The Lognormal Lung:

Quantifying Inhomogeneity of Ventilation

James Mountain

Lady Margaret Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary 2016

Acknowledgements

First and foremost I would like to thank my supervisors, Peter Robbins and Jonathan Whiteley, for all their enthusiasm and guidance throughout my doctoral studies. Also to Andrew Farmery for his early input. Thanks to everyone involved in the development of the LGA, without their hard work and dedication over a number of years, this work would not have been possible.

I am extremely grateful to all the volunteers who gave up their time to participate in the pilot study. Thanks to David O'Neill and Peter Santer for their help conducting the pilot study.

Thank you to the Doctoral Training Centre for providing a change of direction and many friends. Also to the members of the Computational Biology group for interesting presentations, useful discussions, and good cake! To the Oxford University Boat Club where I learnt so much.

To my friends and family for all their patience and encouragement. In particular, to Maxie for all her support.

Finally to my parents, for so many things, but in particular for giving me a curiosity about the world around me, a gift I value above all others.

The Lognormal Lung:
Quantifying Inhomogeneity of Ventilation

James Mountain

Lady Margaret Hall

A thesis submitted for the degree of Doctor of Philosophy

Hilary 2016

Abstract

Chronic obstructive pulmonary disease is a significant cause of morbidity and mortality worldwide. Despite this, the most common lung function testing techniques are not sensitive enough to detect the disease in its early stages, limiting opportunities for secondary prevention. It has been shown that measures of inhomogeneity detect signs of the disease earlier than spirometry, the most common form of lung function test. There is however, currently no "gold standard" method for assessing the inhomogeneity of the lungs.

A common technique for investigating inhomogeneity of ventilation is the multi-breath nitrogen washout. During this procedure a patient is switched from breathing room air to pure oxygen, and the concentration of nitrogen in the expired gas is measured as it washes out of the lungs. In this thesis we propose a novel method of assessing the inhomogeneity of the lungs from high quality nitrogen washout data.

This method involves the estimation of parameters from a mathematical model of the lungs which may then be reported as measures of inhomogeneity. The mathematical model we develop aims to represent the underlying physiology well; while having a small parameter set, and governing equations that may be solved quickly and accurately.

Careful attention is also paid to the parameter estimation procedure. We develop a practical approach which uses the available data, and considers computational constraints. We test our method on synthetic data to investigate the uncertainty of the parameters we estimate.

Finally, we present results from a small experimental trial designed to test our method in fifteen volunteers. The results appear promising, with our method distinguishing between the young, the old, and those with respiratory disease. However, there are some inconsistencies in our results. We highlight these, and suggest some areas for future investigation that might improve the work presented in this thesis.

Contents

Chapter 1

Introduction

In the European Union chronic obstructive pulmonary disease (COPD) is the third most common cause of death, and in North America it is the fourth [104]. Strategies for dealing with this burden focus on either primary prevention, aiming to reduce risk factors such as exposure to tobacco smoke, or tertiary prevention, reducing symptoms in patients where the disease is fully developed [65].

Currently the most common method for assessing lung function is spirometry [72]. To detect lung disease spirometry measurements are compared to 'normal' values for patients of similar physical characteristics. However, a report by the U.S. Preventive Services Task Force found that spirometry is not sensitive enough to make its use as a screening tool for COPD worthwhile [63]. Thus there is a need to develop more sensitive measures of lung function that could be used to test for mild COPD.

One concept that has long been considered an important indicator of respiratory health is the degree of inhomogeneity in the lungs. Inhomogeneity in the lungs generally refers to a mismatch between either ventilation and perfusion, or ventilation and alveolar volume. In a homogeneously ventilated lung each unit of volume would the same amount of fresh gas and perfusion. In practice however, this is not the case. Even healthy lungs are slightly inhomogeneous, and it has long been recognised that the amount of inhomogeneity appears to increase with disease [11, 22, 31, 32]. It has been suggested recently that indicators of inhomogeneity show abnormality earlier than spirometry in mild COPD patients [97].

A common technique for investigating inhomogeneity of ventilation is the multi-breath nitrogen washout (MBNW). During this procedure a patient is switched from breathing room air to pure oxygen, and the concentration of nitrogen in the expired gas is measured as it washes out of the lungs [95]. A number of measures of inhomogeneity can be derived from this technique, with the most popular being the lung clearance index (LCI) [11, 50]. This index (which will be discussed in more detail in Section 2.6.4) has been shown to be able to separate healthy controls from people with both cystic fibrosis and COPD [13, 38, 51, 54] demonstrating the potential of the nitrogen washout as a method for detecting lung disease.

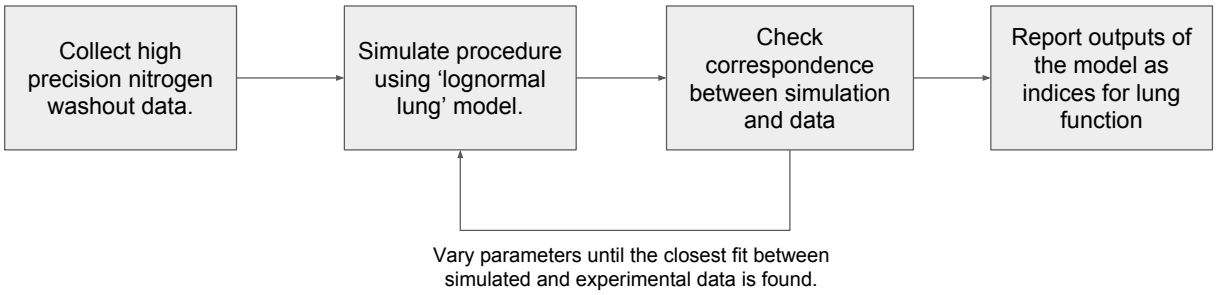

Figure 1.1: An outline of the process we use to investigate a subject's lung function.

1.1 Central research question

This thesis is based on the premise that there is information in the transient expired CO_2, O_2, and N_2 concentrations (both during air breathing and the nitrogen washout) that is currently not used in clinical assessment of lung function, and that use of such information might provide a clinically useful measure of the inhomogeneity of the lungs. Our proposed lung function test that makes use of this information is outlined in Figure 1.1.

This test involves estimating the parameters of a novel lung model ("the lognormal lung") from experimentally-measured nitrogen washout procedures, and using those parameters as as measures of inhomogeneity in the lungs. For this test to work we require the following.

- A device that can accurately monitor the concentrations and flows of CO_2, O_2, and N_2 at the mouth with high time resolution, so we can apply the principle of conservation of mass.
- A mechanistic model of the lungs, whose parameters provide a good measure of lung function (and in particular inhomogeneity in the lungs).
- A procedure for estimating the parameters of that model based on experimental data.

While the focus of this thesis is theoretical, the work detailed here was performed in close collaboration with Professor Robbins' group in the Department of Physiology, Anatomy and Genetics. Professor Robbins was a co-supervisor of this thesis, providing significant guidance on the physiological content and direction of the work.

As part of their work to measure oxygen consumption in intensive care Professor Robbins' group have developed an experimental device for high precision in-airway molecular flow sensing [17]. This device integrates the measurement of CO_2, O_2, and H_2O using laser absorption spectroscopy, with a novel high precision flow sensor, which is designed to improve on specifications for such devices recently laid out by the European Respiratory Society (ERS) [96] by more than an order of magnitude.

When using the device a patient breathes through a mouth-piece and the concentrations of CO_2, O_2, and H_2O and the flow rate of gas are measured every 10ms. Both the flow and concentration measurements are performed in the airway, and with a very fast response time, removing many of the problems associated with alignment between flow and concentration signals that are seen in data collected with many devices [96]. In this thesis this device is referred to as the laser gas analyser (LGA).

This thesis focuses on the last two challenges listed in the bullet points above: (i) the development of a novel mathematical model of the lungs; and (ii) a robust procedure for estimating the parameters of this model. Thus the central research question of this thesis is: "Can we develop a novel mathematical model of the lungs, whose parameters we can estimate from high quality nitrogen washout data provided by the LGA, to provide a sensitive measure of the inhomogeneity of the lungs?"

1.2 Thesis outline

After reviewing relevant literature in Chapter 2, we describe the novel mathematical model of the lung we have developed in Chapter 4. This model makes two big assumptions. First, we assume that the lung can be modelled with two types of volume; deadspace compartments in which all gas transport is via convection, and perfectly mixed alveolar compartments. Thus we ignore serial inhomogeneity or stratification within the alveoli. Secondly, we assume that the inhomogeneity of ventilation, perfusion, and volume in the alveoli can be parameterised with a bivariate lognormal distribution. This chapter also explains how we model gas flow in the deadspace, and how we solve the governing equations in a computationally efficient manner.

In Chapter 5 we develop a method to estimate the parameters of our model from experimental data (the second of the challenges identified above). We describe the difficulties of estimating the venous gas concentrations (which we do not measure) and develop a method for doing this. We then test the performance of our parameter estimation method on synthetic data (data created using our model), both with and without a realistic amount of experimental error (for which we develop a detailed model). We consider the uncertainty of our parameter estimates and find that our method requires a very accurate measuring device, such as the LGA, to give accurate parameter estimates.

Next, in Chapter 6, we discuss the results of estimating the parameters of our model from six washout procedures performed on 15 (six young healthy, six older healthy, and three with moderate COPD) volunteers. Here we assess: (i) how well our model fits the experimental data; (ii) the repeatability of our parameter estimates; (iii) how these estimates compare to measurements made with other modalities; and (iv) the differences between the three groups.

The experimental data discussed in Chapter 6 was collected in conjunction with Dr O'Neill and Dr Santer. Dr O'Neill ensured the device was correctly calibrated for each visit. Dr Santer recruited the COPD patients and older volunteers; performed the lung function testing; and provided medical support. I recruited the young healthy volunteers and operated the LGA during the washout procedures.

Finally in Chapter 7 we summarise the work we have performed and offer some suggestions for future work in this area.

Chapter 2

Literature review

In the previous chapter we established that the goal of this thesis is to develop a novel method of quantifying inhomogeneity in the lungs, by combining a novel mathematical model of the lungs with high quality measurements of gases entering and leaving the lungs. In this chapter we will review the literature that forms the background to the work presented in the rest of the thesis.

First we give a brief introduction to the structure and function of the lungs. We then give some details about COPD, the lung disease that motivated this work.

Next we discuss several different types of mathematical model of the lungs and their uses. In this section we focus particularly on models with governing equations that may be solved in a reasonable amount of time, and that are characterised by limited parameter sets, as these are likely to be the category of models that can be parameterised for an individual patient using experimental data. This section explains in some detail how we apply the principle of conservation of mass in the lungs, as this principle is used repeatedly in this thesis. We then discuss the 'inverse' problem that we must solve to find parameter estimates for individuals from experimental data.

We introduce a conceptual model of the lung, the functional lung unit [128], and use it to explain a concept of central importance in this thesis, inhomogeneity. In this section we also use a very simple model to illustrate the effects of inhomogeneity on lung function, and discuss the experimental evidence that led us to ignore one type of inhomogeneity in the lungs, serial inhomogeneity.

Finally, we discuss existing pulmonary function tests (PFTs). In particular we discuss spirometry, the most common PFT, and its limitations that make the development of alternative tests attractive. We also discuss a number of methods that have been used to investigate inhomogeneity in the lungs. Many of these methods of assessing inhomogeneity have found significant correlation between estimates of inhomogeneity and lung disease, demonstrating the potential of reliable measures of inhomogeneity as biomarkers of lung function. In keeping with the goal of this thesis we pay particular attention to the N_2 washout, and methods used to analyse this procedure.

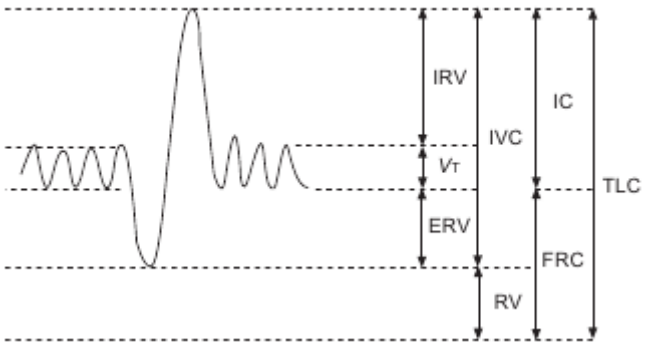

Figure 2.1: The definition of static lung volumes based on a volume-time spirogram of an Inspiratory Vital Capacity manoeuvre. **Inspiratory Reserve Volume (IRV)**: Volume of gas that can be further inspired from the end of inspiration in conventional tidal breathing. **Tidal Volume ($\mathbf{V_T}$)**: Volume of gas expired/inspired during normal breathing. **Expiratory Reserve Volume (ERV)**: Volume of gas that can be further expired from the end of expiration during normal tidal breathing. **Inspiratory Vital Capacity (IVC)**: Maximum volume of gas that can be inspired after a maximal expiration. **Residual Volume (RV)**: Volume of gas remaining in the lungs after a maximal expiration. **Inspiratory Capacity (IC)**: Volume of gas that can be maximally inspired from the end of tidal expiration. **Functional Reserve Capacity (FRC)**: Volume of gas remaining in the lungs at the end of tidal expiration. **Total Lung capacity (TLC)**: Total volume of gas in the lungs at the end of a maximal inspiration. Figure from [126].

In this thesis I will refer to several abbreviations for different lung volume measures, these are defined with the aid of a spirogram in Figure 2.1.

2.1 Lung physiology

The primary function of the lung is to exchange gas with the blood, and in particular, eliminate the CO_2 produced by the metabolism, and take up the O_2 required for the metabolism to continue [128]. The underpinning physiology that allows the lungs to perform this function is extraordinarily intricate.

Here we provide a brief overview of some of the key features of the structure and function of the lungs, before discussing the pathology of COPD, the disease that motivated our work.

2.1.1 Lung structure

Morphometric studies reveal that the airways of the lung form a dichotomously branching structure, often known as the airway tree, that terminates in alveoli [127]. There are approximately 23 generations of branches in the airway tree. These generations are often split into three zones: (i) the conducting zone, from generation 0 to approximately generation 16; (ii) the transitory zone, from approximately generation 17 to 19; and (iii) the

respiratory zone, from approximately generation 20 to 23. The conducting zone comprises the conducing airways, which are not thought to expand significantly during inspiration and do not have alveoli. Alveoli begin to appear in the transitory zone of respiratory bronchioles, and therefore gas exchange starts to occur in this zone. The respiratory zone consists of alveolar ducts connecting the bronchioles to alveolar sacs. This collection of structures is sometimes known as the acinus. This is where the majority of alveoli are situated, and the majority of gas exchange occurs. For an average size adult, whose lung is inflated to three quarters of TLC, the volume of the conducting zone is typically around 150 ml, the transitory zone 1500 ml, and the respiratory zone 3150 ml [127].

The alveoli, where the airway tree terminates, are hollow elastic structures that expand during inspiration and contract during expiration. Gas diffuses from the airspace into the blood (and vice versa) across the alveolar membrane. There are between 400 and 700 million alveoli in a typical human lung (the number varies with body size). In total the alveoli comprise a surface area of about 70 to 100 m^2, approximately 70% of which is covered by pulmonary capillaries [80, 127].

As even this short overview makes clear, the lung is an enormously complicated structure with many physical processes occurring. We do not attempt to model the lung in this level of anatomical detail, but a discussion of some models that do is presented in Section 2.2.4. Next we give a brief overview of the mechanics of breathing.

2.1.2 Mechanics of breathing

Air enters or leaves the lungs because of small pressure gradients (approximately 1 cmH_2O or 98 Pa) between the air in the alveoli and the atmosphere [60]. Here we give a brief explanation of what causes these pressure gradients during spontaneous breathing, based on the overviews presented in [60, 129], and introduce the concept of compliance.

At FRC, the elastic recoil of the alveoli (which tends to cause them to contract) is balanced by the elastic recoil of the chest wall (which tends to pull outwards). This causes the intrapleural pressure (the pressure in the space between the chest wall and the lungs) to be very slightly negative, equivalent to approximately 5 cmH_2O [60]. At FRC the alveolar pressure is equal to the atmospheric pressure, and there is no airflow.

At the start of inspiration, the respiratory muscles contract, expanding the chest wall. This reduces the intrapleural pressure, causing the alveoli to expand, increasing their elastic recoil. This expansion causes the alveolar pressure to drop (by approximately 1 cmH_2O [60]), establishing a pressure gradient between the alveoli and the atmosphere. This drives airflow into the lungs until the alveolar pressure equilibrates with the atmospheric pressure. If the airways connecting the alveoli had no resistance to airflow, then the flow would compensate for this pressure drop instantaneously, and the drop in intrapleural pressure would be smaller [129].

Expiration starts when the respiratory muscles relax. This allows the chest wall to contract, increasing the intrapleural pressure, reducing the 'pull' on the alveoli. Thus, due to their elastic recoil, the alveoli contract, increasing the alveolar pressure (by approximately $1\,\mathrm{cmH_2O}$ [60]), and driving air out of the lungs until the pressure equilibrates.

Expansion and contraction of the lung is typically characterised by considering the compliance of the lung, which is the change in volume of the lung divided by the change in transpulmonary pressure (the difference between the intrapleural pressure and atmospheric pressure) associated with the volume change [60]. Compliance is a function of lung volume, typically decreasing with increasing lung volume (after a small increase at the at the start of inspiration), and is different on expiration and inspiration (i.e. demonstrates hysteresis) [60].

The compliance of the lung (which is inversely related to its elastic properties) is a combination of the compliance of the chest wall and of the alveoli [129]. If we consider an individual alveolus, the volume change associated with it is the product of its compliance (which is determined by its elastic properties), and the change in pressure that caused the volume change. COPD leads to increased compliance, and fibrosis (a condition caused by a proliferation of connective tissue in the lungs) decreased compliance [60].

Before discussing COPD, the disease that motivated this work, in more detail, we briefly discuss gas exchange in the lung.

2.1.3 Gas exchange in the lung

As we explained previously, gas exchange in the lungs occurs by diffusion across the alveolar membrane. An assumption of many mathematical models of the lungs is that this process occurs instantaneously. Theoretical analysis suggests that under normal physiological conditions this is a reasonable assumption [136]. Thus we also adopt this assumption, and assume that the partial pressures of $\mathrm{CO_2}$, $\mathrm{O_2}$, and $\mathrm{N_2}$ equilibrate across the alveolar membrane.

The gas concentration in the blood is then a function of alveolar partial pressure. For inert gases (such as $\mathrm{N_2}$) this is simply a linear function and,

$$C = S_b P, \tag{2.1}$$

where C is the concentration of the inert gas in the blood, P the partial pressure, and S_b the solubility.

For $\mathrm{CO_2}$ and $\mathrm{O_2}$ the situation is significantly more complicated. As well as dissolving, these gases react with many of the proteins (principally haemoglobin) and ions (for example bicarbonate) that are present in the blood [81]. This is an intricate system, and means that the concentrations of these gases are a coupled non-linear function of their partial pressures. In this thesis we use the detailed model of the biochemistry of the

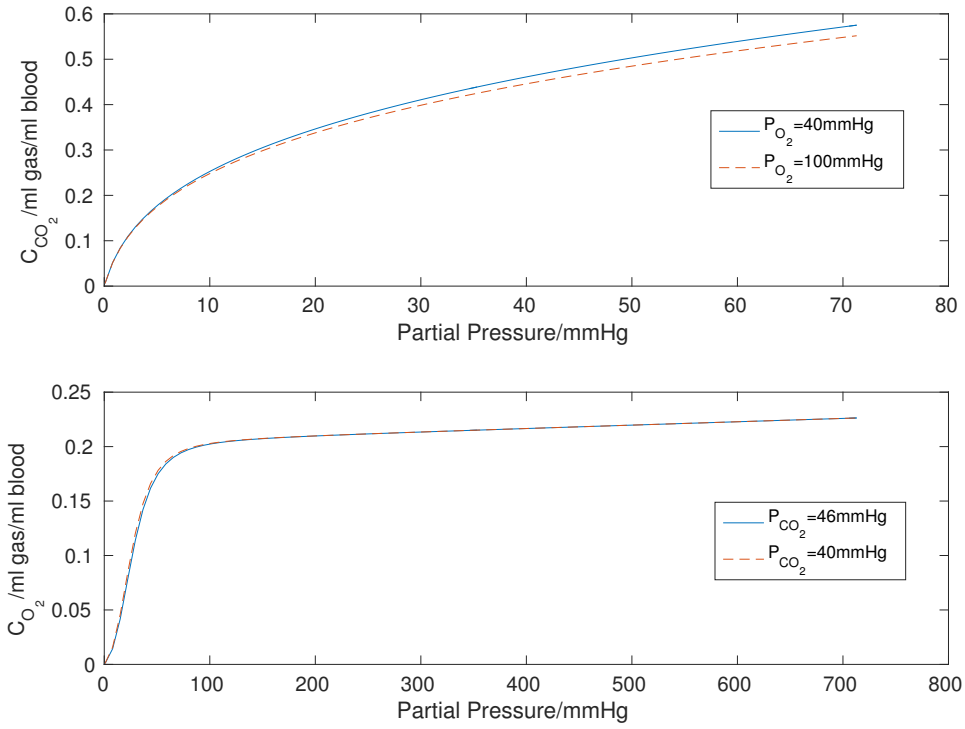

Figure 2.2: A plot of the CO_2 (top plot) and O_2 (bottom plot) concentrations predicted by the blood-gas model used in this thesis at a range of different partial pressures [81].

blood developed by O'Neill and Robbins [81] to convert partial pressures of CO_2 and O_2 into concentrations. Plots of the concentrations of CO_2 and O_2 as a function of partial pressure predicted by this model are shown in Figure 2.2.

In most of this thesis we refer to dry gas fractions for simplicity (as it allows us to write $F^{CO_2} + F^{O_2} + F^{N_2} = 1$[1]), rather than partial pressures as is often done in the literature. To convert between the two we assume that air everywhere within the lungs is at body temperature, pressure, and is saturated (BTPS), and thus

$$F^g = \frac{P^g}{P_b - P_{H_2O}}, \tag{2.2}$$

where F^g is the dry fraction of gas g, P^g its partial pressure, P_b atmospheric pressure, and P_{H_2O} the saturated partial pressure of H_2O.

2.1.4 The expirogram

In this thesis we will refer frequently to expirograms, which are plots of expired gas fraction measured at the mouth against cumulative expired volume. An example of one of these plots is shown in Figure 2.3.

[1]We include the volume of the other inert gases present in air, principally Ar, in the volume of N_2.

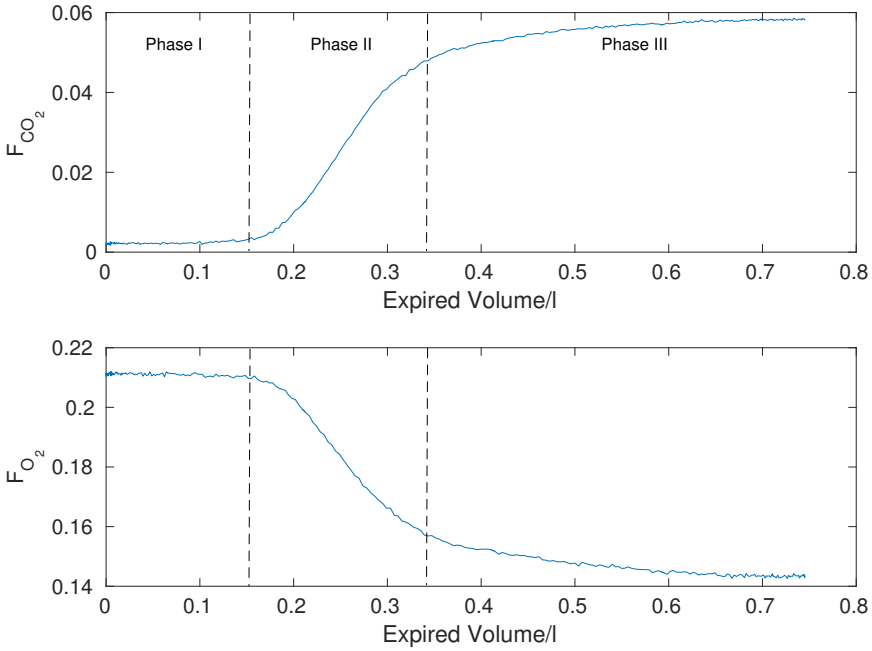

Figure 2.3: Examples of experimentally measured expirograms for CO_2 (top plot) and O_2 (bottom plot), showing the fraction of the gas of interest measured at the mouth against expired volume. The three phases, whose locations we approximately identify in this plot, are discussed in more detail in the main body of the text.

Three phases of this diagram are commonly identified. Phase I is typically believed to consist of gas expired from the deadspace (the volume of the lungs that does not participate in gas exchange), and thus gas fractions are very similar to their inspired values during this phase. In phase II gas from the alveoli begins to appear at the mouth, and thus there is a rapid change in gas fraction. Phase III, which is known as the alveolar phase, is assumed to consist of alveolar gas. During maximal expirations a rapid rise at the end of this graph is sometimes observed, this is referred to as phase IV [96].

2.1.5 COPD

In the European Union chronic obstructive pulmonary disease (COPD) is the third most common cause of death, and in North America the fourth [104]. Strategies for dealing with this burden focus on either primary prevention, i.e. aiming to reduce risk factors such as exposure to tobacco smoke, or tertiary prevention, i.e. reducing symptoms (such as coughing, and severe breathlessness) in patients where the disease is fully developed [65].

While COPD is an extremely heterogeneous disorder [53], there are two main components of the disease. These are narrowing of the small conducting airways, leading to their increased resistance to gas flow [67], and emphysema, destruction of the lung parenchyema (tissue within the respiratory zone of the lung), leading to a significant loss

of tissue elasticity [48]. Both this loss of elasticity (and thus increased compliance), and the increased resistance of the small airways contribute to the slow emptying of the lung during forced expiration which traditionally characterises COPD [48].

Based on this slow emptying of the lung, the disease is generally classified into four stages of severity, known as GOLD (Global Initiative for Chronic Obstructive Lung Disease) stages 1-4, by using a patients measured FEV_1 (the maximum volume of air that a person can expire in one second, for more details on this measure see Section 2.4). Gold stage 1 (mild COPD) is classified as a COPD patient with an FEV_1 that is above 80% of their predicted value, while in GOLD stage 4 (very severe) the FEV_1 is below 30% of the predicted value [122]. Despite this being the dominant form of assessing and classifying the disease, it is well recognised that FEV_1 is not a sensitive measure of the severity of the disease, or its progression [63, 65, 122].

As well as the difficulties of detecting COPD during its early stages, and the difficulty of tracking the progression of the disease that this lack of sensitivity causes, it also means that very large numbers of patients must be recruited to perform clinical trials into treatments for COPD. This makes the clinical trials very expensive, and has been suggested as one of the reasons why only one new class of medication, roflumilast [79], has been approved to manage COPD in the last 25 years [71].

It has has been observed using the multiple inert gas elimination technique (see Section 2.6.3 for more details on this technique), that the mismatch of ventilation and perfusion increases in COPD [123], and that detecting this mismatch may be a better method of detecting the early stages of the disease than spirometry [97]. When tested using this technique COPD patients exhibit regions of very high ventilation-perfusion ratio, which, it has been suggested, is due to regions of the lung in which emphysematous destruction of a significant number of pulmonary capillaries has occurred [128]. This, and the changes in elasticity of lung tissue, suggest that COPD may have large effects on the inhomogeneity of both ventilation and perfusion.

2.2 Mathematical models of the lungs

As we explained in Chapter 1 the goal of this thesis is to develop a mechanistic model of inhomogeneity in the lungs that can be fitted to high quality experimental data, to provide useful measures of lung function. This requires a model with a limited parameter set, governed by equations that can be solved rapidly (parameter recovery procedures typically involve solving the equations describing a model a large number of times), but which still represents the underlying physiology.

Our model is presented in Chapter 4, but here we discuss a number of different models of the lungs that have been developed previously. For the purposes of the following discussion we classify these models into three categories.

1. *Compartmental models.* These models approximate the alveolar region of the lung with perfectly mixed alveolar compartments. The conducting airways are approximated with deadspace compartments in which no diffusion, or gas exchange with the blood, occurs.

2. *Trumpet models.* These models still use a simplified geometry, but consider transport in the conducting airways via both convection and diffusion.

3. *Biophysical models.* These models generally involve the use of imaging methods to derive a detailed anatomical geometry on which to solve the governing equations. These governing equations aim to represent the underlying physics of the processes occurring within the lung in more detail than the other classes of model.

We will now discuss each of these in turn, aiming to give an idea of their assumptions, and the problems they have been used to investigate.

2.2.1 Compartmental models

In Section 2.1.1 we described the branching structure of the airway tree which terminated in alveoli. These alveoli appear from approximately generation 17 to generation 23. However, the simplest form of lung models ignore this complexity, and assume that the lung can be split into regions in which gas transport is only by convection, and there is no gas exchange with the blood (the deadspace), and perfectly mixed alveolar compartments in which all the gas exchange between the air and the blood occurs.

Before describing the two main classes of such models, continuous ventilation models and tidal ventilation models, we state the principle of the conservation of mass that will be used repeatedly in this thesis, and explain how it applies to the functional lung unit presented in Figure 2.7.

2.2.1.1 Conservation of mass

The principle of conservation of mass that we use in this thesis states that the rate of change of volume of a particular gas within a region is equal to the net flux of that gas into that region. We refer here to volume rather than mass as we assume that gas in the lungs is at constant pressure and temperature, and thus the two measures are equivalent.

In the case of the functional lung unit shown in Figure 2.7 (this conceptual model is introduced properly in Section 2.5) this means that the rate of change of the volume of a given gas within the unit is equal to the sum of the net flux of that gas from the airways due to ventilation, and across the alveolar membrane due to exchange with the blood. We will discuss the net flux due to ventilation in the following sections, for both continuous ventilation and tidal ventilation models. However, first we will briefly explain the flux due to gas exchange with the blood, which is common to both classes of model.

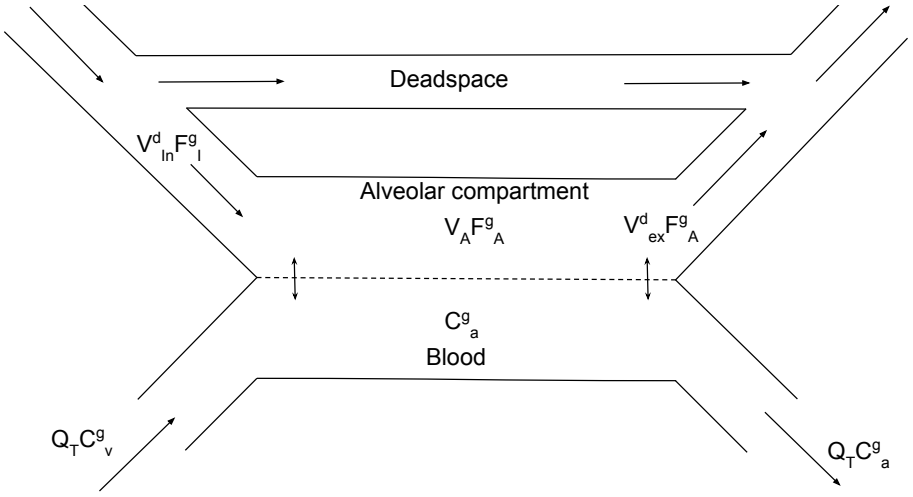

Figure 2.4: A schematic diagram of a continuous ventilation model, showing gas entering and leaving the alveoli via both ventilation and exchange between the gas and blood volumes. All quantities are defined in the main body of the text. Volumes are not to scale. Note for ease of illustration, the ventilation entering the alveolar compartment, $\dot{V}_{in}$, and the ventilation leaving the alveolar compartment, $\dot{V}_{ex}$, are shown as $\mathrm{V}^{\mathrm{d}}_{in}$ and $\mathrm{V}^{\mathrm{d}}_{ex}$ respectively in this figure.

Consider the blood compartment shown in Figure 2.4. In this figure we assume that all the pulmonary capillaries in the lung form a single compartment. The flux of a particular gas flowing into the compartment is $Q_T C_v^g$, where Q_T is the cardiac output (which we assume is constant) and C_v^g is the mixed venous gas concentration of gas g entering the compartment. The flux of the gas leaving the compartment is $Q_T C_a^g$, where C_a^g is the arterial concentration of that gas. Since the total pulmonary blood volume is small (approximately 70ml across the whole lung [128]), any change in the volume of gas in the pulmonary capillaries (and thus the blood compartment) is ignored, and by conservation of mass, the flux of the gas crossing the alveolar membrane is,

$$Q_T \left(C_v - C_a \right). \tag{2.3}$$

2.2.1.2 Continuous ventilation

An assumption included in many compartmental models is that of continuous ventilation [42]. Such models do not have separate inspiratory and expiratory phases. Instead, ventilation continuously enters at one end of the compartment and continuously leaves from the other. A further assumption is that the volume of the compartment is assumed to be constant. We will discuss the implications of this assumption later in this section. Figure 2.4 shows the key features of such a model. To model inhomogeneity, more alveolar compartments can be added with different ventilations, volumes, and perfusions [39, 61, 124].

It is assumed that the alveolar compartment is well mixed, and the volume of the gas of interest in the alveoli is $F_A^g(t)V_A$, where V_A is the alveolar volume, and $F_A(t)$ is the fraction of that volume occupied by the gas of interest. As can be seen in Figure 2.4, the fluxes of gas entering and leaving the alveolar compartment via ventilation are: $\dot{V}_{in}F_I^g$, where $\dot{V}_{in}$ is the ventilation entering the alveolar compartment and F_I^g the inspired fraction of gas g; and $\dot{V}_{ex}F_A^g$, where $\dot{V}_{ex}$ is the ventilation leaving the alveolar compartment. The net flux of gas g into the compartment via ventilation is then, $\dot{V}_{in}F_I^g - \dot{V}_{ex}F_A^g$. Combining this with the net flux of gas g into the alveolar compartment via exchange with the blood, (2.3), and applying conservation of mass to the alveolar volume, we have

$$\frac{d\left(F_A^g(t)V_A\right)}{dt} = \dot{V}_{in}F_I^g - \dot{V}_{ex}F_A^g + Q_T\left(C_v^g - C_a^g\right). \tag{2.4}$$

Summing this equation for CO_2, O_2, and N_2, and remembering that $F^{CO_2} + F^{O_2} + F^{N_2} = 1$, we have

$$\frac{dV_A}{dt} = \dot{V}_{in} - \dot{V}_{ex} + \sum_g Q_T\left(C_v^g - C_a^g\right). \tag{2.5}$$

Since the alveolar compartment has a constant volume in this model, the left hand side of this expression will be zero. Thus we have

$$\dot{V}_{ex} = \dot{V}_{in} + \sum_g Q_T\left(C_v^g - C_a^g\right). \tag{2.6}$$

As the net rate of gas exchange with the blood, given by $\sum_g Q_T\left(C_v^g - C_a^g\right)$, is typically small in comparison to the ventilation, it is often assumed that

$$\dot{V}_{ex} = \dot{V}_{in} = \dot{V}_A, \tag{2.7}$$

where $\dot{V}_A$ is the alveolar ventilation. If we make this assumption then (2.4) can be written in the more familiar form [42],

$$\frac{d\left(F_A^g(t)V_A\right)}{dt} = \dot{V}_A\left(F_I^g - F_A^g\right) + Q_T\left(C_v^g - C_a^g\right). \tag{2.8}$$

Another assumption common when using a continuous ventilation model is that we are in a steady state. In this case derivatives with respect to time are zero, and from (2.4) we have

$$\dot{V}_{in}F_I^g - \dot{V}_{ex}F_A^g + Q_T\left(C_v^g - C_a^g\right) = 0, \tag{2.9}$$

which we note does not depend on the alveolar volume, V_A. Hence, methods, such as the multiple inert gas elimination technique (MIGET) [124], that analyse data with a continuous ventilation model in steady state will not be sensitive to alveolar volume. MIGET will be discussed in more detail in Section 2.6.3.

Despite the simplicity of this model it has been used in a number of different settings: to determine deadspace volumes [30]; as the basis for the MIGET technique [124, 125];

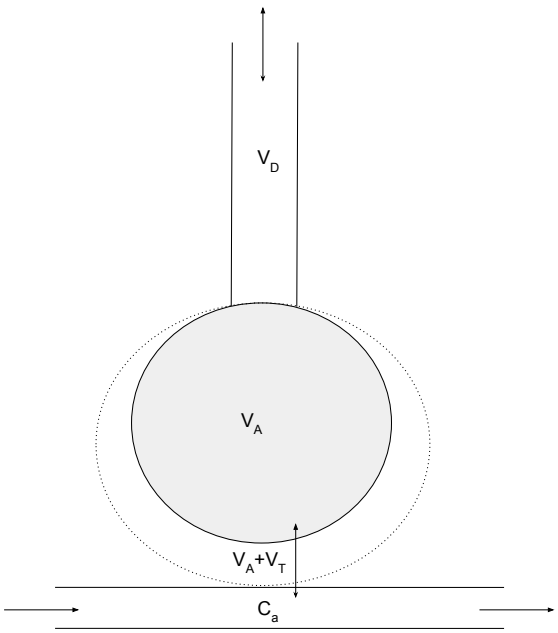

Figure 2.5: A schematic diagram of a single compartment tidal ventilation model, showing gas entering and leaving the mouth, and exchange between the gas and blood volumes. The dotted line indicates expansion of the alveoli during inspiration. All quantities are defined in the main body of the text. Volumes are not to scale.

in the sinewave technique [34, 143]; and in analysis of N_2 washouts [32, 39, 61]. There are two reasons for the wide usage of this model. First, the use of more complicated models is only justified if the data quality is better than much of what has been available until now [39, 95]. Secondly, these equations can either be solved analytically, or are computationally tractable with the computational resources that were available when the experimental techniques were first developed [42]. These experimental techniques, and their successes and failures, will be discussed in Section 2.6.

2.2.1.3 Tidal ventilation

A step towards a more realistic model of the lungs is to assume tidal, rather than continuous, ventilation. A simple schematic diagram of a single compartment of a tidal ventilation model is shown in Figure 2.5. In this model the alveolar compartment is ventilated via a series deadspace volume, and its volume expands and contracts during inspiration and expiration respectively.

Applying conservation of mass to the (well mixed) alveolar compartment during inspiration, we have that the rate of change of volume of the gas of interest is the sum of the flux due to gas exchange with the blood (explained in Section 2.2.1.1) and the flux due to ventilation, which is the flow rate entering the compartment multiplied by the inspired fraction of the gas of interest. The governing ordinary differential equation (ODE) for the

alveoli on inspiration is then [35, 42],

$$\frac{d\left(F_A(t)V_A(t)\right)}{dt} = u(t)F_{IA}(t) + Q_T\left(C_v(t) - C_a(t)\right), \tag{2.10}$$

where $u(t)$ is the flow rate entering the compartment, and $F_{IA}(t)$ is the fraction of the gas of interest inspired into the alveolar compartment at time t. In this case, the alveolar volume, $V_A(t)$, is also a function of time. Most models [14] have assumed that since the volume of gas exchanged with the blood is small in comparison to the flux due to ventilation,

$$\frac{dV_A(t)}{dt} = u(t). \tag{2.11}$$

However, as we will explain further when developing our model in Chapter 4 we have not done this. The reason for this is that we use experimentally measured flow patterns, and while on any one breath the difference in expiratory and inspiratory volumes due to gas exchange is small, over the 15 minutes of breathing we consider, this difference becomes significant.

The situation during expiration is very similar, except the flux due to ventilation is now the flow rate (which is now negative as gas is leaving the compartment) multiplied by the fraction of the gas of interest in the alveoli. Thus we have

$$\frac{d\left(F_A(t)V_A(t)\right)}{dt} = u(t)F_A(t) + Q_T\left(C_v(t) - C_a(t)\right). \tag{2.12}$$

As with the continuous model, inhomogeneity in a tidal ventilation model may be introduced by considering a number of compartments with different ventilations, volumes, and perfusions.

Such models have been developed for a number of the techniques discussed in Section 2.6, including the sinewave technique [35, 134], MIGET [133, 141], and the N_2 washout [141]. Rather than using the flow measured at the mouth (as we do in Chapter 4) standard breathing patterns have often been assumed, for example, constant flow during inspiration and an exponentially declining flow during expiration, although it has been shown that the behaviour of the models, specifically the elimination of CO_2 from the lungs, is sensitive to the form of breathing pattern assumed [135].

In Chapter 4 we develop a novel mathematical model of the lungs. We extend the material presented in this section significantly in that chapter.

2.2.2 Trumpet model

The compartmental models discussed above assume (on the basis of experimental evidence, see Section 2.5.1.1) that the volume where both convection and diffusion are important is minimal. However, 'Trumpet models' assume that gas transport in the conducting airways is dependent on both convective and diffusive processes.

These models assume that the cross-sectional area of the airways ($A_{aw}(x)$) can be treated as a smooth function of 'distance' (x) into the lungs from the mouth [82, 102]. The form of $A_{aw}(x)$, which increases rapidly at the end of the respiratory tree, gives this class of models their name [127, 131].

To derive the governing partial differential equation (PDE) several further assumptions are made.

- Mixing perpendicular to the axis of the respiratory tree is instantaneous.
- The left and right lungs are identical and all branching is symmetric.
- The volume (and cross-sectional area) of the airways is constant over time.
- The total volume of the alveoli expands by the tidal volume, V_T, during each breath.
- Alveoli are present from a distance of x_{alv} into the lungs and are the only place where gas exchange with the blood occurs.
- The length of the airways, L, is constant and all gases are incompressible.

The governing equation of the model is then [82, 102],

$$(A_{aw} + A_{alv})\frac{\partial F}{\partial t} + vA_{aw}\frac{\partial F}{\partial x} = D\frac{\partial}{\partial x}\left(A_{aw}\frac{\partial F}{\partial x}\right) + S(x,t), \tag{2.13}$$

where A_{aw} and A_{alv} are the cross-sectional areas of the airways and alveoli respectively, v is the gas velocity, D is the diffusion coefficient of the tracer gas, and $S(x,t)$ is a source (or sink) term for gas diffusing from (or into) the blood.

Solving this equation accurately is computationally intensive [131]. Thus, in general, these models have been employed to gain mechanistic insight into the behaviour of the lungs rather than used in the estimation of parameters from experimental techniques. For example, they have been used to look at: stratification of gas in the lung [103]; to investigate the sources of the alveolar slope during the N_2 washout [82, 83], arguing that the alveolar slope cannot solely be caused by stratification within the lung; and the effect of different breathing patterns on CO_2 transfer during normal breathing [135], where it was shown that the breathing pattern (even for the same total alveolar ventilation) can affect the elimination of CO_2 from the lungs.

Before proceeding with our discussion of the different mathematical models of the lungs, we use the Trumpet model to compare the magnitude of diffusive and convective transport in the conducting airways.

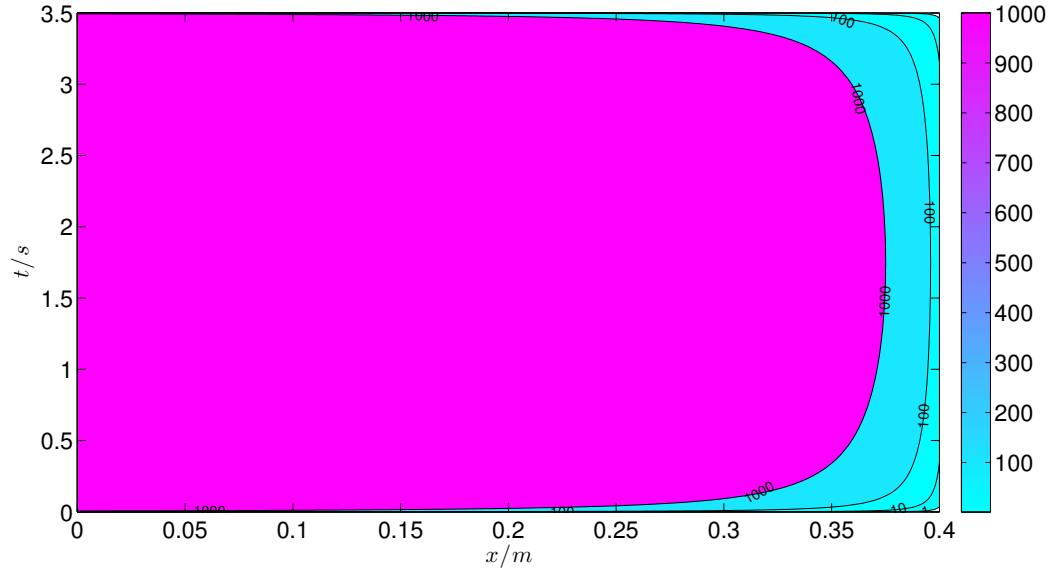

Figure 2.6: A contour plot showing the variation of the Peclet number (Pe) in space and time in the conducting airways. The only regions in which diffusion is an important method of gas transport (where $Pe \ll 1$) are the small white regions in the top and bottom right hand corners of the plot.

2.2.3 Diffusion in the conducting airways

In the conducting airways we assume that there are no alveoli, and thus A_{alv} and $S(x,t)$ are zero. Thus (2.13) becomes,

$$A(x)\frac{\partial F}{\partial t} + v(x,t)A(x)\frac{\partial F}{\partial x} = D\frac{\partial}{\partial x}\left(A(x)\frac{\partial F}{\partial x}\right), \tag{2.14}$$

where we have dropped the subscript $_{aw}$ from the cross-sectional area of the airways for simplicity of notation.

If we non-dimensionalise (2.14) we have

$$\frac{1}{\tilde{t}}\hat{A}\frac{\partial F}{\partial \hat{t}} + \frac{\tilde{v}}{\tilde{x}}\hat{v}\hat{A}\frac{\partial F}{\partial \tilde{x}} = D\frac{1}{\tilde{x}^2}\frac{\partial}{\partial x}\left(\hat{A}\frac{\partial F}{\partial \hat{x}}\right), \tag{2.15}$$

where symbols with a tilde denote typical dimensions of the parameters, and those with a hat the non-dimensionalised variables. Re-arranging and simplifying we have,

$$\frac{\tilde{x}\tilde{v}}{D}\left(\hat{A}\frac{\partial \hat{F}}{\partial \hat{t}} + \hat{v}\hat{A}\frac{\partial \tilde{F}}{\partial \tilde{x}}\right) = \frac{\partial}{\partial x}\left(\hat{A}\frac{\partial \hat{F}}{\partial \hat{x}}\right), \tag{2.16}$$

where $\frac{\tilde{x}\tilde{v}}{D}$ is known as the Peclet Number, Pe, which represents the ratio of convective to diffusive transport. $Pe \gg 1$ indicates diffusion is negligible and convection dominant, while $Pe \ll 1$ indicates that diffusion is important and convection may be neglected [119]. To estimate the value of Pe in the conducting airways we need appropriate values for $\tilde{x}$, for which we use the length of the conducting airways (L), D for which we use the diffusion

coefficient of N_2, and $\tilde{v}$. However, the gas velocity, $v(x,t)$, and therefore the appropriate value of $\tilde{v}$, varies significantly both in time (due to changes in the flow rate), and in space (due to the large variation in A down the respiratory tree). Therefore, we do not report a single number for Pe in the conducing airways. Instead, we consider the variation of $\tilde{v}$ with x and t, and use this to estimate how Pe varies in space and time in the conducting airways.

$v(x,t)$ is related to $u(t)$ (the flow rate measured at the mouth) by the relationship

$$v(x,t) = \frac{u(t)}{A(x)}. \tag{2.17}$$

Thus to find the appropriate value of $\tilde{v}$ at particular values of x and t we assume,

$$\tilde{v} = \frac{1}{A(x,t)} u(t). \tag{2.18}$$

One form of $A(x)$ used in the literature based on anatomical data is [131]

$$A(x) = \frac{\alpha}{(\beta - x)^2}. \tag{2.19}$$

Values for α and β, parameters that govern the shape of $A(x)$, are shown in Table 2.1. We assume that the flow rate varies sinusoidally during breathing [102], and thus during expiration

$$u(t) = -\frac{\pi V_T}{2t_e} \sin\left(\pi \frac{t}{t_e}\right), \tag{2.20}$$

where V_T is the tidal volume and t_e is the length of expiration.

Then using (2.19) and (2.20), we have on expiration

$$\tilde{v} = \frac{\pi V_T}{2t_e} \sin\left(\pi \frac{t}{t_e}\right) \frac{(\beta - x)^2}{\alpha}. \tag{2.21}$$

Using this expression we can then calculate Pe at different points in time and space in the conducting airways. The results for expiration are shown in Figure 2.6; the plot for inspiration would look very similar. The regions of the plot where $Pe \ll 1$ are almost impossible to see (they are at the small white regions at the bottom and top right of the plot, representing the end of the deadspace at the beginning and end of expiration respectively). This indicates that diffusion is probably only important at the end of the deadspace and when $u(t)$ is very close to zero (i.e. at the end of inspiration/expiration).

Thus our analysis indicates that convective transport is probably dominant in the deadspace. This supports the experimental evidence we will present in Section 2.5.1.1, that the volume of the lung in which both convection and diffusion are important is small.

Symbol	Value	Parameter
L	40 cm	Length of the conducting airways
α	100 cm^4	Shape parameter
β	40.55 cm	Shape parameter
V_T	600 ml	Tidal volume
t_e	3.5 s	Duration of expiration
D	0.25 cm^2	N_2 diffusion coefficient

Table 2.1: Parameter values used to consider the importance of diffusion in the conducting airways. These values are taken form [131].

2.2.4 Biophysical models

Biophysical models are considerably more complex than any of the models presented thus far. A good review of key considerations when setting up models of this form is presented in [111]. Briefly however, they involve using a space filling branching algorithm and imaging data to define a finite-element mesh for the airway (and capillary) structure [110, 112]; and the solution of coupled tissue mechanics, airflow, and gas transfer models on this mesh [14, 84, 108].

These equations combine non-linear elasticity and fluid dynamics in three dimensions, which are very complicated problems, even before their coupling is considered, and thus they are very intensive to solve computationally [42]. Along with their large parameter sets, and the difficulty of determining boundary conditions from experimental data, this means they are probably not appropriate for the kind of parameter estimation we attempt in this thesis. However, they can be extremely useful in considering mechanistic questions about lung function.

For example, they have been used to study: the influence of airway asymmetry and bronchoconstriction on the alveolar slope during the N_2 washout [73, 113], where they have shown that patchy bronchoconstriction may be insufficient to affect the value of one of the indices used to interpret the N_2 washout, S_{cond} (see Section 2.6 for more details on these indices); how the location of particle deposition after inhaler use (important for the deliver of medicine to asthma sufferers) depends on particle size [75]; and mechanisms relating to emphysema and its progression [107].

2.3 The inverse problem

To use a mathematical model to gain understanding of a physical process there are two classes of problem,: (i) the 'forward' problem, in which a model with known structure and parameters is used to predict outputs that can be measured experimentally; and (ii) the 'inverse' (or 'backward') problem, in which parameters are found to ensure that the models output to match experimental observations. Defined most broadly the inverse problem involves defining both the structure of the model ('model selection') and its parameter

values ('parameter recovery'). In this thesis, model selection is performed mostly using qualitative comparisons with experimental data, however we note more formal methods exist, see for example [114].

Parameter recovery typically involves finding the parameter values which minimise an objective function R,

$$\min R(\hat{y}(\theta), y), \tag{2.22}$$

where $\hat{y}(\theta)$ is data simulated with parameters θ, and y is the experimentally measured data. We consider objective functions which are sum of squared errors (SSE),

$$R = \sum_i w_i \left(\hat{y}_i(\theta) - y\right)^2, \tag{2.23}$$

where i are the data-points and w_i weighting factors. Typically we are interested in three things.

1. Determining which parameters are identifiable from the experimental data.

2. Finding the values of the parameters that minimise (2.23).

3. Estimating the uncertainty of the recovered parameters.

A parameter of a model is identifiable if its value can be determined from data experimental data. Non identifiability may be structural, in which case even noise-free data does not give unique parameter estimates, or practical, if unique estimates are only possible using data with a higher signal-to-noise ratio than is available [46]. Such problems can sometimes be dealt with by changing either the model or the experimental design [62]. Two of the most common methods of investigating these issues are maximum likelihood (ML) methods and Bayesian inference.

In ML methods the parameters, θ, that minimise (2.23) are found (often using non-linear least squares algorithms), and the curvature of R around this minimum is used to derive confidence intervals which are consistent with the level of experimental noise [1, 20, 46].

Another approach to finding confidence intervals is bootstrapping [16]. This involves calculating the residuals between the output of a model (with fitted parameters), and the experimental data. It is assumed that the size of the residuals are independent of their position in the dataset, and thus they can be randomly re-ordered. These residuals are added to the model's output to form a synthetic dataset that is assumed to be representative of a noisy experimental dataset. A number of different synthetic datasets (with different, random ordering of the residuals) are produced and the variation in the estimated parameters is used to estimate of the uncertainty in the estimated parameters.

The approach in Bayesian inference is different. We treat the parameters θ, as well as the data y, as random variables, and aim to calculate the posterior probability distribution

(over all of parameter space) of a particular θ given the data y_i, $p(\theta|y_i)$. To estimate this Bayes' rule is used,

$$p(\theta|y) \propto p(y_i|\theta)p(\theta), \tag{2.24}$$

where $p(y_i|\theta)$ is the probability of observing the data given the parameters (which given a model and observed data is relatively straightforward to calculate), and $p(\theta)$ the prior probability of observing the parameters, which is set based on the prior knowledge (or lack thereof) of the system. The advantage of Bayesian techniques is that the posterior distribution allows rigorous estimation of identifiability, best fit parameters, and uncertainty. However, calculation of the posterior distribution, even when using efficient numerical methods such as Markov chain Monte Carlo (MCMC) or approximate Bayesian computation (ABC) can be very computationally expensive, particularly for high dimensional models [1, 45, 46, 114].

2.4 Spirometry

There are a number of different Pulmonary Function Tests (PFTs), for example: plethysmography [19, 126]; CO diffusing capacity [64]; and helium dilution [69]. More details of each of these techniques can be found in the references, but briefly: Plethysmography and helium dilution involve estimating a patient's FRC, and then its comparison with a normal value for their height, weight, age, and gender; while measurement of the CO diffusing capacity involves measurement of the uptake of CO, to quantify factors limiting the rate of gas exchange. However, by far the most common lung function test is spirometry, because it is the easiest and cheapest to perform.

During spirometry a subject breaths through a mouthpiece and the flow (and therefore the volume) is measured as a function of time. The standard procedure for spirometry involves measurement of the Forced Vital Capacity (FVC), the maximum volume of air that a subject can expire with a forced effort after a maximal inspiration; and measurement of the Forced Expiratory Volume in one second (FEV_1), the maximum volume of air expired in the first second of an FVC manoeuvre. Subjects performing this manoeuvre generally require coaching to perform it correctly. Even after coaching, significant variation in a subject's performance is seen. The full protocol and criteria for acceptability of measurements are set out in [72].

Spirometry may be used to distinguish between obstructive and restrictive abnormalities. Obstructive abnormalities are normally characterised by a disproportionate reduction in maximal airflow (FEV_1) compared to FVC. Restrictive abnormalities, on the other hand, are characterised by a severe reduction in TLC despite a normal FEV_1 to FVC ratio. Other indices that are sometimes extracted from spirograms are peak expiratory flow (PEF), and the mean forced expiratory flow between 25% and 75% of

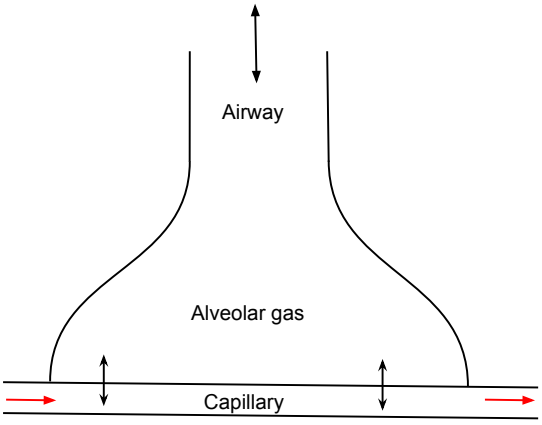

Figure 2.7: An idealised functional unit of the lung, this presentation follows the concept discussed by West [128].

FVC ($FEF_{25\%-75\%}$), which may be useful in detecting the presence of upper airway obstructions [85].

The high level of variation in performance seen between subjects, and between different laboratories conducting these tests, means that a single spirometry test, viewed in isolation, cannot reliably confirm the presence of a particular pathology [12, 21, 85]. In fact, in tests on healthy subjects 10% of readings are abnormal [36], and, as we noted in the previous chapter, the U.S. Preventive Services Task Force does not recommend screening for COPD with spirometry [63].

2.5 The functional lung unit and inhomogeneity

One important conceptual model for understanding the factors that impair the ability of the lungs to exchange gas with the blood is that of the functional lung unit [128]. A schematic diagram of a functional lung unit is shown in Figure 2.7. We consider the lung to be made up of a large number of such units.

In this functional unit gas enters, and leaves, the unit via the airway, and via exchange with the blood across the alveolar membrane. A perfect lung, one which maximises the amount of gas exchange for a given blood flow and ventilation, would be one in which each of these lung units (assuming they are of equal volume) receives an equal fraction of the blood flow (perfusion) and ventilation [128]. However, as we will see in Section 2.6, even for someone with healthy lungs this is not the case, and different parts of the lung receive different amounts of ventilation and perfusion.

We define this difference of ventilation and perfusion between different lung units as inhomogeneity of ventilation and perfusion respectively. One concept that has often been used to investigate this inhomogeneity is the ventilation-perfusion distribution. This is the

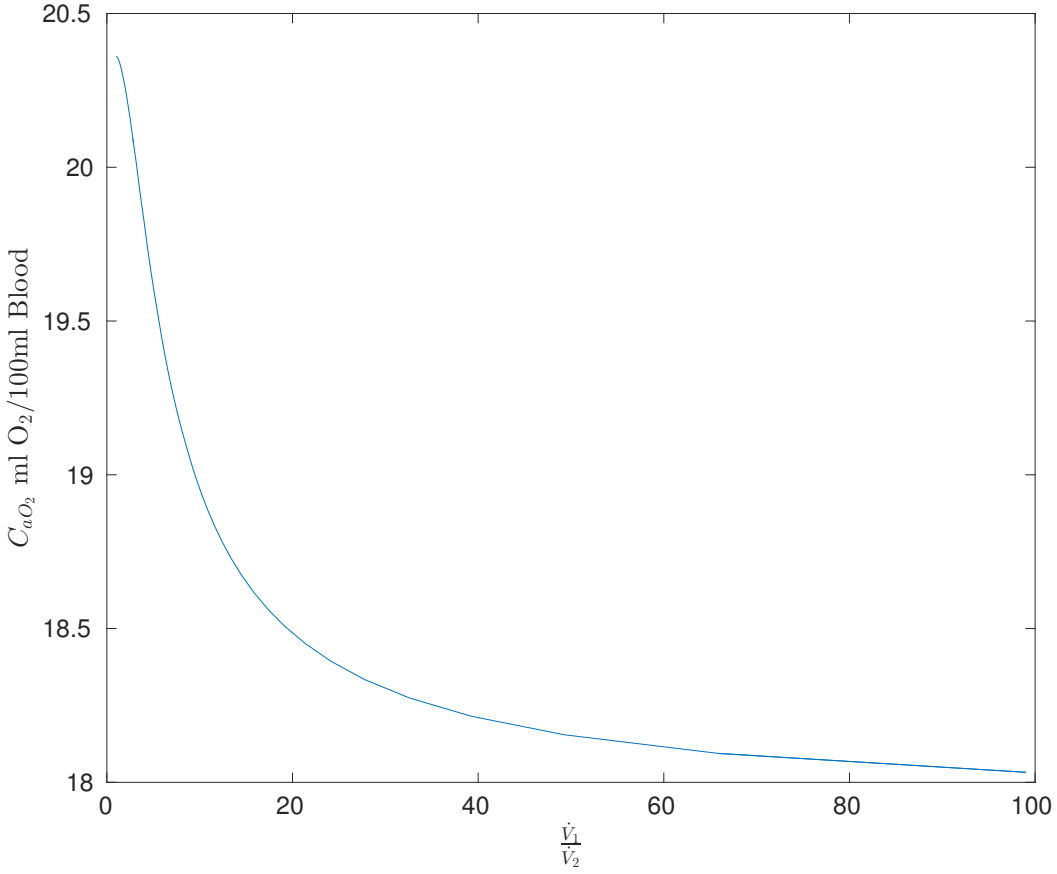

Figure 2.8: A simple illustration of the effects of inhomogeneity on gas exchange in the lung based on a simple two compartment model. This plot shows the variation of the arterial concentration of O_2 as the ratio of ventilation between the two compartments is varied. In this figure the perfusion is equally split between the two compartments.

distribution of ventilation-perfusion ratios between different lung units. We will discuss a number of methods for estimating this distribution in Section 2.6.

In Chapter 4 we present a novel lung model which incorporates inhomogeneity of ventilation, perfusion, and volume of the airways leading to lung units. These are all forms of parallel inhomogeneity [101]. However, one form of inhomogeneity we do not consider in our model is serial inhomogeneity. We define serial inhomogeneity as any inhomogeneity, due to diffusion limitation, which leads to concentration differences within a lung unit, rather than just between units. This is often referred to as stratification within the lung [101, 119]. Before explaining the experimental evidence which motivated the decision to ignore serial inhomogeneity, we will give an illustration of the consequences of parallel inhomogeneity using a simple two compartment continuous ventilation model.

2.5.1 Consequences of inhomogeneity

The importance of inhomogeneity of ventilation and perfusion on the lung's ability to take up O_2 has long been understood [90, 93]. Riley and Cournand [93] showed that the effect on the steady-state arterial O_2 and CO_2 partial pressures of any distribution of ventilation and volume within the lung can be represented by a three compartment model. This model is comprised of one deadspace compartment, one "ideal" alveolar compartment, and one shunt (venous blood that by-passes the alveoli and thus does not participate in gaseous exchange) compartment. Here we use a simple two compartment continuous ventilation model to see what the effect is on the arterial O_2 concentration, and thus the uptake of O_2, of introducing mismatch between ventilation and perfusion.

Consider a model, identical to that shown in Figure 2.4, but with two alveolar compartments, with volumes, V_{A_1} and V_{A_2}, alveolar ventilations, $\dot{V}_1$ and $\dot{V}_2$, and perfusions Q_1 and Q_2. Here we assume for simplicity that the ventilations entering and leaving a compartment are constant and equal. In steady state (when all derivatives with respect to time are zero), from (2.8) we have that

$$\dot{V}_i \left(F_I^g - F_{A_i}^g \right) + Q_i \left(C_v^g - C_{a_i}^g (F_{A_i}^{\mathrm{CO}_2}, F_{A_i}^{\mathrm{O}_2}) \right) = 0, \tag{2.25}$$

where the subscript i denotes the compartment and the superscript g either O_2 or CO_2 as relevant. We have written the arterial gas concentrations as $C_{a_i}^g(F_{A_i}^{\mathrm{CO}_2}, F_{A_i}^{\mathrm{O}_2})$ to make clear that they are functions of the alveolar gas fractions, $F_{A_i}^g$. As described in Section 2.2 we use the blood-gas model to convert from alveolar gas fractions to concentrations in the blood. We can then solve these equations simultaneously in the two compartments, for both CO_2 and O_2, to get the arterial gas concentrations in each compartment. The mixed arterial gas concentration is the perfusion weighted sum of arterial concentration from each compartment.

