Table 7. Pseudo-code for the GA.

- (1) create a random initial population (first generation)
- (2) repeat (until terminated)
 - evaluate each individual's fitness
 - select individuals to form parents
 - create children for the next generation by using:
 - elite children: individuals with best-fitness values
 - crossover children: created by combining the vectors of a pair of parents
 - mutation: created by randomly changing the genes of a single parent
 - check for termination criteria:
 - amount of time elapsed
 - fitness limit
 - function tolerance
- (5) go back to (2)

tolerance, (iii) limits relative to time or fitness are exceeded, or (iv) any other user-defined criterion is violated. For this study, the population size is fixed at 20 with the number of generations allowed set to 30. A roulette wheel selection scheme is used in which the probability of a parent candidate being chosen for reproduction is proportional to its fitness. A simple scattered crossover scheme is selected which decides the manner in which two individuals are selected to form a new candidate. A Gaussian mutation scheme is employed, which makes small random changes in the individuals based on a Gaussian curve to ensure adequate coverage of the entire parameter space. From the present generation, the best two move into the next generation (elite count). The rest of the population is derived from crossover (80% of the new population, crossover fraction = 0.8) and mutation. Table 7 provides a pseudo-code which summarizes the implemented algorithm.

5.3. Mesh-adaptive direct search

MADS (also known as pattern search) is another global optimization algorithm that does not require derivatives of the objective function to determine a solution. The pattern search algorithm searches for the best point around a current point based on the evaluation of the objective function. At each step, a pattern search algorithm searches a set of points called a mesh, formed by adding the current point to a scalar multiple (mesh size) of a set of vectors called a pattern. If the algorithm finds a point in the mesh that improves

Table 8. Pseudo-code for the MADS algorithm.

MADS algorithm with $2N$ poll method, where N = number of variables. If $N = 2$:

- (1) select an initial point $A = (a_1, a_2)$ in the search space
- (2) set initial value for mesh size, M , and a pattern vector, V_{2N} , which is randomly chosen as follows: $V_1 = [x \ 0]$, $V_2 = [0 \ y]$, $V_3 = [-x \ 0]$ and $V_4 = [0 \ -y]$. x and y are random vectors limited by the upper and lower bounds of the search space
- (3) obtain the mesh using: $\bar{A} = a + M \cdot V$
- (4) evaluate the points in the mesh \bar{A} using the objective function. In the case of a successful poll (a better objective function value has been found), set this point as A , double the mesh size and go back to step 2
- (5) in the case of an unsuccessful poll, reduce the mesh size and go back to step 2
- (6) stop if termination criteria satisfied

the value of the objective function (a successful *poll*), then the new point becomes the current point and the mesh size increases. In the current study, the set of vectors in a pattern are chosen at random. If the current point gives a better value for the objective function than the other vectors in the mesh (an unsuccessful *poll*), then the mesh size is reduced and the method is repeated for the new points. Table 8 gives the pseudo-code for the algorithm; for further details, the reader is referred to [18,21].

5.4. Monte Carlo simulation

MCS refers to the repeated evaluation and statistical analysis of a given objective function for randomly chosen values of the uncertain parameters. At its simplest, no prior knowledge of the system is required. Assuming uniform distribution in the parameter search space, random parameter combinations, within the specified bounds, are generated using the random number generator available in MATLAB and repeated simulations are carried out for all these combinations. Statistical information such as sample averages and variances are then computed to assess the performance of the system. The simplicity of this approach has led to it being widely used in industrial applications. A serious disadvantage of MCS is the exponential increase in computational effort with the desired statistical confidence level of the results [8,23]. Also, recent studies have clearly demonstrated that MCS can fail to adequately approximate the worst-case behaviour of complex systems even when large numbers of simulations (providing high statistical confidence) are used [10,11].

The authors would like to thank the anonymous reviewers for their helpful suggestions which have significantly improved the paper. This research was carried out under EPSRC Research Grant EP/F057059/1.

