Modelling and Optimisation of Mechanical Ventilation for Critically Ill Patients

Submitted by **Anup Das**, to the University of Exeter as a thesis for the degree of
Doctor of Philosophy in Engineering,
March 2012.

This thesis is available for library use on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgment.

I certify that all material in this thesis which is not my own work has been identified and that no material has previously been submitted and approved for the award of a degree by this or any other University.

Signature: ----------------------------------

UNIVERSITY OF
EXETER

Abstract

This thesis is made up of three parts: i) the development of a comprehensive computational model of the pulmonary (patho)physiology of healthy and diseased lungs, ii) the application of a novel optimisation-based approach to validate this computational model, and iii) the use of this model to optimise mechanical ventilator settings for patients with diseased lungs.

The model described in this thesis is an extended implementation of the Nottingham Physiological Simulator (NPS) in MATLAB. An iterative multi-compartmental modelling approach is adopted, and modifications (based on physiological mechanisms) are proposed to characterise healthy as well as diseased states.

In the second part of the thesis, an optimisation-based approach is employed to validate the robustness of this model. The model is subjected to simultaneous variations in the values of multiple physiologically relevant uncertain parameters with respect to a set of specified performance criteria, based on expected levels of variation in arterial blood gas values found in the patient population. Performance criteria are evaluated using computer simulations. Local and global optimisation algorithms are employed to search for the worst-case parameter combination that could cause the model outputs to deviate from their expected range of operation, i.e. violate the specified model performance criteria. The optimisation-based analysis is proposed as a useful complement to current statistical model validation techniques, which are reliant on matching data from *in vitro* and *in vivo* studies.

The last section of the thesis considers the problem of optimising settings of mechanical ventilation in an Intensive Therapy Unit (ITU) for patients with diseased lungs. This is a challenging task for physicians who have to select appropriate mechanical ventilator settings to satisfy multiple, sometimes conflicting, objectives including i) maintaining adequate oxygenation, ii) maintaining adequate carbon dioxide clearance and iii) minimising the risks of ventilator associated lung injury (VALI). Currently, physicians are reliant on guidelines based on previous experience and recommendations from a very limited number of *in vivo* studies which, by their very nature, cannot form the basis of personalised, disease-specific treatment protocols. This thesis formulates the choice of ventilator settings as a constrained multi-

objective optimisation problem, which is solved using a hybrid optimisation algorithm and a validated physiological simulation model, to optimise the settings of mechanical ventilation for a healthy lung and for several pulmonary disease cases. The optimal settings are shown to satisfy the conflicting clinical objectives, to improve the ventilation perfusion matching within the lung, and, crucially, to be disease-specific.

Acknowledgements

First and foremost, I would like to thank Professor Declan Bates, whose vision and boundless enthusiasm continues to be the driving force behind our project and my work. I am thankful to Professor Jonathan Hardman from University of Nottingham, whose work and knowledge in the field of physiology and physiological modelling has been the cornerstone of this thesis. I would like to kindly thank Dr Prathyush Menon from the University of Exeter, for his ever-reliable and extensive guidance. All my supervisors have shown incredible patience and support towards me, for which I will be forever grateful.

I would like to thank my friends and family for their understanding and care. I would like to thank my colleagues and the staff at University of Exeter for the numerous instances they have provided me help.

Finally, I would like to dedicate this thesis to my father, mother and my sister, whose love and encouragement have been a source of inspiration throughout my life.

Table of Contents

List of Figures

List of Tables

Abbreviations and Nomenclature

BIR: Bronchial Inlet Resistance, airway resistance of assigned to each alveolar compartment.

C: content (concentration) [ml l^{-1}] e.g. Ca_{O2}: ml of O_2 per litre of arterial blood.

CO: cardiac output [ml min^{-1}]

CO_2: Carbon dioxide

FiO_2: Fraction of inspired air constituting of oxygen.

FRC: Functional residual Capacity [ml].

$FRC_{nominal}$: Nominal value of FRC used in the model = 3000ml.

Hb: Haemoglobin content in blood [gm l^{-1}]

HPV: Hypoxic Pulmonary Vasoconstriction

I:E : Inspiratory to Expiratory Ratio

N_2: Nitrogen

NPS: Nottingham Physiology Simulator

O_2: Oxygen

P: Partial pressure [kPa] e.g. P_{O2} – Partial pressure of oxygen, P_{alvO2}: Partial pressure of oxygen in alveolar compartment, P_{vO2}: Partial pressure of oxygen in venous compartment.

PEEP: Positive End Expiratory Pressure [cmH_2O]

PVR: Pulmonary Vascular Resistance, blood flow resistance assigned to each pulmonary capillary compartment.

Q: Perfusion distribution of the lung [ml]

RQ: Respiratory Quotient

S: Saturation (0-100%) e.g. S_{O2} – arterial saturation of haemoglobin with oxygen.

T: Temperature [°c]

TIR: the combined resistance of all alveolar compartments

TLC: Total Lung Capacity [ml].

TOP: Threshold Opening Pressure. Extra pressure within the collapsed alveolar compartment.

V/Q: Ventilation Perfusion Ratio

V: Ventilation distribution in the lung [ml]

VALI: Ventilator Associated Lung Injury

VentRate: Respiration rate set by the ventilator [breaths per minute, bpm]

V_{O2}: oxygen consumption [ml min^{-1}]

Vtidal: Tidal Volume [ml]

α, β: Resistance multiplier of size *n*, where *n* is the total number of diseased compartments in the model. α is used to create ventilation distribution and β is used to create perfusion distribution.

pH: pH level in blood.

v: volume of alveolar compartment, [ml]

t: time, most commonly used to represent the duration of a simulation. It can also describe the 'time slice', depicting the duration of real time of one model iteration.

Special subscripts-

i: inspired,   e: expired,   alv: alveolar, a: arterial,   v: venous

Chapter 1: Introduction

An explanation of the motivation for the research described in this thesis is provided in this chapter. Key concepts are introduced, pertaining to the modelling and optimisation of pulmonary (patho) physiology, which are discussed in more detail later in this thesis. Finally, the structure of the thesis, the focus of each chapter and the main contributions of the research are described.

1.1   Motivation

Mechanical ventilation of the lungs is a commonly-used, life-saving procedure. The majority of critically ill patients in Intensive Therapy Units (ITU) spend some time with their lungs ventilated by a mechanical ventilator. However, mechanical ventilation also exposes patients' lungs to potentially damaging pressures and flows. Studies have shown that 2.9% of patients receiving mechanical ventilation suffer from Ventilator Associated Lung Injury (VALI) (Anzueto et al 2004), necessitating prolonged stays in the ITU, and potentially causing pneumonia, lifelong pulmonary scarring, and even multiple organ failure (Dreyfuss 1992, Gammon 1992).

The main objectives of mechanical ventilation are to maximise gas exchange and minimise the risk of VALI (Slutsky 1993). The selection of appropriate settings of mechanical ventilators in an ITU is a significant challenge for the physician, due mainly to the heterogeneity of the pulmonary disease states, as well as the non-uniform distribution of ventilation in the lungs (Villar 2004, Gattinoni 2006). Currently, physicians rely on "arbitrarily fixed" guidelines

(DeProst 2011) based on their previous experience and recommendations from a limited number of *in vivo* studies (ARDS Network 2000).

The mortality rate (30%-40%) of mechanically ventilated patients suffering from VALI still remains high under current guidelines (Mercat 2008). Various studies (Keenan et al 2004, Morris 2002) have suggested that incorporation of patient specific treatment in the ITU could significantly reduce this high mortality rate and also improve patient care. However, the lack of available bedside data to clearly determine the state of patients, together with the complexity of the relevant internal physiology, which often has large variance, has hindered the incorporation of patient specific information into physician decisions (DeProst 2011).

Computational modelling has increasingly been proposed as a useful and reliable tool with which physicians could be able to gain a deeper understanding of the underlying physiology and make more informed decisions for patients in the ITU (McCahon 2008, Kathirgamanathan 2009). Even though modelling approaches can be restricted by the detail and complexity of the underlying physiology, computational modelling is not limited by patient compliance, ethical considerations and financial constraints, in contrast to the traditional *in vivo* clinical studies, (Hardman 2006). The use of modelling in this area of clinical medicine has thus become more common, with several successful research studies reported in the literature (see Weibel 1963, West 1969, Hinds 1984, Joyce 1996, Hahn 2003, Yem 2006 and references therein) that have had significant impact on the field of pulmonary physiology.

However, physiological models intended for the use in clinical environments can often face harsh expectations from physicians; the models must cope with significant levels of uncertainty arising from assumed values of parameters (Harrell 1996, Chase 2011), and population and disease heterogeneity, and their validation must provide convincing evidence of their applicability and reliability in the clinical arena (Hardman 2008). Current validation of

physiological models is typically limited to comparison of the performance of the model against the associated empirical data (Hardman 1998, Moppett 2008, Wilson 2009, Chase 2010), and this can potentially leave the model vulnerable to variability in parameters which are found within the patient population but are not adequately represented by the available empirical data.

By adapting and augmenting established systematic model validation approaches that are currently available in the area of systems and control engineering, additional levels of formalism and reliability can be brought to the problem of validating physiological models. The concept of 'robustness' is widely used in engineering contexts to denote the requirement that certain desired characteristics are maintained by a system despite fluctuations in the behaviour of its component elements or its environment (internal/external parameters and disturbances) (Doyle 2002). In systems and control engineering, many different analytical and computational approaches have been developed to assess the robust stability and performance of computer simulation models of safety-critical applications (Menon 2007, Martin 2008). In this thesis, we show how similar approaches can be utilised for validating computational models of pulmonary (patho)physiology – these validated models can then be further developed and employed in designing *in silico* experiments for exploring current practice for choosing the settings of mechanical ventilator for individual patients and disease states.

1.2   Thesis organisation

The thesis begins with a discussion of the necessary background in pulmonary (patho)physiology. Subsequently the development of a model of pulmonary (patho)physiology is presented. In the second part of the thesis, an optimisation approach is employed to validate the robustness of this model. Validated pulmonary physiology models are used to represent various diseased lung states *in silico*, and an optimisation-based methodology is used to

determine patient specific optimal mechanical ventilator settings for these *in silico* disease states. The thesis has five chapters, which are outlined below:

Chapter 2 introduces relevant background on the relevant physiology for the model considered in this thesis. The structure of the lungs is discussed, explanations of the important physiological terms are provided, and the most important mathematical equations describing arterial blood gas calculations are also given. The proposed extended implementation of the Nottingham Physiological Simulator (NPS) (Hardman 1998), which consists of different modules including the respiratory system, the cardiovascular system, the renal system, the acid base mechanism and a mechanical ventilator, is described in detail. Modifications (based on physiological mechanisms) are proposed to characterise healthy as well as diseased states. These modifications address two distinct faults within the previous version of the model: 1) the inability of the model to accurately represent *collapse* within the lungs under disease states and 2) the incomplete response of the model to positive end expiratory pressure (PEEP). These faults were identified during the model development process and the solutions which were developed to address them are described in detail.

Chapter 3 begins with a discussion of current practices followed in physiological model validation and the need for improved techniques. Optimisation algorithms such as Sequential Quadratic Programming (SQP), Mesh Adaptive Direct Search (MADS), Genetic Algorithms (GA) and Differential Evolution (DE) are described. The general underlying principles of optimisation-based model validation, and previous efforts to employ optimisation algorithms to similar problems, are also described. In the process of model validation, the pulmonary physiology model is subjected to simultaneous variations in the values of multiple physiologically relevant uncertain parameters with respect to a set of specified performance

criteria. These performance criteria are based on expected levels of variation in arterial blood gas values found in the patient population and they are evaluated using numerical simulations.

Chapter 4 begins with a brief introduction to modern mechanical ventilators and the associated problem of VALI. Current ITU strategies for deciding on the values of key mechanical ventilator settings are presented. The problem of choosing optimal mechanical ventilator settings for diseased lungs is formulated as a constrained multi-objective optimisation problem. Two multi-objective optimisation approaches are presented – a) an approach using an aggregated objective function with user defined weights (WAOF) and a hybrid optimisation algorithm and b) a Pareto dominance-based approach using a nondominated sorting genetic algorithm (NSGA-II). The NSGA-II result is used to illustrate the conflicting objectives of the optimisation problem considered here. The relative merits of the two approaches are compared using an example *in silico* disease case. A range of different pulmonary disorders are then created *in silico* by manipulating the ventilation and perfusion distribution within the validated computational model. Resistances to airflow and blood movement are introduced which simulate a wide spectrum of ventilation and perfusion abnormalities representative of pulmonary disease scenarios in the published literature. The results of applying the proposed optimisation-based approach for choosing ventilator settings for these different disease states are discussed.

Finally in chapter 5, the main conclusions of the work in this thesis are presented and possible areas for future research are proposed.

1.3   Main contributions of this study

1. An improved version of one of the most complex and detailed pulmonary physiological models currently reported in the literature is implemented in MATLAB following an iterative multi-compartmental modelling approach.

2. To augment traditional model validation methods typically used in clinical settings, an advanced optimisation-based approach is adapted from the engineering literature and applied to validate the robustness of a pulmonary (patho)-physiological model.

3. A new approach for representing pulmonary disease states *in silico* using a (patho)-physiological model is presented. Significant modifications, based on a detailed consideration of the underlying physiological mechanisms, are incorporated into the model in order to allow it to accurately represent a number of different disease states.

4. A novel optimisation-based strategy is proposed to determine the disease-specific optimal settings for the ventilator which satisfy all clinical objectives of mechanical ventilation, many of which are in conflict with each other. The optimal settings are shown to significantly improve the ventilation-perfusion matching within the diseased lung, and to be highly disease specific.

The key publications resulting from this thesis are listed below:

Das A, Menon PP, Bates DG, Hardman JG. Optimisation of mechanical ventilator settings: towards personalised management of disease states, submitted to *IEEE Transactions on Biomedical Engineering*, 2012.

Das A, Gao Z, Menon PP, Hardman JG, Bates DG. A systems engineering approach to validation of a pulmonary physiology simulator for clinical applications. *Journal of Royal Society Interface* 2011; 8(54):44-55.

Das A, Gao Z, Menon PP, Hardman JG, Bates DG. Optimisation of mechanical ventilator settings. In proceedings of the *International Federation of Automatic Control (IFAC) World Congress*. Milan, Italy. 2011.

Das A, Bates, DG, Hardman JG. An engineering approach to validation of a pulmonary physiology simulator. In proceedings of the *International Congress on Computational Bioengineering (ICCB)*. Bertinoro (Forli), Italy. 2009.

Chapter 2: Modelling of pulmonary physiology

This chapter presents a brief introduction to the principles of pulmonary physiology that pertain to the modelling work described in this thesis. The structure of the lungs is discussed, an explanation of the important terms is provided and the most important mathematical equations describing arterial blood gas calculations are also given. The development of the extended implementation of the Nottingham Physiological Simulator (NPS) (Hardman 1998), is described in detail. Modifications (based on physiological mechanisms) are proposed to characterise healthy as well as diseased states.

2.1   Pulmonary physiology – key concepts

2.1.1   Lungs

The lungs primarily have two functions – the movement of oxygen (O_2) and carbon dioxide (CO_2) to and from the alveoli (through convection) and the provision of a surface for gas exchange (through diffusion). The lungs consist of approximately 300 million alveoli, surrounded by a network of capillaries, providing a large surface area for the transfer of gases. These gases reach the alveoli though a complex tree like structure of airways, starting from the trachea, which divides into the bronchi. These further divide into bronchioles, terminal bronchioles, alveolar ducts and finally the alveoli. Gas flow to and from the alveoli depends on the airway resistances and the pressure gradient (between the mouth and the lungs) created by the respiratory muscles, which causes the expansion and deflation of the lungs. The pulmonary artery brings deoxygenated blood to the lungs from the heart and divides into the pulmonary

capillaries. Here O_2 is diffused to the blood and CO_2 is diffused to the alveolar units. The pulmonary capillaries containing the oxygenated blood converge into the pulmonary vein. This blood is transported to the heart which pumps the oxygenated blood into systemic circulation.
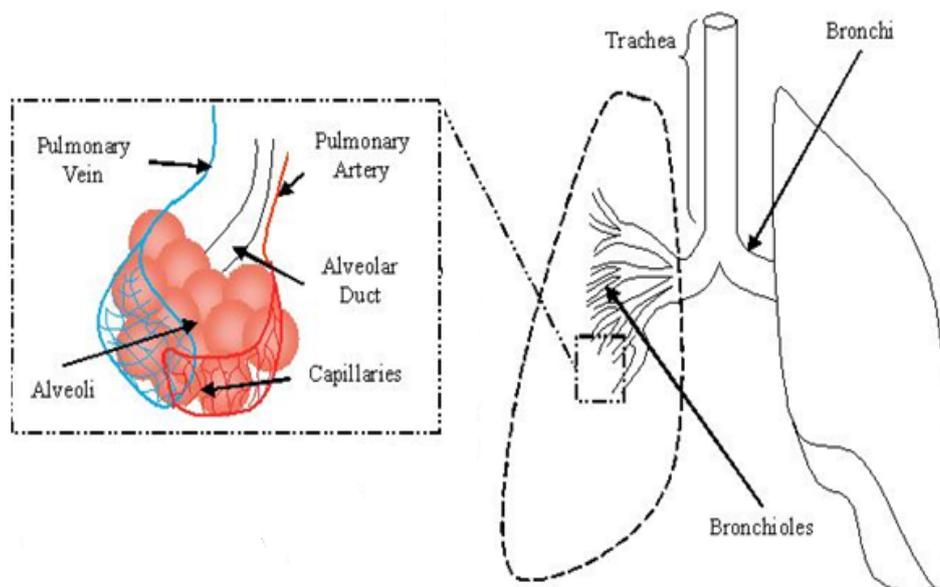
Figure 2.1: Lung structure

2.1.2 Lung volumes and capacities

The amount of air that the lung inhales, exhales or holds under different conditions can be described by specific lung volumes and capacities. Figure 2.2 introduces various lung volumes and capacities which are important clinically as their values can be affected by pathological processes. When compared to normal physiological ranges, they can help physicians in making diagnoses regarding the underlying pathological conditions. Their descriptions along with their approximate values (for a normal 70kg young male) are given below.
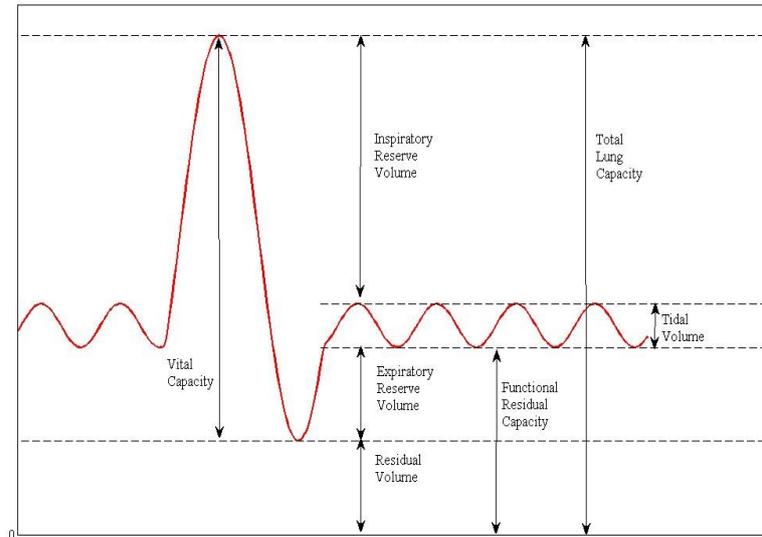
Figure 2.2: Lung volumes and capacities

Tidal Volume (Vtidal): this is the volume of air that moves in and out of the lungs in a single breath. It normally has a value in the range 400ml – 500ml but it is also dependent on other variables such as response to disease or exercise.

Residual Volume (RV): The volume remaining in the lung after a forced expiration. In general, the value of RV is found to be approximately equal to 1.5 litres.

Inspiratory Reserve Volume (IRV): The maximum volume that can be inspired after Vtidal. In general, IRV is approximately equal to 3 litres.

Expiratory Reserve Volume (ERV): The maximum volume that can be expired at the end of normal expiration. In general, ERV is approximately 1.5 litres.

Functional Residual Capacity (FRC): The total volume of gas left in the lung at the end of normal expiration. In general, the value of FRC is approximately 3 litres.

Total Lung Capacity (TLC): The maximum volume that the lung can stretch to while limited by the chest wall and the force of the inspiratory muscles. TLC in general is approximately 6 litres.

Vital Capacity (VC): The maximum volume that can be inhaled from Residual Volume. In general, the VC is approximately equal to 4.5 litres.

Lung function is most commonly tested using a spirometer which measures the volume and flow of air entering and leaving the lung during a breath. The spirometer is thus used to determine the lung volumes of Vtidal, IRV, ERV and VC. Other tests of lung functionality such as the nitrogen washout technique (NWT) can be used to approximate FRC. In NWT, a patient breathes air with 100% oxygen while the concentration of nitrogen is monitored in expired air using a nitrogen analyzer. When the nitrogen is completely 'washed out' of the expired air, the amount of nitrogen at the beginning of the test can be approximated by multiplying total volume of expired air with percentage of nitrogen in the expired air. As the air in the lung at the beginning of the test consisted of 80% nitrogen, the total volume at the beginning of the test can be approximated. Using FRC, RV can simply be calculated as RV = FRC - ERV. For detailed information on the spirometer, NWT and other tests for lung volumes/capacities the reader is referred to (Cotes 2006).

Estimates of lung volumes aid physicians in making pathological diagnoses. For example, restrictive diseases such as fibrosis reduce the compliance of the lung, which reduces these volumes as the lung struggles to expand under normal breathing patterns. Obstructive diseases such as chronic bronchitis and emphysema increase resistance to airflow. To overcome this resistance, high pressures are generated which can cause a decrease in the recoil of the lungs at expiration. RV, FRC and therefore TLC tend to be higher than normal under such conditions.

2.1.3 O_2 and CO_2 movement across the alveolar membrane

The exchange of gases between blood and ventilated air occurs by diffusion across the alveolar membrane. The O_2 partial pressures are higher in the alveoli than the capillaries and CO_2 partial pressure is higher in the capillaries, which creates a gradient for the diffusion of gases. The rate of diffusion depends on the surface area (large in the lungs due to the number of alveoli), the concentration gradients of gases (large between pulmonary capillary blood and alveolar air) and the individual properties of a gas. For example the rate of CO_2 diffusion across the alveolar membrane is dependent on the partial pressure of CO_2 molecules (lower than O_2) and the solubility of CO_2 in the plasma (CO_2 is more soluble in liquid than O_2).
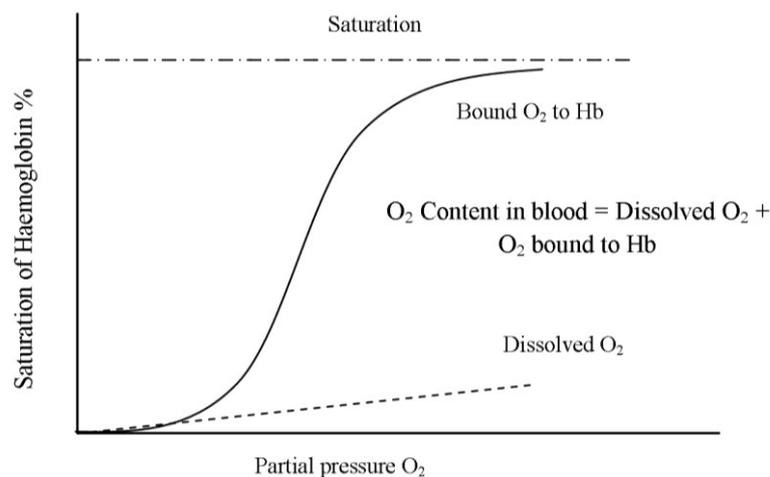
Figure 2.3: O_2 dissociation curves adapted from (West 2004, pg 76). O_2 = Oxygen, Hb = Haemoglobin.

Blood carries O_2 to the tissues in two forms: dissolved and combined with haemoglobin (Hb). As indicated by the dissociation curve in Figure 2.3, the dissolved O_2 constitutes a relatively small portion of the total blood oxygen content. This occurs due to the reduced solubility of O_2 in blood relative to CO_2. To satisfy the O_2 requirements of the body, O_2 is bound chemically to haemoglobin.

The plateau of the oxygen dissociation curve at high O_2 partial pressure (P_{O2}) allows the body to maintain oxygen saturation in the event of a sudden fall in alveolar P_{O2}. At the peripheral tissues, P_{O2} is low, so oxygen is released quickly. The curve can shift to the right under the following conditions:

1. Under the Bohr effect, wherein haemoglobin has less affinity to O_2 as H^+ ions in the blood increase. For example, in the peripheral tissues the increase in CO_2 as a result of metabolism increases the number of H^+ ions bound to haemoglobin and promotes the dissociation of O_2 at the peripheral tissue.

2. Under the Haldane effect, wherein haemoglobin has less affinity to CO_2 as the concentration of O_2 rises. For example, the increase in available O_2 promotes dissociation of CO_2 and uptake of oxygen from the blood.

3. As temperature increases - this is usually not very significant except in the case of extreme temperature changes such as hypothermia.

4. As the 2, 3-Disphosphoglycerate (2, 3 DPG – a regulatory organophosphate produced by red blood cells) concentration increases.

CO_2 is present in a higher quantity in blood than oxygen as it is more soluble and more chemically reactive. CO_2 has a major regulatory contribution in controlling blood pH levels. In plasma, the dissolved CO_2 reacts with water to form H^+ and HCO_3^- ion (bicarbonate) in the following reversible reaction.

$$CO_2 + H_2O \Leftrightarrow H_2CO_3 \Leftrightarrow H^+ + HCO_3^-$$

[2.1]
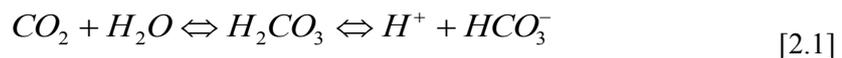
This equation represents the acid base equilibrium in the blood. If CO_2 increases, H^+ ions in the blood increase, which reduces the pH (making it more acidic). If HCO_3 increases, the available H^+ ions in the blood decrease and the pH increases (more alkaline). The reaction helps to

maintain the pH balance in the blood. In the human body, the kidneys also regulate the acid base balance in blood by absorbing HCO_3^- ions and producing H^+ ions. CO_2 also binds directly to proteins, known as carbamino compounds which contribute to CO_2 transfer around the body. In blood, CO_2 is dissolved in plasma (P_{CO_2} about 8%), in the form of H_2CO_3 (about 80%) and as carbamino compounds (about 11%).

2.1.4   The airways and the blood vessels
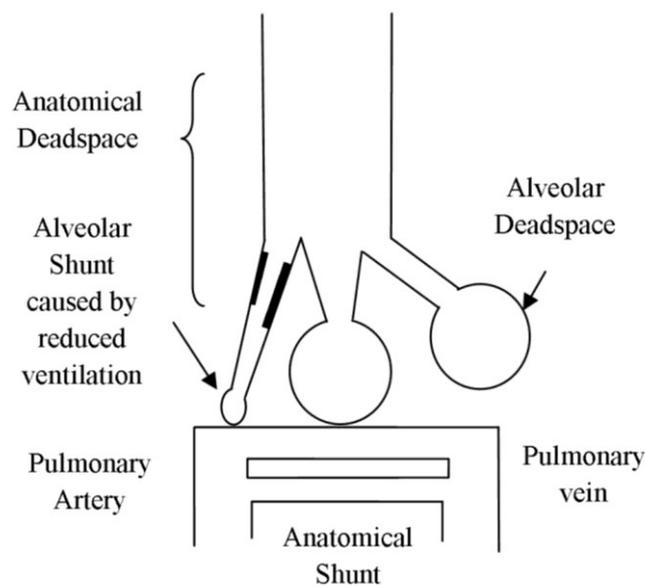
Figure 2.4: Deadspace and shunt.

The region of the lungs that does not participate in gas exchange is called the physiological deadspace. It consists of the series and the parallel deadspace.

- The series (anatomical) deadspace (SD) consists of the conducting airways such as the trachea, bronchi and the bronchioles.

- The parallel (alveolar) deadspace (PD) are alveoli which are well ventilated but poorly perfused by blood through the pulmonary capillaries.

The fraction of deoxygenated blood from the heart (cardiac output) that does not participate in gas exchange is called the physiological shunt. It is composed of the anatomical shunt and alveolar shunt.

- The anatomical shunt (AS) is the shunting of blood directly from the pulmonary artery to the pulmonary vein without participating in gas exchange.
- The alveolar shunt (AlvS) is the blood passing through pulmonary capillaries with poorly ventilated alveoli caused by obstructed or collapsed airways.

2.1.5   Distribution of ventilation and perfusion

An impairment of gas exchange functions in the lung can be indicated by poor oxygenation (arterial hypoxemia) and CO_2 retention (hypercapnia) or both. Mechanisms mentioned in Section 2.1.4 dictate the movement of the partial pressure of O_2 and CO_2 across the alveolar capillary membrane. The movement of O_2 from alveoli into the blood is determined by saturation of the haemoglobin (Figure 2.3). This is affected by the supply of O_2 to the tissues and extraction of CO_2 from the tissues as a result of metabolic processes.

The distribution of gas exchange can provide significant information regarding the pathology of the lung. Figure 2.5 shows the typical ventilation perfusion distribution of a healthy lung. The healthy lung is relatively heterogeneous, i.e. the ventilation and perfusion distribution is generally uniform across the lung. This conforms closely to the data reported in (Wagner 1974), in which 95% of ventilation and perfusion in a healthy lung is confined to within the V/Q ratio of 0.3 and 2.1.
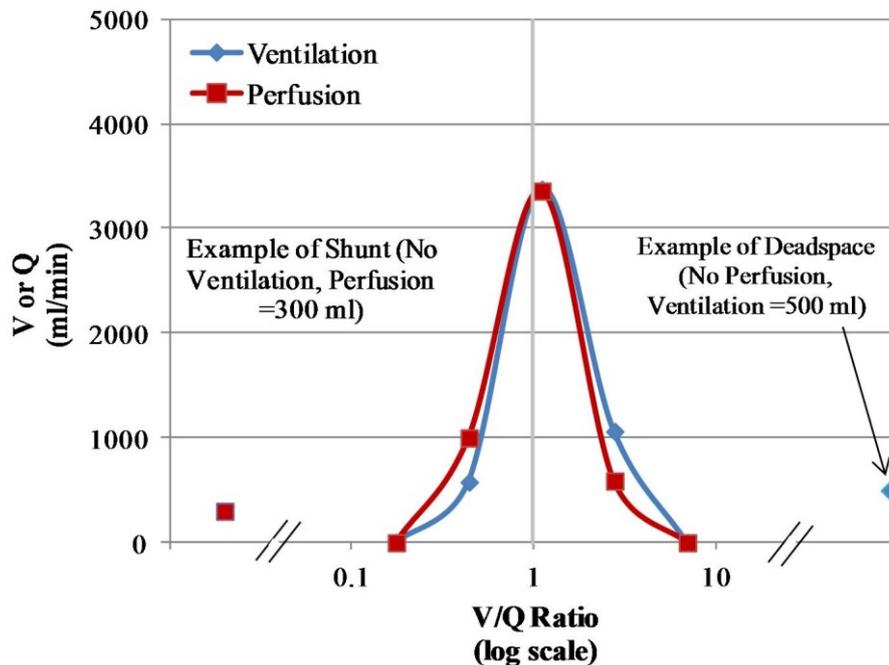
Figure 2.5: Ventilation (V) perfusion (Q) distribution of a healthy lung. The lines are drawn as a visualisation aide only.

Ventilation (V) in the lungs is influenced by the variation in distribution of air across the lung. The perfusion (Q) is dependent on the cardiac output, the arterial-venous pressure gradient and arterial-alveolar pressure gradient. Efficient gas exchange occurs when ventilation is adequately matched to perfusion in the gas exchanging lung unit. This matching can be represented by the V/Q (ventilation/perfusion) distribution, where a V/Q ratio of 1 would indicate that the ventilation and perfusion is matched and the lung unit is normal. At one extreme of V/Q distribution, there is *deadspace* (Section 2.1.3) which is an area of the lung that is normally ventilated but not perfused with blood. At the other extreme of the V/Q distribution is the area of the lung with normal perfusion where the alveoli is blocked or collapsed and no ventilation occurs (Section 2.1.3). All three characteristics contribute to the V/Q distribution.

Many pathological anomalies can affect the airways and the blood vessels, creating an imbalance and resultant abnormal V/Q ratio. High V/Q contributes to a higher P_{O2} and increases the amount of wasted ventilation. However at high P_{O2}, the O_2 dissociation is robust to sudden decrease in P_{O2}, as indicated by the plateau of the haemoglobin (Hb) saturation curve of the O_2 dissociation (Figure 2.3). Under low V/Q, the effect of shunt is evident as poorly ventilated areas create low P_{O2} and low Hb saturation which lead to reduced oxygen content and hypoxemia. Therefore, decrease in P_{O2} is not followed by a proportional decrease in haemoglobin saturation.

V/Q imbalance holds key information regarding the state of the lung, but it is hard to measure accurately as the relationship can be heavily influenced by time varying parameters of cardiac output and minute ventilation (amount of air in breaths per minute). For a more detailed description of the relevant physiology, the reader is referred to (Cotes 2006) and (West 2004).

2.2 Modelling and implementation of a pulmonary physiology simulator

2.2.1 Modelling of pulmonary systems - background

Since the early 20th century, many efforts (Henderson 1920, Van Slyke 1921) have been made to explain the underlying theoretical principles of the relationships between ventilation, blood flow and gas exchange in pulmonary physiology. The outcome of this research has been the derivation of mathematical equations which have become fundamental to the study of pulmonary physiology in recent years. The earliest analysis of these equations is reported in (Rahn 1949), which determines the alveolar gas concentrations using several analytical techniques. The detailed insight into the structure of the human lung as a geometric model of the airways consisting of 24 generations was originally proposed in (Weibel 1963). The reducing

ratio of radii of the airway as the generations progressed was another important outcome of this pioneering research which laid the foundation for several other follow-up studies as in (Peterman 1984, Wiggs 1992, Tsoukias 1998 , Yem et al 2003).