In Figure 2.8 we plot the arterial O_2 concentration as the ratio of the ventilations ($\frac{\dot{V}_1}{\dot{V}_2}$) of the compartments is varied while the total ventilation and perfusion are kept constant (at 4.2 l/min and 5 l/min respectively). In this figure the perfusion is equally split between the compartments, although a similar figure where we varied the perfusion and kept the ventilation equally split could be produced. In this figure we have assumed venous partial pressures of 46 mmHg for CO_2 and 40 mmHg for O_2.

Even though the total ventilation and perfusion is always the same in this figure, it is evident that increasing the inhomogeneity of ventilation reduces the arterial concentration of O_2 leaving the lungs. This demonstrates that inhomogeneity impairs the ability of the lungs to perform their primary function, exchange gas.

2.5.1.1 Diffusion limitation

One method for studying gas mixing in the lungs is to perform simultaneous washouts of He and SF_6 [119]. The reason these gases are used is that they are: inert; almost insoluble, and thus not affected by perfusion; and have very different diffusivities.

Kawashiro *et al.* [59] performed simultaneous multi-breath washouts of these two gases. These experiments involve the subject breathing the gas through a closed ventilatory system, until the concentration of the tracer gas in their lungs equilibrates with that in the rest of the system. They are then switched to breathing room air and the concentration of the tracer gas is measured as it leaves the lung.

Despite the extremely different diffusivities of He and SF_6, the profile of the washouts were very similar. Specifically they estimated the dilution factor ratio r_n, where

$$r_n = \frac{\omega_n^{\mathrm{He}}}{\omega_n^{\mathrm{SF_6}}}, \tag{2.26}$$

and ω_n is the dilution factor,

$$\omega_n = \frac{F_{ET}(n+1)}{F_{ET}(n)}, \tag{2.27}$$

where $F_{ET}(n)$ is the end tidal gas fraction on breath n.

Although there was a statistically significant difference between the two gases, it was small, with the value of r remaining close to unity (0.98-1.00). This suggests that the region in which the rate of gas transfer via diffusion is limiting is small (compared to the volume in which the gas being washed out is diluted in). Thus partitioning the lungs into a region in which only convection occurs, the deadspace, and a perfectly mixed alveolar volume seems reasonable.

There is however some evidence that the magnitude of diffusion affects the value of the normalised alveolar slope in humans [87] (although this result was not found in dogs [70]), and so models that neglect diffusion may fail to reproduce this feature. The concept of normalised alveolar slope is discussed in detail in Section 2.6.4.2.

2.6 Measuring inhomogeneity

In this section we discuss a number of experimental methods that have been used to investigate inhomogeneity in the lungs. We saw in Section 2.5.1 that inhomogeneity is a major cause of the impairment of gas exchange in the lungs. Despite this, no single 'gold standard' method for assessing inhomogeneity has been developed. However, the sheer number of methods, and the correlations with lung disease, discussed below demonstrate the potential of a measure of inhomogeneity as an indicator of lung function.

2.6.1 Hyperpolarised gas magnetic resonance imaging

In conventional proton MRI it is not possible to image the airspaces of the lung, due to their low proton density, and the high magnetic susceptibility of the surrounding lung tissue. This can be overcome by adding hyperpolarised gas (either ^{3}He or ^{129}Xe) into the inspired gas [27, 77]. The magnetic moment of these gases is typically two orders of magnitude larger than that of the water in the surrounding tissue, allowing the airways and tissue to be distinguished.

^{3}He was initially the most popular gas for hyperpolarised gas magnetic resonance imaging (HP MRI), as its gyromagnetic moment is three times larger than ^{129}Xe , and it is therefore easier to polarise. However, it is significantly more expensive and difficult to acquire. In fact, in the US its supply is currently regulated by the government due to its use in security scanners. The medical community does receive an allocation for research, but it is likely that if the technique were adopted more widely ^{129}Xe would have to be used [27, 77].

We will now consider four applications of HP MRI that have been used to study the inhomogeneity of the lungs.

Ventilation imaging: In ventilation imaging a stack of 2D images are acquired during a breath hold following an inspiration of HP gas. Regions of high signal intensity in the resulting images correspond to normally ventilated areas, and regions of low intensity to ventilation defects (totally or partially obstructed regions). Ventilation defects have been detected in patients with pulmonary disease using both ^{3}He [56] and ^{129}Xe [3].

There is however some difference between images taken with the two gases, with ^{129}Xe images showing larger and more prevalent defects. It has been argued that this is due to the large differences in molecular weights between the two gases, and that ^{129}Xe is more sensitive to mild obstructions [77]. However, respiratory gases have different molecular weights to either of these gases, and it has not yet been determined how the respiratory gases are affected by these mild obstructions that ^{129}Xe detects but ^{3}He does not.

^{3}He washouts have been imaged with HP-MRI, allowing calculation of the ratio of fresh gas entering a unit of the lung to the total end inspiratory volume of the unit. This allows a regional map of the ventilation-volume ratio in the lung to be produced [49].

Diffusion MRI: By carefully measuring the signal attenuation under different conditions an Apparent Diffusion Coefficient (ADC), which varies inversely with how much the diffusion of the HP gas is restricted in different areas of the lung, can be calculated [77]. Both ^{3}He [99] and ^{129}Xe [57] have been used to show that the ADC is larger in emphysema, due to the enlargement of airspaces caused by tissue destruction that is characteristic of the disease. It has also been shown that the technique is sensitive enough to detect gravitational gradients and the variation of alveolar size with age [57].

Estimation of partial pressure of O_2: The longitudinal relaxation rate of a HP gas is directly proportional to the local O_2 concentration [77]. By mapping this relaxation rate, the partial pressure of O_2 (P_{AO_2}) in the lung can also be mapped [77]. If the composition of mixed venous blood is measured, and the P_{AO_2} values are assumed to represent steady state values, then corresponding ventilation-perfusion ratios ($\frac{\dot{V}}{Q}$ ratio) for regions of the lung can be estimated using conservation of mass relationships [94]. Using this approach the authors report a $\frac{\dot{V}}{Q}$ distribution, a plot of the perfusion or ventilation to regions with a given $\frac{\dot{V}}{Q}$ ratio, which is approximately lognormal in a group of seven healthy Yorkshire pigs [94].

Gas exchange and uptake: ^{129}Xe (unlike ^{3}He) is soluble in both tissue and blood. In the lungs it exists in dynamic equilibrium between the gas and dissolved phase (approximately 2% is in the dissolved phase at equilibrium), and this, coupled with the large chemical shift of the spectroscopic peaks between the two phases, allows study of its exchange and uptake [77]. Although this method requires reasonable care to implement, it has been shown to be sensitive enough to detect the heterogeneity of gas exchange in normal subjects caused by gravitational gradients [58].

2.6.2 Positron emission tomography:

In this technique a radiotracer, originally $^{13}N_2$, is introduced into the lungs either by infusion into the blood, or via the inspired air. The regional activity of the radiotracer is then mapped using a positron emission tomography (PET) scanner. Depending on the technique, the ventilation, perfusion, or ventilation-perfusion distributions of the lungs may be estimated [4, 78, 91, 115].

Using a combination of bolus infusion, continuous infusion, and ventilation based techniques, Treppo *et al.* [115] estimated the specific ventilation (the ventilation per unit volume), perfusion, and ventilation-perfusion distribution of ten dogs. They found that the distributions were well approximated by lognormal distributions (as has been theoretically suggested [139]).

Musch *et al.* [78] used a bolus injection, followed by a breath hold (during which period the activity of the radiotracer was assumed to be proportional to the perfusion), and then monitored the activity as the tracer washed out the lungs (assuming a mono-exponential decay in intensity in a voxel, with a time constant inversely proportional to the ventilation) to map these quantities in humans.

More recently ^{68}Ga has been used as the radiotracer, as it has increased sensitivity, allowing reduction of the radiation dose of this procedure by up to 50% [4].

2.6.3 Multiple inert gas elimination technique

The multiple inert gas elimination technique (MIGET) is used to investigate the inhomogeneity of ventilation-perfusion matching in the lung. Six inert gases with differing solubilities are dissolved into a saline solution and infused into a patient at a constant rate. After 20 minutes (the delay is to allow the inert gas concentrations to settle down to a steady state), retention values (the arterial partial pressure divided by the mixed venous partial pressure) are calculated for each gas [47, 125].

It is then assumed that these values represent steady state values (although even at 'steady state' the partial pressures will vary due to the tidal ventilation of the lung [133]), and that conservation of mass can be applied to the exchange of gas with the blood. This allows the derivation of an equation which relates gas retention to local ventilation-perfusion ratio. Six gases with different solubilities are used because each one is sensitive to different values of this ratio. A simple 50 compartment lung model, in which each compartment has a different ventilation-perfusion ratio, is fitted (assuming that the ventilation-perfusion ratios are subject to a smoothing function) to the data to estimate the ventilation perfusion distribution of the lungs [124]. However, given this effectively involves fitting 50 unknowns to six data points, the information content of this estimation procedure has been questioned [132].

Healthy subjects appear to have narrow unimodal distributions, older subjects broader unimodal distributions, and subjects with COPD may have multi-modal distributions [125]. It has also been suggested that indicators of inhomogeneity derived from MIGET display abnormality earlier than spirometry in mild COPD patients [97].

2.6.4 N_2 washout

Since the late 1940s the multi-breath N_2 washout test (MBNW) has been used to investigate inhomogeneity of ventilation [31]. In this procedure a subject is switched from breathing room air to breathing pure O_2, and the expired concentration of N_2 is monitored as it washes out of the lungs. A consensus statement from the European Respiratory Society (ERS) and American Thoracic Society (ATS) gives a full outline of the experimental procedure, and specifications for measuring devices for collecting data from this technique [95].

Below we discuss a number of ways in which the data from this technique has been analysed, and what it has revealed about the relationship between inhomogeneity of ventilation and respiratory disease.

2.6.4.1 Lung clearance index

A number of indices of the inhomogeneity of the distribution of ventilation have been based on the observation that a N_2 washout proceeds most quickly if the lung is homogeneously

ventilated [31]. One such index is the lung clearance index (LCI), which is defined as

$$\frac{\text{Litres of ventilation required to reduce the } N_2 \text{ concentration to } \frac{1}{40} \text{ of its original value}}{\text{FRC}}. \tag{2.28}$$

The LCI has been shown to separate subjects with both emphysema and cystic fibrosis (CF) from healthy subjects [6, 11, 38, 51, 95]. Although there is variation in the values reported in the literature, LCI appears to take values of around 6.5 in healthy adults and values over 8 in those with known lung disease [95].

2.6.4.2 Normalised alveolar slope

Another approach to estimating inhomogeneity of ventilation has been to look at the slope of phase III from N_2 expirograms (see Figure 2.3) from each breath of the washout. This slope is then normalised by dividing it by the mixed expired N_2 concentration of the breath under consideration. It should be noted that this slope is non-zero even in healthy subjects [31].

It has been shown using the trumpet model (see Section 2.2.2) that this slope does not exist in homogeneous lungs [82], and experimentally that it is not caused by gravity [88]. Thus the source is probably inhomogeneity [26]. Two mechanisms have been suggested to explain this: (i) diffusion convection dependent interactions (DCDI); and (ii) convection dependent interactions (CDI) [116].

DCDI result from interactions between convection and diffusion at asymmetric branch points, located in a region where transport by convection and diffusion are of equal magnitude (the diffusion front). This is generally believed to be in the acinus of the lungs. A two branch trumpet model has been used to show that this can occur even if the two branches' emptying is synchronised and the ventilation volume ratios are identical [26]. This contribution to the normalised alveolar slope (S) is considered to be breath independent [116].

CDI result from sequential emptying between unequally ventilated lung units proximal to the diffusion front. Thus this interaction is believed to arise from inhomogeneity in conducting airways, or different pressure-volume (P-V) characteristics of large lung units. If poorly ventilated units contribute relatively more at the end of expiration, a positive slope will be generated. This effect increases in magnitude with breath number [116].

Thus Verbanck [117] established the following procedure for determining indices relating to the two interactions from a N_2 washout procedure (with a target tidal volume of 1 l on each breath).

1. Estimate FRC from total volume of N_2 collected during the washout (until the point where the concentration of N_2 is 1.5%). This is done by assuming that the

N_2 fraction is uniform everywhere within the lungs, and thus

$$\text{FRC} = \frac{\text{Cumulative expired volume of } N_2}{F_{\text{initial}}^{N_2} - F_{\text{final}}^{N_2}}. \tag{2.29}$$

2. Evaluate alveolar slope on each breath (between expired volumes of 0.65l and 1l), and normalise with the mixed expired concentration for that breath.

3. Plot this against lung turnover (total expired volume of gas divided by FRC).

4. The slope of this curve between turnovers of 1.5 and 6 (where DCDI is negligible) is defined as S_{cond}, which is considered an index of CDI.

5. S_{acin}, the index of DCDI, is then evaluated by taking the value of the normalised slope on the first breath, and subtracting the influence of inhomogeneity in conducting airways (S_{cond} multiplied by the lung turnover in breath one).

These indices have been shown to have abnormally high values after bronchoprovocation [117] and in COPD patients [118], and have been used to study the changes in lung structure in these cases.

2.6.4.3 Ventilation-volume distribution

The N_2 washout has also been used to attempt to assess the distribution of ventilation within the lung. Over sixty years ago Fowler *et al.* [32] fitted a simple three compartment model to the mean expired N_2 fractions during the washout. They found some inhomogeneity even in young, healthy subjects, which increased slightly in older subjects and markedly in those with lung disease. More recently Gustafsson *et al.* [39] fitted slow and fast ventilated lung compartments to mean expired N_2 fractions in both healthy controls and CF patients. They found the presence and size of the slowly ventilated compartment was an indicator of the severity of the disease (as assessed by other respiratory markers).

Another method of analysing the distribution of ventilation has been to attempt to recover a ventilation-volume distribution. This involves fitting a 50 compartment model (49 alveolar and one deadspace), with each compartment having a different specific ventilation (the ratio of the tidal volume received by a compartment to its volume), to the data. It has been found experimentally that the recovered distribution is narrow and unimodal for young healthy subjects, broader for older subjects, and may be multi-modal for those with respiratory disease [61]. This approach has also been taken with O_2 washin data [74].

The N_2 washout procedures used in this approach typically involve 17 breaths, and a 50 compartment model is fitted to these 17 data points. Thus, similar to MIGET, the estimation procedure suffers from under-determination, and a smoothing function must be used [61]. It has been shown, for the standard choice of smoothing function, that the

width of the recovered distribution is only weakly dependent on the width of the true distribution [132].

2.6.5 Inspired gas sinewave technique

Introduced in the mid seventies [143], the inspired gas sinewave technique involves introducing a gas with a sinusoidally varying concentration into the inspirate. This perturbation is designed to be relatively small, so that the inspired concentrations of all gases are clinically acceptable, and unlikely to affect the subject's physiology. The mean concentration in the airways and tissues of the body is allowed to equilibrate, and the partial pressure of the expired gas is continuously monitored. It is expected theoretically, and found experimentally, that the end and mixed-expired partial pressures of the tracer gases also vary sinusoidally, but with a different amplitude and phase to the forcing signal.

A mathematical model is then used to relate the amplitude and phase of the partial pressure of the expired gas to the inspired forcing signal. Inversion of this model then provides an estimate for the deadspace volume, FRC, and, if a soluble gas is used, the cardiac output [34, 35, 138].

Zwart *et al.* [143] used a simple mathematical model to establish conditions for the solubility of the tracer gas, and the period of the inspired sinewave, that allow recirculation of the signal in the mixed venous blood returning to the lungs to be ignored. This led to the use of halothane (a very potent anaesthetic agent) in extremely low concentrations (0.03%) over a period range of 0.5 to 4 minutes.

However, such low concentrations require extremely accurate concentration measurements, thus the significantly less potent anaesthetic agent nitrous oxide was used in an animal study, despite this gas not fulfilling Zwart's solubility condition. This allowed a mean value for the signal of 6% without causing any physiological effects [41]. Theoretical support for ignoring recirculation for a wider range tracer gas solubilities than previously suggested was provided by [34].

Another very attractive tracer gas is O_2, as this avoids the need for an anaesthetic agent and can be used at a high mean concentration. However, there is the added complication that O_2 is consumed by the body, and this rate of consumption can vary significantly between individuals. It should be noted, however, that there is some theoretical evidence that, in steady state, this consumption can easily be compensated for, and that the amplitude and phase of the expired sinewave for the insoluble gas Ar, and O_2 (after adjustment for the consumption) should be identical [40]. This has been investigated experimentally with varying conclusions [44, 137].

Ventilation volume distributions have also been investigated in the context of the sinewave technique. Whiteley *et al.* [134] simulated data using a model based on a continuous ventilation volume distribution. They then attempted to recover the alveolar volume using a tidal, single compartmental model at a range of frequencies (as a recovery

procedure for inhomogeneous lungs in the sinewave technique is yet to be defined). It was observed that the more inhomogeneous the distribution, the greater the difference between the values at high and low forcing frequencies. Thus they defined the following index of inhomogeneity (I),

$$I = \frac{V_A(5) - V_A(1)}{V_A(1)}, \tag{2.30}$$

where $V_A(T)$ is the estimated alveolar volume at a period of T minutes. This index was shown to be able to distinguish between broad and narrow unimodal, and bimodal distributions, and to be robust to the magnitude of error expected in practice [134].

2.6.6 Deadspace volume

An important determinant of alveolar ventilation is the volume of deadspace, the volume of the lungs that does not participate in gas exchange with the blood, and how it is distributed throughout the lungs. One of the most common methods for estimating the volume of deadspace is the single breath test for CO_2 (SBT-CO_2). During this test the fraction of CO_2 (F^{CO_2}) and the volume of air expired are measured at all times during a single expiration.

It can be assumed that the atmospheric value of F^{CO_2} is approximately zero, thus the only source of CO_2 in the lungs is diffusion across the alveolar membrane from the blood. As gas exchange only occurs in the alveoli, it is assumed that these contain CO_2 at the end of inspiration but the conducting airways do not. Thus considering the deadspace and alveoli as separate, well mixed, compartments we assume the deadspace has a F^{CO_2} of zero, whereas the alveolar compartment has a F^{CO_2} equal to the end tidal value we measure (air from the alveoli is expired last). Applying conservation of mass we can then derive the simple Bohr formula to estimate the volume of deadspace [30],

$$V_D = V_T \left(1 - \frac{F_E^{CO_2}}{F_{ET}^{CO_2}}\right), \tag{2.31}$$

where V_D and V_T are the deadspace and tidal volumes respectively, and $F_E^{CO_2}$ and $F_{ET}^{CO_2}$ are the mixed expired and end tidal CO_2 fractions respectively.

A modified 'functional' deadspace has also been estimated from the MBNW technique and its difference from the Bohr deadspace used as an index of ventilatory efficiency for these manoeuvres [5].

2.6.7 Summary

The sheer variety of methods used to investigate inhomogeneity in the lungs presented in this section demonstrate the interest in a method of quantifying inhomogeneity suitable for widespread clinical use. So far none of the methods presented in this section has

gained widespread adoption. There are a variety of reasons for this, for example, the complication of performing them (MIGET [125]), or their expense (HP MRI [27]).

Of the methods presented here the N_2 washout is perhaps the most straightforward, as rather than exotic tracer gases such as those used in MIGET or HP MRI, it simply involves a subject breathing high levels of O_2 for a sustained period. In the rest of this thesis we present a novel method of analysing N_2 washout data that may provide a mechanistic measure of inhomogeneity in the lungs.

Chapter 3

Methods Summary

This chapter summarises the mathematical model (the "lognormal lung") of the lungs and how its parameters are estimated from experimental data. This is designed to give the reader an overview of these areas, so that the full details of the model (presented in Chapter 4) and the parameter estimation procedure (presented in Chapter 5) may be more readily understood. References to these more detailed discussions are made throughout this chapter. Some of the text and figures from this chapter are repeated in those chapters for clarity.

3.1 Structure of the model

Our model is designed to simulate a patient breathing on the LGA; it takes the flow rate and compositions of the inspired gas and mixed venous blood as inputs, and then returns the compositions of the expired gas and arterial blood as a function of time.

A schematic diagram of the components of our model is displayed in Figure 3.1. In our model we assume that all gas exchange and volume change within the lung occurs in perfectly mixed alveolar compartments. Each alveolar compartment has an associated volume of personal deadspace, and each of these personal deadspace compartments is ventilated via the common deadspace compartment. We assume that the common deadspace compartment (V_{DC}) is the known volume of deadspace introduced by the measuring apparatus, while the total volume of the alveolar compartments when the lung is at FRC (V_A), and of the personal deadspace compartments (V_{DP}) are parameters which must be determined for each test subject.

Not shown in this diagram is our model of the recirculation of arterial nitrogen, for which we use a simple four compartmental model from the literature [7], or the recirculation of arterial O_2 for which we use a single body compartment of volume V_L. These are discussed further in Section 3.2.2

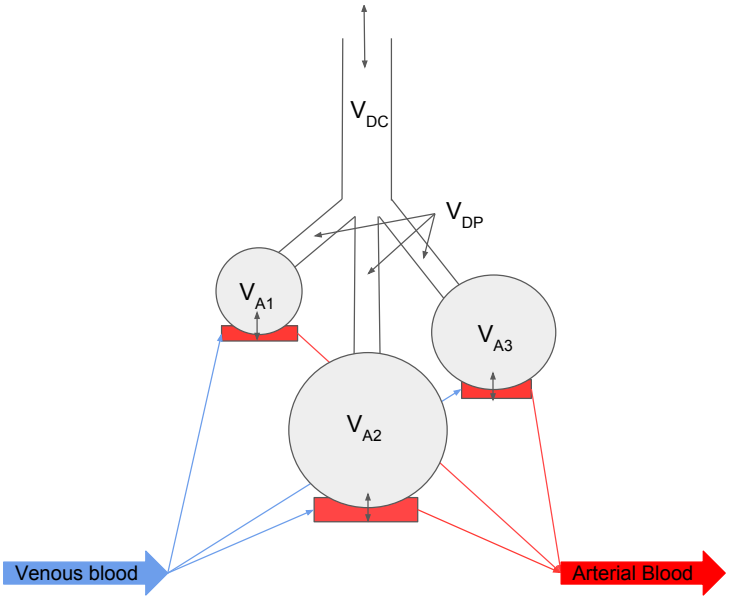

Figure 3.1: A schematic diagram of our model. We consider the lung to be composed of three different types of volume. N (here $N = 3$) perfectly mixed alveolar compartments (grey circles), N personal deadspace compartments (V_{DP}) that connect each alveolar compartment to the common deadspace compartment (V_{DC}), which we assume represents the deadspace introduced by the measurement apparatus.

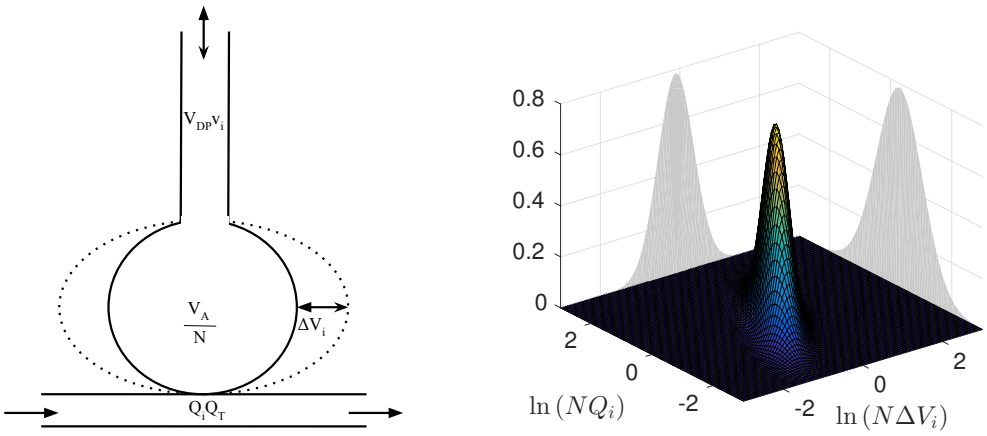

(a) (b)

Figure 3.2: (a) A schematic diagram of an alveolar unit. V_{DP} is the total volume of personal deadspace of which a fraction v_i is connected to compartment i. V_A is the total alveolar volume, and ΔV_i is the fractional alveolar volume change ("compliance"). Q_T the total perfusion, of which a fraction Q_i perfuses compartment i. (b) The bivariate lognormal distribution of compliance and perfusion.

Symbol	Description
V_A	Total alveolar volume when the lungs are at FRC
V_t	CO_2 equivalent tissue volume
Q_T	Cardiac output
σ_x	Width of the compliance distribution
σ_y	Width of the perfusion distribution
ρ	Correlation of compliance and perfusion
V_{DC}	Apparatus deadspace volume
V_{DP}	Total personal deadspace volume
σ_z	Width of the deadspace distribution

Table 3.1: A summary of the parameters that describe our model. V_{DC} is set equal to the measured apparatus deadspace volume of 155 ml, but all the other parameters must be estimated.

3.1.1 Inhomogeneity

To model inhomogeneity in the lung we consider the alveoli as N ($N = 3$ in Figure 3.1) perfectly mixed compartments in which gas exchange with the blood occurs, and which expand during inspiration and contract during expiration. Each alveolar compartment has equal volume at FRC, but differs in its fractional share of the total perfusion (Q_i), compliance (ΔV_i), and personal deadspace volume (v_i). A diagram of a single alveolar unit in which each of these fractional shares is illustrated is shown in Figure 3.2a.

If each of these fractional shares were a unique parameter, our model would have an extremely large parameter count. Thus instead a key assumption of our model is that the fractional compliance and perfusion are governed by a bivariate lognormal distribution, illustrated in Figure 3.2b and outlined in detail in Section 4.1. This distribution has three parameters that must be determined for each subject: (i) σ_x, the width of the compliance distribution; (ii) σ_y, the width of the perfusion distribution; and (iii) ρ, the correlation between compliance and perfusion.

The fractional perfusion and compliance of the N compartments are then determined by discretising this continuous distribution. This process is explained in detail in Section 4.2.1. Similarly, the fractional deadspace volume is determined by discretising the deadspace distribution (which has one parameter, the width of the deadspace distribution σ_z) as described in Section 4.2.2. Thus the structure of the model is defined by the nine parameters summarised in Table 3.1.

3.2 Governing Equations

3.2.1 Alveolar compartments

We assume that the alveolar compartments are perfectly mixed, and that the gas exchange between the blood and the alveolar compartment reaches equilibrium. Since N_2 and O_2

are almost insoluble in lung tissue we assume that the volume of these gases dissolved in lung tissue will be negligible. By conservation of mass the rate of change of the volume of gas in a lung unit is equal to the sum of the flux entering via ventilation from the personal deadspace, and the net flux due to exchange with the blood across the alveolar membrane. Thus for N_2 and O_2 during inspiration we have [35, 42],

$$\frac{d\left(F_i^g(t)V_{Ai}(t)\right)}{dt} = \dot{V}_i^{in}u(t)F_{Ii}^g(t) + Q_T Q_i \left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right), \tag{3.1}$$

where $F_i^g(t)$ is the gas fraction g (where g may be O_2 or N_2) in compartment i at time t, $\dot{V}_i^{in}$ is the fraction of the flow measured at the mouth that enters compartment i during inspiration (the calculation of this quanitity is discussed in detail in Section 4.3.1), $F_{Ii}^g(t)$ is the fraction of gas g inspired into compartment i from the connected personal deadspace compartment, Q_T is the cardiac output, $C_{\bar{v}}^g(t)$ is the mixed venous (pulmonary arterial) gas concentration of gas g entering the compartments, $C_{ai}^g(t)$ the blood gas concentration of gas g leaving compartment i, $V_{Ai}(t)$ the volume of compartment i, and $u(t)$ is the flow rate of gas (measured by the LGA). It should be noted that since we consider gas exchange with the blood, $\frac{dV_{Ai}(t)}{dt}$ is not equal to $\dot{V}_i^{in}u(t)$, as is often assumed [14], but is given by (3.5), which is derived below.

However, CO_2 is significantly more soluble than N_2 or O_2, and thus is present in appreciable volumes in lung tissue [100]. We assume that the tissue volume is a constant fraction (b) of the total alveolar volume at FRC (V_A) and distributed among the compartments in proportion to their volume at FRC. We assume that the partial pressures of CO_2 in the tissue and the alveoli are at equilibrium, and thus the volume of CO_2 dissolved in the tissue associated with compartment i is

$$V_i V_A b\lambda_t \left(P_b - P_{H_2O}\right) F_i^{CO_2}(t), \tag{3.2}$$

where λ_t is the solubility of CO_2 in lung tissue. We then define the effective CO_2 tissue volume, V_t, as

$$V_t = b\lambda_t \left(P_b - P_{H_2O}\right). \tag{3.3}$$

Therefore, including this volume, for CO_2 on inspiration, by analogy with (3.1), we have

$$\frac{d\left(F_i^{CO_2}(t)\left(V_{Ai}(t) + V_i V_t V_A\right)\right)}{dt} = \dot{V}_i^{in}u(t)F_{Ii}^{CO_2}(t) + Q_T Q_i \left(C_{\bar{v}}^{CO_2}(t) - C_{ai}^{CO_2}(t)\right). \tag{3.4}$$

Summing (3.1) (for both N_2 and O_2) and (3.4), re-arranging, and remembering that $F^{CO_2} + F^{N_2} + F^{O_2} = 1$, we have

$$\frac{dV_{Ai}(t)}{dt} = \dot{V}_i^{in}u(t) + Q_T Q_i \sum_g \left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right) - V_i V_t V_A \frac{dF_i^{CO_2}(t)}{dt}, \tag{3.5}$$

where g denotes one of the three respiratory gases. Thus as we noted earlier $\frac{dV_{Ai}(t)}{dt}$ is not equal to $\dot{V}_i^{in}u(t)$.

On expiration the flux due to ventilation is leaving rather than entering the compartment. Thus (3.1) becomes [35, 42]

$$\frac{d\left(F_i^g(t)V_{Ai}(t)\right)}{dt} = \dot{V}_i^{ex}u(t)F_i^g(t) + Q_TQ_i\left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right), \tag{3.6}$$

where $\dot{V}_i^{ex}$ is the fraction of the flow measured at the mouth that leaves compartment i during expiration (the calculation of this quanitity is discussed in detail in Section 4.3.1), and g is O_2 or N_2. Similarly for CO_2 we have

$$\frac{d\left(F_i^{CO_2}(t)\left(V_{Ai}(t) + V_iV_tV_A\right)\right)}{dt} = \dot{V}_i^{ex}u(t)F_i^{CO_2}(t) + Q_TQ_i\left(C_{\bar{v}}^{CO_2}(t) - C_{ai}^{CO_2}(t)\right). \tag{3.7}$$

Summing (3.6) (for both N_2 and O_2) and (3.7), and re-arranging, gives the governing equation for compartmental volumes on expiration,

$$\frac{dV_{Ai}(t)}{dt} = \dot{V}_i^{ex}u(t) + Q_TQ_i\sum_g\left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right) + V_iV_tV_A\frac{dF_i^{CO_2}(t)}{dt}. \tag{3.8}$$

The method used for solving these equations is described in detail in Section 4.6.

3.2.2 The body's gas stores

Previously we derived the governing equations for the alveolar compartments, and saw that they depend on the venous gas concentrations $C_{\bar{v}}^g(t)$. During the air breathing phase of our measurement protocol, these values are relatively stable and can be calculated from the data without the need for a model of recirculation. During the high O_2 breathing phase, the concentrations of both N_2 and O_2 in the pulmonary venous blood change substantially, and so models for recirculation must be employed.

To model the body's N_2 stores, and their exchange with the blood, we adopt the model (and parameters) of Baker and Farmery [7]. This model, illustrated in Figure 3.3, assumes that the body is made up of four compartments perfused in parallel. The venous partial pressure of nitrogen leaving the l-th body compartment is found by solving the following conservation of mass equation ((21) from [7]),

$$\frac{dP_{vl}^{N_2}}{dt} = \frac{Q_TQ_l}{\lambda_lV_{tl}}\left(P_a^{N_2}(t) - P_{vl}^{N_2}(t)\right), \tag{3.9}$$

where $P_{vl}^{N_2}$ is the venous (which is assumed to be equal to the tissue) nitrogen partial pressure in body compartment l, Q_l is the fractional perfusion of body compartment l, λ_l is the blood-tissue partition coefficient of compartment l, V_{tl} is the tissue volume of

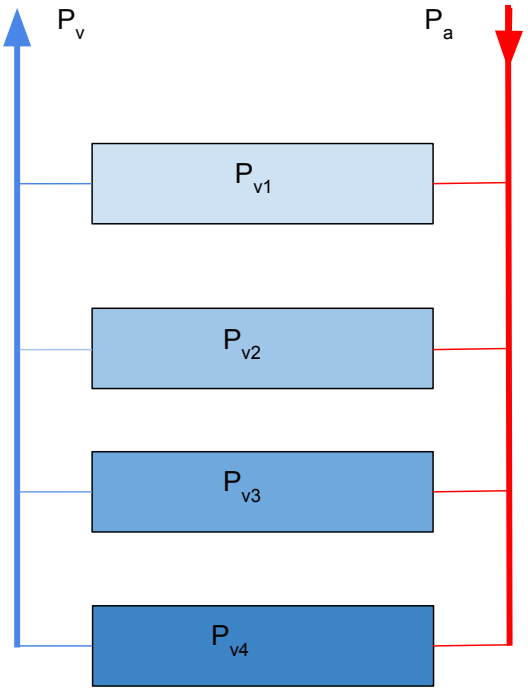

Figure 3.3: A schematic diagram of the four compartment model of the recirculation of arterial nitrogen [7]. Arrows show the arterial blood leaving the lungs and mixed venous blood returning to the lungs.

compartment l, and $P_a^{\mathrm{N}_2}(t)$ is the perfusion weighted average arterial partial pressure of N_2 in each alveolar compartment. $P_a^{\mathrm{N}_2}(t)$ is given by

$$P_a^{\mathrm{N}_2}(t) = \sum_i Q_i F_i^{\mathrm{N}_2}(t) \left(P_b - P_{\mathrm{H_2O}}\right). \tag{3.10}$$

The mixed venous nitrogen concentration entering the lung is the perfusion weighted concentration leaving the body compartments,

$$C_{\bar{v}}^{\mathrm{N}_2}(t) = S_b \sum_l Q_l P_{vl}^{\mathrm{N}_2}(t), \tag{3.11}$$

where S_b is the solubility of N_2 in the blood.

Unlike N_2, O_2 is consumed by the body, and thus if we were to adopt the four compartment model for O_2, we would have to determine how this consumption was distributed among the compartments. We lack information with which to do this, and thus for the O_2 stores we assume a single body compartment. Further, we assume there is a constant rate of O_2 consumption, $\dot{V}_{\mathrm{O}_2}$, which is the same as that we measure at the mouth during the air breathing period. Thus the governing differential equation for $C_{\bar{v}}^{\mathrm{O}_2}$ in the high O_2 breathing period is

$$\frac{dC_{\bar{v}}^{\mathrm{O}_2}}{dt} = \frac{1}{V_L}\left(Q_T\left(C_a^{\mathrm{O}_2}(t) - C_{\bar{v}}^{\mathrm{O}_2}(t)\right) - \dot{V}_{\mathrm{O}_2}\right), \tag{3.12}$$

where V_L is the volume of the compartment and $C_a^{O_2}$ is the perfusion weighted average arterial concentration. $C_a^{O_2}$ is given by

$$C_a^{O_2}(t) = \sum_i Q_i C_{ai}^{O_2}(t). \tag{3.13}$$

The O_2 consumption, $\dot{V}_{O_2}$, is assumed to be equal to that measured at the mouth during the air breathing period, and thus

$$\dot{V}_{O_2} = \frac{1}{t_{air}} \int_0^{t_{air}} F_M^{O_2}(t) u(t) \, \mathrm{d}t, \tag{3.14}$$

where $F_M^{O_2}$ is the O_2 fraction measured at the mouth and t_{air} is the duration of the air breathing period. This leaves one additional parameter, V_L, to be estimated during the parameter recovery procedure.

3.2.3 Deadspace compartments

On the basis of both theoretical, see Section 2.2.2, and experimental, see Section 2.5.1.1, evidence we neglect gas transport via diffusion in the conducting airways (which we approximate with cylindrical deadspace compartments) and assume all transport is via convection.

However, even in a simplified geometry such as a cylindrical pipe, many different patterns of flow are observed [106]. In Section 4.4 we consider two of these flow patterns: plug flow, representative of turbulent flow seen at high speeds; and Poiseuille Flow, a fully developed laminar flow.

Plug flow is the simplest flow pattern in which gas fronts entering the deadspace propagate down the deadspace unaltered. In Section 4.4 we show (by comparing the output of our model to experimental data) that this simple flow pattern (often used by other authors e.g. [133]) is a valid and sufficiently accurate way to model flow in the conducting airways.

3.3 Parameter estimation

In this section we describe the minimisation problem that we solve to estimate the parameters of the model (Table 3.1) from data collected during the measurement protocol. For our purposes a measurement protocol comprises ten minutes of air breathing followed by five minutes of breathing pure O_2.

The objective function, R, that we minimise to estimate the parameters of our model is

$$R = \sum_{\text{air}-\text{breathing}} u(t_l)^2 \left(\left(F_M^{CO_2}(t_l) - F_S^{CO_2}(t_l) \right)^2 + \left(F_M^{O_2}(t_l) - F_S^{O_2}(t_l) \right)^2 \right) + \sum_{O_2\text{breathing}} u(t_l)^2 \left(F_M^{N_2}(t_l) - F_S^{N_2}(t_l) \right)^2, \tag{3.15}$$

where F_M^g and F_S^g are the measured and simulated fractions in the measurement plane of the LGA respectively. The 'air-breathing period' for the objective function above is defined as the five minutes of normal breathing immediately preceding the high O_2 breathing phase. The first two minutes of normal breathing are discarded to allow the subject to adjust to breathing on the LGA, and the following three minutes are used to allow the model to settle into a quasi steady state.

In (3.15) we consider the CO_2 and O_2 fluxes during the air-breathing period, and N_2 flux during the high O_2 breathing period. The reason for not using N_2 during the air breathing period is that during this period the variation of N_2 is small, and thus the signal is relatively noisy.

During the O_2 breathing phase, we have only an approximate idea of how the venous O_2 and CO_2 concentrations behave. As the N_2 exchange with the blood is relatively limited, the influence of perfusion on its signal is much lower, and thus during the washout we only consider the N_2 flux.

To evaluate the objective function we must solve the forward problem, where we solve the model with known parameter values. However, when we solve the forward problem, as well as values for the parameters, we also need to know initial conditions for the differential equations, and the composition of the inspired gas and mixed venous blood. The initial conditions required are the gas fractions in every deadspace and alveolar compartment, and the starting volumes of the alveolar compartments. The calculation of the initial conditions is described in detail in Section 5.2.1.

We estimate the mixed venous concentrations of O_2 and CO_2 using a combination of our model, and data collected from the LGA. Although highly damped, these venous concentrations are not entirely constant because volunteers may hyperventilate to some degree when breathing on the LGA. For this reason we allow for a linear trend during the air breathing period. During the air breathing period, we use a continuous flow version of our model to provide initial estimates of the venous CO_2 and O_2 concentrations on every breath from the volumes of these gases exchanged at the mouth. We then fit straight lines to these values. Finally, we use the full version of our model to estimate the starting values for the venous concentrations that ensure the estimated and experimentally measured CO_2 and O_2 consumption rates are equal. This process is described in detail in Section 5.2.2.

During the high O_2 breathing phase, the recirculating arterial O_2 will lead to a rise in the venous O_2 concentrations. We used a simple model of the body's oxygen stores with a single parameter, V_L. To estimate the behaviour of the mixed venous O_2 concentration during the high O_2 breathing period, a value for this parameter must be found. This is achieved by finding the value of V_L that ensures the alveolar volume remains stable during the high O_2 breathing phase, as described in Section 5.2.3.

Since we estimate the venous concentrations with our model, they will vary as the parameters of the model change in each objective function evaluation. Thus every time we evaluate the objective function, we must: establish initial conditions; estimate the venous gas concentrations during air breathing; and find an appropriate value of V_L.

Chapter 4

The lognormal lung

In this chapter we discuss the development of a novel mathematical model of the lungs, whose parameters are intended to provide measures of the inhomogeneity of the lungs. The model will require, as inputs, various parameters that describe the lung, the total flow rate of gas at the mouth, and inspired gas fractions. Outputs will be the simulated expired gas fractions at the mouth.

We will discuss the inverse problem of estimating the parameters of our model from experimental data in Chapter 5, but we note here that almost all methods for solving the inverse problem (see Section 2.3) involve solving the forwards problem a large number of times. Thus to be suitable for our purposes the model must:

1. represent the underlying physiology, and produce an output that is a good approximation to that seen experimentally;

2. be described by a small parameter set that can be estimated from experimental data;

3. be governed by equations that may be solved quickly and accurately (conserving mass with high precision) using reasonable computational resources.

Before discussing the development of the model in detail we present an outline of the argument that this chapter will follow.

We explained in Section 2.5 that our conceptual model of the lung is that it is composed of a large number of functional lung units (Figure 2.7) connected to the mouth via the conducting airways. Each one of these units may have a different ventilation, perfusion, and volume. If, in our mathematical model, each of these were represented by an independent parameter, this would represent an enormous parameter set. Such a large parameter set would be very difficult, if not impossible, to estimate from experimental data.

Thus instead we follow Wilson and Beck [139], and assume that the volume changes, perfusions, and volumes at FRC of the lung units follow an underlying continuous distribution. In this chapter we introduce this distribution, show that it is described by only

three parameters, and compare it to the output of a more complicated biophysical model of the lung.

We assume that we can approximate this continuous distribution with a finite number of perfectly mixed alveolar compartments. We outline this, and then explain how we use a change of variables to approximate the governing continuous distribution with discrete compartments in an efficient way. This section concludes with an explanation of how the deadspace, which we assume represents the conducting airways, is discretised.

Next we present the governing conservation of mass equations for the alveolar and body compartments. Here we explain an important distinction in our model, between the fractional ventilation of a lung unit and the fractional volume change of a lung unit, and why we introduce it.

The governing equations for the deadspace compartments, which we assume represent the conducting airways of the lung, require careful consideration. On the basis of both theoretical and experimental evidence we neglect diffusion in the conducting airways, and assume mass transport is by convection alone. Having done this we derive the advection equation that governs gas transport in the deadspace compartments. Next we compare two different flow profiles, plug flow and Poiseuille flow, to experimental data. We find that plug flow provides a better approximation to the signal seen at the mouth, and thus this flow pattern is adopted.

Finally, we explain how we solve the governing equations. In the case of the deadspace compartments, we explain how to apply a change of variables to solve the governing advection equation in a manner that conserves mass with a high degree of precision. For the governing differential equations of the alveolar compartments we compare a number of different MATLAB ODE solvers, but find that we can develop a method which is faster and conserves mass with a higher degree of precision.

4.1 The governing distribution

At the start of this chapter we explained that our conceptual model of the lung was one of functional lung units (Figure 2.7) whose ventilations, perfusions, and volumes were determined by a continuous distribution. In this section we introduce the distribution, show that it is described by only three parameters, and compare the form of the distribution to the output of a biophysical model of the lung.

Wilson and Beck [139] established a theoretical framework for a distribution, encompassing both ventilation and perfusion, to allow them to reconcile experimentally measured ventilation-volume and perfusion-volume distributions, with ventilation-perfusion distributions measured using MIGET. The governing distribution of our model is similar, except we consider fractional distributions (so the distribution is independent of, for

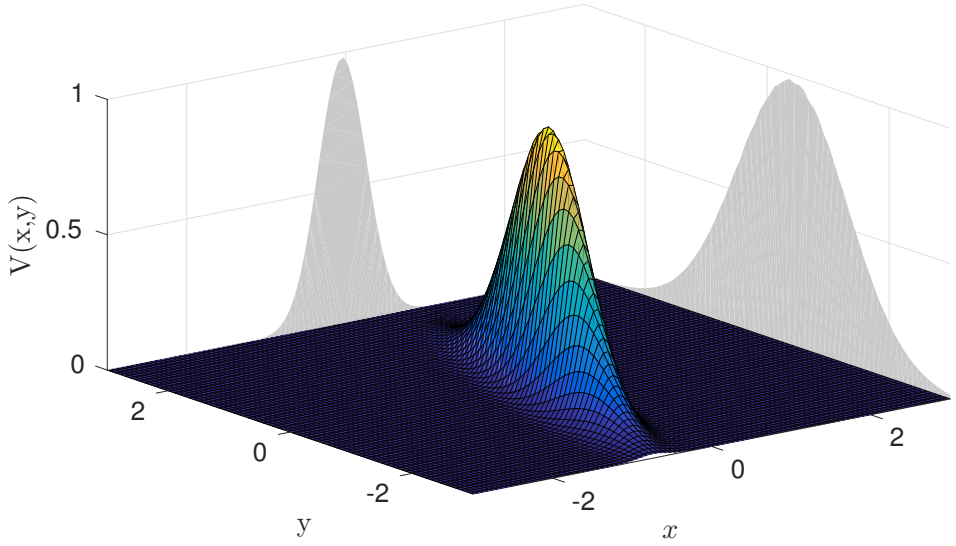

Figure 4.1: An example of a bivariate normal distribution in x and y ($\log\left(\frac{\Delta V}{V}\right)$ and $\log\left(\frac{Q}{V}\right)$ respectively). The functional form of this distribution is given in (4.2). The shadows show the projections onto the x and y axes. As shown in Section 4.1, these projections are normal distributions with, variances σ_x^2 and σ_y^2, and means $\bar{x}$ and $\bar{y}$ respectively.

example, tidal volume), and we assume that fractional volume change rather than ventilation is lognormally distributed. As we show in Section 4.3.1, this distinction is made because in our model volume change is caused by both ventilation and gas exchange.

In Section 2.1.2 we also introduced the concept of compliance, and noted that the change in volume of an alveolar unit was the product of its compliance and the pressure change it experiences. In this thesis we therefore refer to the distribution of fractional volume change as the distribution of compliance, as we feel this terminology is more intuitive. We note however, that the two are not entirely equivalent, as the pressure changes experienced by different lung units may also vary within the lung due to, for example, the effects of airway resistance or gravity [129].

Having explained this important distinction, we now describe the governing bivariate distribution. For ease of notation we define

$$\begin{aligned} x &= \log\left(\frac{\Delta V}{V}\right), \\ y &= \log\left(\frac{Q}{V}\right), \end{aligned} \tag{4.1}$$

where ΔV is the fractional alveolar volume change, V is the fractional alveolar volume at FRC and Q is the fractional alveolar perfusion.

The governing distribution of our model ($V(x, y)$) is defined with the statement that $V(x, y)\mathrm{d}x\mathrm{d}y$ is the fraction of the alveolar volume with $x = \log\left(\frac{\Delta V}{V}\right)$ and $y = \log\left(\frac{Q}{V}\right)$ values that lie in small intervals of length $\mathrm{d}x$ and $\mathrm{d}y$ around the values of x and y.

We assume that V is a bivariate lognormal distribution of $\frac{\Delta V}{V}$ and $\frac{Q}{V}$, thus

$$V(x,y) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}} \exp\left(\frac{-1}{2\left(1-\rho^2\right)}\left(\frac{\left(x-\bar{x}\right)^2}{\sigma_x^2} - \frac{2\rho\left(x-\bar{x}\right)\left(y-\bar{y}\right)}{\sigma_x\sigma_y} + \frac{\left(y-\bar{y}\right)^2}{\sigma_y^2}\right)\right), \tag{4.2}$$

where σ_x^2 and σ_y^2 are the variances of the volume change-volume and perfusion-volume distributions respectively, $\bar{x}$ and $\bar{y}$ are their means, and ρ the coefficient of correlation between x and y. An example of a distribution of this form is plotted in Figure 4.1. Evidence for the assumption of lognormality from a number of modalities (HP MRI, PET, and MIGET) is discussed in Section 2.6, and is examined for data taken from a biophysical model of the lungs in Section 4.1.1.

The volume change-volume ($V(x)$) and perfusion-volume ($V(y)$) distributions are the projections of $V(x,y)$ onto the x and y axes respectively. Thus,

$$V(x) = \int_{-\infty}^{\infty} V(x,y)\ \mathrm{d}y, \tag{4.3}$$

and

$$V(y) = \int_{-\infty}^{\infty} V(x,y)\ \mathrm{d}x. \tag{4.4}$$

Substituting (4.2) into these expressions we have

$$V(x) = \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(\frac{-\left(x-\bar{x}\right)^2}{2\sigma_x^2}\right), \tag{4.5}$$

and

$$V(y) = \frac{1}{\sqrt{2\pi}\sigma_y} \exp\left(\frac{-\left(y-\bar{y}\right)^2}{2\sigma_y^2}\right). \tag{4.6}$$

Thus $V(x)$ and $V(y)$ are univariate lognormal distributions.

We now show that this distribution is described by only three independent parameters. From (4.1),

$$\Delta V(x) = V(x)e^x, \tag{4.7}$$

substituting this into (4.5) and rearranging we have,

$$\Delta V(x) = \frac{1}{\sqrt{2\pi}\sigma_x} \exp\left(\frac{-\left(x-\left(\bar{x}+\sigma_x^2\right)\right)^2}{2\sigma_x^2}\right) \exp\left(\frac{\left(-\sigma_x^2-2\bar{x}\right)}{2}\right). \tag{4.8}$$

Since we require that

$$\int_{-\infty}^{\infty} \Delta V(x)\ \mathrm{d}x = 1, \tag{4.9}$$

we must have

$$\exp\left(-\frac{\sigma_x^2+2\bar{x}}{2}\right) = 1, \tag{4.10}$$

and so

$$\bar{x} = -\frac{\sigma_x^2}{2}. \tag{4.11}$$

Since from (4.1),

$$Q(y) = V(y)e^y, \tag{4.12}$$

we can apply an analogous procedure to obtain

$$\bar{y} = -\frac{\sigma_y^2}{2}. \tag{4.13}$$

In contrast to previous work [139] the means of the distribution, $\bar{x}$ and $\bar{y}$, are not independent parameters, and the distribution is governed by three parameters: σ_x, σ_y, and ρ. This is because we use fractional quantities, and assume our distribution is independent of the tidal volume and the cardiac output. Before discussing how we discretise this continuous distribution into the compartments that make up our model, we examine the ventilation-volume and perfusion-volume distributions of a biophysical model of the lungs.

4.1.1 Estimates of distributions from a biophysical model of the lungs

In this subsection we show that the ventilation-volume and perfusion-volume distributions taken from a biophysical model of the lungs are well approximated by a lognormal one. We used data taken from a biophysical model of the lungs [18, 108] (kindly provided by Dr A Clarke, University of Auckland), which we refer to as the "New Zealand lung". More details on biophysical lung models can be found in Section 2.2.4.

The data we have available from the New Zealand lung is the acinar volume, ventilation, and perfusion of each voxel in the lung at functional reserve capacity (FRC). We calculated the values of $\log\left(\frac{\dot{V}}{V}\right)$, and $\log\left(\frac{Q}{V}\right)$ for each voxel, and then calculated the standard deviations of these as estimates for σ_x and σ_y respectively. The correlation between the two was used as an estimate for ρ. The values we found were 0.30, 0.57, and 0.87 for σ_x, σ_y, and ρ respectively[1].

In Figure 4.2 we plot the distributions $V(x)$ and $V(y)$ from the simulated data and the distributions given by (4.5) and (4.6) using the estimated values of σ_x and σ_y. We can see in this plot that although there may be some skew in the distributions from the biophysical model the lognormal distribution is a good approximation for both distributions.

Figure 4.3 displays a comparison of the full bivariate distribution of the data from the New Zealand lung and that estimated from (4.2) using our estimated values of σ_x, σ_y, and ρ. In this case the correlation between the two variables is captured well (as seen by the similar slope and shape of the distribution). However, as we might expect, our estimated distribution is significantly smoother and more regular than the one taken from the biophysical model.

[1]For this comparison we have assumed that we can ignore the small difference between volume-change and ventilation, as our interest here is seeing if the distributions of the New Zealand lung are approximately lognormal.

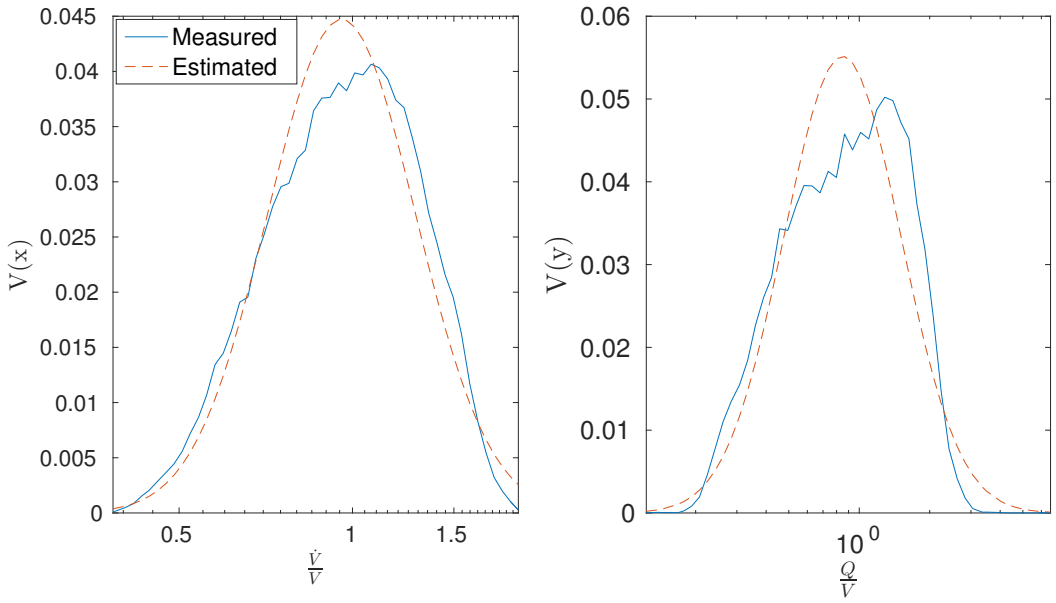

Figure 4.2: The ventilation-volume distribution (left) and perfusion-volume (right) distributions of the "New Zealand lung". On both plots we display the actual distributions from the "New Zealand lung", and the estimated distributions from (4.5) and (4.6) using values for σ_x and σ_y of 0.30 and 0.57 respectively. Although there appears to be some skew in the distributions from the biophysical model of the lungs, they are very close to lognormal.

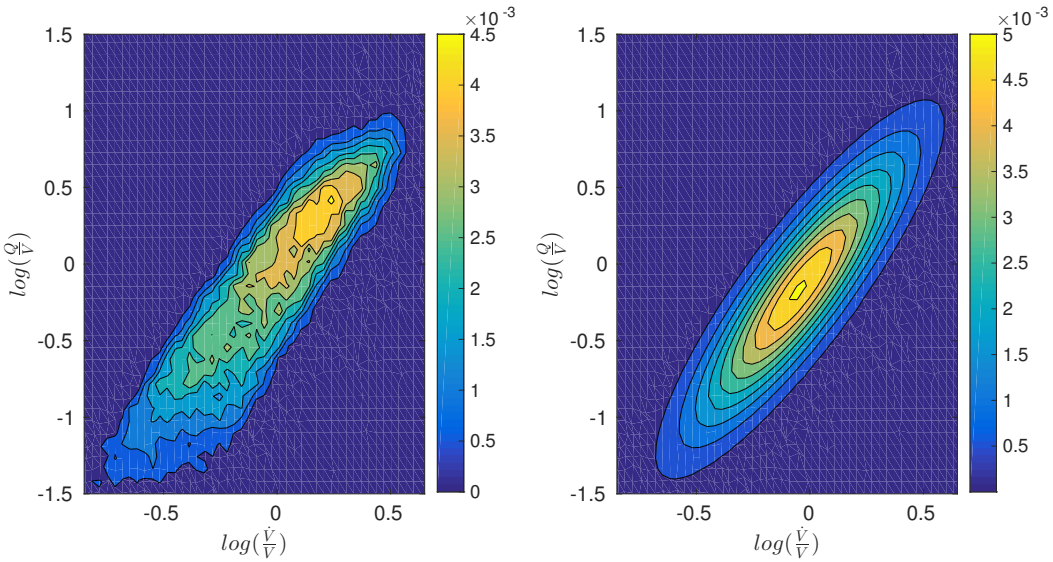

Figure 4.3: The full bivariate distribution of the "New Zealand lung" (left plot) compared to that estimated from (4.2) using values of σ_x, σ_y, and ρ of 0.30, 0.57 and 0.87 respectively (right plot). The values in the colour bars represent $V(x, y)$.

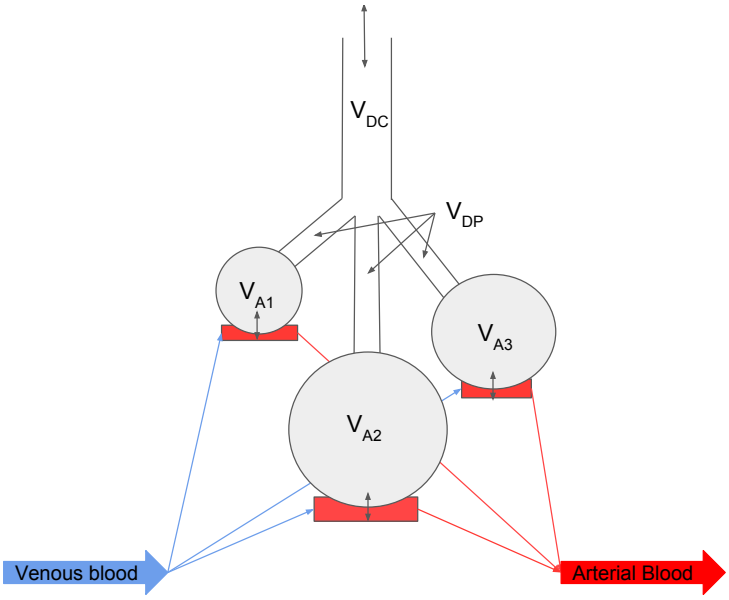

Figure 4.4: A schematic diagram of our model. We consider the lung to be composed of three different types of volume. N (here $N = 3$) perfectly mixed alveolar compartments (grey circles), N personal deadspace compartments (V_{DP}) that connect each alveolar compartment to the common deadspace compartment (V_{DC}), which we assume represents the deadspace introduced by the measurement apparatus.

4.2 A compartmental model

Our model is designed to simulate a patient breathing on the LGA; it takes flow and inspired gas (O_2 and CO_2) fractions measured every 10ms as inputs, and then returns the expired CO_2, O_2, and N_2 fractions in the measurement plane of the LGA (the point at which the LGA measures the flows and gas fractions) as a function of time.

In Section 2.2 we discussed three different classes of model. Here we develop a 'compartmental model'. We do not use a trumpet or biophysical model because, as discussed in Section 2.2, these models require precise, subject specific knowledge of the geometry of the airways (which we will not have for subjects on the LGA), and are computationally intensive to solve.

A schematic diagram of the components of our model is displayed in Figure 4.4. In our model we assume that all gas exchange and volume change within the lung occurs in perfectly mixed alveolar compartments. Each alveolar compartment has an associated volume of personal deadspace, and each of these personal deadspace compartments is ventilated via the common deadspace compartment. We assume that the common deadspace compartment (V_{DC}) is the known volume of deadspace introduced by the measuring apparatus, while the total volume of the alveolar compartments when the lung is at FRC (V_A), and of the personal deadspace compartments (V_{DP}) are parameters which must be determined for each test subject.

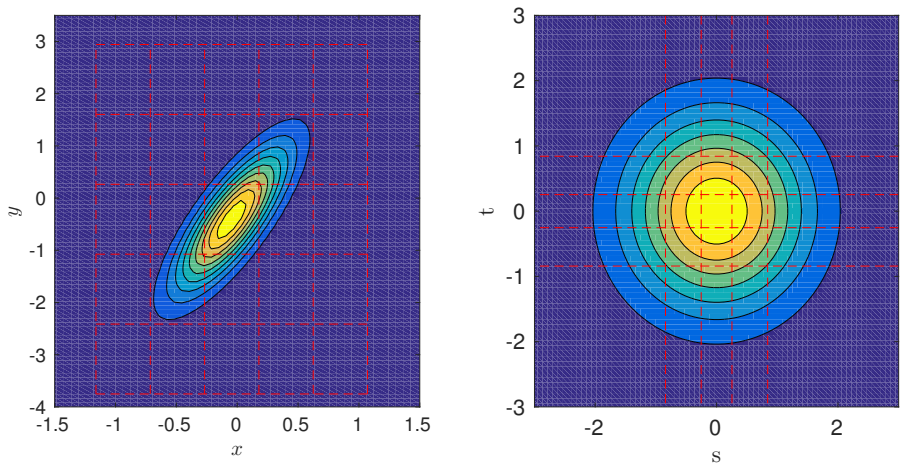

Figure 4.5: A plot comparing the two methods of discretising the distribution discussed in Section 4.2.1.1. The contour plots represent the bivariate lognormal distribution either as a function of x and y (left hand plot), or as a function of s and t (right hand plot). The dashed-lines represent the limits on the integrals that define the compartments (here $N_x = N_y = N_s = N_t = 5$), and their position is found as described in the text (the upper and lower limits on the right hand plot are at $+\infty$ and $-\infty$ respectively, and so not visible here). In the left hand plot, many of the compartments are located in points of very low probability density.

Not shown in this diagram is our model of the recirculation of arterial nitrogen, for which we use a simple four compartmental model from the literature [7], or the recirculation of arterial O_2 for which we use a single body compartment of volume V_L. These are discussed further in Section 4.3.2

In the following sections we explain in more detail how the model is parameterised.

4.2.1 Alveolar compartments

We consider the alveoli as N ($N = 3$ in Figure 4.4) perfectly mixed compartments in which gas exchange with the blood occurs, and which expand during inspiration and contract during expiration.

A key assumption of our model is that the fractional volume change-fractional volume and fractional perfusion-fractional volume ratios are governed by the underlying bivariate lognormal distribution that we outlined in Section 4.1. Thus the values of these ratios in the N compartments are determined by discretising this continuous distribution. This process is explained in more detail below.

4.2.1.1 A naive discretisation

We now explain how we approximate the continuous governing distribution of our model with $N_x N_y$ unique pairs of volume change-volume and perfusion-volume ratios. Each of these pairs is indexed with the subscripts $_{ij}$, where $i = 1, 2, \ldots N_x$ and $j = 1, 2, \ldots N_y$.

We first discretise x and y into $N_x + 1$ and $N_y + 1$ equally spaced points. Since these points are equally spaced we must choose finite upper and lower bounds, and thus x_0 is set such that

$$\int_{-\infty}^{x_0} V(x)\ \mathrm{d}x = 0.0001, \tag{4.14}$$

and x_{N_x} such that

$$\int_{x_{N_x}}^{\infty} V(x)\ \mathrm{d}x = 0.0001. \tag{4.15}$$

The values y_0 and y_{N_y} are defined by analogous expressions involving the relevant univariate distribution $V(y)$.