REFERENCES

- 1 Hardman, J. G. & Ross, J. J. 2006 Modelling: a core technique in anaesthesia and critical care research. *Br. J. Anaesth.* **97**, 589–592.
- 2 Hardman, J., Bedforth, N., Ahmed, A., Mahajan, R. & Aitkenhead, A. 1998 A physiology simulator: validation of its respiratory components and its ability to predict the patient's response to changes in mechanical ventilation. *Br. J. Anaesth.* **81**, 327–332.
- 3 Mccahon, R. A., Columb, M. O., Mahajan, R. P. & Hardman, J. G. 2008 Validation and application of a high-fidelity, computational model of acute respiratory distress syndrome to the examination of the indices of oxygenation at constant lung-state. *Br. J. Anaesth.* **101**, 358–365. (doi:10.1093/bja/aen181)
- 4 Severinghaus, J. W. 1979 Simple, accurate equations for human blood O_2 dissociation computations. *J. Appl. Physiol.* **46**, 599–602.
- 5 Morohashi, M., Winn, A. E., Borisuk, M. T., Bolouri, H., Doyle, J. & Kitano, H. 2002 Robustness as a measure of plausibility in models of biochemical networks. *J. Theor. Biol.* **216**, 19–30. (doi:10.1006/jtbi.2002.2537)
- 6 Kim, J., Bates, D. G., Postlethwaite, I., Ma, L. & Iglesias, P. A. 2006 Robustness analysis of biochemical network models. *IET Syst. Biol.* **153**, 96–104.
- 7 Kim, J.-S., Valeyev, N. V., Postlethwaite, I., Heslop-Harrison, P., Cho, K.-H. & Bates, D. G. 2009 Analysis and extension of a biochemical network model using robust control theory. *Int. J. Robust Nonlinear Control* **20**, 1017–1026. (doi:10.1002/rnc.1528)
- 8 Fielding, C., Varga, A., Bennani, S. & Selier, M. 2002 *Advanced techniques for the clearance of flight control laws*. Berlin, Germany: Springer Verlag.
- 9 Menon, P. P., Bates, D. G. & Postlethwaite, I. 2006 Computation of worst-case pilot inputs for nonlinear flight control system analysis. *J. Guid. Control Dyn.* **29**, 195–199. (doi:10.2514/1.14687)
- 10 Menon, P. P., Bates, D. G. & Postlethwaite, I. 2007 Non-linear robustness analysis of flight control laws for highly augmented aircraft. *Control Eng. Pract.* **15**, 655–662. (doi:10.1016/j.conengprac.2006.11.015)
- 11 Menon, P. P., Postlethwaite, I., Bennani, S., Marcos, A. & Bates, D. G. 2009 Robustness analysis of a reusable launch vehicle flight control law. *Control Eng. Pract.* **17**, 751–765. (doi:10.1016/j.conengprac.2008.12.002)
- 12 Hardman, J. G. 2001 Respiratory physiological modelling—the design, construction, validation and application of a set of original respiratory physiological models. PhD thesis, Division of Anaesthesia and Intensive Care, University of Nottingham, UK.
- 13 Hardman, J. & Aitkenhead, A. 1999 Estimation of alveolar deadspace fraction using arterial and end-tidal CO_2 : a factor analysis using a physiological simulation. *Anaesth. Intensive Care* **27**, 452–458.
- 14 Hardman, J. G. & Bedforth, N. 1999 Estimating venous admixture using a physiological simulator. *Br. J. Anaesth.* **82**, 346–349.
- 15 Wagner, P. D., Laravuso, R. B., Uhl, R. R. & West, J. B. 1974 Continuous distributions of ventilation-perfusion ratios in normal subjects breathing air and 100% O_2 . *J. Clin. Invest.* **54**, 54–68. (doi:10.1172/JCI107750)

- 16 Fletcher, R. & Powell, M. 1963 A rapidly convergent descent method for minimization. *Comp. J.* **6**, 163.
- 17 Mitchell, M. 1998 *An introduction to genetic algorithms*. Cambridge, MA: MIT Press.
- 18 Audet, C. & Dennis Jr, J. 2007 Mesh adaptive direct search algorithms for constrained optimization. *SIAM J. Optim.* **17**, 188–217. (doi:10.1137/040603371)
- 19 Vidyasagar, M. 1998 Statistical learning theory and randomized algorithms for control. *IEEE Control Syst. Mag.* **18**, 69–85. (doi:10.1109/37.736014)
- 20 Mathworks. 2002 Optimization toolbox user's guide.
- 21 Mathworks. 2005 MATLAB genetic algorithm and direct search toolbox. See <http://www.mathworks.com/access/helpdesk/help/toolbox/gads/>.
- 22 Marshall, B. E., Clarke, W. R., Costarino, A. T., Chen, L., Miller, F. & Marshall, C. 1994 The dose–response relationship for hypoxic pulmonary vasoconstriction. *Resp. Physiol.* **96**, 231–247. (doi:10.1016/0034-5687(94)90129-5)
- 23 Williams, P. S. 2001 A Monte Carlo Dispersion Analysis of the X-33 Simulation Software. In *Proc. AIAA Conf. on Guidance, Navigation and Control, Montreal, Canada*.
- 24 Goldberg, D. E. 1989 *Genetic algorithms in search and optimization*. Reading, MA: Addison-wesley.