To the best of the author's knowledge, the first ever published dynamical representation of the respiratory regulation is by Grodins and colleagues in 1954 (Grodins 1954, 1967). In (Grodins 1967) the proposed model characterized the relationship among lung, blood and tissue with differential equations, in which the time delays of the gas exchange system were accounted for. The model included the effect on O_2-CO_2 interaction by the chemoreceptors, chemical buffering etc, and has formed the basis of many other theoretical works (Saunders 1980, Ursino 2001, Stuhmiller 2005). However, the model developed here is extremely simplified, with only three compartments representing the lung, the brain and the tissues. The simplified lung is modelled as a compartment of constant volume, uniform content and no deadspace, ventilated by a continuous unidirectional stream of air. The model was utilised to understand the gas transport and acid-base buffering in the respiratory system.

Research works by John B West (West 1969) and colleagues helped to develop deeper insight into pulmonary ventilation, perfusion and gas exchange. To achieve this, computational models of the lungs were used and the effects of varying ventilation perfusion inequality were analysed. West studied the effects of varying numbers of alveolar compartments (between 3-1000) and recommended that at least 100 alveolar compartments were needed for an optimal compromise between the computational complexity and accuracy of pulmonary models. Furthermore, the effect of variation of ventilation perfusion ratios on O_2 and CO_2 transfer into the blood was also investigated. However, individual compliances, inflow resistances and vascular resistances that are associated with alveolar compartments were not accounted for in these models.

Chiari (Chiari 1997) and colleagues presented a model of the respiratory system incorporating O_2-CO_2 transport and exchange. The model constitutes three compartments: the lung, the brain and the body. The basic representation was an adoption of the model of Grodins (Grodins 1967) with the added implementation of nonlinear O_2-CO_2 gas movement and a ventilation controller that adjusted alveolar ventilation and cardiac output. The gas flows are assumed to be continuous and unidirectional. The lungs as well as the body tissue were represented as single units and hence effects of variable ventilation, perfusion and metabolic rates were not considered.

The number of distinct alveolar compartments considered and the motivation for choosing this number has varied hugely across the literature. Single/bi-compartmental models offer simplicity in calculations and allow for better understanding of some inadequately understood physiological behaviour. Swanson (Swanson 1983) used a bi-compartment model to analyse the alveolar gas exchange during exercise relative to a normal lung. Vidalmelo (Vidalmelo 1993, 1998) used a bi-compartmental model to depict the alveolar-arterial gas transport through analytically describing the gas diffusion, alveolar-arterial pressure gradient, CO_2 and O_2 dissociation curves and the behaviour due to physiological shunt and cardiac output. Hotchkiss (Hotchkiss 1994) investigated the effects on ventilation distribution and alveolar pressures of mechanical ventilation inputs such as frequency and duty cycle. Other examples of research studies using single and bi-compartmental alveolar models of the lung include Farmery (1996), Joyce (Joyce 1996), Liu (Liu 1998). However, single and bi-compartmental models of the lung can be inadequate in representing the effect of inhomogeneties arising from the distribution of ventilation / perfusion, variations in time constants associated with respiratory frequency, tidal breathing, series/parallel resistances etc, especially in different pathological states.

Many models (e.g. (West 1969)) have thus incorporated multiple gas exchanging compartments while also attempting to capture other characteristics of the respiratory system in more detail. These models have been used for the analysis of distributions of gases, resistances, and compliances across the whole lung that are important, but generally ignored in the more simplistic models. Hinds (Hinds 1980, 1984) proposes one such model, accounting for ventilation perfusion distribution, and looking at variations of volume and pressure within the lung. Petrini (Petrini 1983) measured the effect of series deadspace created by the conducting airways and the role of this in effective ventilation. Hickling (Hickling 1998) looked at the combined effect of recruitment and derecruitment of 27000 alveoli in parallel that offered an insight into the pressure volume relationship of these alveolar units.

Rutledge (Rutledge 1993, 1994) proposed two versions of a respiratory model within his proposed ventilator administrator strategies, VentSim and VentPlan. VentPlan is a simplified physiological model with uniform airway resistance, fixed compliance, no ventilation perfusion abnormalities apart from the anatomical shunt, and a simple input as the driving pressure of the lungs. VentSim is a relatively more thorough representation, but so far remains unvalidated relative to VentPlan. Rutledge's simplistic model as used in VentPlan may compromise its applicability on patients with more complex pathophysiological abnormalities.

Rees (2006) presented a model of O_2 and CO_2 transport in the lungs, blood and the tissues and applied it for the selection of optimal mechanical ventilator settings. The lungs consist of two compartments each with different ventilation perfusion ratios. The circulation of blood is represented by a mixed venous compartment and an arterial compartment, storing corresponding blood gas parameters. The pulmonary model used is extremely simplified and is unlikely to accurately represent pathophysiological conditions due to the relatively inadequate ventilation-perfusion distribution considered.

Yem (Yem 2003) present a model of a comprehensive cardio respiratory system. This system includes compartments for cardiovascular circulation, brain, peripheral shunt, liver, kidneys, muscles and fat. The series deadspace incorporates a compartmental approximation to a lung model proposed by Weibel (Weibel, 1963). The ventilation perfusion distribution is limited to four alveolar compartments (one unperfused, three perfused). The model has been applied in determining the sources of error in non invasive blood flow measurements (Yem 2003) and analysing ventilation perfusion abnormalities in diseased lungs (Yem 2006).

In the literature, the respiratory system has often been depicted as a number of lumped gas exchanging units. Simplistic models, i.e. single or bi-compartmental compartmental representations, have proved to be useful in improving the understanding of pulmonary physiology. However, difficult problems such as understanding and treating lung injury under different pathological conditions require the incorporation of much more detailed models of pulmonary physiology. For this purpose, this thesis presents the development of a comprehensive physiological model, specifically an improved implementation of the Nottingham physiology simulator.

The Nottingham physiology simulator (NPS)

The NPS is a set of physiological models representing the pulmonary system, gas transport in the blood and gas exchange at the peripheral tissues. The main features of the simulator are shown in Figure 2.6.

The models use mass-conserving algebraic equations of pulmonary dynamics and blood gas transport. Beginning with default initial settings, the equations are solved iteratively in a sequence, where each iteration represents a 'time slice' t of real physiological time (Hardman 2001). Theoretically, the shorter the time slices are the more accurate the representation of the

physiological dynamics will be, but this advantage is offset by the corresponding increase in computation times.
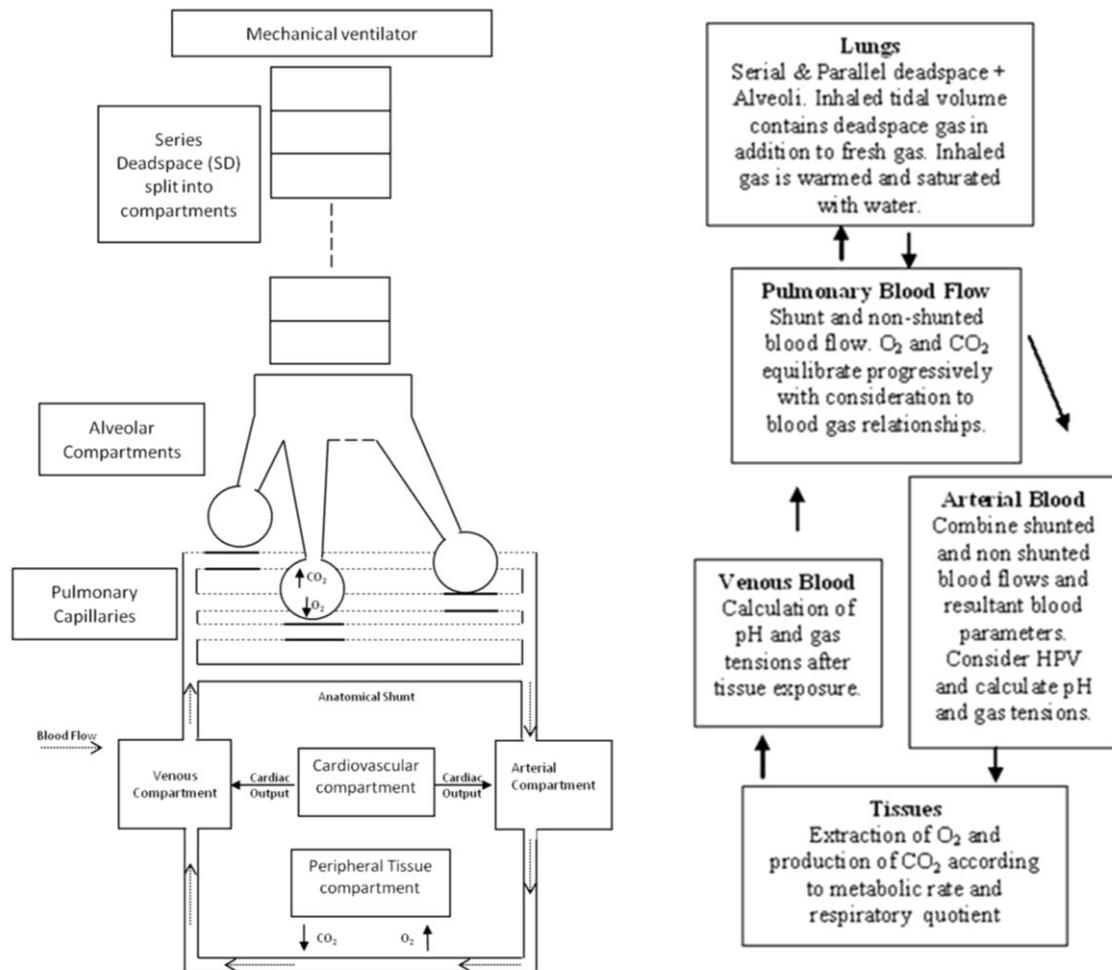
Figure 2.6: Diagrammatic representation of important features of the respiratory model in the NPS (adapted from Hardman et al 1998)

The pulmonary model consists of the mechanical ventilation equipment, anatomical and alveolar deadspace, anatomical and alveolar shunts, ventilated alveolar compartments and corresponding perfused capillary compartments. To recap from Section 2.1.4, the series deadspace (SD) consists of the trachea, bronchi and the bronchioles where no gas exchange occurs. It is located between the mouth and the alveolar compartments. Inhaled gases pass

through the SD during inspiration and alveolar gases pass through the SD during expiration. In the model, an SD of volume 60ml is split into 50 stacked compartments of equal volumes (see Figure 2.6). No mixing between the compartments of SD is assumed.

Any residual alveolar air in the SD at the end of expiration is re-inhaled as inspiration is initiated. This residual air is composed of gases exhaled from both perfused alveolar compartments (normal perfusion) and the parallel deadspace (PD) (alveolar compartments with limited perfusion). Therefore, the size of deadspace (SD and PD) can have a significant effect on the gas composition of the alveolar compartments. Compared to analytic techniques representing 'single step' tidal movement of gases through the SD (e.g. Stuhmiller 2005), the iterative modelling technique has the important advantage of producing realistic and gradual changes in gas compositions of the SD and the alveolar compartments and accounting for the residual gases from the previous breath.

The inhaled air consists of five gases: oxygen, nitrogen, carbon dioxide, water vapour and a 5th gas used to model additives such as helium or other anaesthetic gases. During inspiration, the flow (f_i) of air to or from an alveolar compartment i in a time slice t is determined by the following equation:

$$f_i = \frac{(P_v - P_{Ai})}{(R_u + R_{Ai})} \qquad [2.2]$$

where P_v is the pressure supplied by the mechanical ventilator, P_{Ai} is the pressure in the alveolar compartment i at t, R_u is the constant upper airway resistance and R_{Ai} is the bronchial inlet resistances of the alveolar compartment i. The total flow of air entering the SD at time t is calculated by

$$f_{SD} = \sum_{i=1}^{100} f_i \qquad [2.3]$$

During inspiration, the fractions of gases (F^x) in each compartment of the SD are updated based on the composition of f_i. The volume of gas x in alveolar compartment (v_A^x) is given by:

$$v_A^x = \begin{cases} v_A^x + f_i \cdot \frac{v_A^x}{v_A} & \textit{Exhaling} \\ v_A^x + f_i \cdot F^x & \textit{Inhaling} \end{cases} \qquad [2.4]$$

Each alveolar compartment can be assigned a bronchial inlet resistance and compliance. Compliance specifies the volume gradient with respect to pressure, and represents a measure of the 'stiffness' of the lung. Ageing and pulmonary disorders such as emphysema increase the compliance of the lung and reduce the elastic recoil of the lung. Other pulmonary disorders such as fibrosis and pneumonia decrease the compliance of the lung and make the lung 'stiffer'.

Following (Macklem 1967), a fixed resistance to the flow was introduced. The tension at the centre of the alveolus and at the alveolar capillary border is assumed to be equal. 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). This is modelled as a simple function, resembling the stimulus response curve suggested by Marshall (Marshall 1994), and incorporated into the simulator to gradually constrict the blood vessels as a response to low alveolar oxygen tension. The atmospheric pressure is fixed at 101.3kPa and the body temperature is fixed at 37.2°C.

During each time slice t, equilibration between the alveolar compartment and the corresponding capillary compartment is achieved iteratively by moving small volumes of each gas between the compartments until the partial pressures of these gases differ by <1% across the alveolar-capillary boundary. The process includes the nonlinear movement of O_2 and CO_2 across the alveolar capillary membrane during equilibration.

In blood, the total O_2 content (C_{O2}) is carried in two forms, as a solution and as oxyhaemoglobin (saturated haemoglobin in Figure 2.3):

$$C_{O2} = S_{O2} \cdot Huf \cdot Hb + P_{O2} \cdot O2sol \qquad [2.5]$$

In this equation, S_{O2} is the haemoglobin saturation, Huf is the Hufner constant, Hb is the haemoglobin content and O2sol is the O_2 solubility constant. The following pressure-saturation relation, as suggested by (Severinghaus 1979) to describe the O_2 dissociation curve, is used in this model:

$$S_{O2} = \left(\left(\left(P_{O2}^3 + 150 \cdot P_{O2}\right)^{-1} \times 23400\right) + 1\right)^{-1} \qquad [2.6]$$

S_{O2} is the saturation of the haemoglobin in blood and P_{O2} is the partial pressure of oxygen in the blood. Eqn. 2.6 applies to a standard O_2 dissociation curve such as that shown in Figure 2.3 (at pH=7.4 and temperature=37°C). As suggested by (Severinghaus 1966), P_{O2} has been determined with appropriate correction factors in base excess BE, temperature T and pH (7.5005168 = pressure conversion factor from kPa to mmHg):

$$P_{O2} = 7.5006168 \cdot P_{O2} \cdot 10^{[0.48(pH\text{-}7.4)-0.024(T\text{-}37)-0.0013\cdot BE]} \qquad [2.7]$$

The CO_2 content of the blood (C_{CO2}) is deduced from the plasma CO_2 content ($C_{CO2}plasma$) by the following equation (Douglas 1988):

$$C_{CO2} = C_{CO2plasma} \cdot \left[1 - \frac{0.0289 \cdot Hb}{(3.352 - 0.456 \cdot S_{O2}) \cdot (8.142 - pH)}\right] \qquad [2.8]$$

where S_{O2} is the O_2 saturation, Hb is the haemoglobin concentration and pH is the blood pH level. The coefficients were determined by Douglas as a standardised solution to the McHardy

version of Visser's equation (McHardy 1967) by iteratively finding the best fit values to a given set of clinical data.

The value of $C_{CO2plasma}$ in Eqn. 2.8 can be deduced by the Henderson-Hasselbach logarithmic equation for plasma C_{CO2} (Kelman 1966):

$$C_{CO2plasma} = 2.226 \cdot s \cdot P_{CO2plasma} \left(1 + 10^{(pH - pK')}\right) \qquad [2.9]$$

where s is the plasma CO_2 solubility coefficient and pK' is the apparent pK (acid dissociation constant of the CO_2 bicarbonate relationship). $P_{CO2plasma}$ is the partial pressure of CO_2 in plasma and '2.226' refers to the conversion factor from miliMoles per litre to ml/100ml. (Kelman 1966) gives the equations for s and pK' as:

$$s = 0.0307 + 0.0057 \cdot (37 - T) + 0.00002 \cdot (37 - T)^2 \qquad [2.10]$$

$$pK' = 6.086 + 0.042 \cdot (7.4 - pH) + (38 - T) \cdot \left(0.00472 + \left(0.00139 - (7.4 - pH)\right)\right) \qquad [2.11]$$

For pH calculation, the Henderson Hasselbach and the Van Slyke equation (Siggard-Andersen 1977) are combined. Below is the derivation of the relevant equation. The Henderson-Hasselbach equation (governed by the mass action equation (acid dissociation)) states that (Cotes 2006, pg 70):

$$pH = pK + \log\left(\frac{bicarbonate\ concentration}{carbonic\ acid\ concentration}\right) \qquad [2.12]$$

Substituting pK=6.1 (under normal conditions) and the denominator ($0.225 \cdot P_{CO2}$) (acid concentration being a function of CO_2 solubility constant 0.225 and P_{CO2} (in kPa)) gives:

$$pH = 6.1 + \log\left(\frac{HCO_3}{0.225 \cdot P_{CO2}}\right) \qquad [2.13]$$

For a given pH, base excess (BE), and haemoglobin content (Hb), HCO_3 is calculated using the Van-Slyke equation as given by (Siggard-Andersen 1977):

$$HCO_3 = \left((2.3 \times Hb + 7.7) \times (pH - 7.4)\right) + \frac{BE}{(1 - 0.023 \times Hb)} + 24.4 \qquad [2.14]$$

The capillary blood is mixed with arterial blood using Eqn. 2.15 which considers the anatomical shunt (Sh) with the venous blood content of gas x (C_v^x), the non-shunted blood content from the pulmonary capillaries (C_{cap}^x), arterial blood content (C_a^x), the arterial volume (v_a) and the cardiac output (CO).

$$C_a^x = \frac{CO \cdot \left(Sh \cdot C_v^x + (1 - Sh) \cdot C_{cap}^x\right) + C_a^x \cdot (v_a - CO)}{v_a} \qquad [2.15]$$

The peripheral tissue model consists of a single tissue compartment, acting between the peripheral capillary and the *active* tissue (undergoing respiration to produce energy). The consumed O_2 (V_{O2}) is removed and the produced CO_2 (V_{CO2}) is added to this tissue compartment. Similarly to alveolar equilibration, peripheral capillary gas partial pressures reach equilibrium with the tissue compartment partial pressures, with respect to the nonlinear movement of O_2 and CO_2. Metabolic production of acids, other than carbonic acid via CO_2 production, is not modelled.

A simple equation of renal compensation for acid base disturbance is incorporated. The base excess (BE) of blood under normal conditions is zero. BE increases by 0.1 per time slice if pH falls below 7.36 (to compensate for acidosis) and decreases by 0.1 per time slice if pH rises above 7.4 (under alkalosis). The cardiovascular model considered here simply sets the cardiac output (blood flow entering the pulmonary model) per time slice.

NPS has been applied previously in the study of various physiological phenomena and as a training simulator. A simulated patient is assumed to be under complete mechanical ventilation. Consequently, the effects of ventilatory autoregulation by the patient have not been incorporated into the models. Currently, it also does not take into consideration any metabolic, myogenic or neurogenic autoregulation of the physiological parameters.

It should be noted that the NPS is effectively solving a set of ordinary differential equations numerically using an Euler method over a time step t. A large value of t can produce inaccurate results. The value of t is fixed at 5 ms for the all the results that were produced for this thesis, which is considerably shorter than earlier investigations where t of up to 200ms were regularly employed (Hardman 2000). The implementation of more stable numerical method can have significant advantages with respect to faster computational speeds and availability of tools for various computational analyses. However the thesis is focused towards analysing and implementing the current model. The processes implemented in the NPS have been described and validated successfully against clinical data in various research publications (Hardman 1998, Hardman 1999a, Hardman 1999b, Hardman 2000, Hardman 2003).

The NPS has been implemented into the MATLAB environment (Matlab R2006a on a PC running Microsoft Windows XP version 2002). The new MATLAB implementation of the simulator has several advantages over the previous PASCAL based implementation of the NPS:

- The model is easily integrated with an extensive library of MATLAB toolboxes (such as the global optimisation toolbox, robust control toolbox) for computational analysis & design. This is useful for the work on model validation and the identification of optimal ventilator settings to be performed in the subsequent stages of this project. The easy availability of systems and control toolboxes provide an easy interface for the wider control community.

- The MATLAB simulator allows easy storage and access to all variables involved for the duration of the simulation, with available computer memory being the only constraint. This makes small variations in any selected variable straightforward to record and monitor for further analysis.

- The explicit high-level MATLAB code is easy to amend and available to scrutiny, hence any modifications are easy to administer for the user, unlike the executable file only approach of NPS where access to the code is only available to the developer.

- The results produced from running simulations are saved in a MAT format which makes it easier to display the data and to transfer to spreadsheet software such as MICROSOFT EXCEL.

MATLAB is an interpreted language, i.e. an interpreter program executes the source code. This makes program execution slower compared to compiled languages where the source code is converted directly to machine code. However the current applications of the physiology simulator are not real time dependent, hence the present simulation times should be acceptable. Furthermore, by linking C files with MATLAB, individual sections of the code are optimised for rapid execution and *bottlenecks* in the code are reduced.

Since the MATLAB simulator is not GUI (graphical user interface) based, it is not as easy to manually change values and settings during simulations and to consequently monitor the effects. However the current scope of the simulator is for theoretical and computational research and not for real time ITU training, and thus this is not expected to be a significant issue.

Finally, it is important to point out that both simulators can work well alongside each other. For example, to match the models to a specific clinical scenario, the GUI based NPS can be used to determine the matching parameters such as airway resistances etc which would then make

simulation in MATLAB very simple. Conversely, if a user finds the NPS GUI restricting, the MATLAB simulator will allow for an easier access to any variable within the code.

2.2.2 Simulation model initialisation

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
Number of alveolar compartments	100
Respiratory quotient	0.8
Oxygen consumption	250 ml min^{-1}
Cardiac output	5 litres min^{-1}

Table 2.1: Simulation model default configuration (unless otherwise stated)

The lungs are modelled as a system comprising external equipment (e.g. a mechanical ventilator), anatomical and alveolar deadspace, and ventilated and perfused alveoli. In this study, 100 individual alveolar compartments have been incorporated into the model. The initial configuration of the model is given in Table 2.1. The inspired air consists of oxygen (19.6%), nitrogen (74%), carbon dioxide (0.1%). The balance is made up of water vapour (6.3%). Each

alveolar compartment can be attributed a specific bronchial inlet resistance to create a desired ventilation-perfusion distribution (see Section 2.1.6).

2.2.3   Program Structure

Figure 2.7 uses a sequence diagram to formally describe the structure of the code. A sequence diagram (Fowler 2003) shows the interaction between a group of objects. The vertical information in the diagram represents the sequence (order) of messages/calls as they occur while the horizontal information from left to right shows the objects and the messages that are sent between objects. Each object has a *lifeline* (a box with vertical dashed lines) and an *activation bar*, that shows whether the object is active or not. Thus the sequence diagram demonstrates the construction of functional modules in the simulator. The diagram further explains the flow of the information among the various functional modules of the software. The compartmentalisation of the model allows for easier validation and testing of individual components during implementation and the direct access to the structures such as the cardiovascular system (CVS) for future modifications.
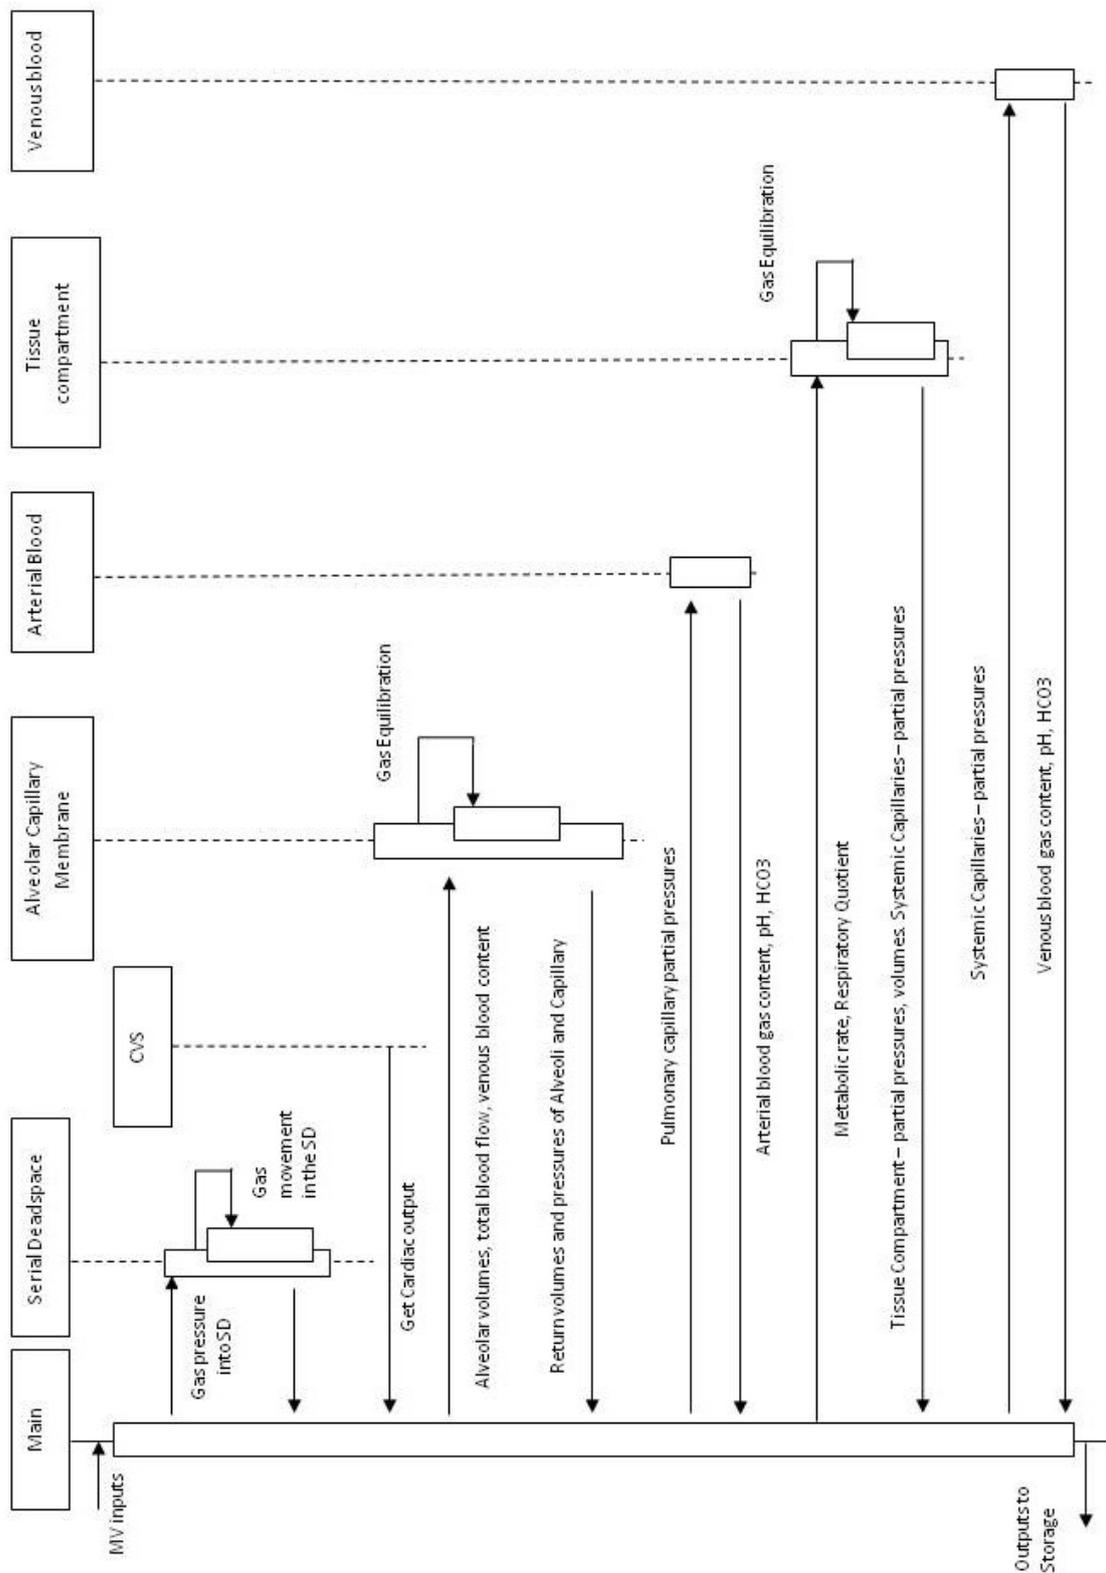
Figure 2.7: Sequence diagram of the objects and their relationships in the MATLAB implementation of NPS.

2.3   Model development and utilisation

The simulated lung in the model consists of 100 parallel alveolar compartments, each with three major parameters: volume of the alveolar compartment (v), pressure within the alveolar compartment (P) and bronchiolar inlet resistance (BIR). The flow of air to and from the alveolar compartment can be observed by the change in v and represents the ventilation (V) of the alveolar compartment. Each alveolar compartment has a corresponding pulmonary capillary compartment with a pulmonary vascular resistance (PVR). The flow of blood (Q) within the pulmonary compartment represents the perfusion of the pulmonary capillary compartment. The ventilation and perfusion (V/Q) distribution of the simulated lung can be modified by adjusting the BIR and PVR of the alveolar compartment.

The mechanism for the movement of air from the mechanical ventilator, through the serial deadspace and to the alveolar compartments can be summarised as follows:

1) The lung is at rest at the end of expiration. The volume v of each alveolar compartment at the end of expiration is given by the FRC (functional residual capacity) of the compartment. The FRC of a healthy alveolar compartment ($FRC_{nominal}$) is set at a default value of 30ml. For a healthy lung, the pressure P_{lung} of the alveolar compartment at FRC is zero.

2) At the initiation of inspiration, mechanical ventilator settings of tidal volume and PEEP are converted to positive pressure P_{MV} at the top of the airways, developing a pressure gradient between the mechanical ventilator and alveolar compartments.

3) The pressure gradient causes a flow of air into the alveolar compartments. This is equivalent to a flow of current between two points in an electric circuit, where a potential difference exists across an ohmic resistance. Extending the standard electric circuit analogy, the flow of air

between the ventilator and the lung can be determined using the equivalent of "Ohm's Law": the difference in pressures ($P_{MV} - P_{lung}$) divided by the total resistance of the conducting airways.

4) The flow contributes to the existing volume of the alveolar compartments,

5) The instantaneous pressure in the alveolar compartment is calculated using the static piecewise linear volume-pressure relationship (Figure 2.8). The expansion of the lung above $FRC_{nominal}$ represents the compliance of the alveolar compartment. The maximum volume that the alveolar compartment can achieve is the total lung capacity (TLC). During inspiration, the pressure P of the alveolar compartment rises as v increases.
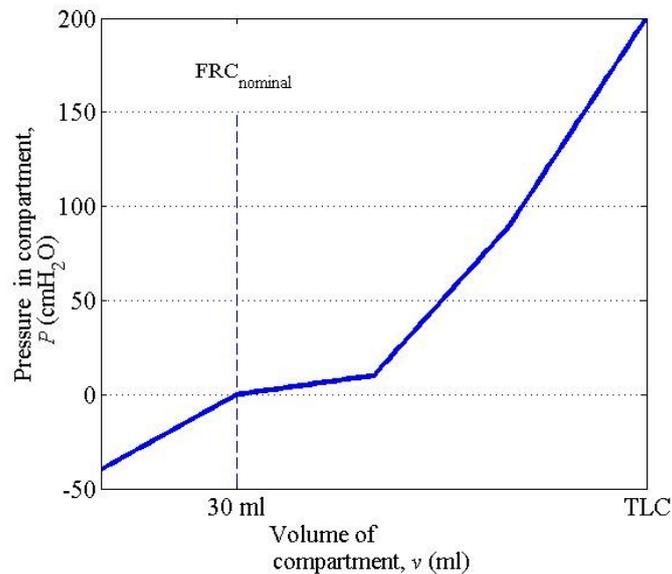
Figure 2.8: The static volume-pressure relationship used in the model

6) At the end of inspiration the ventilator pressure is reduced to zero as the ventilator output is removed. The lung exhales passively along the pressure gradient until the pressure in the lung equilibrates (returns to zero) at resting volume (FRC).