We can then define the fractional volume of compartment (ij) when the lung is at FRC as,

$$V_{ij} = \int_{y_{j-1}}^{y_j} \int_{x_{i-1}}^{x_i} V(x, y)\ \mathrm{d}x\mathrm{d}y. \tag{4.16}$$

Using the definitions of x and y given by (4.1), the fractional volume change (ΔV_{ij}) is

$$\Delta V_{ij} = \int_{y_{j-1}}^{y_j} \int_{x_{i-1}}^{x_i} \exp(x) V(x, y)\ \mathrm{d}x\mathrm{d}y, \tag{4.17}$$

and the perfusion Q_{ij},

$$Q_{ij} = \int_{y_{j-1}}^{y_j} \int_{x_{i-1}}^{x_i} \exp(y) V(x, y)\ \mathrm{d}x\mathrm{d}y. \tag{4.18}$$

This method of discretising the continuous distribution works well if the major and minor axes of the elliptical probability density function lie along the x and y axes (i.e. if $\rho = 0$). However if $\rho \neq 0$ the probability density function is an ellipse whose major and minor axes do not lie along the x and y axes, as was the case for the distribution estimated from the biophysical lung model, shown in Figure 4.3. In this case the approach described above, in which we discretise the distribution into equal sized rectangles along the x and y axes, will be quite inefficient; since, as shown in the left hand plot of Figure 4.5, there will be a number of compartments with almost zero volume, and one or two containing the majority of the volume.

Thus in the next section we take advantage of a geometrical property of the normal distribution that allows us to change variables so that this discretisation can be performed more efficiently.

4.2.1.2 Discretisation by change of variables

To discretise the governing continuous distribution more efficiently, we take advantage of a useful property of the bivariate normal distribution. This property is that any bivariate normal distribution can be represented as a transformation of the standard bivariate normal distribution (zero mean, zero covariance, and unit variances). In this section we follow the outline given in [24] and explain how this change of variables may be made

by formulating the bivariate normal distribution in vector notation. We then set up the integrals that we must solve to define our compartments in these new variables.

The bivariate distribution, (4.2), can be written in the following notation [24],

$$V(\mathbf{x}) = \frac{1}{2\pi \left|\boldsymbol{\Sigma}\right|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}\left(\mathbf{x}-\boldsymbol{\mu}\right)^{T} \boldsymbol{\Sigma}^{-1}\left(\mathbf{x}-\boldsymbol{\mu}\right)\right) \tag{4.19}$$

where

$$\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}, \tag{4.20}$$

$$\boldsymbol{\mu} = \begin{pmatrix} \bar{x} \\ \bar{y} \end{pmatrix}, \tag{4.21}$$

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{pmatrix}. \tag{4.22}$$

The contour lines of constant probability are ellipses centred on the mean, with major and minor axes defined by the eigenvectors of $\boldsymbol{\Sigma}$ [24].

To perform a change of variables we first diagonalise $\boldsymbol{\Sigma}$,

$$\boldsymbol{\Sigma} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^{\mathbf{T}}, \tag{4.23}$$

where $\boldsymbol{\Lambda}$ is a diagonal matrix with, as $0 \leq \rho \leq 1$, positive entries on the diagonal and $\mathbf{U}$ is a rotation matrix, and so

$$\mathbf{U} = \mathbf{U}^{\mathbf{T}}. \tag{4.24}$$

We can re-write (4.23) as

$$\boldsymbol{\Sigma} = \left(\mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\right)\left(\mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\right)^{\mathbf{T}}, \tag{4.25}$$

where

$$\boldsymbol{\Lambda}^{\frac{1}{2}} = \begin{pmatrix} \sqrt{\Lambda_{11}} & 0 \\ 0 & \sqrt{\Lambda_{22}} \end{pmatrix}. \tag{4.26}$$

If we substitute (4.25) into (4.19) and use the fact that the determinant of a matrix's transpose is equal to the determinant of the matrix, we have

$$V(\mathbf{x}) = \frac{1}{2\pi \left|\left(\mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\right)\right|} \exp\left(-\frac{1}{2}\left(\mathbf{x}-\boldsymbol{\mu}\right)^{T}\left(\mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\right)^{-\mathbf{T}}\left(\mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\right)^{-1}\left(\mathbf{x}-\boldsymbol{\mu}\right)\right). \tag{4.27}$$

To simplify this expression we define a new set of axes $\mathbf{s}$,

$$\mathbf{s} = \left(\mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\right)^{-1}\left(\mathbf{x}-\boldsymbol{\mu}\right), \tag{4.28}$$

where

$$\mathbf{s} = \begin{pmatrix} s \\ t \end{pmatrix}, \tag{4.29}$$

and thus multiplying (4.28) by $\mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}$ we have

$$\mathbf{x} = \boldsymbol{\mu} + \mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}\mathbf{s}. \tag{4.30}$$

The Jacobian of this transformation is $\mathbf{J}$, where

$$\boldsymbol{J} = \begin{pmatrix} \frac{\partial x}{\partial s} & \frac{\partial x}{\partial t} \\ \frac{\partial y}{\partial s} & \frac{\partial y}{\partial t} \end{pmatrix}. \tag{4.31}$$

Calculating this explicitly from (4.30) we have that,

$$\boldsymbol{J} = \mathbf{U}\boldsymbol{\Lambda}^{\frac{1}{2}}. \tag{4.32}$$

To change the variables of a probability density function (PDF), the PDF in the old variables is multiplied by the determinant of the Jacobian for the transformation between the two sets of variables [24]. Thus in our case,

$$V(\mathbf{s}) = V(\mathbf{x})\,|\mathbf{J}|\,. \tag{4.33}$$

Using this and substituting for $\mathbf{s}$ in (4.27), we obtain

$$V(\mathbf{s}) = \frac{1}{2\pi}\exp\left(-\frac{1}{2}\mathbf{s}^{\mathbf{T}}\mathbf{s}\right). \tag{4.34}$$

This is the standard bivariate normal distribution (zero mean, zero covariance, and unit variances) defined on $\mathbf{s}$. Thus any bivariate normal distribution defined on $\mathbf{x}$ can be represented as the standard bivariate normal distribution on $\mathbf{s}$, using a suitable change of variables. The advantage of using the standard bivariate normal distribution is that the major axes of the ellipses that form the probability density function are along the axes on which we define the distribution (see the right hand plot in Figure 4.5).

We can then define a rectangular grid on s and t, which samples the distribution effectively, and transform the integrals discussed above to find the volume changes, volumes, and perfusions associated with these (N_sN_t) rectangles (compartments).

In the new variable s and t we can re-define (4.16) as,

$$V_{ij} = \int_{t_{j-1}}^{t_j}\int_{s_{i-1}}^{s_i} V(s,t)\;\mathrm{d}s\mathrm{d}t, \tag{4.35}$$

where $i = 1\ldots N_s$, $j = 1\ldots N_t$, and

$$V(s,t) = \frac{1}{2\pi}\exp\left(\frac{-s^2-t^2}{2}\right). \tag{4.36}$$

We can factorise (4.35) into the product of two univariate integrals. If we do this we have

$$V_{ij} = \int_{t_{j-1}}^{t_j}\frac{1}{\sqrt{2\pi}}\exp\left(\frac{-t^2}{2}\right)\;\mathrm{d}t\int_{s_{i-1}}^{s_i}\frac{1}{\sqrt{2\pi}}\exp\left(\frac{-s^2}{2}\right)\;\mathrm{d}s. \tag{4.37}$$

In our new variables we can then define the limits on the integrals such that each compartment represents an equal fraction of the alveolar volume at FRC (i.e. $V_{ij} = \frac{1}{N_sN_t}$), rather

than the compartments defined previously which represented equally sized rectangles on x and y. To do this we define s_i such that

$$\int_{s_{i-1}}^{s_i} \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-s^2}{2}\right) \, \mathrm{d}s = \frac{1}{N_s}. \tag{4.38}$$

In the new variables the limits are no longer equally spaced, and so $s_0 = -\infty$ and $s_{N_s} = \infty$. Similarly we define t_j such that

$$\int_{t_{j-1}}^{t_j} \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-t^2}{2}\right) \, \mathrm{d}t = \frac{1}{N_t}, \tag{4.39}$$

where $t_0 = -\infty$ and $t_{N_t} = \infty$. A comparison between the two methods of discretising the distribution is shown in Figure 4.5.

Using (4.17), (4.30), and (4.31), the fractional volume change defined on the new axes is then

$$\Delta V_{ij} = \int_{t_{j-1}}^{t_j} \int_{s_{i-1}}^{s_i} \exp\left(\bar{x} + J_{11}s + J_{12}t\right) V(s,t) \, \mathrm{d}s\mathrm{d}t. \tag{4.40}$$

Similarly using (4.18), (4.30), and (4.31), the fractional perfusion is

$$Q_{ij} = \int_{t_{j-1}}^{t_j} \int_{s_{i-1}}^{s_i} \exp\left(\bar{y} + J_{21}s + J_{22}t\right) V(s,t) \, \mathrm{d}s\mathrm{d}t. \tag{4.41}$$

Our method for evaluating (4.40) and (4.41), which factorises these bivariate integrals into the product of two univariate integrals, again taking advantage of the change of variables, is briefly explained in the next section.

4.2.1.3 Evaluation of ΔV_{ij} and Q_{ij} in MATLAB

Re-arranging (4.40) we have,

$$\Delta V_{ij} = \frac{\exp\left(\bar{x}\right)}{2\pi} \int_{t_{j-1}}^{t_j} \exp\left(\frac{-t^2 + 2J_{12}t}{2}\right) \, \mathrm{d}t \int_{s_{i-1}}^{s_i} \exp\left(\frac{-s^2 + 2J_{11}s}{2}\right) \, \mathrm{d}s, \tag{4.42}$$

thus,

$$\Delta V_{ij} = \frac{1}{\sqrt{2\pi}} K_1 \int_{t_{j-1}}^{t_j} \exp\left(\frac{-(t - J_{12})^2}{2}\right) \, \mathrm{d}t \frac{1}{\sqrt{2\pi}} \int_{s_{i-1}}^{s_i} \exp\left(\frac{-(s - J_{11})^2}{2}\right) \, \mathrm{d}s, \tag{4.43}$$

where K_1 is given by

$$K_1 = \exp\left(\bar{x} + \frac{(J_{11})^2 + (J_{12})^2}{2}\right). \tag{4.44}$$

Analogously,

$$Q_{ij} = \frac{1}{\sqrt{2\pi}} K_2 \int_{t_{j-1}}^{t_j} \exp\left(\frac{-(t - J_{22})^2}{2}\right) \, \mathrm{d}t \frac{1}{\sqrt{2\pi}} \int_{s_{i-1}}^{s_i} \exp\left(\frac{-(s - J_{21})^2}{2}\right) \, \mathrm{d}s, \tag{4.45}$$

where K_2 is given by

$$K_2 = \exp\left(\bar{y} + \frac{(J_{21})^2 + (J_{22})^2}{2}\right). \tag{4.46}$$

Thus these integrals are the product of the integrals of two normal distributions with standard deviation equal to one and means given by the appropriate entry of $\boldsymbol{J}$. We can calculate these integrals in MATLAB using the inbuilt function `normcdf`, which calculates the cumulative distribution function of the normal distribution, and is defined by the expression [66]

$$\texttt{normcdf}(a, \bar{x}, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{a} \exp\left(-\frac{(x-\bar{x})^2}{2\sigma^2}\right) \, \mathrm{d}x. \tag{4.47}$$

Thus substituting into (4.43) we have

$$\begin{aligned}\Delta V_{ij} = K_1 \Big(&\texttt{normcdf}(t_j, J_{12}, 1) - \texttt{normcdf}(t_{j-1}, J_{12}, 1)\Big) \\ &\Big(\texttt{normcdf}(s_j, J_{11}, 1) - \texttt{normcdf}(s_{j-1}, J_{11}, 1)\Big).\end{aligned} \tag{4.48}$$

We can use an analogous approach to calculate the perfusion, Q_{ij}. Calculating the integrals using this approach is approximately 3,000 times faster than calculating (4.40) and (4.41) directly. During parameter estimation, when these integrals are performed repeatedly, this reduction in computation time becomes significant.

4.2.2 Discretisation of the deadspace

The total volume of personal deadspace associated with each of the $N_s N_t$ unique pairs of volume change-volume and perfusion-volume ratios is $V_{ij} V_{DP}$. However, units with the same volume change-volume and perfusion-volume ratios may be located in different regions of the lung. Thus they will be associated with different volumes of deadspace. We assume that rather than forming a single compartment, this deadspace is distributed according to a personal deadspace distribution.

This distribution is discretised into N_d compartments for each pair of volume change-volume and perfusion-volume ratios; thus the total number of both alveolar and deadspace compartments, N, is $N_s N_t N_d$. The deadspace distribution is assumed to be identical for each pair of volume change-volume and perfusion-volume ratios.

In the case of the governing distribution of the alveolar compartments, presented in Section 4.1, we had both theoretical and experimental evidence to suggest that the appropriate distribution was lognormal. In the case of the deadspace distribution, we lack such evidence. Thus here we define both a lognormal and a normal distribution, and compare the two distributions in Section 4.4.4.

4.2.2.1 Lognormal distribution

To set up a lognormal distribution of deadspace we first define,

$$z = \log\left(\frac{v_{ij}}{V_{ij}}\right), \tag{4.49}$$

where v_{ij} is the fractional personal deadspace volume. The distribution of fractional alveolar volume is then

$$V_{ij} = \frac{v_{ij}}{\exp(z)}. \tag{4.50}$$

Assuming $\frac{v_{ij}}{V_{ij}}$ is lognormally distributed we have

$$v_{ij}(z) = \frac{V_{ij}}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-\left(z-\bar{z}_{ij}\right)^2}{2\sigma_z^2}\right), \tag{4.51}$$

where the factor of V_{ij} is introduced because we assume that the total volume of personal deadspace associated with each of the $N_s N_t$ unique pairs of volume change-volume and perfusion-volume ratios is $V_{ij}V_{DP}$. Substituting (4.50) into this expression we have,

$$V_{ij}(z) = \frac{V_{ij}}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-\left(z-(\bar{z}_{ij}-\sigma_z^2)\right)^2}{2\sigma_z^2}\right) \exp\left(\sigma_z^2 - \frac{2\bar{z}_{ij}}{2}\right). \tag{4.52}$$

Since the total fractional volume associated with each unique pair of ij of unique $\frac{\Delta V_{ij}}{V_{ij}}$ and $\frac{Q_{ij}}{V_{ij}}$ ratios is V_{ij}, we require that

$$\int_{-\infty}^{\infty} V_{ij}(z)\ \mathrm{d}z = V_{ij}, \tag{4.53}$$

and thus we must have

$$\exp\left(\frac{\sigma_z^2 - 2\bar{z}_{ij}}{2}\right) = 1, \tag{4.54}$$

and so

$$\bar{z}_{ij} = \frac{\sigma_z^2}{2}. \tag{4.55}$$

Thus the deadspace distribution introduces one more parameter, σ_z.

We are now going to discretise the continuous distribution, (4.51), into discrete compartments. As each of the $N_s N_t$ compartments (denoted with indices ij) is split into N_d compartments, we introduce a third index k. We again assume each alveolar compartment has equal volume, and thus the fractional volume of compartment ijk, V_{ijk} ,is

$$V_{ijk} = \int_{z_{k-1}}^{z_k} V_{ij}(z)\ \mathrm{d}z = \frac{V_{ij}}{N_d}, \tag{4.56}$$

with $z_0 = -\infty$ and $z_{N_d} = \infty$.

We can then define the fractional deadspace volume (v_{ijk}) associated with compartment ijk as

$$v_{ijk} = \int_{z_{k-1}}^{z_k} v_{ij}(z)\ \mathrm{d}z. \tag{4.57}$$

The perfusion-volume and volume change-volume ratios of the $N_s N_t$ compartments (denoted with indices ij) do not change when each of them is split into N_d compartments (as these are set by the governing alveolar distribution), and thus we have

$$Q_{ijk} = Q_{ij}\frac{V_{ijk}}{V_{ij}}. \tag{4.58}$$

Similarly the fractional volume change is

$$\Delta V_{ijk} = \Delta V_{ij} \frac{V_{ijk}}{V_{ij}}. \tag{4.59}$$

For clarity we note here that

$$\begin{aligned} i =& 1, 2, \ldots, N_s, \\ j =& 1, 2, \ldots, N_t, \\ k =& 1, 2, \ldots, N_d. \end{aligned}$$

4.2.2.2 Normal distribution

As has already been explained we will present two different deadspace distributions, having presented the lognormal deadspace distribution, we now present the normal deadspace distribution. In the case of the normal distribution we define,

$$z = \frac{v_{ij}}{V_{ij}}, \tag{4.60}$$

and thus

$$v_{ij} = zV_{ij}. \tag{4.61}$$

Assuming the fractional volume is normally distributed, we have

$$V_{ij}(z) = \frac{V_{ij}}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-\left(z - \bar{z}_{ij}\right)^2}{2\sigma_z^2}\right). \tag{4.62}$$

Substituting (4.61) into this expression we have

$$v_{ij}(z) = \frac{V_{ij} z}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-\left(z - \bar{z}_{ij}\right)^2}{2\sigma_z^2}\right). \tag{4.63}$$

The total fractional deadspace volume associated with each unique pair of $\frac{\Delta V_{ij}}{V_{ij}}$ and $\frac{Q_{ij}}{V_{ij}}$ ratios is V_{ij}. Thus we require

$$\int_{-\infty}^{\infty} v_{ij}(z)\, \mathrm{d}z = V_{ij}. \tag{4.64}$$

If we let $Z = z - \bar{z}_{ij}$ we have,

$$\int_{-\infty}^{\infty} \frac{V_{ij}\left(Z + \bar{z}_{ij}\right)}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-Z^2}{2\sigma_z^2}\right)\, \mathrm{d}Z = V_{ij}, \tag{4.65}$$

hence

$$V_{ij} = \bar{z}_{ij} V_{ij} + \int_{-\infty}^{\infty} \frac{V_{ij} Z}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-Z^2}{2\sigma_z^2}\right)\, \mathrm{d}Z, \tag{4.66}$$

and since $Z \exp\left(\frac{-Z^2}{2\sigma_z^2}\right)$ is an odd function we have,

$$\bar{z}_{ij} = 1. \tag{4.67}$$

Thus, as for the lognormal deadspace distribution, the normal deadspace distribution introduces one additional parameter, σ_z.

To allocate compartments we again apply (4.56), (4.57), (4.58), and (4.59). However, the normal deadspace distribution (4.62) is defined from $-\infty < z < \infty$. From (4.60) we see that $z < 0$ gives negative values of $\frac{v}{V}$.

Thus if the value of σ_z is sufficiently large (the larger the value of σ_z the greater the density of the distribution away from its mean of one), then (4.57) may give a deadspace compartment with a negative volume, which is clearly not physically realistic. Thus we establish a maximum value for σ_z that prevents (4.57) giving a compartment with a negative volume.

To do this we find the volume of the smallest compartment as an (implicit) function of σ_z, and find (numerically) the value of σ_z which makes this compartment's volume zero. This value is set as the maximum value for σ_z.

Substituting (4.63) into (4.57), and using $\bar{z}_{ij} = 1$ we have

$$v_{ijk} = \int_{z_{k-1}}^{z_k} \frac{V_{ij} z}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-(z-1)^2}{2\sigma_z^2}\right) \, \mathrm{d}z. \tag{4.68}$$

If we let $Z = z - 1$, then we have that

$$v_{ijk} = \int_{z_{k-1}-1}^{z_k-1} \frac{V_{ij}(Z+1)}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-Z^2}{2\sigma_z^2}\right) \, \mathrm{d}Z, \tag{4.69}$$

thus

$$v_{ijk} = \int_{z_{k-1}-1}^{z_k-1} \frac{V_{ij}}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-Z^2}{2\sigma_z^2}\right) \, \mathrm{d}Z + \int_{z_{k-1}-1}^{z_k-1} \frac{V_{ij} Z}{\sqrt{2\pi}\sigma_z} \exp\left(\frac{-Z^2}{2\sigma_z^2}\right) \, \mathrm{d}Z. \tag{4.70}$$

From (4.56) we can see the first integral is equal to $\frac{V_{ij}}{N_d}$, and integrating the second we have

$$v_{ijk} = \frac{V_{ij}}{N_d} + \frac{V_{ij}\sigma_z}{\sqrt{2\pi}} \left(\exp\left(\frac{-(z_{k-1}-1)^2}{2\sigma_z^2}\right) - \exp\left(\frac{-(z_k-1)^2}{2\sigma_z^2}\right)\right). \tag{4.71}$$

This will have its lowest value when $k = 1$, as $z_{k-1} = z_0 = -\infty$ and thus the positive exponential term in this expression will be zero. In this case,

$$v_{ij1} = \frac{V_{ij}}{N_d} + \frac{V_{ij}\sigma_z}{\sqrt{2\pi}} \left(- \exp\left(\frac{-(z_1-1)^2}{2\sigma_z^2}\right)\right). \tag{4.72}$$

Since the volume of a compartment cannot be less than zero, this expression must be greater than or equal to zero. Thus a maximum value for σ_z can be found by solving the following equation,

$$\frac{V_{ij}}{N_d} - \frac{V_{ij}\sigma_z}{\sqrt{2\pi}} \exp\left(\frac{-(z_1-1)^2}{2\sigma_z^2}\right) = 0. \tag{4.73}$$

For $N_d = 5$ this gives a maximum value of σ_z of 0.7144.

4.2.3 Note on indexing of compartments

We began this section referring to compartments with two indices, ij, due to our discretisation of the bivariate lognormal distribution. We then introduced a third index, k, as each one of these ij compartments was split into N_d compartments due to the discretisation of the deadspace distribution. Thus each compartment had three indices, ijk.

This indexing was convenient for explaining the discretisation of the distributions. However, for simplicity of notation, in the rest of the thesis we uniquely identify each compartment with a single index $\hat{i}$, where

$$\hat{i} = i + N_s\,(j-1) + N_s N_t\,(k-1)\,, \tag{4.74}$$

where

$$\begin{aligned} i =& 1, 2, \ldots, N_s, \\ j =& 1, 2, \ldots, N_t, \\ k =& 1, 2, \ldots, N_d. \end{aligned}$$

In the rest of the thesis we will drop the $\hat{}$ and i (given by (4.74)) will be the sole index we use to refer to the compartments.

Having explained how we define the volumes of our model lung we now describe the governing equations.

4.3 Governing equations for the alveolar compartments

In this section we describe the governing equations for the alveolar compartments in our model. For a more detailed discussion of how they are derived from the principle of conservation of mass see Section 2.2.

We assume that the alveolar compartments are perfectly mixed, and that the alveolar and arterial gas partial pressures are in equilibrium. Since N_2 and O_2 are almost insoluble in lung tissue we assume that the volume of these gases dissolved in lung tissue will be negligible. By conservation of mass the rate of change of the volume of gas in a lung unit is equal to the sum of the flux entering via ventilation from the personal deadspace, and the net flux due to exchange with the blood across the alveolar membrane. Thus for N_2 and O_2 during inspiration we have [35, 42],

$$\frac{d\left(F_i^g(t)V_{Ai}(t)\right)}{dt} = \dot{V}_i^{in} u(t) F_{Ii}^g(t) + Q_T Q_i\left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right), \tag{4.75}$$

where $F_i^g(t)$ is the gas fraction g (where g may be O_2 or N_2) in compartment i at time t, $\dot{V}_i^{in}$ is the fraction of the flow measured at the mouth that enters compartment i during

inspiration (the calculation of this quantity is discussed in Section 4.3.1), $F_{Ii}^{g}(t)$ is the fraction of gas g inspired into compartment i from the connected personal deadspace compartment, Q_T is the cardiac output, $C_{\bar{v}}^{g}(t)$ is the venous gas concentration of gas g, $C_{ai}^{g}(t)$ the arterial gas concentration of gas g leaving compartment i, $V_{Ai}(t)$ the volume of compartment i, and $u(t)$ is the flow rate of gas (measured by the LGA). It should be noted that since we consider gas exchange with the blood, $\frac{dV_{Ai}(t)}{dt}$ is not equal to $\dot{V}_i^{in}u(t)$, as is often assumed [14], but is given by (4.79), which is derived below.

However, CO_2 is significantly more soluble than N_2 or O_2, and thus is present in appreciable volumes in lung tissue [100]. We assume that the tissue volume is a constant fraction (b) of the total alveolar volume at FRC (V_A) and distributed among the compartments in proportion to their volume at FRC. We assume that the partial pressures of CO_2 in the tissue and the alveoli are at equilibrium, and thus the volume of CO_2 dissolved in the tissue associated with compartment i is

$$V_i V_A b \lambda_t \left(P_b - P_{H_2O}\right) F_i^{CO_2}(t), \tag{4.76}$$

where P_b is barometric pressure, P_{H_2O} is the saturated partial pressure of water at body temperature and pressure (BTPS), λ_t is the solubility of CO_2 in lung tissue. We then define the effective CO_2 tissue volume, V_t, as

$$V_t = b\lambda_t \left(P_b - P_{H_2O}\right). \tag{4.77}$$

Therefore, including this volume, for CO_2 on inspiration, by analogy with (4.75), we have

$$\frac{d\left(F_i^{CO_2}(t)\left(V_{Ai}(t) + V_i V_t V_A\right)\right)}{dt} = \dot{V}_i^{in}u(t)F_{Ii}^{CO_2}(t) + Q_T Q_i \left(C_{\bar{v}}^{CO_2}(t) - C_{ai}^{CO_2}(t)\right). \tag{4.78}$$

Summing (4.75) (for both N_2 and O_2) and (4.78), re-arranging, and remembering that $F^{CO_2} + F^{N_2} + F^{O_2} = 1$, we have

$$\frac{dV_{Ai}(t)}{dt} = \dot{V}_i^{in}u(t) + Q_T Q_i \sum_g \left(C_{\bar{v}}^{g}(t) - C_{ai}^{g}(t)\right) - V_i V_t V_A \frac{dF_i^{CO_2}(t)}{dt}, \tag{4.79}$$

where g denotes one of the three respiratory gases. Thus as we noted earlier $\frac{dV_{Ai}(t)}{dt}$ is not equal to $\dot{V}_i^{in}u(t)$.

On expiration the flux due to ventilation is leaving rather than entering the compartment. Thus (4.75) becomes [35, 42]

$$\frac{d\left(F_i^{g}(t)V_{Ai}(t)\right)}{dt} = \dot{V}_i^{ex}u(t)F_i^{g}(t) + Q_T Q_i \left(C_{\bar{v}}^{g}(t) - C_{ai}^{g}(t)\right), \tag{4.80}$$

where $\dot{V}_i^{ex}$ is the fraction of the flow measured at the mouth that leaves compartment i during expiration (the calculation of this quantity is discussed in Section 4.3.1), and g is O_2 or N_2. Similarly for CO_2 we have

$$\frac{d\left(F_i^{CO_2}(t)\left(V_{Ai}(t) + V_i V_t V_A\right)\right)}{dt} = \dot{V}_i^{ex}u(t)F_i^{CO_2}(t) + Q_T Q_i \left(C_{\bar{v}}^{CO_2}(t) - C_{ai}^{CO_2}(t)\right). \tag{4.81}$$

Summing (4.80) (for both N_2 and O_2) and (4.81), and re-arranging, gives the governing equation for compartmental volumes on expiration,

$$\frac{dV_{Ai}(t)}{dt} = \dot{V}_i^{ex} u(t) + Q_T Q_i \sum_g (C_{\bar{v}}^g(t) - C_{ai}^g(t)) + V_i V_t V_A \frac{dF_i^{CO_2}(t)}{dt}. \tag{4.82}$$

4.3.1 Fractional ventilation

In this section we first explain why we assume that fractional volume change of compartment i, ΔV_i, rather than the fractional ventilation of that compartment, $\dot{V}_i$, is constant. We then describe the practical approach we developed to establish the fractional ventilations during inspiration, $\dot{V}_i^{in}$, and expiration, $\dot{V}_i^{ex}$.

Consider a breath which starts (at $t = 0$) and finishes (at $t = t_b$) at the same total lung volume. This implies that

$$\int_0^{t_b} \frac{dV_A(t)}{dt} \, \mathrm{d}t = 0. \tag{4.83}$$

If we assume that $\dot{V}_i$ is constant and equal for both expiration and inspiration, sum (4.82) over all compartments, and integrate we have

$$\int_0^{t_b} u(t) + \sum_i \left(Q_T Q_i \sum_g (C_{\bar{v}}^g(t) - C_{ai}^g(t)) - V_i V_t V_A \frac{dF_i^{CO_2}(t)}{dt} \right) \, \mathrm{d}t = 0. \tag{4.84}$$

If the total lung volume is the same at the start and end of this breath, then we also require that the volume of each compartment is the same at the start and the end of the breath. If we assume that the fractional ventilation to compartment i is constant and equal to $\dot{V}_i$, then integrating (4.82) we have

$$\int_0^{t_b} \frac{dV_{Ai}(t)}{dt} \, \mathrm{d}t = \int_0^{t_b} \dot{V}_i u(t) + Q_T Q_i \sum_g (C_{\bar{v}}^g(t) - C_{ai}^g(t)) - V_i V_t V_A \frac{dF_i^{CO_2}(t)}{dt} \, \mathrm{d}t. \tag{4.85}$$

This should be equal to zero if the compartment's volume is the same at the start and the end of the breath. However, in general, it will not be. In fact, (4.84), which holds because the total lung volume is the same at the start and end of the breath, only implies that (4.85) is equal to zero if $\dot{V}_i = Q_i = V_i$ (i.e. if we have perfect ventilation perfusion matching). This is not true in every compartment of our model, and thus the fractional ventilation to each compartment cannot be constant over the whole breath. Instead, we define (from the lognormal distribution) a constant fractional volume change ΔV_i, and calculate the ventilations of each compartment on inspiration and expiration from this. We discuss the exact definition of the fractional volume change, and how we estimate the fractional ventilations in each compartment below.

The most natural definition of ΔV_i is

$$\frac{dV_{Ai}(t)}{dt} = \Delta V_i \frac{dV_A(t)}{dt}, \tag{4.86}$$

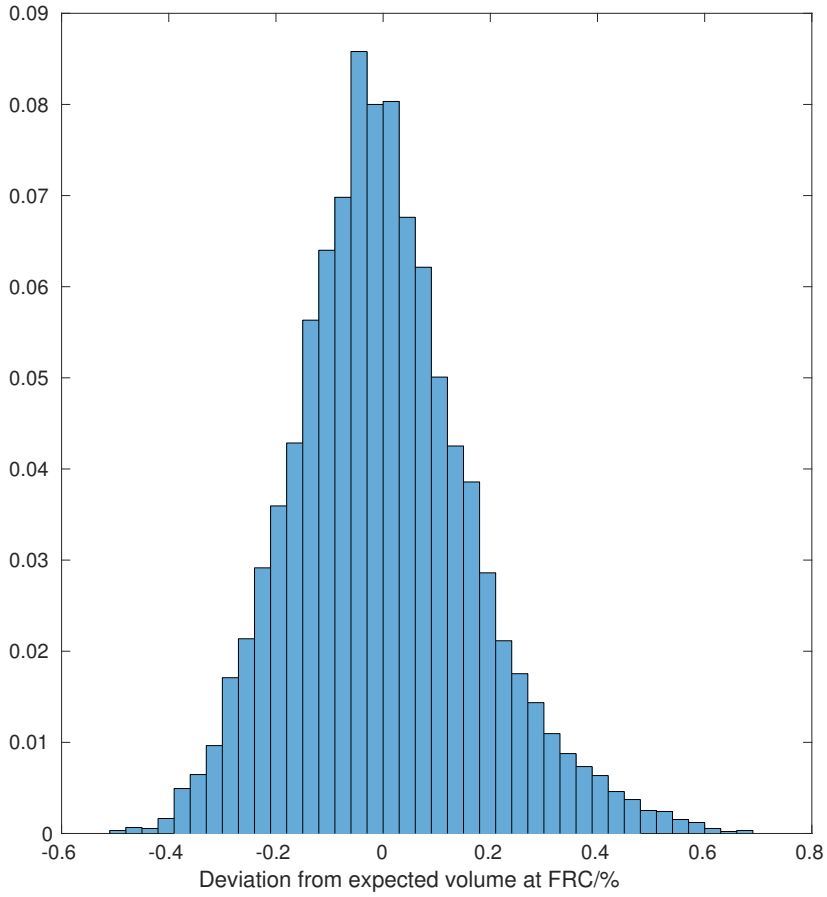

Figure 4.6: Histogram of the deviations of compartments' volumes from $V_A V_i$ when the total alveolar volume is within 0.5 ml of V_A during 13 minutes of breathing. The data used to generate this figure is taken from one of the parameter fits to experimental data discussed in Chapter 6. The histogram is normalised so the sum of the heights of the bars is one.

where $V_A(t)$ is the total alveolar volume at t. Substituting (4.79) into (4.86) we have

$$\dot{V}_i = \frac{1}{u(t)} \left(\Delta V_i \frac{dV_A}{dt} - Q_T Q_i \sum_g (C_{\bar{v}}^g(t) - C_{ai}^g(t_t)) + V_i V_t V_A \frac{dF_i^{\mathrm{CO}_2}(t)}{dt} \right). \tag{4.87}$$

However, if (4.87) applies, then when $u(t) = 0$, $\dot{V}_i$ is undefined. The physical reason for this is that when the flow is zero the only cause of alveolar volume change is gas exchange with the blood, and thus there is no value of the ventilation that can ensure the fractional volume change remains constant. Thus (4.86) cannot hold.

Instead of requiring a completely constant value of the ratio $\frac{dV_{Ai}(t)}{dt} / \frac{dV_A(t)}{dt}$ we take a more approximate approach. We assume constant values for $\dot{V}_i$ during both expiration and inspiration, although the values for expiration and inspiration will be different, and are updated on a breath by breath basis as explained below.

We initialise the values for $\dot{V}_i$ by setting them equal to the values of ΔV_i, since the volume change during a single breath due to ventilation is much larger than that due to gas exchange with the blood. On subsequent breaths, the value of $\dot{V}_i$ during inspiration ($\dot{V}_i^{in}$) is the value that would have led the following integral to hold on the previous breath,

$$\int_0^{t_I} \frac{dV_{Ai}(t)}{dt}\ \mathrm{d}t = \int_0^{t_I} \Delta V_i \frac{dV_A(t)}{dt}\ \mathrm{d}t, \tag{4.88}$$

where t_I is the duration of that inspiration. Substituting in (4.79), integrating, and re-arranging we have

$$\dot{V}_i^{in} = \frac{\Delta V_i \left(V_A(t_I) - V_A(0)\right) - \int_0^{t_I} Q_T Q_i \sum_g \left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right) + V_i V_t V_A \frac{dF_i^{\mathrm{CO}_2}(t)}{dt}\ dt}{\int_0^{t_I} u(t)\ dt}. \tag{4.89}$$

Analogously on expiration we have

$$\dot{V}_i^{ex} = \frac{\Delta V_i \left(V_A(t_b) - V_A(t_I)\right) - \int_{t_I}^{t_b} Q_T Q_i \sum_g \left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right) + V_i V_t V_A \frac{dF_i^{\mathrm{CO}_2}(t)}{dt}\ dt}{\int_{t_I}^{t_b} u(t)\ dt}. \tag{4.90}$$

The values of these expression are estimated at the end of each breath and the calculated values of $\dot{V}_i^{in}$ and $\dot{V}_i^{ex}$ used on the subsequent inspiration/expiration.

To verify this approach we checked that when the lung was at FRC, the volume of each compartment was $V_A V_i$. Deviations from this were typically less than 1% of the compartment's volume, and no trends in the deviations over the nitrogen washout procedures we simulated (up to 13 minutes of breathing) were observed. Figure 4.6 shows a histogram of the deviations of the compartments' volumes from their expected value ($V_A V_i$) when the lung was at FRC. For this figure FRC was defined as when the total volume of the alveolar compartments was within 0.5 ml of V_A. The data used to generate this figure is taken from one of the parameter fits to experimental data discussed in Chapter 6.

4.3.2 Governing equations for the body compartments

In Section 4.3, we derived the governing equations for the alveolar compartments, and saw that they depend on the venous gas concentrations $C_{\bar{v}}^g(t)$. These are not measured for subjects breathing on the LGA, and thus we must estimate them.

Venous CO_2 and O_2 concentrations (the latter only during air breathing) are estimated from experimental data, using estimates of the gas exchange on each breath. This is described in Section 5.2.2.3. However, N_2 is very insoluble in blood, and hence the volume of this gas exchanged with the blood during normal breathing is almost zero.

Thus to model the body's N_2 stores, and their exchange with the blood, we adopt the model (and parameters) of Baker and Farmery [7]. This model, illustrated in Figure 4.7, assumes the body is made up of four compartments perfused in parallel. The venous

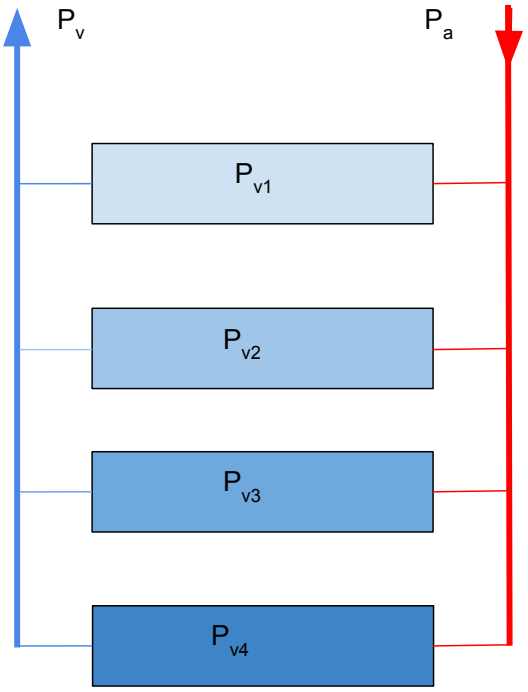

Figure 4.7: A schematic diagram of the four compartment model of the recirculation of arterial nitrogen leaving the lungs through the body[7]. Arrows show the arterial blood leaving the lungs and mixed venous blood returning to the lungs.

partial pressure of nitrogen leaving the l-th body compartment is found by solving the following conservation of mass equation ((21) from [7]),

$$\frac{dP_{vl}^{\mathrm{N}_2}}{dt} = \frac{Q_T Q_l}{\lambda_l V_{tl}} \left(P_a^{\mathrm{N}_2}(t) - P_{vl}^{\mathrm{N}_2}(t) \right), \tag{4.91}$$

where $P_{vl}^{\mathrm{N}_2}$ is the venous (which is assumed to be equal to the tissue) nitrogen partial pressure in body compartment l, Q_l is the fractional perfusion of body compartment l, λ_l is the blood-tissue partition coefficient of compartment l, V_{tl} is the tissue volume of compartment l, and $P_a^{\mathrm{N}_2}(t)$ is the perfusion weighted average arterial partial pressure of N_2 in each alveolar compartment. $P_a^{\mathrm{N}_2}(t)$ is given by

$$P_a^{\mathrm{N}_2}(t) = \sum_i Q_i F_i^{\mathrm{N}_2}(t) \left(P_b - P_{\mathrm{H}_2\mathrm{O}} \right). \tag{4.92}$$

The mixed venous nitrogen concentration entering the lung is the perfusion weighted concentration leaving the body compartments,

$$C_{\bar{v}}^{\mathrm{N}_2}(t) = S_b \sum_l Q_l P_{vl}^{\mathrm{N}_2}(t), \tag{4.93}$$

where S_b is the solubility of N_2 in the blood.

As mentioned above, we estimate the venous concentrations of CO_2 and O_2 during air breathing from experimental data. As we will discuss in Section 5.2.2, we cannot

continue to apply this method during the high O_2 breathing period (the nitrogen washout procedures we simulate consists of a ten minute period of breathing normal air followed by five minutes of breathing pure O_2). We assume that the behaviour of venous CO_2 is the same in both the air breathing and high O_2 breathing phase, as the inspired CO_2 concentration does not change.

However, the arterial O_2 concentration during the high O_2 breathing period rises, and we need a model of how this O_2 in the arterial blood will recirculate through the body, and into the venous blood. If we ignore this recirculation, then, as shown in Section 5.2.3, we simulate a very large O_2 exchange with the blood, and the alveolar volume shows a sharp decline during the O_2 breathing phase.

Unlike N_2, O_2 is consumed by the body, and thus if we were to adopt the four compartment model for O_2, we would have to determine how this consumption was distributed among the compartments. We lack information with which to do this, and thus for the O_2 stores we assume a single body compartment. Further, we assume there is a constant rate of O_2 consumption, $\dot{V}_{O_2}$, which is the same as that we measure at the mouth during the air breathing period. Thus the governing differential equation for $C_{\bar{v}}^{O_2}$ in the high O_2 breathing period is

$$\frac{dC_{\bar{v}}^{O_2}}{dt} = \frac{1}{V_L}\left(Q_T\left(C_a^{O_2}(t) - C_{\bar{v}}^{O_2}(t)\right) - \dot{V}_{O_2}\right), \tag{4.94}$$

where V_L is the volume of the compartment and $C_a^{O_2}$ is the perfusion weighted average arterial concentration. $C_a^{O_2}$ is given by

$$C_a^{O_2}(t) = \sum_i Q_i C_{ai}^{O_2}(t). \tag{4.95}$$

The O_2 consumption, $\dot{V}_{O_2}$, is assumed to be equal to that measured at the mouth during the air breathing period, and thus

$$\dot{V}_{O_2} = \frac{1}{t_{air}}\int_0^{t_{air}} F_M^{O_2}(t)u(t)\ \mathrm{d}t, \tag{4.96}$$

where $F_M^{O_2}$ is the O_2 fraction measured at the mouth and t_{air} is the duration of the air breathing period. We now discuss the governing equation of the deadspace compartments.

4.4 Governing equations for the deadspace compartments

In Chapter 2 we discussed both theoretical, see Section 2.2.2, and experimental, see Section 2.5.1.1, evidence that gas transport via diffusion in the conducting airways (which we approximate with deadspace compartments) was negligible compared to that by convection. In this section we begin by deriving an advection equation for the deadspace compartments, which we assume are cylinders of uniform cross sectional area.

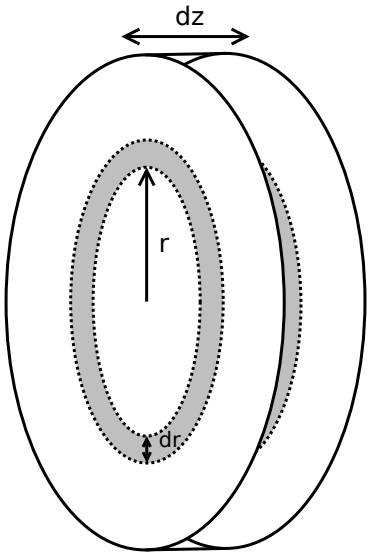

Figure 4.8: An illustrative diagram of a section of a cyclindrical deadspace volume. The shaded volume of width dr and axial length dz is used in the derivation of (4.100)

However, even in a simplified geometry such as a cylindrical pipe, many different patterns of flow are observed [106]. We consider two of these flow patterns, plug flow, representative of turbulent flow seen at high speeds, and Poiseuille Flow, a fully developed laminar flow.

Rather than trying to determine the appropriate flow pattern from theoretical analysis, we implement both the plug and Poiseuille flows in our model, and compare the results to experimental data. We find that plug flow model produces a qualitatively better fit to the data (although both models show reasonable agreement). Finally, we find that this fit can be improved by introducing the distribution of deadspace we outlined in Section 4.2.2. We briefly compare the lognormal and normal deadspace distributions, and discuss why we adopt the normal distribution.

In both this section and Section 4.5, we refer to the flow rate in all the deadspace compartments as $u(t)$, the flow rate measured at the mouth. This will be the flow rate in the common deadspace volume. However, by conservation of flow, in the personal deadspace compartments the flow rate is the same as that entering/leaving the alveolar compartments ($\dot{V}_i^{in}u(t)$ and $\dot{V}_i^{ex}u(t)$ during inspiration and expiration respectively). We use this slight abuse of notation for simplicity, and note that it does not affect the governing equations or their solution.

4.4.1 The advection equation

In this section we derive the governing advection equation for a deadspace compartment. We assume the deadspace compartment is a straight cylindrical pipe of radius R and length L. We also assume that we can neglect gravity, and that the flow is: incompressible;

independent of z; axial; and independent of the polar co-ordinate θ (i.e. we assume axial symmetry).

Consider a region of the pipe at a radius r from the centre, radial width dr and axial width dz (the shaded volume in Figure 4.8). The volume of gas g in that region at time t will be

$$2\pi r dr dz F^g(r, z, t), \tag{4.97}$$

where F^g is the fraction of gas g. Since we ignore diffusion, the net volume of gas g entering that region at time t is,

$$2\pi r dr v(r, t) \left(F^g(r, z, t) - F^g(r, z + dz, t)\right), \tag{4.98}$$

where $v(r, t)$ is the gas velocity. By conservation of mass the change of volume of gas in this shaded region from t to $t + dt$ is equal to the net volume of gas entering that region over the period and thus

$$dz \left(F^g(r, z, t) - F^g(r, z, t + dt)\right) = dt v(r, t) \left(F^g(r, z, t) - F^g(r, z + dz, t)\right), \tag{4.99}$$

where we have cancelled $2\pi r dr$ which appears in every term. Dividing by $dzdt$ and taking limits as they tend to zero, we have the governing advection equation for the deadspace compartments,

$$\frac{\partial F^g(r, z, t)}{\partial t} = -v(r, t) \frac{\partial F^g(r, z, t)}{\partial z}. \tag{4.100}$$

In Section 4.5 we present a method for solving this equation based on the method of characteristics [76]. We leave a full discussion of that method until there, but here we note that it is fast (approximately 5,000 times faster than real time for the problems we consider), and conserves mass to almost machine precision.

When we include these deadspace volumes in our full model, we are interested in the flow weighted average gas fractions leaving the deadspace compartments, rather than the fraction at each point along the radius. This is because these are the gas fractions that we will measure at the mouth or that will enter the alveolar compartments. The flow weighted average gas fraction, $F^g(z, t)$, is given by

$$F^g(z, t) = \frac{\int_0^R 2\pi r v(z, t) F^g(r, z, t)\ dr}{u(t)}, \tag{4.101}$$

where $u(t)$ is the flow rate (measured in units of volume per unit time), and is given by

$$u(t) = \int_0^R 2\pi r v(r, t)\ dr. \tag{4.102}$$

We measure $u(t)$ (in units of litres per second) at the mouth, and not the gas velocity, $v(r, t)$. Below we present two different flow patterns that are observed in cylindrical pipes, plug flow and Poiseuille flow, that give different forms of $v(r, t)$. Later we compare the solution of our model with these two flow patterns to experimental data, and find that plug flow gives better agreement.

4.4.2 Plug flow

Empirically it is found that for turbulent flow in a cylindrical pipe the variation of $v(r,t)$ with r is approximately given by [106]

$$v(r,t) \propto (R-r)^{\frac{1}{m}}, \tag{4.103}$$

where m is a constant. At high flow rates, m typically becomes very large [106]. In the case of plug flow we assume that $m=\infty$, and so the flow profile (and thus F^g) is uniform across the pipe. Hence from (4.102),

$$v(r,t) = \frac{u(t)}{\pi R^2}. \tag{4.104}$$

Substituting this into (4.100), we have

$$\frac{\partial F^g(z,t)}{\partial t} = -\frac{u(t)}{A}\frac{\partial F^g(z,t)}{\partial z}, \tag{4.105}$$

where $A=\pi R^2$, and we have used that $F^g(z,t) = F^g(r,z,t)$ for plug flow, as the gas fraction is uniform across the pipe. Thus to solve for $F^g(z,t)$ in the deadspace compartment when we assume plug flow, we solve this equation once per deadspace compartment using the method described in Section 4.5. Before considering Poiseuille flow, we briefly give a change of variables for this equation that we will use when deriving our method of solution in that section.

Since the deadspace has constant cross-sectional area, we can define

$$V_D = Az, \tag{4.106}$$

where $V_D(z)$ is the cumulative volume of deadspace at a distance z into the deadspace volume. We can then re-write (4.105) as

$$\frac{\partial F^g(V_D,t)}{\partial t} = -u(t)\frac{\partial F^g(V_D,t)}{\partial V_D}. \tag{4.107}$$

We note here that this depends only on V_D and $u(t)$ (which we measure), and not the dimensions of the cylinder, R and L.

4.4.3 Poiseuille flow

If we assume that the flow is laminar and fully developed, then applying no slip boundary conditions, an exact solution to the Navier Stokes equations for the axial velocity of the flow v can be derived. The velocity is given by

$$v(r,t) = \frac{2u(t)}{\pi R^4}(R^2-r^2), \tag{4.108}$$

which is the famous Hagen-Poiseuille flow equation [130]. Poiseuille flow is often assumed in biophysical models of the lung [18, 84].

If we assume a Poiseuille flow pattern, the value of the gas velocity is different at every value of r. Thus to calculate the value of $F^g(z,t)$, we have to solve (4.100) at a number of points along the radius and then apply (4.101). As we want our methods of solution to conserve mass with a very high degree of precision, this would involve solving (4.100) at a very large number of points, which would be computationally expensive.

Instead, we discretise the radius into $N_r + 1$ points, denoted r_l, and assume that the value of $v(r,t)$ is constant between any two adjacent points. We select r_l such that each annulus between two adjacent points (the shaded area in Figure 4.8) represents an equal fraction of the cross-sectional area of the cylinder. Thus

$$\pi \left(r_{l+1}^2 - r_l^2 \right) = \frac{\pi R^2}{N_r}, \tag{4.109}$$

with

$$r_1 = 0. \tag{4.110}$$

A scheme in which the values of r_l were chosen such that the flow rather than the cross sectional area was equally divided was also tried, but this made no appreciable difference to the solution. The gas velocity between r_l and r_{l+1}, v_l, is the average value of the velocity between those points, thus

$$v_l = \frac{\int_{r_l}^{r_{l+1}} v(r,t) 2\pi r \ \mathrm{d}r}{\int_{r_l}^{r_{l+1}} 2\pi r \ \mathrm{d}r}. \tag{4.111}$$

Substituting (4.108) into this expression we have

$$v_l = \frac{4u(t) \int_{r_l}^{r_{l+1}} r(R^2 - r^2) \ \mathrm{d}r}{\pi R^4 (r_{l+1}^2 - r_l^2)}, \tag{4.112}$$

and thus

$$v_l = \frac{2u(t)}{\pi R^4} \left(R^2 - \frac{r_{l+1}^2 + r_l^2}{2} \right). \tag{4.113}$$

Thus (4.101) becomes,

$$F^g(z,t) = \frac{\sum_l \pi \left(r_{l+1}^2 - r_l^2 \right) v_l F^g(r_l, z, t)}{u(t)}. \tag{4.114}$$

Hence to find $F^g(z,t)$ we solve (4.100) for $l = 1 \ldots N_r$, and find the flow weighted average of the results. This is done using the same approach, based on the method of characteristics [76], used for solving the plug flow pattern. Thus the computation time for solving the deadspace assuming this flow pattern is approximately N_r times longer than that for the plug flow pattern. In the next section we determine an appropriate value for N_r by comparing this numerical approach to an analytical solution of the Poiseuille flow model that we derive for a simple set of initial and boundary conditions.

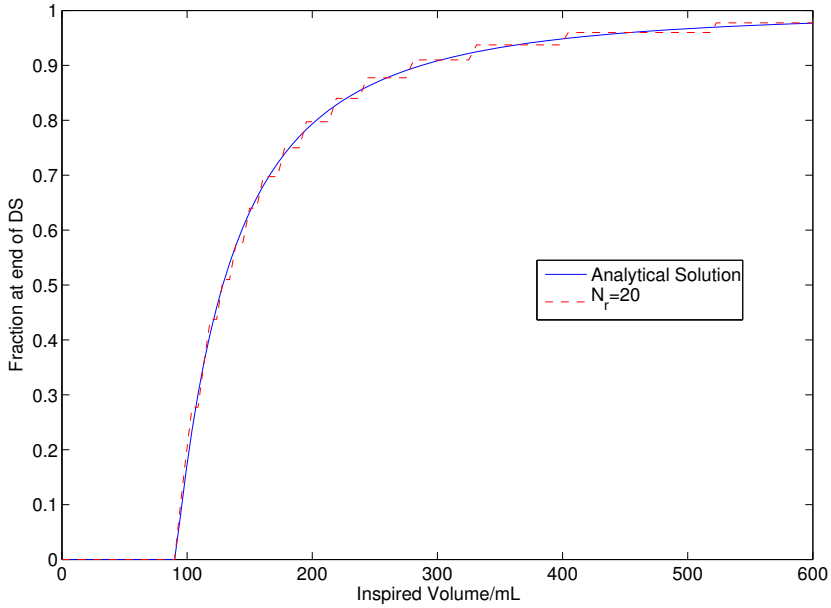

Figure 4.9: A comparison between the analytical solution at the end of the deadspace calculated using (4.125) and a numerical solution found using the compartmental approach for a single inspiration with parameters and breathing patterns from Section 2.2.2. The discrete steps in the numerical solution are due to the finite number of compartments used. In this diagram $F_0^g = 0$ and $F_I^g(t) = 1$.

4.4.3.1 An analytical solution

Consider a deadspace volume of radius R, and length L (thus total volume $V_{DT} = \pi R^2 L$), with conditions such that Poiseuille flow is established. We assume $F^g(r, z, 0) = F_0^g$ and $F(r, 0, t) = F_I^g$, i.e. the initial gas fraction in the volume is F_0^g, then gas at F_I^g is washed in at $z = 0$. As we ignore diffusion, the gas fraction as a function of the radius at the end of the deadspace will be

$$F^g(r, L, t) = \begin{cases} F_0^g & \text{for } t \leq t_L(r) \\ F_I^g & \text{for } t > t_L(r), \end{cases} \tag{4.115}$$

where $t_L(r)$ is the time taken for a signal entering the deadspace at r at $t = 0$ to reach the other end of the deadspace. L is defined by the equation,

$$L = \int_0^{t_L(r)} v(r, t)\ dt. \tag{4.116}$$

Since $v(r, t)$ is largest at $r = 0$, the signal from the mouth will first reach the end of the deadspace at this radius. This will happen at $t_L(0)$, which substituting (4.108) into (4.116) at $r = 0$ is given by,

$$L = \int_0^{t_L(0)} \frac{2u(t)}{\pi R^2}\ dt. \tag{4.117}$$

As $V_{DT} = \pi R^2 L$ we can rewrite this expression as

$$V_{DT} = \int_0^{t_L(0)} 2u(t)\ dt. \tag{4.118}$$

We note that $t_L(0)$ depends on the total volume, and not the dimensions, of the cylinder. The flow weighted average nitrogen fraction at the end of the deadspace is then

$$F^g(L,t) = \begin{cases} F_0^g & \text{for } t \le t_L(0) \\ \int_0^{r_{max}} 2\pi r F_I^g \frac{v(r,t)}{u(t)}\ dr + \int_{r_{max}}^{R} 2\pi r F_0^g \frac{v(r,t)}{u(t)}\ dr & \text{for } t > t_L(0), \end{cases} \tag{4.119}$$

where r_{max} is the radius at time t at which the following is true,

$$L = \int_0^t v(r_{max}, t')\ dt'. \tag{4.120}$$

Substituting (4.108) into this expression we have

$$\frac{V_{DT} R^2}{2} = \left(R^2 - r_{max}^2\right) \int_0^t u(t')\ dt', \tag{4.121}$$

and hence

$$r_{max} = R\sqrt{1 - \frac{V_{DT}}{2\int_0^t u(t')\ dt'}}. \tag{4.122}$$

Substituting (4.108) into (4.119) we have, for $t > t_L(0)$,

$$F^g(L,t) = \frac{4F_I^g}{R^4} \int_0^{r_{max}} \left(R^2 - r^2\right) r\ dr + \frac{4F_0^g}{R^4} \int_{r_{max}}^{R} \left(R^2 - r^2\right) r\ dr. \tag{4.123}$$

If we perform these integrals we then have

$$F^g(L,t) = \frac{4}{R^4}\left(F_I^g - F_0^g\right)\left(\frac{R^2 r_{max}^2}{2} - \frac{r_{max}^4}{4}\right) + F_0^g. \tag{4.124}$$

Substituting (4.122) into this expression we obtain

$$F^g(L,t) = 2\left(F_I^g - F_0^g\right)\left(1 - \frac{V_{DT}}{2\int_0^t u(t')\ dt'}\right)\left(1 - \frac{1}{2}\left(1 - \frac{V_{DT}}{2\int_0^t u(t')\ dt'}\right)\right) + F_0^g. \tag{4.125}$$

We note that $F^g(L,t)$ depends on the total volume and not the dimensions of the cylinder. When we tried this procedure numerically, with a variety of more general (i.e. non uniform) initial and boundary conditions, we found that the flow weighted average fraction did not depend on R, only on the total volume of the deadspace compartment.

A comparison between this method of solution, at $N_r = 20$, and the analytical solution is shown in Figure 4.9. In this figure we see that the compartmental approach tracks the analytical solution well, although discrete steps due to the compartments are visible.

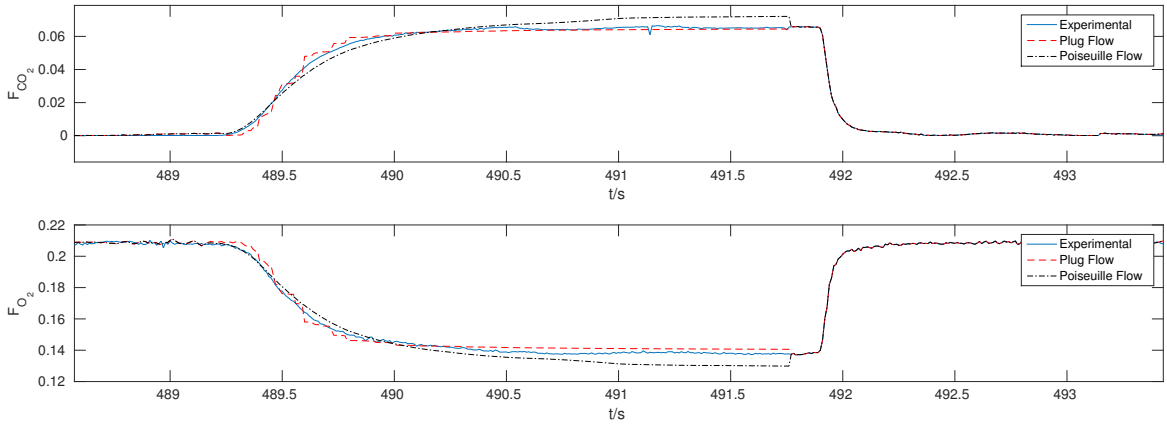

Figure 4.10: A comparison of the plug flow and Poiseuille flow patterns with experimental data for a single expiration. In this plot $N_d = 1$ (i.e. there is no deadspace distribution).

4.4.4 Deadspace in the full model

Both Poiseuille flow and plug flow patterns were implemented within the full model, with no deadspace distribution (i.e. $N_d = 1$), and fitted to experimental data taken from the LGA (using the process that will be described in Chapter 5). A comparison of the two different flow patterns for a single expiration during normal breathing is plotted in Figure 4.10. Here we consider the qualitative form of the behaviour.

For the Poiseuille flow pattern we see that the slope of phase II is too shallow (compared to that seen experimentally), but phase III is too steep. The plug flow pattern produces a phase II slope that is perhaps slightly too steep, but the alveolar plateau is similar to that observed experimentally. Thus it appears that the qualitative fit of the plug flow pattern is better, and as we noted earlier, it is faster to solve. This is a useful mathematical result, that the simplest flow pattern (often used by other authors [133]) is a valid and sufficiently accurate way to model flow in the conducting airways. Thus plug flow, rather than Poiseuille flow is the flow pattern we assume in the rest of this thesis.

4.4.4.1 Distribution of deadspace

In Figure 4.10 we considered a single personal deadspace compartment associated with each of the $N_s N_t$ unique pairs of volume change-volume and perfusion-volume ratios (i.e. $N_d = 1$). However, a further refinement can be made to the model if we consider that lung units with the same volume change-volume and perfusion-volume ratios may be located in different regions of the lung, and thus associated with different volumes of deadspace. As outlined in Section 4.2.2, we assume this can be treated by assuming the deadspace associated with each pair of volume change-volume and perfusion-volume ratios is distributed according to a continuous distribution that we approximate with N_d compartments.

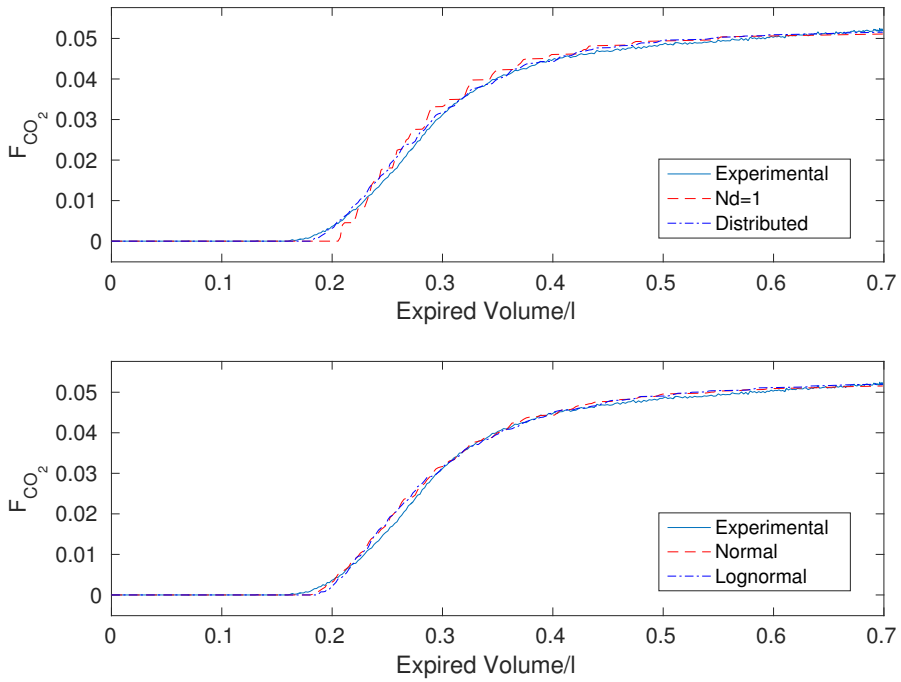

Figure 4.11: The top plot compares an experimentally measured expirogram with simulated expirograms with $N_d = 1$ (no deadspace distribution) and a normally distributed deadspace. In the bottom plot normal and lognormal deadspace distributions are compared to experimental data.

The top plot in Figure 4.11 shows the improved fit that is achieved by introducing the deadspace distribution. In this plot we assumed a normal deadspace distribution. Not only was the qualitative fit improved, but σ_z, the width of the deadspace distribution, was found to be well identified in both synthetic data (discussed in Chapter 5) and experimental data (discussed in Chapter 6). Thus this model of distributed deadspace was adopted.

In Section 4.2.2 we also described two candidates for this underlying continuous distribution of deadspace, a normal and a lognormal one. Fits assuming both of these distributions are shown in Figure 4.11. This figure shows that the output of the model is very similar independent of which distribution we assume (the estimated values of the other parameters of the model were also very similar).

This is not surprising, as at the width of the distributions shown here, $\sigma_z = 0.40$ (the mean value of σ_z we found in healthy volunteers was 0.41), the normal and lognormal distributions are very similar. Once discretised into five compartments, there is virtually no difference in the deadspace volumes that these two distributions produce.

Since we had no theoretical reason to assume a lognormal distribution, and as shown in Figure 4.11, it does not appear to make any difference to the output, we adopted the normal distribution. This will be the distribution of deadspace used in the rest of this thesis. We now describe in some detail the method we use to solve the governing equations

of the deadspace compartments.

4.5 Solving the governing equations in the deadspace

In Section 4.4 we determined that the appropriate governing equation for the deadspace compartments in our model was the advection equation,

$$\frac{\partial F^g(V_D,t)}{\partial t} = -u(t)\frac{\partial F^g(V_D,t)}{\partial V_D}, \tag{4.126}$$

where $u(t)$ is the flow rate (measured in units of unit volume per unit time) in that deadspace compartment, and $V_D(z)$ is the cumulative volume of deadspace at a distance z into the deadspace compartment.

For the purposes of this section we consider a deadspace compartment with the end closest to the mouth at $V_D = 0$, and the other end at $V_D = V_{\text{end}}$. The distribution of gas in the deadspace at $t = 0$ will be specified as the initial condition, i.e. $F^g(V_D, 0)$ is known for $0 < V_D < V_{\text{end}}$. The boundary conditions at one end of the deadspace, $F^g(0,t)$ during inspiration and $F^g(V_{\text{end}},t)$ during expiration, will also be specified. In this section we describe our method for solving (4.126).

When we solve (4.126) in the deadspace compartments we solve for two things: (i) the gas fraction leaving the deadspace (at $V_D = V_{\text{end}}$ during inspiration and $V_D = 0$ during expiration); and (ii) the distribution of gas left in the deadspace at the end of inspiration/expiration, which we use as the initial condition for the subsequent expiration/inspiration.

Our method of solution of (4.126) is based on the method of characteristics [76], which we explain in the next section. We then explain a change of variables that we apply to ensure our method conserves mass.

4.5.1 The method of characteristics

A common method for solving the advection equation is the method of characteristics [76]. In this section we briefly explain how this method can be applied to solve (4.126).

Using the chain rule we have that the total derivative of $F^g(V_D,t)$ with respect to t is

$$\frac{dF^g(V_D,t)}{dt} = \frac{\partial F^g(V_D,t)}{\partial t} + \frac{\partial F^g(V_D,t)}{\partial V_D}\frac{dV_D}{dt}, \tag{4.127}$$

and thus from (4.126) we can see that

$$\frac{dF^g(V_D,t)}{dt} = 0 \quad \text{on} \quad \frac{dV_D}{dt} = u(t), \tag{4.128}$$

i.e. F^g is constant on lines of $\frac{dV_D}{dt} = u(t)$. These lines are known as the characteristic lines.

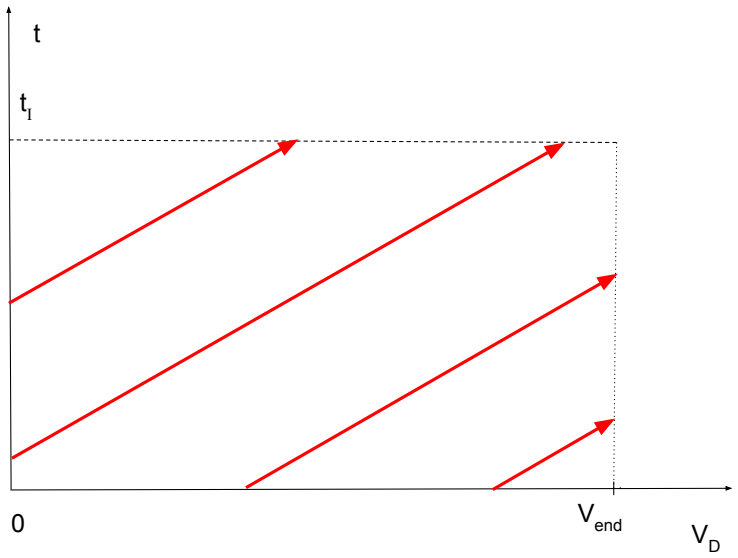

Figure 4.12: A graphical representation of the method of characteristics, as applied to a deadspace compartment during inspiration. The red arrows represent characteristic lines. In this diagram $u(t)$ is constant, and thus these lines are straight. Arrows connected to the vertical dotted line show how the solution at the end of the deadspace as a function of t is found. Arrows connected to the horizontal dashed line show how the solution in the deadspace at the end of inspiration is found.

Figure 4.12 illustrates the application of the method of characteristics to a deadspace compartment during inspiration. As discussed previously, we solve for F^g at $V_D = V_{\text{end}}$ as a function of time during inspiration, and at every point in compartment at the end of inspiration ($t = t_I$). Thus we want to know the value of F^g on the vertical dotted line and the horizontal dashed line in Figure 4.12. The values of F^g on the horizontal and vertical axes are the known initial and boundary conditions respectively.

In this figure characteristic lines are illustrated with red arrows. For ease of illustration we have drawn them as straight lines, which they will be if $u(t)$ is a constant. If $u(t)$ is not constant, then they will be parallel curves not straight lines. As discussed above, F_g is constant on the characteristic lines. Thus to find the value of F^g at a given point on the vertical dotted/horizontal dashed line, we integrate the characteristic equation, $\frac{dV_D}{dt} = u(t)$, from this point (the tip of the red arrows in the figure) to find the characteristic curve on which F^g is constant. This allows us to locate the point in the initial or boundary conditions (the base of the red arrows in the figure) where the (known) value of F^g is the same as at the point of interest. In the next section we explain how we do this in practice.

4.5.2 Implementation for inspiration

In the previous section we explained how the solution to the advection equation can be found from the initial and boundary conditions using the method of characteristics. In this section we explain how we 'extend' the initial conditions to include the boundary

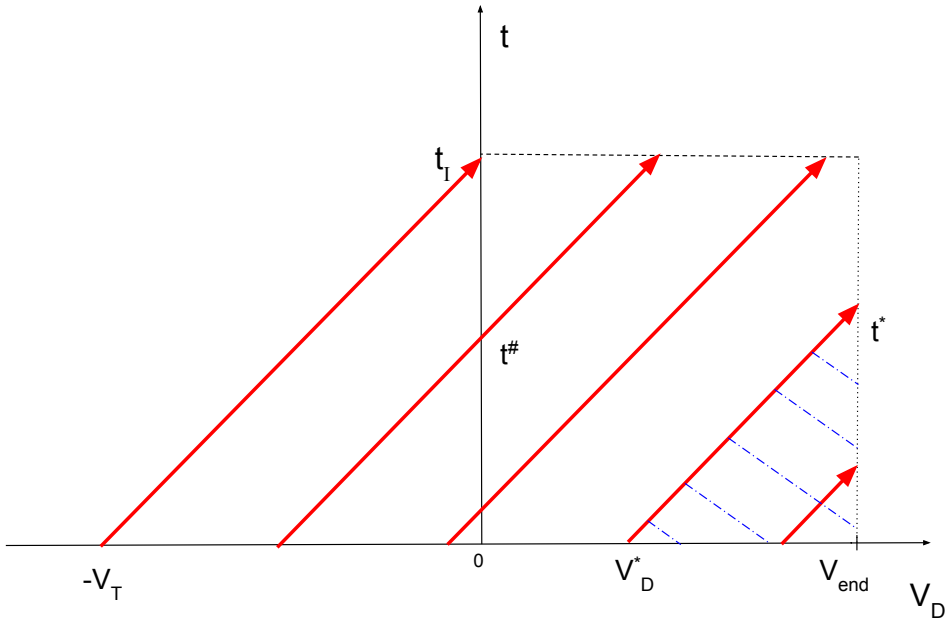

Figure 4.13: A graphical representation of our computational implementation of the method of characteristics. The red arrows represent characteristic lines. In this diagram $u(t)$ is constant, and thus these lines are straight. Arrows connected to the vertical dotted line show how the solution at the end of the deadspace as a function of t is found. Arrows connected to the horizontal dashed line show how the solution in the deadspace at the end of inspiration is found. Here we have included the boundary conditions in the initial conditions, so that the solution can be found in a computationally efficient manner.

conditions, as this simplifies the computational implementation of the calculations that have to be performed. We also explain a change of variables that ensures conservation of mass, even when solving (4.126) at a finite number of time points. We explain our method for solving the advection equation during inspiration in detail in this section, and then briefly discuss how this method is applied during expiration in the next section.

We now describe how we 'extend' the initial conditions to include the boundary conditions. Consider the characteristic curve in Figure 4.13 which intersects the vertical axis at $t = t^{\#}$. We want to find the value that V_D would have if we extended the characteristic curve back to $t = 0$. If we integrate the characteristic equation of this curve we have

$$\int_0^{t^{\#}} \frac{dV_D}{dt} \, \mathrm{d}t = \int_0^{t^{\#}} u(t) \, \mathrm{d}t, \tag{4.129}$$

and thus

$$V_D(t^{\#}) - V_D(0) = \int_0^{t^{\#}} u(t) \, \mathrm{d}t. \tag{4.130}$$

Since V_D is zero at $t = t^{\#}$, dropping the $^{\#}$, we then have

$$V_D(0) = -\int_0^{t} u(t') \, \mathrm{d}t', \tag{4.131}$$

for $t < t_I$, where t_I is the length of the inspiration. The value of $F^g(V_D, 0)$ in this region is $F_I^g(t)$, the fraction of the gas of interest inspired into the deadspace compartment at time t (i.e. the boundary condition). This 'extension' is illustrated in Figure 4.13, where characteristic lines that would have terminated on the vertical axis in Figure 4.12 are now extended so that their base is on the horizontal axis.

As mentioned above, we do not solve for the fractions of the gas of interest, but rather cumulative volumes of the gas of interest. Before describing how we do this, we briefly explain why we make this change of variables.

Consider the two red arrows connecting the horizontal axis to the vertical dotted line in Figure 4.13. These arrows represent the characteristic lines, with the value of $F^g(V_D, 0)$ at the base of the arrow being equal to the value F^g takes at every point on the line, including at the tip of the arrow. If we assume that the two arrows represent neighbouring points in a temporal discretisation, then we can see that the solution (the values of F^g at the tip of the arrows) will be insensitive to the values of F^g between the bases of the arrows. Thus the precision with which we are able to conserve mass will depend on the temporal discretisation. However, if we work with cumulative volumes rather than fractions, then the value at the base of the second arrow will contain information about all the points in between the two arrows, and thus, as we show in Section 4.5.5, we can conserve mass with very high precision. Below we describe how we do this in practice.

We define the cumulative volume of the gas of interest in the deadspace at $t = 0$, $V_{Dg}^g(V_D)$, as

$$V_{Dg}^g(V_D) = \int_{V_D}^{V_{\text{end}}} F^g(V_D', 0) \, \mathrm{d}V_D', \tag{4.132}$$

where we have defined zero of this quantity at $V_D = V_{\text{end}}$, and V_D extends to $-\int_0^{t_I} u(t)\ \mathrm{d}t$ to include the boundary conditions. Rather than solving for the gas fraction leaving the deadspace at time t, $F^g(V_{\text{end}}, t)$, we solve for the cumulative volume of the gas of interest that has left the deadspace by time t, $V_{in}^g(t)$. This is given by

$$V_{in}^g(t) = \int_0^t F^g(V_{\text{end}}, t')u(t')\ \mathrm{d}t'. \tag{4.133}$$

We now explain how V_{in}^g can be calculated from the initial conditions V_{Dg}^g. Consider the hatched region in Figure 4.13. As the characteristic lines never cross (even if they are not straight they will be parallel), the volume of gas g that has left the deadspace by $t = t^*$ (i.e. the quantity $V_{in}^g(t^*)$) will be equal to the volume of gas g that started inspiration in the region of the compartment between $V_D = V_D^*$ and $V_D = V_{\text{end}}$ (i.e. the quantity $V_{Dg}^g(V_D^*)$), and thus we have

$$V_{in}^g(t^*) = V_{Dg}^g(V_D^*). \tag{4.134}$$

By integrating the characteristic equation, $\frac{dV_D}{dt} = u(t)$, we have that

$$V_D^* = V_{\text{end}} - \int_0^{t^*} u(t)\ \mathrm{d}t, \tag{4.135}$$

and thus we have that

$$V_{in}^g(t) = V_{Dg}^g\left(V_{\text{end}} - \int_0^t u(t')\ \mathrm{d}t'\right), \tag{4.136}$$

where we have dropped the * for ease of notation. Thus $V_{in}^g(t)$ can be calculated from the initial condition, which we extended to include the boundary condition for convenience, without finding the solution at intermediate points in either time or space. When we implement this in MATLAB, we calculate vectors representing $V_{Dg}^g(V_D)$ and $\left(V_{\text{end}} - \int_0^t u(t')\ \mathrm{d}t\right)$ at the start of inspiration, and then vectorise the calculation of V_{in}^g using (4.136). As we show in Section 4.5.5 this is a very computationally efficient approach.

We can demonstrate that (4.136) is true more formally. Let Ω be the hatched region in Figure 4.13, with boundary $\partial\Omega$. Integrating (4.126) over this region we have that

$$\iint_\Omega \frac{\partial F^g(V_D, t)}{\partial t} + u(t)\frac{\partial F^g(V_D, t)}{\partial V_D}\ \mathrm{d}V_D\mathrm{d}t = 0. \tag{4.137}$$

If we apply Green's theorem in a plane [92] to the surface integral, then we have that

$$\iint_\Omega \frac{\partial F^g(V_D, t)}{\partial t} + u(t)\frac{\partial F^g(V_D, t)}{\partial V_D}\ \mathrm{d}V_D\mathrm{d}t = \oint_{\partial\Omega} u(t)F^g(V_D, t)\ \mathrm{d}t - F^g(V_D, t)\ \mathrm{d}V_D = 0. \tag{4.138}$$

The boundary $\partial\Omega$ is made up of three lines: the horizontal line from $V_D = V_D^*$ to $V_D = V_{\text{end}}$, at $t = 0$; the vertical line from $t = 0$ to $t = t_I$, at $V_D = V_{\text{end}}$; and the

characteristic curve from $t = t^*$ to $V_D = V_D^*$. Writing the line integral around this boundary out in full we then have that

$$-\int_{V_D^*}^{V_{\text{end}}} F^g(V_D, 0)\ \mathrm{d}V_D + \int_0^{t^*} u(t)F^g(V_{\text{end}}, t)\ \mathrm{d}t + \int_{t^*}^{0} u(t)F^g(V_D, t) - u(t)F^g(V_D, t)\ \mathrm{d}t = 0, \tag{4.139}$$

where we have used that $\mathrm{d}V_D = u(t)dt$ on the characteristic curve. From (4.132) and (4.133) the first two integrals are equal to $V_{Dg}^g(V_D^*)$ and $V_{in}^g(t^*)$ respectively. Since the third is zero, we have, as required,

$$V_{in}^g(t^*) = V_{Dg}^g(V_D^*). \tag{4.140}$$

As characteristic lines do not cross, the distribution of gas in the deadspace compartment at the end of inspiration, V_{DgEx}^g, will be the same as the distribution in the part of the (extended) initial condition which has not left the deadspace compartment by the end of inspiration, and thus

$$V_{DgEx}^g(V_D) = V_{Dg}^g\left(V_D - \int_0^{t_I} u(t)\ \mathrm{d}t\right) - V_{Dg}^g\left(\int_0^{t_I} u(t)\ \mathrm{d}t\right), \tag{4.141}$$

where we have subtracted the second term so that $V_{DgEx}^g(0) = 0$. This definition simplifies our calculations on expiration, when gas will be leaving the deadspace at $V_D = 0$, as we now explain.

4.5.3 Implementation for expiration

During expiration (when $u(t)$ is negative) the method of solution is similar, as shown in Figure 4.14, except that to include the boundary conditions in the initial condition we now extend V_D to values greater than V_{end} (rather than less than zero). Thus we define

$$V_D = V_{\text{end}} - \int_0^t u(t)\ \mathrm{d}t, \tag{4.142}$$

for values of V_D greater than V_{end}.

During expiration we solve for the cumulative volume of gas g that has emerged from the deadspace (at $V_D = 0$) by time t, $V_{ex}^g(t)$. This is defined by the equation

$$V_{ex}^g(t) = -\int_{t_I}^{t} F_D^g(0, t')u(t')\ \mathrm{d}t'. \tag{4.143}$$

Applying an analogous argument to that used in inspiration this is found from the cumulative volume of gas g in the deadspace (extended to include the boundary condition) at the start of expiration using

$$V_{ex}^g(t) = V_{DgEx}^g\left(-\int_{t_I}^{t} u(t')\ \mathrm{d}t'\right), \tag{4.144}$$

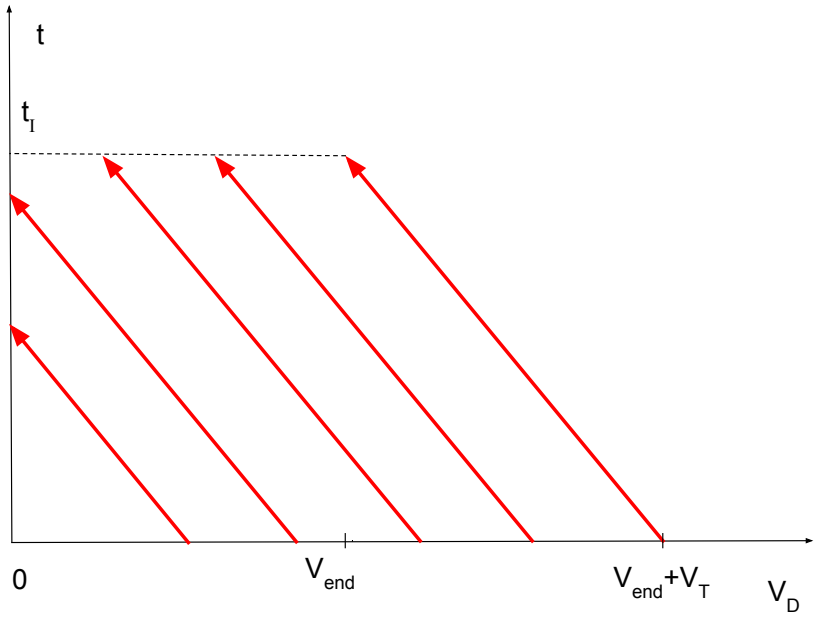

Figure 4.14: A graphical representation of our computational implementation of the method of characteristics. The red arrows represent characteristic lines. In this diagram $u(t)$ is constant, and thus these lines are straight. Arrows connected to the vertical dotted line show how the solution at the end of the deadspace as a function of t is found. Arrows connected to the horizontal dashed line show how the solution in the deadspace at the end of inspiration is found. Here we have included the boundary conditions in the initial conditions, so that the solution can be found in a computationally efficient manner.

where V^g_{DgEx} is defined such that its zero is at $V_D = 0$, i.e.

$$V^g_{DgEx}(V_D) = \int_0^{V_D} F^g(V'_D, t_I)\ \mathrm{d}V'_D. \tag{4.145}$$

The distribution of gas in the end of the deadspace at the end of expiration ($V^g_{Dg}(V_D)$) is the same as the distribution of gas in the (extended) initial condition that does not leave the deadspace compartment by the end of expiration. To match the definition of the initial conditions for inspiration, (4.132), we ensure that $V^g_{Dg}(V_D)$ is zero at $V_D = V_{\text{end}}$, and increases with decreasing V_D, and thus we have

$$V^g_{Dg}(V_D) = V^g_{DgEx}\left(V_{\text{end}} - \int_{t_I}^{t_E} u(t)\ \mathrm{d}t\right) - V^g_{DgEx}\left(V_D - \int_{t_I}^{t_E} u(t)\ \mathrm{d}t\right), \tag{4.146}$$

where t_E is the end of expiration. We assume this equation applies in the common deadspace volume. However, in the personal deadspace volume we apply a simple correction to approximate the behaviour of the branching structure of the airway tree. This is described in more detail in the next section.

4.5.4 Treatment of deadspace at the end of expiration

At the end of inspiration we assumed that the initial condition for the subsequent expiration was the distribution of gas in the deadspace at the end of the previous inspiration. At the end of expiration we assume the same in the common deadspace compartment, and thus the initial condition for the subsequent inspiration is given by (4.146). However, we treat the personal deadspace compartments differently as we now explain.

Our model assumes that deadspace is either entirely personal (associated with a single alveolar compartment), or entirely common. However, the deadspace is a dichotomously branching structure, with truly personal branches at the level of the alveoli progressively joining until they reach the common deadspace of the trachea [127].

During inspiration, where gas with a common fraction is splitting as it goes down the respiratory tree, this branching structure can be well approximated by the independent volumes. Similarly, for gas which completely transits the personal deadspace compartments during expiration, the distinction is not important (it is mixed as it enters the common deadspace anyway). However, for gas which is left in the personal deadspace compartments at the end of expiration, our assumption that each personal deadspace compartment is independent will not be a good one.

As an example of the errors caused by this assumption, consider a poorly ventilated compartment during the nitrogen washout. If the compartment is poorly ventilated then at the end of expiration the N_2 fraction in its personal deadspace compartment, that it inspires at the start of inspiration, will be relatively high. However, in reality, its inspirate should contain gas (with a lower N_2 fraction) from better ventilated compartments that progressively joined the gas from this compartment as gas moved up the branching airway

structure at the end of expiration. Thus, in this case, the model would predict too high an inspired N_2 fraction into the compartment.

To solve this problem rigorously would need a full model of the branching deadspace. However, we have neither the geometrical information, nor the computational resources (given our constraints on computation time) to do this. Instead, we assume that at the end of expiration the personal deadspace compartments 'mix' such that the gas fraction at the end of the personal deadspace compartments joining the common deadspace compartment is the same across all compartments. The fraction at the alveolar end of the personal deadspace compartment is assumed to be equal to the fraction in the connected alveolar compartment. The gradient of the gas fractions within each personal deadspace compartment is assumed to be linear.

To find the common fraction, F^g_{DP}, we apply conservation of mass. First we calculate the total volume of gas in all the personal deadspace compartments at the end of expiration, V_g. This is done using (4.146). Since we assume a linear gradient in the personal deadspace compartments after mixing, the average fraction in a personal deadspace compartment will be $\frac{F^g_{DP}+F^g_{Ai}}{2}$. Thus by conservation of mass,

$$V_g = \sum_i v_i V_{DP} \frac{F^g_{DP} + F^g_{Ai}}{2}, \tag{4.147}$$

where F^g_{Ai} is the fraction of gas g in compartment i at the end of expiration, and v_i is the fractional deadspace volume of compartment i. Re-arranging this expression we have

$$F^g_{DP} = \frac{2V_g}{V_{DP}} - \sum_i v_i F^g_{Ai}. \tag{4.148}$$

We note that this fraction will not necessarily be the same as the fraction at the end of the common deadspace.

Having found the common fraction we can apply (4.132) and define the initial condition for the subsequent inspiration in each personal deadspace compartment. In the notation used in the previous two sections this is given by

$$V^g_{Dg}(V_D) = \frac{F^g_{DP} + F^g_{Ai}}{2} \left(v_i V_{DP} - V_D\right). \tag{4.149}$$

4.5.5 Summary

To test our implementation of this method of solving the governing advection equation of the deadspace we confirmed that it does indeed conserve mass with very high precision. More specifically, we confirmed that the volume of each gas in a deadspace compartment at the end of inspiration/expiration was equal to the volume of gas that was in that compartment at the start of inspiration/expiration, plus the net volume of gas that entered the compartment during the period. This was found to be true to within 10^{-15}l of gas,

demonstrating the precision with which this method conserves mass. Solving one inspiration/expiration for one compartment takes approximately 0.0004 s, showing that this is also a computationally efficient approach.

4.6 Solving the governing equations for the alveoli

In this section we discuss the solution of the governing ODEs for the alveolar compartments that were derived in Section 4.3. In that section we derived ODEs for the alveolar fractions of CO_2, O_2, and N_2, and for the alveolar volume. We solve for the CO_2 and O_2 fractions, and the alveolar volume, and then calculate the N_2 fractions from these ($F^{N_2} = 1 - F^{O_2} - F^{CO_2}$). This means that we solve $3N+5$ ODEs, one for the fraction of CO_2 in each compartment; one for the fraction of O_2 in each compartment; one for the compartment's volume; and a total of five for the body compartments. In this section we first try solving these equations with a number of different MATLAB ODE solvers, and then explain why we developed our own method.

At the start of this chapter we stated that for us to solve the inverse problem our method of solving the forward problem had to be fast (because solving the inverse problem typically involves solving the forwards problem a large number of times) and conserve mass with high precision. Conservation of mass is very important because errors in the conservation of mass are cumulative. Thus even relatively small errors in conservation of mass on a single breath may be significant over the 15 minutes of experimentally measured breathing that we use to solve the inverse problem. Another reason for the importance of conservation of mass is that we calculate N_2 fractions using $F^{N_2} = 1 - F^{O_2} - F^{CO_2}$. At the end of the washout F^{N_2} becomes very low, and thus if we do not conserve the masses of CO_2 and O_2, F^{N_2} may become negative.

Before proceeding we discuss an important issue that we have thus far ignored, the temporal resolution of our experimental data.

4.6.1 Temporal resolution of experimental data

All the fractions, flows, and concentrations we described in Section 4.3 were instantaneous values, i.e. we wrote down differential equations that assume we know the flow at every time t. However, the fractions and flows measured at 10ms intervals at the mouth using the LGA are not of this form. They are calculated based on a number of measurements taken during each 10ms time period, and thus are average values over the 10ms period.

Thus the flow reported by the LGA at time point t_l, $\bar{u}(t_l)$, is related to the instantaneous flow $u(t)$, as follows,

$$\bar{u}(t_l) = \frac{\int_{t_{l-1}}^{t_l} u(t)\ \mathrm{d}t}{\Delta t}, \tag{4.150}$$

where $\Delta t = t_l - t_{l-1}$=10ms. Similarly the average fraction is

$$\bar{F}_M^g(t_l) = \frac{\int_{t_{l-1}}^{t_l} F_M^g(t)\ \mathrm{d}t}{\Delta t}. \tag{4.151}$$

As the method of solution for the deadspace compartments we described in Section 4.5 depended only on integrals of the flow, these integrals can be replaced with discrete sums using (4.150) without affecting that method. However, as we will see in the next section, if we do not take account of this finite resolution of $u(t)$ when solving the governing equations of the alveolar compartments, and simply use inbuilt MATLAB ODE solvers, quite unexpected behaviour can result. Thus in Section 4.6.3 we develop our own method for solving the alveoli that treats this explicitly.

4.6.2 Solving the governing equations of the alveolar compartments using inbuilt ODE solvers in MATLAB

If our model has N alveolar compartments then, as we discussed at the start of this section, solving it requires the solution of $3N$+5 ODEs. MATLAB has a number of different inbuilt functions to solve such ODEs [66]. In this sub-section we compare the different solvers, and explain why we developed our own method for solving the ODEs (outlined in Section 4.6.3).

The inbuilt solvers in MATLAB can be broadly split into two categories adaptive and non-adaptive. Non-adaptive solvers take time-steps of fixed size. In MATLAB there are five such solvers: `ode1`; `ode2`; `ode3`; `ode4`; and `ode5`, with the number corresponding to the order of the solver.

Adaptive solvers estimate the error at each time point, and vary the size of the subsequent time-step such that the estimated error is kept within a prescribed limit. This has the advantage that when the solution varies quickly the solver can take small time-steps, and when it varies more slowly it can take larger ones. However, as we shall see, this is associated with considerable computational overheads. Adaptive solvers may be further sub-divided into, stiff and non-stiff solvers. Stiff solvers are designed to deal with problems where the system has (at least) one very rapidly changing component, and (at least) one much more slowly varying one. In MATLAB these are denoted with an `s` at the end of their name.

At the start of this chapter we established that our method of solving the governing equations must be fast and conserve mass with a high degree of precision. In the rest of this section we test how the various inbuilt solvers in MATLAB perform when judged against these criteria.

AbsTol	ode23/s	ode45/s	ode113/s	ode15s/s	ode23s/s
10^{-4}	10.4	11.1	10.7	35.5	429
10^{-5}	14.5	12.7	16.5	49.1	1.07×10^3
10^{-6}	29.9	17.9	30.2	85.7	2.73×10^3
10^{-7}	70.2	35.2	68.0	226	6.87×10^3
10^{-8}	152	79.0	139	513	1.43×10^4
10^{-9}	257	167	256	965	*

Table 4.1: Computation times for a 100 breath washout simulated using the model described in Section 4.2 (with 125 compartments), using a number of different adaptive ODE solvers, over a range of error tolerances, to solve the governing equations of the alveolar compartments. Values shown here are the average of three runs. *The computation time at AbsTol=10^{-9} with ode23s was not measured, as even at the previous error tolerance, the solver took over 4 hours to do 100 breaths.

4.6.2.1 Computation time

Table 4.1 shows the computation time [2] for a 100 breath washout simulated with our model using a number of different adaptive solvers at a variety of different error tolerances. It is clear from the very long computation times that using a stiff solver (ode15s or ode23s) is not appropriate in this case. At all but the largest error tolerances, the fastest non-stiff solver is ode45, thus this is the only adaptive solver we will consider further.

Table 4.2 shows these computation times for the non-adaptive solvers, there are no error tolerances here as the timesteps are fixed. As the order of the solver increases, the number of times the derivative is evaluated at each timestep increases, and so the computation time increases. In the table we can see that the range of computation times of the non-adaptive solvers is similar to that of ode45 with AbsTol values between 10^{-5} and 10^{-8}.

Before considering how well each solver conserves mass it is worth considering how this computation time is spent. Table 4.3 summarises the results of 'profiling' our solution of these washouts. Before we discuss the contents of this table, a brief note on profiling in MATLAB.

Profiling code in MATLAB produces a report summarising the number of times all the lines within our code were run, how long was spent doing this, and the overheads associated with each function. This is an extremely useful tool, both in debugging and in speeding up code. However, the act of profiling introduces its own overheads that are included in the estimates of a function's overheads. Thus the estimates of overheads are larger than the true overheads when the code is run without profiling. Hence we do not refer to computation times in these tables, and it should be remembered that the estimates of overheads are likely to be over-estimates.

[2] All times in this section are measured on a standard PC (Dell Optiplex, Intel® Quad Core™ i5-3570 CPU @ 3.40GHz, 3.7GB RAM)

Solver	Computation time/s
`ode1`	15.1
`ode2`	24.2
`ode3`	37.9
`ode4`	48.5
`ode5`	72.8

Table 4.2: Computation times for a 100 breath washout simulated using the model described in Section 4.2 (with 125 compartments), using a number of different non-adaptive ODE solvers to solve the governing equations of the alveolar compartments. Values shown here are the average of three simulations.

Solver	% time solving ODE	% overheads	Calls of Insp. DE	Calls of Exp. DE
10^{-4}	60%	36%	7,708	7,660
10^{-5}	63%	31%	14,098	7,954
10^{-6}	74%	26%	29,992	13,450
10^{-7}	87%	22%	73,818	43,346
10^{-8}	94%	19%	191,098	115,462
10^{-9}	97%	18%	418,936	273,160
`ode1`	63%	14%	33,239	17,908
`ode2`	78%	15%	66,378	35,716
`ode3`	86%	13%	99,517	53,254
`ode4`	89%	14%	132,656	71,332
`ode5`	93%	18%	198,934	106,948

Table 4.3: Results of profiling the solution of a 100 breath washout (comprising 51,147 time-points), with 125 compartments. Overheads (expressed as a percentage of the time spent solving the ODE) include those resulting from the operation of the code-profiler as well as those from the ode solver itself.

The top six entries in Table 4.3 show the results of profiling our solution of a 100 breath washout, where we have used `ode45` at a variety of error tolerances to solve the differential equations of the alveolar compartments. As we might expect, the percentage of the total computation time spent solving the ODEs, and the number of times the functions which calculate the derivatives are called (at each solution point `ode45`, which is based on an explicit Runge-Kutta method, evaluates the derivatives 6 times [66]), increase as we make the error tolerance smaller. The second point worth noting is that, particularly at large error tolerances, a significant amount of the ODE solvers time is spent on overheads.

The bottom five entries in Table 4.3 show the equivalent results for the non adaptive solvers. The percentage of time spent on overheads is much smaller here than it is for `ode45`, which explains why the direct solvers can perform a larger number of evaluations of the derivatives, and still solve the equations more quickly (compare the computation times for `ode45` at `AbsTol`=10^{-6} and `ode1`). Nevertheless, it still represents approximately 15% of computation time.

Solver	Error in conservation of mass/l
`ode1`	2.6×10^{-5}
`ode2`	2.5×10^{-4}
`ode45`(10^{-6})	1.9×10^{-4}
`ode45`(10^{-8})	1.9×10^{-4}

Table 4.4: Departure from conservation of mass from a single inspiration for a selection of the ODE solvers. The figures in brackets for `ode45` represent the values of `AbsTol`.

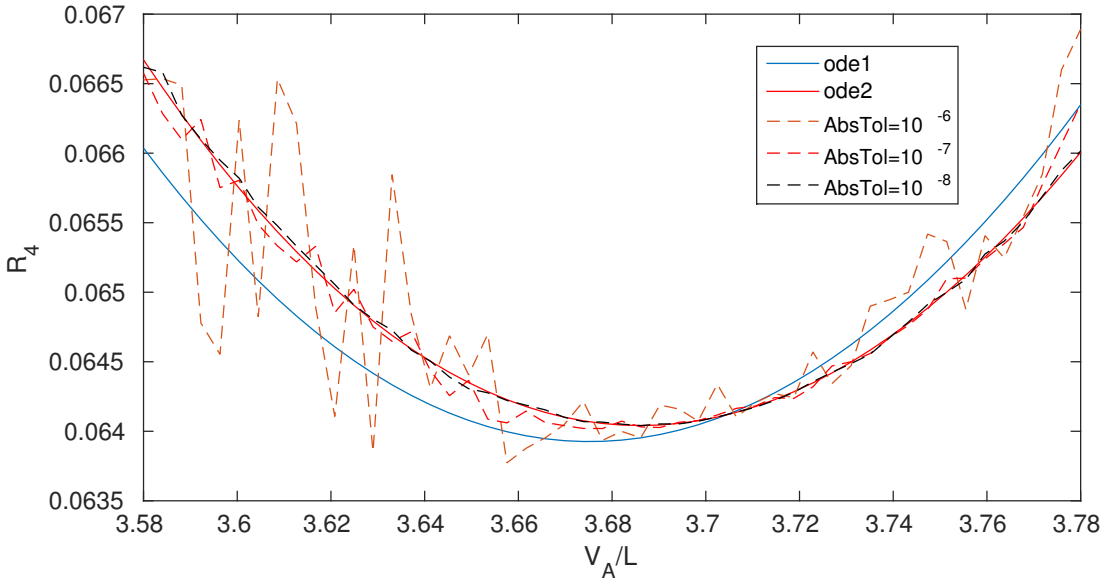

Figure 4.15: Plot of the variation of R_4 with V_A, where we have solved the ODEs in our model with `ode1`, `ode2` (solid curves), and `ode45` with several different values of `AbsTol` (dashed curves).

4.6.2.2 Conservation of mass and parameter estimation

Table 4.4 shows the error in the conservation of mass (the difference between the change of mass of gas in the lungs and the sum of the net mass of gas entering the lungs via both ventilation and perfusion) from a single inspiration for a selection of these solvers. We note that in all cases these errors are a number of orders of magnitude larger than the method of solution we developed for the deadspace.

What effect do the different solvers, and their different accuracies have on parameter estimation? While the full parameter estimation procedure is yet to be established (we will discuss this in Chapter 5) we focus here on the effect of the ODE solver choice on the variation with the total alveolar volume (V_A) of the sum of the squares of the differences between measured and simulated nitrogen flux during the washout, R_4,

$$R_4 = \sum_l \left(u(t_l) \left(F_{\mathrm{N}_2}^{exp}(t_l) - F_{\mathrm{N}_2}^{sim}(t_l) \right) \right)^2, \tag{4.152}$$

where $F_{\mathrm{N}_2}^{exp}$ and $F_{\mathrm{N}_2}^{sim}$ are the experimentally measured and simulated nitrogen fractions respectively. The variation of R_4 with different values of V_A, for a variety of different ODE

Solver	Recovered V_A/l	Standard Deviation/l
`ode1`	3.675	2×10^{-5}
`ode2`	3.686	2×10^{-5}
`ode45` (10^{-6})	3.68	0.013
`ode45` (10^{-7})	3.69	0.0074

Table 4.5: Estimated values of V_A and their standard deviations calculated from 10 repeats with varying initial guesses. Where the solver is `ode45`, the value in brackets is the value of `AbsTol`.

solvers, is plotted in Figure 4.15. This plot has a number of features which are worth further consideration.

Consider the curves produced by the adaptive solver, `ode45` (dashed lines). Each curve seems to follow the same overall shape, with the same global minimum. However, even at `AbsTol`=10^{-8}, the curve is not smooth, and becomes increasingly bumpy as the value of `AbsTol` is increased. The problem with this is that it makes it difficult for minimisation algorithms to find the global minima of these curves. Any algorithms that involve derivatives cannot be used, and even those that do not may find local minima.

To examine this problem quantitatively we used `fminsearch` (a derivative free minimisation function in MATLAB) to try and find the value of V_A which minimises R_4. Table 4.5 summarises the results of ten attempts with different initial guesses (ranging between 3l and 4l). As expected, the smaller error tolerance leads to a smaller standard deviation, although looking at the curves it is perhaps surprising `fminsearch` performs as well as it does.

If we look at the curves produced using the non-adaptive solvers they, at least when viewed on this scale, appear to be smooth. However, in this case the global minima for the two solvers are different. Similarly in Table 4.5 we see that the mean of the estimated values for the two non-adaptive solvers are slightly different, with the mean of the values found using `ode2` being the same as that of the adaptive solvers. This difference is down to the difference between `ode1`, which only evaluates derivatives once every time step, and assumes the values of these derivatives are the same for the entire time step, and the higher order methods which evaluate the derivatives multiple times.

In the analysis in Table 4.5 we used `fminsearch` to estimate the value of V_A, because the variation of R_4 with V_A (when we solve the ODEs with an adaptive solver) is not smooth enough to allow a derivative based minimisation method. However, if we use a non-adaptive solver, then this restriction does not apply and we can use `lsqnonlin` (unlike `fminsearch`, `lsqnonlin` is a derivative based method). As shown in Table 4.6, `lsqnonlin` finds the minimum significantly quicker than `fminsearch` does.

In many cases adaptive ODE solvers may be more efficient than those with a fixed time step, but that does not appear to be the case here. These solvers take variable time steps, assuming that $u(t)$ is a continuous function that can be sampled at any t. However,

Solver	Recovered V_A/l	Standard Deviation/l	Computation Time/s
`lsqnonlin`	3.686	4×10^{-6}	183
`fminsearch`	3.686	2×10^{-5}	615

Table 4.6: A comparison of using `fminsearch` and `lsqnonlin` for recovering V_A with the ODEs solved with `ode2`. These are the results of 10 repeats with differing initial guesses.

as we noted above, we do not know instantaneous values of $u(t)$, and it appears that the variability introduced by the variable step size prevents the use of efficient derivative based minimisation algorithms. This slowed the estimation of V_A almost three times, and this is only likely to worsen as the number of parameters we want to estimate increases. Thus it appears a non adaptive solver is more appropriate for our needs.

However, we saw in this section that up to 15% of the non adaptive solver's computation time was spent on overheads, and the precision with which mass was conserved was not as good as our method of solving the deadspace. The advantage of using inbuilt solvers is that they are well validated, and can be used for a wide range of problems. However, we only want to solve a small number of known equations. Thus in the next section, we develop our own method for solving the governing ODEs of the alveolar compartments. This method conserves mass with a very high degree of precision, and is (slightly) faster than `ode1`.

4.6.3 A better method for solving the alveolar compartments

In this subsection we develop our own method for solving the governing equations of the alveoli that were presented in Section 4.3. As we describe below, we assume we start at a time point where we know the volumes of CO_2 and O_2 in each compartment, and total volume of each compartment (thus we can infer the volume of N_2). We then integrate the relevant governing equations to find the values of these volumes at the next time point. We use the same approach to solve the governing equations of the body compartments.

To find the volumes of CO_2 and O_2 during inspiration we integrate (4.75) from t_{l-1} (where we assume we know the solution) to the next time point, t_l,

$$\int_{t_{l-1}}^{t_l} \frac{dV_i^g(t)}{dt}\,\mathrm{d}t = \int_{t_{l-1}}^{t_l} \dot{V}_i^{in} u(t) F_{Ii}^g(t) + Q_T Q_i \left(C_{\bar{v}}^g(t) - C_{ai}^g(t)\right)\,\mathrm{d}t, \tag{4.153}$$

where,

$$V_i^{\mathrm{O}_2}(t) = F_i^{\mathrm{O}_2}(t) V_{Ai}(t), \tag{4.154}$$

and

$$V_i^{\mathrm{CO}_2}(t) = F_i(t)\left(V_{Ai}(t) + V_i V_t V_A\right). \tag{4.155}$$

The first term on the right hand side of (4.153) represents the volume of gas that enters the alveolar compartment from the connected personal deadspace compartment. When we solve the personal deadspace compartments (as described in Section 4.5), we calculate

$V^g_{in(i)}(t)$, the cumulative volume of gas that has left personal deadspace compartment i by time t. Since this gas must go into the connected alveolar compartment we have,

$$\int_{t_{l-1}}^{t_l} \dot{V}_i^{in} u(t) F^g_{Ii}(t)\ \mathrm{d}t = V^g_{in(i)}(t_l) - V^g_{in(i)}(t_{l-1}). \tag{4.156}$$

The second term on the right hand side is the volume of gas exchanged with the blood over the period in compartment i, $V^g_{bi}(t_l)$, and is given by,

$$V^g_{bi}(t_l) = \int_{t_{l-1}}^{t_l} Q_T Q_i \left(C^g_{\bar{v}}(t) - C^g_{ai}(t)\right)\ \mathrm{d}t. \tag{4.157}$$

We approximate this integral with the first order Euler expression,

$$V^g_{bi}(t_l) = \Delta t Q_T Q_i \left(C^g_{\bar{v}}(t_{l-1}) - C^g_{ai}(t_{l-1})\right). \tag{4.158}$$

As previously described, we use the blood-gas model developed by O'Neill and Robbins [81] to convert alveolar CO_2 and O_2 fractions to concentrations. We assume that nitrogen concentrations are given by,

$$C^{\mathrm{N}_2} = S_b \left(P_b - P_{\mathrm{H_2O}}\right) F^{\mathrm{N}_2}, \tag{4.159}$$

where S_b is the solubility of nitrogen in the blood, P_b is barometric pressure, and $P_{\mathrm{H_2O}}$ is the saturated partial pressure of water at BTPS.

Substituting (4.156) and (4.157) into (4.153) we then have

$$V^g_i(t_l) = V^g_i(t_{l-1}) + \left(V^g_{in(i)}(t_l) - V^g_{in(i)}(t_{l-1})\right) + \Delta t Q_T Q_i \left(C^g_{\bar{v}}(t_{l-1}) - C^g_{ai}(t_{l-1})\right). \tag{4.160}$$

To find the volume of each compartment, we use a two stage approach. First we calculate $V_{\mathrm{tp}(i)}$, which is the value $V_{Ai}(t_l)$ would have if the fraction of CO_2 in the tissue stayed constant. This is found by integrating (4.79) assuming the last term is zero, thus

$$V_{\mathrm{tp}(i)} = V_{Ai}(t_{l-1}) + \Delta t \bar{u}(t_l) + \sum_g V^g_{bi}(t_l), \tag{4.161}$$

where we have used the definition of $V^g_{bi}(t_l)$ from (4.157).

Next we apply conservation of mass to work out how much CO_2 must move from the gas phase into the tissue to equalise the gas fractions between the two phases (we defined V_t in Section 4.3 such that these two fractions would be equal). We then subtract this volume of CO_2 from $V_{\mathrm{tp}(i)}$ to find the new alveolar volume.

By definition,

$$F_i^{\mathrm{CO}_2}(t_l) = \frac{V_i^{\mathrm{CO}_2}(t_l)}{V_{Ai}(t_l) + V_i V_A V_t}. \tag{4.162}$$

By conservation of mass, $V_{Ai}(t_l)$ will be $V_{\mathrm{tp}(i)}$ less the net change in CO_2 dissolved in the tissue from t_{l-1} to t_l, thus

$$V_{Ai}(t_l) = V_{\mathrm{tp}(i)} - V_i V_A V_t \left(F_i^{\mathrm{CO}_2}(t_l) - F_i^{\mathrm{CO}_2}(t_{l-1})\right). \tag{4.163}$$

Substituting (4.162) into (4.163) we have,

$$V_{Ai}(t_l) = V_{\mathrm{tp}(i)} - V_i V_A V_t \left(\frac{\bar{V}_i^{\mathrm{CO}_2}(t_l)}{V_{Ai}(t_l) + V_i V_A V_t} - F_i^{\mathrm{CO}_2}(t_{l-1}) \right), \tag{4.164}$$

thus following some rearrangement we can write

$$V_{Ai}(t_l)^2 + bV_{Ai}(t_l) = c, \tag{4.165}$$

where

$$b = -\left(V_{\mathrm{tp}(i)} + V_i V_A V_t \left(F_i^{\mathrm{CO}_2}(t_{l-1}) - 1 \right) \right), \tag{4.166}$$

$$c = V_i V_A V_t \left(V_i V_A V_t F_i^{\mathrm{CO}_2}(t_{l-1}) + V_{\mathrm{tp}(i)} - V_i^{\mathrm{CO}_2}(t_l) \right). \tag{4.167}$$

This quadratic could be solved using the quadratic formula. However, given this involves numerically estimating a square root, it can be solved more quickly and accurately using the following iterative scheme [98],

$$V_{Ai}(t_l)(j+1) = \frac{V_{Ai}(t_l)(j)^2 + c}{2V_{Ai}(t_l)(j) + b}, \tag{4.168}$$

where j is the index of iteration (we used ten iterations setting $V_{Ai}(t_l)(1) = V_{\mathrm{tp}(i)}$ and found this sufficient to solve the equation to machine precision).

Finally to solve the body compartments we numerically integrate (4.91), using Euler's method, to obtain

$$P_{vl}^{\mathrm{N}_2}(t_l) = P_{vl}^{\mathrm{N}_2}(t_{l-1}) + \frac{\Delta t Q_T Q_l}{\lambda V_{tl}} \left(P_a^{\mathrm{N}_2}(t_{l-1}) - P_{vl}^{\mathrm{N}_2}(t_{l-1}) \right), \tag{4.169}$$

and (4.94) which gives

$$C_{\bar{v}}^{\mathrm{O}_2}(t_l) = C_{\bar{v}}^{\mathrm{O}_2}(t_{l-1}) + \frac{1}{V_L} \left(V_b^{\mathrm{O}_2}(t_l) - \dot{V}_{\mathrm{O}_2} \Delta t \right), \tag{4.170}$$

where $V_b^{\mathrm{O}_2}(t_l) = \sum_g V_{bi}^{\mathrm{O}_2}(t_l)$.

On expiration our approach is the same, except to find the volumes of CO_2 and O_2, we integrate (4.80) to obtain

$$V_i^g(t_l) = V_i^g(t_{l-1}) + \Delta t \dot{V}_i^{ex} \bar{u}(t_{l-1}) F_i^g(t_{l-1}) + V_{bi}^g(t_l). \tag{4.171}$$

4.6.4 Performance of our solution method

At the start of this section we noted that our method of solving the governing equations of our model (for it to be useful in the inverse problem) should be fast and conserve mass with high precision. In this section we briefly discuss how our implementation in MATLAB performs relative to these two criteria.

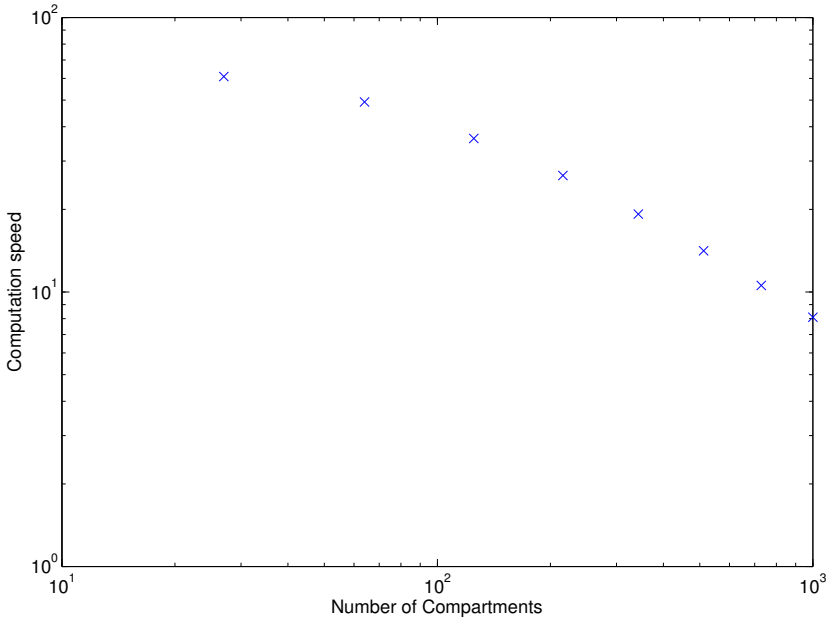

Figure 4.16: A plot showing the variation of the computation speed (defined as the duration of the measured procedure divided by the duration of the simulation) with the number of compartments simulated, performed on a standard desktop PC. Each point is the average of ten runs.

Figure 4.16 shows the variation of computation speed (defined as the duration of the measured procedure divided by the duration of the simulation) with the number of compartments simulated. The issue of the appropriate number of compartments is addressed in Section 5.3.3, but here we note that with $N = 125$ the model runs approximately 40 times faster than real time.

If we apply conservation of mass to the whole of our model, we have that the change of the volume of a gas in the lungs (which is the sum of its changes in the common deadspace compartment, N personal deadspace compartments, and N alveolar compartments) over a period is equal to the net volume entering at the mouth plus the net volume entering the alveolar compartments from the blood. This was checked for both inspiration and expiration for each gas, the mean departure from conservation of mass over 72 breaths was 4×10^{-15}l. Thus our implementation conserves mass with almost machine precision.

4.7 Summary

In this chapter we have developed a novel mathematical model of the lungs, with a small parameter set, that can be solved quickly, while conserving mass with a high degree of precision. Our conceptual picture of the lungs is of a large number of functional lung units, whose fractional volume changes, perfusions, and volumes are determined by a continuous bimodal lognormal distribution.

We presented an efficient method of discretising this distribution into a finite number of alveolar compartments. We assume these compartments are perfectly mixed, and are where all gas exchange and volume change in the lungs occurs. We wrote down the governing equations for the alveolar compartments, and explained that, since we include the volume change due to gas exchange with the blood in our equations for alveolar volume, the fractional ventilation of the alveolar compartments cannot be constant. Instead, we defined the fractional volume change of a compartment, which ensures the compartment's volume behaves in a physically realistic manner.

We assumed the conducting airways that connect the alveoli to the mouth can be represented by deadspace compartments in which all gas transport is by convection. We derived a governing advection equation for these volumes, and compared two possible flow patterns, plug flow and Poiseuille flow, to experimental data. We found that plug flow matched the data better, and thus this flow pattern was adopted. We also compared two distributions of the deadspace volume, and explained why we adopted a normal distribution.

Finally we presented efficient methods for solving the governing equations of the deadspace and alveolar compartments that conserved mass with a high degree of precision. In the next chapter we discuss how we estimate the parameters of this model from experimental data.

Chapter 5

The inverse problem

In this chapter we focus on the 'inverse' problem, how we estimate the parameters required by our model (outlined in Chapter 4) from experimental data. To be useful as part of the lung function test outlined in Chapter 1, we need parameter estimates that are reliable and that can be obtained within a reasonable amount of time (hours not days). This chapter details how we estimate the parameters from experimental data taken with the LGA, and our investigation of the uncertainty of these parameter estimates.

First we describe the minimisation problem we solve to estimate these parameters. We define the objective function and explain our approach to finding the set of parameters that minimise its value. We then discuss in some detail how the objective function is evaluated, paying particular attention to how we estimate the venous concentrations and initial conditions required by our model.

We then use noise-free synthetic data to consider which parameters we can estimate from the data, and the influence of those parameters we cannot estimate. We consider a model for experimental error, both for the LGA and for a device built to the ERS specifications [95]. We use this error model to add noise to our synthetic data, and consider the uncertainty in our parameter estimates due to experimental noise. Finally, we suggest a procedure for estimating one of the parameters that is not identifiable from a nitrogen washout procedure.

5.1 Minimisation problem

In this section we define the minimisation problem that we solve to estimate the parameters of the model we described in Chapter 4 from a nitrogen washout procedure. For our purposes a 'nitrogen washout procedure' comprises ten minutes of air breathing followed by five minutes of breathing pure O_2.

The objective function, R, that we use to compare simulated and experimental data is

$$R = \sum_{\text{air}-\text{breathing}} u(t_n)^2 \left(\left(F_M^{\text{CO}_2}(t_n) - F_S^{\text{CO}_2}(t_n) \right)^2 + \left(F_M^{\text{O}_2}(t_n) - F_S^{\text{O}_2}(t_n) \right)^2 \right) + \sum_{\text{O}_2\text{breathing}} u(t_n)^2 \left(F_M^{\text{N}_2}(t_n) - F_S^{\text{N}_2}(t_n) \right)^2, \quad (5.1)$$

where t_n are the discrete measurement times, and F_M^g and F_S^g are the measured and simulated fractions in the measurement plane of the LGA respectively. The 'air-breathing period' for the objective function above is defined as the five minutes of normal breathing immediately preceding the washout. The first two minutes of normal breathing are discarded to allow for the subject adjusting to breathing on the LGA, and the following three minutes are used to allow the model to settle into a quasi steady state (as described in Section 5.2.2).

In (5.1) we consider the CO_2 and O_2 fluxes during the air-breathing period, and N_2 flux during the high O_2 breathing period. The reason for not using N_2 during the air breathing period is that during this period the variation of N_2 is small, and thus the signal is relatively noisy (see the expirograms from normal breathing in Appendix B).

During the O_2 breathing phase, we have, as we will explain later, only an approximate idea of how the venous O_2 and CO_2 concentrations behave. As N_2 is insoluble, the influence of perfusion on its signal is much lower, and thus during the washout we only consider the N_2 flux.

The evaluation of (5.1) for a given parameter set, discussed in the next section, is a multi-stage process that takes approximately three minutes. Thus Bayesian methods for solving the inverse problem, discussed in Section 2.3, which typically require very large numbers of function evaluations, are not suitable here. Instead we use `fmincon`, a MATLAB non-linear least squares optimisation function, to find the set of parameters (the parameters required by our model are listed as a reminder in Table 5.1) that minimise our objective function. This is referred to throughout this thesis as: "The parameter estimation procedure".

The reason for using this optimisation function is that it is a derivative based method (we saw in Section 4.6.1 that derivative based methods appear to be faster for this problem), that can easily be implemented in parallel. This means that when the minimisation algorithm estimates derivatives of the objective function numerically it can evaluate the objective function on different CPUs simultaneously. For an objective function as computationally expensive as ours, this provides a significant speed up. All of the parameter estimates in this thesis were performed using `fmincon` implemented in parallel on the Arcus-B high performance computer[1].

[1]I would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work. http://dx.doi.org/10.5281/zenodo.22558

Symbol	Description
V_A	Total alveolar volume when the lungs are at FRC
V_t	CO_2 equivalent tissue volume
Q_T	Cardiac output
σ_x	Width of the compliance distribution
σ_y	Width of the perfusion distribution
ρ	Correlation of compliance and perfusion
V_{DC}	Apparatus deadspace volume
V_{DP}	Total personal deadspace volume
σ_z	Width of the deadspace distribution

Table 5.1: A summary of the parameters that describe our model. V_{DC} is set equal to the measured apparatus deadspace volume of 155 ml, but all the other parameters must be estimated.

In the next section we describe in some detail how (5.1) is evaluated.

5.2 Evaluation of the objective function

In the previous section we explained that we use the MATLAB least-squares solver, `fmincon`, to find the parameters, listed in Table 5.1, that minimise the objective function, (5.1). To evaluate the objective function we must solve the forwards problem, where we solve the model with known parameter values. We described a computationally efficient means of doing this in Sections 4.5 and 4.6. However, when we solve the forwards problem, as well as values for the parameters, we also need to know initial conditions for the differential equations, and the composition of the inspired gas and mixed venous blood.

The initial conditions required are the gas fractions in every deadspace and alveolar compartment, and the starting volumes of the alveolar compartments. The gas fractions will be inhomogeneous across different compartments (see Figure 5.1), and there is no guarantee that the first inspiration will start from FRC. Thus, as we will describe in Section 5.2.1, the initial conditions must be carefully calculated.

N_2 is very insoluble and not metabolically active. Thus we adopt the simple four compartment model of Baker and Farmery [7] (described in Section 4.3.2) to estimate the venous N_2 concentration. However, since CO_2 and O_2 are exchanged in much greater volumes, and are metabolically active, such a model is not appropriate.

We estimate the venous concentrations of O_2 and CO_2 using a combination of our model, and data collected from the LGA. To estimate the variation of the venous concentrations during the air breathing period we solve a simplified version of our model to estimate the venous CO_2 and O_2 concentrations on every breath; and then fit straight lines to these values. The starting values for the venous concentration are set so that the estimated (with the full model) and experimentally measured CO_2 and O_2 consumption rates are equal.

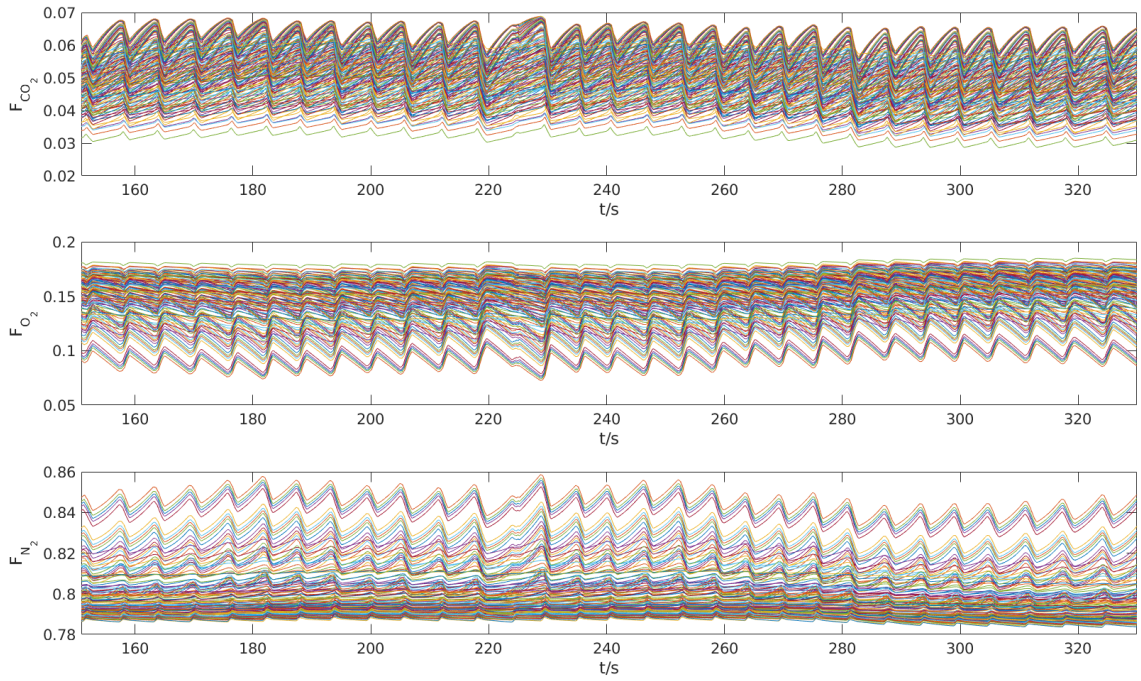

Figure 5.1: An example of the behaviour of the simulated alveolar gas fractions during the burn in period. Each line represents the alveolar gas fraction in one compartment of the model.

However, during the high O_2 breathing phase, the recirculating arterial O_2 will lead to a rise in the venous O_2 concentrations. We introduced a simple model of this in Section 4.3.2 that had a single parameter, V_L, the volume of distribution for venous O_2. To estimate the behaviour of the venous O_2 concentration during the high O_2 breathing period, a value for this parameter must also be found.

Since we estimate the initial conditions and venous concentrations with our model, they will vary as the parameters of the model change in each objective function evaluation. Thus every time we evaluate the objective function, we must: estimate the venous gas concentrations during air breathing; establish initial conditions; and find an appropriate value of V_L. Each of these steps is discussed in more detail in the following sections.

5.2.1 Initial conditions

To initialise the model we must specify the initial conditions in each one of the alveolar and deadspace compartments at the start of the first inspiration (the first inspiration we simulate is two minutes after we have started collecting data). We cannot measure these initial conditions, and so we use a continuous model to provide guesses for them. We then simulate three minutes of normal breathing using the full model before performing any analysis. This three minute 'burn in' period is to prevent differences between the full and continuous model influencing the value of the objective function, and thus our parameter

estimation.

The initial conditions we set are listed below.

- The alveolar gas fractions in a given compartment are set equal to the value found for that compartment in the equivalent continuous ventilation model (see Section 5.2.2.2).
- The nitrogen partial pressures in the body compartments are set equal to the perfusion weighted average nitrogen partial pressure of the alveolar compartments.
- The gas fractions in each personal deadspace volume are set equal to those of the connected alveolar compartment.
- The gas fraction in the common deadspace volume is set equal to the end tidal value from the previous breath.

Figure 5.1 shows an example of the behaviour of the alveolar gas fractions in each compartment during the burn in period. In this figure we can see that the alveolar gas fractions are approximately steady during this burn in period. It should be noted that we would not expect them to be exactly steady, both because of the hypocapnia that we will discuss in Section 5.2.2.3, and the irregularity of spontaneous breathing. We now discuss how we set the initial volumes of the alveolar compartments.

5.2.1.1 Initialisation of lung volume

When a subject breathes spontaneously their expected end expiratory lung volume is, by definition, FRC. However, there is significant breath to breath variation in this quantity. If we plot the integral of the flow we measure with the LGA ("The tidal volume record"), then, as shown in Figure 5.2, the value of the integral at the local minima increases with time. This is due to the inequality of O_2 consumption and CO_2 production (the rate of the former is greater than the rate of latter), and is thus an expected feature. However, it makes it difficult to identify when the lungs are at FRC. The correct identification of FRC is important, because our definition of V_i is that it is the fractional volume of compartment i when the lungs are at FRC.

To estimate where FRC is on the tidal volume record, we fit a cubic smoothing spline to the local minima of the tidal volume record using the MATLAB function `csaps`. This function has one adjustable parameter that sets the degree of smoothing. The value of this parameter used in all cases was 10^{-7}, with zero representing a straight line and one a spline going through every point. Our value was set by inspection of a number of datasets, where we aimed to find a value that responded slowly to trends in the minima. Both smaller values and larger values of p were tried, but the choice of p had little effect on

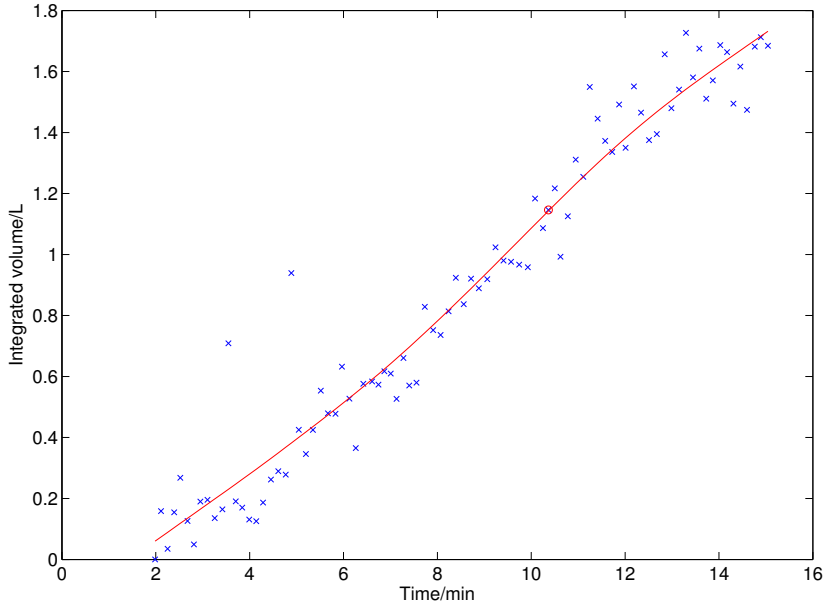

Figure 5.2: A plot of the troughs (the end of expiration) in the tidal volume record (the integral of the flow measured at the mouth with respect to time) of a healthy volunteer, and a cubic smoothing spline with smoothing parameter, $p = 10^{-7}$. The red circle is the expiration prior to the start of the washout.

parameter estimation. An example plot of the local minima in the tidal volume record, and a spline fitted through them is shown in Figure 5.2.