2.3.1 Simulating a healthy lung

For a healthy lung, the original model of 100 alveolar compartments (Hardman 2001) incorporated a fixed BIR (BIR$_{nominal}$) of 0.001 kpa per ml per minute and fixed PVR (PVR$_{nominal}$) of 16 ml per minute. The current model implementation randomly distributes BIR and PVR with a uniform distribution of $U(0.75x, 2x)$ where x is BIR$_{nominal}$ and PVR$_{nominal}$ respectively. This enables the simulation of heterogeneous ventilation perfusion (V/Q) distribution in the simulated lung as proposed by (Wagner 1974) and given in Figure 2.9.
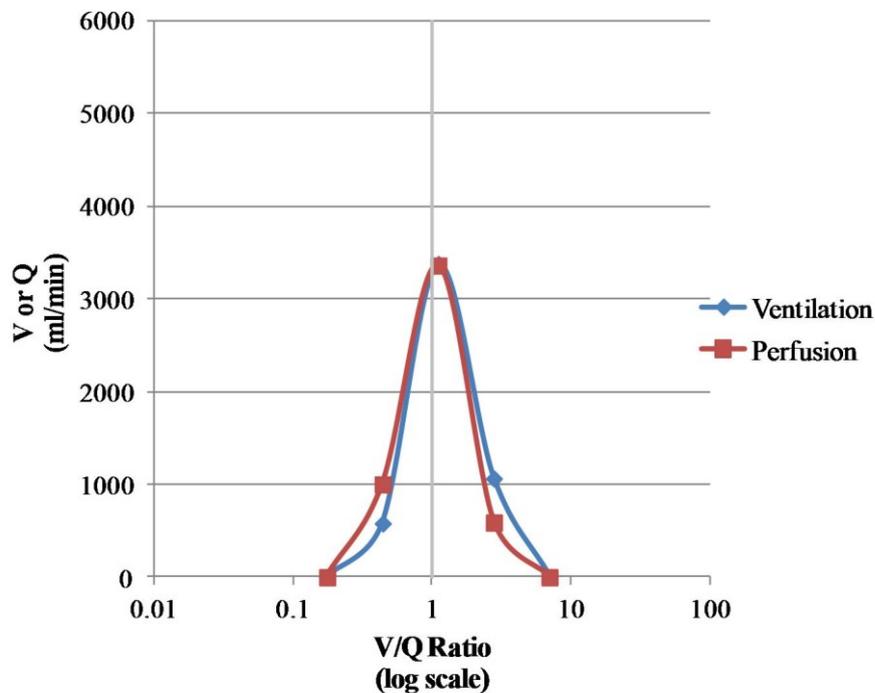
Figure 2.9: V/Q distribution of a healthy lung

Figure 2.10 depicts a five minute simulation of the change in volume of 100 alveolar compartments for a healthy simulated lung. The plot was produced using a tidal volume of 400

ml, ventilator rate (VentRate) of 12 breaths per minute and an inspiratory to expiratory ratio (I:E) of 1/3. These settings correspond to the typical breathing pattern of a healthy human.
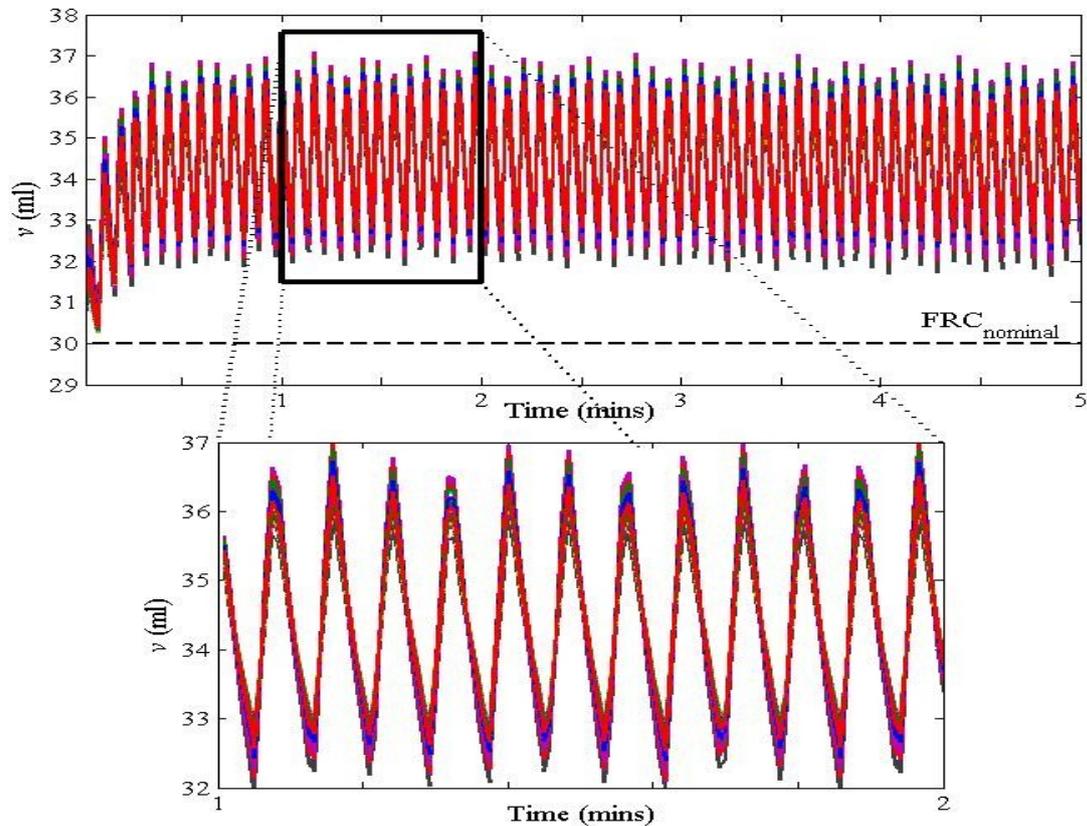
Figure 2.10: Change in volume (v) of healthy alveolar compartments. Each alveolar compartment has a unique BIR, resulting in the small variations in v shown in the projection. Vtidal = 400ml, VentRate = 12 bpm and I:E = 1/3.

2.3.2   Simulating diseased lungs

Abnormal V/Q distribution in the lung results from the obstruction of air flow and blood flow. Abnormal V/Q distribution has been shown to exist in pulmonary diseases such as chronic obstructive pulmonary disease (COPD) (Wagner 1977), asthma (Wagner 1978), pulmonary embolism (Elliott 1992), and acute respiratory distress syndrome (ARDS) (Pappert 1994).

Virtual diseased lungs can be simulated in the model by creating abnormal V/Q distributions. These distributions are configured by increasing the magnitudes of the BIR and PVR of the alveolar and capillary compartments. This method has previously been applied in (Kathirgamanathan 2009) to study responses of lungs in patients suffering from ARDS. In the Kathirgamanathan study, the lungs were divided into 3 regions (A_l, B_m, C_n) where l, m and n are the number of alveolar compartments in each region such that $l+m+n = 100$. Each compartment of each region was initially configured with BIR = $BIR_{nominal}$ and PVR = $PVR_{nominal}$. The BIR of the compartments of region A were then adjusted to the values of αx where $x= BIR_{nominal}$ and α is an integer vector of size l such that:

$$\alpha = [\alpha_{max}... \ 1], \ \alpha_{max} >= 1 \qquad [2.16]$$

The most severe obstruction to ventilation in the Kathirgamanathan study was created with α_{max} = 900. The PVR of the compartments of the region C were adjusted to the values of βy where $y= PVR_{nominal}$ and β is an integer vector of size n such that:

$$\beta = [1... \ \beta_{max}], \ \beta_{max} >= 1 \qquad [2.17]$$

The most severe obstruction to perfusion in the Kathirgamanathan study was created with β_{max} = 16. The alveolar compartments in the remaining region B remained configured as BIR = $BIR_{nominal}$ and PVR = $PVR_{nominal}$. This created a ventilation perfusion mismatch in the lung such that region A corresponds to alveolar shunt (low ventilation, high perfusion); region C corresponds to alveolar deadspace (high ventilation, low perfusion) with the region B corresponding to a healthy V/Q distribution within the lung. It should be noted that unlike the Kathirgamanathan study, we assign a randomly selected BIR and PVR to each alveolar and capillary compartment with a uniform distribution of $U(0.75x, 2x)$ where x is $BIR_{nominal}$ and $PVR_{nominal}$ respectively.

2.3.3 *In silico* behaviour of diseased lungs

The relationship between changes in the pressure distending the alveoli and the change in the lung volume dictates the inflation and deflation of a lung during a breath. Figure 2.11 depicts a typical pressure volume (PV) curve. An important feature of this curve is the difference between PV curves of inflation and deflation. This difference is known as hysteresis and is attributed to the opening of alveoli (recruitment) at different stages of inspiration and closing (derecruitment) during different stages of expiration. The pressure that is needed to recruit derecruited regions is called the Threshold Opening Pressure (TOP).
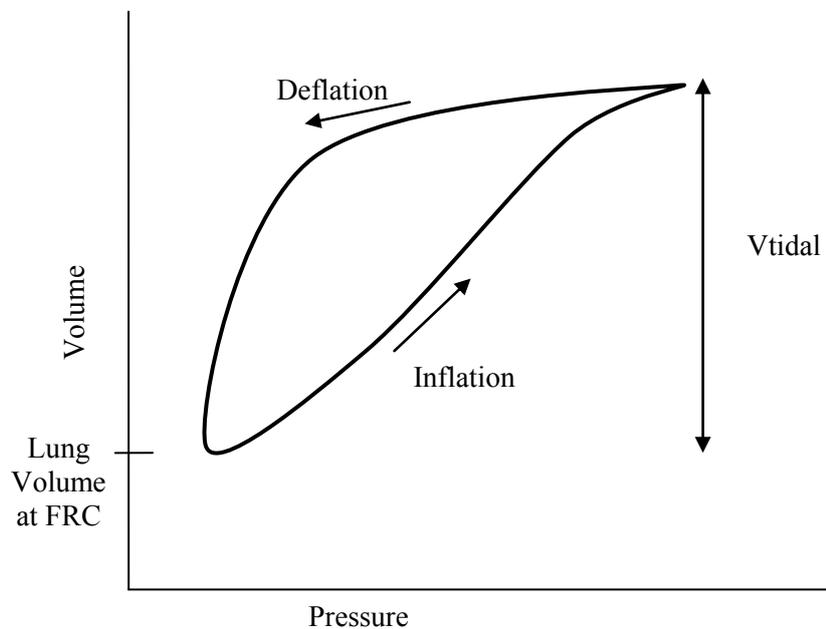
Figure 2.11: The pressure volume curve of a lung. A transpulmonary pressure inflates the lung which could be created under the negative pressure of normal breathing or positive pressure applied by a mechanical ventilator.

Patients suffering from severe respiratory disease and requiring mechanical ventilation, such as Acute Respiratory Distress Syndrome (ARDS), are prone to alveolar and airway collapse. Such collapse is common and found to be present in 90% of subjects undergoing mechanical ventilation (Lundquist 1995, Hedenstierna 2000] and is referred to as *atelectasis*. In such cases,

size of effective lung (the normally aerated region of the lung) is dramatically reduced, the volume of the total lung being considerably smaller at the same pressure (Gatinoni 2005). Therefore, TOP is found to be widely distributed along the PV curve where the values of TOP are uniquely determined by the patient physiology [Cheng et al. 1995].

Airways, when collapsed, obstruct the flow of air and thus impair ventilation to the alveoli resulting in inadequate pulmonary gas exchange. This creates alveolar shunt, resulting in arterial hypoxemia (low partial pressures of O_2 in the blood). To address the risk of atelectasis, it is common during mechanical ventilation to maintain a positive pressure in the lungs at the end of exhalation to maintain ventilation into regions of the lung that are prone to collapse. This positive pressure at the end of exhalation is known as positive end expiratory pressure (PEEP).

During simulation of a diseased lung in the model, alveolar compartments will have reduced ventilation because of extreme narrowing of the inlet (bronchial) airways. These correspond to alveolar compartments with high BIR with reduced ventilation and low alveolar volumes. To display the effective response of PEEP in the model, it is imperative therefore that the simulated diseased lungs include areas of reduced ventilation and low volumes. The major modifications implemented in the model are described below. The modifications are intended to align the simulated behaviour of the diseased lung with physiological reality by simulating the reduction in alveolar volumes and response to PEEP. The modifications were developed after detailed discussions with clinicians with the objective of closely matching the behaviour of the simulation model to *in vivo* observations.

2.3.4 Simulating alveolar volumes

Realistic *in silico* simulation of diseased lungs under pathological conditions such as ARDS requires 1) a reduction in ventilation to the alveolar compartments and 2) a reduction in volume

of the alveolar compartments (Pappert 1994, Hedenstierna 2000). In the model, reduction of ventilation can be configured by increasing the BIR of the individual alveolar compartments by a factor α. An example of a simulated diseased lung is given in Figure 2.12.
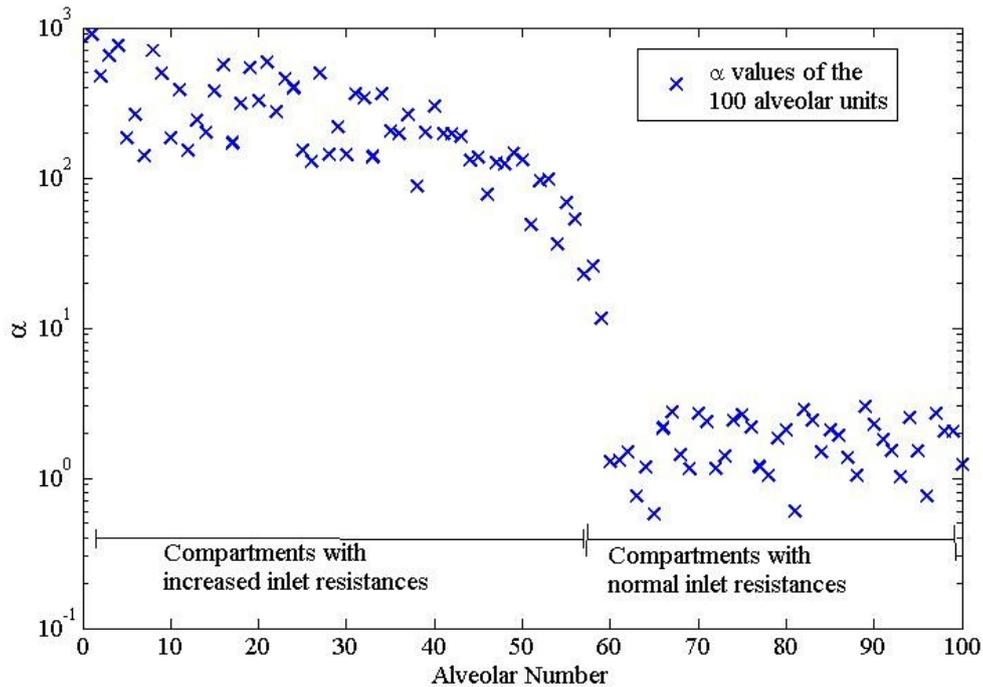
Figure 2.12: The distribution of α for alveolar compartments in an example of a simulated diseased lung.

Figure 2.13(upper) plots the change in volume of three selected compartments from the disease configuration given in Figure 2.12 – a compartment with high BIR (α = 900), a compartment with low BIR (α = 100) and a healthy lung (α=1). Figure 2.13 is plotted for a single breath, produced using identical ventilator settings to those employed in Figure 2.10. Relative to the healthy compartment; it is evident that the diseased compartments show a reduction in ventilation (equivalent to the reduction in change in volume (v) of the compartment).
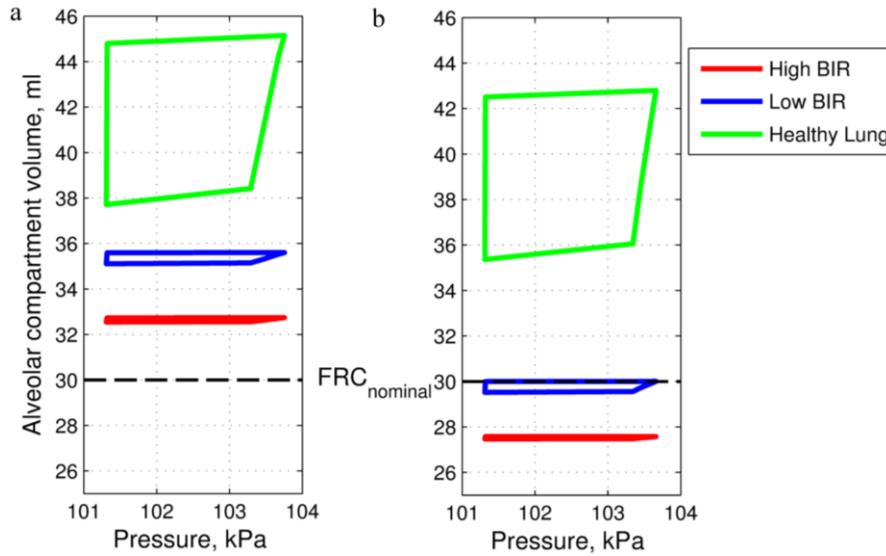
Figure 2.13: Pressure volume curves for individual alveolar compartments. (a) prior to modification and (b) post modification. The horizontal axis corresponds to the applied pressure to the compartment under the settings of Vtidal = 400ml, VentRate = 12 bpm and I:E = 1/3. The plots are given for alveolar compartments in one simulated lung (divided into 100 alveolar compartments) with different bronchial inlet resistances associated with – severe pulmonary obstruction (High BIR), mild pulmonary obstruction (Low BIR) and healthy lung.

Alveoli are generally considered as spherical shaped units (Levitzky 2007). The mechanism governing the change in size of alveolar compartments at small volumes can be attributed to Laplace's law (Lachmann 1992, Hickling 1998). According to Laplace's law the size of an alveolus is inversely proportional to the pressure inside the alveolar compartment. In addition (Lachmann 1992, Hickling 1998) suggest a threshold opening pressure (TOP) at low lung volumes, i.e. a minimum pressure that needs to be attained for collapsed alveoli (derecruitment) to open (recruitment). To incorporate the mechanism of Laplace's law and TOP, we propose a pressure p that is added to the pressure calculation (Figure 2.8) where

$$p = k \cdot P_{TOP} \cdot \left(1 - v_i / TLC_i\right)^n \ for \ i = 1, \dots, 100 \qquad [2.18]$$

P_{TOP} is typical threshold opening pressure of a diseased lung and is set to 20 cmH$_2$O (Lachmann 1992), v is the volume of the alveolar compartment i and TLC is the total lung capacity of the alveolar compartment i.

Eqn. 2.18 creates a distribution of pressures based on the volumes of the compartment (v). The pressure p would need to be overcome by the applied pressure during breathing so that the alveolar compartment can return to normal breathing. The parameters $k(=3)$ and $n(=4)$ can be used as further 'tuning' parameters that modify the volume-pressure relationship (Figure 2.14 red line). These parameters effectively control the compliance of the alveolar compartment at low lung volumes. The change in alveolar compartment volumes as a result of this modification is plotted in Figure 2.13(right). It can be observed that the new equation causes a further reduction in alveolar volume in the diseased lung.
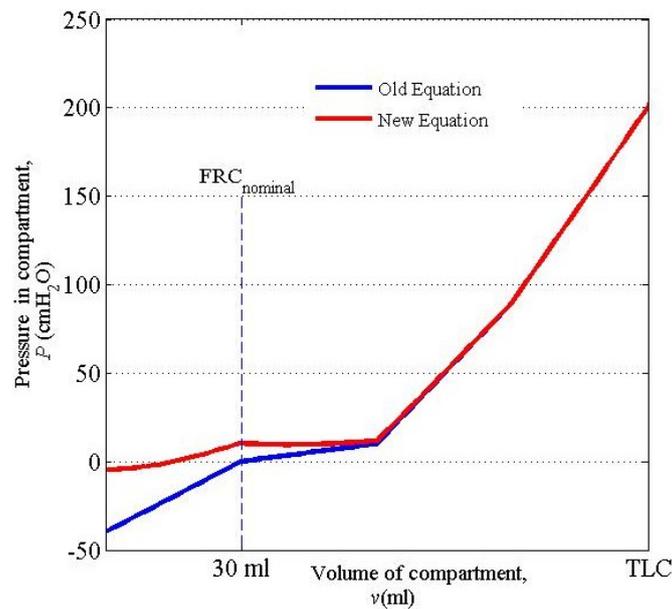
Figure 2.14: The modified static volume-pressure relationship in the model

The parameters of $FRC_{nominal}$, TOP, TLC, k and n, are independent variables that can be selected by the user to manipulate the individual volume-pressure relationship of any number of alveolar compartments in the model.

2.3.5 The effect of PEEP

In diseased lungs, the pressure in the airways can fall below the critical pressure required to keep them open. Positive End Expiratory Pressure (PEEP) is a commonly used mechanical ventilator setting that prevents the airways from falling to these critical low pressures.

In the model, the total inlet resistance (TIR) of an alveolar unit is determined by the following equation:

$$TIR_i = \alpha_i \cdot Rm_i \qquad\qquad [2.19]$$

where α is the integer vector of size n number of alveolar compartments, from Eqn. 2.16. Each element of α corresponds to the increase in bronchial inlet resistance of an alveolar compartment and is used to create an abnormal ventilation distribution. Rm is a vector of multipliers based on the volume of the alveolar compartment. Rm is determined by the equation:

$$Rm_i = 1 + 0.25 \cdot \left(\frac{FRC_{nominal}}{v_i} - 1\right) \qquad\qquad [2.20]$$

where v is the volume of the alveolar compartment i. Nominal resting volume ($FRC_{nominal}$) of each alveolar compartment is set to a value of 30ml, corresponding to the cumulative FRC of 3 litres for a lung with 100 compartments.
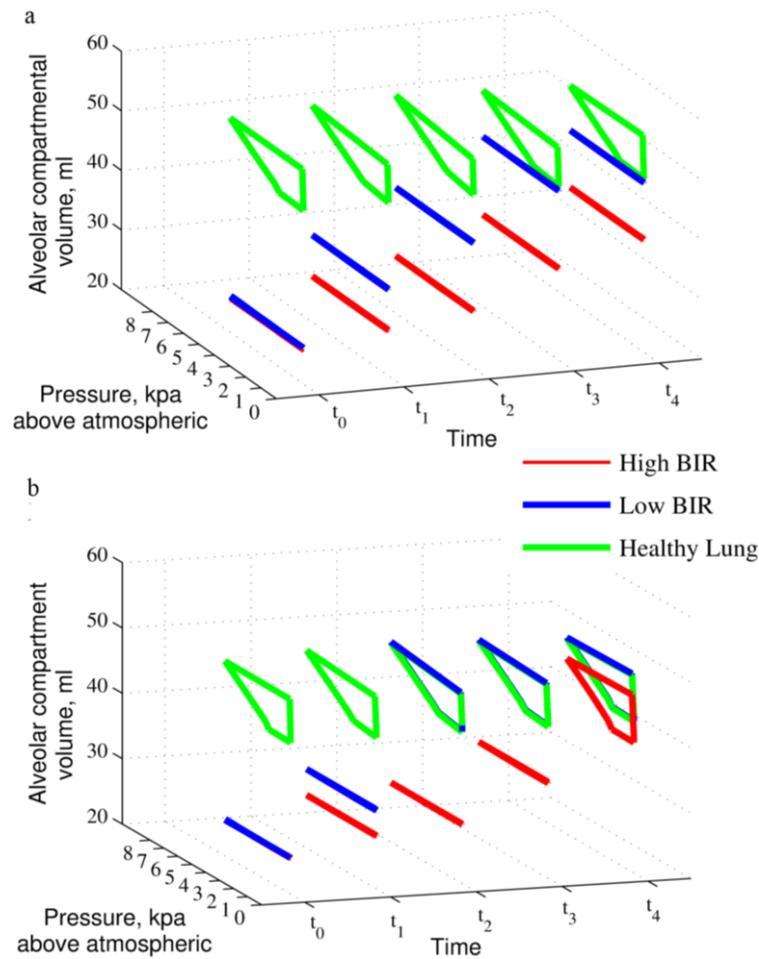
Figure 2.15: The effect of high PEEP on the alveolar compartment volume in (upper) under Eqn. 2.20 and (lower) after the modification. The upper figure clearly shows the diseased compartments (High BIR and Low BIR) fail to open with high PEEP. The simulated data is recorded from initial time, t_0, to t_4 (30 mins).

Figure 2.15 (upper) displays the effect of increasing PEEP from zero to 10 cmH_2O over a 30 minute period in the diseased lung. It can be observed that even though there is an increase in the volume of the alveolar compartment (v), diseased compartments remain poorly ventilated. This suggests that an increase in the volume of the compartment through high PEEP was not accompanied by an improvement in ventilation in lungs, contradicting clinical experience.

Further investigation reveals that under Eqn. 2.20, there is minimal change in total inlet resistance (TIR) (Figure 2.16 (upper)) with increasing PEEP.

To model the improvement in ventilation observed under PEEP, a mechanism needs to be introduced which would improve ventilation proportional to the alveolar volume. We propose a parameter iV, which represents an intermediate volume and is equal to $(\frac{4}{3} \times \text{FRC}_{\text{nominal}})$. Between $\text{FRC}_{\text{nominal}}$ and iV TIR passively decreases to be approximately equal to TIR of a healthy lung ($Rm_i = {}^{1}/{\alpha_i}$). Eqn. 2.20 is retained at alveolar volumes below $\text{FRC}_{\text{nominal}}$. The change is illustrated in Figure 2.16 (lower).

$$Rm_i = \begin{cases} 1 + 0.25 \cdot \left(\frac{\text{FRC}_{\text{nominal}}}{v} - 1 \right) & v \leq \text{FRC}_{\text{nominal}} \\ \dfrac{(iV \text{-} v)}{(iV \text{ - } \text{FRC}_{\text{nominal}})} & \text{FRC}_{\text{nominal}} < v < iV \\ \dfrac{1}{\alpha_i} & v \geq iV \end{cases} \qquad [2.21]$$

The effect can be summarised as follows: as the volume of the compartment increases above $\text{FRC}_{\text{nominal}}$, the bronchial inlet resistance of the compartment is reduced passively to the resistance found in a healthy lung, which improves the ventilation to the compartment. Figure 2.15 (lower) displays the results of changing PEEP and in comparison to Figure 2.15 (upper); there is now clear improvement in ventilation with a higher PEEP.
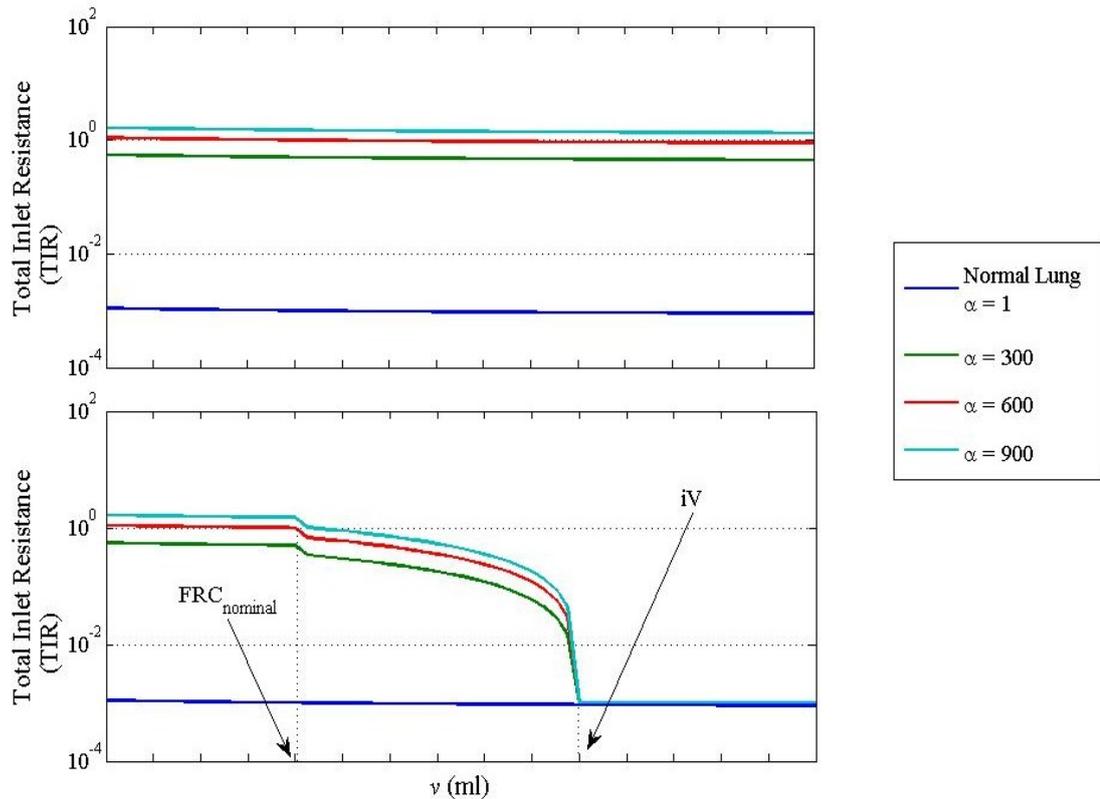
Figure 2.15: The change in total inlet resistance (TIR) as volume (*v*) of the alveolar compartment increases. TIR=α.*Rm* .The upper figure shows change in TIR due to Eqn. 2.20 while the lower figure displays the result of the modification (Eqn. 2.21).

2.3.6   Clinical validation of the modified model

The modified model was validated against clinical data previously used to validate the original model in (McCahon 2006). Using the "dynamic validation" dataset of (McCahon 2006), the modified model was matched to initial patient settings. The model was then subjected to the same change in inspired oxygen fraction, ventilatory rate and tidal volume as the patients (Appendix A). The error in predicting the change in arterial oxygen tension ($|\Delta PO_2|$) was used to compare the two models and the data. Table 2.2 summarises the patient and model response. The results show a strong correlation between the original model's response and the modified

model's response (correlation coefficient $r = 0.98$) and patient data against modified model's response ($r = 0.93$). It should be noted that the proposed modifications on the response to PEEP would not be observable here as PEEP was fixed during the clinical trial.

The small differences between the two model predictions could be due to the implementation of the collapsed regions in the lung in the modified model, and/or the difference in the duration of time post-intervention at which the model response is recorded. The modified model response is set to be recorded when mean PaO_2 reaches equilibrium (i.e. variation in PaO_2 is less than 0.5kpa over one minute period) or at the end of 25 minutes of simulation time. The time taken before recording the post-intervention data for the patient or the original model has not been mentioned in the paper. This highlights a common problem with current model validation approaches based on comparing one set of model responses to another. In the next chapter, we explore more sophisticated computational approaches to the model validation problem.

Patient	Pre-intervention	Post -intervention						
		Clinical data			Model Predictions			
					Original Model		Modified Model	
	PaO$_2$ measured, [kPa]	PaO$_2$ measured, [kPa]	$|\Delta PO_2|$, [kPa]		PaO$_2$, [kPa]	$|\Delta PO_2|$, [kPa]	PaO$_2$, [kPa]	$|\Delta PO_2|$, [kPa]
1	10.23	13.78	3.55		12.51	2.28	12.90	2.67
2	19.38	10.97	8.41		11.44	7.94	11.04	8.34
3	9.6	11.84	2.24		11.22	1.62	11.14	1.54
4	22.23	16.92	5.31		15.62	6.61	15.52	6.71
5	21.66	16.8	4.86		15.40	6.26	14.59	7.07
6	12.6	23.15	10.55		24.15	11.55	22.62	10.02
	Mean		5.82			6.04		6.06
	Standard deviation		3.10			3.69		3.29

Table 2.2: Comparison table of original model predictions and modified model predictions with clinical data.

Chapter 3: Validation of Pulmonary Physiology Models

This chapter begins with a discussion of the current practices followed in physiological model validation and motivates the need for improved techniques. The optimisation algorithms which form the basis of the proposed validation approach, namely Sequential Quadratic Programming (SQP), Mesh Adaptive Direct Search (MADS), Genetic Algorithm (GA) and Differential Evolution (DE), are described in detail. In the process of validation, the pulmonary physiology model is subjected to simultaneous variations in the values of multiple physiologically relevant uncertain parameters with respect to a set of specified performance criteria. These performance criteria are based on expected levels of variation in arterial blood gas values found in the patient population and are evaluated using numerical simulations.

3.1   Current methods

Mathematical modelling and computational simulation are becoming increasingly important tools in many fields of medicine where *in vivo* studies are expensive, difficult, or impractical (Hardman 2006).

Many applications of modelling (Joyce 1996, Hotchkiss 1994, Farmery 2002) involve the investigation of specific scenarios involving relatively few equations, mostly based on basic, recognised physical laws and relationships. As such their validation is also simplistic, where model output is assessed in comparison to established knowledge and expert opinion. As the models increase in scale, by integrating multiple smaller models (e.g. a training simulator representing several organ systems), expert opinion alone cannot deem the model valid. In such

scenarios it is common to compare the model predictions to clinical data (Liu 1998, Hardman 1998, Kwok 2003, Chiari 1997). Current physiological validation approaches can be divided broadly into the following categories:

- Visual Inspection: At its most basic, model validation can simply be the demonstration of the ability of the model to represent standard physiological behaviour (Swanson 1983, Joyce 1996, Farmery 2002). For example, (Farmery 2002) presents a validation of a model of oxyhaemoglobin desaturation, by simply plotting the effect of several relevant parameters and disease conditions on the desaturation curves.

- Qualitative Analysis: Model outputs are plotted alongside the measured clinical data, where qualitative observations are made regarding the agreement between model predictions and the measured outputs (Liu 1998, Hahn 1993, Athansiades 2000). Simple observational assessment such as "the predictions are found to be in good agreement with the data" or "as can be seen, the model closely matches the empirical results" (Norman 2006) is common.

- Quantitative analysis: In this case, the model prediction is numerically compared to the measured clinical data. Statistical indices such as prediction error (e) and correlation coefficient (r) are computed to describe the quality of the model prediction. A brief introduction to these indices is given below. The simplest comparison between a vector of measured values (x_m) and a vector of predicted values (x_p), both of size n, could be stated as the error (e) (Chiari 1997, Suki 1998):

$$e = \left| x_m - x_p \right| \hspace{4cm} [3.1]$$

Accurate model prediction would ideally seek to make the magnitudes of the vector e go to zero, i.e. $\|e\| \rightarrow 0$. The use of linear regression analysis has become common in assessing the

validity of the predictions of a model (Hotchkiss 1994, Neumann 2001). The first step in linear regression is to calculate the values of the vector e, also called the residuals, i.e. difference between the measured and the predicted data. By minimising the sum of the squares of the *residuals,* a line of *best fit* can be drawn representing a *least squares fit.*
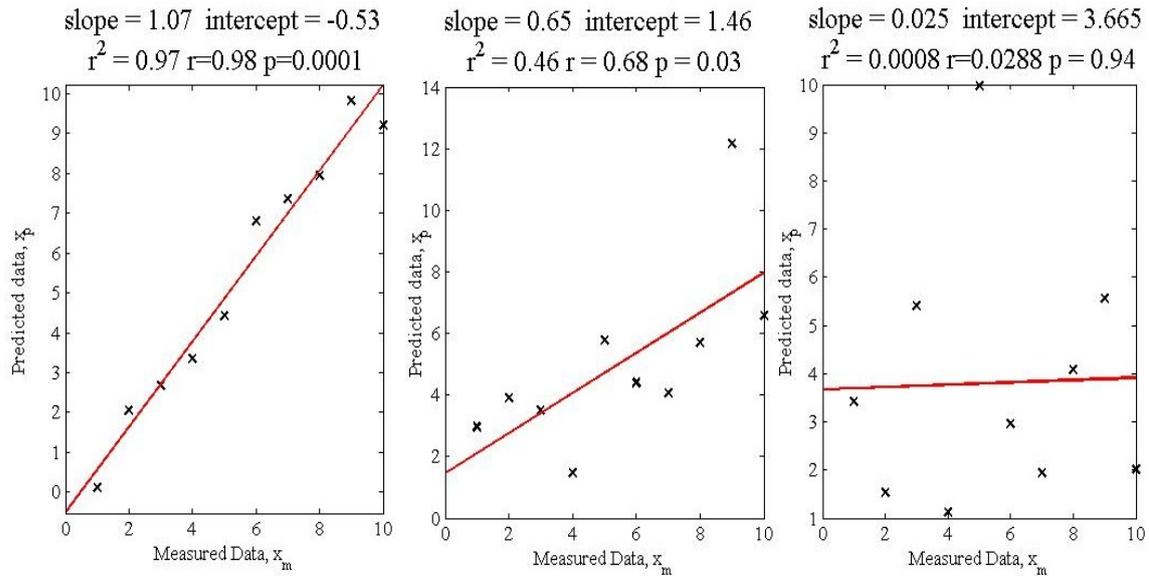
Figure 3.1: Indices of linear regression for three different examples.

A measure of the *goodness* of the fit, also known as the correlation coefficient (or regression coefficient) is denoted with the symbol r. The value of r indicates the strength and direction of the relationship between x_m and x_p. It is common to give the coefficient of determination r^2 instead of r, which expresses the proportion of variance between x_m and x_p. The value of r^2 is a fraction between 0 and 1. A lower value of r^2 indicates that the data is further scattered away from the least square fit. An r^2 value of 1 along with line of best fit with *slope* =1 and *intercept* =0, would indicate a perfect agreement between the predicted and measured data. Figure 3.1 gives three examples with values of r^2, *slope*, and *intercept*. Another value to be considered along with the correlation coefficient r is the probability p where p describes the reliability of r, i.e. if p is very small, then the correlation r is said to be significant.

In (Bland 1989, 1999), Bland noted that a high correlation coefficient does not necessarily indicate a strong agreement between x_m and x_p. Perfect correlation could be along any line of best fit but perfect agreement has to be along the line of equality ($x_m = x_p$). Factors such as *bias,* where there is a consistent difference between measured data and model predictions, could still show a strong correlation but ignore the bias. This led to Bland proposing the Bland-Altman plots (example in Figure 3.2), which plot the difference between x_m and x_p against their mean values ($\frac{x_p - x_m}{2}$). This method has since become a common index of validation. (Hardman 1998, Hardman 1998, Kwok 2003, Wilson 2009).
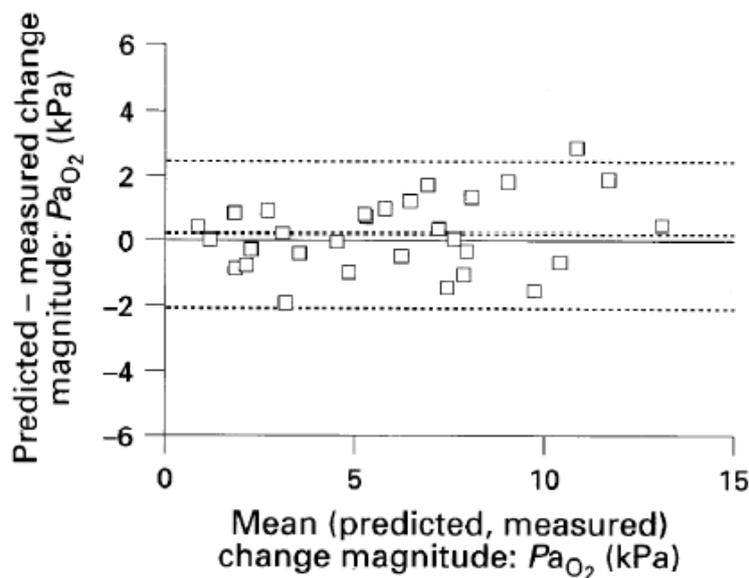
Figure 3.2: Comparison of measured vs predicted data in outputs. Reproduced from (Hardman 1998).

Figure 3.2 shows an example of a Bland-Altman plot, presented in (Hardman 1998) showing the errors in predicting the magnitude of changes in Pa_{O2}. The solid line in the figure represents the ideal line of error (no error). The dotted lines represent *bias* and 95% limits of agreement of predicted and measured values for magnitude of change in Pa_{O2}. The *bias* represents the consistent difference found between x_p and x_m. The 95% limits of agreement are calculated

assuming a normal distribution with *bias* ± k × SD, where SD is the standard deviation of $(x_m - x_p)$. As the number of samples is low, the value of k is determined using the *t*-distribution table giving k = 2.04.

3.1.1 Disadvantages of current methods

The previous section shows that the current methods of validation of models in many fields of medicine are heavily reliant on the expertise of the model designer and on comparisons against clinical or lab data. However, there are numerous unresolved concerns with the current methods, which we attempt to summarise below.

When computational models of single patients have been validated against heterogeneous data sets using methods such as Band Altman plots, model validity has been generally accepted if the model responses lie within 95% of population observations. For any given section of the data which is outside the 95% limits of agreement (i.e. outliers), the plot gives little information regarding the nature of the error. The outliers could represent significant sub populations unrepresented by the model and ignoring these outliers can risk inaccurate future predictions from the model.

Apart from the possible inaccuracies in observing the outputs of a model, errors can also exist in the values of internal variables. Non-measurable parameters in physiological simulators can be set at values derived as certain nominal conditions within some uncertainty bounds (Eldridge 1996, Hardman 2001). For example, the simple representation of the cardiovascular system in the Nottingham Physiological Simulator (NPS) uses a fixed value of cardiac output and oxygen consumption during a simulation. However these values are known to vary significantly within a given patient population, even under normal conditions (Stead 1945, Hickam 1948). Variations in these parameters naturally exist in the human population and as such the

performance of the model can be dependent on these parameters. Small deviations between the assumed values of these parameters and the true values could potentially exaggerate errors between model predictions and measured data.

As an example, consider Eqn. 3.2 used within the implementation of NPS (Hardman 1998) to describe the pressure (P_{O2})-saturation (S_{O2}) relationship for the O_2 dissociation curve (described earlier in Eqn. 2.2 and 2.3)

$$S_{O2} = \left(\left(\left(P_{O2}^3 + 150 \cdot P_{O2} \right)^{-1} \times 23400 \right) + 1 \right)^{-1} \qquad [3.2]$$

S_{O2} is the saturation of the haemoglobin in blood and P_{O2} is the partial pressure of oxygen in the blood. Eqn. 3.2 applies to a standard O_2 dissociation curve in Figure 2.3 (at pH=7.4 and temperature=37°C). As suggested by (Severinghaus 1966), P_{O2} has been determined with appropriate correction factors in Base Excess (BE), temperature (T) and pH (7.5005168 = pressure conversion factor from kPa to mmHg):

$$P_{O2} = 7.5006168 \cdot P_{O2} \cdot 10^{[0.48(pH\text{-}7.4)-0.024(T\text{-}37)-0.0013 \cdot BE]} \qquad [3.3]$$

Multiple uncertainties in many interconnected sub-systems of a complex system could propagate through the system, the cumulative effect of which is difficult to predict. A level of uncertainty of only ±5% in each of the input parameters of Eqn. 3.2-3.3 could lead to a possible ±14% variation in the resulting prediction of S_{O2}. It should also be noted that the calculation of BE currently also relies upon determining the pH. Clearly, if Eqn. 3.2 was used as part of a more complex nonlinear simulation model involving many other such equations, each with its own set of uncertain parameters, the possible effects on overall model predictions could be highly significant.

In the case of unmeasured parameters, the simplest strategy involves manually uploading nominal values to these parameters (Hinds 1984, Hardman 1998, Eldridge 1996). Some studies use parameter estimation techniques to identify the values of the unknown parameters. In other cases where model parameters are directly unavailable from the patient data, studies have used statistical techniques such as Bayesian estimation (Rutledge 1993), regression (Liu 1998) etc to match their model parameters to steady state patient conditions. However these methods are insufficient when representing dynamic models with time varying parameters.

Without other more rigorous procedures for model validation and verification, the wider clinical exploitation of *in silico* simulation models will continue to be strongly limited, since doubts about model validity negatively affect the clinical applicability of any predictions arising from simulation studies. 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.

The next section introduces optimisation based methods that have shown considerable success in assessing the robust stability and performance of computer simulation models of safety-critical applications (Menon 2007, Martin 2008). The underlying principles of optimisation-based model validation and structure of some common optimisation algorithms are discussed in detail. This should enable the reader to easily interpret the final section of the chapter where the pulmonary physiology model is subjected to simultaneous variations in the values of multiple

physiologically relevant uncertain parameters with respect to a set of specified performance criteria.

3.2 Model validation using optimisation methods

Recent research in the field of simulation-based analysis of flight control systems (Fielding 2002, Menon 2006) and systems biology (Kim 2006, Kim 2010, Morohashi 2002) has highlighted the usefulness of the engineering concept of *robustness* for validating complex *in silico* simulation models. 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 (Doyle 2002). This concept may be applied for the purposes of physiological model validation as follows.

Consider a situation where the outputs of a model (or of several different models) show a good match to the available experimental data for a given set of model parameters. With respect to expected levels of uncertainty in the model parameters, the robustness of the responses of the model could be evaluated by using analytical or computational approaches 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.

Robustness analysis is now a well-established research area in the field of systems and control engineering, and many approaches have been developed in recent years to quantify the robustness of complex nonlinear systems (Fielding 2002, Morohashi 2002, Menon 2006). For example, prior to flight tests, flight control engineers dedicate a significant amount of resources

to simulation-based analysis of the effect of uncertainty in an aircraft's mass, inertia, centre of gravity positions etc on the performance of flight control system designed using a "nominal" model. This type of analysis is expensive but crucial to establishing the safety of the system, which has provided strong motivation for the development of novel techniques to conduct the analysis in the most reliable yet efficient way possible. Recent research in this area has proposed the use of global optimisation algorithms for a number of different applications of robustness analysis, such as aircraft stability (Menon 2006), robustness of cellular level biochemical models (Kim 2006), model validation of a jet engine (Martin 2008), and attitude control of satellites (Wang 2010).

The first step in the optimisation-based model validation approach is to formulate the optimisation problem (OP). A general formulation of an OP is described as follows:

$$\text{Maximise} \quad f(x) \qquad [3.4]$$
$$\text{subject to} \quad x \in S$$

The objective function $f : \mathbb{R}^n \to \mathbb{R}$, is a scalar function that we wish to maximise. The decision vector $x = (x_1, x_2, ..., x_n)^T \in \mathbb{R}^n$ is a vector of independent variables $(x_1, x_2, ..., x_n)$ known as the *decision variables*. S is the set of all possible solutions of x such that $S \subseteq \mathbb{R}^n$. The OP can be stated as finding a vector x from S that returns the maximum value of the objective function f. If $S = \mathbb{R}^n$ then the OP is referred to as an unconstrained maximisation problem. In a constrained optimisation problem, OP of Eqn. 3.4 is subject to the following constraints:

$$\text{Subject to} \quad g_j(x) \leq 0, \quad j = 1, ..., v \qquad [3.5]$$
$$g_j(x) = 0, \quad j = v + 1, ..., w \qquad [3.6]$$
$$x^L \leq x \leq x^U \qquad [3.7]$$

The decision vector x is restricted to be within a lower bound x^L and upper bound x^U which defines the *decision space*. The function $g_j(x)$ is called the constraint function. A solution x that does not lie in the decision space or satisfy $g_j(x_i)$ is called an *infeasible solution*. Conversely if a solution x exists in the decision space and satisfies $g_j(x_i)$ then it is called a *feasible solution*. S is thus the set of feasible solutions.

In model validation problems, the objective function $f(x)$ is chosen to find the worst-case combination of parameters x^* that would maximise the difference between model's response and the nominal response. The objective space may contain multiple local optima but the true worst-case combination is given by the global optimal of the objective function $f(x)$. The next step is the selection of an optimisation algorithm that can be used to drive the search in the decision space towards x^*.

Many different classes of optimisation algorithms can be found in literature. Some of these algorithms use the gradient information of the objective function to decide the search direction, while others rely only on the objective function value. The search space can be convex or non-convex which can decide if the optimisation algorithm would provide a local or a global solution, respectively. A local optimisation algorithm will find the global solution in a convex space but may only find a local solution in a non-convex search space. Most real world problems have a search space which is non-convex. Therefore, global search algorithms are required to reliably find global solutions. The performance of any given optimisation algorithm is dependent on the problem and no unique algorithm can guarantee convergence to the global solution with a reasonable amount of computational expense. Therefore, a number of different algorithms are applied and compared in this study, in order to discover the most promising methods for the problem of model validation of physiological systems, and to cross-validate the worst-case analysis results. All the algorithms considered in this section are subject to general

termination conditions on improvement in objective function value (TolFun) and change in decision vector x (TolX), which are explained in Appendix B. The algorithms terminate if TolFun$<10^{-6}$ or TolX$<10^{-6}$. The individual algorithms and any other specific termination conditions associated with them are described below.

3.2.1 Sequential Quadratic Programming

Sequential Quadratic Programming (SQP) is one of the most effective classes of local algorithms for addressing nonlinear constrained optimisation problems (Han 1976, Powell 1978). It can be considered as a generalisation of Newton's method for unconstrained optimisation in that it decides the next step from the current point by minimising a sequence of quadratic programming sub-problems. The SQP algorithm requires first and second derivative information of the objective and the constraints function. With problems where the local gradient information of the objective function cannot be determined analytically, SQP algorithm defines a quasi-Newton matrix by making a quadratic approximation of the Lagrangian function and approximating the Hessian matrix. This enables the SQP algorithm to replace the objective function with the quadratic approximation and replace the constraints with linear approximations. The Hessian matrix is updated every iteration using the standard Broyden-Fletcher-Goldfarb-Shanno (BFGS) formula. The present study uses the SQP implementation provided in the optimisation toolbox of MATLAB (Mathworks 2000).

3.2.2 Mesh Adaptive Direct Search (MADS)

MADS is a relatively new local optimisation algorithm (Audet 2006) that addresses the problem of minimising a non smooth function where derivate information of the objective function is unavailable. It extends the Generalised Pattern Search (GPS) algorithm (Lewis 2000), by enabling the exploration of the search space in an infinite number of directions. MADS has been

shown to be very effective relative to GPS in reaching global optima (Audet 2006). The algorithm works by searching for the best point around a current point, within a *mesh*, based on the evaluation of the objective function. If the algorithm finds a point in the mesh that improves the value of the objective function (*successful*), then the new point becomes the current point and the mesh size increases. If the current point does not improve the objective function (*unsuccessful*) then the mesh size is reduced and the method is repeated. The description of the method as implemented in this study is presented below.

MADS is an iterative algorithm that attempts to locate a minima of the function *f*(x) over its feasible region by evaluation 2*n* trial points (where *n* is the number of parameters being considered) at every iteration. Each of these trial points lies on the current *mesh*, constructed from a vector D_k of 2*n* direction vectors, scaled by a mesh size parameter m_k.
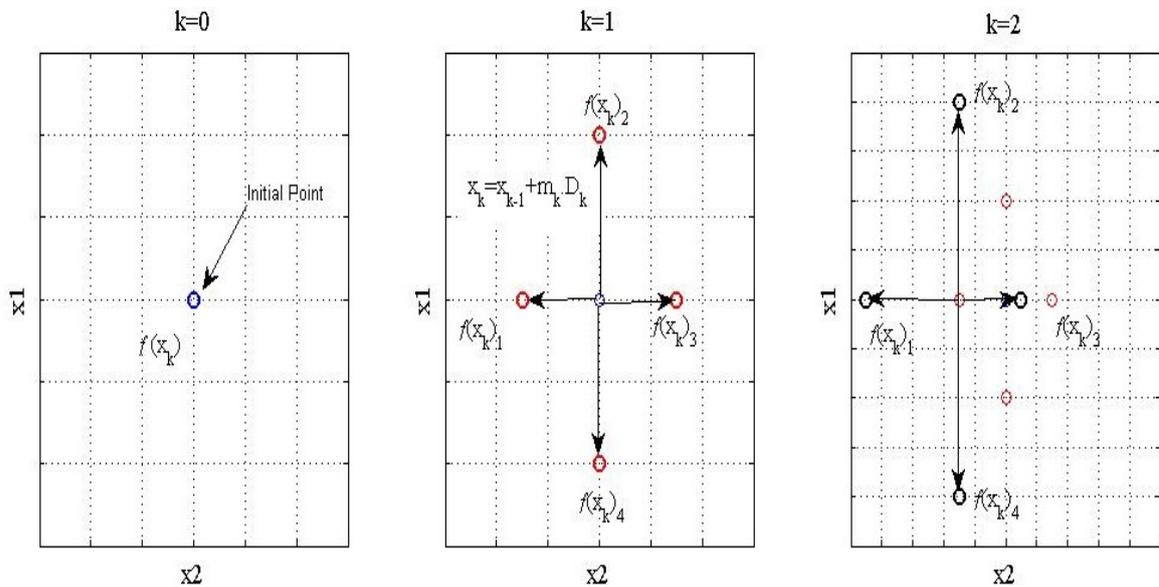
Figure 3.3: MADS poll step example for number of parameters n=2. The third plot illustrates the process at k=2 if at k=1, $f(x_k)_1 < f(x_k)$ for all trial points x.

1) Initial point x_o is provided to the algorithm and the objective function value *f*(x_0) is evaluated. At iteration k, the current mesh is defined by

$$M_k = \bigcup_{x \in S_k} \{x + m_k \cdot D_k\} \qquad\qquad [3.8]$$

where S_k is the set of points where the objective function f has been evaluated by the start of iteration k. Defining the mesh as a union, all visited points lie on the mesh and new trial points can be selected around any of them using directions in D.

2) *Search*: From the current best point, a search is conducted. If the search returns a better value of the objective function value, than this becomes the centre of the new mesh. The *poll* step is ignored and a new mesh is created. The current implementation conducts the SQP based local search, mentioned in the previous section.

3) *Poll:* If the *search* step fails to improve the current mesh, then a poll is run around the current best mesh point. The mesh points are evaluated. D_k is a vector of size 2n and is selected randomly by the algorithm. If an improved mesh point is found then the poll is *successful* and m_k is doubled for the next iteration (Figure 3.3). If the poll fails to improve the objective function value then the poll is *unsuccessful* and m_k is halved.

4) Termination: The algorithm is terminated if the function fails to improve by 10^{-6} between iterations, if the mesh size becomes smaller than 10^{-3} or if the number of total function evaluations exceeds 600.

3.2.3   Genetic Algorithms (GA)

GA's are general purpose stochastic search and optimisation procedures based on evolutionary principles (Goldberg 1989). Evolutionary algorithms aim to generate a population of fittest candidates by implementing evolutionary concepts such as selection, mutation, recombination etc. Due to their ease of application in problems with large and small parameter search spaces, GA's have become a popular search and optimisation technique in many fields of science and engineering (Mitchell 1998, Menon 2006). With their probabilistic and parallel nature, GA's are

generally capable of converging to the global optimum even in highly non-convex parameter spaces – convergence can be slow, however, and global algorithms like GA's typically require much longer computation times than local gradient-based methods (Fleming 2002).

In GA's, a randomly selected population of candidates (1st generation) undergoes a repetitive process of reproduction, where selection is based on the value of the objective function (also called the fitness). Every generation, the best candidates from the previous generation (elitism) and candidates obtained through mutation and crossover, recombine to form a new population. The average fitness of the individuals in the population is expected to increase as strong individuals are protected and combined with one other and weak individuals are discarded.

The specific method implemented here is described in detail below.

1) Encoding: A double vector encoding scheme (Mathworks 2005) is used to encode the real values of the parameters. Each parameter vector is called a chromosome and each parameter (element) in the vector is called a gene.

2) Initialisation: The GA search starts from an initial random number of candidates for a population size n. The fitness of each chromosome is evaluated using the objective function. The population loses diversity over the search space if the population size is too low, which can reduce the quality of the final solution and increase the duration of algorithm computation (Goldberg 1989). For the work conducted in this thesis, the population size n is fixed at 30.

3) Selection: The candidates in the current generation are selected as parents to form future generations based on the selection scheme and their fitness. The work done in this thesis utilises a roulette wheel selection scheme, where each segment of the notional roulette is assigned a chromosome, with the size of the segment proportional to the fitness of the

chromosome. This ensures that the stronger candidates have a higher probability to influence the next generation.

4) Elitism: The GA implementation incorporates an elitist strategy which selects the two chromosomes with the best fitness from the population into the next generation. This ensures that the average fitness of the population will not degrade over the generations.

5) Crossover: During crossover (Figure 3.4), a new individual is created by combining the genes from two parent chromosomes which are selected at random from the parent population. A simple single point crossover is employed where at a randomly selected length; the two chromosomes are split and merged with each other. The crossover fraction is set to 0.8, i.e. 80% of the new population is generated through crossover.
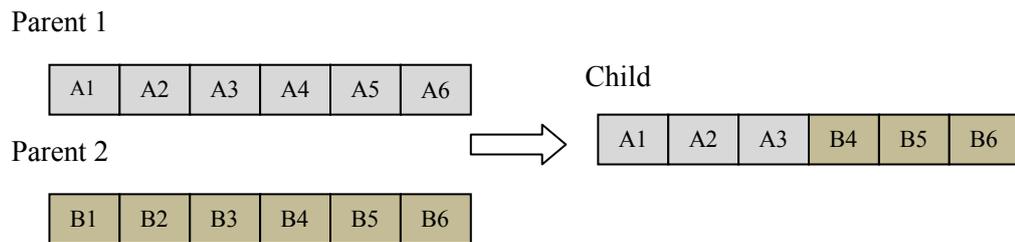
Figure 3.4: The GA crossover operation example

6) Mutation: Mutation introduces random variation in the chromosome of the mutation population (Figure 3.5). After the elitism and the crossover operators, the remaining members are populated using the mutation operator.
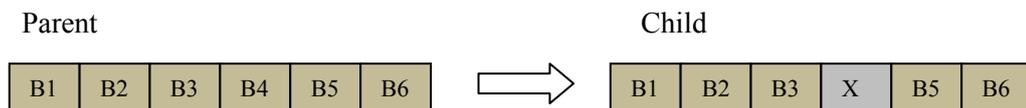
Figure 3.5: The GA mutation operation example

The problems considered in this thesis have search spaces that are either constrained, bounded or both. Therefore, an adaptive mutation scheme (from Mathworks 2005) is utilised that randomly generates the direction and the size vectors of the mutation that would satisfy the constraints and the bounds.

7) Termination: The maximum number of generations is set to 20. In addition, an adaptive criterion is used to terminate the algorithm if the improvement in the solution over successive generations is less that 10^{-6} (see appendix B).

3.2.4 Differential Evolution (DE)

The second evolutionary optimisation algorithm considered in this study is DE, introduced by Storn and Price in (Storn 1997). This method belongs to the same class of evolutionary techniques as GA's, but unlike GA's it does not require a selection operator or an encoding scheme. The DE method has recently been applied in several problems in different fields of engineering design with promising results (Vesterstrom 2004, Menon 2007), and was found to provide better performance when compared with other evolutionary algorithms in solving different numerical problems.

DE generates new parameter vectors by adding the weighted difference between two vectors to a third vector, to form a mutated vector $\bar{x}_{new}^{G+1}$. Then under crossover, a target vector is mixed with this mutated vector to form a trial vector. Finally during selection, if the trial vector improves the objective function value compared to the target vector, then it replaces the target vector in the next generation. The process is continued until any of the termination criteria is met (see below). The specific implementation of DE utilised here is the "DE/rand/1/bin" (Storn 1997). The steps are described are described in detail below.

1) Initialisation: The initial vector population is

$$x_i^G, i = 1,2, \dots, N \qquad [3.9]$$

where N is the population for each generation G. The initial vector population is chosen randomly, based on the assumption of a uniform probability distribution to cover the entire parameter space. Each x_i^G is an *n*-dimension vector equal to the number of parameters being considered in the objective function. The population size N is fixed at 30 in this implementation.
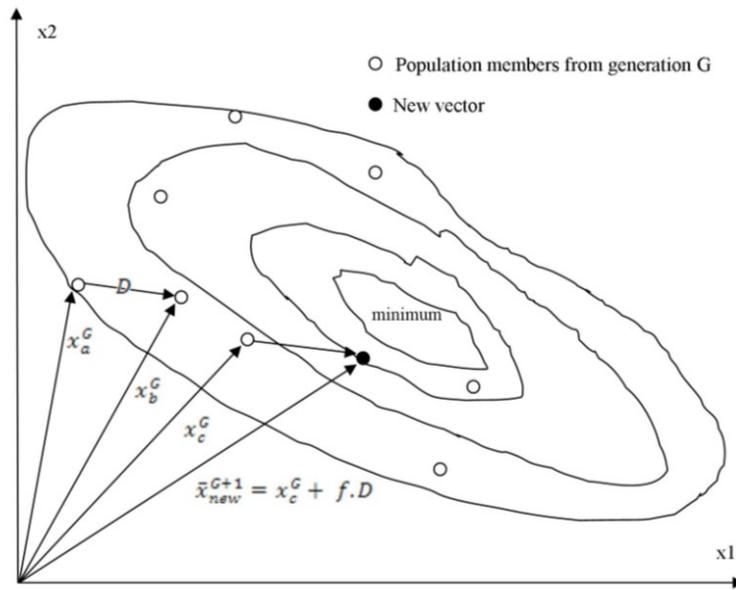
Figure 3.6: Example of two dimensional objective function showing the contour lines and process for generation of the mutated vector x_{new}^{G+1}

2) Mutation: For each target vector x_i^G, a mutant vector is generated according to

$$\bar{x}_{new}^{G+1} = x_i^G + f \cdot D \quad where\ D = (x_b^G - x_a^G) \qquad [3.10]$$

Figure 3.6 illustrates a simple two dimensional example of the process. The difference vector D determines the search direction. The mutation scale factor f is a real valued

number in the range [0, 1], (fixed at 0.8 in this study) and determines the step size in that direction.

3) Crossover: Using a crossover factor CF , the mutant vector x_{new}^{G+1} and a target vector x_i^G a trial vector is formed where

$$x_{trial}^{G+1} = \begin{cases} x_i^G & if\ R > \text{CF} \\ \bar{x}_{new}^{G+1} & otherwise \end{cases} \qquad [3.11]$$

In Eqn, 3.11, R is a random number generated between [0, 1] by the user. Here, CF is fixed at 0.8.

4) Selection: To decide whether or not the trial vector x_{trial}^{G+1} becomes a member of generation G+1, it is compared to the target vector x_i^G. If the trial vector improves the objective function value then it replaces the target vector in the next generation.

5) Termination: The above steps will continue until the termination requirement is met. The same adaptive termination criterion mentioned in the previous section is utilised.

Other variants and applications of DE can also be found in (Storn 1997, Price 2005).

3.3 Validation of pulmonary physiology simulator for clinical applications

3.3.1 Introduction

In this section, a robustness analysis approach is used to compute worst-case deviations of arterial blood gas predictions from physiologically realistic values due to the effects of realistic levels of parametric uncertainty in a high fidelity simulation model of cardio-pulmonary gas

dynamics (Hardman 1998, Hardman 2001, McCahon 2008). The effect of simultaneous variations in multiple uncertain parameters can be analysed by formulating the problem as a multivariable constrained optimisation problem. A detailed comparison of the effectiveness and computational efficiency of a number of different global optimisation algorithms in solving this problem has been carried out, and results are compared with standard statistical approaches (Monte Carlo sampling).

Unlike the traditional approach to model validation, the aim is not to validate a particular setting of the model against data from an individual patient, but to quantify its robustness (i.e. its ability to always produce realistic results for all expected levels of uncertainty in internal model parameters). The intent is to produce a simulator that has been validated across the expected range of variation that would be found in multiple patients, thus representing 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).

3.3.2   Modelling of uncertain parameters

There are a large number of uncertain parameters that could be considered in a complete analysis of a pulmonary system model which includes both static and time-varying parameters arising from uncertainty within and between individual subjects. To illustrate the proposed approach, a representative subset of these parameters has been chosen, representing the haemoglobin level (Hb), cardiac output (CO), oxygen consumption (VO_2), respiratory quotient (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 3.1. 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 amongst the general patient population.

3.3.3 Formulation of the model validation problem

The model validation problem considered in this study is the following: for the realistic levels of uncertainty on key model parameters defined in Table 3.1, guarantee that the model responses for PO_2' (current arterial O_2 pressure, kPa), PCO_2' (current arterial CO_2 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 PO_2', PCO_2' and pH' were defined based on clinical experience and are shown in Table 3.2.

Uncertain Parameters	Nominal value and uncertainty range	[lower bound (x^L), upper bound (x^U)]
Haemoglobin content in blood, Hb [gm l^{-1}]	$140 \pm 15\%$	[119,161]
Cardiac output, CO [ml min^{-1}]	$5000 \pm 10\%$	[4500, 5500]
Oxygen consumption, VO$_2$ [ml min^{-1}]	$250 \pm 20\%$	[200, 300]
Respiratory quotient, RQ	$0.8 \pm 12\%$	[0.704, 0.896]
Temperature, T [ºc]	$37.2 \pm 0.5\%$	[37.0, 37.4]

Table 3.1: Nominal values, allowable uncertainty ranges and the resultant lower and upper bound for the selected uncertain parameters in the model.

Predicted Parameters	[lower bound, upper bound]
Current arterial O_2 pressure, PO_2' [kPa]	[9,15]
Current arterial CO_2 pressure, PCO_2' [kPa]	[3.5,7.5]
Current plasma, pH'	[7.3,7.5]

Table 3.2: Allowable upper and lower bounds for model predictions

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

$$\max_x \quad F(x) \quad where \; F(x) = \|\Delta_{O2}\|_\infty + \|\Delta_{CO2}\|_\infty + \|\Delta_{pH}\|_\infty \qquad [3.12]$$

$$subject \; to \; x^L \le x \le x^U$$

$$\Delta_{O2} = \frac{PO_2 - PO_2'}{PO_2}, \quad \Delta_{O2} = \frac{PCO_2 - PCO_2'}{PCO_2}, \quad \Delta_{O2} = \frac{pH - pH'}{pH}, \qquad [3.13]$$

In Eqn. 3.12, x is the vector of uncertain model parameters, x^L and x^U are vectors of lower and upper bounds, respectively, for the uncertain parameters (see Table 3.1) and $F(x)$ is the *objective function* which is maximised by the optimisation algorithm. In Eqn. 3.13, PO_2, PCO_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 of size n (i.e. $\max_{1<i<n}|a_i|$), commonly known as the infinity norm. Thus, $\|\Delta\|_\infty$ measures the normalised absolute magnitude of the largest difference between the "nominal" model prediction and the current model prediction over the course of the simulation time window, normalised using the nominal output.

After a few preliminary tests, a simulation time-window of 10 minutes was chosen, in order to ensure that all model outputs had settled down to steady-state values before the objective function was evaluated. This is decided to be an adequate choice, but it cannot be guaranteed that all the model outputs have indeed settled down for all possible combinations of inputs. In this case we have chosen to normalise 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 weight the different variables in situations where deviations in one are considered more problematic or interesting than another. With the above formulation, the optimisation 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 3.7.
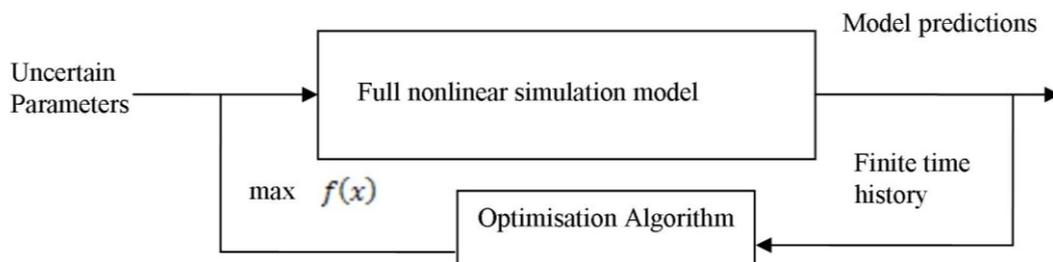
Figure 3.7: Optimisation-based model validation framework

In order to ensure that the globally optimal solution to the above problem has been found, several different optimisation algorithms are employed. Once the optimal solution has been found, the corresponding "worst-case" model predictions are considered with respect to the allowable ranges specified in Table 3.2. 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, if it corresponds to a particular condition which would be likely to produce unexpected responses, etc.

In this study, a gradient-based local optimisation algorithm, Sequential Quadratic Programming (SQP), a derivative free direct search method Mesh Adaptive Direct Search (MADS) and two evolutionary algorithms - Genetic Algorithms (GA) and Differential Evolution (DE) were used to solve the problem defined in Eqn. 3.12. For the purposes of comparison, a standard statistical method (Monte Carlo simulation (Fielding 2002, Vidyasagar 1998) was also implemented. All algorithms, along with the simulation model, were coded in the MATLAB® high-level modelling and computation environment (Mathworks 2002, Mathworks 2005).