The first breath we simulate may not be at FRC (see Figure 5.2), thus when initialising the compartmental volumes, we set

$$V_{Ai}(t_0) = V_i V_A + \Delta V_i V_{\text{diff}}, \tag{5.2}$$

where V_A is the total volume of the alveolar compartments at FRC, and V_{diff} is the deviation of the starting lung volume from FRC. V_{diff} is estimated as follows.

Initially, we assume that V_{diff} is equal to the difference between the smoothing spline and the local minimum of the tidal volume record at the start of the first breath of the air breathing period. However, it is very important for the estimation of a number of parameters that the difference from FRC of the lung volume on the first breath of the washout is correct. Therefore, once we have established the slopes and starting values of the venous concentrations during normal breathing (see Section 5.2.2), we run the full model, and update our value of V_{diff} by the difference between: (i) the expected deviation of the lung volume from FRC at the start of the washout (the difference between the local minimum from the tidal volume record and smoothing spline at that point); and (ii) the difference between the model's alveolar volume at that point and V_A.

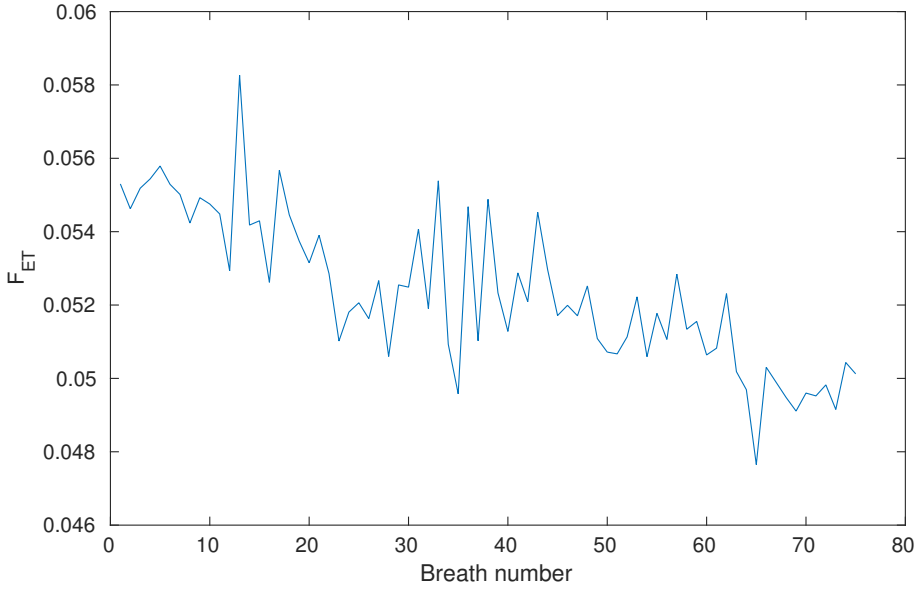

Figure 5.3: An example of the declining end tidal CO_2 fractions of a healthy volunteer breathing on the LGA.

5.2.2 Venous concentrations during air breathing

We previously stated that the venous concentrations during the air breathing period were estimated to match the net volumes of CO_2 and O_2 passing through the mouth in the air breathing period. If the venous concentrations were constant during the air breathing period, then this would be a relatively simple condition to apply. However, as we now explain, this is not the case.

During our data collection we observed that when breathing through a mouthpiece subjects tended to breath more deeply than they would normally. This caused a gradual fall in $F_{ET}^{CO_2}$ (see Figure 5.3) and hence $C_{\bar{v}}^{CO_2}(t)$. This phenomenon is known as hypocapnia. In this subsection we explain how we include this behaviour in our model.

First we describe how we estimate pulmonary gas exchange using data from the LGA. Next we describe a continuous ventilation version of the full model we described in Section 4.2. This model is used to estimate how the venous concentrations vary during the air breathing period. Finally, we describe how we use the full model to find the starting venous concentrations that ensure the net volumes of CO_2 and O_2 passing through the mouth in the air breathing period in the model match those measured using the LGA.

5.2.2.1 Estimation of gas exchange at the pulmonary capillary

Before applying the continuous model we must first estimate the volume of each gas exchanged with the blood for each breath using data from the LGA. By conservation of mass, the volume of a gas exchanged with the blood at the pulmonary capillary on a breath is equal to the difference between: (i) the net volume of that gas entering the

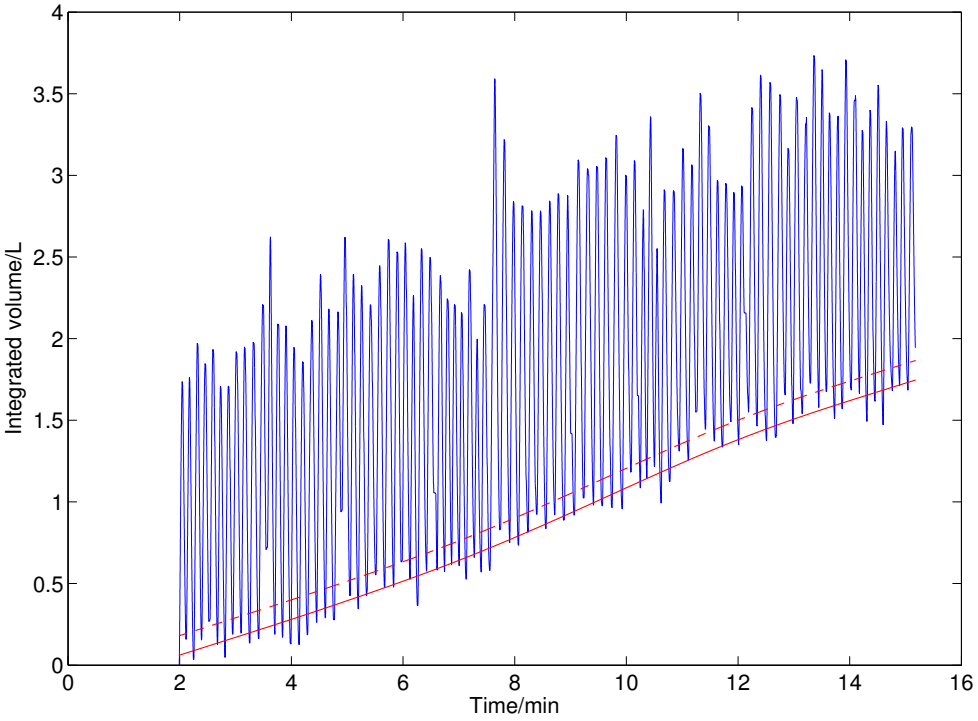

Figure 5.4: A plot of the integral of the flow measured using the LGA, the 'tidal volume record'. The red line is a smoothing spline fitted to the local minima (see Section 5.2.3). The smoothing spline represents points on the tidal volume record at which the lung is at the same volume. The dashed-red line is this line shifted up by the standard deviation of the difference of local minima from this line.

lungs through the mouth (measured with the LGA); and (ii) the change in the volume of that gas in the lungs. Given that the end expiratory volume varies significantly during spontaneous breathing (see Figure 5.4), the volume of the respiratory gases in the lungs at the end of a breath will also vary significantly, even if the concentrations of the gases do not change. Thus to estimate the breath-by-breath pulmonary gas exchange, we change our definition (just for the purposes of this calculation) of a breath to try and ensure each 'modified breath' finishes at the same lung volume.

Figure 5.4 shows an example of a tidal volume record, with both a spline (solid red line) and a spline shifted up by the standard deviation of the difference of the spline from the local minima (dashed red line) also illustrated. As explained in Section 5.2.1.1, we assume that the smoothing spline represents the location of FRC on the tidal volume record (i.e. if the local minimum of the tidal volume record coincides with this line the subject has expired to FRC on that breath). Further, we assume that if we shift this spline up by a fixed amount then every point at which the tidal volume record crosses the spline represents a time when the lung is at the same volume. For the purposes of estimating the pulmonary gas exchange, we define a 'modified breath' as from a point where the adjusted smoothing spline, i.e. the red dashed in line in Figure 5.4, intersects the tidal volume record during a conventionally defined expiration to the next such point. We note that such a modified breath may include multiple inspirations/expirations.

We then estimate the pulmonary gas exchange of gas g (neglecting any change in the volume of that gas stored in the lungs) on that modified breath, $V_p^g(m)$, using

$$V_p^g(m) = \int_{\hat{t}_{m-1}}^{\hat{t}_m} F_M^g(t)u(t)\ \mathrm{d}t, \tag{5.3}$$

where $\hat{t}_m$ is the time at which modified breath m ends, $F_M^g(t)$ is the fraction of gas g measured at the mouth, and $u(t)$ is the measured flow rate. As this calculation is independent of our model it is only done once during parameter estimation, prior to the first evaluation of the objective function.

5.2.2.2 The continuous model

To model the variation of the venous gas concentrations during the air breathing period, we apply a continuous ventilation model to each modified breath. Using this model we estimate the venous gas concentrations on each modified breath, and then fit a straight line to these estimated venous concentrations.

The reason for using a continuous ventilation model rather than the full model is that its governing equations (listed below) can be solved very quickly. We describe the continuous model we use in more detail below, but the guiding principle is that it is designed to correspond as closely as possible to the full model. Thus, for example, when calculating

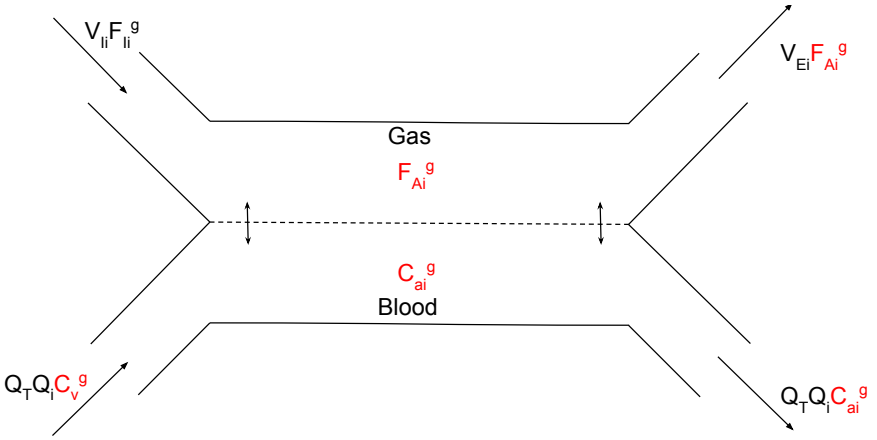

Figure 5.5: A schematic diagram of a compartment in the continuous model, showing gas entering and leaving the compartment, and exchange between the gas and blood volumes. The model is solved to estimate the quantities in red.

the inspired ventilation into a compartment we consider the volume of a particular gas that would be inspired on that breath into that compartment in the full model.

One compartment of this continuous ventilation model is shown in Figure 5.5. We assume that each of these compartments has the same fractional volume (V_i), perfusion (Q_i), ventilation ($\dot{V}_i$), and associated volume of deadspace ($v_i V_{DP}$) as the equivalent compartment in the full model.

Governing equations

In Chapter 2 we derived the governing equation for gas fractions in the alveoli in a continuous ventilation model,

$$\frac{d\left(F^g_{Ai}(t)V_{Ai}\right)}{dt} = \dot{V}_{in}F^g_{Ii}(t) - \dot{V}_{ex}F^g_{Ai}(t) + Q_i Q_T \left(C^g_v - C^g_{ai}\right), \tag{5.4}$$

where F^g_{Ai} and F^g_{Ii} are the alveolar and inspired fractions of gas g in compartment i, C^g_v and C^g_{ai} are the arterial and venous concentrations of that gas, V_{Ai} is the volume of the compartment, and $\dot{V}_{in}$ and $\dot{V}_{ex}$ are the ventilations entering and leaving the compartment respectively.

We assume there is no change in the stored volume of gas in the compartment over the modified breath, and thus the integral of the hand side of this expression over a modified breath is zero. Integrating (5.4) over the modified breath m in compartment i we have

$$\int_{\hat{t}_{m-1}}^{\hat{t}_m} \dot{V}_{in}F^g_{Ii}(t) - \dot{V}_{ex}F^g_{Ai}(t)\ \mathrm{d}t = Q_i Q_T T_m \left(C^g_a\left(F^{\mathrm{CO}_2}_{Ai}(m), F^{\mathrm{O}_2}_{Ai}(m)\right) - C^g_v(m)\right), \tag{5.5}$$

where $C^g_a\left(F^g_{Ai}(m), F^g_{Ai}(m)\right)$ is the average arterial concentration of gas g in the compartment (which we assume can be calculated from the average gas fractions in the compartment, $F^g_{Ai}(m)$, using the blood model), $C^g_v(m)$ is the average venous concentration of gas g on modified breath m, and $T_m = \hat{t}_m - \hat{t}_{m-1}$.

The volume of gas entering compartment i during modified breath m, $V_{Ii}(m)$, is

$$V_{Ii}(m) = \int_{\hat{t}_{m-1}}^{\hat{t}_m} \dot{V}_{in} \, \mathrm{d}t. \tag{5.6}$$

We also define the total volume of all gases leaving the compartment during modified breath m, $V_{Ei}(m)$, where

$$V_{Ei}(m) = \int_{\hat{t}_{m-1}}^{\hat{t}_m} \dot{V}_{ex} \, \mathrm{d}t. \tag{5.7}$$

Using these definitions we re-write (5.5) in the following way,

$$V_{Ii}(m)F_{Ii}^g(m) - V_{Ei}(m)F_{Ai}^g(m) = Q_i Q_T T_m \left(C_a^g \left(F_{Ai}^{\mathrm{CO}_2}(m), F_{Ai}^{\mathrm{O}_2}(m) \right) - C_v^g(m) \right), \tag{5.8}$$

where $F_{Ii}^g(m)$ is the average inspired gas fraction into compartment i during modified breath m. The reason for writing the equation in this form will become apparent below.

The total volume of gas exchanged with the blood during modified breath m, $V_p^g(m)$, is equal to sum of the volume of gas exchanged with the blood in each compartment during modified breath m. Thus we have

$$\sum_i Q_i Q_T \left(C_a^g \left(F_{Ai}(m)^{\mathrm{CO}_2}, F_{Ai}(m)^{\mathrm{O}_2} \right) - C_v^g(m) \right) = V_p^g(m), \tag{5.9}$$

where $V_p^g(m)$ is estimated as described in the previous section. Below we describe how we calculate $V_{Ii}(m)F_{Ii}^g(m)$ and $V_{Ei}(m)$ to correspond as closely as possible to their equivalents in the full model using experimental data. Once this is done we have $2(N+1)$ equations, (5.8) for CO_2 and O_2 in each compartment and (5.9) for both gases, that we solve to find the $2(N+1)$ unknowns ($F_{Ai}^g(m)$ and $C_v^g(m)$) on each modified breath.

Inspired gas volume

Consider the gas inspired into compartment i of the full model on modified breath m. The total volume of this gas, $V_{Ii}(m)$, is $\sum_j \dot{V}_i V_T(j)$, where $V_T(j)$ is the inspired tidal volume on breath j, and we sum over all the j inspirations within the modified breath m. We consider each of the j inspirations separately.

The first $v_i V_{DP}$ (as long as this is smaller than $\dot{V}_i V_T(j)$) of the inspired volume will have started inspiration in the personal deadspace, where the average fraction of gas g is $\frac{F_{Ai}^g(m-1)+F_{DP}^g(m-1)}{2}$ (where F_{DP}^g is the common fraction at the end of the personal deadspace at the end of expiration, see Section 4.5.4). Since we assume plug flow this will still be the average fraction of this volume of gas when it enters the alveolar compartment. The next $\dot{V}_i V_{DC}$ (as long as this is smaller than $\dot{V}_i V_T(j) - v_i V_{DP}$) of the inspired tidal volume into the compartment will have started the breath in the common deadspace volume, where we assume the fraction of gas g is equal to the end tidal fraction of that gas from the previous expiration, $F_{ET}^g(j-1)$. Finally, the remainder of the tidal volume, which

started inspiration outside the lungs, will enter the alveolar compartment. We measure the fraction of this volume of gas as it passes through the measurement plane of the LGA.

We assume the gas inspired into the equivalent compartment in the continuous ventilation model has the same composition, and thus summing over the j inspirations in modified breath m, we have that

$$\begin{aligned} V_{Ii}(m)F^g_{Ii}(m) = \sum_j \Big(& \frac{F^g_{Ai}(m-1) + F^g_{DP}(m-1)}{2} \min\left(v_i V_{DP}, \dot{V}_i V_T(j)\right) \\ & + F^g_{ET}(j-1) \min\left(\dot{V}_i V_{DC}, \max\left(\dot{V}_i V_T(j) - v_i V_{DP}, 0\right)\right) \\ & + \int_0^{t_{ID(i)}} \dot{V}_i u(t) F^g_M(t)\, \mathrm{d}t \Big), \end{aligned} \tag{5.10}$$

where $t_{ID(i)}$ is the time, during breath j, at which the inspired volume into compartment i is $\dot{V}_i V_{DC} + v_i V_{DP}$ less than its maximum value, $\dot{V}_i V_T(j)$. Thus $t_{ID(i)}$ is defined by the expression

$$\dot{V}_i V_T(j) - (\dot{V}_i V_{DC} + v_i V_{DP}) = \int_0^{t_{ID(i)}} \dot{V}_i u(t) \mathrm{d}t, \tag{5.11}$$

here we assume $\dot{V}_i = \Delta V_i$. If $\dot{V}_i V_T(j) < (\dot{V}_i V_{DC} + v_i V_{DP})$, then $t_{ID(i)} = 0$. The j we sum over in (5.10) are all the conventionally defined inspirations within modified breath m (typically one but sometimes more).

Expired gas volume

To estimate $V_{Ei}(m)$ we apply the Haldane transformation to each compartment [43], and assume that the volume of nitrogen leaving each compartment over the period is equal to the volume entering. Thus,

$$V_{Ei}(m)(1 - F^{\mathrm{CO}_2}_{Ai}(m) - F^{\mathrm{O}_2}_{Ai}(m)) = V_{Ii}(m)(1 - F^{\mathrm{O}_2}_{Ii}(m) - F^{\mathrm{O}_2}_{Ii}(m)). \tag{5.12}$$

We note that the Haldane relation will not apply during the high O_2 breathing phase, and thus this model cannot be applied to that phase of the procedure.

Having calculated $V_{Ii}(m)F^g_{Ii}(m)$ and $V_{Ei}(m)$ we can solve (5.8) (in each compartment) and (5.9) for both gases to get estimates for the venous concentrations and average gas fractions in each compartment for each breath. This is done using `fsolve` in MATLAB.

5.2.2.3 Variation of venous concentration

Figure 5.6 shows an example of the behaviour of the venous gas concentrations, estimated as described in Section 5.2.2.2, during the air breathing period. In this case, and in many other data sets we collected, $C^{\mathrm{CO}_2}_{\bar{v}}(t)$ shows a steady decline over the period, and $C^{\mathrm{O}_2}_{\bar{v}}(t)$ a very slow rise. This is the hypocapnia we described in Section 5.2.2.

To model this we fit a smoothing function (a straight-line) to the estimates of $C^{\mathrm{CO}_2}_{\bar{v}}(t)$ and $C^{\mathrm{O}_2}_{\bar{v}}(t)$ found using the continuous model. The slope of this line then defines the rate

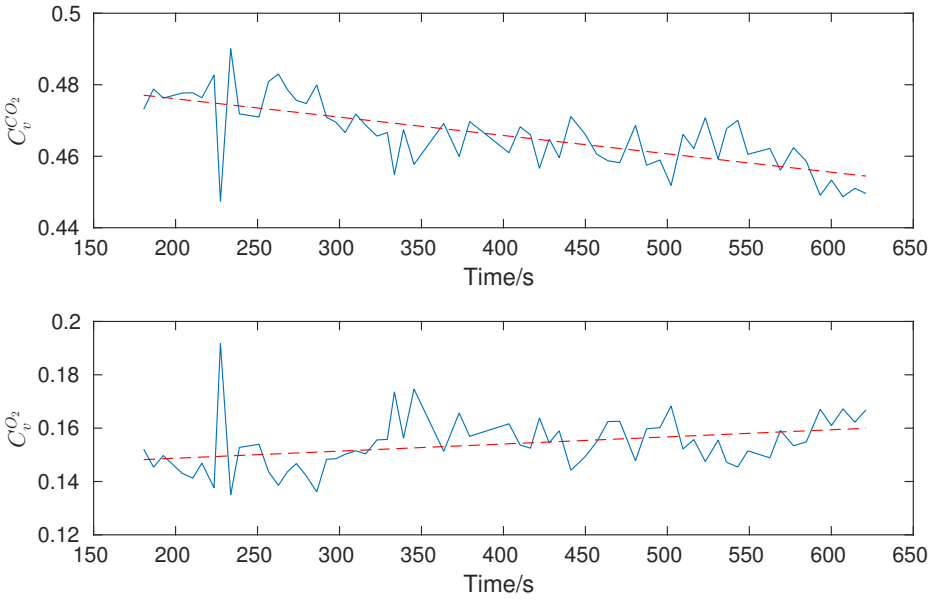

Figure 5.6: An example of the variation of the venous gas concentrations estimated using the continuous ventilation model, as described in Section 5.2.2.2. The dashed line is a straight line fitted to this data.

of change of the venous gas concentrations during the air-breathing period. To validate this method we simulated data using the full model, with known slopes of $C_{\bar{v}}^{\mathrm{CO}_2}(t)$ and $C_{\bar{v}}^{\mathrm{O}_2}(t)$, and found this method estimated these slopes correctly.

As there is no change in inspired CO_2 fraction, we assume the same slope for $C_{\bar{v}}^{\mathrm{CO}_2}(t)$ for the entire nitrogen washout procedure. How we estimate $C_{\bar{v}}^{\mathrm{O}_2}(t)$ during the high O_2 breathing phase is discussed in Section 5.2.3. We next discuss how we use the full model to get initial values for $C_{\bar{v}}^{\mathrm{CO}_2}(t)$ and $C_{\bar{v}}^{\mathrm{O}_2}(t)$.

5.2.2.4 Estimation of starting venous concentrations

During the air breathing period we assume that $C_{\bar{v}}^{\mathrm{O}_2}(t)$ and $C_{\bar{v}}^{\mathrm{CO}_2}(t)$ change linearly at the rate found in Section 5.2.2.3, but we have not yet described how we find their initial values. To find the initial values of $C_{\bar{v}}^{\mathrm{O}_2}(t)$ and $C_{\bar{v}}^{\mathrm{CO}_2}(t)$ we find the values that match the experimentally measured and simulated (using the full model) net volumes of O_2 and CO_2 passing through the mouth. We use the full rather than continuous model, because values for these concentrations estimated with the two models differ slightly; and if we use the continuous model values when we simulate data with the full model, the simulated and measured volumes of CO_2 and O_2 passing through the mouth do not quite match.

Formally we use `fsolve` to find the initial values of $C_{\bar{v}}^{\mathrm{CO}_2}(t)$ and $C_{\bar{v}}^{\mathrm{O}_2}(t)$ which give zeros in the vector function,

$$R_1 = \left[\Delta t \sum_l u(t_l)\left(F_M^{\mathrm{CO}_2}(t_l) - F_S^{\mathrm{CO}_2}(t_l)\right), \Delta t \sum_l u(t_l)\left(F_M^{\mathrm{O}_2}(t_l) - F_S^{\mathrm{O}_2}(t_l)\right)\right], \qquad (5.13)$$

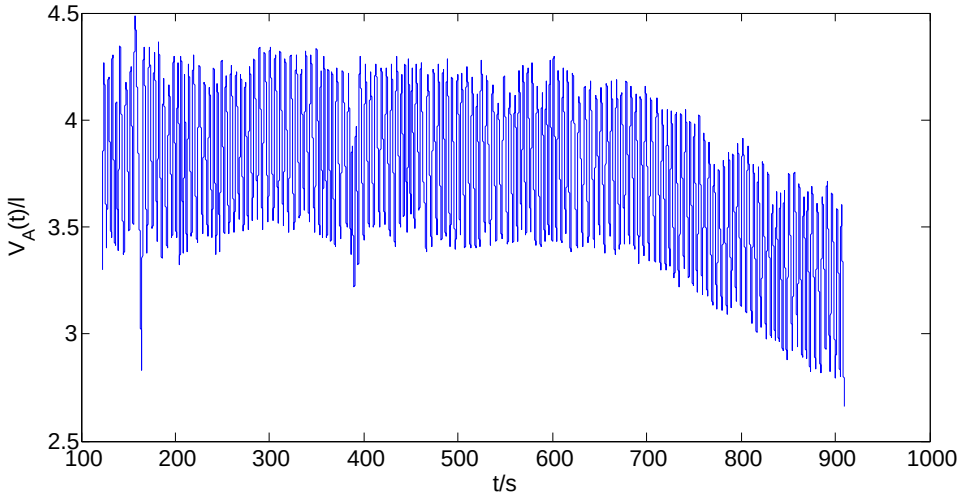

Figure 5.7: A plot of the behaviour of the alveolar volume if the venous O_2 concentration is assumed to continue varying linearly during the high O_2 breathing phase (which starts at $t = 600s$). Following the start of the washout the alveolar volume declines sharply.

where we sum over each time point (t_l) in the last five minutes of the air breathing period. To check the robustness of this procedure we tried a variety of initial guesses for the venous gas concentrations and ensured the same initial values of $C_{\bar{v}}^{CO_2}(t)$ and $C_{\bar{v}}^{O_2}(t)$ were found.

5.2.3 Venous O_2 concentration during high O_2 breathing

In Section 5.2.2.4 we assumed that the $C_{\bar{v}}^{O_2}(t)$ varied linearly during the air breathing period. However, if we continue this assumption into the high O_2 breathing phase, the rise in the arterial O_2 concentration, and consequent increase in O_2 exchange with the blood, causes the lung volume to decline rapidly (see Figure 5.7). Instead we use the simple model described in Section 4.3.2 to allow $C_{\bar{v}}^{O_2}(t)$ to rise during the high O_2 breathing phase, so that the volume of O_2 exchanged with the blood falls sufficiently to achieve a stable alveolar volume.

As we do not have a direct experimental measurement of alveolar volume, we assume that it continues to track the smoothing spline during the high O_2 breathing phase. We find the value of V_L (the volume of distribution for venous O_2) that minimises the objective function,

$$R_2 = \sum_j \left(V_{\text{sp}}(t_{Ej}) - (V_A(t_{Ej}) - V_A)\right)^2, \tag{5.14}$$

where t_{Ej} are the end expiratory times during the high O_2 breathing phase; $V_{\text{sp}}(t_{Ej})$ is the difference between the local minimum at t_{Ej} on the tidal volume record and the smoothing spline at that point; $V_A(t_{Ej})$ is the total alveolar volume of the model at t_{Ej}; and V_A is the total alveolar volume at FRC.

Performing this minimisation using the full model takes approximately four minutes. As this is performed every time we evaluate the overall objective function, (5.1), this

would quickly become very computationally expensive.

However, the model for $C_{\bar{v}}^{O_2}(t)$, introduced in Section 4.3.2, is a low order one, and our objective function, (5.14), is based on a qualitative idea of what the behaviour of the alveolar volume should be rather than explicit measurements. Thus we use the more approximate method explained below, and demonstrate that it produces qualitatively correct behaviour.

The first step in our approach is to run the full model with $V_L = \infty$ to estimate $V_A(t)$ when $C_{\bar{v}}^{O_2}(t)$ is constant during the high O_2 breathing phase. Next we use a simple, single compartment model, to estimate the volume of O_2 exchanged in the alveoli on each breath. We then find the value of V_L which reduces the volume of O_2 exchanged sufficiently to give the correct behaviour for $V_A(t)$.

In this simple model we assume that the average CO_2 and O_2 fractions in the alveoli during expiration are equal to the end tidal values for that breath, and that the values during inspiration ($F_{Ain}^g(j)$) are,

$$F_{Ain}^g(j) = \frac{1}{2}(F_{ET}^g(j) + F_{ET}^g(j-1)). \tag{5.15}$$

Thus integrating (2.3) separately for both expiration and inspiration, and summing the results, the volume of O_2 exchanged during breath j, $V^{O_2}(j)$, is

$$\begin{aligned} V^{O_2}(j) = Q_T \Big(& T_{I(j)} \left(C_a^{O_2}(F_{Ain}^{CO_2}(j), F_{Ain}^{O_2}(j)) - C_{\bar{v}}^{O_2}(t_I) \right) \\ & + T_{E(j)} \left(C_a^{O_2}(F_{ET}^{CO_2}(j), F_{ET}^{O_2}(j)) - C_{\bar{v}}^{O_2}(t_E) \right) \Big), \end{aligned} \tag{5.16}$$

where $T_{I(j)}$ and $T_{E(j)}$ are the lengths of the j^{th}inspiration and expiration respectively.

This is first calculated with constant $C_{\bar{v}}^{O_2}(t)$ to give an estimate for the volume of O_2 consumed by this model with $V_L = \infty$. Next we solve (4.170), the governing equation for venous O_2 concentration during the high O_2 breathing phase, using the above assumptions for average alveolar fractions, to estimate the venous O_2 concentrations with finite values of V_L. We can then calculate the volume of O_2 exchanged on each breath, $V^{O_2}(j)$, for finite V_L.

The reduction in the volume of O_2 exchanged on breath j for a given value of V_L, $V^{O_2}(j, V_L)$, is equal to the difference between: (i) $V^{O_2}(j)$ calculated with $V_L = \infty$; and (ii) $V^{O_2}(j)$ calculated with the given value of V_L. We assume that the alveolar volume of the full model during the high O_2 breathing phase is the sum of the value found with $V_L = \infty$, $V_A^\infty(t_{Ej})$, and the cumulative estimated reduction in O_2 exchange for a given value of V_L, $V^{O_2}(j, V_L)$ from the start of the washout to breath j, thus

$$V_A(t_{Ej}) = V_A^\infty(t_{Ej}) + \sum_{k=1}^{j} V^{O_2}(k, V_L). \tag{5.17}$$

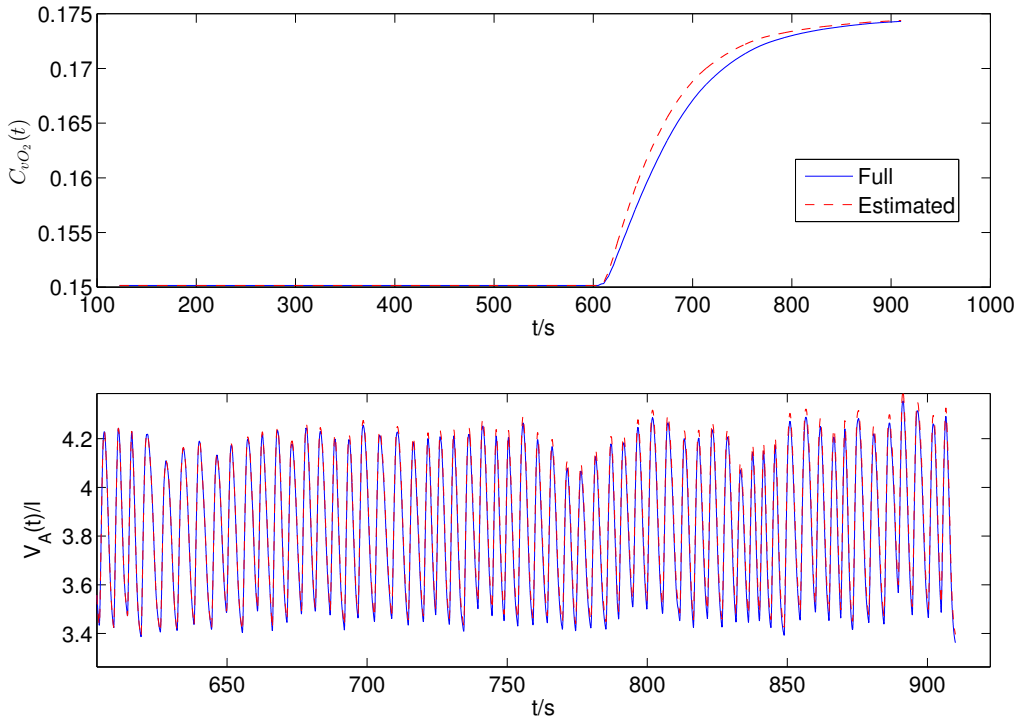

Figure 5.8: The effect on $C_{\bar{v}}^{\mathrm{O}_2}(t)$ (top plot) and $V_A(t)$ (bottom plot) of the method by which find an appropriate value of V_L. 'Full' refers to use of the full model and 'Estimated' to the method using the single compartment model.

Substituting this expression into (5.14) we have,

$$R_2 = \sum_j \left(V_{\mathrm{sp}}(t_{Ej}) - (V_A^{\infty}(t_{Ej}) + \sum_{k=1}^{j} V^{\mathrm{O}_2}(k, V_L) - V_A) \right)^2 . \tag{5.18}$$

We can then find the value of V_L that minimises this modified objective function. This method is approximately ten times faster than using the full model, and gives slightly different values for V_L. Thus, as shown in the top plot of Figure 5.8, the form of $C_{\bar{v}}^{\mathrm{O}_2}(t)$ during the pure O_2 breathing phase is slightly different. However, when we compare the behaviour of the alveolar volume in the full model with values of V_L found using the full and estimated methods, we find it is very similar, as shown in the bottom plot of Figure 5.8. Thus given the difference in computation times, we adopt the approximate approach.

Finally, having found good initial conditions, venous concentrations during the air breathing period, and a value for V_L, we can simulate the full nitrogen washout procedure, and evaluate (5.1). In total one evaluation of the objective function takes approximately two and a half minutes.

5.3 Synthetic data

In Chapter 1 we outlined our proposal for a novel lung function test, which involves estimating the parameters of the lung model we developed in Chapter 4 from experimentally measured nitrogen washout procedures. In the previous section we outlined, in some detail, our parameter estimation procedure. In this section we test this procedure using data created with our model ("synthetic data"). The advantage of using synthetic data is that we know the values of the underlying parameters (Table 5.2). We are interested in several questions.

1. Are the parameters of our model estimable "in principle" (if there was no experimental error) from a nitrogen washout procedure?

2. If some of our parameters are not estimable, can we set them to physiologically reasonable values without influencing the estimation of the parameters that are identifiable?

3. What number of compartments should we use when estimating parameters?

4. If we include realistic experimental noise what is the level of uncertainty in our parameter estimates?

In this section we consider the first three of these questions using error free synthetic data. The incorporation of realistic experimental noise and its consequences are discussed in Section 5.4. First we describe how we created the synthetic data. Next we look at the variation of the objective function with the parameters of our model. We then consider what this suggests about which parameters are identifiable from a nitrogen washout procedure, and what the influence of the parameters that are not identifiable is on the parameters that are identifiable. Finally, we test the performance of the full parameter estimation procedure and briefly consider the effect of the number of compartments that we use.

5.3.1 Creation of synthetic data

To create noise free synthetic data we assume a full set of parameters for the lung model (Table 5.2). We then simulate a full nitrogen washout procedure, using an experimentally measured flow pattern and inspired gas fractions (venous gas concentrations are estimated as described in Section 5.2). The flow pattern and inspired gas fractions of the synthetic data are then set equal to the measured values used as inputs to the model, and the expired fractions to the simulated values. Here we used $N_s = N_t = N_d = 5$ (a total of 125 compartments). The issue of the number of compartments that we use in our model is explored in Section 5.3.3.

Symbol	Value	Description
V_A	2.85 l	Total alveolar volume when the lungs are at FRC
V_{DC}	0.155 l	Common deadspace volume
V_{DP}	0.135 l	Total personal deadspace volume
σ_x	0.45	Width of the compliance distribution
σ_y	1.01	Width of the perfusion distribution
ρ	0.8	Covariance of the compliance and perfusion
σ_z	0.40	Width of the deadspace distribution
Q_T	$5.76\,\mathrm{l\,min^{-1}}$	Cardiac output
V_t	0.25	CO_2 equivalent tissue volume

Table 5.2: A summary of the parameters used in the creation of synthetic data. They are similar to those found in Section 4.1.1, and thus are assumed to be representative of healthy subjects. Producing one such estimate takes approximately five hours.

5.3.2 Parameter identifiability in noise free data

A necessary (but not sufficient) condition for estimating the correct values of the parameters is that the objective function has a minimum value at the true values of the parameters. Figure 5.9 shows the variation of the objective function with each of the underlying parameters (the value of the common deadspace is set equal to the measured deadspace of the LGA, and so we do not consider it as a parameter we have to estimate) while the parameters not being varied are held at their true values.

Each plot in the figure shows a clear minimum at the true value of the parameter, although there appear to be a number of local minima in the plot of varying ρ. We also note that the minima at the true values of the parameters are not exactly at zero. The reason for this is that when we produced the synthetic data we estimated venous concentrations using experimentally measured data, but in the simulations shown in Figure 5.9 we use the synthetic data for this (as part of the evaluation of the objective function). This leads to small differences in the venous gas concentrations and their slopes, which cause the small non zero value of the objective function at the true values of the parameters.

The size of the variation of the objective function with the parameters in the top four plots (V_{DP}, σ_z, V_A, and σ_x) is more than an order of magnitude larger than the size of the variation with the parameters in the bottom four plots (σ_y, ρ, V_t, and Q_T). Thus, as we will show below, small errors in the estimation of one of the former four parameters, or experimental noise, can obscure this small signal and prevent the accurate estimation of the latter four parameters.

In the rest of this thesis we will refer to the first four parameters as the "airway parameters", as they determine the volumes and ventilations of the alveolar and deadspace compartments. The latter four are referred to as the "circulatory parameters", as they relate to perfusion or, in the case of V_t, to CO_2 in lung tissue, rather than directly to the airspaces of the lungs.

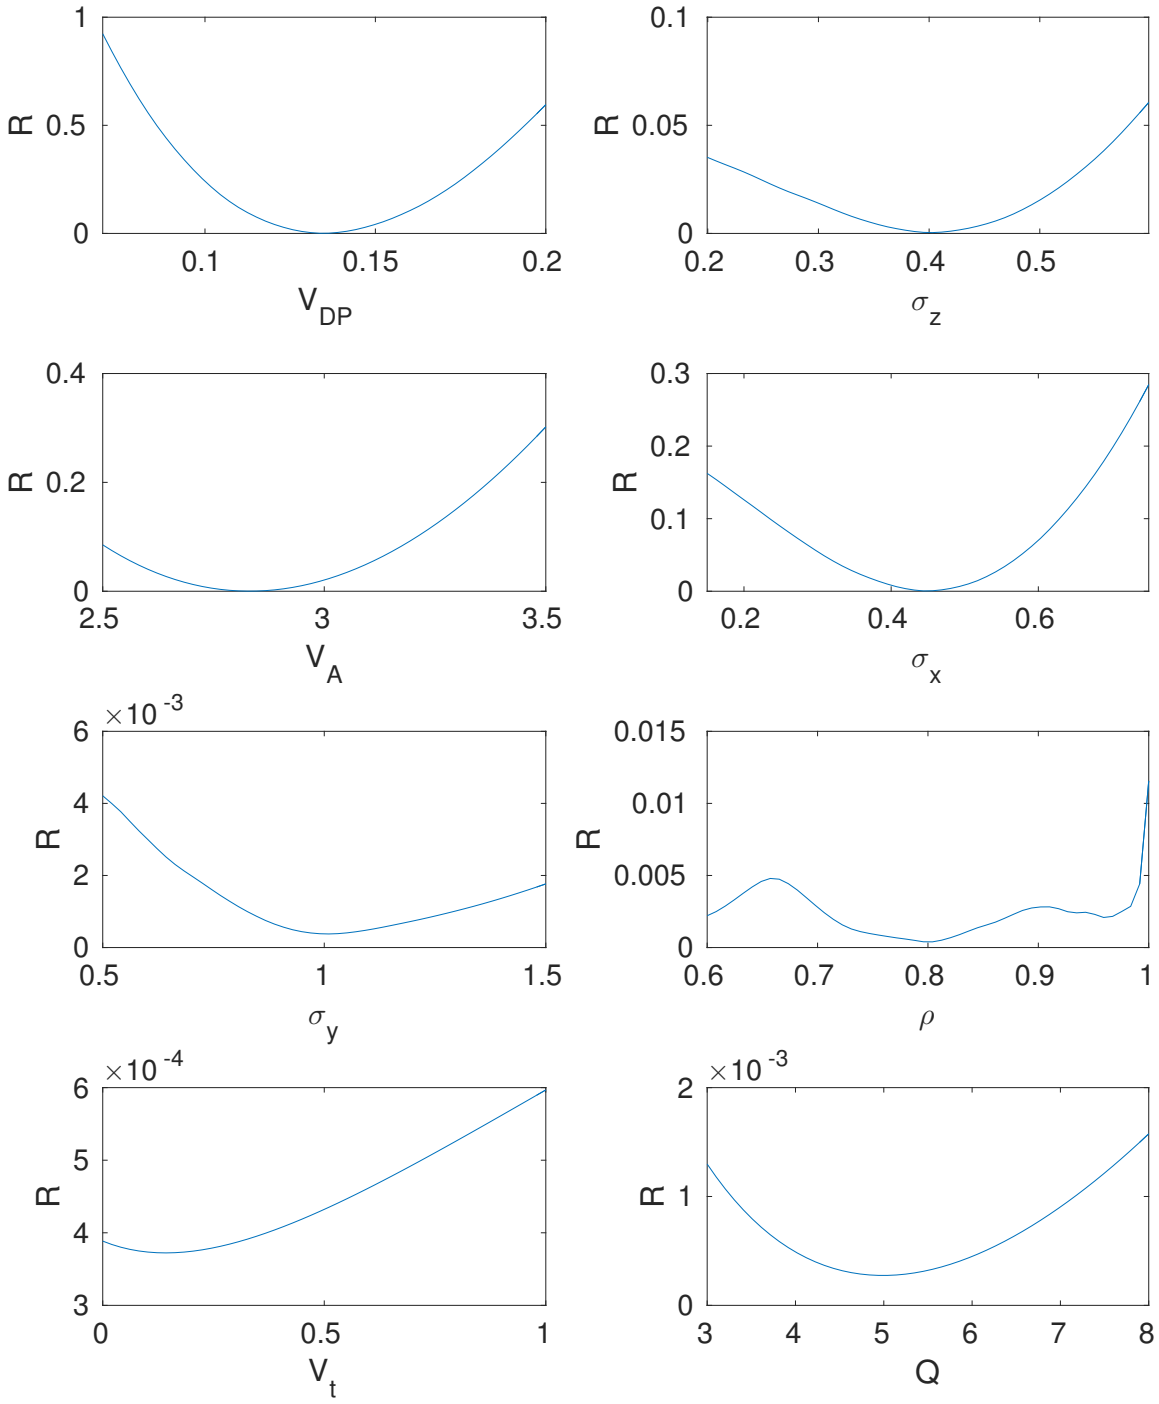

Figure 5.9: The variation of the value of the objective function (calculated as described in Section 5.2) with each parameter. In each case, one parameter is varied and the values of the others are set to the values given in Table 5.2.

When estimating the parameters, however, we are not estimating one parameter at a time, but all eight simultaneously. Thus it is useful to try and understand the dependency of the parameters. Here we define the dependence of two parameters as how the variation of one parameter changes the value of another parameter at which the objective function has its minimum value.

With eight parameters it is extremely difficult to visualise, and computationally intensive to explore, the parameter space comprehensively. Thus to explore the dependence of the parameters, we plot the variation of the objective function as two parameters are varied simultaneously in Figure 5.10. Even this figure is computationally intensive to produce; each of the 28 plots comprises 2,500 points each of which take 2.5 minutes to evaluate. Thus if we performed the simulations in series on one computer, it would take approximately 17 weeks to produce this figure. Evaluated on 28 nodes of Arcus-B it takes approximately 15 hours, demonstrating the value of using this resource.

Examining this figure we again see that V_{DP} has the largest influence on the value of the objective function, and almost irrespective of the value of the other parameters, the minimum for the objective function is at the same value of V_{DP}. σ_x, σ_z and V_A all show more dependence, for example a higher value of σ_z means the minimum in σ_x tends to be at a lower value (increasing either parameter increases the inhomogeneity of the lungs). Nevertheless, although not visible on this scale, the global minimum when any two of these parameters are varied is located at the true value of the parameters.

However, the most important thing we note here is that the values of the circulatory parameters (σ_y, ρ, V_t, and Q_T), which as we noted previously may be difficult to estimate, do not seem to affect the location of the minimum in the objective function as the airway parameters (V_{DP}, σ_z, V_A, and σ_x) are varied. This suggests that even if we can't correctly estimate the values of the circulatory parameters, this will not affect our ability to estimate the values of the airway parameters.

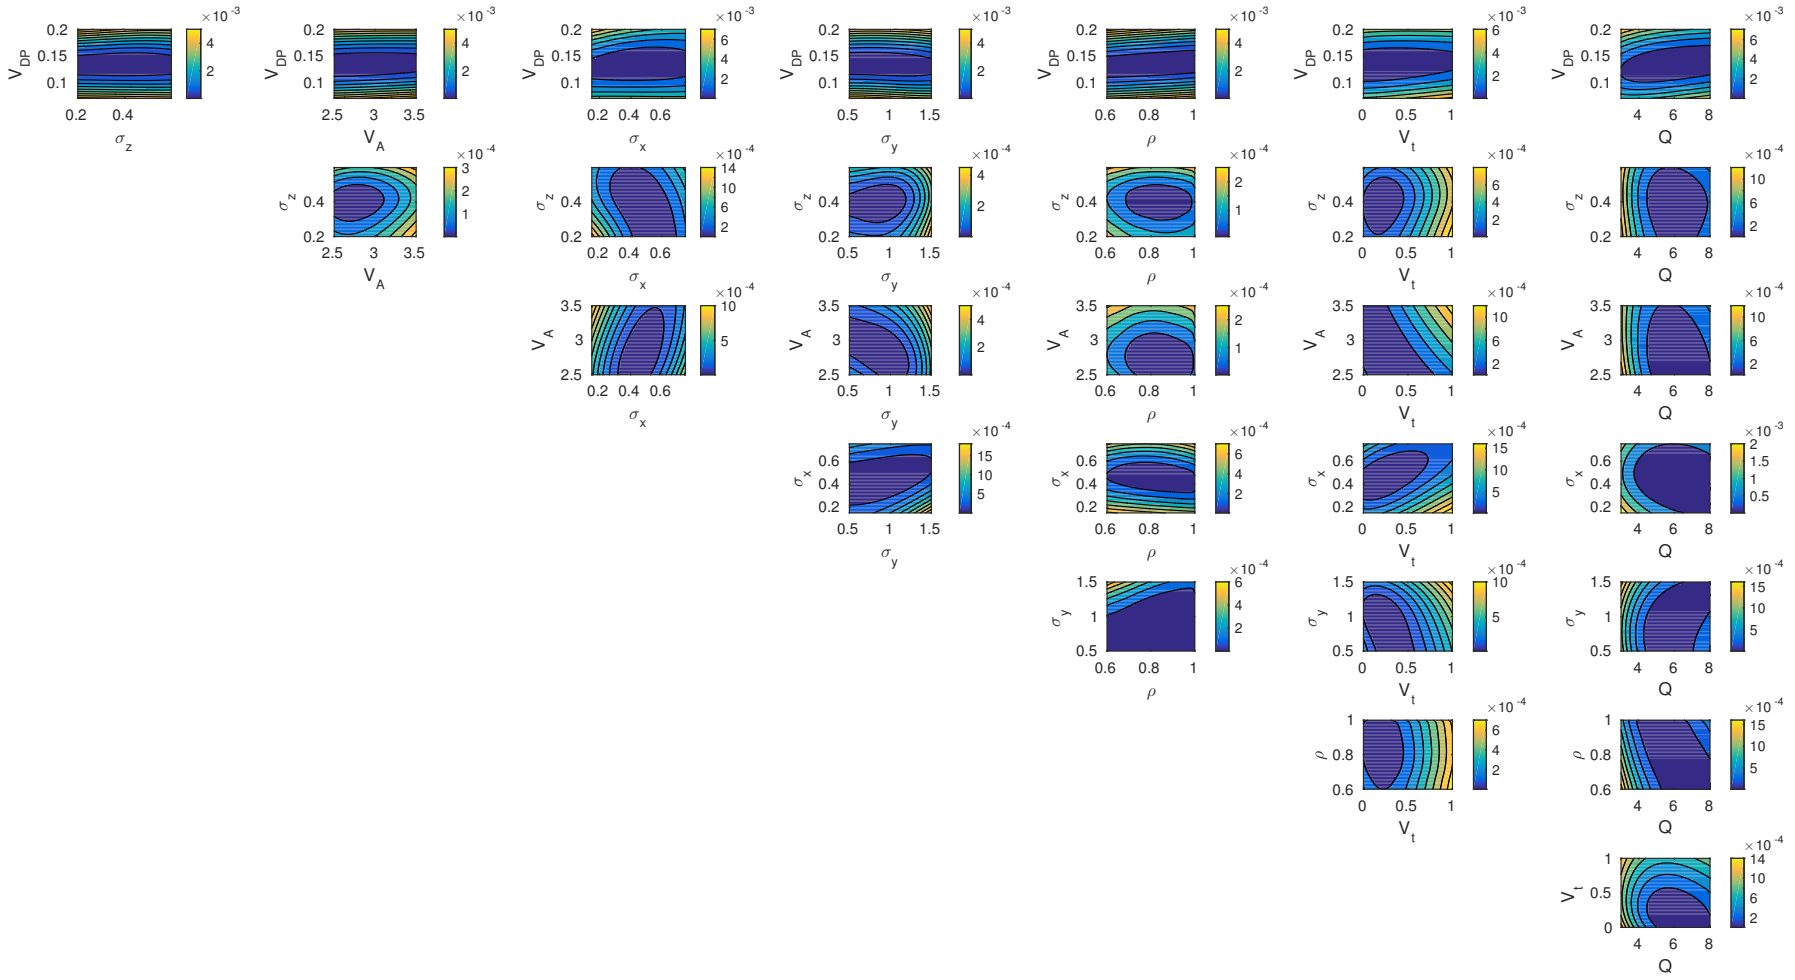

Figure 5.10: The variation of the value of the objective function (calculated as described in Section 5.2) while two parameters are varied simultaneously. In each case, while pairs of parameters are varied the values of the others are set to the values shown in Table 5.2. In this figure blue represents low values of the objective function and yellow high values.

Parameter	True Value	Mean	Standard Deviation	Coefficient of variation
V_{DP}/L	0.135	0.135	2×10^{-4}	0.15%
σ_z	0.40	0.40	0.004	0.97%
V_A/L	2.85	2.84	0.01	0.35%
σ_x	0.45	0.44	0.007	1.48%
σ_y	1.01	0.92	0.09	9.50%
ρ	0.8	0.89	0.08	9.08%
Q_T/Lmin^{-1}	5.76	6.57	0.64	9.74%
V_t	0.25	0.27	0.1	37.33%

Table 5.3: A summary of the results of ten attempts at parameter estimation on noise free synthetic data. The range of initial guesses for the parameters was the range over which they were varied in Figure 5.9.

To test our hypothesis that the airway parameters should all be estimable from noise free data, irrespective of whether correct values are found for the circulatory parameters, we performed ten parameter estimation procedures on the synthetic data, with varying initial guesses for the parameters. The range of initial guesses for the parameters was the range over which they were varied in Figure 5.9.

The results are displayed in Table 5.3. The values of V_A, V_{DP}, σ_x, and σ_z are estimated with a high degree of precision. However, while the true values of σ_y, ρ, V_t, and Q_T are (just) within the mean estimate plus the standard deviation of the relevant parameter, the standard deviations of the estimates are about an order of magnitude larger than for the airway parameters suggesting that the circulatory parameters cannot be estimated as precisely.

The analysis in Table 5.3 was performed with the default value (10^{-6}) for the stopping criteria `TolFun` on the optimisation function (`fmincon`). This criteria is the minimum change in the value of the objective function that must occur between successive iterations for the optimisation function to continue iterating. Lowering this value did not meaningfully change the above results (although it did increase the computation time), and thus in the rest of this thesis we use the default value.

5.3.3 Number of compartments

As we described in Chapter 4, our model involves approximating continuous distributions with $N_s \times N_t \times N_d$ compartments. Thus an obvious question is, how many compartments do we need to use to get a good approximation?

We saw in Figure 4.16 that increasing the number of compartments decreased the computation speed. Since most of the computation time during the parameter estimation procedure is spent evaluating the model, halving the computation speed will roughly double the length of time it takes to estimate parameters. With $5 \times 5 \times 5$ compartments

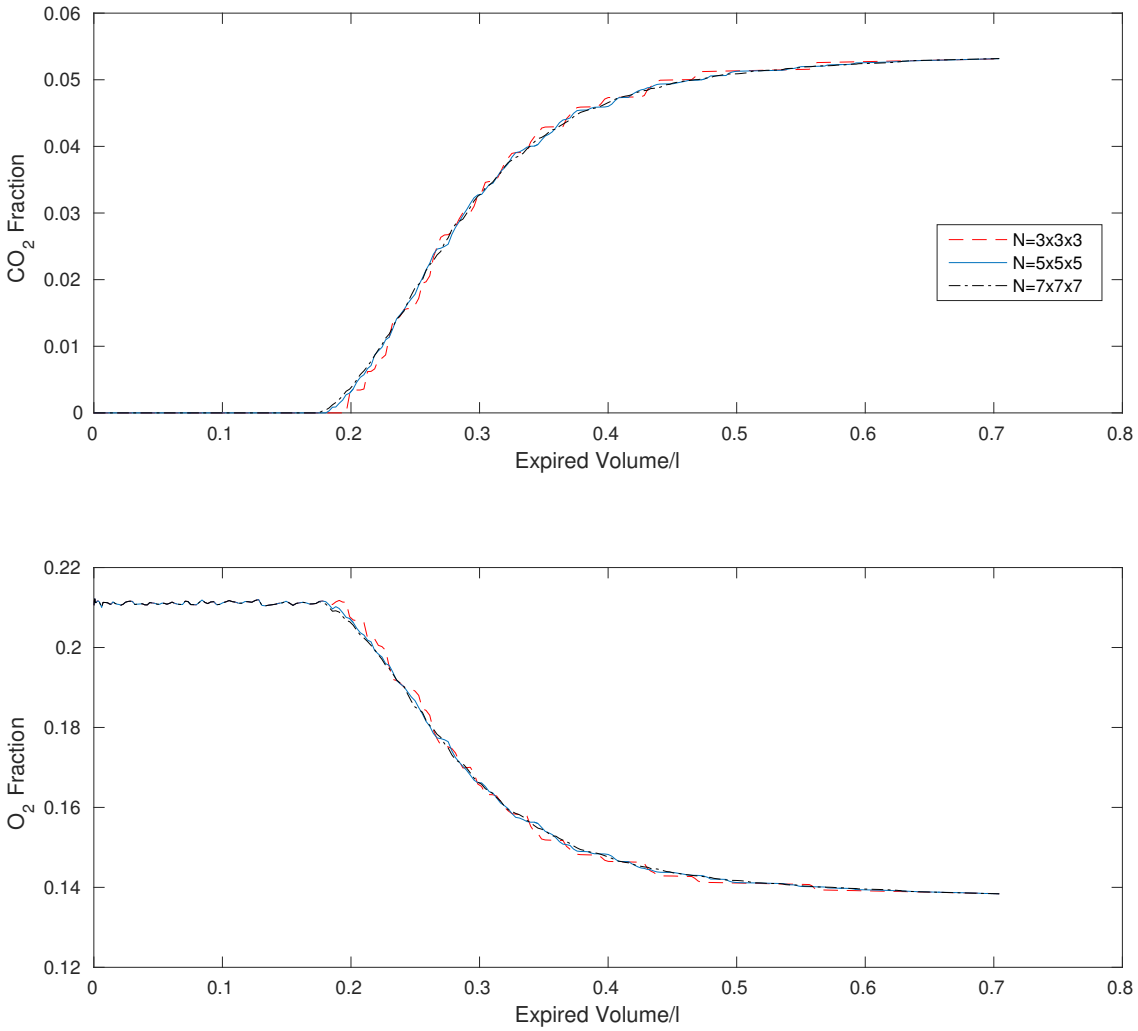

Figure 5.11: A simulated expirogram of a single breath during normal breathing with varying number of compartments.

Parameter	True Value	$5 \times 5 \times 5$	$7 \times 7 \times 7$	$9 \times 9 \times 9$
V_{DP}/L	0.135	0.134	0.135	0.135
σ_z	0.40	0.40	0.42	0.42
V_A/L	2.85	2.84	2.82	2.81
σ_x	0.45	0.44	0.44	0.44
σ_y	1.01	0.92	0.95	0.99
ρ	0.8	0.89	0.86	0.86
Q_T/Lmin^{-1}	5.76	6.67	7.79	8.07
V_t	0.25	0.22	0.35	0.30

Table 5.4: A summary of results of simulating data with $5 \times 5 \times 5$, $7 \times 7 \times 7$, and $9 \times 9 \times 9$ compartments, and estimating parameters with $5 \times 5 \times 5$ compartments.

parameter estimation takes around five hours, and thus there is a computational incentive to minimise the number of compartments we use.

Figure 5.11 shows the effect of using different numbers of compartments on the expired CO_2 and O_2 traces produced by our model. In this figure we can see that while moving from $3 \times 3 \times 3$ to $5 \times 5 \times 5$ compartments makes a noticeable difference to the solution, the difference between the $5 \times 5 \times 5$ solution and the $7 \times 7 \times 7$ solution is relatively small.

To check the effect of this small difference in the forward solution on our parameter estimation procedure, we simulated synthetic data using our model with $7 \times 7 \times 7$ and $9 \times 9 \times 9$ compartments. We then attempted to estimate the parameters using $5 \times 5 \times 5$ compartments. Table 5.4 shows the results of these parameter estimations.

Here we can see that we can recover the parameters that we believe are estimable from our procedure (V_A, V_{DP}, σ_x, and σ_z) accurately using a $5 \times 5 \times 5$ model even when the underlying data was produced with a larger number of compartments. This suggests that even though our conceptual model of the lungs is that they are made up of a very large number of units that form a continuous distribution, we can use a model with $5 \times 5 \times 5$ (125) compartments to estimate the parameters of the underlying continuous distribution. Thus in the rest of this thesis, unless otherwise specified, we adopt this number of compartments.

5.4 Experimental error

In the previous section we found that our parameter estimation routine provided good estimates of the airway parameters (V_{DP}, σ_x, σ_z, and V_A) from noise free data. However, the circulatory parameters (σ_y, ρ, V_t, and Q_T) could not be reliably found, although this did not affect the estimation of the airway parameters. In this section we explore the effects of experimental noise on our parameter estimation procedure. The steps we take to do this are outlined below.

1. Establish detailed noise models for the LGA and for a theoretical device representative of current (non LGA) devices.

2. Apply these noise models to create "noisy synthetic data".

3. Attempt to estimate parameters from a number of these noisy synthetic datasets.

4. Compare the results of these estimates to the true parameter values (shown in Table 5.2) to get an idea of the uncertainty in our parameter estimates due to experimental noise.

5.4.1 Noise model

In this section we consider appropriate models for experimental noise for both data measured by the LGA, and data measured by a device representative of those currently available. The European Respiratory Society (ERS) recently established criteria that a device for measuring nitrogen washouts should conform to [96]. The authors of the paper note: "It is unlikely that all individual criteria outlined will be fulfilled by any one system", and so we use the maximum error tolerances from this paper as representative of a current (non LGA) device ("ERS Device"). We note, however, several devices exist that exceed a number of these criteria, both for measuring SF_6 [52] and N_2 washouts [142], have been reported in the literature.

The accuracy of experimental data may be limited by [96]:

1. the accuracy of the measured gas fractions;

2. the finite response times of the gas sensors;

3. the accuracy of the flow measurement;

4. temporal misalignment between the flow and gas fraction measurements.

We will discuss how each of these sources of error (noise) is added to noise-free data to produce 'noisy synthetic data' for both devices. Apart from error in the gas fraction measurements, which we consider explicitly, the values used for the LGA were supplied by those developing the system.

5.4.1.1 Accuracy of measured gas fractions

To estimate the error in gas fractions measured with the LGA, we consider the residuals between the expired gas fractions simulated with our model (with fitted parameters) and the experimentally measured fractions. We then find a continuous distribution to approximate these residuals, and use this as a model for the distribution of experimental error. These residuals will include error due to model misfit (the quality of the fit will be

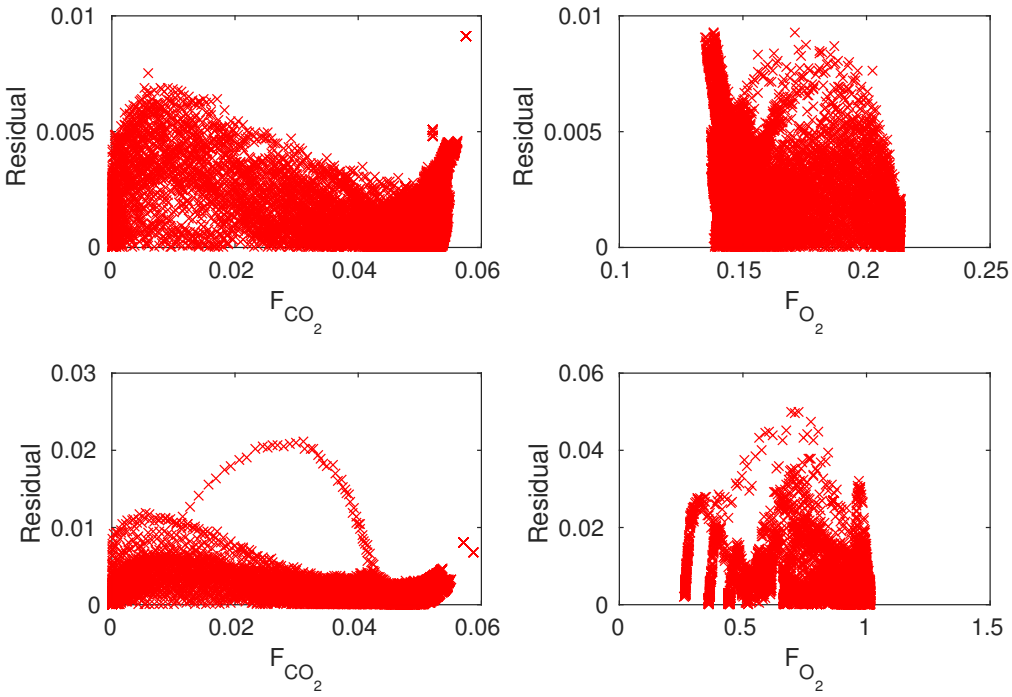

Figure 5.12: Plots of the magnitude of the residuals against the measured gas fraction. The top two plots are from the air breathing period and the bottom two from the high O_2 breathing phase.

discussed in detail in Chapter 6), as well as measurement error, but this will be the case when we run our parameter estimation procedure on experimental data, and so it seems reasonable to include it here.

Figure 5.12 shows the relationship between the size of the measured gas fraction and the residual size for one representative dataset. While there is some structure here, this relationship does not appear to be very strong, and thus we assume the magnitude of the error is independent of the measured fraction.

Histograms of these residuals from five datasets are plotted in Figure 5.13. The air breathing and washout phases are plotted separately. We assume that we can approximate these histograms with normal distributions with zero mean, and standard deviation equal to the standard deviation of the residuals. This assumption, while clearly a simplification, looks reasonable, although there appears to be a larger number of small residuals than you would expect from a normal distribution (particularly for CO_2 where there are a large number very close to zero). There is also some skew in the O_2 residuals during the washout, however the degree and direction of this skew varied depending on the datasets we examined and thus we do not include it here. The standard deviation of the CO_2 residuals (0.0020) is the same during normal breathing and the high O_2 breathing phase, however for O_2 it is about two to three times bigger during the high O_2 breathing phase (0.0026 and 0.0068 respectively). Thus to include the error in the measured gas fractions, we sample from a normal distribution with the appropriate standard deviation, and add

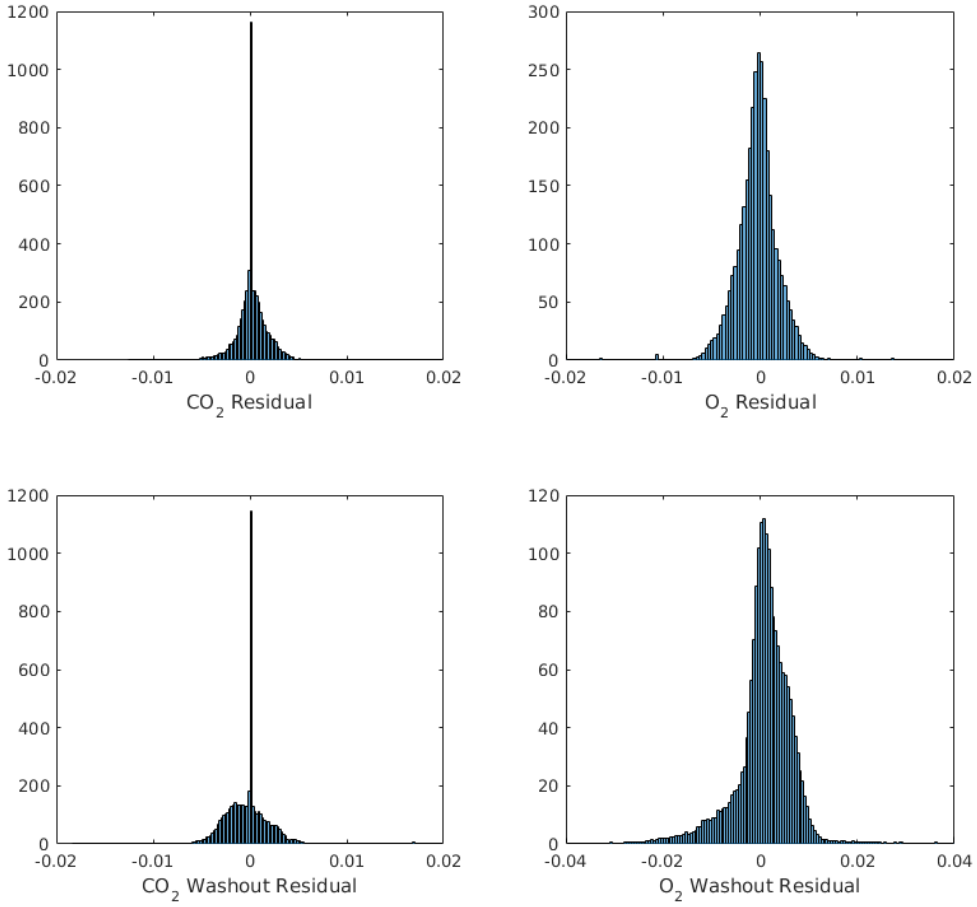

Figure 5.13: Histograms of differences between the fitted synthetic data and the experimentally measured gas fractions. The top two histograms are from the air breathing period and the bottom two from the washout.

the value to the noise free gas fraction.

The ERS device is assumed to be composed of a N_2 detector rather than CO_2 and O_2 detectors. The criteria for the accuracy of such a device is [96]: "Linearity within 1% relative of full scale (e.g. 0-80% is ±0.8% at 80% N_2) to ensure appropriate assessment of starting concentration, and within 5% relative of any lower value (e.g. ±0.25% at 5% N_2)". We note the magnitude of these errors are similar to those we found for the LGA and for simplicity assume the same error model for the fractions measured with both devices.

5.4.1.2 Finite response time

When there is a step change in gas fraction a sensor will take a finite amount of time to respond to that signal [8]. The LGA measures the spectra, from which the gas concentrations are estimated, several times in each 10ms measurement window, and thus this response time should be very small. Here we use 10ms. The ERS specification requires the response time to be at most 100 ms [95], thus we assume this value for the ERS device.

To add this error to the experimental signal we assume a single exponential rise time model. Thus following a step of magnitude F_s the measured fraction, $F(t)$, is [8, 28]

$$F(t) = F_0 + F_s\left(1 - \exp\left(-\frac{t}{t_0}\right)\right), \tag{5.19}$$

where F_0 is the initial gas fraction. We assume that the 100 ms response time is the time taken for the signal to reach 90% of its final value, thus

$$t_0 = \frac{100\text{ms}}{\log(10)}. \tag{5.20}$$

The instrument response function $g(t)$ is the derivative of the step response [8, 28], and thus

$$g(t) = \frac{\exp\left(-\frac{t}{t_0}\right)}{t_0}. \tag{5.21}$$

The measured signal ($F_M(t)$) is then a convolution of the instrument function and the true signal ($f(t)$) [28], thus

$$F_M(t) = \int_{-\infty}^{t} f(u)g(t-u)\ \mathrm{d}u. \tag{5.22}$$

5.4.1.3 Errors in the flow measurements

There are two types of error in the flow measurements, Gaussian noise on the instantaneous flow measurements, and errors in the recorded volume.

The first of these is assumed to have a magnitude of 5% of the measured flow for the ERS device [96] and 0.5% for the LGA. To include this form of noise at each time point we sample from a normal distribution with a mean of one and standard deviation equal to 5% for the ERS device and 0.5% for the LGA, and multiply the noise free flow by this value.

The ERS specifications state that recorded volumes should be accurate to within 3% [96], for the LGA the equivalent specification is 0.25%. We assume that this error is the result of drift in the flow measurement (volumes are the integral of the flow), and that this drift may occur throughout the washout period.

To model this drift we add a Gaussian random walk to the flow measurements with standard deviation (σ). To set the value of σ we set the expected size of the error in the recorded volume (the expected length of the Gaussian random walk, $\sigma\sqrt{n}$ [15], multiplied by the duration of the experiment) equal to that specified by the error tolerance. Thus,

$$\sigma\Delta t n\sqrt{n} = \text{Error Level} \times V_{\text{cell}}, \tag{5.23}$$

where n is the number of time steps, and V_{cell} is the the total volume of gas that has passed through the flow sensor during the experiment.

Type of Noise	ERS Level	LGA Level
Gaussian noise on F^{CO_2}	0.002	0.002
Gaussian noise on F^{O_2} (air breathing)	0.0026	0.0026
Gaussian noise on F^{O_2} (washout)	0.0068	0.0068
Finite response time	100 ms	10ms
Gaussian noise on flow	5%	0.5%
Error on recorded volumes	3%	0.25%
Signal misalignment	10 ms	5 ms

Table 5.5: A summary of the different types of noise on the data, and their magnitude for the LGA and the ERS device. With the exception of the noise on the fractions which we estimate, values for the ERS device are taken from [96] and those for the LGA from our experimental collaborators who developed the device.

5.4.1.4 Temporal misalignment

Finally, as many devices measure flow and concentration in different locations, there may be some temporal misalignment of the signals. The ERS suggest that this should be at most 10 ms, while in the LGA the signals are measured in the same place and thus co-incident within 5 ms. This error was applied by shifting the flow and concentration signals relative to each other by the stated amount.

5.4.1.5 Summary

Table 5.5 gives a summary of the different types of noise and their magnitudes for both the LGA and the ERS device.

Figure 5.14 shows a comparison of the experimentally measured gas fractions and the noisy synthetic version of the same breath. Our decision to use a normal distribution, despite the true residuals having a larger density close to zero than would be expected for such a distribution, means that although the residuals and the experimental noise have the same standard deviation, the synthetic noisy data is more noisy than the experimental data. A second point worth noting is that while the ERS and LGA signal are similar for most of the breath, the greater response time of the ERS device means that sharp changes in the gas fraction (for example just after 524 s in the plot) happen slightly later, and more gradually for this device.

Figure 5.15 compares the effect of the different error sizes in the flow measurements for the LGA and the ERS device. While the flow pattern on any one breath is not that different between the two devices (top plot), the errors on the volume measurements (that we simulated with a Gaussian random walk) can cause large divergences between the ERS and LGA data in the tidal volume record (bottom plot).

It is worth noting that not all devices show volume drift, see for example [142]. In devices where this drift does occur, corrections are sometimes applied [9, 96]. However, we cannot apply such corrections, as we need accurate determination of the O_2 and CO_2

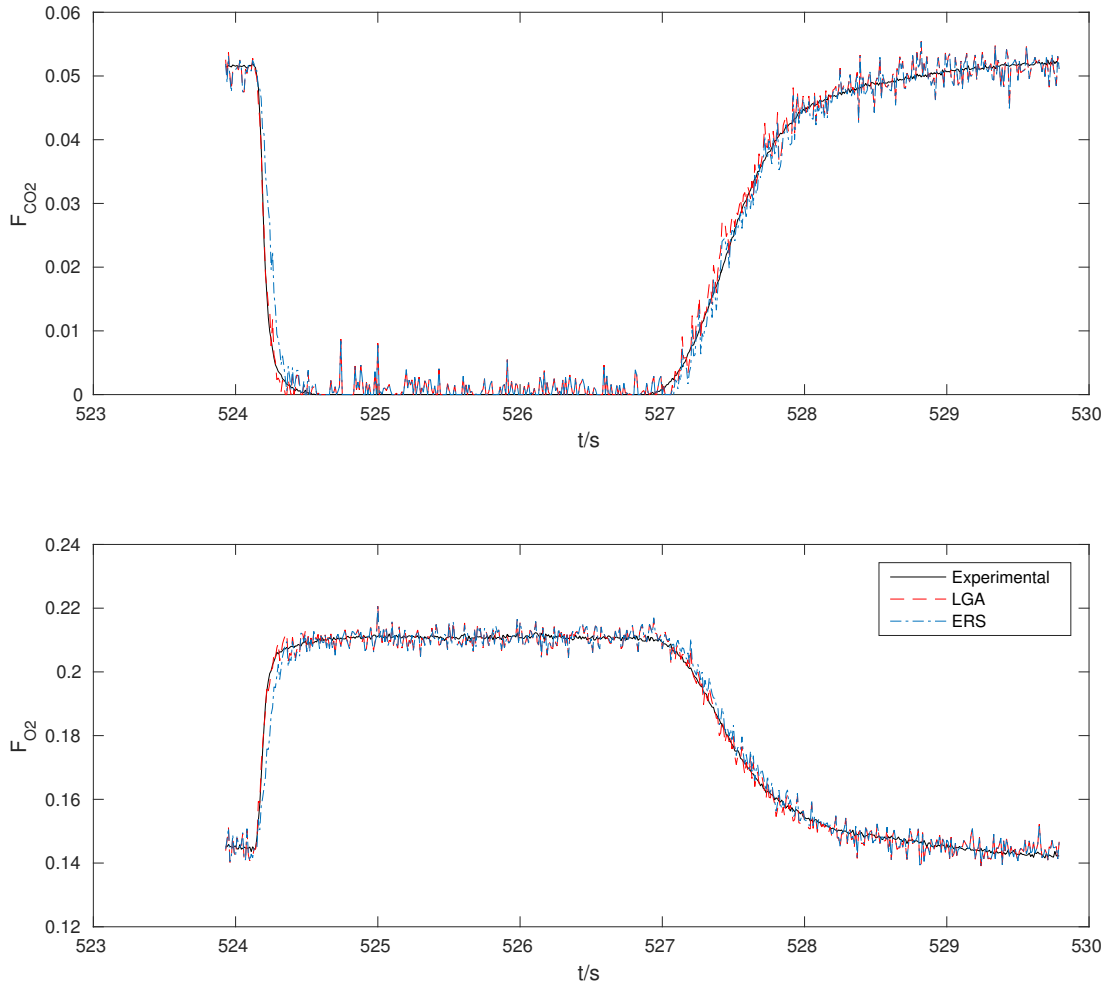

Figure 5.14: A comparison of the experimentally measured data and noisy synthetic fractions for the ERS device and the LGA for a single breath during the air breathing period.

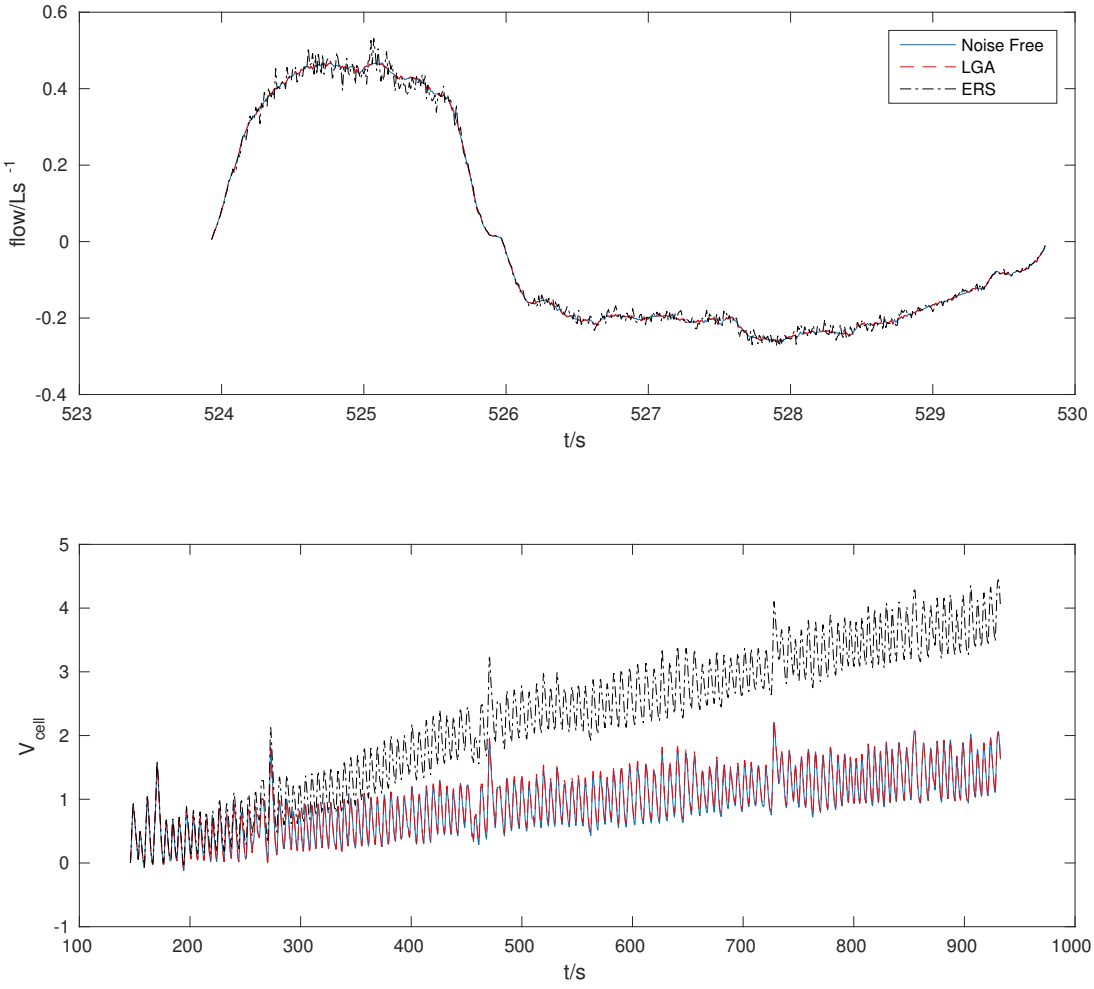

Figure 5.15: A comparison of the effect of the levels of noise on the flow for the LGA and the ERS device. In the top plot, the flow pattern for a single breath is plotted. In the bottom plot, the integral of the flow, the net volume of gas that has passed through the mouth, is shown.

consumption rates (which depend on the difference between inspiratory and expiratory volumes) to estimate the venous gas concentrations (see Section 5.2.2).

In the next section we consider the effect of these two noise models on parameter estimation.

5.4.2 Uncertainty of estimated parameters

In Section 5.3 we found that we could estimate the airway parameters (V_{DP}, σ_x, σ_z, and V_A) from noise free data. The circulatory parameters (σ_y, ρ, V_t, and Q_T) could not be reliably estimated, but this did not prevent the estimation of the airway parameters. In this section we test whether that conclusion holds once we include the effect of realistic experimental error.

To investigate this we created ten datasets (with the same underlying noise free data)

Parameter	True Value	Mean	Standard Deviation	Coefficient of variation
V_{DP}/L	0.135	0.134	8×10^{-4}	0.5%
σ_z	0.40	0.40	0.004	1.39%
V_A/L	2.85	2.81	0.02	0.64%
σ_x	0.45	0.43	0.006	2.93%
σ_y	1.01	1.00	0.16	17.81%
ρ	0.8	0.84	0.11	14.4%
Q_T/Lmin^{-1}	5.76	3.37	1.5	44.51%
V_t	0.25	0.27	0.17	62.96%

Table 5.6: A summary of parameter estimations performed on ten different synthetic datasets with the LGA noise model.

with each of the two noise models. We then attempted to estimate the parameters from each of these datasets. The mean parameter estimates, standard deviations of those estimates, and coefficients of variation (the standard deviation divided by the mean) for the LGA noise model are shown in Table 5.6.

Here we see that we can estimate V_{DP}, σ_x, σ_z, and V_A to within a physically meaningful precision. While the true values of σ_y, ρ, and V_t are within one standard deviation of the mean of the estimates, these standard deviations are large (over 60% of the size of the true value for V_t for example), and thus we conclude that they cannot be estimated accurately from our data. The true value of Q_T is not within the (large) estimated standard deviation of the mean estimated value, and thus we conclude that this parameter also cannot be estimated from our data.

The standard deviations and coefficients of variation give an idea of the uncertainty of our parameter estimates in the presence of experimental error. For the volumes, V_{DP} and V_A, the standard deviations are 0.8 ml and 20 ml respectively. These are small in an absolute sense, and smaller than the variability we might expect from a subject changing position between runs. For σ_x and σ_z, it is difficult to have a physical insight into the size of the standard deviation, but we note that for these parameters the coefficients of variation are 2.9% and 1.4% respectively.

One limitation, as well as the accuracy of the noise model, in the strength of these conclusions is that it is difficult to determine the number of datasets required to quantify the uncertainty accurately. To check the effect of the number of datasets we estimated parameters from five more noisy datasets, and found coefficients of variation of 0.3%, 0.5%, 1.9%, and 1.3% for the estimated values of V_{DP}, V_A, σ_x, and σ_z respectively from these datasets. The similarity of these numbers to those shown in Table 5.6 suggests that our conclusion, that the uncertainty of these parameter estimates due to experimental noise is relatively small (and probably smaller than the biological variability), is valid even if we cannot be certain of the precise size of the uncertainties.

For the ERS noise model our parameter estimation procedure was only able to find a minimum on three of the ten datasets (and the estimated parameters from these three datasets were not close to the true values). This suggests that the criteria set out by the ERS for a nitrogen washout device are not stringent enough to provide data accurate enough for our method of estimating parameters. This of course assumes that our interpretation of these criteria, used to build the noise model, is valid.

5.4.3 Summary

At the start of the section on synthetic data we set out four questions.

1. Are the parameters of our model estimable "in principle" (if there was no experimental error) from a nitrogen washout procedure?

2. If some of our parameters are not estimable can we set them to a physically reasonable value without influencing the estimation of the parameters that are identifiable?

3. What number of compartments should we use when estimating parameters?

4. If we include realistic experimental noise what is the level of uncertainty in our parameter estimates?

In this section we have found (in answer to the first two questions) that the airway parameters of our model (V_{DP}, σ_x, σ_z, and V_A) are identifiable from a nitrogen washout procedure, while the circulatory parameters (σ_y, ρ, V_t, and Q_T) are not. As we only measure data at the mouth, and have no direct information about the concentrations or flows in the blood, this is perhaps not surprising.

Fortunately however, the values of the circulatory parameters (as long as they are in a physiologically reasonable range) do not prevent accurate estimation of the airway parameters, and thus we conclude the circulatory parameters can be set to physiologically reasonable values without preventing the estimation of the airway parameters.

An obvious question, given this lack of identifiability of the circulatory parameters is whether the model of perfusion should be simplified to reduce the complexity (and computation time) of the overall model. However, as discussed in Chapter 7 there are plans to use a more soluble gas (acetylene), which hopefully will improve the identifiability of the circulatory parameters in the model. Thus the decision was made to retain the full model of perfusion to allow the flexibility to attempt to determine these parameters in the future.

Next we found that we could use a 125 compartment model to accurately estimate the parameters of data produced with a much larger number of compartments. Thus we adopt this number of compartments.

Finally, we considered the uncertainty of our parameter estimates in the presence of experimental noise. To do this we developed detailed noise models for a device conforming to the ERS specifications and for the LGA. Our analysis suggested that a device conforming to the ERS specifications would not be sufficiently accurate for our purposes. In the case of the LGA, our analysis suggested that the uncertainty due to experimental error of our estimates of the airway parameters was small, and, at least for the two volumes, smaller than the degree of uncertainty we would expect from biological variability (for example, the variation of position of the subject when breathing on the device).

In the next section we investigate a novel experimental procedure which, in Chapter 6, we use to try and set a physiologically reasonable value for the equivalent CO_2 tissue volume, V_t.

5.5 Tissue volume

We showed in the previous section that we cannot estimate a value for the equivalent CO_2 tissue volume, V_t, from the nitrogen washout procedure. Fortunately, we also showed that the uncertainty in the value of this parameter does not prevent the accurate estimation of the airway parameters. Nevertheless, it would be useful to have a rough idea of the size of this parameter. In this section we investigate an experimental technique we developed to estimate a physiologically reasonable value for this parameter.

Since we assume only CO_2 is soluble enough to be present in significant volumes in the lung tissue, the only equations in which this parameter appears are the governing ODEs for CO_2 in the alveoli, (4.78), and (4.81). Thus to get a signal from this parameter, we require a procedure which causes a large change in the alveolar CO_2 fractions. The obvious thing to try is a CO_2 washin. However, unlike O_2, CO_2 cannot be delivered in large quantities for long periods of time without risk to the subject [23]. Even relatively low quantities delivered for long periods of time have been found to induce anxiety [140].

Thus our CO_2 washin protocol is: ten minutes of breathing room air, followed by approximately 40 s when the inspired CO_2 fraction is 6%, and finally 2 minutes of room air breathing. Other than a slight deepening of breathing following the CO_2 step no ill effects were observed.

5.5.1 The objective function

Since the goal of this procedure is to estimate V_t, we find the parameters that minimise the sum of the squared differences between the measured and simulated CO_2 fluxes during the CO_2 washin. Quantitatively we minimise the objective function, R_{V_t}, where

$$R_{V_t} = \sum_{\mathrm{CO_2\ washin}} u(t_l)^2 \left(F_M^{\mathrm{CO_2}}(t_l) - F_S^{\mathrm{CO_2}}(t_l) \right)^2 . \tag{5.24}$$

As for the nitrogen washout function the evaluation of the objective function is a multi step procedure. To evaluate (5.24) we carry out the following steps.

1. We discard the first two minutes of the air breathing period, during which we assume the subject is settling down to breathing on the LGA.

2. We use the continuous model to estimate the slope of the venous gas concentrations and the starting gas fractions (as in Section 5.2.2).

3. We find starting values of the venous gas concentrations, with the full model, that match the net volumes of CO_2 and O_2 volumes passing through the LGA during the second five minutes of the air breathing period (as in Section 5.2.2.4).

4. We find the correct starting lung volume such that, at the start of the CO_2 washin, the difference between the lung volume and FRC is correct (as judged by a smoothing spline fitted to the tidal volume record, see Section 5.2.1.1).

5. Finally, we simulate the CO_2 washin procedure with the full model and calculate (5.24).

5.5.2 Parameter identifiability and uncertainty

To study parameter identifiability and uncertainty due to experimental noise, we created synthetic data (using the parameters shown in Table 5.2) for this procedure. Figure 5.16 shows the variation of the objective function, (5.24), with each parameter while the others are held constant at the values shown in Table 5.2.

As we saw with the nitrogen washout procedure (Figure 5.9), each curve has a minimum at the true value of the parameter. A difference from Figure 5.9 is the relative magnitudes of the variation with the different parameters. While V_{DP} still seems to have the largest influence on the value of the objective function, the difference between the magnitudes of the variation with the rest of the parameters is not as pronounced. In the nitrogen washout procedure, the effect of V_t on the value of the objective function was negligible in comparison to a number of the other parameters. That is not the case here, suggesting V_t may be identifiable from this procedure.

Figure 5.17 shows the variation of (5.24) with two parameters, while the rest are held constant at their true values. It is again instructive to compare this plot with the equivalent from the nitrogen washout procedure, Figure 5.10. In Figure 5.10 we saw that the values of the four circulatory parameters (σ_y, ρ, V_t, and Q_T) had very little influence over where the minima in the objective function was for the airway parameters (V_{DP}, σ_x, σ_z, and V_A). In the case of the CO_2 washin, it is again true that the values of the other parameters have little effect on the location of the minimum with varying V_{DP}, but all the other parameters show some degree of dependence. This is not surprising given

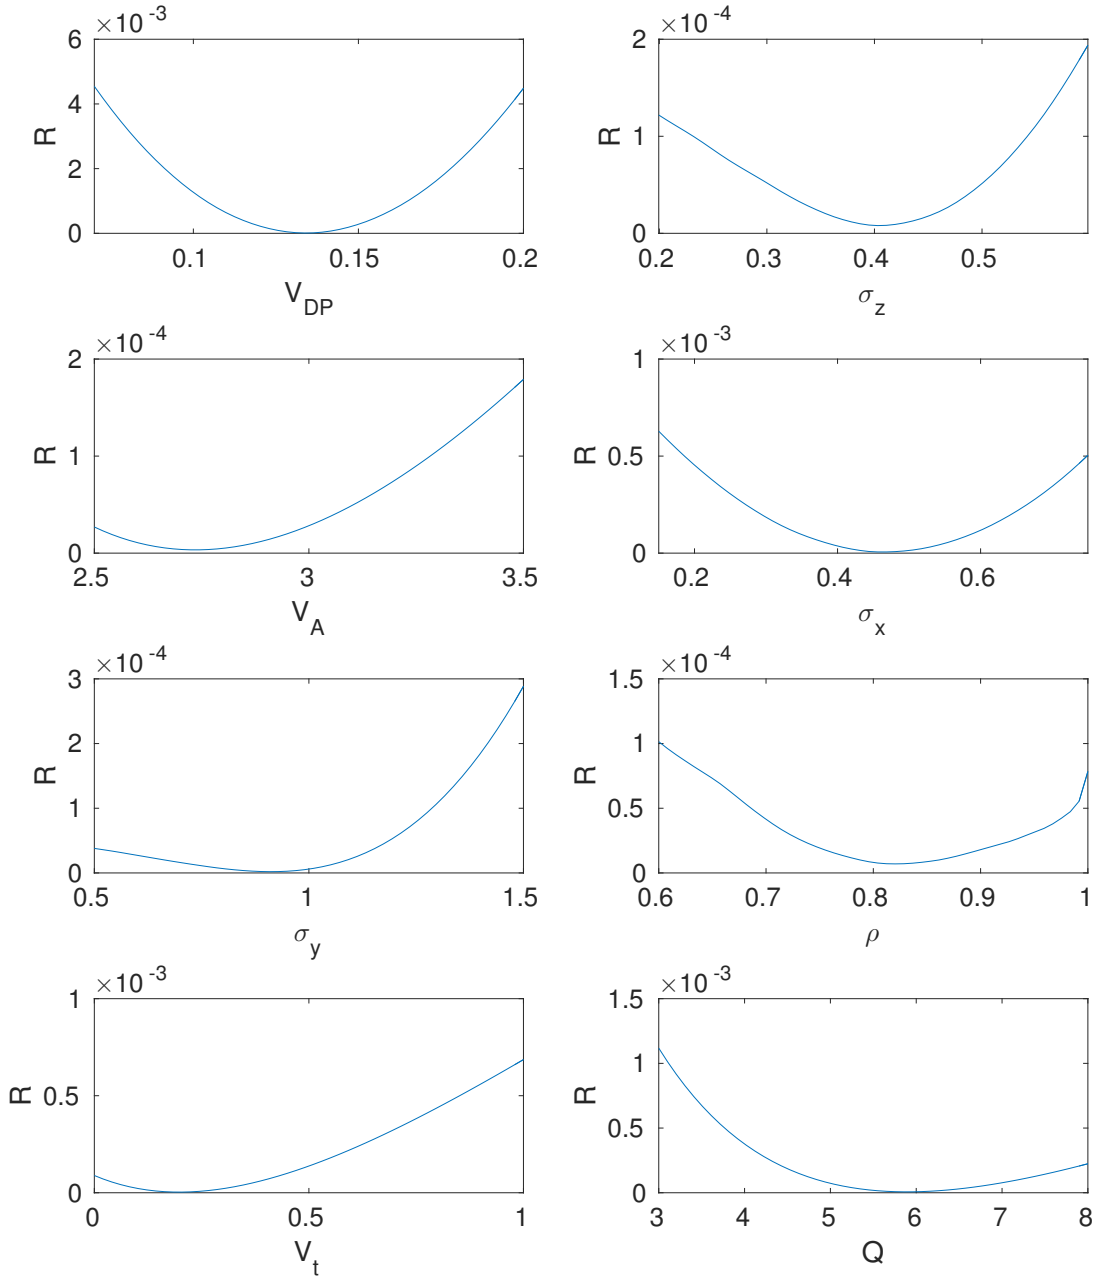

Figure 5.16: The variation of the value of the CO_2 tissue volume objective function, (5.24), with each parameter. In each case one parameter is varied while the values of the others are set to those shown in Table 5.2.

Parameter	True Value	Mean estimate	Standard Deviation
Q_T/lmin^{-1}	5.76	5.57	0.17
V_t	0.25	0.23	0.025

Table 5.7: A summary of parameter estimations performed on noisy synthetic CO_2 washin data.

the similarity of the magnitudes of the variation in the objective function we observed in Figure 5.16.

To test whether this dependence was a problem for parameter identification we attempted to find the set of airway and circulatory parameters that minimised (5.24) using `fmincon` (implemented in parallel on Arcus-b). Unfortunately in ten attempts with varying initial guesses, none of the parameters were estimated correctly. Thus we decided to try a different approach.

Instead of trying to estimate all the parameters, we assume that we know the value of the airway parameters (V_{DP}, σ_x, σ_z and V_A), ρ, and σ_y from washout procedures performed on the same subject (how this is done is discussed in Chapter 6), and then attempt to estimate V_t and Q_T from the CO_2 washin data.

To see if Q_T and V_t were identifiable in principle under these conditions, we attempted 10 parameter estimates on a single set of noise free data with varying initial guesses. This gave us values of $5.81 \pm 1 \times 10^{-4}$ l/min and $0.25 \pm 1 \times 10^{-4}$ for Q_T and V_t respectively, compared to the true values of 5.76 l/min and 0.25 respectively. This suggests that V_t can be estimated from error free data, if we know the values of the other lung parameters. The estimated value of Q_T was close, but not equal to the true value, which suggests that this procedure may not estimate an accurate value of Q_T.

Next to assess the level of uncertainty in these parameter estimates due to experimental noise, we attempted to estimate the values of the two parameters on ten noisy synthetic data sets, created using the LGA noise model we established in Section 5.4. The results are displayed in Table 5.7.

The mean estimated value of Q_T is slightly over one standard deviation from the true value, thus we again conclude that this technique may not give an accurate value for this parameter (although it may give a rough value). The mean estimated value of V_t is within one standard deviation of the true value. This suggests that we may be able to estimate the value of V_t, albeit with a relatively high degree of uncertainty (the standard deviation is approximately 10% of the estimated value of V_t), if we know the true values of the airway parameters (V_{DP}, σ_x, σ_z and V_A), ρ, and σ_y. However, as we will see when we apply this method to experimental data in Section 6.6, in practice we cannot be certain of the true values of these parameters, and the true uncertainty in estimates of V_t is considerably higher.

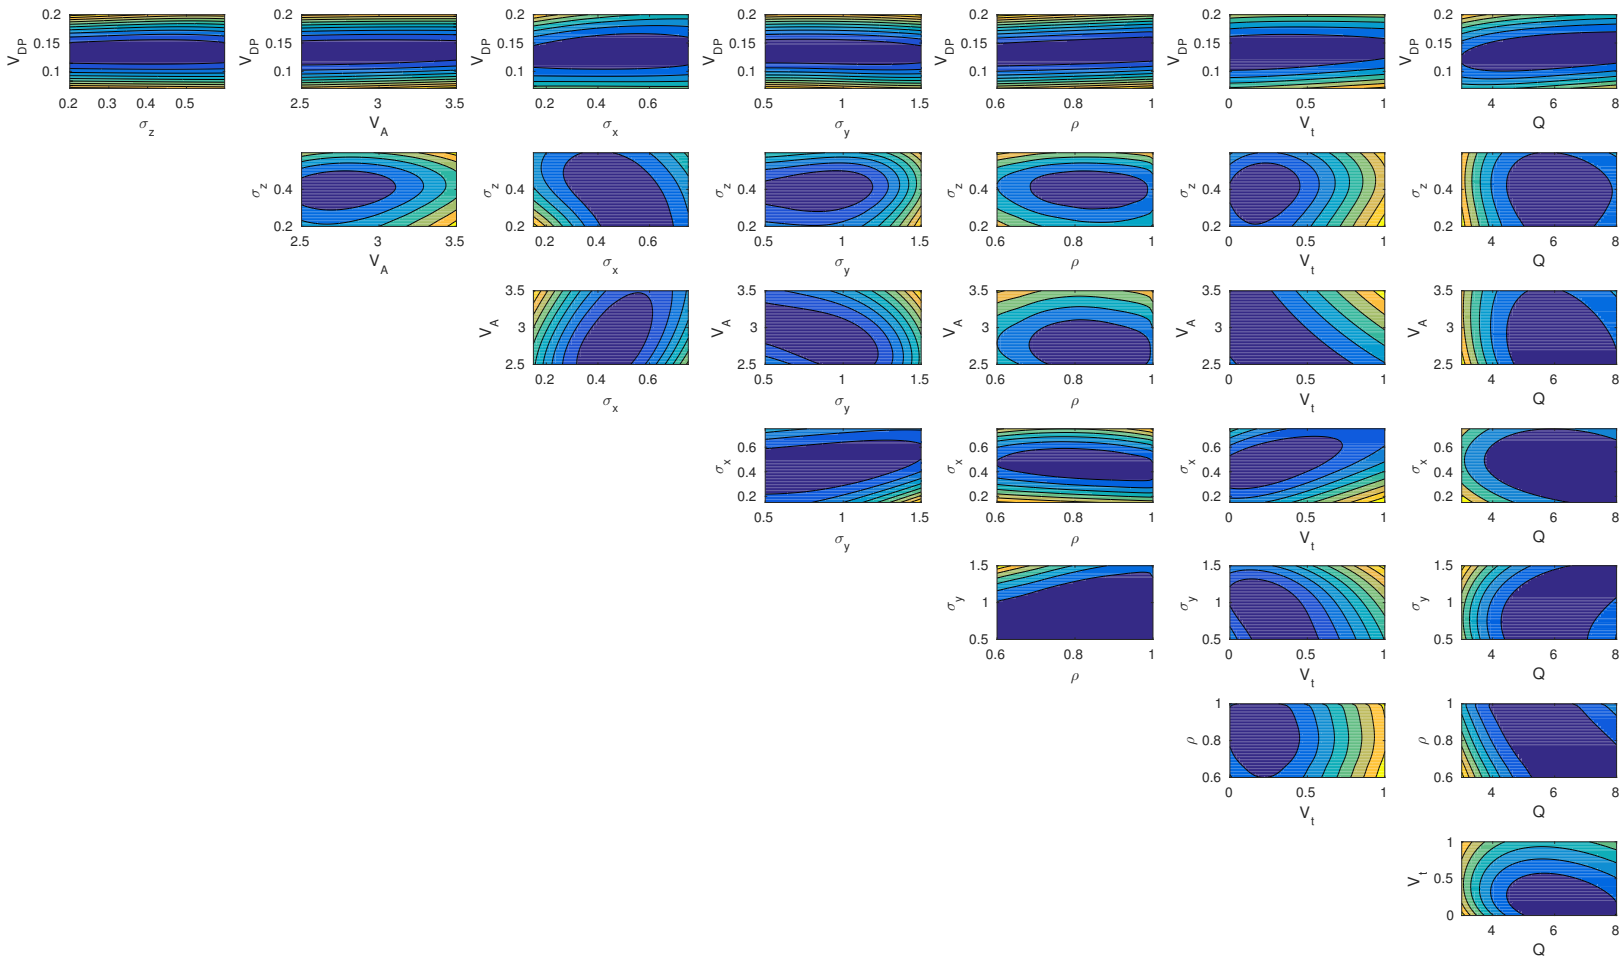

Figure 5.17: The variation of the value of the CO_2 tissue volume objective function, (5.24), as two parameters are varied. In each case two parameters are varied, while the values of the others are set to the values shown in Table 5.2. In this figure blue represents low values of the objective function and yellow high values.

5.6 Summary

The lung function test we proposed in Chapter 1 requires us to be able to estimate the parameters of our lung model from high quality nitrogen washout data. In this chapter we have explained how we estimate those parameters, discussing in some detail the steps we go through to calculate the objective function, and how we find the set of parameters that minimise it. Next, we used synthetic data noise free data to determine that the airway parameters (V_{DP}, σ_x, σ_z, and V_A) are identifiable in principle from the nitrogen washout, while the circulatory parameters (σ_y, ρ, V_t, and Q_T) are not. However, the circulatory parameters could be set to physiologically reasonable values without influencing the estimates of the airway parameters.

We developed a detailed noise model, which allowed us to show that a device which conformed to the ERS recommendations for nitrogen washout measuring devices [96] was not accurate enough for our purposes, but the LGA was. We argued that the uncertainty in the parameters we could estimate due to experimental noise was almost certainly less than the degree of biological variability we expect from patients breathing on the device. Finally, we suggested a technique, which while it could not give as good estimates as the nitrogen washout for most parameters, might allow us to estimate a physiologically reasonable value for one parameter, V_t, that we could not estimate from the nitrogen washout.

Chapter 6

Pilot study results

In this chapter we discuss the results of estimating the parameters of the mathematical model described in Chapter 4 from experimental data. Specifically, we consider data taken from 15 volunteers. Six of the volunteers were young and healthy, while nine were older volunteers. Three of the older volunteers had moderate (GOLD stage 2) COPD.

For the proposed pulmonary function test we outlined in Chapter 1 to give a mechanistic measure of inhomogeneity in the lungs, our mathematical model must represent the physiology of a subject performing a nitrogen washout sufficiently well. This is difficult to assess, but in this chapter we do so in three ways.

First we check how well the output of our model (with fitted parameters) agrees with the experimental data we collect. Next we assess how similar the estimated parameters are between repeat measurements on the same volunteer; as if a subject's physiology does not change between repeat measurements, then neither should our parameter estimates. Finally, we compare our estimated parameters (and measures derived from them) to measurements taken with other modalities. This chapter describes the results of assessing our model by these criteria in each of the groups, and comparing the results of the groups to each other.

At the end of the chapter we discuss the results of the CO_2 washin procedure which we use to estimate a representative value for V_t, the effective CO_2 tissue volume.

6.1 Study details

Our study comprised six young healthy volunteers (all less than 25 years old), six older healthy volunteers (all over 70 years old), and three volunteers with moderate COPD. Some of the physical characteristics of the volunteers are displayed in Table 6.1. A healthy volunteer was classified as someone with no reported lung disease or history of smoking (volunteer E06 smoked as a teenager, but gave up in 1964), while a moderate COPD patient was one with an FEV_1 of between 50% and 80% of their predicted value (GOLD stage 2).

ID	Gender	H/cm	W/kg	A/yrs	FEV_1		FVC		FEV_1/FVC	
Y03	Female	171	65	22	109%	110%	112%	113%	85%	85%
Y04	Female	180	66	24	85%	71%	94%	80%	79%	75%
Y05	Female	173	55	23	110%	114%	110%	110%	88%	90%
Y06	Female	180	63	23	80%	79%	75%	75%	94%	93%
Y07	Male	193	83	20	95%	95%	114%	116%	69%	68%
Y08	Male	177	78	20	142%	141%	142%	144%	84%	83%
E02	Male	183	93	70	87%	87%	104%	105%	65%	64%
E03	Male	169	72	79	65%	58%	73%	82%	86%	60%
E04	Male	176	80	76	129%	123%	130%	127%	75%	72%
E05	Female	160	50	71	92%	93%	114%	115%	67%	67%
E06	Male	185	69	70	99%	103%	114%	117%	67%	67%
E07	Female	165	49	71	113%	114%	131%	140%	72%	68%
C01	Male	178	101	63	87%	73%	128%	104%	53%	55%
C02	Male	163	85	72	75%	75%	124%	116%	46%	50%
C03	Male	170	60	90	58%	49%	100%	93%	41%	38%

Table 6.1: A summary of some of the physical characteristics of the volunteers who participated in this study. Note identifiers (ID) in the healthy volunteers do not start at one as these identifiers were used for members of the team during development of the device. Datasets from these individuals were only taken while the device was under development, and thus they are not included here. FEV_1 and FVC are given as percentage of predicted values, while FEV_1/FVC is the ratio of the quantities expressed in l/s and l. For each of the spirometry measures values shown are those taken on the first and second visit in that order.

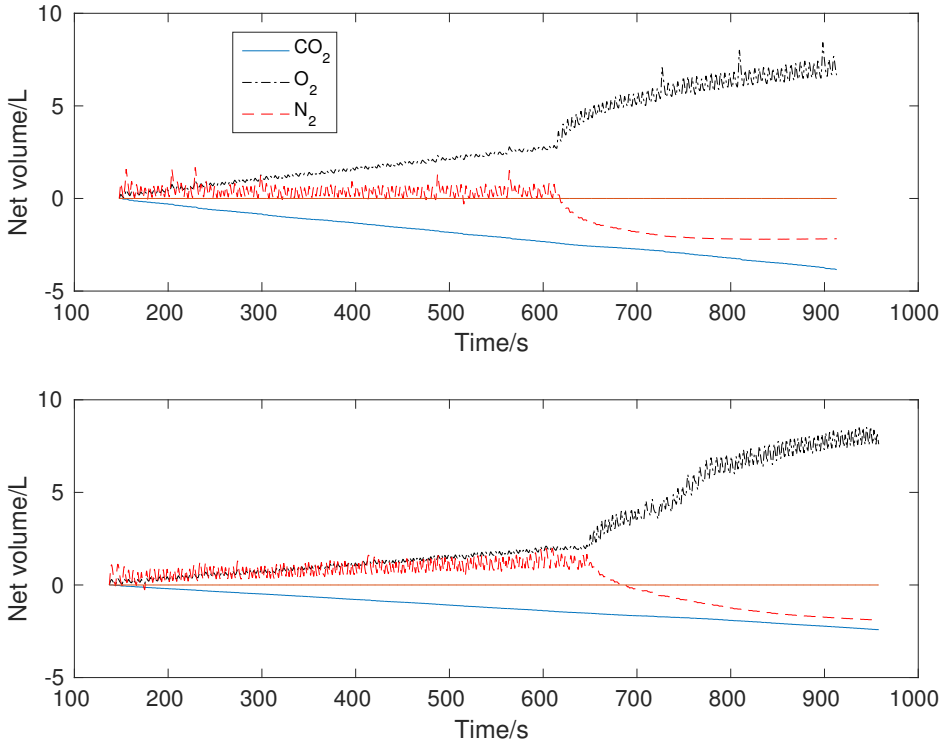

Figure 6.1: Example washout datasets taken from volunteer Y04 (top plot) and C03 (bottom plot). In the plots the net volume of each gas that has entered through the mouth is shown. The high O_2 breathing phase starts after approximately 600 s. The horizontal line is a visual aid to help with the identification of any slope in the N_2 record.

Each volunteer visited the laboratory on two occasions. On each visit three washout procedures (ten minutes air breathing, followed by five minutes of pure O_2 breathing) and one CO_2 washin (ten minutes of air breathing followed by 40 s of breathing 6% CO_2) were performed, and data collected using the LGA. Volunteers breathed through a mouthpiece and air was prevented from leaving or entering their nose with a nose clip (and for those volunteers for whom a nose clip was insufficient, tape as well).

During all procedures the arterial O_2 saturation was monitored via a peripheral probe (Datex-Ohmeda 3900 Tru-Trak+ ®). Spirometry and blood tests were also performed on each volunteer on each visit, except for volunteers Y05 and Y06 who declined to have blood tests. Full ethical approval for this study was received from the Scotland A Research Ethics Committee (REC Reference Number: 11-SS-0091).

6.1.1 Data quality

Example data from nitrogen washout procedures performed by two volunteers, Y04 (top plot) and C03 (bottom plot), is shown in Figure 6.1. These plots show the net volume of each gas that has entered through the mouth as a function of time. The horizontal line is a visual aid to help with the identification of any slope in the N_2 record.

Consider first the top plot. The high O_2 breathing phase starts after approximately 600 s, when large volumes of N_2 start to leave the lung. Note the lack of slope in the N_2 record during the air breathing phase, and the slopes of the CO_2 and O_2 records, associated with CO_2 production and O_2 consumption respectively. An example of the variability of spontaneous breathing can be seen in the large inspirations that correspond to the occasional spikes that are visible in the N_2 record.

In the bottom plot however, the N_2 record is not flat, and shows an apparent N_2 consumption. Since N_2 is almost insoluble in blood and not produced or consumed by the body, we expect that over any extended period inspired and expired volumes of N_2 will be equal (the Haldane effect [43]). If this is not the case then either: the volunteer must have changed position during the experiment (a large decline in lung volume due to a progressive slump would lead to an apparent production of N_2), there must be a leak (for example through the nose where we do not measure the flows or concentrations of gases), or there must be errors in the data caused by the LGA.

Thus before attempting to estimate parameters, the N_2 records from each nitrogen washout procedure were examined. If there was evidence of a significant error in nitrogen balance, as in the bottom plot in Figure 6.1, then this run was excluded from further analysis. Datasets that showed either an error in N_2 balance (difference between inspired and expired nitrogen volumes) of over 45 ml/min over an eight minute air breathing period, or where the N_2 record contained significant sloping periods, were excluded. None of the 36 repeats from the healthy volunteers had to be excluded. However, all the repeats from C03 (who struggled to hold the mouthpiece in his mouth) and E02 (who struggled to maintain his posture for 15 minutes) were excluded because of poor nitrogen balance. Of the remaining 42 repeats from the older volunteers, only one was excluded on the grounds of poor nitrogen balance.

This difference between the groups suggests that the problems are related to volunteer performance rather than the measuring device. While efforts were made to ensure the volunteers were seated as comfortably as possible, and that there were no leaks from either the nose or the mouthpiece, there is clearly a need for further development in this area.

6.2 Values for fixed parameters

In Chapter 5 we showed that the airway parameters (V_{DP}, σ_x, σ_z, and V_A) are estimable to a good degree of accuracy in the presence of realistic experimental noise from a nitrogen washout procedure, while the circulatory parameters (σ_y, ρ, V_t, and Q_T) are not. However, the values of the circulatory parameters did not seem to influence the estimation of the airway parameters suggesting the circulatory parameters could be set to physiologically reasonable values. Thus we attempted to recover values of V_{DP}, σ_x, σ_z, and V_A in each volunteer, our approach to estimating values for the other parameters is discussed below.

Symbol	Value	Description
V_{DP}	Estimated	Total personal deadspace volume
σ_z	Estimated	Width of the deadspace distribution
V_A	Estimated	Total alveolar volume when the lungs are at FRC
σ_x	Estimated	Width of the compliance distribution
σ_y	Estimated	Width of the perfusion distribution
ρ	0.8	Correlation of compliance and perfusion
Q_T	$\frac{\dot{V}_{O_2}}{50}$	Cardiac output
V_t	0.25	CO_2 equivalent tissue volume
V_{DC}	0.155 l	Apparatus deadspace volume

Table 6.2: A summary of the parameters of our model and how their values are obtained.

In Chapter 5 we found that V_t was the most poorly identified parameter, with a coefficient of variation of 37% in estimates of this parameter even from noise free data (Table 5.3). In addition, as shown in Figure 5.9, V_t had by an order of magnitude the smallest effect on the value of the objective function. Given this, V_t was set equal to a single value, 0.25, based on [25], although we show in Section 5.5 that 0.4 may be a more reasonable estimate. However, as we saw in Chapter 5, the choice of this parameter has minimal effect on the estimation of the other parameters.

Cardiac output was also particularly poorly estimated in the presence of noise. Biologically, cardiac output is closely related to metabolic rate, so that venous CO_2 and O_2 concentrations are relatively constant under resting conditions. A typical value for O_2 extraction from the blood at rest is 25% [60]. For the Hb concentration used in the blood gas model (149 g/l), this corresponds to an arteriovenous O_2 concentration difference of approximately 50 ml of O_2/ l of blood. As we explained in Section 2.2 the O_2 consumption ($\dot{V}_{O_2}$) is equal to the cardiac output multiplied by the arteriovenous O_2 concentration difference, thus we set the cardiac output such that

$$Q_T = \frac{\dot{V}_{O_2}}{50}, \tag{6.1}$$

where $\dot{V}_{O_2}$ is measured in units of ml/min. We also explored cardiac outputs set to produce arteriovenous O_2 concentration differences of 40 ml/l and 60 ml/l to investigate the importance of this parameter.

The coefficients of variation for the estimation of both ρ and σ_y were also large in the presence of noise (see Table 5.6), suggesting that they could not be successfully estimated from the data. However, Figure 5.10 suggests that there may be some dependency between these parameters. If so, then it may be possible that if one value is fixed, then a value for the other may be estimated. To see if this was true, we chose to try and estimate σ_y while fixing ρ. We chose to set ρ at 0.8, based on high-resolution micro-sphere studies in animals [2, 68], which give values of 0.80 and 0.81 respectively. Similar values are

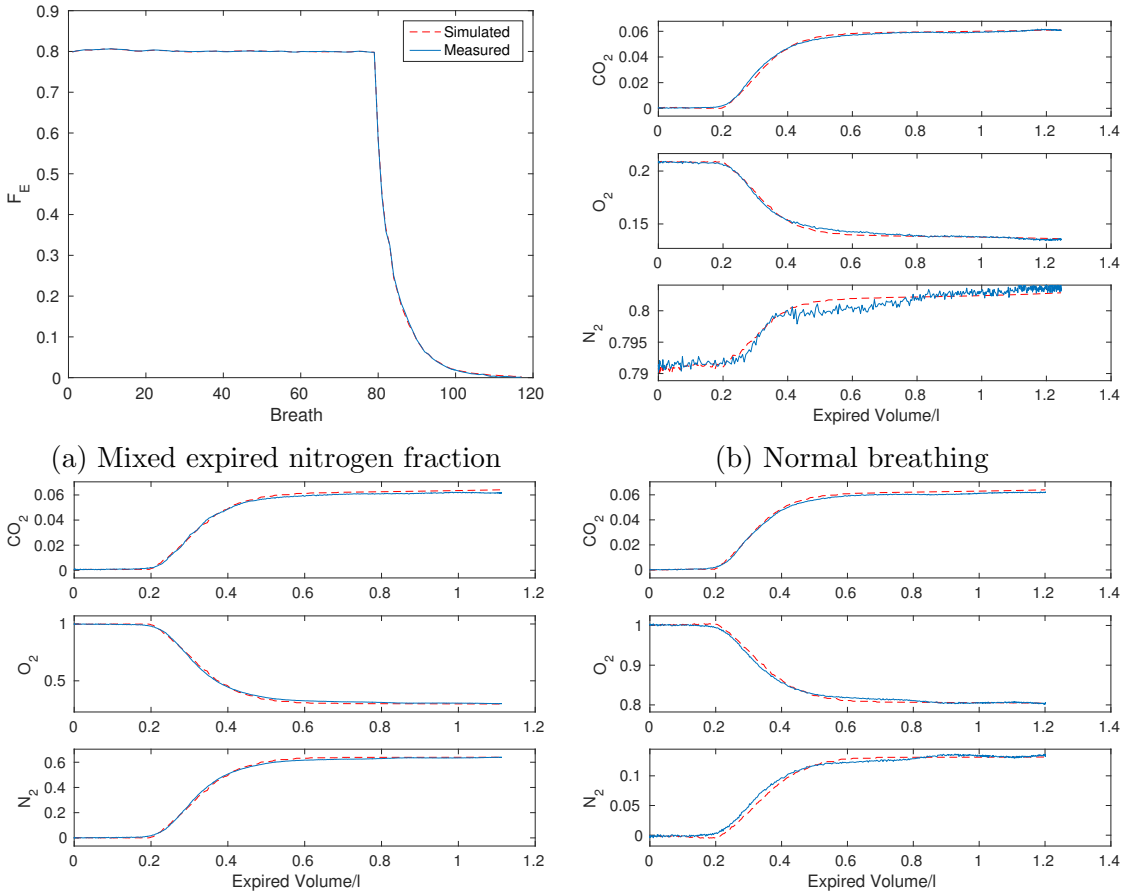

(a) Mixed expired nitrogen fraction

(b) Normal breathing

(c) Washout start

(d) Mid washout

Figure 6.2: These plots compare the simulated data (with fitted parameters) to the measured data for one nitrogen washout procedure performed in volunteer Y07. (a) shows the mixed expired nitrogen fraction for the whole procedure, (b) expirograms during the air breathing phase for all three gases, (c) a breath at the start of the washout, and (d) a breath midway through the washout.

suggested from more derived studies in humans [10]. We also explored values of ρ of 0.65 and 0.95 to investigate the importance of this parameter.

A summary of the parameters we estimate in each of the volunteers, and the assumed values of those that are fixed, is shown in Table 6.2. In the rest of this thesis, unless otherwise stated, the parameters which we do not estimate have the values shown in this table.

6.3 Young healthy volunteers

Having established which parameters we estimate, and values for those we do not, we estimated parameters for the six healthy volunteers from each of the washouts they performed. Figure 6.2 shows a comparison of the simulated output of the model (with fitted parameters) and experimental data for one volunteer. In Appendix B equivalent plots for all the volunteers are shown.

These plots show the mixed expired (average) N_2 fraction over the whole procedure. Also shown are expirograms from the air breathing period (10 breaths before the start of the washout), the start of the washout, and midway through the washout (10 breaths after the start of the washout). In each figure in Appendix B these expirograms are from the same breaths (relative to the start of the washout) to be as representative of the fit quality as possible.

The measured and simulated mixed expired nitrogen fractions are, in almost all cases, very similar. This suggests that the model describes the nitrogen wash out process very well for these volunteers. The agreement between the measured and simulated expirograms is more variable, but in most cases is still very good. This suggests the model describes the within breath morphology well, although there is some variability here. This is perhaps not surprising given the amount of variability there is in spontaneous breathing (see for example the very large inspirations visible in Figure 6.1). Other sources of variability are volunteers changing position, coughing, or swallowing. Nevertheless, we conclude that the model appears to describe the nitrogen washout fairly well for these healthy volunteers, and move on to examining the estimated parameters.

6.3.1 Estimated parameters

Figure 6.3 shows the parameter values estimated from the six repeats in each of the six healthy volunteers. With the exception of σ_y, which as we saw in Chapter 5 was poorly determined even from synthetic data, estimated parameter values were very similar for repeated washout measurements from the same volunteer. Since in most cases the two visits were separated by a couple of days, and we had no reason to suspect that any of the volunteers' lung functions had changed in this period, this is an extremely promising result.

In the following subsections we will discuss the repeatability in more detail, the correlation of the estimated parameters with lung volume, and the influence of the assumed values of ρ and Q_T on the estimated parameters.

6.3.1.1 Repeatability

The coefficients of variation for the estimation of each parameter in each volunteer, and the mean of these coefficients across all volunteers are shown in Table 6.3. These coefficients of variation are all significantly larger than those we estimated from noisy synthetic data, shown in Table 5.6. This is not surprising, as these results include the variation in volunteer performance (for example, variation in posture, coughing etc.) as well as experimental and fitting errors.

Nevertheless, for all parameters except σ_y, the average value of the coefficients of variation are less than 7%. Given that this involves fitting a model as well as errors in

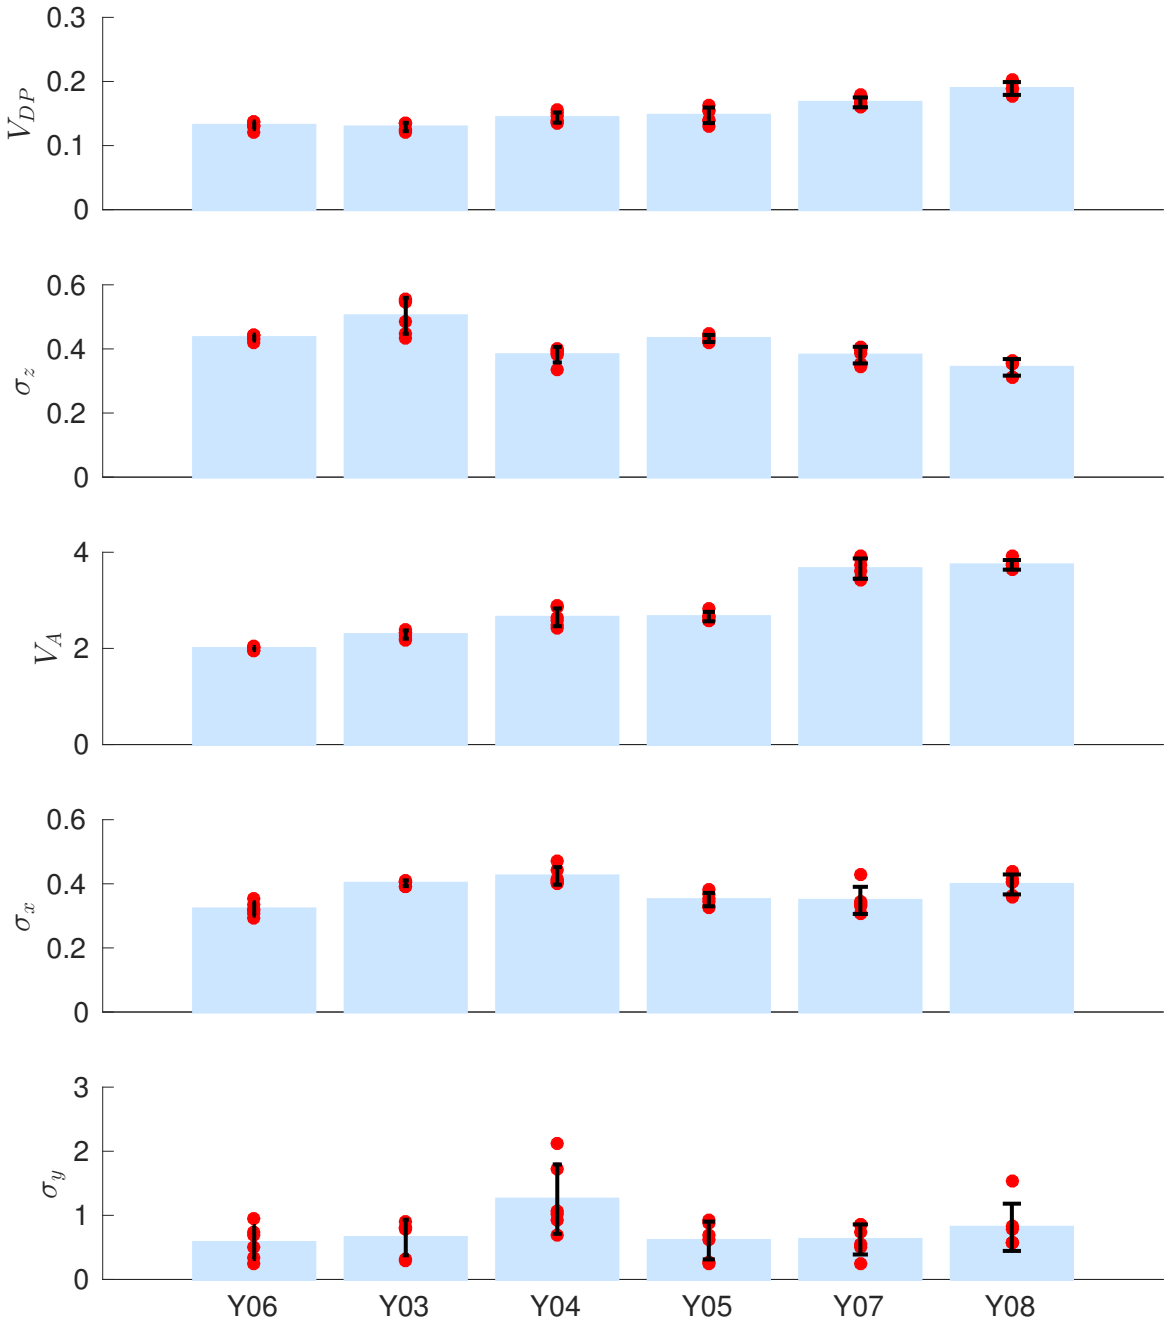

Figure 6.3: Plots showing the estimated parameters for each of the young, healthy volunteers. The red spots show the values for each of the 6 repeats of the washout procedure. The height of the bars show the mean values in each volunteer, and the error bars represent ± one standard deviation. The volunteers are ordered from smallest to largest alveolar volumes. For all the estimates in this plot V_{DC} =0.155 l, $\rho = 0.8$, $V_t = 0.25$, and Q_T was set to give an arteriovenous O_2 concentration difference of 50 ml O_2/l of blood.

Identifier	V_{DP}	σ_z	V_A	σ_x	σ_y
Y03	5.0%	11.1%	3.7%	2.2%	42%
Y04	5.6%	6.4%	7.0%	6.5%	43%
Y05	8.2%	2.5%	3.7%	6.0%	49%
Y06	4.4%	2.2%	1.7%	6.5%	45%
Y07	4.6%	6.8%	5.8%	12.1%	38%
Y08	5.3%	7.6%	2.7%	7.8%	45%
Average	5.5%	6.1%	4.1%	6.9%	44%

Table 6.3: Coefficients of variation (standard deviations divided by the mean) from 6 repeats on each of the six healthy volunteers. For all the estimates in this table V_{DC} =0.155 l, $\rho = 0.8$, $V_t = 0.25$, and Q_T was set to give an arteriovenous concentration difference of 50 ml of O_2/L of blood. The values in the final row are the averages for the whole group.

data collection and volunteer performance, this compares reasonably well with the intra-visit repeatability of lung clearance index (LCI) measurements from SF_6 washouts of 3.6% and the inter-visit reproducibility of 8.9% (0.6 on a mean value of 6.7 for healthy adults) reported in [51].