3.3.4 Results

Figures 3.8-3.12 show the results of the application of the different optimisation algorithms to the problem defined in Eqn. 3.12. In each Figure, the graph on the left shows the improvement in objective function value as the algorithm proceeds (i.e. the rate of convergence to the optimal solution). Each algorithm has been run five times and the results of the best run are given. For SQP, each run is started at a randomly selected initial point in the search space. 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 normalised between -1 and 1 (lower bound and upper bound respectively, see Table 3.2, with a value of zero corresponding to its nominal value. The bar graphs show the optimal values of each of the 5 uncertain parameters in the order given in Table 3.1 (i.e. Hb, CO, VO_2, RQ and Temperature).
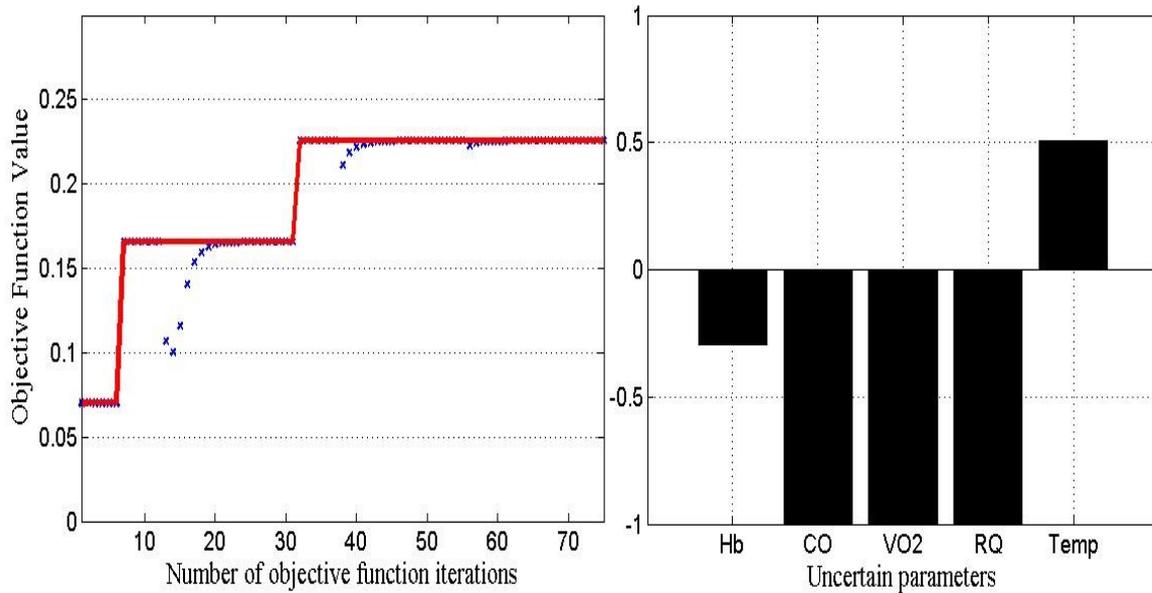
Figure 3.8: SQP result (a) SQP performance, (b) Optimal values of the uncertain parameters

Figure 3.8 displays the progression and termination of the best run of a local search SQP algorithm. Local optimisation algorithms like 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 multi-dimensional space defined by the uncertain parameters is likely to be non-convex, due 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 3.8(b) gives the optimal parameter combination found by the algorithm.
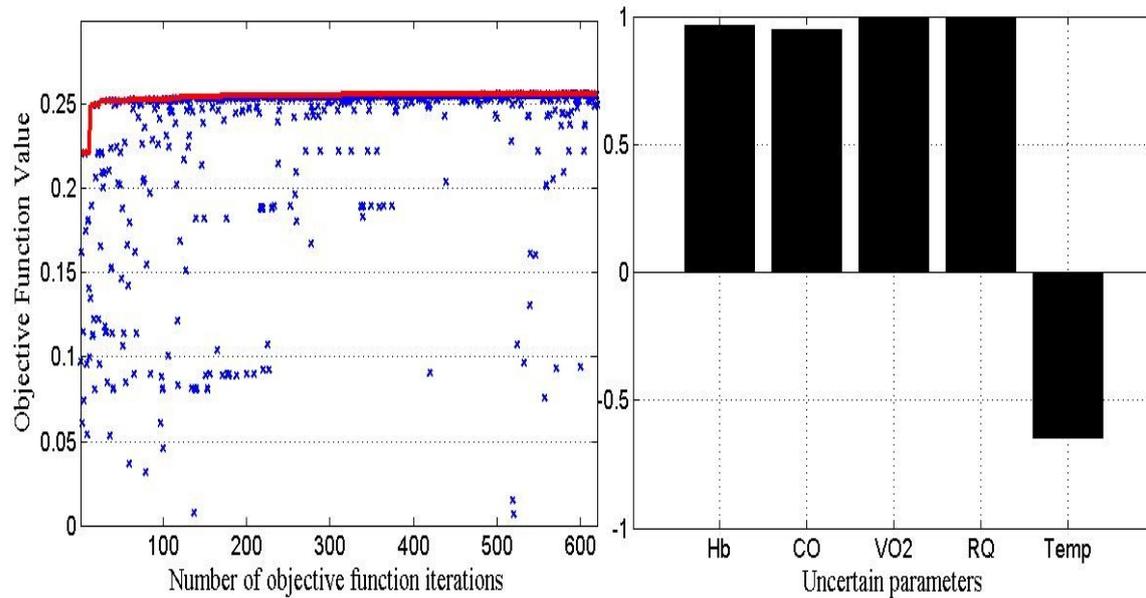
Figure 3.9: GA result (a) GA performance, (b) Optimal values of the uncertain parameters

Global optimisation algorithms use randomisation, 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 3.9 shows the results of the application of the global GA optimisation method to the same problem. The GA starts with an initial randomly chosen population of candidate solutions. Figure 3.9(b) gives the final optimal parameter combination found by the algorithm. Note that the global optimisation algorithm finds a completely different worst-case parameter combination to the optimal solution found by the SQP algorithm, with an improved maximum value of the objective function, although with a corresponding increase in the number of simulations required.
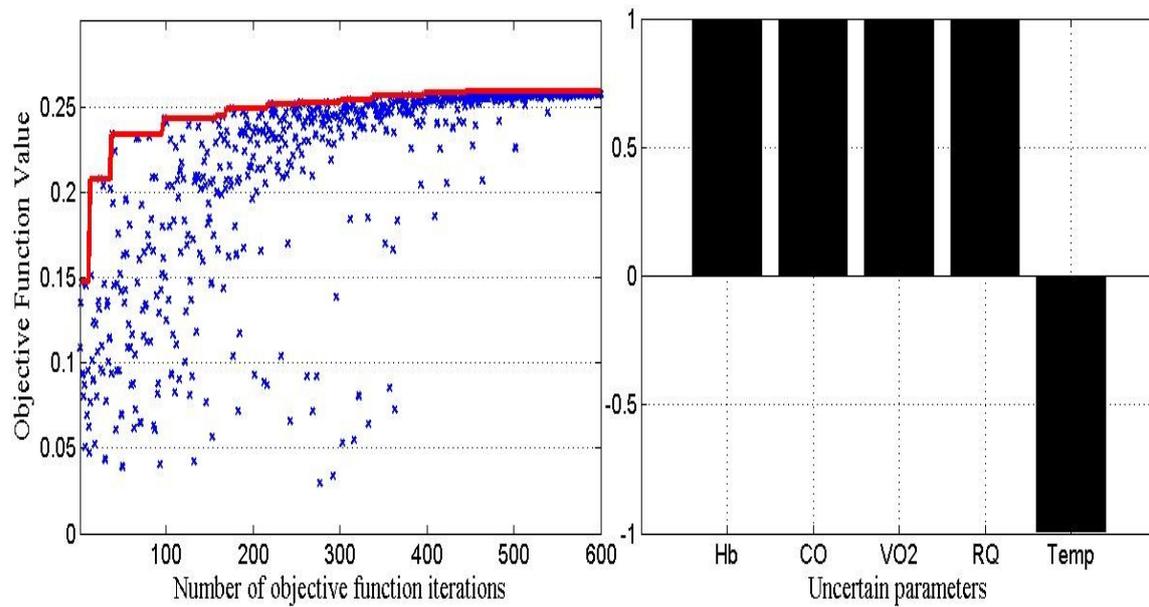
Figure 3.10: DE result a) DE performance, (b) Optimal values of the uncertain parameters

Figure 3.10 shows the results of the application of a second evolutionary algorithm, DE. From Figure 3.10 (b), it can be seen that the DE algorithm converges to a similar worst-case parameter combination and optimal value of the objective function, although with a slower convergence to the global optimal in comparison to the GA.

Figure 3.11 shows the result of the application of a direct search algorithm known as MADS. The mesh size tolerance for the algorithm was set to 10^{-6}. Figure 3.11(b) gives the final optimal parameter combination computed by the algorithm. It can be observed that this algorithm, which is based on completely different search principles to evolutionary algorithms, has varied each of the uncertain parameters in the same direction as 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 randomised nature and lack of gradient information, often arrive in the vicinity of the global solution but then have difficulty in converging precisely to it. In this case, the MADS algorithm proved to be significantly more efficient than the GA, requiring about a sixth as many simulations. Note, however, that results obtained via a local

algorithm like MADS would always need to be cross-validated against those obtained with a global algorithm, since its performance is completely dependent on a fortuitous choice of initial starting point.
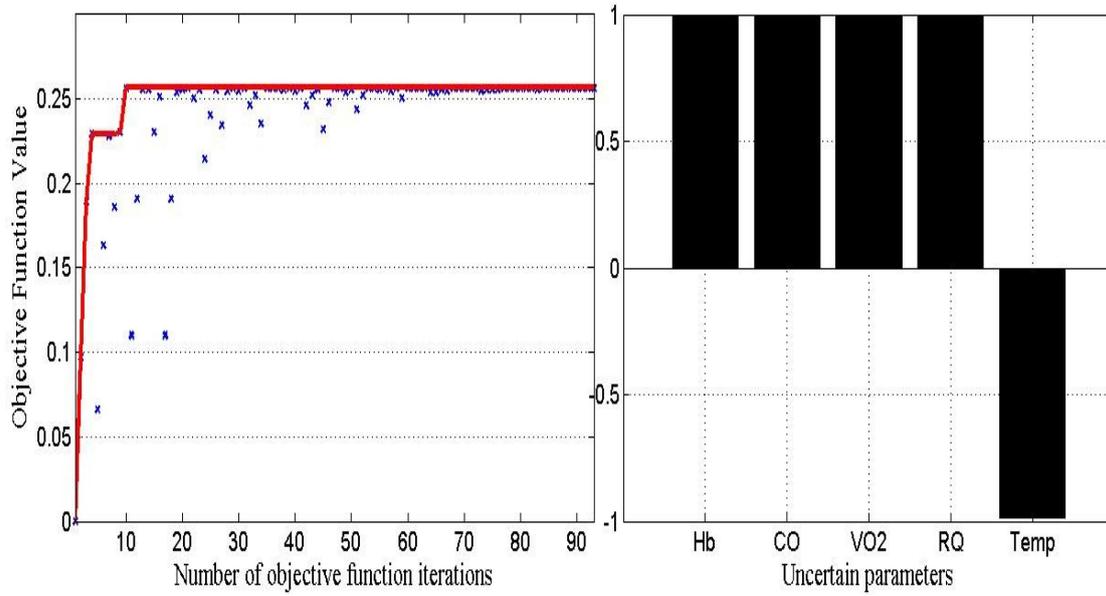
Figure 3.11: MADS result (a) MADS performance, (b) Optimal values of the uncertain parameters

For the purposes of comparison, Figure 3.12 gives the results of a standard Monte Carlo simulation (MCS) campaign using 1500 simulations. Although this number of simulations is fifteen times the number required by the MADS algorithm, and provides strong statistical confidence that the true worst-case behaviour will be found (see Table 3.4), it is apparent from Figure 3.12 that only a local optimum for the problem has in fact been identified, with the corresponding values of the uncertain parameters being significantly different from their true worst-cases.
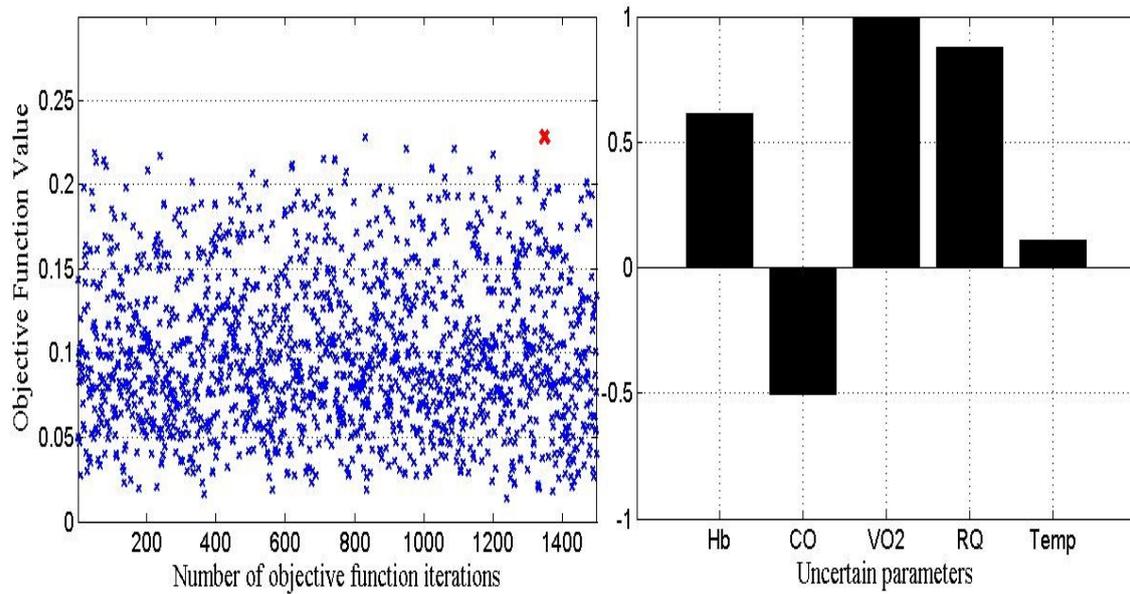
Figure 3.12: Monte Carlo result a) Monte Carlo performance, (b) Optimal values of the uncertain parameters. Note- x denotes the best value found over 1500 function evaluations.

Table 3.3 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.

Optimisation Algorithm	Optimal Parameters [Hb, CO, VO$_2$, RQ, T]	Number of simulations	Optimal value of the objective function
SQP	[134, 4500, 200, 0.704, 37.29]	75	0.2256
GA	[160, 5493, 300, 0.895, 37.08]	620	0.2561
MADS	[161, 5500, 300, 0.896, 37.02]	100	0.2565
DE	[161, 5500, 300, 0.896, 37.01]	600	0.2570
MCS	[153, 4749, 300, 0.884, 37.22]	1500	0.2285

Table 3.3: Comparison of the results of the different optimisation algorithms

Figure 3.13a-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 3.2. The shaded cubical region in Figure 3.13a represents the allowable ranges of each output as given in Table 3.2. It is clear that the worst-case model predictions never violate the boundaries of the region over the course of the simulation. Figure 3.13b-c displays the worst case result as given by each of the different optimisation algorithms.
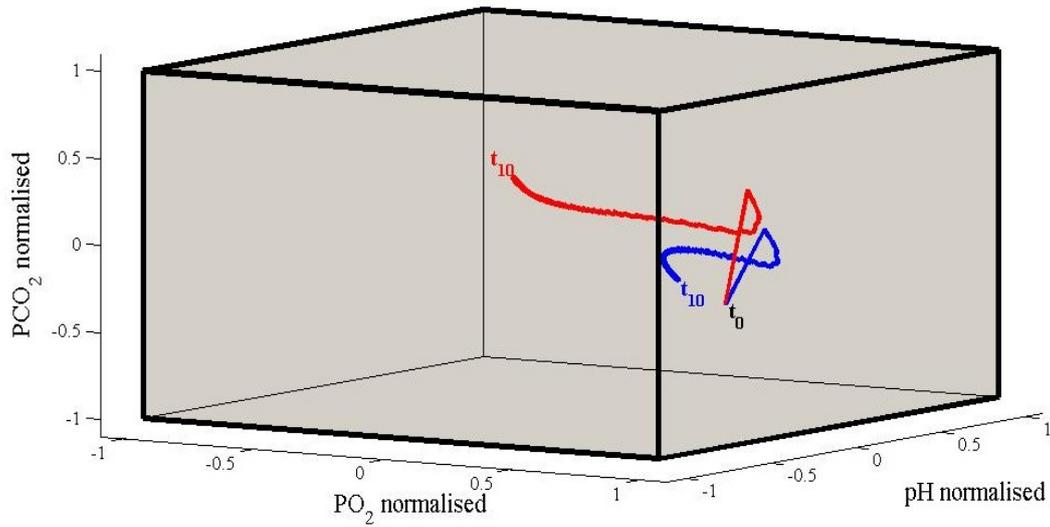
Figure 3.13a: Nominal (blue) and worst-case (red) model predictions for normalised PO_2, PCO_2 and pH. t_0 – simulation starting time, t_{10} – simulation end time, *box* – allowable values
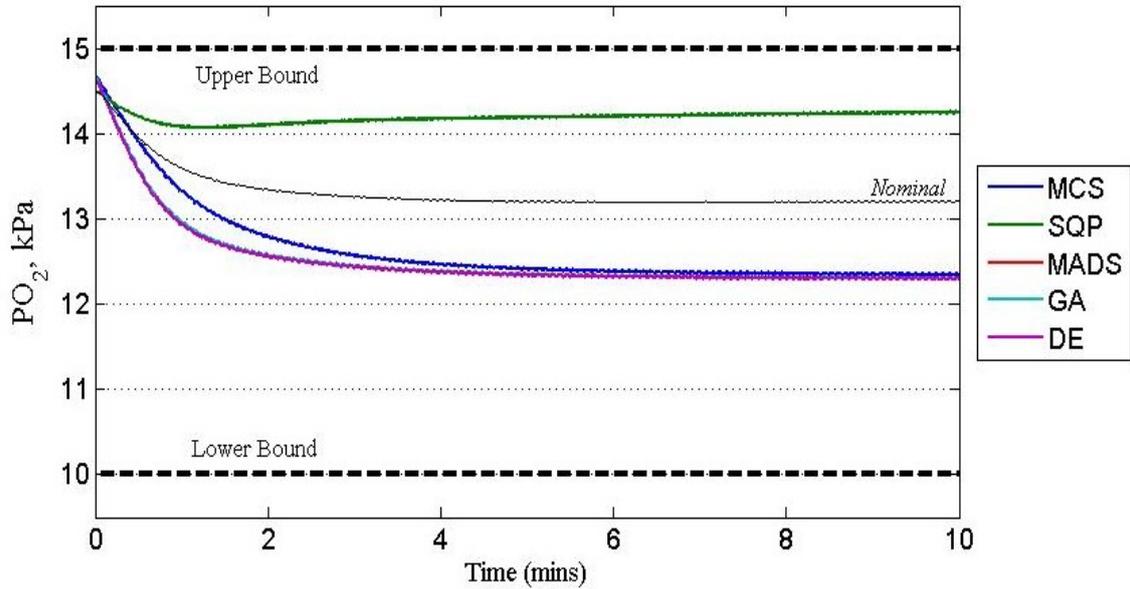
Figure 3.13b: Nominal and worst-case model predictions for PO_2
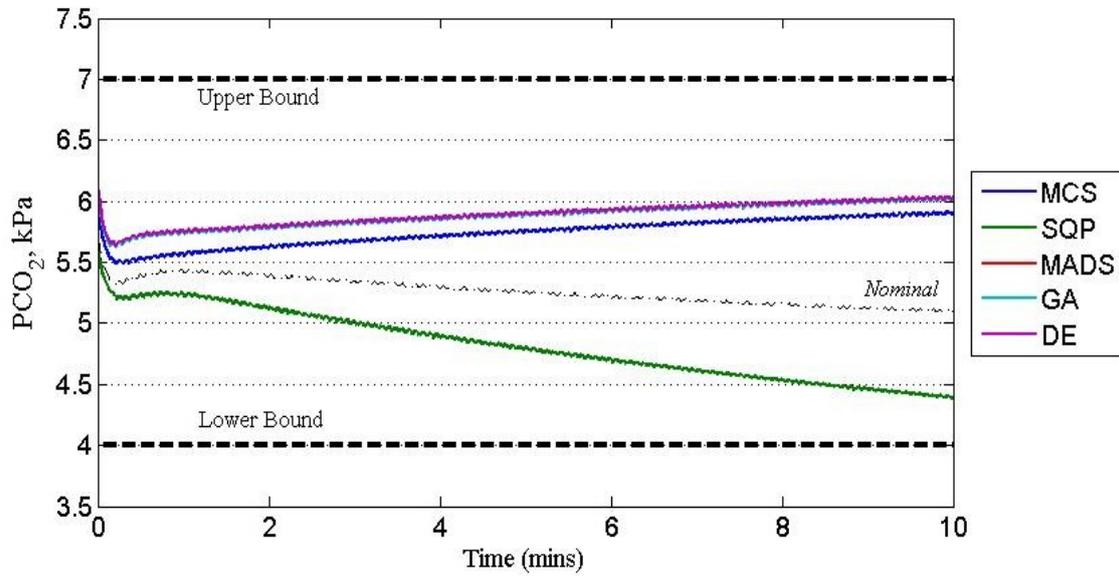
Figure 3.13c: Nominal and worst-case model predictions for PCO_2
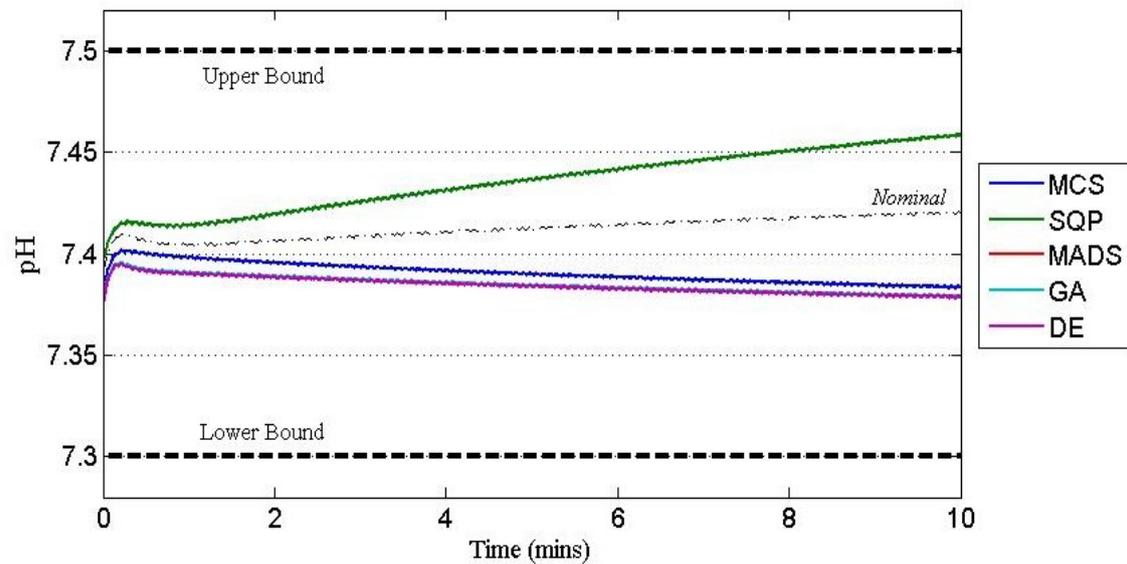
Figure 3.13d: Nominal and worst-case model predictions for pH

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 3.14(a)) as well as at the worst-case point (Figure 3.14(b)) in the uncertain parameter space. As shown in the figures, VO_2 has the strongest affect on the objective function relative to the other

parameters. Reassuringly, this is to be expected from physiological considerations as, of all the parameters, VO_2 will be the one which has the most direct effect on the quantity of oxygen and carbon dioxide in the blood. VO_2 is the oxygen taken away from the blood (reduction in PO_2), which would result in higher carbon dioxide production at the peripheral tissues (higher PCO_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 PO_2, PCO_2 and pH values, i.e. it will have a large influence on the value of the objective function.
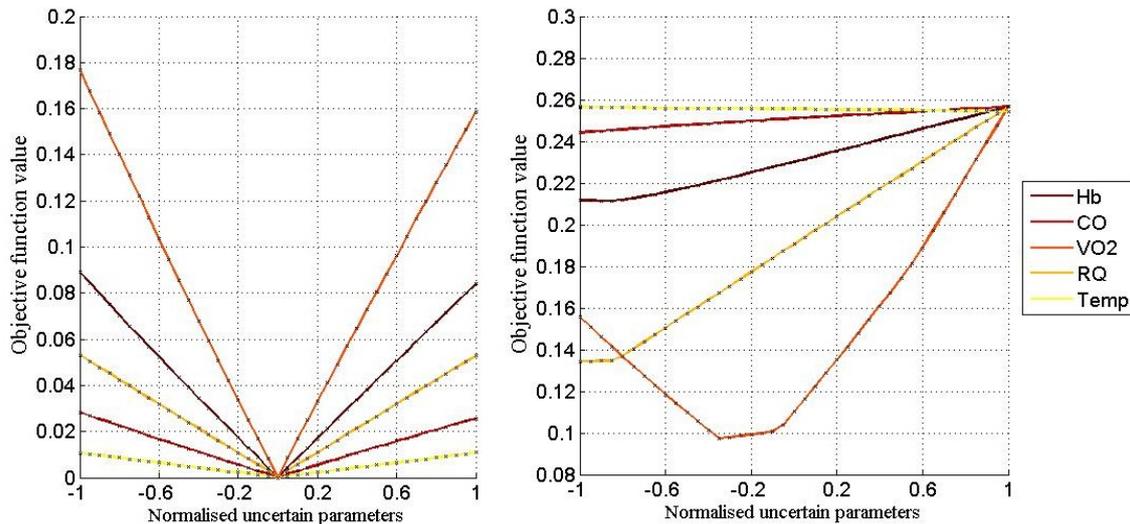
Figure 3.14: Sensitivity analyses at the (a) nominal case and (b) worst case

3.3.5   Discussion

Figure 3.15 shows the box plot of the distribution of optimal parameters as given by the algorithms over several executions. It is a useful comparison tool for the different algorithms. It clearly shows that a result produced by a single execution of the SQP strategy would be unreliable. An interesting result is the comparison between GA vs MADS where even though

Validation of Pulmonary Physiology Models

MADS produced a better single worst case over several executions (Table 3.3), the GA consistently reached closer to the global optimum than MADS.
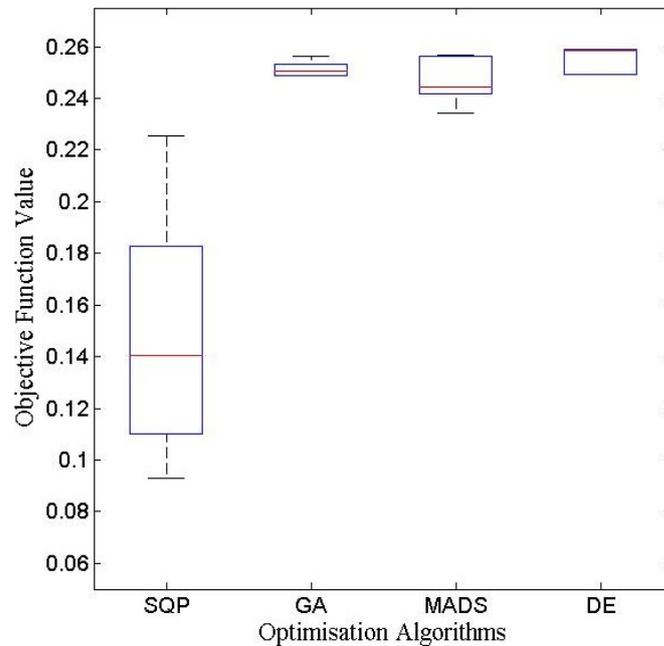
Figure 3.15: Box plot showing the performance of different algorithms. The red line is the median value. The edges of the boxes are the 25[th] and 75[th] percentile while the whiskers extend to the most extreme data points.

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 and accuracy for the statistical analysis are increased. As shown in Table 3.4, (taken from (Williams 2001)) to be within a probability range of 0.75 to 0.95 of estimating the worst-case behaviour of a system to within ±20%, only 25 simulations are required. If, on the other hand we need to be within a probability range of 0.94 to 0.999 of estimating the worst-case behaviour of a system to within ±1%, then 40,000 simulations are required.

Uncertainty probability range	Number of Monte Carlo trials based on percent 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	1,600	40,000

Table 3.4: Number of Monte Carlo trials required to achieve a desired estimation uncertainty with known probability, (Williams 2001) derived using the Chernoff bound (Vidyasagar 1998).

Since this number of simulations is usually 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 chapter, 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 recognised throughout the aerospace industry, which spends vast sums each year on Monte Carlo simulation campaigns to validate new safety-critical flights systems. This has led to a drive to investigate alternative validation strategies based on the use of global optimisation algorithms to intelligently search for worst-case behaviour, and promising results in this direction have appeared in the research literature (Fielding 2002, Menon 2006, Menon 2009, Kim 2006).

It should be noted that a principal motivation for the described attempts to 'break' the model is to uncover unexpected conditions or responses, which could need further investigation, and may (or may not) point to limitations or weaknesses in the formulation of the model. Expert analysis

of the model cannot be replaced by an automated computer program – rather an analysis method is proposed 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 valuable insight into the reliability and limitations of any new therapeutic strategies arising from the model predictions.

Simulation models which are intended for direct application in healthcare research and clinical practice should 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 irrefutable 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. 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 reliable predictions upon which novel therapeutic strategies may be based. As demonstrated by the results presented in this chapter, the time is now ripe for such optimisation-based approaches to play an important role in transforming physiological simulators into powerful biomedical engineering tools for direct application in clinical practice.

Chapter 4: Optimising Mechanical Ventilator Settings

This chapter considers the problem of optimising the settings of mechanical ventilation in an Intensive Therapy Unit (ITU) for patients with diseased lungs. The chapter begins with an introduction to the principles of mechanical ventilation and its application in the ITU. A review of current expert opinion on the incidence of Ventilator Associated Lung Injury (VALI) in mechanically ventilated patients is presented. The inherent trade-offs associated with satisfying the objectives of mechanical ventilation while minimising the risk of VALI in the modern Intensive Therapy Unit (ITU) are discussed. Current ITU approaches for determining mechanical ventilator settings are also reviewed.

Motivated by the limitations of these approaches, the selection of optimal mechanical ventilator settings is formulated as a constrained multi-objective optimisation problem. A validated physiological simulation model is used to optimise the settings of mechanical ventilation for a healthy lung and for several pulmonary disease cases. The constrained multi-objective optimisation problem is solved using a weighted aggregated objective function (WAOF) approach and as well as via a non-dominated sorting genetic algorithm NSGA-II. The optimal solutions obtained by the WAOF approach are compared with the set of optimal solutions obtained by NSGA-II.

The optimal settings computed for several simulated disease scenarios are shown to satisfy the conflicting clinical objectives, and improve the ventilation perfusion matching within the lung. Crucially, they are also shown to be highly disease-specific. The proposed combined modelling and optimisation approach offers physicians a tool with which to conduct *in silico* trials on

strictly-compliant virtual patients, in order to more easily assess and manage the trade-offs involved in determining appropriate patient- and disease-specific ventilator settings.

4.1   Introduction

4.1.1   Mechanical Ventilation

Mechanical Ventilation (MV), in its numerous forms, is a frequently used procedure in the ITU, and is usually required when a patient's natural respiration is inadequate on its own to maintain life. The majority of critically ill patients in Intensive Therapy Units (ITU) spend some time with their lungs ventilated with a mechanical ventilator. It is estimated that MV is required by nearly 1.5 million patients in the United States every year (Hill 2001) and indications are that this figure is set to increase (Rubenfield 2005).

The earliest ventilators worked by creating a negative pressure around the chest, allowing flow of air into the lungs, known as negative pressure ventilation (Woollam 1976). Modern mechanical ventilators use positive pressure ventilation, achieved by increasing the pressure in the patient's airways and thus creating a pressure gradient for the air to flow into the lungs. Mechanical ventilation can be invasive (e.g. tracheal intubation) as well as non-invasive (bag valve / PAP ventilators) and the applications ranges from short term (ITU / Operating room) to long term (chronic lung illnesses).

The purpose of the mechanical ventilator is to support a patient's breathing. Breathing is a periodic activity which can be simplified into two stages: airflow into the lungs (inspiration) and airflow out of the lungs (expiration), which together constitute a *breath*. Mechanical ventilators are machines that create a flow of air into and out of the lungs to create a *breath*, working alongside or replacing the work done by the patients' respiratory muscles.

The desired output provided by a positive pressure mechanical ventilator is simply the *pressure* that would be needed to create a pressure gradient between the ventilator and the lungs that will create the *flow* of air which will produce a change in the *volume* of the lung (Chatburn 2006). Mechanical ventilator outputs can be classified in the following ways:

Ventilator decided control variables: The breath produced by the ventilator is controlled by the magnitude of four variables – *pressure, volume, flow* and *time*. Between flow, volume and pressure, the selection of one variable by the mechanical ventilator can determine the characteristics of the other two. For example, during *pressure controlled ventilation*, the pressure characteristics produced by the ventilator are predetermined, while the flow and volume characteristics are dependent on this pressure (along with compliance and resistance to the ventilator associated with the patient and ventilator equipment).

Patient dependent variables: A breath that is based on the patient parameters can be classified into the following types:

Trigger: The patient's next inspiration is started (*triggered)* when a preset requirement is met on measured values of pressure, flow, volume or time. For example, if the patient airway pressure falls below a base value, the next inspiration is triggered.

Cycle: A measured variable of the patient indicates the end of inspiration so that the expiration stage can be started. For example, under pressure cycle, a flow is delivered to the lung until preset pressure is reached, after which expiration begins.

Limit: If a measured variable reaches a preset *limit*, then the current stage of breathing is held. The change of phase to expiration is dependent on the *cycle* settings mentioned above. For

example, if an upper limit is set on the pressure, then during inspiration, when the pressure in the lung reaches that pressure, the inspiration is held so that the pressure is not exceeded.

Baseline: This is a variable that is controlled by the ventilator during expiration. For example, under pressure control, the pressure in the lung at the end of expiration is often controlled. If this pressure falls to atmospheric pressure, then the baseline pressure is zero. If the pressure in the lung is maintained above a certain value, this baseline pressure is known as positive end expiratory pressure (PEEP).

Ventilation Modes:

Types of Breath: A *spontaneous breath* is a breath which is determined completely by the patient. A *mandatory breath* is determined completely by the mechanical ventilator. During *assisted breath*, the ventilator plays a part in the breathing (*e.g. Pressure Support Ventilation, Continuous Positive Airway Pressure)*.