6.3.1.2 Correlation with lung volume

In the Figure 6.3 the volunteers are ordered from smallest to largest lung volume. There is a clear relationship between the volume of deadspace and lung volume, and perhaps an inverse one between σ_z and lung volume. Analysis with linear regression suggested both of these relationships were statistically significant.

We will later compare the ratio $\frac{V_{DP}}{V_A}$ between the different groups of volunteers, and so values of this ratio for the healthy volunteers are plotted in Figure 6.6. Also plotted in this figure are the values of $\sigma_z \frac{V_{DP}}{V_A}$, which is the width of the deadspace distribution in ml deadspace/ ml alveolar volume. We refer to this as the scaled width of the deadspace distribution.

However there does not appear to be any relationship between lung volume and σ_x. The adjusted R^2 for a regression of σ_x against V_A was 0.01. If this holds true once measurements are made in a larger number of people, this is a promising finding; since it suggests that, unlike FEV_1, σ_x may be an intensive (a measure that does not vary with body size) rather than an extensive measure (a measure that varies with body size).

6.3.1.3 Influence of ρ and Q_T

The results displayed in Figure 6.3 were obtained using a value of Q_T set to give an arteriovenous O_2 concentration difference of 50 ml O_2/l blood, and a value of 0.8 for ρ. To test the effect of these assumptions on the estimated values, we ran a grid of nine different fits with assumed cardiac outputs of 40, 50 and 60 ml O_2/l blood and values of ρ of 0.65, 0.8, and 0.95 on each of the 36 datasets from this group.

Parameter	Q_T effect size	p value	ρ effect size	p value
V_{DP}	0.09%	0.8	1.06 %	4×10^{-4}
σ_z	1.40%	0.2	6.10%	7×10^{-7}
V_A	3.53%	0	0.01%	0.93
σ_x	3.63%	0.03	-7.06%	8×10^{-5}
σ_y	-19.88%	0.006	-47.81%	5×10^{-9}

Table 6.4: A summary of the univariate regression performed to assess the influence of the assumed values of ρ and the arteriovenous O_2 concentration difference (using which we set Q_T). The effect size gives an estimate of the percentage change in the value of the parameter for a doubling of the assumed values of ρ or the arteriovenous O_2 concentration difference as relevant.

Table 6.4 shows the results of a regression analysis to determine the influence of each these parameters. Here we applied a linear model to each of the estimated parameters in turn. We assumed that each run on each volunteer was a random variable (i.e. had its own mean), and performed a regression with ρ and the arteriovenous O_2 concentration difference as independent variables, and the estimated parameter as the dependent variable. As we have parameters with a variety of magnitudes these results are expressed as 'effect sizes'. This is an estimate of the percentage change in the estimated parameter value for a doubling in the assumed value of either ρ or the arteriovenous O_2 concentration difference as relevant. This heuristic quantity was calculated by dividing the regression co-efficient by the mean of the estimated parameter across all the volunteers, and multiplying by the average value of ρ or the arteriovenous O_2 concentration difference as relevant.

As shown in this table, we found for all parameters except σ_y, the choice of Q_T or ρ had little effect on the value of the estimated parameters; equivalent to a 7% or less change in parameter value for a doubling of either Q_T or ρ. This agrees very well with our findings in Chapter 5 that our inability to estimate Q_T and ρ did not affect our ability to estimate values for the airway parameters, and validates our decision to use physiologically reasonable values for these parameters.

In the case of σ_y, doubling the value of ρ had an influence equivalent to a 48% reduction in the value of σ_y (and doubling the arteriovenous concentration difference a 20% reduction). This is consistent with the dependence we noted in Figure 5.10. Along with the large standard deviation between repeats in the same volunteer, this suggests that our method is not a robust means of estimating σ_y (again consistent with our findings in Chapter 5). However, the large variations in σ_y did not prevent consistent values for the other parameters being estimated.

6.3.2 Comparison with existing modalities

In the previous section we demonstrated that our model fits data from normal healthy volunteers well, and that the parameter estimates that it gives are repeatable (with the

Identifier	LCI
Y03	5.82 ± 0.19
Y04	6.06 ± 0.34
Y05	6.65 ± 0.28
Y06	6.65 ± 1.8
Y07	6.29 ± 0.72
Y08	6.00 ± 0.4
Average	6.24 ± 0.8

Table 6.5: Estimated values of the LCI (means $\pm$ standard deviations) from 6 repeats on each of the six healthy volunteers. The values in the final row are the averages for the whole group.

exception of σ_y). In this section we compare these estimated parameters, and other outputs from our model, with existing modalities.

Before doing this we briefly comment on the values of LCI we estimated for the healthy volunteers from the nitrogen washouts we performed on them.

6.3.2.1 Lung clearance index

We calculated LCI for each of our volunteers using equation (2.28), where, following the recommendations of Robinson *et al.* [95], we removed the volume of apparatus deadspace from the ventilation and estimated FRC using the following formula

$$\text{FRC} = \frac{\text{net volume of inert gas exhaled}}{F_{ET}(1) - F_{ET}(n)}, \tag{6.2}$$

where $F_{ET}(1)$ and $F_{ET}(n)$ are the end tidal N_2 fraction on the first and last breath under consideration.

Means and standard deviations of the LCI from the six runs we performed on each volunteer are shown in Table 6.5. The mean values agree reasonably well with those reported by Husemann *et al.* [54] and Bell *et al.* [13], 6.94 ± 0.61 and 7.0 ± 0.7 respectively, for nitrogen washouts on healthy volunteers. However, in a number of volunteers the repeatability is significantly worse than that reported for LCI values calculated from SF_6 washouts [51]. This is probably because we made no effort to control tidal volume (as this is not required for our technique) unlike the authors of that paper.

6.3.2.2 Agreement with predicted FRC

Figure 6.4 compares the estimated value of FRC from our technique ($V_{DP} + V_A$) and predicted values based on the volunteer's age, gender and height. The predicted values were calculated using formulas taken from [105], which are commonly used in lung function testing [126]. For males the predicted value of FRC in litres is given by,

$$\text{FRC}_{\text{predicted}} = 2.34H + 0.01A - 1.09, \tag{6.3}$$

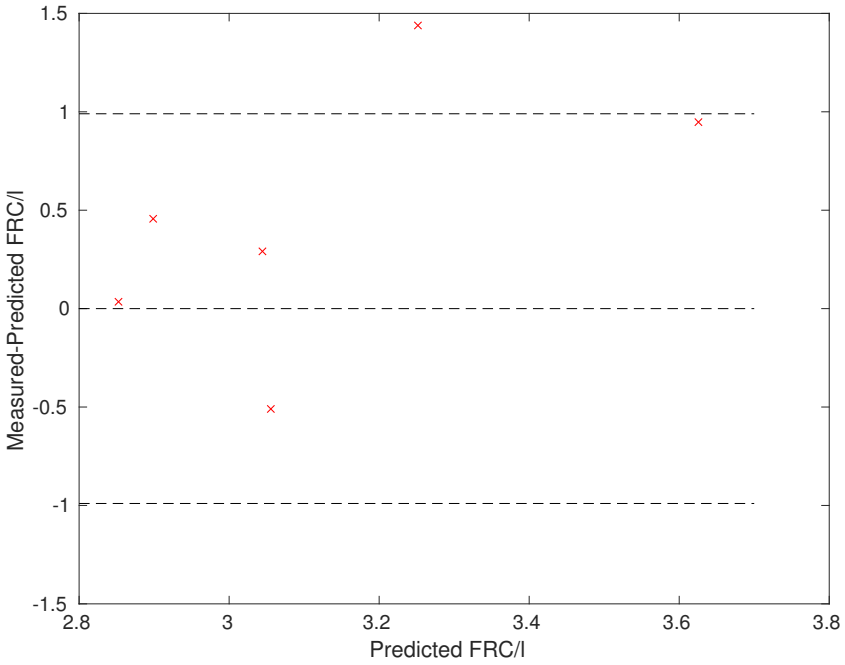

Figure 6.4: A comparison of the value of FRC predicted by (6.3) or (6.4) as relevant, and the value we estimate using our technique. The dotted lines represent 95% confidence intervals.

where H is the height measured in metres and A the age in years. For females the equivalent formula is,

$$\mathrm{FRC}_{\mathrm{predicted}} = 2.24H + 0.001A - 1. \tag{6.4}$$

The 95% confidence intervals on these values are ± 0.99l and ± 0.82l for the males and females respectively.

In this figure we have also converted V_{DP} and V_A from volumes at STPD (as discussed in Chapter 4, we solve for dry fractions and volumes at STPD in our model) to BTPS, which is the conditions under which plethysmographic volumes are reported, using the following formula (derived from the ideal gas law) [129],

$$V_{\mathrm{BTPS}} = V_{\mathrm{STPD}} \frac{T_{\mathrm{B}}}{T_{\mathrm{S}}} \frac{P_{\mathrm{S}}}{P_b - P_{\mathrm{H_2O}}}, \tag{6.5}$$

where $T_{\mathrm{B}} = 310.15$K and $T_{\mathrm{S}} = 273.15$K are the body and standard temperatures respectively, $P_{\mathrm{S}} = 100$kPa and $P_b = 101.3$kPa are the standard and barometric pressures respectively, and $P_{\mathrm{H_2O}} = 6.266$kPa is the saturated partial pressure of water at BTPS.

All volunteers except one (Y08) had estimated values of FRC that were within the 95% confidence interval of the predicted values. This volunteer had both FEV_1 and FVC values in excess of 140% of their predicted values, and thus they appear to have exceptionally large lungs.

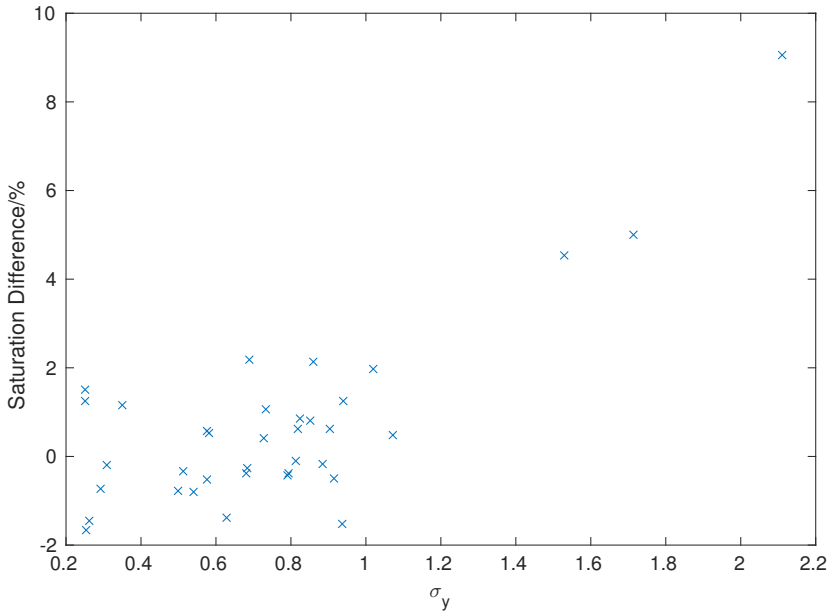

Figure 6.5: A plot of σ_y against the difference between the measured and estimated O_2 saturation for each run on each volunteer.

The formulas for predicted FRC are based on data taken from plethysmographic measurements, and so include non communicating regions of gas within the thorax. Thus similar to the reported differences between gas dilution and plethysmographic values [126], we might guess our values would be lower (on average) than the predicted values. However, five of the six volunteers had estimated values of FRC that were larger than their predicted values.

This is not necessarily incompatible with previous findings, as it is possible the true values of FRC for these volunteers are larger than those predicted by (6.3) and (6.4). As can be seen in Table 6.1, four of the six had FVC values larger than predicted, and thus they may as a group have large lungs. This would not be surprising as they were a group whom we might expect to have very good lung function (four of the six were university level athletes and one a trumpet player). Full lung function testing is being organised to obtain measured values for FRC.

6.3.2.3 Arterial oxygen saturation

As we saw in Chapter 5, values for the circulatory parameters (Q_T, ρ, V_t, and σ_y) were not estimable from our data. This is not surprising, as we only have data taken at the mouth, and thus have to infer what is happening in the blood. However, as we mentioned previously, the volunteers' arterial O_2 saturations (S_pO_2) were monitored during the washout procedures with a peripheral probe. The O_2 saturation of the blood is defined as the percentage of haemoglobin in the blood that is saturated with O_2 [60], and thus this measure gives us some access to what is happening in the blood.

Parameter	Before exclusion	After exclusion
V_{DP}/L	0.144 ± 0.008	0.139 ± 0.005
σ_z	0.38 ± 0.02	0.38 ± 0.03
V_A/L	2.65 ± 0.19	2.72 ± 0.18
σ_x	0.42 ± 0.03	0.43 ± 0.03
σ_y	1.25 ± 0.54	0.92 ± 0.17

Table 6.6: A summary of the effect of excluding the set of parameter estimates that give a difference between the measured and simulated S_pO_2 of greater than 2% on volunteer Y04. In all cases values given are means plus or minus standard deviations.

This probe has a stated accuracy of ±2%, and thus was not felt to be precise enough to be included as part of the objective function used in parameter estimation. However, it is useful to consider how well our estimates of S_pO_2 from the model agree with the experimentally measured values over the final five minutes of the air breathing period.

To estimate S_pO_2 from the model we calculate the blood flow weighted average arterial O_2 concentration during the final five minutes of the air breathing period ($C_a^{O_2}$),

$$C_a^{O_2} = \frac{1}{t_a - t_s} \int_{t_s}^{t_a} \sum_{i=1}^{N_s N_t N_d} Q_i C_{ai}^{O_2}(t_i) \, \mathrm{d}t, \tag{6.6}$$

where t_s and t_a are the start and end times of the air breathing period respectively, $C_{ai}^{O_2}(t_i)$ is the arterial O_2 concentration in compartment i at time t, and Q_i is the fractional perfusion. From this the average S_pO_2 during this period can be found using the blood-gas model [81].

If we compare the S_pO_2 value estimated from our model to the measured value, then only five of the 36 parameter estimates give an estimated value that is more than ±2% away from the measured value. Of these the difference in estimated and measured S_pO_2 is greater than 2.2% in only three cases.

Figure 6.5 shows the difference between the measured and estimated S_pO_2 values for each run on each volunteer plotted against σ_y. From this figure it is clear that the large saturation differences are associated with large values of σ_y.

Two of the three runs with large differences in the estimated and measured saturations were from volunteer Y04. Thus we considered the effect on the parameter estimates in this volunteer of excluding runs with a saturation difference greater than 2.2%. Table 6.6 shows the effect on the means and standard deviations of the parameter estimates for volunteer Y04 when these two runs are ignored. The only parameter for which the mean before exclusion was more than one post exclusion standard deviation away from the new mean was σ_y. This was also the only parameter whose standard deviation was significantly altered by the exclusion of these runs.

Thus it appears the only parameter estimate that is meaningfully changed by this exclusion is σ_y. This supports the finding in Chapter 5 that the uncertainty in σ_y does

not influence the estimation of the airway parameters (V_{DP}, σ_z, V_A, and σ_x). It also suggests that if our interest is in these airway parameters then our conclusions will not be altered by whether or not we exclude points where the estimated and measured SpO_2 differ significantly. However, it may allow us to exclude extreme values of σ_y, and thus improve our estimation of that parameter.

6.3.2.4 Ventilation-perfusion distributions

In Chapter 2 we introduced the multiple inert gas technique (MIGET) which has been used to study the ventilation-perfusion distribution within the lung. In this section we investigate the ventilation perfusion distributions that we estimate from our technique. We then compare them to ventilation-perfusion distributions found in the literature.

When we introduced the governing bivariate lognormal distribution in Section 4.1, we explained that Wilson and Beck [139] introduced this distribution to try to reconcile experimentally measured ventilation-volume ($\frac{\dot{V}}{V}$) and perfusion-volume ($\frac{\dot{Q}}{V}$) distributions with ventilation-perfusion ($\frac{\dot{V}}{\dot{Q}}$) distributions estimated using MIGET.

They assumed that the ventilation-volume and perfusion-volume distributions are lognormal, with standard deviations $\sigma_{\frac{\dot{V}}{V}}$ and $\sigma_{\frac{\dot{Q}}{V}}$ respectively, and then showed that since

$$\log\left(\frac{\dot{V}}{\dot{Q}}\right) = \log\left(\frac{\dot{V}}{V}\right) - \log\left(\frac{\dot{Q}}{V}\right), \tag{6.7}$$

the ventilation-perfusion distribution is also lognormal with standard deviation $\sigma_{\frac{\dot{V}}{\dot{Q}}}$, where

$$\sigma^2_{\frac{\dot{V}}{\dot{Q}}} = \sigma^2_{\frac{\dot{V}}{V}} + \sigma^2_{\frac{\dot{Q}}{V}} - 2\rho\sigma_{\frac{\dot{V}}{V}}\sigma_{\frac{\dot{Q}}{V}}, \tag{6.8}$$

where ρ is the correlation between the ventilation and perfusion. This formula gives us an insight into the behaviour of the width of the ventilation-volume distribution. For example, we can see that as either the ventilation-volume or perfusion-volume distributions get broader, or their correlation, ρ, decreases, the width of the ventilation-perfusion distribution will increase.

However, we cannot use (6.8) to estimate the width of the ventilation-perfusion distribution from our parameter estimates, as in our model ventilation is not exactly lognormally distributed. There are two reasons for this. First, the ventilation of the deadspace must be subtracted from the total ventilation to get the alveolar ventilation, and secondly, we assume that fractional-volume change ("compliance") rather than ventilation is lognormally distributed.

Thus to estimate a ventilation-perfusion distribution comparable with MIGET we use our model (with the parameter values set to those estimated from the run under consideration) to estimate the alveolar ventilation ($\dot{V}_{Ai}$) over the final five minutes of the air breathing period. This is the volume of fresh (i.e. starting inspiration outside the

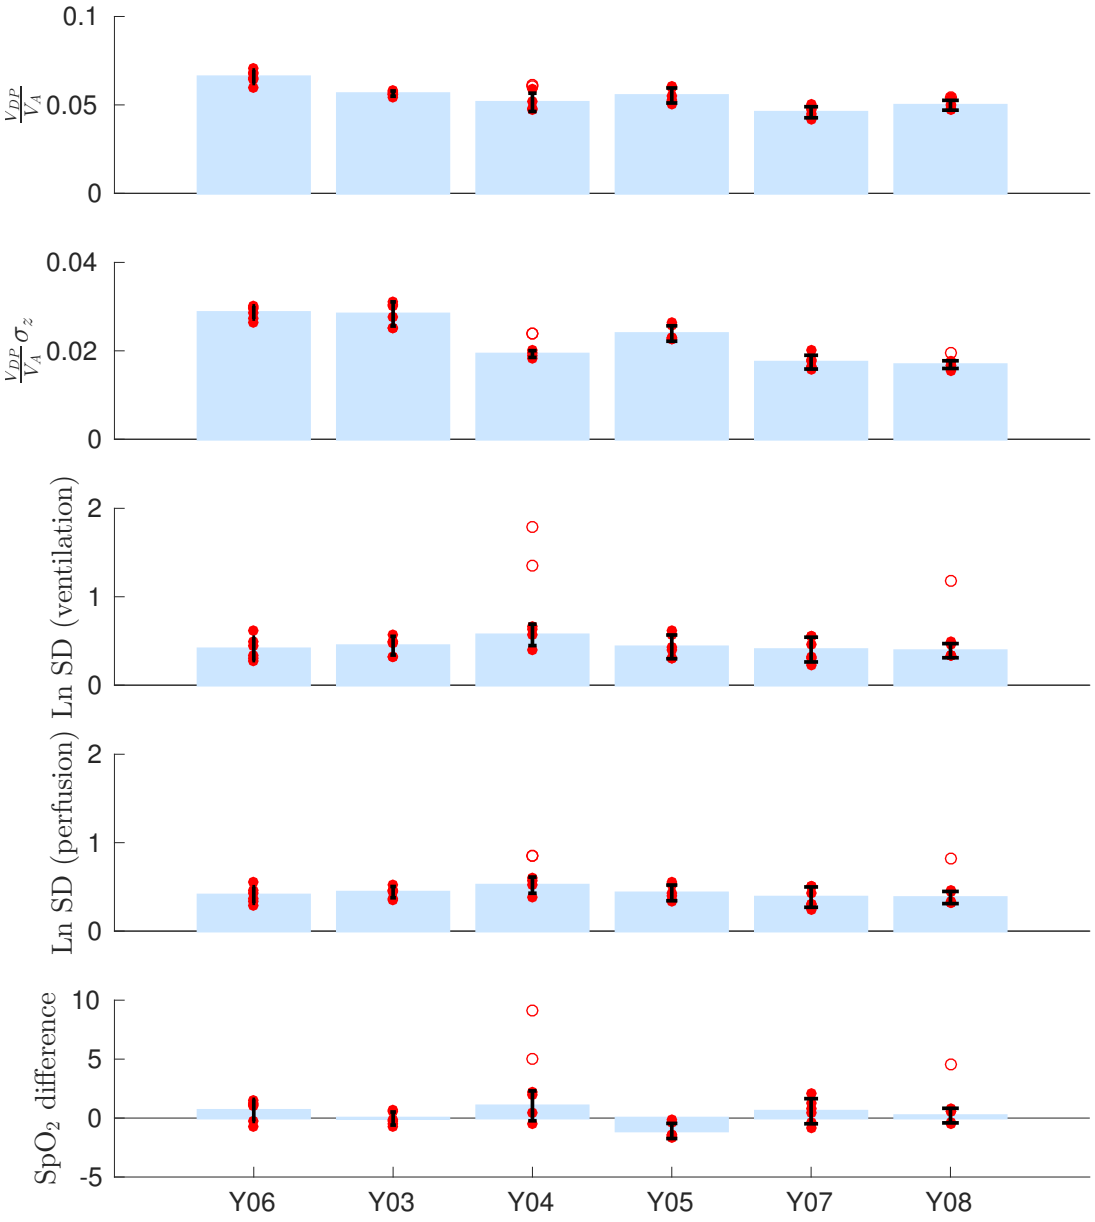

Figure 6.6: Plots showing a number of measures derived from our estimated parameters for the young healthy volunteers. The meaning of these measures and how they are derived is discussed in the main body of the text. The red spots show the values for each of the 6 repeats of the washout procedure, the height of the bars their means, and the error bars are $\pm$ one standard deviation from the mean. The volunteers are ordered from smallest to largest alveolar volumes. For all the estimates in this plot V_{DC} =0.155 l, $\rho = 0.8$, $V_t = 0.25$, and Q_T was set to give an arteriovenous O_2 concentration difference of 50 ml of O_2/l of blood. Red open circles were excluded from the estimates of the mean based on the unphysical values of SpO_2 predicted by the model, see Section 6.3.2.3.

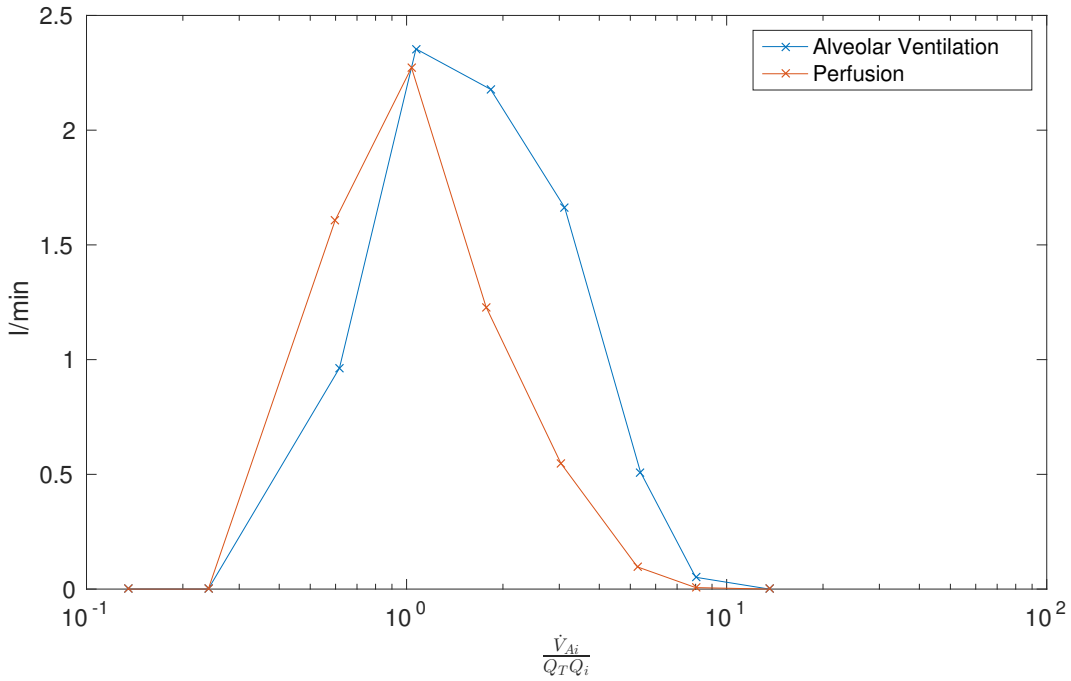

Figure 6.7: An example plot of the ventilation-perfusion distribution in the lung. This shows the ventilation or perfusion going to a region of the lung with a particular value of the ventilation-perfusion ratio. To produce this graph compartments were grouped into nine groups, each with average ventilation-perfusion ratio indicated by a cross.

deadspace) gas that enters each alveolar compartment during the period, divided by the duration of the period. On any one breath the alveolar ventilation is,

$$\dot{V}_{Ai} = \sum_m \left(\frac{\dot{V}_i^m \left(V_T^m - V_{DC} \right) - v_i V_{DP}}{t_b^m} \right), \tag{6.9}$$

where V_T^m and t_b^m are the tidal volume and duration of the breath m, and $\dot{V}_i^m$ is the fractional ventilation entering compartment i on breath m. The perfusion to each compartment is simply $Q_i Q_T$, and thus we can define the ventilation-perfusion ratio of the compartments, $\frac{\dot{V}_{Ai}}{Q_i Q_T}$. An example histogram of the distribution of $\frac{\dot{V}_{Ai}}{Q_i Q_T}$ with both $\dot{V}_{Ai}$ and $Q_i Q_T$, similar to those produced by MIGET (see for example [124]), is shown in Figure 6.7.

The most commonly reported parameters from the MIGET technique are the Ln SD (perfusion) and Ln SD (ventilation) [125]. We refer to Ln SD rather than Log SD to make clear it is the natural logarithm that is used. These are the standard deviations of the perfusion and ventilation distributions of $\log\left(\frac{\dot{V}_{Ai}}{Q_i Q_T}\right)$ respectively. Values from each repeat in each volunteer are plotted in Figure 6.6. Means and standard deviations for each volunteer are also shown in this figure. In the calculations of these means we have excluded the runs for which the saturation difference is greater than 2.2% (denoted with open circles in the figure). The difference between Ln SD (perfusion) and Ln SD

	0.65	0.8	0.95
40	0.64 ± 0.18	0.45 ± 0.12	0.33 ± 0.11
50	0.59 ± 0.15	0.41 ± 0.13	0.26 ± 0.05
60	0.65 ± 0.16	0.41 ± 0.09	0.26 ± 0.05

Table 6.7: Estimated values of Ln SD (ventilation) (means $\pm$ standard deviations) from 6 repeats on volunteer Y06 at a range of values of ρ (listed along the top row) and arteriovenous O_2 concentration differences (listed in the first column in ml O_2/l blood).

(ventilation) is small as observed elsewhere [124], and would be zero if both ventilation and perfusion were exactly lognormally distributed [10].

The standard deviations of the estimates of the Ln SD are high (typically around 25% of the mean value) in individual volunteers, and thus we do not advocate this technique as a method for estimating ventilation-perfusion distributions in healthy young volunteers. This variability is not surprising given the coefficients of variation of our estimates of σ_y, shown in Table 6.3. Nevertheless, it is re-assuring that the mean estimated value of Ln SD for each volunteer lies within the normal range indicated by Wagner of 0.4-0.6 [125].

These values are all calculated with $\rho = 0.8$ (and an arteriovenous O_2 concentration difference of 50 ml O_2/l blood). From (6.8), we can see that increasing ρ will, if no other parameter changes, tend to decrease the variance of the ventilation-perfusion distribution. In fact, as we noted in Section 6.3.1, the value of σ_y (the width of the perfusion distribution) decrease as ρ increases. This will further reduce the width of the ventilation-perfusion distribution as ρ increases.

Table 6.7 shows the effect of the choice of ρ and the arteriovenous O_2 concentration difference on the estimated values of Ln SD (ventilation) for volunteer Y06 (the patterns for other volunteers and Ln SD (perfusion) were similar). Given the wide standard deviations of the estimated values, the effect of the arteriovenous concentration difference is not material. However, as we would expect from the discussion above, the higher the correlation between the compliance and perfusion distributions, the lower the Ln SD (ventilation) of the ventilation-perfusion distribution. At $\rho = 0.95$, values of the Ln SD (ventilation) are slightly below the normal range defined by Wagner, and at at $\rho = 0.65$ slightly above.

Thus our estimates of the Ln SD depend on a parameter that we fix at a constant value because we cannot estimate its value from our data. This further supports our conclusion that our technique is not an appropriate method for exploring ventilation-perfusion distributions. However, that the values of Ln SD are within the normal range at $\rho = 0.8$ provides some further support for our use of this value of ρ.

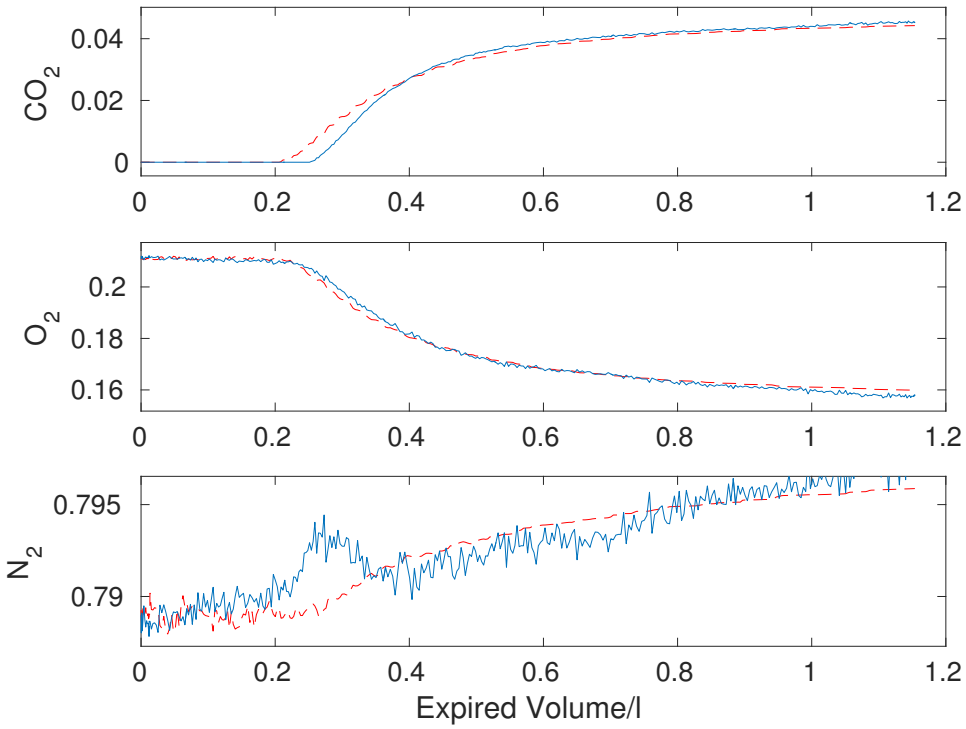

Figure 6.8: An example of an expirogram from volunteer E03 during the air breathing period comparing the simulated output of the model with fitted parameters (dashed red line) and experimental data (solid blue line). A noticeable feature of this expirogram is that alveolar CO_2 appears to emerge later than alveolar O_2.

6.4 Older volunteers

We also collected data from nine older volunteers, six of whom were healthy (no reported history of lung disease) and three of whom had moderate (GOLD stage 2) COPD. As described in Section 6.1.1, while the quality of the data collected from this group was generally very good, two volunteers (C03 and E02), struggled to perform the technique. Their data showed significant deviations from nitrogen balance during the air breathing period, and was excluded from further analysis. Other than this, only one dataset (from volunteer E05) was excluded on nitrogen balance grounds.

Comparisons of the simulated output of the model (with fitted parameters) and experimental data for one run on each older volunteer (except for volunteer C01, where we show one plot from each visit) are shown in In Appendix B (Figures B.7-B.14). As for the healthy volunteers, the simulated mixed expired nitrogen fractions are in good agreement with the measured values, suggesting we describe the washout process well in these volunteers. In general, the model also appears to describe the intra-breath morphology well. See the expirograms shown in, for example, Figures B.9, B.10, B.10, and B.13).

However, there does appear to be more variability in the agreement between the simulated and measured expirograms in the older volunteers than in the younger volunteers. In Figure 6.8 we show an example of a feature that we observed in a number of volunteers

in this group. In the experimental data shown in this figure phase II (in which alveolar gas begins to emerge) appears to start earlier for O_2 than CO_2. In our model, however, this does not occur and phase II begins at the same time for both gases. This feature is apparent in a number of the older volunteers (for example E03, E04, C01, and C02). On closer inspection of the data collected from the younger volunteers it is also visible in volunteer Y04.

We can appreciate why this feature is not shown by our model if we consider an important difference between convective and diffusive processes. Convective processes depend on gradients of air pressure, and thus affect all gases equally, while diffusive processes depend on gradients of partial pressures/dry fractions of particular gases. Thus any feature which appears to show asymmetry between different gases may well be a signature of a diffusive process [119]. In our model diffusion is either complete and instantaneous (in the perfectly mixed alveolar compartments) or ignored (in the deadspace compartments). Thus this feature may be an indication that our assumption that we could ignore any diffusion limitation (stratification) in the alveoli is, at least in some volunteers, invalid. We will discuss possible methods of incorporating stratification in the next chapter.

6.4.1 Estimated parameters

Figure 6.9 shows the estimated parameters from all the viable repeats in the older volunteers. The volunteers are ordered in increasing lung volume with the two volunteers with moderate COPD on the far right. With the exception of C01, in whom we estimated quite different values from his two visits, the measurements appear repeatable (repeatability will be discussed in more detail in the next section).

Also shown in this figure are the range of means we found in the healthy volunteers. All but the two female volunteers (E05 and E07) show values for the widths of the compliance and perfusion distributions (σ_x and σ_y respectively), and deadspace volume (V_{DP}) that are elevated compared to the young healthy volunteers. In a number of cases σ_z takes lower values than those seen in the young volunteers, although as we will see in Section 6.4.2, when we normalise this parameter with $\frac{V_{DP}}{V_A}$ this is not the case.

6.4.1.1 Repeatability

Coefficients of variation for each parameter in each of the older volunteers are shown in Table 6.8, along with means from both groups. In general, the values from the two groups appear similar suggesting that repeatability is similar across the two groups, although it appears the value of σ_y, while still not as precisely determined as the airway parameters, may be more estimable in this group.

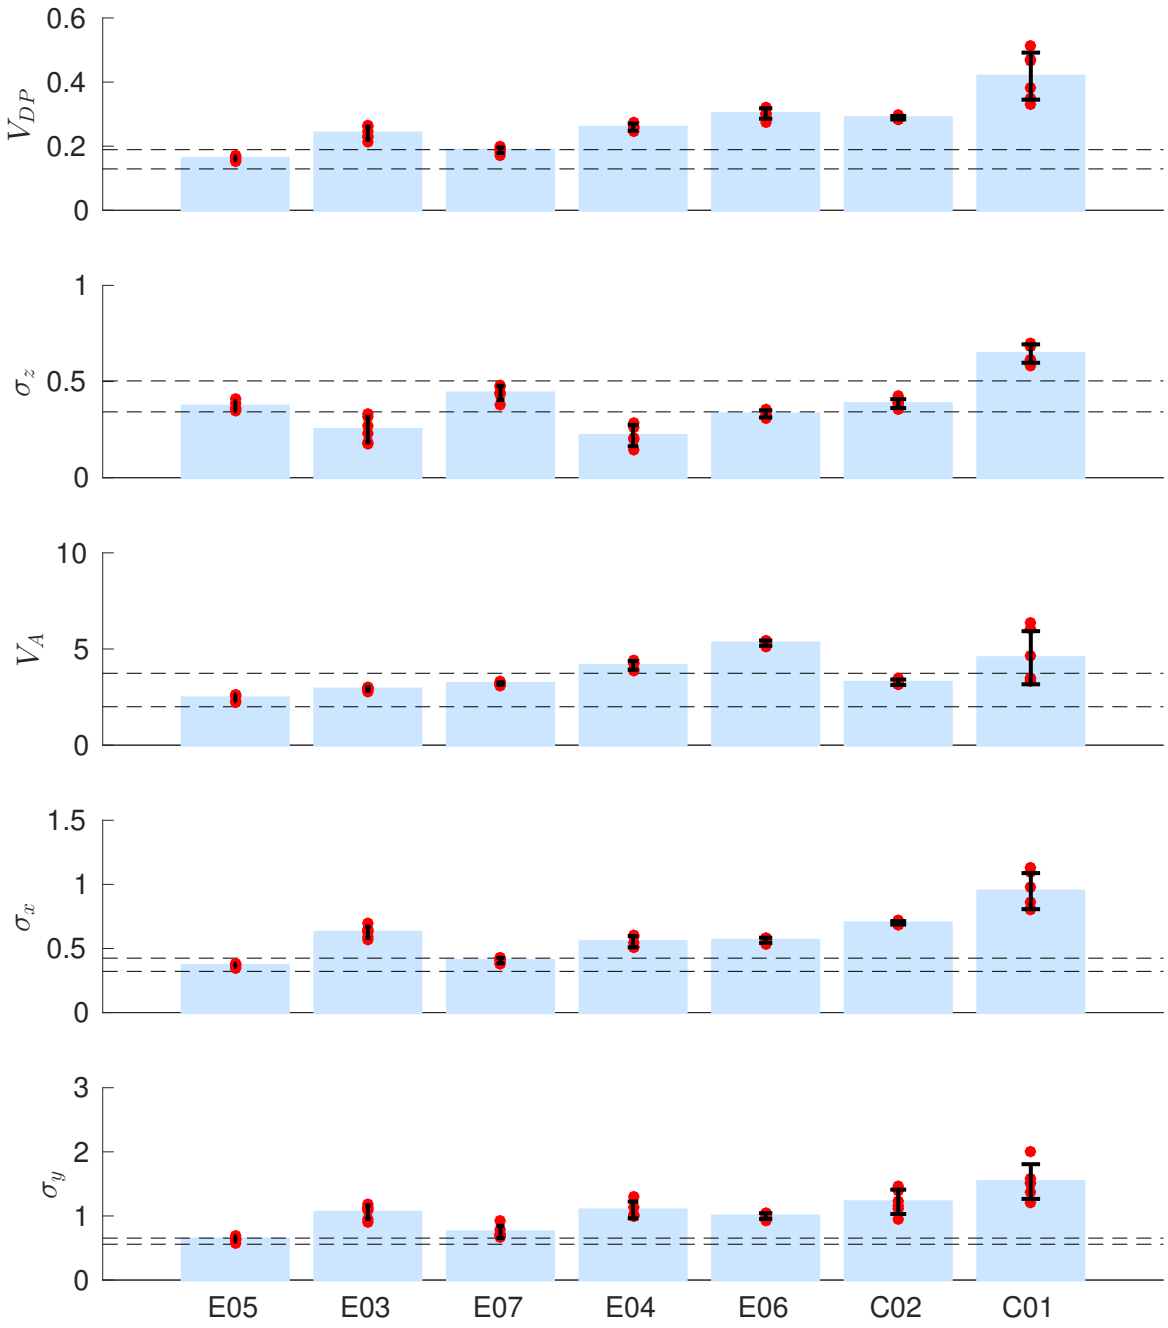

Figure 6.9: Plots showing the estimated parameters for each of the older volunteers. The red spots show the values for each of the 6 repeats of the washout procedure, the height of the bars their means, and the error bars are ± one standard deviation from the mean. The volunteers are ordered from smallest to largest alveolar volumes, with the two volunteers with moderate (GOLD stage 2) COPD on the right of the figure. For all the estimates in this plot V_{DC} =0.155 l, $\rho = 0.8$, $V_t = 0.25$, and Q_T was set to give an arteriovenous O_2 concentration difference of 50 ml O_2/l of blood. The horizontal dashed lines show the largest and smallest mean value from the healthy volunteers.

Identifier	V_{DP}	σ_z	V_A	σ_x	σ_y
E03	8.4%	26.4%	3.4%	7.1%	10.1%
E04	4.4%	25.3%	5.5%	8.0%	11.9%
E05	3.6%	7.1%	7.2%	3.5%	6.5%
E06	5.4%	6.0%	4.3%	2.0%	15.6%
E07	4.6%	8.4%	2.2%	5.1%	12.8%
C01	17.5%	7.4%	30.4%	14.9%	17.6%
C02	1.9%	6.0%	4.3%	2.0%	15.6%
Average	6.5%	12.3%	7.9%	6.3%	11.3%
Young Average	5.5%	6.1%	4.1%	6.9%	44%

Table 6.8: Coefficients of variation (standard deviations divided by the mean) from 6 repeats on each of the older volunteers. For all the estimates in this table V_{DC} =0.155 l, $\rho = 0.8$, $V_t = 0.25$, and Q_T was set to give an arteriovenous O_2 concentration difference of 50 ml of O_2/l of blood. The values in the final two rows are the averages for this group and the healthy volunteers in that order.

Two exceptions to this similar repeatability are the values of σ_z in E03 and E04, from whom our estimated values of σ_z are small (this will be discussed further in the next section), and all the estimated parameters in C01.

C01 is the only volunteer in our study for whom we estimated different parameter values from their two visits. With values of V_{DP}, V_A, and σ_x all noticeably larger on the first visit. Figures B.12 and B.13 show data from C01's two visits. From these two plots it is apparent that not only is the quality of the fit better on the second visit, but also the experimentally measured expirograms are quite different. In particular, the difference in emergence times of alveolar CO_2 and O_2, that we noted in the previous section, is very visible on the first visit, but much less so on the second visit. There was no significant difference between the recent symptoms he reported on the two visits. However, his measured FEV_1 and FVC were both markedly lower on the second visit, by 14% and 24% respectively. The results from this volunteer will be discussed further in the next chapter.

6.4.2 Derived parameters

Figure 6.10 shows the values of the derived parameters for the older volunteers compared to the range of means from the young healthy volunteers.

We noted in the previous section that we estimated what looked like relatively large values of V_{DP} in the older group. However, once we normalise this by alveolar volume, only one of the healthy older volunteers (E03) has a value for $\frac{V_{DP}}{V_A}$ which is outside the range of the young healthy volunteers. Both of the volunteers with moderate COPD have values outside the healthy range.

Similarly, when we calculate the values of $\sigma_z \frac{V_{DP}}{V_A}$ (which is the width of the deadspace distribution expressed in units of ml deadspace/ml alveolar volume), all the older healthy

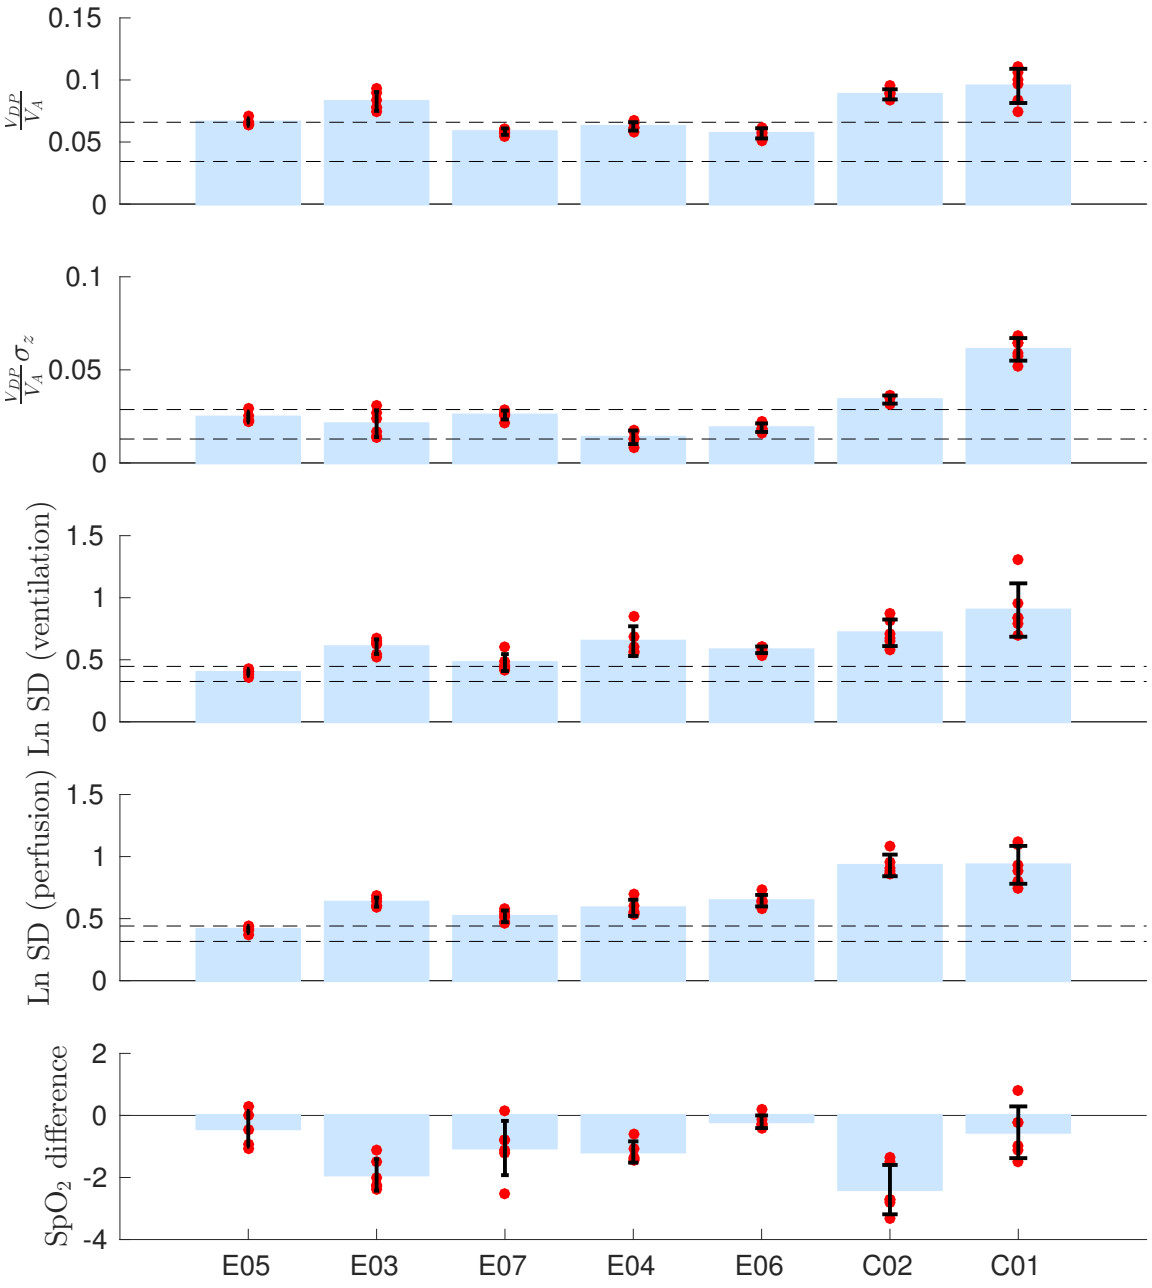

Figure 6.10: Plots showing a number of measures derived from our estimated parameters from the older volunteers. The meaning of these measures and how they are derived is discussed in the main body of the text. The red spots show the values for each of the 6 repeats of the washout procedure, the height of the bars their means, and the error bars are ± one standard deviation from the mean. The volunteers are ordered from smallest to largest alveolar volumes, with the two volunteers with COPD on the right hand side. For all the estimates in this plot V_{DC} =0.155 l, $\rho = 0.8$, $V_t = 0.25$, and Q_T was set to give an arteriovenous O_2 concentration difference of 50 ml of O_2/L of blood.

volunteers lie within the range of the young healthy volunteers (even those with small values of σ_z). However, both the volunteers with moderate COPD appear have larger values for this scaled deadspace width.

Values for Ln SD (ventilation) and Ln SD (perfusion) are, for all but one of the older volunteers (E05), outside the range we found in the normal volunteers. However, all but one (E03) lie within the normal range identified by Wagner [125] of 0.4-0.6. In that paper Wagner gives values of around one for those with mild disease. We find values of around 0.9 (apart from Ln SD (ventilation) in C02 which is 0.72). Similar caveats to those we expressed in the case of the young healthy group, the uncertainty in σ_y (although this is lower here) and the influence of ρ, apply here, but this agreement is promising.

Also plotted in this figure are the differences between measured and estimated S_pO_2 values for this group. Of the 41 repeats in this group, 8 of the estimated S_pO_2 values are more than 2% from the measured values. Only one of those is (just) more than 3% away. However, as can also be seen in this figure, in 38 of the 41 cases the model estimates a value of S_pO_2 that is greater than the measured value. This is not the case in Figure 6.6, where there appears to be no net bias. The overestimates of S_pO_2 suggest that there may be some shunt in this group (mixed venous blood which bypasses the lungs and enters the arterial blood without its concentration being altered). This will be discussed further in the next chapter.

6.4.2.1 Predicted FRC

Figure 6.11 shows the differences in healthy older volunteers between our estimates of FRC, calculated as described for the healthy volunteers, and values predicted (for measurements taken using plethysmography) using formulas taken from [33]. In the male volunteers,

$$\mathrm{FRC}_{\mathrm{predicted}} = -1.215 + 0.0234A + 0.559H^3, \tag{6.10}$$

where H is the height measured in metres and A the age in years. For females the equivalent formula is,

$$\mathrm{FRC}_{\mathrm{predicted}} = 0.563A + 0.4998H^3. \tag{6.11}$$

The 95% confidence intervals on these values are ± 0.81 l and ± 0.67 l for the males and females respectively. All of our estimates for FRC were larger than the predicted values, with three of the five outside the 95% confidence intervals. One volunteer appears to have a very large estimated FRC (over 6.5 l), and thus it appears that our technique may overestimate FRC in this group.

As COPD patients often have enlarged lung volumes due to the emphysematous tissue destruction [109], these formulas do not apply to them. However, we had FRC measurements taken for these volunteers using body plethysmography from a study they had participated in approximately six months before they participated in this study. The

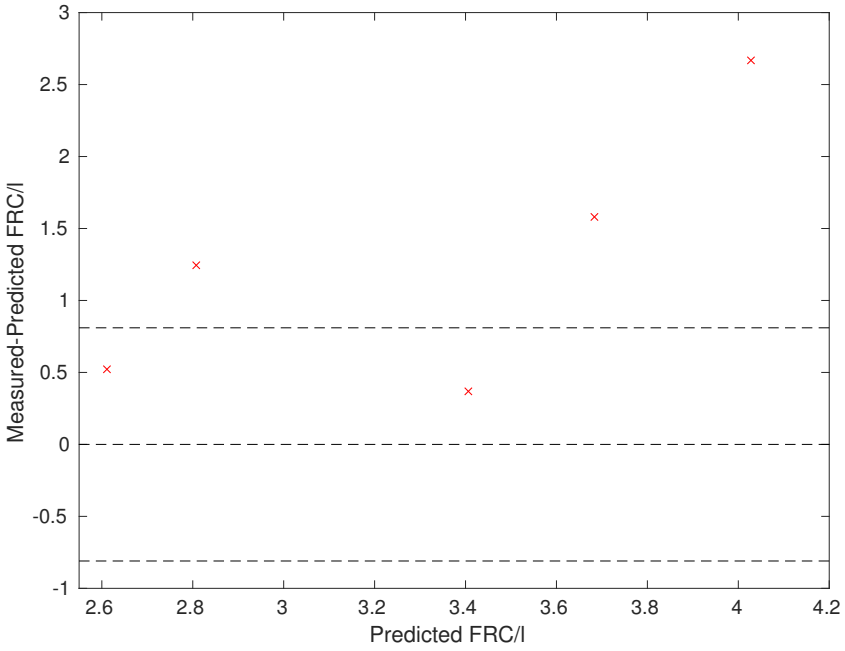

Figure 6.11: A comparison of the value of FRC predicted by (6.10) or (6.11) as relevant, and the value we estimate using our technique. The dotted lines represent 95% confidence intervals.

FRC values measured using plethysmography were 5.08 l for C01, and 4.30 l C02. Our mean estimates for the FRC of these volunteers were 5.93 l (although the mean of the second visit were the model fitted much better was 4.5 l), and 4.26 l respectively. Thus in these patients, where we have a measured rather than predicted value, there does not appear to be the same degree of over-estimation. Measurements of FRC for the rest of this group are currently being organised.

6.4.2.2 LCI values

Table 6.9 shows LCI values calculated for each of the older volunteers. In the healthy older volunteers we find a mean value of 8.2, compared to 6.2 in the young healthy volunteers. In a large study on healthy volunteers from age 25-65 Verbanck *et al.* [120] found an increase of LCI of approximately 0.22 per decade. Our older volunteers are slightly outside this age range, but if this increase is extrapolated, based on our mean value in the young volunteers we would expect a mean value of 7.3 in the older group. But given the low number of volunteers, wide standard deviations, and that we do not control for tidal volume the difference between this number and our mean value for the older group does not seem unreasonable.

Bell *et al.* [13] and Husemann *et al.* [54] report LCI values of 12.23±2.67 and 11.9±1.7 from nitrogen washouts performed on 10 and 20 subjects with COPD respectively. Our estimates for C02 agree well with these values, while C01 appears to fall closer to the

Identifier	LCI
E03	10.0 ± 0.79
E04	5.8 ± 0.69
E05	8.4 ± 1.10
E06	8.6 ± 0.49
E07	8.1 ± 0.7
C01	8.7 ± 1.3
C02	12.4 ± 1.3
Older Average	8.2 ± 1.5
Young Average	6.2 ± 0.8

Table 6.9: Estimated values of the LCI (means $\pm$ standard deviations) from 6 repeats on each of the six healthy volunteers. The values in the final two rows are the averages for the older and young healthy volunteers in that order.

healthy older group. However, this volunteer's tidal volume was around 2 l, significantly in excess of the recommended value of 1 l for measuring LCI [95]. Since the value of LCI has been shown to be dependent on tidal volume [37], this may be a confounding factor here.

Interestingly $E03$ who has the highest value for the LCI among the healthy older group is also most similar to the two volunteers with COPD on a number of our measures, for example: σ_x, $\frac{V_{DP}}{V_A}$, and $\sigma_z \frac{V_{DP}}{V_A}$.

6.5 Discussion

In Section 6.3 we saw that the output of our model (with fitted parameters) agrees well with experimental data taken from young healthy volunteers. These parameter estimates appear very repeatable for repeated nitrogen washout procedures performed on the same volunteer. Furthermore, these parameters, and a number of measures derived from them, show good agreement with measurements taken with other modalities. These results suggest that our model describes this group well.

In older volunteers the agreement between output of our model and experimental data was, in general, still good. However, it showed more variability between volunteers. One feature we noticed in the experimental data was a difference in the emergence times of alveolar CO_2 and O_2. We suggested this might be an indication of serial inhomogeneity (diffusion limitation) in the alveoli. Despite this, the estimated parameters again seemed repeatable, and showed a number of differences from the healthy young volunteers that we discuss in more detail below. It does however appear that, in this group, estimated values of FRC and S_pO_2 may be too high.

To check that this was not an effect of the number of compartments we used approximating the tails of the underlying continuous distributions poorly, we estimated parameters for a number of the datasets from the older group using $7 \times 7 \times 7$ and $9 \times 9 \times 9$

compartments. This did not significantly change the estimated parameters or S_pO_2 values from those found with $5 \times 5 \times 5$ compartments. We also simulated data with $15 \times 15 \times 15$ compartments using parameters estimated with $5 \times 5 \times 5$ compartments. When we did this the estimated S_pO_2 value changed by less than 0.1%, suggesting again that increasing the number of compartments has limited effect. These results suggest that the errors are not due to discretisation of the continuous distribution into a finite number of compartments.

To reduce the estimated value of S_pO_2, a greater fraction of the perfusion needs to go to areas with low ventilation to perfusion ratios. One way this could happen is if correlation between ventilation and perfusion is reduced. Thus we tried estimating parameters at $\rho = 0.6$ in volunteer C02. Reducing the value of ρ reduced the estimated saturation by approximately 0.6%, but it worsened the fit quality (as judged by the minimum value of the objective function). It is then possible that the low S_pO_2 values may be indicative of a deviation from lognormality of the ventilation-perfusion distribution. For example, there may be a second "shunt like" mode with a substantial fraction of perfusion at a very low ventilation-perfusion ratio.

Table 6.10 compares the mean values of a number of the different parameter estimates from the model for the three different groups of volunteers. Given the small number of volunteers in this study it is difficult to draw firm conclusions. Nevertheless, we discuss some of the differences between the groups that appear to be suggested below.

Comparison between the young and older volunteer groups suggested elevated values of the widths of the compliance and perfusion distributions (σ_x and σ_y respectively) in the older group, with the values for the two volunteers with COPD being the largest. As we discussed in Chapter 2, COPD is associated with emphysematous tissue (and pulmonary capillary) destruction [48], leading to reduced elastic recoil, and thus increased compliance in the affected regions [60]. If these changes do not occur completely uniformly across the lung, then both compliance and perfusion would indeed be more inhomogeneous in volunteers with COPD.

As has been frequently observed lung function declines with increasing age [55, 89, 121]. In particular, as shown in recent HP MRI studies [89], alveolar spaces appear to be enlarged with age (without the accompanying destruction of the alveolar wall seen in COPD), leading to a reduction in elastic recoil [121], although these changes are milder than those seen in COPD [55]. Thus widths of the compliance distribution between the healthy young volunteers and the volunteers with COPD seem reasonable for the older healthy group. In our study we observed that the widths of this distribution for the two elderly female volunteers lay within the bounds of the means of the young healthy volunteers, while none of the male volunteers did. However, with such a small number of volunteers it is difficult to comment on the significance of this result.

The other major differences we observed between the groups were in the deadspace parameters. We found that estimated values of both $\frac{V_{DP}}{V_A}$, and $\sigma_z \frac{V_{DP}}{V_A}$ (the width of

	Young	Older	COPD
LCI	6.2 ± 0.8	8.2 ± 1.5	10.6 ± 2.6
FEV_1	$100 \pm 24\%$	$100 \pm 24\%$	64 ± 15
σ_x	0.37 ± 0.04	0.51 ± 0.09	0.82 ± 0.1
σ_y	0.70 ± 0.3	0.90 ± 0.2	1.55 ± 0.2
$\frac{V_{DP}}{V_A}$	0.055 ± 0.006	0.065 ± 0.009	0.092 ± 0.004
$\sigma_z \frac{V_{DP}}{V_A}$	0.023 ± 0.005	0.021 ± 0.005	0.047 ± 0.01
Ln SD (ventilation)	0.50 ± 0.2	0.54 ± 0.07	0.91 ± 0.1
Ln SD (perfusion)	0.43 ± 0.09	0.56 ± 0.08	0.91 ± 0.1

Table 6.10: An illustrative comparison between a number of the different parameters estimates from our model and two existing measures of lung function (LCI and FEV1) for the three groups of volunteers. Quoted numbers are means plus/minus standard deviations.

the deadspace distribution in units of ml deadspace/ml alveolar volume), were elevated in the volunteers with COPD compared to both sets of healthy volunteers, who had similar values. When we developed our model of the deadspace, we assumed that the deadspace represented the conducting airways of the lung. As we discussed in Chapter 2, one of the pathological features of COPD is narrowing of the small airways [67]. Thus if the conducting airways narrow, we might expect the volume of deadspace to decrease, contrary to what we observe in the two volunteers with COPD.

One possible explanation for this, which we discuss further in the next chapter, is that our assumption that we can ignore any diffusion limitation in the alveolar compartments, and assume that they are perfectly mixed is not valid in these volunteers. The closest analogue in our model of serial inhomogeneity in the alveoli is a larger volume of deadspace, and thus enlarged deadspace values may be related to the failure of this assumption. In the next chapter we suggest a possible method for incorporating serial inhomogeneity in our model.

Before presenting a summary of the work discussed in this thesis, and suggesting future avenues for investigation, we briefly discuss our attempts to estimate a representative value for V_t, the equivalent CO_2 tissue volume.

6.6 Tissue volume

As we saw in Chapter 5, V_t, the effective CO_2 tissue volume could not be estimated accurately from nitrogen washout data. However, this did not prevent the accurate determination of the airway parameters. Nevertheless, in Section 5.5 we suggested an experimental procedure, a short CO_2 washin, from which a representative value for this parameter might be set.

On each visit each volunteer performed one of these CO_2 washin procedures. In this section we analyse the results of these procedures for the young volunteers. We consider

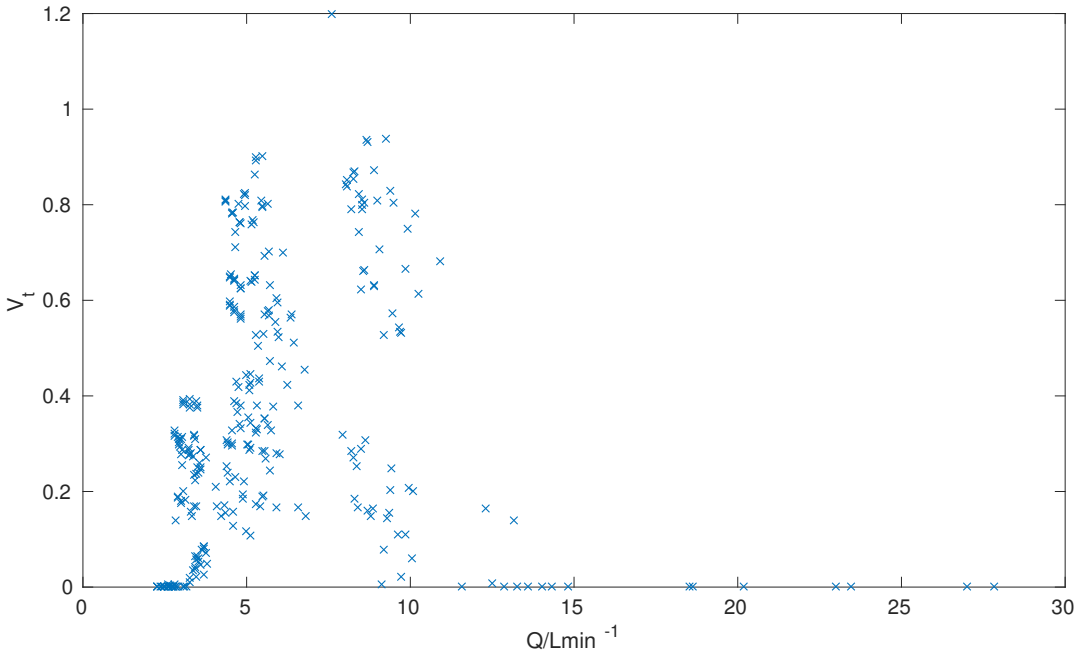

Figure 6.12: Values of the estimated tissue volume and cardiac output for all the healthy volunteers from the CO_2 step procedure.

only the young volunteers as we have more confidence in the values of the parameters that we estimated from the N_2 washout procedure for these volunteers than we do for the older volunteers.