Breath sequences: In *continuous mandatory ventilation* (CMV), all breaths are mandatory and initiated by the ventilator. In *continuous spontaneous ventilation* (CSV), all breaths are spontaneous and initiated by the patient. In *intermittent mandatory ventilation* (IMV), spontaneous breaths are allowed between mandatory breaths. A variant of IMV, synchronised-IMV (SIMV) allows for the mandatory breaths to be started by the patient.

The model in this study simulates a "virtual" patient ventilated under heavy sedation with no effort from the patient. The patient's breaths are determined fully by the mechanical ventilator, i.e. CMV. The breath provided by the mechanical ventilator is set to constant flow.

4.1.2 Ventilator Associated Lung Injury (VALI)

The capacity of modern artificial mechanical ventilation to cause lung injury has been well-known for several decades (Mead 1970). Ventilator Associated Lung Injury (VALI) refers to the injury suffered through the exposure of the patients' lungs to potentially damaging positive pressures and high volumes, which can result in prolonged ITU care, pneumonia, lifelong lung scarring and even multiple organ failure (Dreyfuss 1992, Gammon 1992). More than 40,000 patients undergo mechanical ventilation in the ITU each year in the UK (Nouraei 2009). Of these, approximately 2.9% suffer from VALI (Anzueto 2004), which in the UK alone represents several thousand cases.

Under mechanical ventilation, patients can potentially be subjected to excessively high and possibly damaging ventilation contributions, such as high or highly-variable tidal volumes and pressures (ARDSnet 2000). This could have adverse effects on the operation of the respiratory system and consequently other major organs. Lungs can be exposed to:

- High airway pressures during positive pressure ventilation resulting in barotrauma (Slutsky 1993).

- Over-distension of alveoli can also be caused by excessive volumes of air entering the lung and causing volutrauma (Auten 2001).

- Conversely, low lung volumes and certain ventilation settings expose the lungs to damage caused by repetitive opening and closing of the airways due to shear stresses, causing atelectrauma (Slutsky 1999).

- Intubation used for mechanical ventilators has been shown to release mediators under different mechanisms such as stretching and shear stress. These can cause localised

injury and inflammation (Dreyfuss 2003).. The investigation of this process is presently beyond the scope of this project and hence will not been addressed further here.

There have been numerous clinical trials conducted with the aim of developing guidelines for adjusting mechanical ventilation parameters for ITU patients to reduce the risk of VALI (Amato 1998 Marcy 1993 Laffey 2000 ARDSnet 2000). These studies have focused on proposing ventilation parameters to reduce the exceptionally high mortality rates resulting from VALI. However, both the assessment of current practice in mechanical ventilation and the attempt to provide improved guidelines for physicians have met with several hurdles. The rates of administration of mechanical ventilation vary with geography, hospitals, resources and ITU conditions (Okamoto 2004). Furthermore, there remain differences and ambiguities in the definition of several conditions that lead to administration of mechanical ventilation (Villar 1999). In some cases, even the statistical techniques used in data analysis have been shown not to meet satisfactory standards (Moss et al 2003). All of these issues can leave the outcomes of studies and clinical trials subject to such bias and unreliability that they have the potential to mislead (Concato, 1993).

4.1.3 Current approaches to deciding mechanical ventilation settings

In general, physicians in an ITU aim to maintain the patient's ventilation parameters within regions of 'respiratory comfort' (Dojat 1992, Boudama 2005) based on generic guidelines depending on arterial O_2-CO_2 levels. They are aided in this task by advanced mechanical ventilator systems which provide decision support capabilities and strict command over patient breathing. Early mechanical ventilators were simple feedforward systems (Chatburn 2006) which provided a set volume and pressure of gas to the patient at a predesignated rate. As the machines have evolved the automatic feedback of the patient data in deciding mechanical ventilator settings has become common. At its simplest, these machines regulate the mechanical

ventilator outputs based on the previous outputs and state variables of the system. Modern ventilators with different ventilator support strategies such as Adaptive Support Ventilation (ASV) (Laubscher 1994), Proportional Assist Ventilation (PAV) (Younes 1992), etc utilise feedback control. These strategies are outside the scope of this study, but for a review of the development of closed loop control in the field of mechanical ventilation, the reader is referred to (Tehrani 2008a, 2008b).

It should be noted that most of the systems discussed in (Tehrani 2008a, 2008b) make decisions solely based on patient outputs and protocol based ranges of mechanical ventilator settings. For example, Tehrani (Tehrani 2008c) proposes FLEX, a tool designed to be used as both as an advisory open loop system under physician control and as a closed loop ventilator control system. Its rules (decision making) are adaptive and based on individual patients. However it does not utilise any direct models of gas transport or any relevant physiology but decides mechanical ventilator parameters based solely on a patient's clinical parameters. As such the predicted mechanical ventilator settings provide little physiological insight regarding the underlying disease to the physician. It should also be noted that proposed rule based systems such as FLEX tend to adjust a single mechanical ventilator parameter based on current patient data, and then subsequently adjust other parameters. For example, in FLEX, inhaled oxygen fraction (FIO_2) is decided by the arterial oxygen saturation. The positive end expiratory pressure (PEEP) is then decided based on this FIO_2 and the inspiratory pressure is determined by the PEEP. The applicability of FLEX has to-date only been established in the context of management and weaning of ventilation. The rest of this section will focus on reviewing the state-of-the-art in prediction of mechanical ventilator settings using physiological modelling.

In 1993, Rutledge (Rutledge 1993, 1994) proposed two versions of a respiratory model, VentSim and VentPlan, both designed (alongside a rule based network) to be used in an ITU as

a rapid predictive tool for ventilatory contribution. VentPlan is a simplified three compartment physiological model while VentSim is a relatively more comprehensive representation. The predictive algorithm employs a belief network (an algorithm mapping qualititative data into probability distributed quantitative values) which alongside a mathematical model presents several plans for ventilatory settings for a patient under a given situation. A plan-evaluator then ranks these plans with their predicted outcomes and the most satisfactory solution is displayed. Unlike VentPlan, the results of VentSim have not been reviewed against clinical data. In addition, the simplistic model used in VentPlan could compromise its applicability on patients with more complex and variable pathophysiological abnormalities.

In 2006, Rees proposed a decision support system for optimising mechanical ventilation for patients in an ITU. The optimisation algorithm decides on a ventilation strategy in terms of possible results such as sufficient oxygen, minimising risk of acidosis and alkalosis and ventilator induced lung injury. The system claims the use of actual physiological models instead of relying on heuristic decision based support systems. However the model used is extremely simplistic, and, for example, includes no detailed representation of ventilation-perfusion distribution. No testing of the system has been published so far on patient populations, and hence a valid clinical assessment of its efficacy does not yet exist. Recent work by this group is focussed on developing a more detailed model of the respiratory system for use in VALI assessment (Mogensen 2010).

In general, previous approaches to deciding mechanical ventilation settings have been based on preset guidelines on patient parameters (McKinely 2001, Boudama 2005) with little consideration of individual patient physiology and no attempt to providing settings which optimise the tradeoffs involved in mechanical ventilation. Rule based methods provide little physiological insight into the reasons for their decisions and it is rare that multiple mechanical

variables are adjusted simultaneously. In this thesis, we propose an optimisation based methodology to determine optimal choices for multiple mechanical ventilator settings for specific disease scenarios.

4.2   Problem introduction

The development of a detailed quantitative understanding of the aetiology of VALI poses significant challenges. Previous modelling investigations and several current mechanical ventilators have used only very simplified models, and have not considered crucial factors such as the non-linearity of alveolar compliance and the closure of alveolar airways (Rutledge 1993, Rees 2006).

The combination of high-fidelity, validated, pathophysiological simulation models and global optimisation methods offers a potentially powerful framework with which to address many of the challenges facing clinicians in the ITU, since they enable "virtual" experiments yielding quantitative results to be conducted on an infinitely compliant and strictly controlled *in silico* patient. These experiments may be used to enhance understanding of the causes and the pathophysiology of VALI. They can also be used to interrogate the effectiveness of "typical" ventilator parameter settings on different pathologies in different patients. The experiments can suggest novel, possibly counterintuitive MV settings for subsequent testing in clinical trials. A particular advantage of *in silico* approaches is that they can be used to investigate scenarios where conflicting processes are occurring and the tradeoffs involved in varying multiple parameters are not intuitively obvious. This is clearly the case for the problem considered here, which requires the management of tradeoffs between oxygenation and CO_2 clearance objectives and VALI risk factors, via simultaneous adjustment of five ventilator parameters: tidal volume, ventilatory rate, ratio of inspiratory to expiratory time, positive end expiratory pressure and the

inspired fraction of oxygen. The aim is to optimise these ventilator settings to minimise the risk factors for VALI, while guaranteeing effective control of arterial oxygenation, carbon dioxide clearance and pH maintenance.

In mathematical terms, this problem corresponds to a multivariable and multi-objective constrained optimisation problem. The dynamics of the ventilator-respiratory system are complex, and include significant nonlinear and non-smooth properties. Since the possible parameter space for ventilator settings is almost certain to be non-convex, local gradient-based optimisation algorithms are very unlikely to compute globally optimal results and we instead employ methods that have proven to be effective in solving non-convex problems with non-differentiable functions, nonlinearities and stochastic properties, (Mitchell 1998, Audet 2007, Deb 2002). Although there have been some previous attempts to apply optimisation methods in the context of mechanical ventilation, (Gupta 1978, Rudowski 1991), this study is the first to:

(a) employ a computer simulation model which has been comprehensively validated both from a clinical (Hardman 1998) and engineering point of view (Das 2010),

(b) investigate the potential of global multi-objective optimisation algorithms, and

(c) incorporate specific consideration of VALI risk factors and disease states.

4.3   Multi-objective optimisation

4.3.1   Introduction

A general formulation of a multi-objective optimisation problem with *M* objectives can be described as follows:

$$\text{Minimise} \qquad F_k(x) \qquad\qquad k = 1, 2, \dots, M \qquad\qquad\qquad [4.1]$$

$$\text{Subject to} \qquad g_j(x) \leq 0, \qquad j = 1, \dots, v$$

$$g_j(x) = 0, \qquad j = v + 1, \dots, w$$

$$x^L \leq x \leq x^U$$

x is a vector of n decision variables: $x = (x_1, x_2, \dots, x_n)^T \in \mathbb{R}^n$ and $F_k(x)$ is a scalar objective function such that $F_k : \mathbb{R}^n \to \mathbb{R}$. A solution x can only be selected between the lower and upper bounds x^L and x^U respectively. The set of solutions within these bounds is known as the *decision space*. The constraints that a solution x must satisfy are defined by the constraint function $g_j(x)$. A solution x that does not lie in the decision space or satisfy $g_j(x)$ is called an *infeasible solution*. Conversely if a solution x exists in the decision space and satisfies $g_j(x)$, then it is called a *feasible solution*. The set H of all feasible solutions is called the *feasible region* (Figure 4.1).

In multi-objective optimisation problems, the desired objectives can often be conflicting. In such cases, a set of solutions need to be computed that are *non-inferior* to the other solutions in the decision space. A non-inferior solution is a solution in which an improvement in one objective causes degradation in another objective. The set of non-inferior solutions is called the *Pareto optimal* set of solutions. Consider two candidate solution vectors $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$. The *Pareto* concept may be explained using the *dominance* operator $\prec$ as follows: The vector x dominates vector y ($x \prec y$), if and only if,

$$F_k(x) \leq F_k(y) \quad \forall k = 1, \dots, M \text{ and} \qquad\qquad [4.2]$$

$$F_k(x) < F_k(y), \quad \text{for at least one } k. \qquad\qquad [4.3]$$

When comparing two candidate solution vectors x and y, there are three possibilities: x dominates y ($x \prec y$), x is dominated by y ($x \succ y$) or x and y are non-dominated.
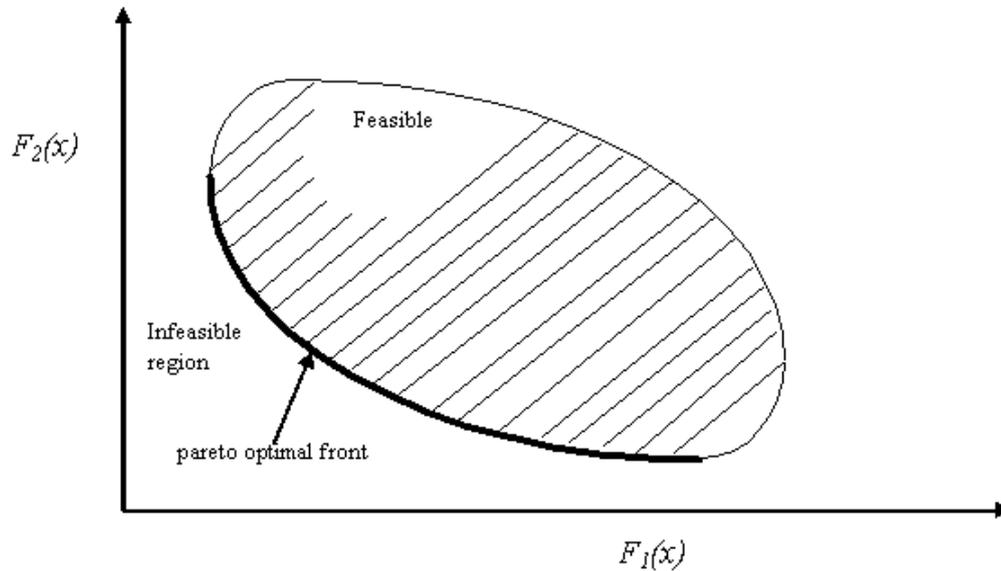
Figure 4.1: Objective space of a multi-objective optimisation problem

A solution vector x is Pareto optimal if and only if there does not exist another solution y such that $(y \prec x)$ $\forall y \in$ H and x $\neq$ y. The resulting solution vector x represents the globally optimal solution to the tradeoff problem defined by the conflicting objectives of $F(x)$. A set of non-dominated solutions x is called the non-dominated front. The set PO of all Pareto optimal solutions is called the *Pareto optimal set*. The set PF given by:

$$PF = (F_k(x)|\ x \in PO, \forall k = 1, \dots, M\ )$$
[4.4]

is known as the *Pareto optimal front*.

In most single objective optimisation problems, the goal is simply to find the globally optimal solution. The optimisation algorithm will accept a new solution if the objective function returned by the new solution is better than that of the previous solution. However, in multi-

objective optimisation, there are two goals. The first goal is to reduce the distance between the non-dominated front and the Pareto optimal front. This can be achieved in several ways, for example by using aggregated weight selection (Roy 1971), switching the search between objectives (Vector Evaluated Genetic Algorithms in Schaffer 1985), ranking solutions according to fitness (Multiple Objective Genetic Algorithm (Fonseca 1993) etc. The second goal is to maintain a diverse set of solutions. This ensures that the non-dominated front has solutions that represent a wide range of trade-off solutions for each objective. The diversity of the two solutions in the objective space can be measured by the Euclidean distance between them in the objective space. This goal can be addressed using several methods such as fitness sharing (Srinivas 1994), crowding (Deb 2002) etc.

4.3.2   Methods for multi-objective optimisation

Methods for multi-objective optimisation methods can be classified broadly into two categories, the transformation of multiple objectives into a single objective function and the search for Pareto optimal solutions. From these categories, the thesis uses the two popular methods:

a) Weighted aggregated objective function

b) Non-dominated sorting genetic algorithm (NSGA –II)

The methods are illustrated and compared by solving a multi-objective optimisation problem of determining the optimal mechanical ventilator settings for a disease case scenario.

a) Weighted aggregated objective function

A simple and intuitive way of addressing the multi-objective optimisation problem is to combine the multiple objectives into a single objective function. In the literature, this function is

often referred to as the aggregated objective function (AOF) (Roy 1971). There are several methods to design an AOF; one of the most popular is the weighted sum approach (Roy 1971, Messac 2000). In the weighted sum approach each objective is multiplied by a user-supplied weight. The composite objective function J is then formed by summing the weighted objectives. In this study, we used the weighted sum approach to obtain a single optimal solution based on the preferential selection of weights provided *a priori* by the user. We refer to this approach as the weighted aggregate objective function (WAOF). The multi-objective optimisation problem of [4.1] is thus converted to a single objective optimisation problem as follows:

$$\text{Minimise} \quad J = \sum_{k=1}^{M} w_k F_k(x) \qquad x = (x_1, x_2, \ldots, x_n)^T \in \mathbb{R}^n \qquad [4.5]$$

In this approach, the multi-objective problem can thus be solved using the single objective optimisation algorithms described in Section 3.2. The optimisation algorithms discussed in Section 3.2 can be broadly classified into two categories: global optimisation methods (GA and DE) and local optimisation methods (SQP and MADS). Global methods based on evolutionary principles are generally accepted as having a high probability of converging to the global or near global solution, if the method is executed for a long enough time with sufficient number of initial candidates. This advantage is offset by the rate of convergence of global methods which can be slow. Local methods have been shown to rapidly converge to local solutions but the quality (i.e. the proximity to the global solution) of those solutions entirely depends on the starting point chosen for the optimisation routine.

To solve Problem [4.5], a hybrid local/global algorithm is employed (Krasnogor 2005), where the starting point of the local method (MADS) will be decided by initially running a global optimisation method (GA). Relative to the use of a global algorithm on its own, this should

provide faster convergence while still avoiding getting trapped in local optima (El-Mihoub 2006). The local search algorithm MADS is used to rapidly refine the "close to optimal" solution returned from the GA. The GA is set up with an initial population size of 30 and number of generations set to 20. MADS is set to terminate if the number of function evaluations exceeds 500 or the objective function fails to improve by 10^{-3} over successive iterations (TolFun<10^{-3}, Appendix B). The maximum number of function evaluations using the hybrid algorithm is thus fixed at 1100.

b) Non-dominated sorting genetic algorithm (NSGA–II)

NSGA-II is an elitist multi-objective evolutionary algorithm (Deb 2002). It is based on the genetic algorithm (GA), utilising evolutionary operators such as crossover and mutation, described previously in Section 3.2 of this thesis. NSGA-II differs from GA in its selection operation. In GA and NSGA-II, the total population consists of the parent population P_{parent} from the previous generation and the offspring population $P_{offspring}$ resulting from crossover and mutation. The selection operator in GA uses the fitness of the individual members of the population to determine the population for the next generation. However, the NSGA-II selection operator differs from GA by incorporating non-dominated sorting, a crowding distance calculation and a crowded tournament selection (described below). These three operators allow NSGA-II to a) efficiently guide the search towards the Pareto optimal set relative to other multi-objective algorithms and b) maintain a diverse population which prevents premature convergence and achieves a well distributed non-dominated front (Deb 2002).

Non-dominated sorting: Using the dominance property, each solution in the population is classified into fronts $\mathcal{F}$ which represent the tradeoff involved in the associated multi-objective

problem. The non-dominated solutions are assigned to front $\mathcal{F}_1$ while the remaining members of the population are classified into the subsequent fronts using the method described in Table 4.1.

Step 1: For each solution p in the population P, two quantities are calculated. i) n_p – the number of solutions that dominate p and ii) S_p – a set of solutions that p dominates. Initially, for all $p \in P$, $n_p = 0$ and $S_p = \emptyset$.

Step 2: For all $q \neq p$ and $q \in P$, if p dominates q, then $S_p = S_p \cup \{q\}$. Otherwise, n_p is calculated as $n_p = n_p + 1$.

Step 3: If $n_p = 0$, p is kept in the first non-dominated front by $\mathcal{F}_k = \mathcal{F}_k \cup \{p\}$, where $k = 1$.

Step 4: While $\mathcal{F}_k \neq \emptyset$, initialise $Q = \emptyset$ where the next set of non-dominated solutions are stored. For each $p \in \mathcal{F}_k$ and for each $q \in S_p$, update $n_q = n_q - 1$. If $n_q = 0$, $Q = Q \cup \{q\}$.

Step 5: Set $k = k + 1$ and $\mathcal{F}_k = Q$. Return to step 4 and repeat until all solutions are classified.

Table 4.1: Algorithm for non-dominated sorting

Crowding distance assignment: As previously mentioned, one of the goals of multi-objective optimisation algorithm is to maintain a diverse set of solutions. For this purpose, NSGA-II calculates a crowding distance metric d_i for each individual i in a front $\mathcal{F}$, which is calculated as shown in Table 4.2.

Step 1: Let the number of solutions in $\mathcal{F}$ be l.

Step 2: Sort the solutions in $\mathcal{F}$ by ascending order of magnitude for each objective function F_k ($k = 1, \dots, M$ where M is the total number of objectives). For each solution i in $\mathcal{F}$, assign $d_i = 0$.

Step 3: For each objective function k, the solutions at the boundaries are assigned a large distance value, $d_1 = d_l = \infty$. For $j = 2, \dots, l-1$, the distance d_j is measured as

$$d_j = d_j + \frac{F_k^{j+1} - F_k^{j-1}}{F_k^l - F_k^1}$$

Step 4: The overall crowding distance value is calculated as the sum of individual distance values corresponding to each objective.

Table 4.2: Crowding distance algorithm

Crowded tournament selection operator: After non-dominated sorting and crowding distance assignment, a solution i is selected over a selection j if the following conditions are true:

1) Solution i has a better rank than solution j. The rank of a solution is decided by the front it belongs to, i.e. solutions in front $\mathcal{F}_k$ have a rank = k.

2) If the solutions have the same rank but the solution i has a better crowding distance d than solution j, i.e. $d_i > d_j$.

Table 4.3: Crowding Tournament Selection

The pseudo code for the NSGA-II search is described in Table 4.4 below:

1) For $gen = 1$, generate a random initial population P_{gen}, of population size N.

2) The individuals in the population P_{gen} are ranked and sorted by non-dominated sorting.

3) An offspring population $P_{offspring}$ is created of size N using crossover and mutation operators.

4) Combine the parent and offspring population: $R = P_{gen} \cup P_{offspring}$.

5) Sort population R into a set of non-dominated fronts $\mathcal{F} = (\mathcal{F}_1, \mathcal{F}_2, ...)$.

6) Generate a new population $P_{gen+1} = \emptyset$ and let $i = 1$.

7) Let p be the number of individuals in P_{gen+1}. Until $p < N$

      Calculate crowding distance in $\mathcal{F}_i$.

      Combine the non-dominated individuals of $\mathcal{F}_i$ with the current population and increment counter i.

$$P_{gen+1} = P_{gen+1} \cup \mathcal{F}_i. \qquad i = i + 1.$$

8) Select $N - P_{gen+1}$ individuals with a large crowding distance in $\mathcal{F}_i$ and add to P_{gen+1}.

9) Generate an offspring population $P_{offspring}$ of size N using crowded tournament selection, crossover and mutation operators. Increment the generation counter $gen = gen + 1$.

10) Repeat step 4-9 until termination criteria are satisfied.

Table 4.4: NSGA algorithm pseudo code

Since its appearance in 2002 (Deb 2002), NSGA II has become a popular tool in addressing multi-objective optimisation problems, with the key reference paper currently reporting in excess of 6000 citations. In multi-objective problems in the field of medicine, multi-objective optimisation techniques such as NSGA-II have shown considerable improvement over traditional methods for problems such as cancer diagnosis and classification (Garcia-Nieto 2009), image reconstruction (Dunn 2004), and radiation therapy (Milckovic 2001) etc. To the author's knowledge, multi-objective optimisation algorithms have so far not been considered in the context of optimisation of mechanical ventilation.

4.4   Problem description

4.4.1   Mechanical ventilator settings

The mechanical ventilator settings considered in this study, assuming the patient to be under complete sedation, are:

1. Positive end-expiratory pressure (PEEP, [cmH_2O]):   The closure of airways and repeated opening (recruitment) and closing (de-recruitment) of alveoli, especially for alveoli with normal or high perfusion of pulmonary blood flow, can cause ventilation perfusion mismatching. The poor ventilation of blood flow can cause poor oxygenation of the blood, while increased *dead space* due to collapsed airways can result in hypercapnia, an abnormally high level of CO_2 in the blood.   PEEP allows for re-establishing normal ventilation to de-recruited alveolar regions (Lachmann 2002). PEEP is the positive pressure in the lungs at the end of exhalation. When provided by an external source such as a ventilator, this is often referred to as the extrinsic PEEP. PEEP settings in the ITU can vary between 5-24 cmH_2O (ARDSnet 2004). Although the

benefits of PEEP in preventing collapse in the lung are apparent, PEEP can also cause hyperinflation resulting in possible volutrauma (Vieira 1999).

2. Fraction of Inspired Oxygen (FIO_2): FIO_2 refers to the fraction of oxygen constituting the inhaled volume of gas as provided by the mechanical ventilator. In normal air, FIO_2 is approximately equal to 0.21. High values of FIO_2 improve oxygen levels in the blood. However a high FIO_2 can result in rapid absorption of oxygen into the blood, potentially reducing alveolar volume below critical collapsing volume and causing collapse in the alveolar regions. Studies (ARDSnet 2000, Villar 2007) have suggested that increased FIO_2 should be accompanied by an increase in PEEP to prevent collapse.

3. Tidal volume (Vtidal, [ml]): This is the volume of air travelling in and out of the patient's lungs during every breath. It is typically set between 10 and 15 ml/kg of body weight (Holets 2006). Vtidal and FIO_2 are the main settings considered by the physician in increasing oxygenation in diseased lungs. However, Vtidal settings of even 10ml/kg have shown to cause over-inflation of lungs (Dreyfuss 1998) while lower volumes can cause atelectasis (Muscedere 1994).

4. Ventilation rate (VentRate, [breaths/min]): VentRate is the number of breaths per minute (also referred to as the frequency of breaths) as decided by the mechanical ventilator. Under normal breathing, this value is approximately 12.5 breaths per minute. Critically ill patients in the ITU are commonly provided breaths with VentRate value between 14 and 30 breaths/min (Holets 2006).

5. Inhale to Exhale (I:E) ratio: The I:E ratio decides the duration of breath that is spent inhaling relative to exhaling. Normal I:E ratio is approximately equal to 1:2, i.e. in a breath lasting six seconds, the duration of inspiration is two seconds while expiration lasts for four seconds. Long inspiration times can favour recruitment of collapsed lung units, and potentially increase oxygenation of the lung units and the pulmonary blood

flow. But as with PEEP, increased inspiration time can lead to hyperinflation. Common ITU settings for I:E ratio can vary between 1:1 and 1:3 (ARDSnet 2004). Inverse Ratio Ventilation (IRV), where inspiration is longer than expiration has been shown to improve gas exchange in some cases. (Cole 1984).

Other settings such as minute ventilation, mean expiratory flow, continuous positive airway pressure (CPAP) etc; provide other options for a physician while setting a mechanical ventilator. However their description and use is outside the scope of this thesis. The reader is referred to *Principals and Practices of Mechanical Ventilation* (Tobin 2006), for a comprehensive discussion on current mechanical ventilation equipment, strategies and settings.

4.4.2   *In silico* setup

The simulation model considered in this study is an extended MATLAB® implementation of several physiological models originally developed within the Nottingham Physiology Simulator (NPS). The model is designed to represent a dynamic in-vivo cardio-pulmonary state using a mass-conserving set of equations based on well-established physiological principles. Some significant parameters indicating the configuration of the model are given in Table 4.5.

The model has already been described in detail in Chapter 2 of this thesis. The model has been comprehensively validated using both standard clinical approaches and formal Systems Engineering methods. In (Hardman 1998, McCahon 2008) the model was validated against patient data using statistical methods, while in Chapter 3 of this thesis, the robustness of the model responses to parametric uncertainty was demonstrated using analysis techniques from the field of control engineering.

Configuration	Value
Weight [kg]	70
Inspired gas	Warmed and humidified air
Inspired flow pattern	Constant flow
Number of alveolar compartments	100
Respiratory quotient	0.8
Oxygen consumption [ml min^{-1}]	250
Cardiac output [liters min^{-1}]	5

Table 4.5: Significant parameters describing model configuration.

In this study, the five key ventilator settings described above are considered as variable parameters that may be adjusted to optimise the trade-off between effective gas exchange and minimising the risk of VALI. Maximum allowable ranges of variation for the values of these parameters have been defined based on clinical experience and published guidelines (ARDSnet 2000); Vtidal is allowed to vary within a range from 280 - 840 ml, corresponding to 4 ml/kg - 12 ml/kg for a body weight of 70 kg. VentRate is bounded within the range 10 - 20 breaths/min, I:E is limited to the interval 0.25 - 0.8 (i.e. a ratio between 1:3 and 4:1), PEEP is constrained within 0 - 10 cmH$_2$O and FIO$_2$ is within 0.21 - 1. All of these parameters are also assigned "typical" nominal values, which are used to derive their respective ventilation perfusion (V/Q) distributions (see Figure 4.2). A summary of these values is provided in Table 4.6.

Parameter (x)	Lower Bound (x^L)	Upper Bound (x^U)	Nominal
Vtidal [ml]	280	840	400
VentRate [bpm]	10	20	12.5
I:E	0.25	0.8	0.33
PEEP [cmH_2O]	0	10	0
FIO_2	0.21	1	0.21

Table 4.6: Mechanical ventilator parameters, their bounds and their nominal values. See text below for explanation.

Physiological parameters	Acceptable values	Desired
P_{alv} [kPa]	< 5	0
P_{aO2} [kPa]	> 8	10
P_{aCO2} [kPa]	> 4, < 8	5

Table 4.7: Selected physiological indicators, their acceptable ranges and desired values. Palv refers to peak pressure in kilopascals (kPa) above atmospheric pressure.

Maximum and/or minimum allowable values are also defined for several key physiological indicators which are returned as outputs by the model (Table 4.7). To monitor effective arterial oxygenation, the partial pressure of oxygen, P_{aO2}, needs to be considered. In order to maintain effective arterial oxygenation, P_{aO2} is constrained to be higher than 8 kpa. Arterial partial pressure of carbon dioxide, P_{aCO2} is another key indicator of alveolar ventilation which also indirectly reflects acid-base balance. P_{aCO2} is constrained to be between 4 kpa and 10 kpa. The

peak alveolar pressure, (P_{alv}, kpa above atmospheric pressure) is the maximum pressure inside an alveolar compartment, and thus is a risk indicator for barotrauma. It is limited to an upper bound of 5 kPa. The values of all these parameters are recorded after a physiological time simulation of 20 minutes. At the end of the 20 minute period, mean values of P_{aO2} and P_{aCO2} are computed over the previous 2 minutes (representing the final 10% of data). P_{alv} is calculated as the average of the peak pressure in the most highly-pressurised 20% of the 100 alveolar compartments over the final 2 minutes of the simulation. Ventilation perfusion plots are drawn with total ventilation and perfusion recorded across the 100 alveolar units at the end of a simulation.

4.5   Modelling healthy and diseased lung states

Prior to investigating the problem of obtaining optimal mechanical ventilator settings, appropriate disease states need to be simulated *in silico* by configuring the model as follows. The primary goal of the respiratory system is the transfer of gases into and from the blood across the alveolar membrane. To satisfy this requirement optimally, the amount of gas (ventilation, V) and blood (perfusion, Q) taking part in the gas exchange needs to be matched. In many pulmonary disorders, alveolar mismatching can be the most influential mechanism for abnormal arterial blood gases.

Figure 4.2 shows the V/Q plot produced by the model when configured to represent a healthy lung (Wagner 1974). Each data point shows a specific amount of ventilation or perfusion at a respective V/Q ratio. Lines are drawn as a visual aide only. Total flows can be calculated by adding the individual data points. As shown, the healthy lung V/Q distribution is unimodal, approximately symmetric, and clustered fairly tightly around a V/Q ratio of 1. To simulate a disease state, the model is configured by altering the inlet resistances of the alveolar

compartments and the vascular resistances of the pulmonary capillary compartments, in order to create a ventilation perfusion mismatch. An example of a resulting disease configuration is given in Figure 4.2.
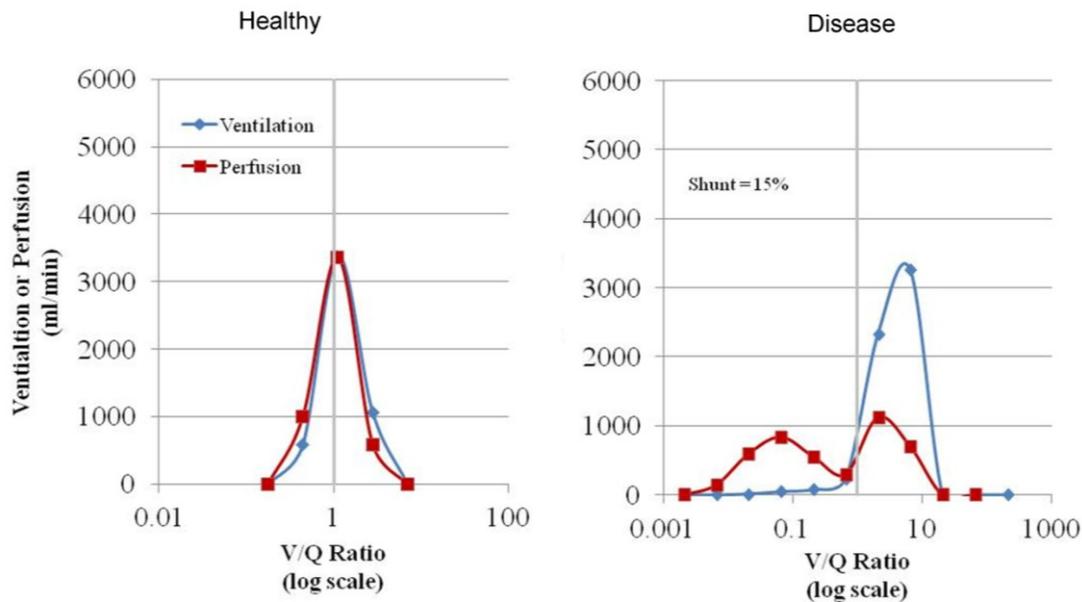
Figure 4.2: Ventilation Perfusion distribution (V/Q) of the simulated healthy lung (left) and an example of a diseased lung (right)

4.6   Multi-objective optimisation of mechanical ventilation

A schematic representation of the proposed *in silico* optimisation framework is shown in Figure 4.3. The simulator represents the combined mechanical ventilator and physiological model described in Chapter 2. The input to the model is vector x = (Vtidal,VentRate, I:E, PEEP, FIO$_2$)$^{\mathrm{T}}$. The output from the model, $y = f(x)$, is the patient data consisting of partial pressure of oxygen in blood (P_{aO2}[kpa]), partial pressure of carbon dioxide in blood (P_{aCO2}, [kpa]) and the alveolar pressure (P_{alv}, [kpa]). The data is collected over a 20 minute simulation period.
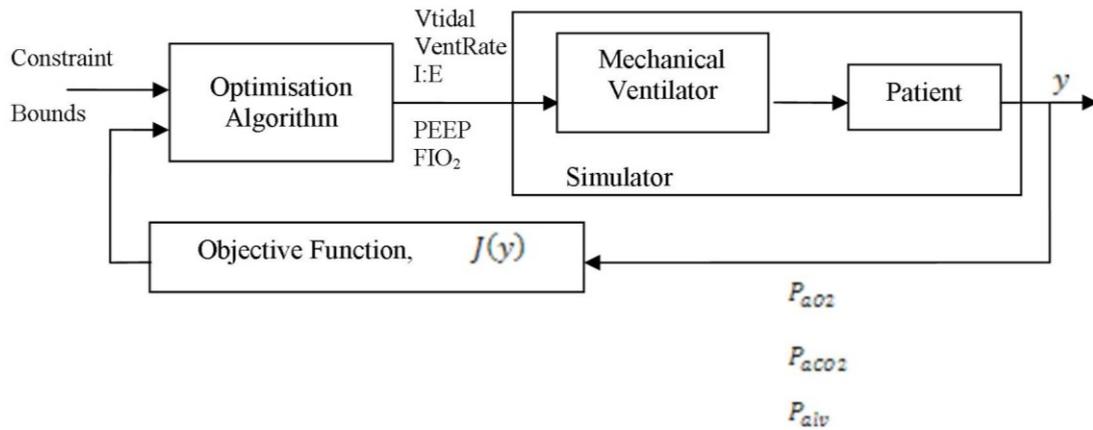
Figure 4.3: Schematic diagram of the optimisation framework.

The optimisation algorithm uses the objective function J to determine the next input vector x_t from the feasible space defined by the bounds and the constraints. The optimisation problem can be formulated mathematically as follows:

$$\text{Minimise} \quad J(f(x)) \quad\quad\quad [4.6]$$

subject to the following constraints (taken from Table 4.3):

$$4 \leq P_{aCO2} \leq 8 \quad\quad\quad [4.7]$$

$$P_{aO2} \geq 8 \quad\quad\quad [4.8]$$

$$P_{alv} \leq 5 \quad\quad\quad [4.9]$$

$$x^L \leq x \leq x^U \quad\quad\quad [4.10]$$

In this formulation, J represents the objective function for the optimisation problem, $f(x)$ represents the simulation model, and x is a vector containing the values of the optimisation parameters (the ventilator settings from Table 4.6). The optimisation algorithm searches for a

solution x, in the decision space defined in Table 4.6 by (x^L) and upper (x^U) bounds, which minimise J. Furthermore, the feasible space is also defined by the constraints [4.7-4.9].

4.6.1   Weighted Aggregate Objective Function approach

To capture the trade-offs inherent in the problem, the objective function J is defined as:

$$J = w_1 J_1 + w_2 J_2 \qquad\qquad [4.11]$$

where
$$J_1 = \frac{\text{Palv}}{r_1} \qquad\qquad [4.12]$$

and
$$J_2 = \left(\frac{|P_{aCO2} - 5|}{r_2} + \frac{|P_{aO2} - 10|}{r_3} \right) \qquad\qquad [4.13]$$

Thus large values of J_1 will be produced by combinations of ventilator settings which cause high peak alveolar pressures (and hence increase the risk of VALI) while large values of J_2 reflect settings that provide poor gas exchange. The parameters r_1, r_2 and r_3 are used to approximately normalise the different output parameters and are determined by the maximum difference between acceptable values and desired values given in Table 4.7. Therefore using Table 4.3, we set $r_1 = 5$, $r_2 = 3$ and $r_3 = 2$. The WAOF approach is reflected by the use of weights w_1 and w_2 whose default values are set equal to one but whose values can be adjusted by the user to explore the effects of prioritising different criteria on the computed ventilator settings. For example, increasing w_2 relative to w_1 causes the optimisation algorithm to search for solutions that minimise J_2 more than J_1 . If the patient is exhibiting lower than required oxygenation, the physician thus has the option to tailor the optimisation towards improving oxygenation at the cost of generating a higher value of P_{alv}.

Since the values of the optimisation parameters vary significantly in magnitude (e.g. Vtidal is several hundred times larger than IE), they must also be normalised to avoid numerical

problems with the chosen optimisation routines. All five parameters are thus allowed to vary within their bounds (see Table 4.6), normalised between -1 and 1.

The WAOF approach allows the user to search for different optimal ventilator settings by adjusting the weights w_1 and w_2 in Eqn. 4.10 and re-running the optimisation algorithm. For example, Table 4.7 shows the effect on the optimal settings for the disease case in Figure 4.2 of making w_1 equal to 10 while leaving w_2 at its nominal value of 1. This has the effect of making the value of P_{alv} the dominant term in the objective function, so that the optimisation algorithm tries to find settings that reduce its value by a larger amount. As shown in Table 4.8, the effect of this is to reduce P_{alv} from 2.989 to 2.28 kpa, at the cost of a small increase in P_{aCO2} and a small decrease in P_{aO2} (although both these parameters still remain within their allowable ranges).

	Optimal input settings					Outputs		
Weight w_1	Vtidal [ml]	VentRate [bpm]	I:E	PEEP [cmH$_2$O]	FIO2	P_{alv} [kpa]	P_{aCO2} [kpa]	P_{aO2} [kpa]
$w_1 = 1$	359	15.07	0.35	4.5	0.32	2.989	4.8958	10.1109
$w_1 = 10$	375	13.75	0.29	3.6	0.25	2.28	6.3608	8.602

Table 4.8: Comparison of results for disease case in Figure 4.1 with different objective function weights using the WAOF approach ($w_2 = 1$).

4.6.2   Comparison of the WAOF approach to NSGA-II

To further investigate the trade-offs involved in this problem, the multi-objective NSGA-II algorithm is executed using a population size of 30, for 100 generations and the objective functions J_1 and J_2 of [4.12-4.13]. The set of 30 solutions returned by the algorithm, together with the optimal solutions of the WAOF approach given in Table 4.8, are shown in Figure 4.4.

It is evident that the previous solutions lie on the appropriate sections of the non-dominated front computed by the NSGA-II approach.

There is a significant general disadvantage of the WAOF method compared to NSGA-II that is clearly illustrated by Figure 4.4. The WAOF approach relies on the user to supply weights whereas the NSGA-II can compute uniformly distributed solutions in the objective space without any extra effort on the part of the user. NSGA-II also offers several solutions with varying trade-off between objectives and clearly illustrates the trade-off needed to achieve the conflicting objectives. Relative to NSGA-II, it is also difficult to obtain optimal solutions in a desired region of the objective space using WAOF approach. For example in Figure 4.4, the point A represents the optimal solution derived using weights $w_1 = w_2 = 1$. With both the weights equal, the point could have been expected to lie nearer the origin than was actually the case in Figure 4.4. This is an unavoidable result of the normalisation of different objective functions in the WAOF approach.

However the WAOF approach also has its advantages. The concept is intuitive, easy to apply and flexible such that any single objective optimisation algorithm can be utilised. If the optimisation algorithm is tailored to the optimisation problem, for example through the use of a hybrid local/global method, the WAOF obtained optimal result can be a better solution than solutions generated from NSGA-II for a lower overall computational overhead. For example, the solution 'A' in Figure 4.4 is a non-dominated solution computed after 1100 function evaluations compared to the set of solutions computed by NSGA-II after 3000 function evaluations. The WAOF approach, with its weights that can be 'tuned' to a user's requirement, provides greater control over the optimisation search than NSGA-II. In an environment where the physician's final decision on the ventilator settings is crucial, the WAOF approach allows the physician to choose and adjust the weights associated with different objectives pertinent to

the specific scenario. In subsequent sections, the search for the global optimal solution is conducted using the WAOF approach.
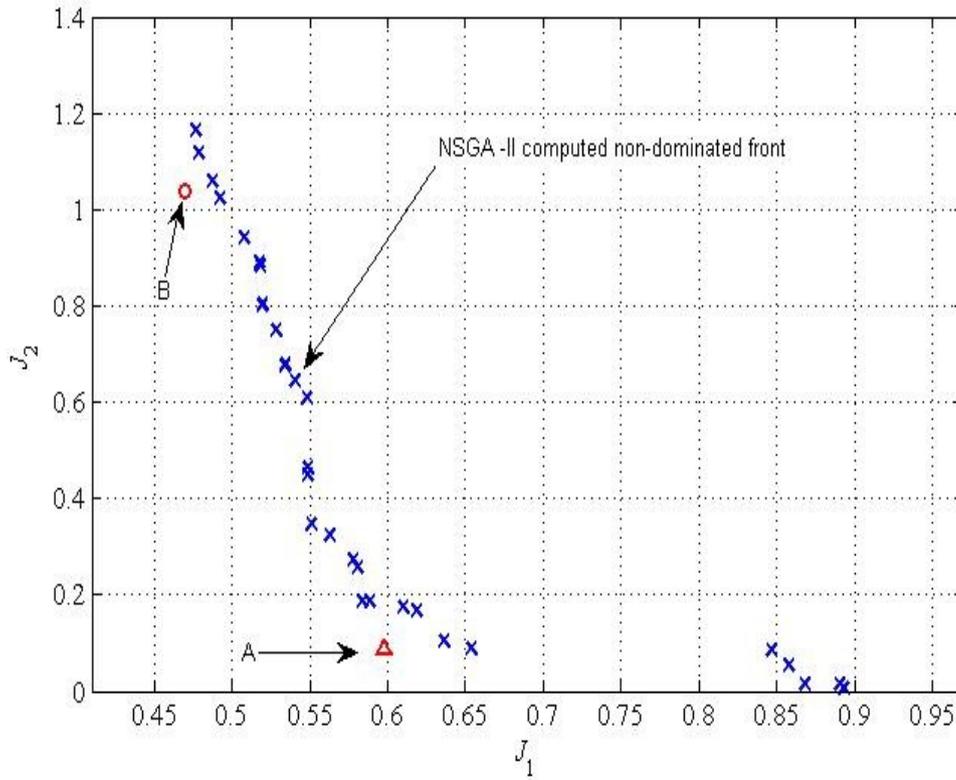
Figure 4.4: Set of optimal solutions generated by NSGA-II compared to the results of the WAOF approach (Table 4.8). Point 'A' is the solution with weights $w_1 = w_2 = 1$ and point 'B' is the solution with weights $w_1 = 1$ and $w_2 = 10$.

4.7   Computing optimised ventilator settings

	P_{alv} [kpa]	P_{aCO2} [kpa]	P_{aO2} [kpa]
Disease A	1.03	5.81	6.94
Disease B	1.16	5.98	4.98
Disease C	1.03	7.05	1.81
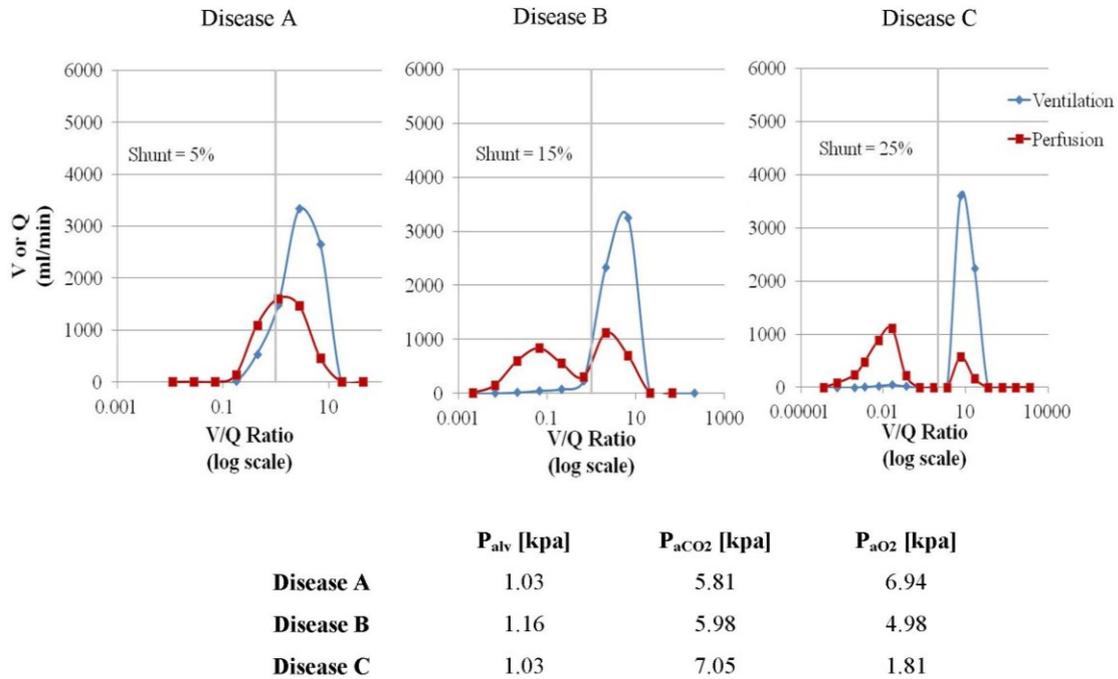
Figure 4.5: Ventilation Perfusion curves of different simulated pulmonary diseases describing variation in distribution of gas exchange across the lung.

Values for shunt refer to total calculated shunt fraction. The respective physiological parameters for the disease cases under nominal ventilator settings are also shown.

To simulate a disease state, the model is configured by altering the inlet resistances of the alveolar compartments and the vascular resistances of the pulmonary capillary compartments, in order to create a ventilation perfusion mismatch. The resulting V/Q mismatch for three different disease states, of progressively increasing severity, is shown in Figure 4.5.

Figure 4.5 also provides the values of the physiological parameters produced under these disease cases when the MV parameters remain at their nominal values. It is clearly evident that as the severity of the disease increases, the arterial oxygen partial pressure falls and arterial

carbon dioxide rises. Indeed, the value of P_{aO2} in all three cases is lower than its minimum allowable value.

	Optimal settings					Outputs		
	Vtidal [ml]	VentRate [bpm]	I:E	PEEP [cmH₂O]	FIO2	P_{alv} [kpa]	P_{aCO2} [kpa]	P_{aO2} [kpa]
Nominal	400	12.50	0.33	0	0.21	0.799	5.440	13.029
Optimised	281	11.6	0.34	0	0.21	0.537	6.152	9.989

Table 4.9: Nominal and optimal settings and the resulting physiological parameters for the healthy lung model

A comparison of the results obtained with the nominal and optimised ventilator parameters for the healthy lung model is shown in Table 4.9. Compared with the nominal settings, the optimisation algorithm has sought to find a combination of ventilator parameters that pushes down the value of P_{alv}, at the cost of slightly increasing P_{aCO2} and reducing P_{aO2} (while still keeping these parameters within their specified bounds). It has achieved this principally by reducing Vtidal, as would be expected from normal clinical practice. Of course, "re-balanced" settings which give a higher priority to gas exchange could easily be computed by adjusting the weights w_1 and w_2 in the objective function [4.11] for the optimisation algorithm.
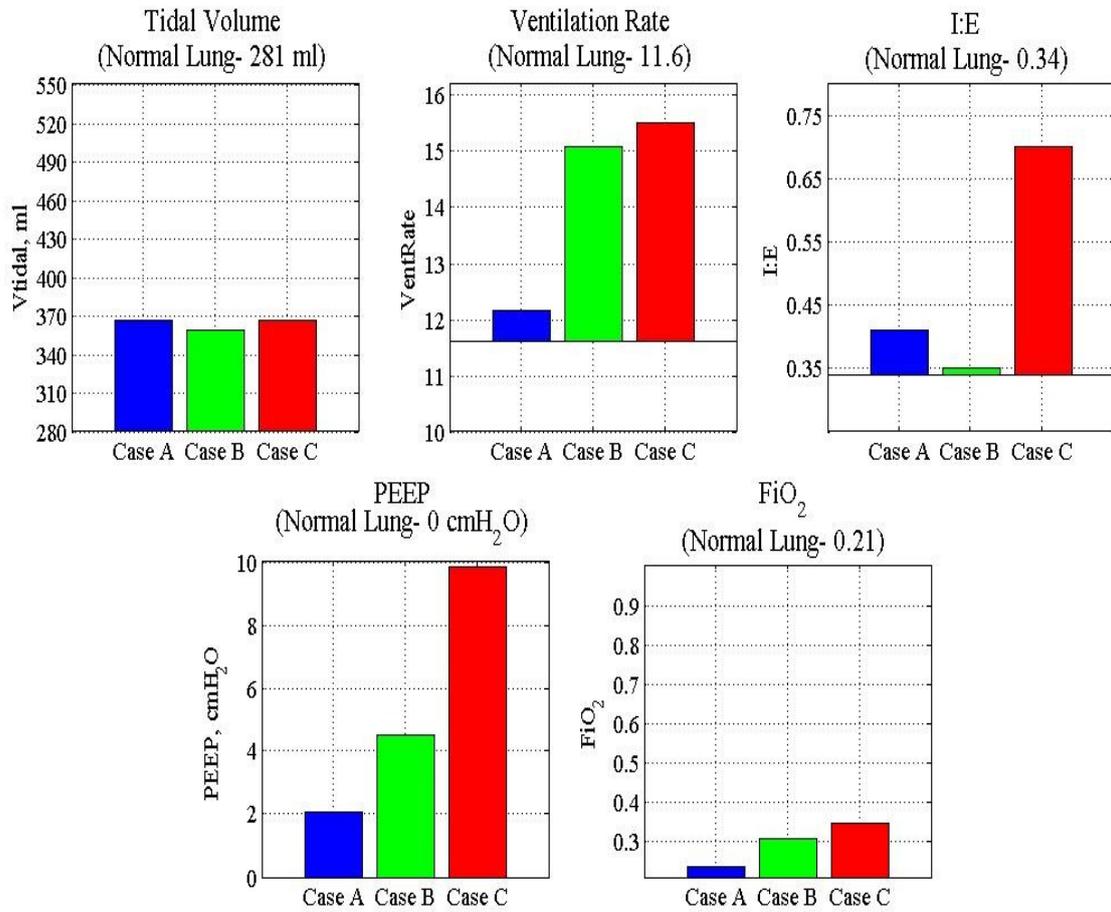
Figure 4.6: Change in optimal MV settings for Disease Cases A, B and C relative to the optimal settings for a healthy lung.

Figure 4.6 shows the different optimal parameter settings computed for each of the three disease cases shown in Figure 4.5. For the purposes of comparison, each plot title gives the optimal value found for that parameter for the healthy lung, and the bar graphs are drawn relative to this value. Table 4.10 summarises the values of the optimal settings for each individual disease case with their corresponding physiological parameters. When compared with the results of applying the nominal settings for each disease state (shown in Figure 4.5), it is clear that the optimised settings provide a marked improvement in the values of the physiological parameters. In particular, note that the value of P_{aO2} has now increased to lie within its allowable range for all

three cases. This improvement in oxygenation has been achieved at the cost of an increase in the value of P_{alv}, but even for the most severe case P_{alv} still remains lower than its maximum limit.

	Optimal settings					Outputs		
	Vtidal [ml]	VentRate [bpm]	I:E	PEEP [cmH$_2$O]	FIO2	P_{alv} [kpa]	P_{aCO2} [kpa]	P_{aO2} [kpa]
Disease A	366	12.50	0.41	2.1	0.24	2.1286	5.0608	9.9999
Disease B	359	15.07	0.35	4.5	0.32	2.9890	4.8958	10.1109
Disease C	366	15.49	0.7	9.8	0.35	4.8874	5.4419	8.4193

Table 4.10: Optimal settings and the resulting physiological parameters for three disease cases

Poor oxygenation of the blood and resultant hypoxia can be fatal to patients in an ITU. Carbon dioxide (CO_2) clearance from the blood is also essential as excessive CO_2 in the blood can lead to acid base imbalance, risking serious damage to the major organ systems. Traditionally, MV operators have focused on ensuring that oxygenation requirements are adequately met and carbon dioxide clearance is maintained. In recent years, however, the increasing incidence of VALI, which has been attributed to the occurrence of high pressures and volumes in the alveoli, has led MV operators be more cautious and to consider other physiological factors which previously tended to be ignored. The objective of minimising the risk of VALI can often be in direct conflict with the traditional aim of ensuring good gas exchange, and optimal settings that satisfy both sets of criteria can be difficult to predict.

The investigation of systematic approaches for choosing ventilator settings in order to maintain desired blood-gas partial pressures while minimising the risk factors associated with VALI has been the subject of intense interest in the medical community in recent years. However, there are many factors hampering traditional approaches to research in this area, including (a) the

difficulty or impossibility of measuring key patient variables reflecting the effects of mechanical ventilation *in vivo*, (b) the inhomogeneity of physiological responses over time and over the patient population, (c) the number of different ventilator parameters which may be adjusted by the clinician, (d) the lack of clarity over which physiological parameters represent the most important risk factors for VALI, and (e) the ethical constraints on research on patients who are unable to give their own consent. VALI itself has been difficult to diagnose, with limited real time data available, and thus it represents a serious risk to patient welfare. MV input settings also tend to be highly correlated, which can make their effect on internal patient dynamics unpredictable. As a result, limited progress has been made, with ventilator parameters still typically being set using heuristic approaches which are heavily influenced by the clinician's ability and experience, and large local, national and international variations existing in treatment protocols.

In this study three different lung conditions were simulated, representing disease states of increasing severity, which could be realistically found in patients in an ICU. The disease states were obtained by altering the ventilation and perfusion to gas exchanging units, and changing the amount of blood flow not involved in gas exchange (shunt). Higher bronchiolar resistance would decrease the flow of air to these sections of the lung and cause an increase in ventilation to the other parts of the lung. Increased vascular resistance would transfer the blood flow to other capillary compartments. Increased shunt would reduce the amount of blood in pulmonary capillaries. These would result in ventilation perfusion (V/Q) mismatch shown through V/Q plots. Many pulmonary disorders have been shown to exhibit unique V/Q behaviour - a small but representative sample of possible V/Q distributions have been considered in this chapter. Interestingly, the optimal settings computed here not only optimise the particular parameters contained in the objective function, but also act to more generally improve the gas exchange within the lung. Figure 4.7 compares the V/Q distribution under nominal settings with the V/Q

distributions under optimal settings for the three disease cases. The improvement is evident in all cases, where the optimal setting causes the gas exchange to shift towards a V/Q ratio of 1. Under nominal settings, there are a large number of perfused alveoli where little gas exchange occurs due to under-ventilated or collapsed airways. From Figure 4.6 and Table 4.10, it is evident that the higher values of ventilator rate, I:E ratio and PEEP act to improve the ventilation to these areas of the lung where there is poor gas flow, which in turn improves the V/Q mismatch and help achieves better gas exchange.

The results of this study highlight the significant differences which can exist between settings that optimally satisfy the conflicting objectives of mechanical ventilation, even for relatively minor differences in lung compositions (see the optimal settings for disease case A relative to those for the healthy lung). The ability to target MV therapy to particular diseases, and to particular severity of diseases, represents significant progress towards the goal of personalised ventilator management for critically ill patients in the modern ITU, where it could help to avoid lung injury, exacerbation of underlying disorders and allow faster weaning from MV.
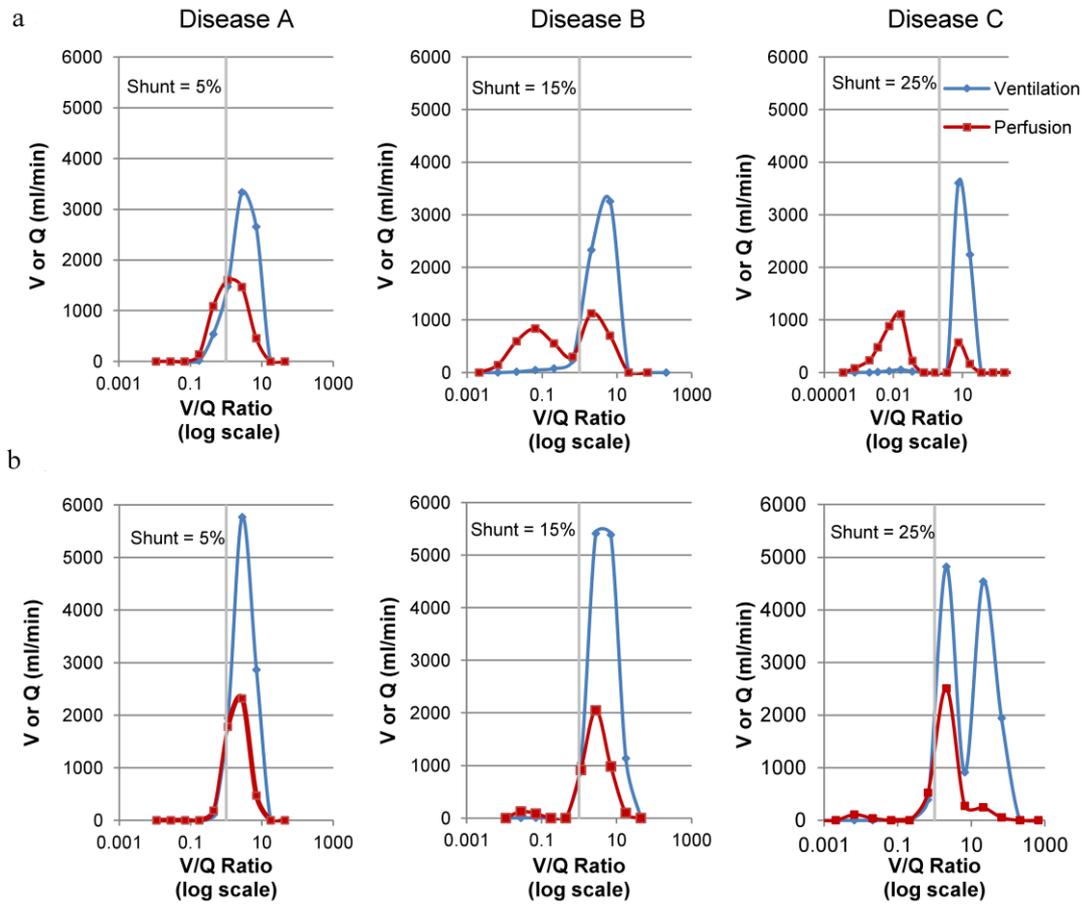
Figure 4.7: Comparison of ventilation perfusion under nominal and optimal settings for each disease state. (a) displays the VQ distribution under nominal settings and (b) displays the VQ distribution under optimal settings.

Chapter 5: Conclusions and Future Work

5.1   Conclusions

This thesis has described the extended implementation and development of a comprehensive computational model of pulmonary (patho)physiology. A detailed description is presented of the mathematical and physiological concepts that are incorporated in the model. The model has been reviewed in comparison with a cross-section of other models found in literature. The model has been newly implemented in the MALTAB environment which provides access to numerous computational toolboxes for further model analysis and novel applications. The thesis also proposes modifications (based on physiological mechanisms) to adequately characterise healthy as well as diseased lung behaviour. These modifications address two distinct faults within the previous model: 1) the inability of the model to accurately represent *collapse* within the lungs under disease states and 2) the incomplete response of the model to positive end expiratory pressure (PEEP). These faults are illustrated, and solutions identified through consultations with clinical practitioners are proposed.

Current model validation strategy in clinical simulators is based on the statistical evaluation of the model responses in comparison to data obtained from *in vivo* investigations. Typically, if model responses are found to lie within the 95% of the population observations, then the model is considered validated. Uncertainty within heterogeneous datasets parameters and un-modelled dynamics can significantly hinder the applicability of the model to accurately generate

predictions across patient populations. Further rigorous procedures for model validation and verification are needed to enable the wider clinical exploitation of *in silico* simulation models.

This thesis proposes an optimisation-based approach for the validation of an *in silico* model of pulmonary physiology. The explicit modelling of uncertainty/variability in the internal parameters of the model is combined with advanced global optimisation methods to demonstrate that the model predictions never deviate from physiologically plausible values for realistic levels of parametric uncertainty. The approach to uncertainty modelling is adapted from current best-practice in the field of systems and control Engineering, and a range of advanced optimisation methods are employed to check the robustness of the model, including Sequential Quadratic Programming, Mesh Adaptive Direct Search, Differential Evolution 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.

Unlike traditional validation of physiological simulators, the aim of this approach is to formally quantify the model's robustness (i.e. its ability to always produce realistic results for all expected levels of uncertainty in internal model parameters). Through this approach, the simulator that has been validated across the expected range of variation that would be found in multiple patients, which represents a much more stringent test of the underlying model assumptions than simple tests on the ability of the model to match a single set of data. If no anomalous model responses exist for all levels of uncertainty in the model parameters, then the model can be expected to be robust to uncertainty in those parameters. If one of 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, if it corresponds to a particular condition which would be likely to produce unexpected responses,

etc. Expert analysis and experience cannot be replaced by an automated computer program – rather a computational analysis method is proposed which can efficiently uncover anomalous model responses for further investigation by the clinical user. The results of the study indicate that NPS provides robust predictions of arterial gas pressures for all realistic ranges of given model parameters, and also demonstrate the general applicability of the proposed approach to model validation for physiological simulation.

The thesis also addresses the problem of selection of appropriate mechanical ventilator settings to attain adequate oxygenation and carbon dioxide clearance while minimising the risk of ventilator-associated lung injury. The problem is a significant challenge for physicians, who currently rely on experience and guidelines based on recommendations from a limited number of *in vivo* studies whose data are typically more applicable to fictitious "average" members of populations than to individual patients. The thesis uses validated computational models of pulmonary pathophysiology to design *in silico* experiments to determine optimal combinations of mechanical ventilator settings for individual patient and disease states. The set of optimal solutions from the multi-objective algorithm NSGA-II clearly illustrate the conflicting objectives that need to be achieved.

Different types of pulmonary disease states were simulated by manipulating gas exchange parameters in a validated computational model to give appropriate ventilation perfusion distributions. Using advanced global optimisation algorithms, optimal ventilator parameters were computed for each case. The results revealed significant differences in optimal settings both between healthy and diseased states, and between different disease states.

The proposed approach offers clinicians a tool with which to conduct *in silico* trials on strictly-compliant virtual patients, in order to more easily assess and manage the trade-offs involved in determining appropriate patient- and disease-specific ventilator settings. The optimal settings

are shown to satisfy the conflicting clinical objectives, to improve the ventilation perfusion matching within the lung, and, crucially, to be disease-specific.

5.2   Future work

Evaluation of predicted optimal mechanical ventilator settings using patient data.

Results in Chapter 4 clearly identify the benefits of the proposed optimisation based strategy in deducing optimal mechanical ventilation strategy for specific disease case scenarios *in silico*. The strategy could be evaluated further by establishing the validity of model responses to optimal ventilation predictions in a clinical setting. Specific patient data with known ventilation perfusion distribution could be collected and matched to a virtual patient *in silico*.   Under careful observation and assessing real time patient data, optimal ventilator settings would be applied to the patient. With the optimisation parameters bounded within the suggested guidelines of mechanical ventilation protocols, the risk to the patient should be minimal. If successful, the optimisation based mechanical ventilation strategy has the potential to significantly improve current critical care in the ITU.

Validation and application of population based models using ensemble forecasting

The use of population-based models to inform and guide clinical practice requires the ability to make reliable predictions of the effect of changing model inputs, initial conditions or parameters on the behaviour of a population of heterogeneous patients, represented by a set (ensemble) of models. This is a directly analogous problem to that which exists in the field of climate science, where researchers are required to make robust predictions about the future climatic conditions based on ensembles of models which are subject to significant levels of uncertainty. Recent years have seen the development of extremely powerful ensemble forecasting techniques

Conclusions and Future Work

allowing for more reliable predictions of mean and extreme behaviour (Joliffe 2003). Techniques such as reliability diagrams, probabilistic ROC scores etc) can help quantify the effects of uncertainty arising from a) unknown model parameters and b) variation of individuals within a population.

The validated physiological model from this thesis can be applied to conduct a formal examination of novel strategies for mechanical ventilation. The strategies to be developed will be required to explicitly consider uncertainty and variability across patient populations, thus motivating the use of statistical prediction techniques from ensemble forecasting.

Investigation of tissue perfusion using cardiopulmonary modelling

Poor tissue perfusion is most often caused by low cardiac output or occasionally by pathological vasodilation; these states are typically referred to as "shock states" (e.g. cardiogenic shock, septic shock). These shock states form a large part of organ level pathology of critical illness. Acute failure of organ perfusion is treated currently using drugs and fluids affecting multiple organ systems, with widespread and often unhelpful side effects. In clinical practice there is significant interaction between cardiac function, pulmonary function and the behaviour of the systemic vasculature, such that treatments intended to improve function in one part of the system may have detrimental effects on others. Clinicians therefore inevitably have to adopt compromised approaches based on theory and experience. The within-organ, real-time perfusion and intracellular oxygenation cannot be quantified using present measurement techniques which limits the potential of clinical or animal-model research, and may only provide surrogate outcomes with limited direct relevance to human pathology. In human studies, the techniques used and outcomes measured have varied considerably, making comparisons difficult (Conway 2002, Wakeling 2005). All the above have seriously hampered the development of protocols

that allow optimal clinical management of shock states. There is increasing interest in protocol-guided management of at-risk patients but the pathology is still poorly understood.

The model of the respiratory system developed in the thesis, when integrated with high-fidelity models of cardiac, vascular and peripheral tissue system, will provide a powerful platform for the investigation of poor tissue perfusion in pathological states. To ensure the applicability of the model to real medical scenarios, the model will be subjected to an extensive programme of verification and validation, using techniques such as the optimisation-based model validation approach presented in this thesis. The integrated model will provide a powerful and cost effective means of investigating clinically applicable therapeutic strategies for managing shock states. Specific issues that need to be explored include:

1. Gas exchange in tissues:

   a. The effect of low cardiac output on oxygen in variable cell-depth tissue.

   b. The effect of flow impairment in pathological states such as septic shock.

   c. The effect on cellular oxygenation due to impaired pulmonary gas exchange during shock states.

2. Ventilatory manipulation in shock states:

   a. The effect of ventilator manipulation on tissue oxygenation.

   b. The effect of ventilation strategies such as permissive hypercapnia on cellular oxygenation and pH levels during low flow conditions.

3. Contents of blood:

   a. The effect of haemoglobin concentration on tissue oxygenation during shock states.

   b. The effect of blood viscosity on tissue oxygenation during shock states.

4. Cardiac performance:

a. The effect of pharmacologically–enhanced myocardial contractility affect tissue oxygenation

b. The effect of infused intravenous fluid filling (crystalloid vs colloid) on tissue oxygenation.

Appendix

A – Pre and post-intervention ventilator settings to validate modified model

Pre - Intervention			Post- intervention		
Inspired oxygen fraction	Tidal Volume, [ml]	Ventilatory rate	Inspired oxygen fraction	Tidal Volume, [ml]	Ventilatory rate
0.45	780	15	0.6	780	15
0.6	430	13	0.4	430	13
0.7	600	15	0.85	600	15
0.5	625	8	0.4	625	8
0.6	450	17	0.5	450	17
0.5	683	12	0.75	729	14

B – Termination conditions for optimisation algorithms - TolFun and TolX
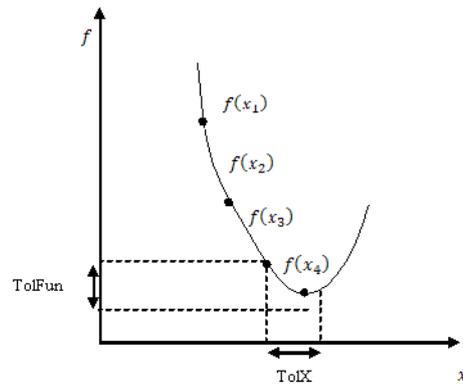
Algorithm terminates if last step is smaller than TolX or the improvement in function is less than TolFun. Specifically:

The algorithm terminates if $|x_{c+1} - x_c| <$ TolX where x_c is the current decision vector.

The algorithm terminates if $|f(x_{c+1}) - f(x_c)| <$ TolFun.

Bibliography

Amato MBP, Barbas CSV, Medeiros DM, Magaldi RB, Schettino GP, Lorenzi-Filho G, Kairalla RA, Deheinzelin D, Munoz C, Oliveira R. Effect of a Protective-Ventilation Strategy on Mortality in the Acute Respiratory Distress Syndrome. *New England Journal of Medicine* 1998; 338(6): 347-354.

Anzueto A, Frutos–Vivar F, Esteban A, Alía I, Brochard L, Stewart T, Benito S, Tobin MJ, Elizalde J, Palizas F, David CM, Pimentel J, González 1M, Soto L, Dempaire G, Pelosi P. Incidence, risk factors and outcome of barotrauma in mechanically ventilated patients. *Intensive Care Medicine* 2004; 30(4): 612-619.

ARDSnet - Acute Respiratory Distress Syndrome Network. Ventilation with lower tidal volumes as compared with traditional tidal volumes for acute lung injury and the acute respiratory distress syndrome *The New England Journal of Medicine* 2000*;* 342(18):1301–1308.

ARDSnet. TheNational Heart, Lung, and Blood Institute ARDS Clinical Trials Network. Higher versus Lower Positive End-Expiratory Pressures in Patients with Acute respiratory Distress Syndrome. *New England Journal Medicine* 2004; 351; 327-36.

Athansiades A, Ghorbel F, Clark JW, Niranjan SC, Olansen J, Zwischenberger JB, Bidani A. Energy analysis of a nonlinear model of the normal human lung. *Journal of Biological Systems* 2000; 8(2): 115-139.

Audet C, Dennis JE. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization* 2006; 17: 188-217.

Auten RL, Vozzelli M, Clark RH. Volutrauma. What is it, and how do we avoid it? *Clin Perinatol* 2001; 28(3): 505–515.

Bland JM, Altman DG. Statistical method for assessing agreement between two methods of clinical measurement. *The Lancet* 1986; i: 307-310.

Bland JM, Altman DG. Measuring agreement in method comparison studies. *Statistical Methods in Medical Research* 1999; 8:135-160.

Bouadma L, Lellouche F, Cabello B, Taillé S, Mancebo J, Dojat M, Brochard L. Computer-driven management of prolonged mechanical ventilation and weaning: a pilot study. *Intensive care medicine* 2005; 31(10): 1446-1450.

Chase JG, Suhaimi F, Penning S, Lin J, LeCompte A, Preiser JC, Pretty CG, Shaw G, Moorehead KT, Desaive T. Validation of a model-based virtual trials methods for tight glycemic control in intensive care. *Biomedical Engineering Online* 2010; 9: 84.

Chase JG, Le Compte AJ, Preiser JC, Shaw GM, Penning S, Desaive T. Physiological modelling, tight glycemic control and the ICU clinician: what are models and how can they affect practice? *Annals of Intensive Care* 2011; 1: 11.

Chatburn RL. Classification of Mechanical Ventilators. Tobin MJ (eds) In Principles and Practice of Mechcanical Ventilation 2nd edition. McGraw-Hill. 2006:37-52.

Cheng W, DeLong DS, Franz GN, Petsonk EL, Frazer DG. Contribution of opening and closing of lung units to lung hysteresis. *Respiratory Physiology* 1995; 102(2-3):205-15.

Chiari L, Avanzolini G, Ursino M. A comprehensive simulator of the human respiratory system: validation with experimental and simulated data. *Annals of Biomedical Engineering* 1997: 25(6); 985-999.

Cole AGH, Weller SF, Sykes MK. Inverse ratio ventilation compared with PEEP in adult respiratory failure. *Intensive Care Medicine* 1984; 10(5): 227-232.

Concato J, Feinstein AR, Holford TR. The risk of determining risk with multivariable models. *Annals Intern Med* 1993; 118: 201–210.

Conway DH, Mayall R, Abdul-Latif MS, Gilligan S, Tackaberry C. Randomised controlled trial investigating the influence of intravenous fluid titration using oesophageal Doppler monitoring during bowel surgery. *Anaesthesia* 2002; 57(9): 845-849.

Cotes JE, Chinn DJ, Miller MR. Lung Function, 6th edition. Blackwell Publishing 2006.

Deb K, Agrawal S, Pratap A, Meyarivan T. A Fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. *IEEE Transactions on Evolutionary Computation* 2002: 6(2); 182-197.

DeProst N, Ricard J, Saumon G, Dreyfuss D. Ventilator-induced lung injury: historical perspectives and clinical implications. *Annals of Intensive Care* 2011; 1:28.

Dojat M, Brochard L, Lemaire F, Harf A. A knowledge-based system for assisted ventilation of patients in intensive care units. *International Journal for Clinical Monitoring Computing* 1992; 9: 239–250.

Douglas AR, Jones NL, Reed JW. Calculation of whole blood CO2 content. *Journal of Applied Physiology* 1988; 65(1): 473-477.

Doyle J, Carlson JM. Complexity and Robustness. *Proceedings of the National Academy of Sciences of the United States of America* 2002; 99: 2538-2545.

Dreyfuss D, Saumon G. Barotrauma is volutrauma, but which volume is the one responsible? *Intensive Care Medicine* 1992; 18(3): 139-141.

Dreyfuss D, Saumon G. From ventilator–induced lung injury to multiple organ dysfunction. *Intensive Care Medicine* 1998; 24(2) 102-104.

Dreyfuss D, Ricard JD, Saumon G. On the physiologic and clinical relevance of lung-borne cytokines during ventilator-induced lung injury. *American Journal of Respiratory and Critical Care Medicine* 2003;167: 1467–1471.

Dunn E, Olague G. Multiobjective sensor planning for efficient and accurate object reconstruction. *Raidl GR et al (eds) In Lecture notes in computer science* 2004; 3005: 312-321. Springer Berlin Heidelberg.

Eldridge F. The North Carolina respiratory model: A multipurpose model for studying the control of breathing, in *Bioengineering Approaches to Pulmonary Physiology and Medicine*. Khoo. MCK. Plenum press, New York 1996: 31-49.

Elliott CG. Pulmonary physiology during pulmonary embolism. *CHEST* 1992;101: 163S-171S.

El-Mihoub TA, Hopgood AA, Nolle L, Battersby A. Hybrid genetic algorithms: a review. *Engineering Letters* 2006; 13(2): EL_13_2_11.

Esteban A, Anzueto A, Frutos F, Alia I, Brochard L, Stewart TE, Benito S, Epstein SK, Apezteguia C, Nightingale P. Characteristics and outcomes in adult patients receiving

mechanical ventilation: a 28-day international study. *Journal of American Medical Association* 2002; 287:345–355.

Esteban A, Ferguson ND, Meade MO, Frutos-Vivar F, Apezteguia C, Brochard L, Raymondos K, Nin N, Hurtado J, Tomicic V, Gonzalez M, Elizalde J, Nightingale P, Abroug F, Pelosi P, Arabi Y, Moreno R, Jibaja M, D'Empaire G, Sandi F, Matamis D, Montanez AM, Anzueto A. Evolution of mechanical ventilation in response to clinical research. American Journal Respiratory and Critical Care Medicine 2008; 177: 170-177.

Farmery A, Roe P. A model to describe the rate of oxyhaemoglobin desaturation during apnoea. *British journal of anaethesia* 1996; 76(2); 284-291.

Fielding C, Varga A, Bennani S, Selier M. Advanced Techniques for the clearance of flight control laws. *Springer Verlag* 2002. ISBN 978-3-540-44054-3.

Fleming P, Purshouse R. Evolutionary algorithms in control systems engineering: a survey. *Control Engineering Practice* 2002; 10(11): 1223-1241.

Fonseca CM, Fleming PJ. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In *Proceedings of the Fifth International Conference on Genetic Algorithms* 1993; 416-423.

Fowler M. UMLDistilled: A brief guide o the standard object modelling language 3[rd] edition. Addison-Wesley Professional 2003.

Gammon R. Pulmonary barotrauma in mechanical ventilation: patterns and risk factors. *Chest* 1992; 102: 568-572.

Garcia-Nieto J, Alba E, Jourdan L, Talbi E. Sensitivity and specificity based multiobjective approach for feature selection: application to cancer diagnosis. *Information Processing Letters* 2009; 109(16): 887-896.

Gatinoni L, Pesenti P. The concept of "baby lung". *Intensive Care Medicine* 2005; 31(6) 776-84.

Gattinoni L, Caironi P, Cressoni M, Chiurnello D, Ranieri VM, Quintel M, Russo S, Patroniti N, Cornejo R, Bugedo G. Lung recruitment in patients with acute distress syndrome. *The New England Journal of Medicine* 2006; 354: 1775-1786.

Goldberg DE. *Genetic algorithms in search and optimization.* Addison-Wesley. 1989.

Grodins FS, Buell J, Bart AJ. Mathematical analysis and digital simulation of the respiratory control system. *Journal of applied physiology* 1967; 22(2): 260-276.

Grodins FS, Gray JS, Schroeder KR, Norins AL, Jones RW. Respiratory responses to CO2 inhalation: a theoretical study of non-linear biological regulator. *Journal of applied physiology* 1954; 7: 283-308.

Gupta AK, Sharma J, Mukhopadhyay P. Optimization method applied to the design of ventilators, *Medical and Biological Engineering and Computing* 1978; 16: 387-396.

Hahn CEW, Black MS, Barton SA, Scott I. Gas exchange in a three-compartment lung model analyzed by forcing sinusoids of N2O. *Journal of Applied Physiology* 1993; 75(4): 1863-1876.

Hahn CEW, Farmery AD. Gas exchange modelling: no more gills please. *British Journal of Anaesthesia* 2003; 91(1): 2-15.

Han S-P, Superlinearly convergent variable metric algorithms for general nonlinear programming problems. *Mathematical Programming* 1976; 11; 260.

Hardman JG, Bedforth N, Ahmed A, Mahajan R, Aitkenhead A. A physiology simulator: validation of its respiratory components and its ability to predict the patient's response to changes in mechanical ventilation. *British Journal of Anaesthesia* 1998; 81(3): 327-332.

Hardman J, Bedforth N. Estimating venous admixture using a physiological simulator. *British Journal of Anaesthesia* 1999b; 82: 346-349.

Hardman JG, Aitkenhead A. Estimation of Alveolar Deadspace Fraction Using Arterial and End-Tidal CO2: A Factor Analysis Using a Physiological Simulation. *Anaesthesia and Intensive Care* 1999a; 27: 452-458.

Hardman JG. Respiratory Physiological Modelling: The design, construction, validation and application of a set of original respiratory physiological models. *Doctoral Thesis*, University of Nottingham, 2001.

Hardman JG, Wills JS, Aitkenhead AR. Investigating Hypoxemia during Apnea: Validation of a Set of Physiological Models. *Anesthesia & Analgesia* 2000; 90(3): 614-618.

Hardman JG, Ross JJ. Modelling: a core technique in anaesthesia and critical care research. *British Journal of Anaesthesia* 2006; 97(5):589-592.

Hardman JG, Wills JS. The development of hypoxaemia during apnoea in children: a computational modelling investigation. *British Journal of Anaesthesia* 2006; 97(4): 564-570.

Hardman JG, Moppett IK, Mahajan RP. Validity, credibility and applicability: the rise and rise of the surrogate. *British Journal of Anaesthesia* 2008; 101(5): 595-596.

Harrell FE, Lee KL, Mark DB. Multivariable prognostic models: Issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in Medicine* 1996; 15: 361–387.

Hedenstierna G, Rothen HU. Atelectasis Formation during Anesthesia: Causes and Measures to Prevent It. *Journal of Clinical Monitoring and Computing* 2000; 16: 329-335.

Heistracher T, Hofmann W. Flow and deposition patterns in successive airway bifurcations. *Annals of Occupational Hygiene* 1997; 41(1): 537-542.

Henderson LJ. The equilibrium between oxygen and carbonic acid in blood. *Journal of Biological Chemistry* 1920; 41(3): 401-430.

Hickam JB, Cargill WH. Effect of exercise on cardiac output and pulmonary arterial pressure in normal persons and in patients with cardiovascular disease and pulmonary emphysema. *Journal of Clinical Investigation* 1948; 27: 10-23.

Hickling KG. The Pressure-Volume Curve Is Greatly Modified by Recruitment: A Mathematical Model of ARDS Lungs. *American Journal of Respiratory and Critical Care Medicine* 1998; 158(1): 194-202.

Hill NS, Levy MM (Eds). Ventilator management strategies for critical care, ISBN: 082470522X, 2001; 158: 46.

Hinds CJ, Ingram D, Adams L, Cole PV, Dickinson CJ, Kay J, Krapez JR, Williams J. An evaluation of the clinical potential of a comprehensive model of human respiration in artificially ventilated patients. *Clinical Science* 1980; 58: 83-91.

Hinds C, Roberts M, Ingram D, Dickinson C. Computer simulation to predict patient responses to alterations in the ventilation regime. *Intensive Care Medicine* 1984; 10(1): 13-22.

Holets S, Hubmayr RD. Setting the Ventilator. Tobin MJ (eds) In Principles and Practice of Mechanical Ventilation 2nd edition. McGraw-Hill. 2006:163-182.

Hotchkiss JR, Crooke PS, Adams AB, Marini JJ. Implications of a biphasic two-compartment model of constant flow ventilation for the clinical setting. *Journal of critical care* 1994; 9(2): 114-123.

Joliffe IT, Stephenson DB. Forecast Verification: A practitioner's guide in atmospheric science. Wiley and Sons.2003.

Joyce C, Hickling K. Permissive hypercapnia and gas exchange in lungs with high Qs/Qt: a mathematical model. *British Journal of Anaesthesia* 1996; 77(5): 678-683.

Kathirgamanathan A, McCahon RA, Hardman JG. Indices of pulmonary oxygenation in pathological lung states: an investigation using high-fidelity, computational modelling. *British Journal of Anaesthesia* 2009; 103(2):291-297.

Keenan SP, Sinuff T, Cook DJ, Hill NS. Does noninvasive positive pressure ventilation improve outcome in acute hypoxemic respiratory failure? A systematic review. *Critical Care Medicine* 2004; 32(12): 2516–2523.

Kelman G, Nunn J. Nomograms for correction of blood Po2, Pco2, pH, and base excess for time and temperature. *Journal of Applied Physiology* 1966; 21(5): 1484-1490.

Kim J-R, Bates DG, Postlethwaite I, Ma L, Igelsias PA. Robustness analysis of biochemical network models. *IET Systems Biology* 2006; 153(3):96-104.

Kim J-S, Valeyev NV, Postlethwaite I, Heslop-Harrison P, Cho KH, Bates DG. Analysis and extension of a biochemical network model using robust control theory. *International Journal of Robust and Nonlinear Control* 2010; 20(9): 1017-1026.

Krasnogor N, Smith J. A tutorial for competent memetic algorithms: model, taxonomy and design issues. *IEEE Transactions on Evolutionary Computation* 2005: 9(5); 474-488.

Kwok H, Linkens D, Mahfouf M, Mills G. Rule-base derivation for intensive care ventilator control using ANFIS. *Artificial Intelligence in Medicine* 2003; 29(3): 185-201.

Lachmann B. Open Lung in ARDS. *Minerva Anestesiol* 2002; 68(9) 637-42.

Lachmann B. Open up the lung and keep the lung open. *Intensive Care Medicine* 1992; 18(6): 319-321.

Laffey JG, Kavanagh BP, Ney L, Kuebler WM, Oba Y, Salzman GA, Brower RG, Matthay MA, Wheeler A. Ventilation with Lower Tidal Volumes as Compared with Traditional Tidal Volumes for Acute Lung Injury. *New England Journal of Medicine* 2000; 343(11): 812-814.

Laubscher TP, Heinrichs W, Weiler N, Hartmann G, Brunner JX. An adaptive lung ventilation controller. *IEEE Transactions in Biomedical Engineering* 1994; 41: 51–59.

Levitzky MG. *Pulmonary Physiology*. 7th edition. McGraw-Hill 2007. pg 25.

Lewis RM, Torczon V. Pattern search algorithms for linearly constrained minimization. *SIAM Journal on Optimization* 2000; 10(3): 917-941.

Liu CH, Niranjan SC, Clark JW, San KY, Zwischenberger JB, Bidani A. Airway mechanics, gas exchange, and blood flow in a nonlinear model of the normal human lung. *Journal of applied physiology* 1998; 84(4): 1447-1469.

Lundquist H, Hedenstierna G, Strandberg A, Tokics L, Brismar B. CT-assessment of dependent lung densities in man during general anaesthesia. *Acta Radiology* 1995; 36(6): 626-632.

Macklem PT, Mead J. Resistance of central and peripheral airways measured by a retrograde catheter. *Journal of Applied Physiology* 1967; 22(3): 395-401.

Marcy TW. Barotrauma: detection, recognition, and management. *Chest* 1993; 104(2): 578-584.

Marshall BE, Clarke WR, Costarino AT, Chen L, Miller F, Marshall C. The dose response relationship for hypoxic pulmonary vasoconstriction. *Respiratory Physiology* 1994; 96(2): 231-247.

Martin S, Wallace I, Bates DG. Development and validation of a civil aircraft engine simulation model for advanced controller design. *ASME Journal of engineering for gas turbines and power* 2008; 130(5): 051601 (15 pages).

Mathworks. *Matlab Genetic Algorithm and Direct Search Toolbox User's Guide*. 2005.

Mathworks. *Optmization Toolbox User's Guide,* 2nd edition 2000.

McCahon RA, Columb MO, Mahajan RP, Hardman JG. 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. *British Journal of Anaesthesia* 2008; 101(3):358-365.

McHardy GJ. The relationship between the differences in pressure and content of carbon dioxide in arterial and venous blood. *Clinical Science* 1967; 32(2): 299-309.

McKinley BA, Moore FA, Sailors RM, Cocanour CS, Marquez A, Wright RK, Tonnesen, AS, Wallace CJ, Morris AH, East TD. Computerized Decision Support for Mechanical Ventilation of Trauma Induced ARDS: Results of a Randomized Clinical Trial. *The Journal of Trauma: Injury, Infection, and Critical Care* 2001; 50(3): 415.

Mead J, Takishima T, and Leith D. Stress distribution in lungs: a model of pulmonary elasticity. *Journal of Applied Physiology* 1970; 28: 596–608.

Menon PP, Bates DG, Postlethwaite I. Computation of worst case pilot inputs for nonlinear flight control system analysis. *Journal of Guidance Control and Dynamics* 2006; 29(1):195-199.

Menon PP, Bates DG, Postlethwaite I. Nonlinear robustness analysis of flight control laws for highly augmented aircraft. *Control Engineering Practice* 2007; 15: 655-662.

Mercat A, Richard JC, Vielle B, Jaber S, Osman D, Diehl JL, Lefrant JY, Prat G, Richecoeur J, Nieszkowska A. Positive end-expiratory pressure setting in adults with acute lung injury and acute respiratory distress syndrome: a randomized controlled trial. *Journal of American Medical Association* 2008; 299**:** 646-655.

Messac A, Puemi-Sukam C, Melachrinoudis E. Aggregate Objective Functions and Pareto Frontiers: required Relationships and Practical Implications. *Optimization and Engineering* 2000; 1:171-188.

Milickovic NB, Lahanas M, Papagiannopulou M, karouzakis K, Baltas D, Zamboglou N. Application of multi-objective genetic algorithms in anatomy based dose optimization in brachytherapy and its comparison with deterministic algorithms. *Engineering in medicine and Biology 2001, Proceedings of the 23 Annual International Conference of the IEEE* 2001; 4: 3919-3922.

Mitchell M. *An Introduction to genetic algorithms.* MIT press. 1998.

Mogensen ML, Thomsen LP, Karbing DS. Steimle KL, Zhao Y, Rees SE, Andreassen S. A Mathematical Physiological Model of the Dynamics of the Pulmonary Ventilation. In proceedings of *UKACC International Conference on CONTROL 2010 September 7-10th,Coventry, UK* 2010; 734-739.

Moppett IK, Hardman JG. Development and validation of an integrated computational model of cerebral blood flow and oxygenation. *Anesthesia and Analgesia* 2008; 105(4): 1094-1103.

Morohashi M, Winn AE, Borisuk MT, Bolouri H, Doyle J, Kitano H. Robustness as a measure of plausibility in models of biochemical networks. *Journal of Theoretical Biology* 2002; 216(1):19-30.

Morris AH. Decision support and safety of clinical environments. *Quality and Safety in Health Care* 2002; 11: 69-75.

Moss M, Wellman DA, Cotsonis GA. An appraisal of Multivariable Logistic Models in the Pulmonary and Critical Care Literature. *Chest* 2003; 123: 923-928.

Muscedere JG, Mullen JB, GAN K, Slutsky AS. Tidal ventilation at low airway pressures can augment lung injury. *American Journal Respiratory and Critical Care Medicine* 1994; 149(5): 1327-1334.

Neumann P, Hedenstierna G. Ventilation-perfusion distributions in different porcine lung injury models. *Acta Anaesthesiologica Scandinavica* 2001; 45:78-86.

Norman RG, Goldring RM, Clain JM, Oppenheimer BW, Charney AN, Rapoport DM, Berger KL. Transition from acute to chronic hypercapnia in patients with periodic breathing: predictions from a computer model. *Journal of Applied Physiology* 2006; 100: 1863-1876.

Nouraei SAR, Battson RM, Koury EF, Sandhu GS, Patel A. Adult post-intubation laryngotracheal stenosis: an underestimated complication of intensive care? *Journal of Intensive Care Society* 2009; 10(3): 229.

Okamoto VN, Rubenfeld GD. Attending to the lightness of numbers: toward the understanding of critical care epidemiology. *Critical Care* 2004; 8: 422-424.

Pappert D, Rossaint R, Slama K, Gruning T, Konrad JF. Influence of positioning on ventilation perfusion relationships in severe adult respiratory distress syndrome. *CHEST* 1994; 106: 1511-1516.

Peterman BF, Longtin A. Multicompartment model of lung dynamics. *Computers and Biomedical Research* 1984; 17: 580-589.

Petrini MF, Robertson HT, Hlastala MP. Interaction of series and parallel dead space in the lung. *Respiration physiology* 1983; 54(1): 121-136.

Powell MJD. A fast algorithm for nonlinearly constrained optimization calculations. In *Numerical Analysis, Lecture Notes in Mathematics* 1978; 630:144-157.

Price KV, Storn RM, Lampinen JA. Differential evolution: a practical approach to global optimization. Springer Verlag 2005.

Rahn H. A concept of mean alveolar air and the ventilation-bloodflow relationships during pulmonary gas. *Americal Journal of Physiology* 1949; 158(1): 21-30.

Rees SE, Allerod C, Murley D, Zhao Y, Smith BW, Kjaergaard S, Throgaard P, Andreassen S. Using physiological models and decision theory for selecting appropriate ventilator settings. *Journal of clinical monitoring and computing* 2006; 20(6): 421-429.

Roy B. Problems and methods with multiple objective functions. *Mathematical Programming 1* 1971; 239-266.

Rubenfield GD, Caldwell E, Peabody E, Weaver J, Martin DP, Neff M, Stern EJ, Hudson LD. Incidence and outcomes of Acute Lung Injury. *New England Journal of Medicine* 2005; 353: 1685–1693.

Rudowski R, Bokliden A, Carstensen A, Gill H, Ludwigs, U, Matell G. Multivariable Optimization of mechanical ventilation. *International Journal of Clinical Monitoring and Computing* 1991; 8:103-115.

Rutledge GW, Thomsen GE, Farr BR, Tovar MA, Polaschek JX, Beinlich IA, Sheiner LB, Fagan LM. The design and implementation of a ventilator-management advisor. *Artificial Intelligence in Medicine* 1993; 5(1): 67-82.

Rutledge GW. VentSim: A simulation model of cardiopulmonary physiology. *Journal of the American Medical Informatics Association, Symposium supplement* 1994; 878-883.

Saunders KB, Bali HN, Carson ER. A breathing model of the respiratory system: the controlled system. *Journal of theoretical biology* 1980: 84 (1): 135-161.

Schaffer JD. Multiple objective optimization with vector evaluated genetic algorithms. In *Proceedings of the First International Conference on Genetic Algorithms* 1985; 93-100.

Severinghaus JW. Blood gas calculator. *Journal of Applied Physiology* 1966; 21(3): 1108-1116.

Severinghaus JW. Simple, accurate equations for human blood O2 dissociation computations. *Journal of Applied Physiology* 1979; 46(3): 599-602.

Siggard-Andersen O. The Van Slyke Equation. *Scandinavian journal of clinical and laboratory investigation* 1977; 37: 15-20.

Slutsky AS. Lung injury caused by mechanical ventilation. *Chest*. 1999;116 (Suppl 1): 9S–15S.

Slutsky AS. Mechanical ventilation: American College of Chest Physicians' Consensus Conference. *Chest* 1993; 104:1833-1859.

Srinivas N, Deb K. Multi-objective function optimization using nondominated sorting genetic algorithms. *Evolutionary Computation Journal* 1994; 2(3): 221-248.

Stead EA, Warren J, Merrill AJ, Brannon ES. The cardiac output in male subjects as measured by the technique of right atrial catheterization. Normal values with observations on the effect of anxiety and tilting. *Journal of Clinical Investigation* 1945; 24: 326.

Storn R, Price K. Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization* 1997; 11(4): 341-359.

Stuhmiller JH, Stuhmiller LM. A mathematical model of ventilation response to inhaled carbon monoxide. *Journal of applied physiology* 2005; 98 (6): 2033-2044.

Suki B, Andrade JS, Coughlin MF, Stamenovic D, Stanley HE, Sujeeer M, Zapperi S. Mathematical modeling of the first inflation of degassed lungs. *Annals of Biomedical Engineering* 1998; 26(4): 608-617.

Swanson GD, Sherrill DL. A model evaluation of estimates of breath-to-breath alveolar gas exchange. *Journal of applied physiology* 1983; 55(6): 1936-1941.

Tehrani FT. Automatic control of mechanical ventilation. Part 1: theory and history of the technology. *Journal of Clinical Monitoring and Computing* 2008a; 22: 409–415.

Tehrani FT. Automatic control of mechanical ventilation. Part 2: The existing techniques and future trends. *Journal of Clinical Monitoring and Computing* 2008b; 22:417–424.

Tehrani FT, Roum JH. Flex: A New Computerized System for Mechanical Ventilation. *Journal of clinical monitoring and computing* 2008c; 22(2): 21-130.

Tobin MJ (eds) In Principles and Practice of Mechcanical Ventilation 2nd edition. McGraw-Hill. 2006:110.

Tsoukias NM, George SC. A two compartment model of pulmonary nitric oxide exchange dynamics. *Journal of Applied Physiology* 1998; 85: 653-666.

Ursino M, Magosso E, Avanzolini G. An integrated model of the human ventilator control system: the response to hypoxia. *Clinical Physiology* 2001; 21 (4): 465-477.

Van Slyke DD. Studies of acidosis XVII. The normal and abnormal variations in the acid-base balance of the blood. *Journal of Biological Chemistry* 1921; 48; 153-176.

Vesterstrom J, Thomsen R. A comparative study of differential evolution, particle swarm optimization and evolutionary algorithms on numerical benchmark problems. In *Proceedings of the 2004 IEEE Congress on Evolutionary Computation* 2004; 1980-1987.

Vidalmelo MF, Loeppky JA, Caprihan A, Luft UC. Alveolar ventilation to perfusion heterogeneity and diffusion impairment in a mathematical model of gas exchange. *Computers and Biomedical Research* 1993; 26(2); 103-120.

Vidalmelo MF. Effect of cardiac output on pulmonary gas exchange: role of diffusion limitation with Va/Q mismatch. *Respiration physiology* 1998; 113(1): 23-32.

Vidyasagar M. Statistical learning theory and randomized algorithms for control. *IEEE Control Systems Magazine* 1998; 18(6): 69-85.

Vieira SR, Puybasset L, Lu Q, Richecoeur J, Cluzel P, Coriat P, Rouby JJ. A scanographic assessment of pulmonary morphology in acute lung injury. Significance of the lower inflection point detected on the lung pressure–volume curve. *American Journal Respiratory and Critical Care Medicine* 1999; 159: 1612–1623.

Villar J, Perez-Mendez K, Kacmarek RM. Current definitions of acute lung injury in acute respiratory distress syndrome do not reflect their true severity and outcome. *Intensive Care Medicine* 1999; 25(9): 930-935.

Villar J, Kacmarek RM, Hedenstierna G. From ventilator-induced lung injury to physician-induced lung injury: Why the reluctance to use small tidal volumes? *Acta Anaesthesiologica Scandinavica* 2004; 48:267–271.

Villar J, Perez-Mendez L, Lopez J, Belda J, Blanco J, Saralegui I, Suarez-Sipmann F, Lopez J, Lubillo S and Kacmarek RM. An Early PEEP/FIO2 Trial Identifies Different Degrees of Lung Injury in Patients with Acute Respiratory Distress Syndrome. *American Journal of Respiratory and Critical Care Medicine* 2007; 176(8): 795-804.

Wagner PD, Laravuso RB, Uhl RR, West JB. Continuous distributions of ventilation-perfusion ratios in normal subjects breathing air and 100% O2. *Journal of Clinical Investigation* 1974; 54: 54-68.

Wagner PD, Dantzker DR, Duek R, Clusen JL, West JB. Ventilation-perfusion inequality in chronic obstructive pulmonary disease. *The Journal of Clinical Investigation* 1977; 59: 203-216.

Wagner PD, Dantzker DR, Iacovoni VE, Tomlin WC, West JB. Ventilation-perfusion inequality in asymptomatic asthma. *American Review of Respiratory Disease* 1978; 118(3): 511-524.

Wakeling HG, MCFall MR, Jenkins CS, Woods WGA, Miles WFA, Barclay GR, Fleming SC. Intraoperative oesophageal Doppler guided fluid management shortens postoperative hospital stay after major bowel surgery. *British Journal of Anaesthesia* 2005; 95: 634-42

Wang W, Menon PP, Bennani S, Bates DG. Robustness analysis of attitude and orbit control systems for flexible satellites. *IET Control Theory and Applications* 2010; 4(12): 2958-2970.

Weibel, E.R. Morphometry of the human lung. 1963. Heidelberg, Springer-Verlag.

West JB. Respiratory Physiology: The Essentials 7th edition. Lippincott Williams and Wilkins 2004

West JB. Ventilation-perfusion inequality and overall gas exchange in computer models of the lung. *Respiration physiology* 1969; 7(1): 88-110.

Wiggs BR, Bosken C, Pare PD, James A, Hogg JC. A model of airway narrowing in asthma and in chronic obstructive pulmonary disease. *American Journal of Respiratory and Critical Care Medicine* 1992; 145(6): 1251-1258.

Williams PS. A monte carlo dispersion analysis of the X-33 simulation software. *In proceedings of the AIAA Conference on Guidance, Navigation and Control, Montreal Canada* 2001.

Wilson AJ, Murphy CM, Brook BS, Breen D, Miles AW, Tilley DG. A computer model of the artificially ventilated human respiratory system in adult intensive care. *Medical Engineering and Physics* 2009; 31: 1118-1133.

Woollam CHM. The development of apparatus for intermittent negative pressure respiration (2) 1919-1976 with special reference to the development and uses of cuirass respirators. *Anaesthesia* 1976; 31:666-685.

Yem JS, Tang Y, Martin JT, Baker AB. Sources of error in noninvasive pulmonary blood flow measurements by partial rebreathing. *Anethesiology* 2003; 98: 881-887.

Yem JS, Turner MJ, Baker AB, Young IH, Crawford ABH. A tidally breathing model of ventilation, perfusion and volume in normal and diseased lungs. *British Journal of Anaesthesia* 2006; 97(5): 718-731.

Younes M. Proportional assist ventilation, a new approach to ventilatory support. *American Review of Respiratory Disease* 1992; 145: 114–120.

\