In this section we discuss the use of these results, and estimates based on data taken from the literature, to set a value of V_t for use in future N_2 washout analysis.

6.6.1 Experimental results

We showed in Section 5.5 that there was a significant degree of dependency between all the different parameters in the CO_2 washin procedure. There we attempted to estimate the values of all eight parameters of our model, which we found was not possible. However, we found that if we set the values of all the parameters except Q_T and V_t we could estimate a value of V_t from noisy synthetic data, albeit with a co-efficient of variation of approximately 10%.

This of course relies on us having values for the other six parameters (V_{DP},V_A, σ_x, σ_y, ρ, and σ_z). While we do not of course know the 'true' values of these parameters for volunteers tested on the LGA, we do have estimates from the nitrogen washout procedures they performed. Thus for each of the CO_2 washin procedures we attempted to estimate values for V_t and Q_T using values of V_{DP}, V_A, σ_x, σ_y, ρ, and σ_z, taken from each of the 27 parameter estimation procedures (the 3×3 grid of varying ρ and arteriovenous O_2 concentration difference on each of the three repeats) performed on a volunteer's nitrogen washout data from the same visit.

Volunteer	Visit	Estimated V_t	Standard deviation	Number of estimates
Y03	1	0.70	0.16	8
Y03	2	0.65	0.48	3
Y04	2	0.21	0.13	6
Y06	1	0	0	1
Y06	2	0.21	0.06	18
Y07	1	0.30	0.18	8
Y07	2	0.80	0.08	11
Y08	1	0.62	0	1
Y08	2	0.54	0.14	11
Average		0.45	0.28	

Table 6.11: The means and standard deviations of the estimated CO2 tissue volumes where the estimated value of Q_T was within 1 l/min of the value used in the estimation of the parameters from the nitrogen washout procedure (set based on the arteriovenous concentration difference).

Figure 6.12 shows all the estimated values of V_t and Q_T. Clearly there is a large degree of variation in the estimated values of V_t, with a large number of the zero values for V_t associated with relatively extreme values for the cardiac output. This suggests that the difference between the estimated cardiac output and that used in the estimation of the parameters from the nitrogen washout procedure might provide a method of filtering some of the more extreme values, and thus reducing the variation in the estimates of V_t.

To produce Table 6.11, we considered only the estimated values of V_t where the estimated value of the cardiac output was within 1 l/min of the value used in the estimation of the other parameters from the nitrogen washout procedure (set based on the arteriovenous concentration difference). Where such estimates existed, the means and standard deviations for each volunteer on each visit are shown. For each volunteer on each visit the maximum possible number of estimates is 27. Thus, as can be seen from the final column, a large number of estimates are rejected by this filter. Despite this, there is still significant variability in the estimated values of V_t (even for the same subject).

This variability is not surprising, as this method relies on us knowing the true values for all the parameters other than Q_T and V_t, but in practice our estimates of these true values have some uncertainty. While the airway parameters (V_{DP}, V_A, σ_x, and σ_z) are relatively precisely estimated from the nitrogen washout procedure, there is some variation in the estimated values of these parameters. In addition, σ_y is not very precisely estimated from the nitrogen procedure, and we use a range of fixed values of ρ as this parameter cannot be determined from the nitrogen washout procedure. Another source of uncertainty is that the "true" value of the parameters may vary between repeats on the same volunteer (for example, V_A vary due to changes in position). Given the dependency between the parameters we saw in Section 5.5, it is not surprising that the uncertainty in the true values of the other parameters leads to variability in the estimated value of V_t.

Authors	Estimated V_t	Citation
Dubois *et al.*	0.2	[25]
Fenn and Dejours	0.13	[29]
Fenn and Dejours	0.56	[29]
Plewes *et al.*	0.53	[86]
Sackner *et al.*	0.41	[100]
Average	0.37	

Table 6.12: Estimated values of V_t based on data in the literature. For more details see the main body of the text. The second value attributed to Fenn and Dejours is based on data taken from unpublished data from Dubois.

As shown in the table the average of the filtered values of V_t across all the volunteers and visits was 0.45, although given the large degree of variability this is at best a very approximate estimate for V_t. In the next section we compare it to values taken from the literature.

6.6.2 Estimated values from the literature

Estimates for the value of the effective CO_2 tissue volume, V_t, calculated from data in the literature are shown in Table 6.12. How the effective lung tissue volume is parameterised varies from study to study, and thus in the rest of this section we explain how we convert these different parametrisations to get a value for V_t.

Dubois *et al.* [25] estimated a value for the ratio $\frac{V_A}{\mathrm{ELV}}$, where ELV is the effective lung volume for CO_2 (in our notation $(1 + V_t)V_A$). They did this by comparing the rates at which CO_2 and O_2 alveolar partial pressures equilibrate with the venous partial pressures during a breath hold. They report an average value for this ratio of 0.83, i.e.

$$\frac{V_A}{(1 + V_t)V_A} = 0.83, \tag{6.12}$$

re-arranging this gives an estimate of $V_t = 0.20$.

Fenn and Dejours [29] use a similar method to get a value of 0.88 for the ratio. Thus from their data $V_t = 0.13$. They also report unpublished data from Dubois which gives a value for this ratio (calculated at the end of an inspiration) of 0.64. This gives $V_t = 0.56$.

Plewes *et al.* [86] studied the CO_2 dissociation curve in excised dog lung tissue. They estimated that for a lung volume of 3l the change in volume of CO_2 per Torr of change in the partial pressure of CO_2 was 3.48 ml CO_2/Torr, of which 1.2 ml CO_2/Torr was in the lung tissue and thus 2.28 ml CO_2/Torr in the air space. Dividing these two values gives $V_t = 0.53$.

Sackner *et al.* [100] estimated the pulmonary tissue volume by studying the uptake of N_2O, and carefully estimated the CO_2 dissociation curves in lung tissue and blood in the same six volunteers. They estimated a pulmonary tissue volume of 438ml, and a volume

for the pulmonary blood of 79 ml. To estimate V_t from their data, we: (i) multiply these volumes by the values from the CO_2 dissociation curve for the change in volume of CO_2 per unit of volume of lung tissue (or blood) per Torr of partial pressure; (ii) sum the resultant values to get an estimate for the change in volume of CO_2 per Torr of change in the partial pressure of CO_2 in the lung tissue; and (iii) divide this by the change in volume of CO_2 for the change in partial pressure of CO_2 in the airspaces of the lung. Using their data, the change in volume of CO_2 per Torr of change in the partial pressure of CO_2 in the lung tissue is

$$438 \times 0.003 + 79 \times 0.0038 = 1.61\text{ml } CO_2/\text{Torr}. \tag{6.13}$$

If we assume a lung volume of 3 l (at BTPS), then from Dalton's Law of partial pressures, the value in the airspaces of the lung is

$$\frac{3000}{760} = 3.95\text{ml } CO_2/\text{Torr}, \tag{6.14}$$

where we have divided by 760 Torr, as this is atmospheric pressure in Torr. Dividing these two values gives $V_t = 0.41$.

6.6.3 Summary

Both the estimates based on the CO_2 washin procedure and those based on data taken from the literature show a large degree of variability, with all the authors cited in the previous section noting, as we do, that their methods give approximate rather than definitive values. Nevertheless, as Sackner notes [100] there is very clear evidence that this parameter is not zero.

As we noted in Chapter 5, the value of V_t does not influence the estimation of the airway parameters, and thus this uncertainty in the true value of V_t is not of concern here. The goal of this section was to estimate a physiologically reasonable value for V_t that could be used in future N_2 washout procedures. Given the means of the estimates from our data and the literature are 0.45 and 0.37 respectively, a value of 0.4 seems reasonable. However, it should be clear from the discussion in this section that this is a very rough estimate.

Chapter 7

Summary and future work

This chapter summarises the work presented in this thesis, and suggests some future directions for work, both experimental and theoretical, in this area.

7.1 Summary

The premise of this thesis is that there is information in the transient expired CO_2, O_2, and N_2 concentrations (both during air breathing and the nitrogen washout) that is currently not used in clinical assessment of lung function, and that the appropriate use of such information could yield a clinically useful measure of the inhomogeneity of the lungs. Measuring these transient concentrations, and the flows of these gases accurately is very difficult [17, 96]. Recently a novel device (the LGA) which measures these concentrations and flows very accurately has been developed [17].

I proposed a method which involves estimating the parameters of a novel mathematical model from high quality nitrogen washout data, and using the parameters of that model as indices of the inhomogeneity of the lungs. Thus the central research question of this thesis was: "Can we develop a novel mathematical model of the lungs, whose parameters we can estimate from high quality nitrogen washout data provided by the LGA, to provide a sensitive measure of the inhomogeneity of the lungs?"

I established in Chapter 4 that the model would need to represent the underlying physiology well, have a limited parameter set, and be described by governing equations that could be solved quickly and accurately. In the rest of that chapter I developed a model with these characteristics.

I explained that our conceptual model of the lung was composed of a large number of functional lung units (described in Chapter 2). I adapted work presented by Wilson and Beck [139], and assumed that the volumes, fractional volume changes, and perfusions of these units were governed by a bivariate lognormal distribution, which I showed was described by only three parameters.

I developed an efficient method for discretising this distribution into a finite number of alveolar compartments. I wrote down the governing equations of these compartments;

and explained that, as I include gas exchange in the equations for alveolar volume change, fractional ventilation cannot be a constant. Instead, I defined the fractional volume change of a compartment, which ensured that all compartments' volumes remained physically realistic.

The conducting airways that connect the alveoli to the mouth were approximated with deadspace compartments that I assumed were uniform cylinders. Further, on the basis of both experimental (Section 2.5.1.1) and theoretical (Section 2.2.2) evidence, I assumed all gas transport in these deadspace compartments was via convection. Even having made these assumptions there are a number of possible flow patterns. I considered two, plug flow and Poiseuille flow, and implemented both in the model. I found that plug flow more closely matched the signal seen at the mouth, and thus this pattern was adopted.

Finally, I presented efficient methods for solving the governing equations of the model. In the case of the deadspace compartments I explained how, after a change of variables, the method of characteristics can be applied to the governing advection equation to conserve mass with very high precision. For the alveolar volumes I developed a method that considered the temporal resolution of our experimental data, which I found was very important in parameter estimation.

In Chapter 5 I investigated the inverse problem of how to estimate the parameters of the model from experimental data. I described the objective function used to compare simulated and experimental data; and explained that because evaluating it takes approximately three minutes, my approach to finding the set of parameters that minimised the objective function was to use a MATLAB least squares minimisation routine, `fmincon`, implemented in parallel on a single node of the Arcus-B super computer.

A particular problem I identified in the evaluation of the objective function was estimating the venous gas concentrations. These are not measured, but they appear in the governing equations for the alveolar compartments. I noted that the hypocapnia we observed, and the high arterial oxygen concentration during the washout phase, meant it could not be assumed that these venous concentrations were constant. During the air breathing phase, the pulmonary gas exchange on each breath was estimated. I then used a simplified version of the model presented in Chapter 4 to estimate the venous gas concentrations on every breath from the estimated pulmonary gas exchange, and assumed the change of the concentrations with time was linear. Finally, I found starting values for the venous concentration so that the full model matched the total oxygen and carbon dioxide consumptions that were measured over the air breathing period. To allow venous oxygen to rise during the washout phase I proposed a computationally efficient method of parameterising a single compartment model of oxygen distribution in the body.

Next I tested our parameter estimation procedure on synthetic data produced with the model. First I examined noise free data, and found that while all parameters could be identified, the circulatory parameters were much more poorly identified than the airway

parameters. However, an analysis of the dependency between parameters suggested that values of the circulatory parameters should have little effect on the estimated values of the airway parameters.

Based on the criteria defined by the ERS [96] I then established a detailed model for the error in experimental data. I did this both for both the LGA and a device built to the ERS specifications.

When attempting to estimate parameters from a number of datasets, I found that the ERS specifications did not produce data accurate enough for our method to estimate the underlying parameters to within an acceptable tolerance. This result demonstrates the reliance of our technique on very high quality data. For error levels consistent with the LGA, I could not determine the circulatory parameters. This did not prevent the estimation of the airway parameters with, I argued, less uncertainty than might be expected from intrinsic biological variability. Thus, I concluded that the values for the circulatory parameters could be set to physiologically reasonable values, and the airway parameters estimated from data taken with the LGA. Finally, I suggested an experimental procedure which might allow estimation of the equivalent CO_2 tissue volume.

In Chapter 6 I attempted to estimate the parameters of the model from experimental data. Specifically, I considered data taken from 15 volunteers who each performed six nitrogen washout procedures on the LGA over two visits. Six of the volunteers were young and healthy, six over the age of seventy with no reported lung disease, and three had moderate (GOLD stage 2) COPD. In general the quality of the datasets taken from these volunteers was extremely high, as indicated by the equality between inspired and expired nitrogen volumes during room air breathing. However, we were unable to obtain data of sufficient quality (as judged by the nitrogen balance) from two of the older volunteers, who had difficulty complying with our protocol. One repeat from one of the other older volunteers was also excluded due to an unexplained departure from nitrogen balance, which occurred halfway through the procedure.

Based on the results of Chapter 5 I discussed which of the models parameters would be estimated, which would be fixed and at what values. I attempted to estimate values for all of the airway parameters, and the width of the perfusion distribution, while the rest of the circulatory parameters were set at physiologically reasonable values.

The simulated output of the model (with fitted parameters) agreed well with the experimental data from the healthy young volunteers. Furthermore, these parameter estimates (with the exception of σ_y) were highly repeatable. The estimates for FRC, arterial oxygen saturation, and the ventilation-perfusion distribution showed good agreement with expected (or, in the case of arterial oxygen saturation, measured) values. This suggests that in young healthy volunteers our method works well.

In the older volunteers (both healthy and those with COPD) the agreement between the simulated output of the model (with fitted parameters) and experimental data, while

still good in most cases, was more variable. In particular, we observed a difference in the emergence times of alveolar carbon dioxide and oxygen during expiration that is not a feature of the model. I will discuss possible origins of this phenomenon, and how it might be incorporated into the model, in the next section.

Nevertheless, estimated parameters appeared to have a similar repeatability to those of the younger volunteers, although values obtained from one of the COPD patients from their two visits were very different. Large differences were also seen in their spirometry from the two visits, but not in their reported symptoms or recent exacerbations.

Comparing parameter values between the groups yielded a number of differences. The widths of both the compliance and perfusion distributions were greater in the older healthy than the younger group, and greater still in the group with COPD. We suggested that this agreed with known changes to lung structure due to both ageing and disease. However, I also observed large deadspace volumes (when normalised with alveolar volume) in this group which I suggested was a result of model misfit rather than physiology, meaning any conclusions about the deadspace in this group are tentative.

The agreement with other modalities was also more variable in the older group. While estimated ventilation-perfusion distributions still agreed well with those measured using MIGET, the model over estimated arterial oxygen saturation during normal breathing (this finding is discussed in the next section), and appeared to estimate values for FRC significantly in excess of the predicted values.

Measured values of FRC were available for the volunteers with COPD from a prior study they had participated in. In one case the agreement of the estimate of FRC to the measured value was very good, while in the case of the other subject (C01 whose parameter estimates different widely between his two visits) there was reasonable agreement from the first visit but values that were significant overestimates from the second. To assess whether my method does indeed overestimate FRC, full lung function testing for all the volunteers is being organised. Before the results of this are available it is difficult to comment further on this issue. However, it appears possible our method may overestimate FRC. In this chapter I also presented the results of our attempts to estimate a representative value for the equivalent tissue volume for carbon dioxide. Estimated values from this procedure were very uncertain, but a value for this parameter that could be used in the future was tentatively suggested.

In conclusion, the premise of this thesis was that there is information in the transient expired CO_2, O_2, and N_2 concentrations, and appropriate use of this information might provide a useful measure of the inhomogeneity of the lungs. Thus the central research question of this thesis was: "Can we develop a novel mathematical model of the lungs, whose parameters we can estimate from high quality nitrogen washout data provided by the LGA, to provide a sensitive measure of the inhomogeneity of the lungs?"

To that end I developed a model, which tried to balance computational considerations (parameterisability, speed, and accuracy of solution) with physiological accuracy. I then developed a practical approach for estimating the parameters of the model and tested it on synthetic data.

The method was tested in a small group of 15 volunteers. The results from the young healthy volunteers suggest that the model describes this group very well. Applying the method to the older group I observed increased inhomogeneity that appears consistent with known physiology. However, there were a number of inconsistencies in the quality of the fit, and the parameter estimates from this group. These inconsistencies suggest that the model may need further development to give truly sensitive measures of the inhomogeneity of the lungs.

Nevertheless, I feel that the results in this thesis demonstrate that there is useful information in the transient expired CO_2, O_2, and N_2 concentrations; and that the approach outlined here shows some promise as a method of using this information. The next section discusses future research that might allow further assessment of the current model's utility, and extend the model to deal with some of the inconsistencies discussed here.

7.2 Future work

There are a number of possible future avenues of research suggested by my work. Here I outline a few.

7.2.1 Experimental work

As we discussed in Chapter 6, we observed a number of inconsistencies in our results for the older volunteers. One of these was an apparent overestimation of FRC; to assess this properly full lung function testing is being organised. It would also be interesting to obtain measurements of other indices of inhomogeneity, for example LCI, from these volunteers to act as comparison with our measures. More generally, a larger number of volunteers, particularly those with COPD, would allow us to make stronger conclusions, and see how widespread the inconsistencies we noted are.

Other possible studies might involve estimating parameters in patients with cystic fibrosis, where LCI has been shown to be a sensitive measure of the disease [51]. One specific idea is to see whether our method can detect changes before and after physiotherapy received by cystic fibrosis patients.

Another possible experiment involves getting healthy volunteers to breathe to different target tidal volumes, for example 1l, 2l, 3l (if achievable), during the nitrogen washout procedure. We would then estimate parameters at each tidal volume. Here we would be interested in how similar the volumes of deadspace we estimated at these different tidal

volumes were. If they were very different, and in particular increased with tidal volume, this might indicate that one of our assumptions, that the deadspace has a constant volume, might need modification. This procedure could also be used to check if the parameters of the lognormal distribution (e.g. σ_x, the width of the compliance distribution) are really independent of the tidal volume. Conducting this experiment requires some modification of the device so that the flow analysis, which is currently performed off-line, is done in real time.

A useful modification to the device would be reducing the amount of deadspace it introduces. The ERS guidelines recommend instrument deadspace should be less than 70 ml for measurements in adults [96], currently the deadspace introduced by the LGA is 155 ml. Reducing this deadspace would make it easier for all subjects to breathe on the device, and allow its use in a wider range of ages (for example, in children).

Work is currently being conducted to assess the feasibility of incorporating a laser to measure methane concentrations into the device. If this does prove feasible, we could try performing a methane washin, rather than a nitrogen washout. A possible protocol for this would be five minutes of air breathing, and then methane would be introduced into the inspirate instead of some of the N_2 (i.e. the inspired fractions of CO_2 and O_2 would remain the same). Methods of analysis would remain similar, except we could consider CO_2, O_2, and CH_4 fluxes simultaneously, rather than having separate air breathing and washout phases. This procedure would have two big advantages. Firstly it would shorten the measurement period from fifteen minutes to ten minutes. Secondly, it would remove the uncertainty of the physiological effects of the high O_2 concentrations on, for example, the venous concentrations.

Another possible modification which is being investigated is the addition of a laser to measure a more soluble gas (acetylene). It is hoped that this will address one of the key weakness identified in Chapter 5, the inability to identify the circulatory parameters.

7.2.2 Model development

As outlined in Chapter 4 a number of assumptions were made during the development of the model. The strength of the conclusions we are able to make about a subject's lungs are limited by the validity of these assumptions, in this section we briefly discuss a couple of the limiting assumptions.

A method of dealing with one of these assumptions, that high O_2 breathing does not have any impact on physiology, was discussed in the previous section. Another limiting assumption is that we assume a unimodal bivariate lognormal distribution of ventilation and perfusion, when other modalities have demonstrated that (at least in disease) bimodal distributions may exist [123]. Thus an obvious extension to this work would be to consider bimodal distributions. The difficulty here would likely be ensuring that parameters remain identifiable.

In Chapter 6 we noted that in some of the subjects alveolar CO_2 appeared to emerge later than alveolar O_2. We argued that this might be evidence of a diffusive process that is not modelled correctly, and thus that our assumption that the alveoli are perfectly mixed (i.e. diffusion is complete) was not valid in these subjects.

One method of incorporating this serial inhomogeneity in the model would be to adapt the approach Scheid *et al.* [101] applied to a continuous ventilation model. In their approach a membrane is introduced into the alveolar compartment, splitting it into two. Flux across the membrane is assumed to be via diffusion (Fick's law is applied), and both compartments are assumed to be well mixed. Careful consideration of the parametrisation of this model, and how to ensure the compartment's volumes remained physically realistic would be required. It is envisaged that this model would require at least two additional parameters, a diffusive conductance for the membrane, and a parameter to determine the fraction of the compartment's volume either side of this membrane. This would be a way of investigating the influence of serial inhomogeneity without using a full convection diffusion equation, which as we have previously discussed is computationally intensive.

Another approach to investigating the importance of diffusion would be to investigate this phenomenon with a biophysical model of the lungs. Although, as we have discussed, this sort of model would not be appropriate for the estimation of parameters, it would be interesting to see whether biophysical models of diseased lungs showed this difference in alveolar gas emergence times.

Finally, we noted that the model's high estimates of S_pO_2 values in some of the older volunteers might be evidence of a deviation from lognormality of the true underlying ventilation-perfusion distributions. Thus other distributions could be investigated. In particular, multi-modal distributions, like those found in some COPD patients using MIGET [123], might be of interest. However, the difficulty here would probably be parameterisation (ensuring that all the parameters remained identifiable in experimental data).

7.3 Concluding remarks

This thesis has demonstrated that there is useful information in the transient expired CO_2, O_2, and N_2 concentrations, and that the approach of estimating the parameters of a mathematical model to take advantage of this information is promising. Of the parameters of the current model, the width of the compliance distribution (σ_x) appears to be the best candidate as a clinically useful index.

While both the model and the experimental procedure presented here need further refinement, we believe that this thesis has shown such work is worthwhile, and may lead to the development of a useful measure of the inhomogeneity of the lungs.

Appendix A

Glossary of key terms

Symbol	Description
V_A	Total alveolar volume when the lungs are at FRC
V_t	CO_2 equivalent tissue volume
Q_T	Cardiac output
σ_x	Width of the compliance distribution
σ_y	Width of the perfusion distribution
ρ	Correlation of compliance and perfusion
V_{DC}	Apparatus deadspace volume
V_{DP}	Total personal deadspace volume
σ_z	Width of the deadspace distribution

Table A.1: A summary of the parameters that describe our model. V_{DC} is set equal to the measured apparatus deadspace volume of 155 ml, but all the other parameters must be estimated.

Symbol	Description
V_i	The fractional volume of compartment i
ΔV_i	The fractional volume change (compliance) of compartment i
v_i	The fractional deadspace volume of compartment i
Q_i	The fractional perfusion of compartment i
$\dot{V}_i$	The fractional ventilation of compartment i

Table A.2: The fractional quantities that define alveolar compartment i.

Acronym	Definition
ATS	American Thoracic Society
BTPS	The saturated partial pressure of water at body temperature and pressure
ERS	European Respiratory Society
FEV_1	Forced expiratory volume in one second
FVC	Forced vital capacity
FRC	Functional reserve capacity
LCI	Lung clearance index
LGA	Laser gas analyser
ODE	Ordinary differential equation
PDE	Partial differential equation
ET	End tidal (referring to gas fractions)

Table A.3: Frequently used acronyms.

Quantity	Definition
$F_i^g(t)$	The gas fraction g in compartment i at time t
$C_{ai}^g(t)$	The arterial gas concentration of gas g leaving compartment i
$C_{\bar{v}}^g(t)$	The venous gas concentration of gas g
$\dot{V}_{\mathrm{g}}$	The consumption rate of gas g
$u(t)$	The flow rate of gas (measured by the LGA)
$V_{Ai}(t)$	The volume of compartment i at time t
$V_{Ai}^g(t)$	The volume of gas g in compartment i at time t
Δ	10ms time interval between LGA measurements
$V_{bi}^g(t_l)$	The volume of gas g exchanged with the blood over Δ t in compartment i
P_b	Barometric pressure
$P_{\mathrm{H_2O}}$	The saturated partial pressure of water at BTPS
P^g	The partial pressure of gas g
λ_t	The solubility of CO_2 in lung tissue
S_b	The solubility of N_2 in the blood
Q_l	Fractional perfusion of body compartment l
V_{tl}	Fractional tissue volume of body compartment l

Table A.4: Key quantities relating to the alveolar compartments.

Quantity	Definition
$v(r,t)$	The gas velocity at a radius r and time t in a deadspace compartment
A	The cross-sectional area of a deadspace compartment
L	Length of a deadspace compartment
R	Radius of a deadspace compartment
z	Distance into a deadspace compartment
$F^g(z,t)$	The flow weighted average gas fraction at a distance z into a deadspace compartment
$V_D(z)$	Cumulative volume of deadspace at a distance z into a deadspace compartment
$V_{Dg}^g(V_D)$	Cumulative volume of the gas of interest in the deadspace at $t = 0$ (during inspiration) at $V_D(z)$
$V_{in}^g(t)$	Cumulative volume of the g that has left the deadspace by time t during inspiration
F_{DP}^g	Common fraction at the end of the personal deadspace volumes at the end of expiration
$F_M^g(t)$	The fraction of gas g measured by the LGA at time t
$F_S^g(t)$	The simulated fraction of gas g at the mouth at time t
$F_{ET}^g(t)$	The end tidal fraction of gas g

Table A.5: Key quantities relating to the deadspace.

Appendix B

Diagnostic plots

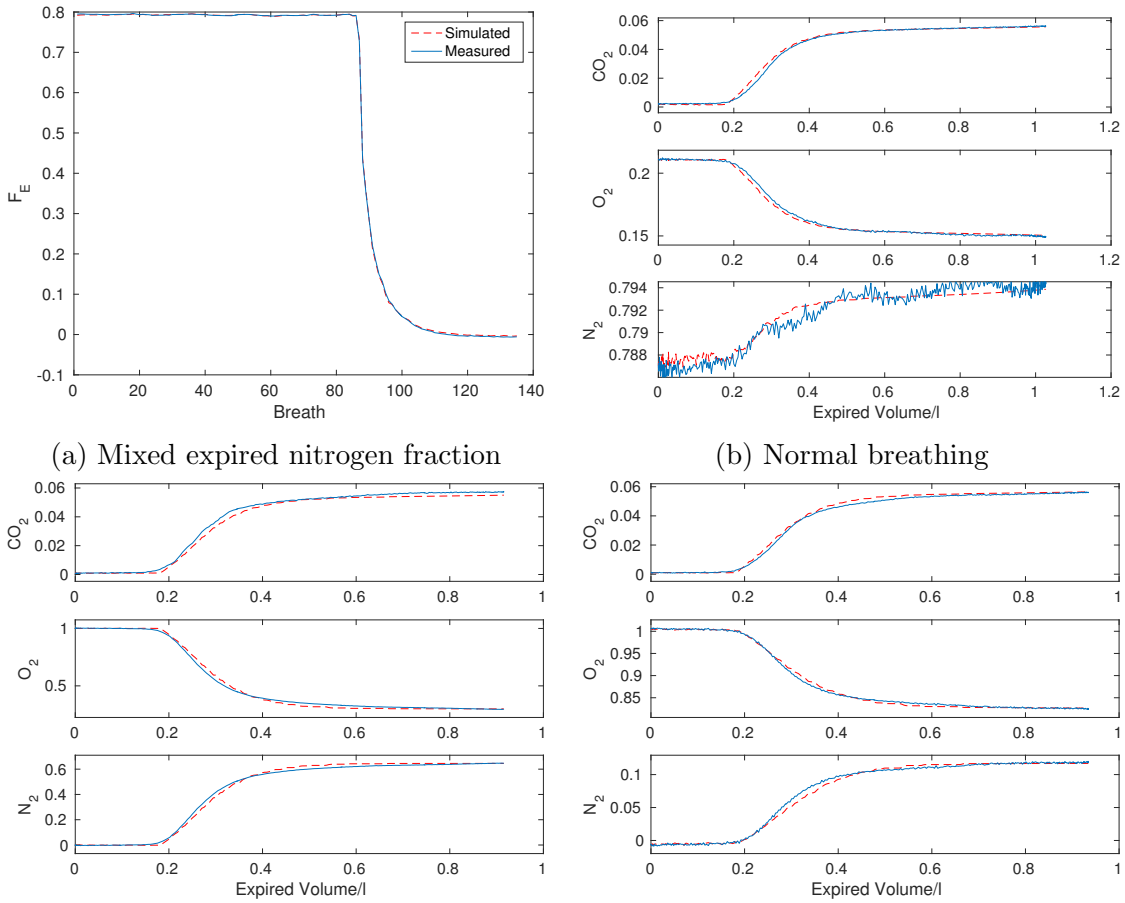

(a) Mixed expired nitrogen fraction

(b) Normal breathing

(c) Washout start

(d) Mid washout

Figure B.1: Example plots of the fit to subject Y03.

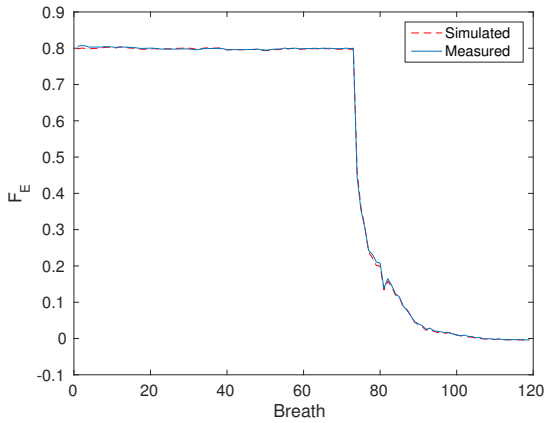

(a) Mixed expired nitrogen fraction

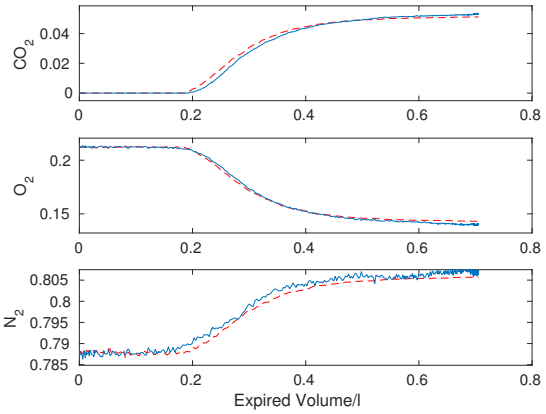

(b) Normal breathing

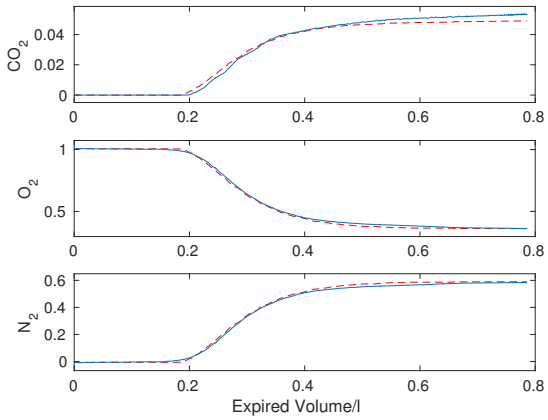

(c) Washout start

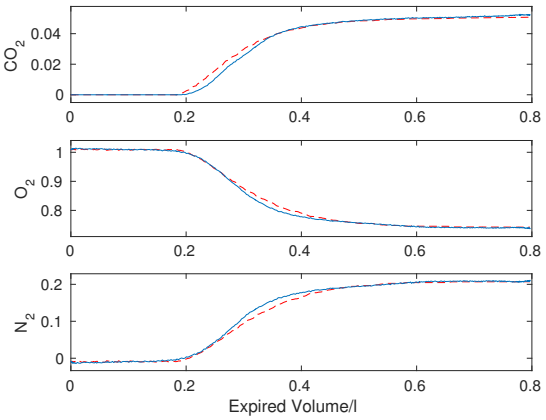

(d) Mid washout

Figure B.2: Example plots of the fit to subject Y04.

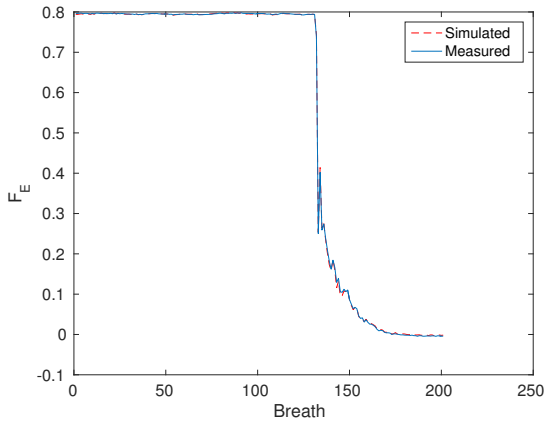

(a) Mixed expired nitrogen fraction

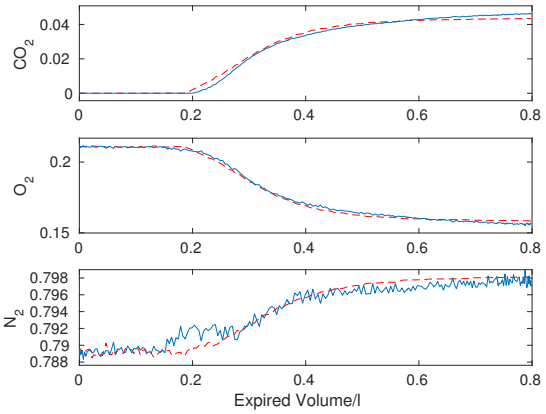

(b) Normal breathing

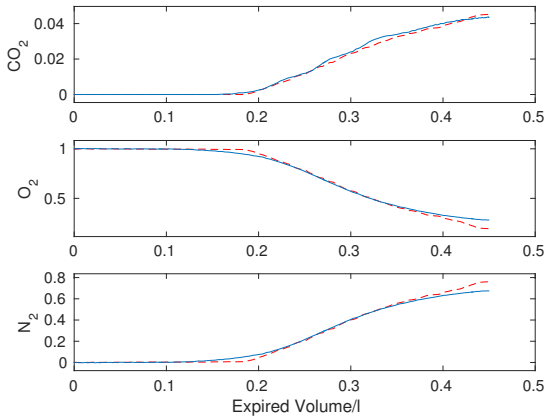

(c) Washout start

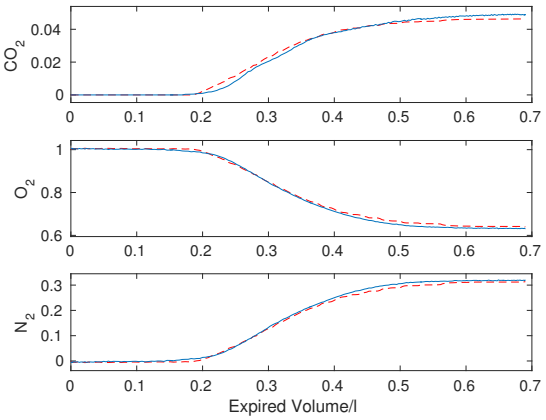

(d) Mid washout

Figure B.3: Example plots of the fit to subject Y05.

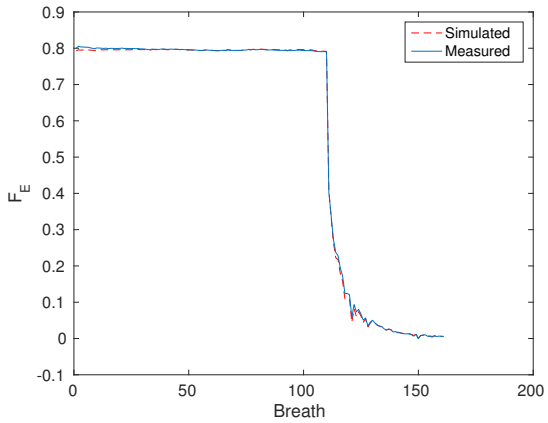

(a) Mixed expired nitrogen fraction

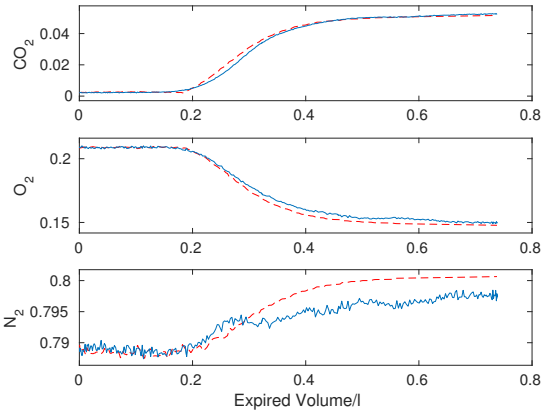

(b) Normal breathing

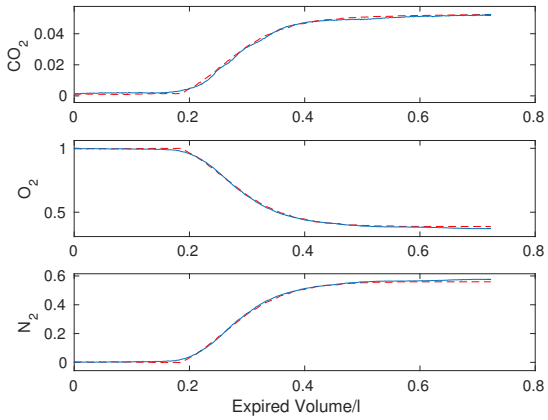

(c) Washout start

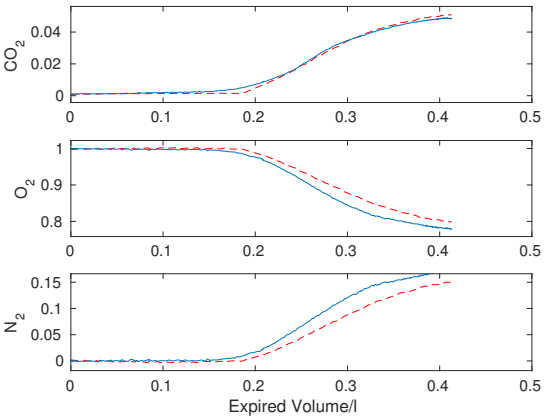

(d) Mid washout

Figure B.4: Example plots of the fit to subject Y06.

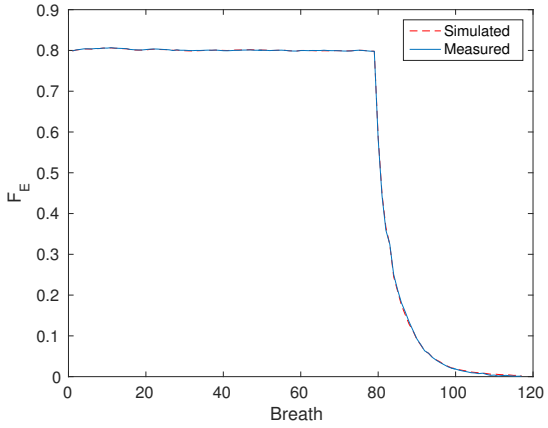

(a) Mixed expired nitrogen fraction

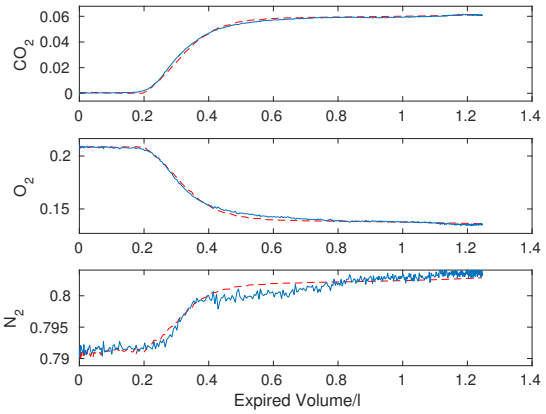

(b) Normal breathing

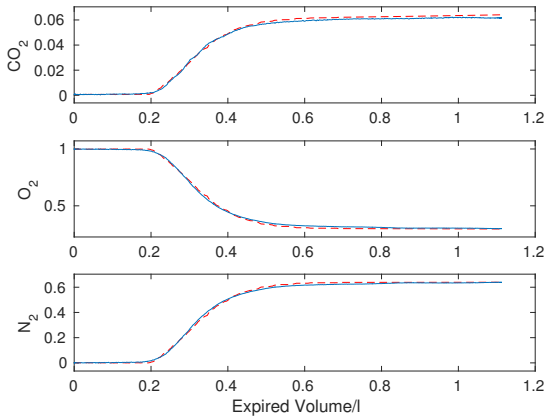

(c) Washout start

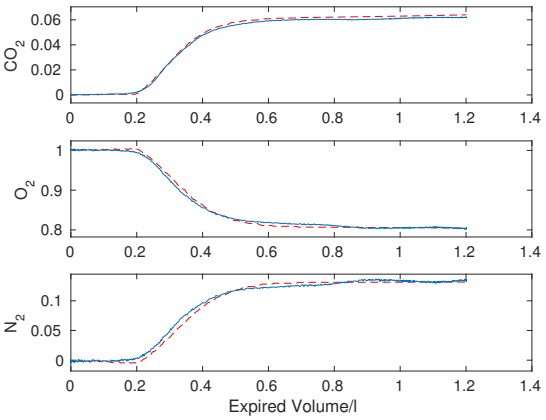

(d) Mid washout

Figure B.5: Example plots of the fit to subject Y07.

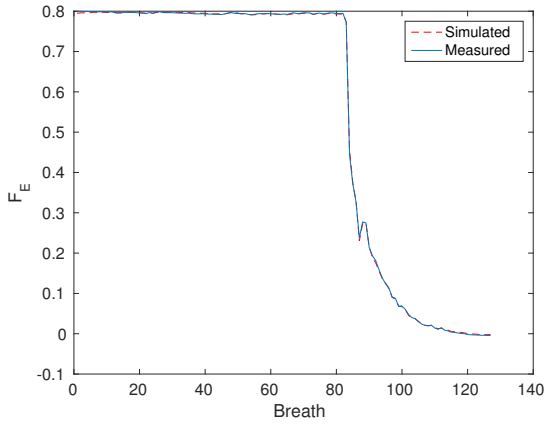

(a) Mixed expired nitrogen fraction

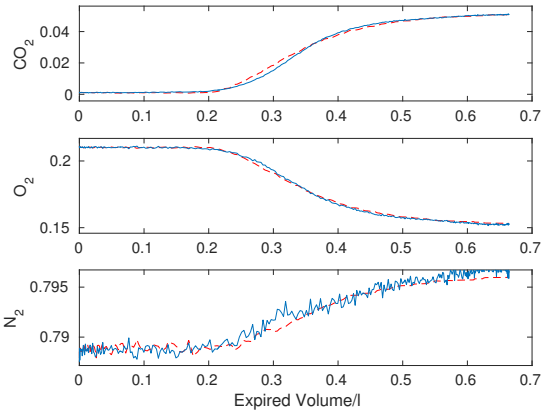

(b) Normal breathing

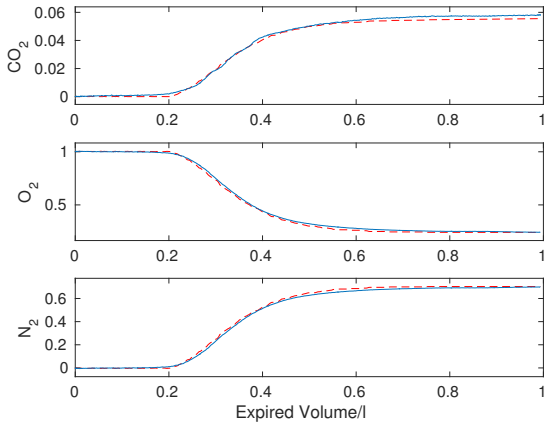

(c) Washout start

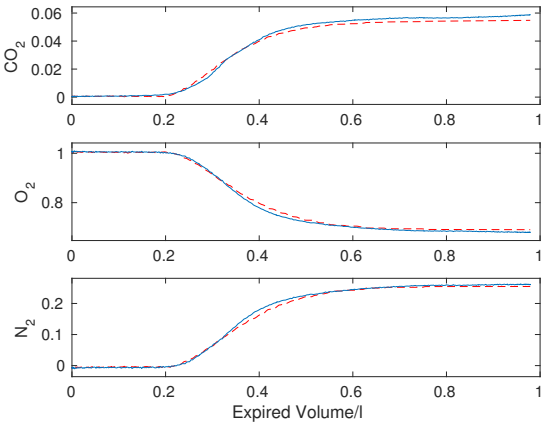

(d) Mid washout

Figure B.6: Example plots of the fit to subject Y08.

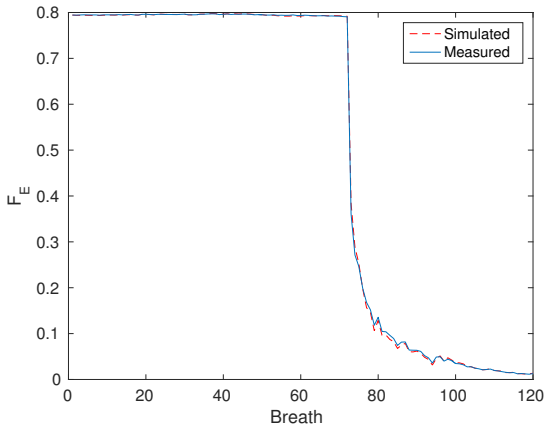

(a) Mixed expired nitrogen fraction

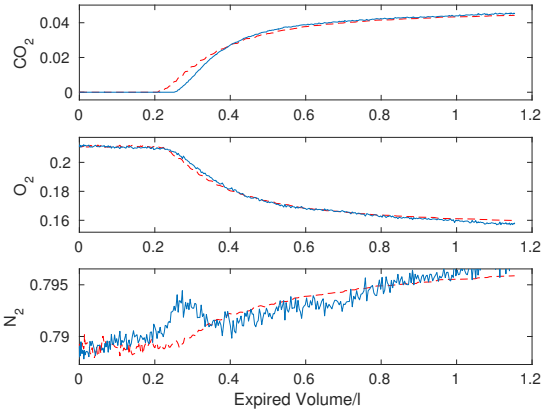

(b) Normal breathing

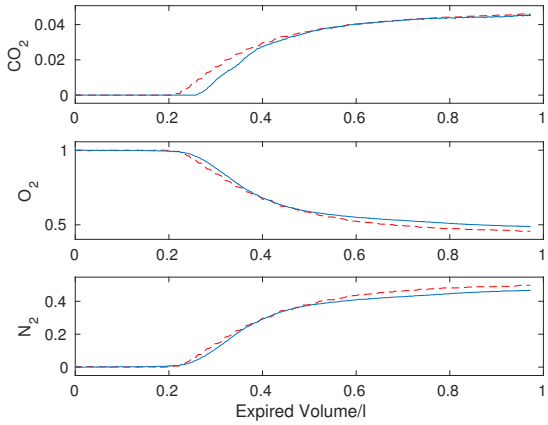

(c) Washout start

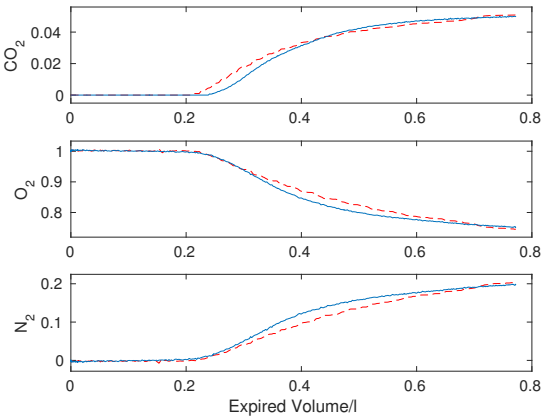

(d) Mid washout

Figure B.7: Example plots of the fit to subject E03.

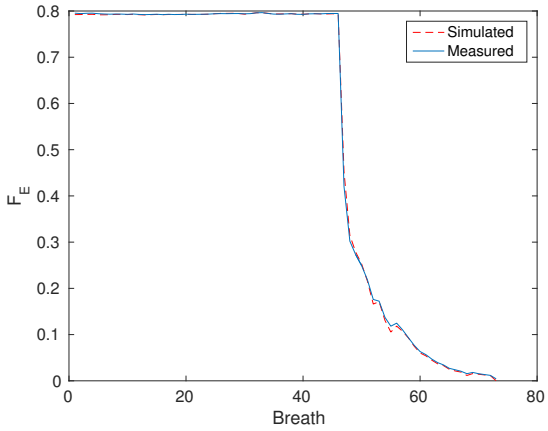

(a) Mixed expired nitrogen fraction

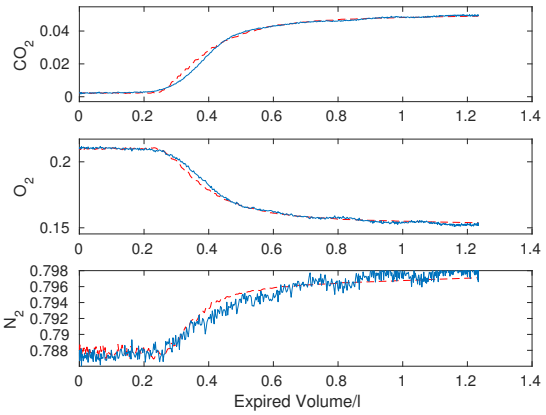

(b) Normal breathing

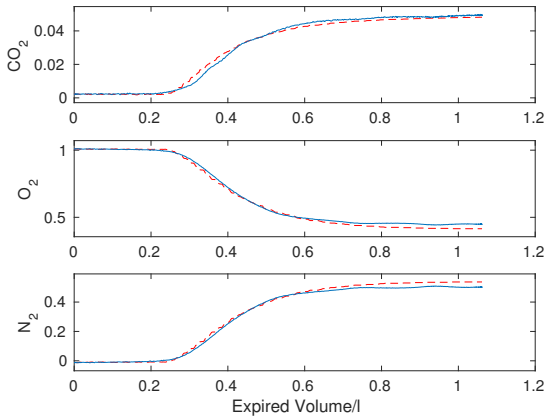

(c) Washout start

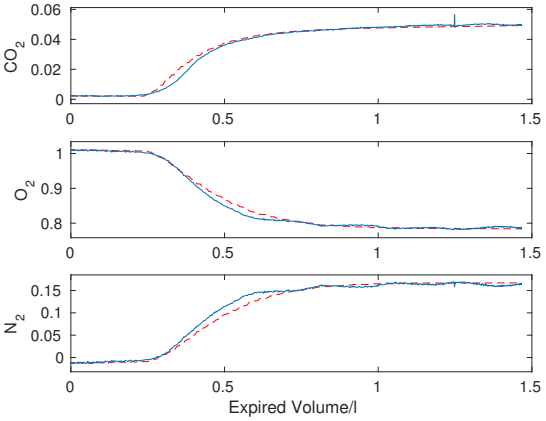

(d) Mid washout

Figure B.8: Example plots of the fit to subject E04.

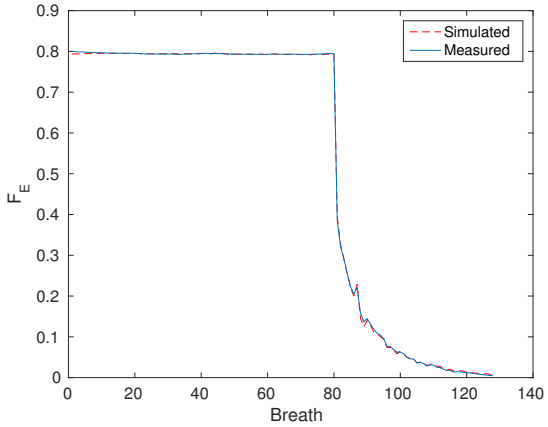

(a) Mixed expired nitrogen fraction

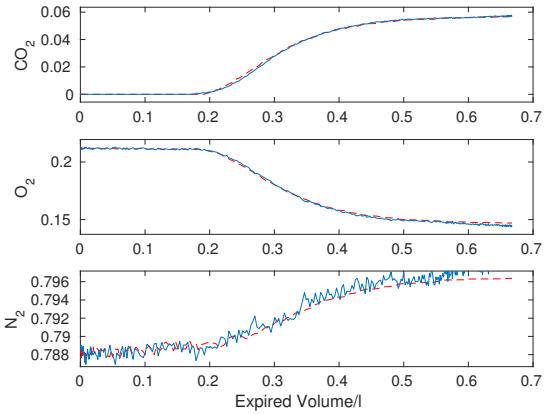

(b) Normal breathing

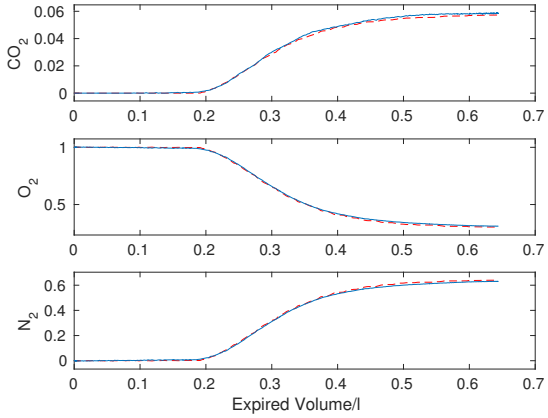

(c) Washout start

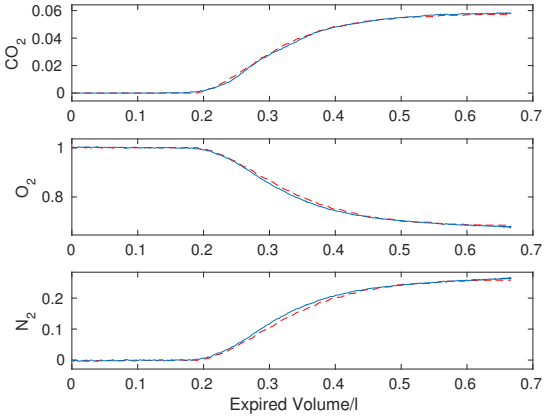

(d) Mid washout

Figure B.9: Example plots of the fit to subject E05.

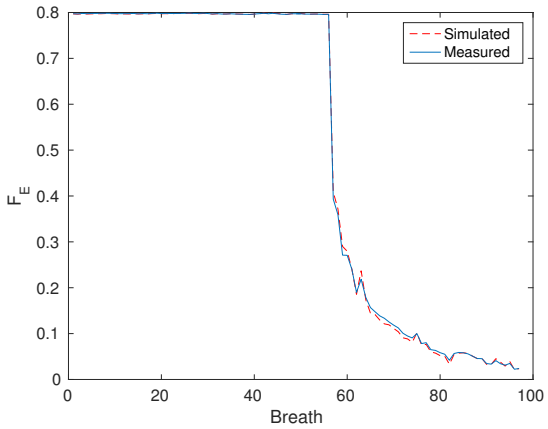

(a) Mixed expired nitrogen fraction

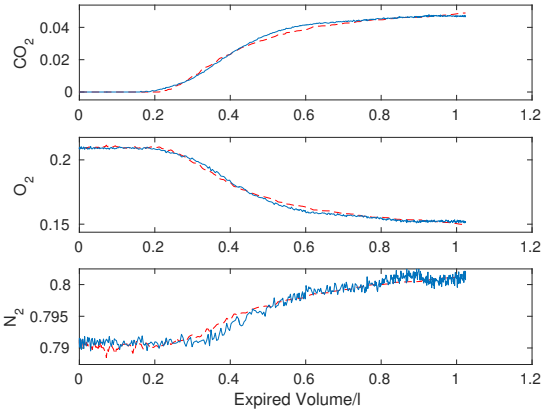

(b) Normal breathing

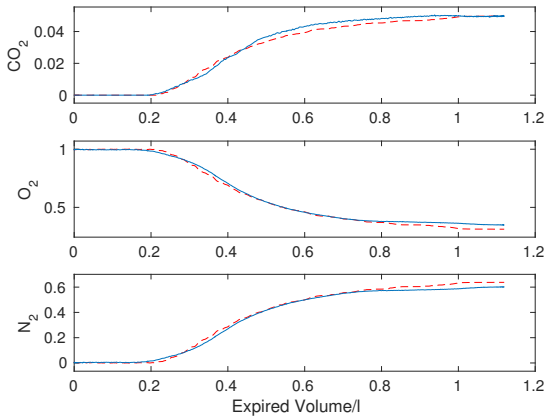

(c) Washout start

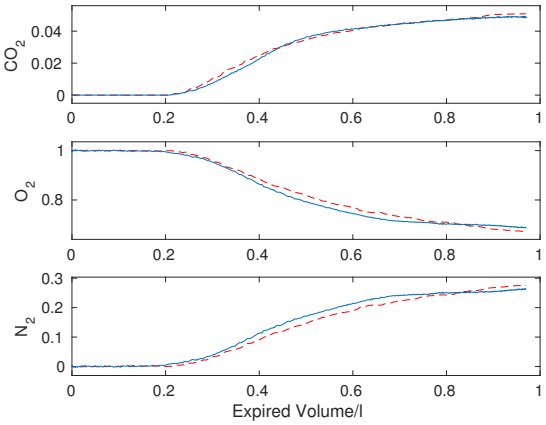

(d) Mid washout

Figure B.10: Example plots of the fit to subject E06.

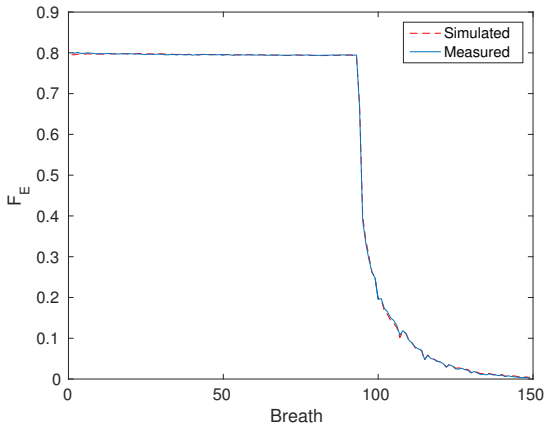

(a) Mixed expired nitrogen fraction

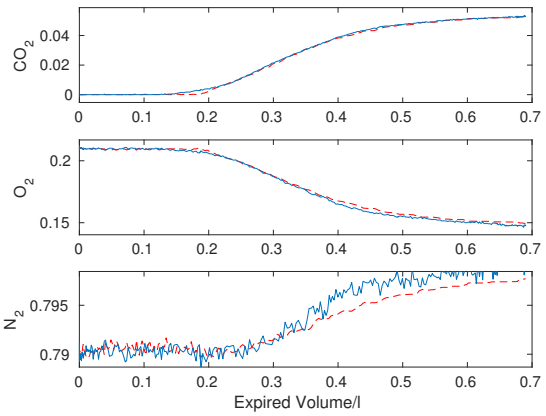

(b) Normal breathing

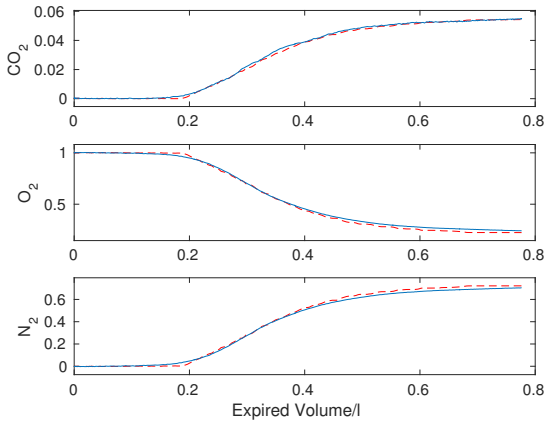

(c) Washout start

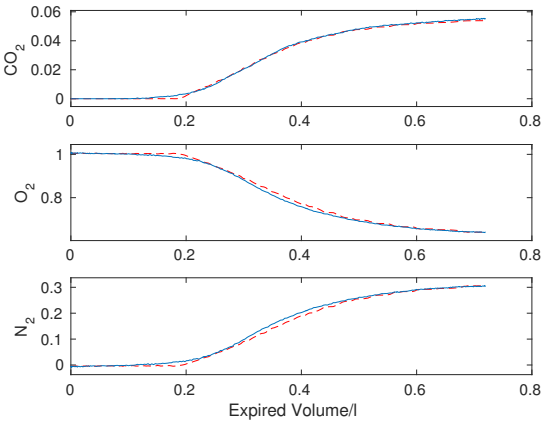

(d) Mid washout

Figure B.11: Example plots of the fit to subject E07.

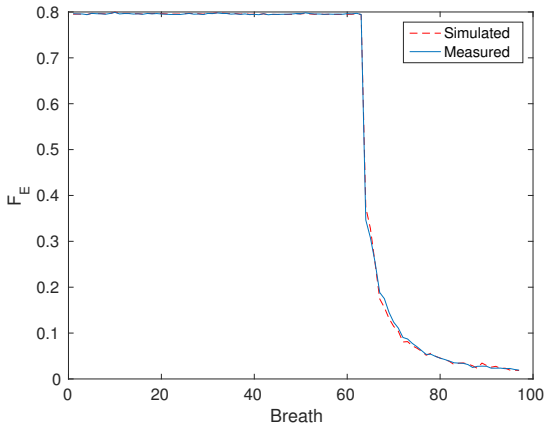

(a) Mixed expired nitrogen fraction

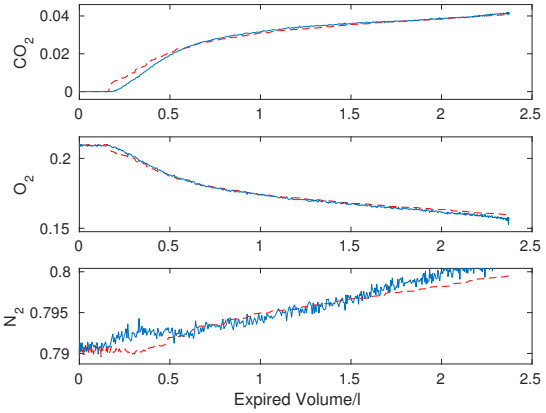

(b) Normal breathing

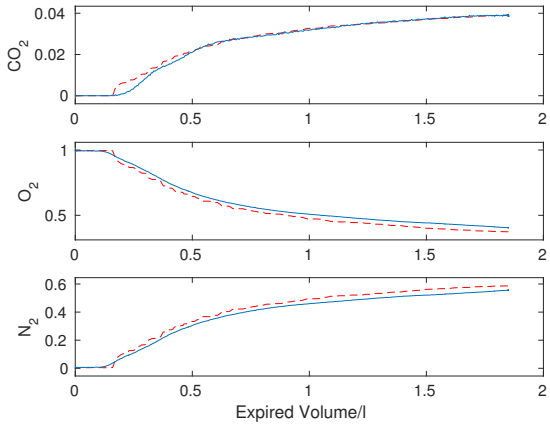

(c) Washout start

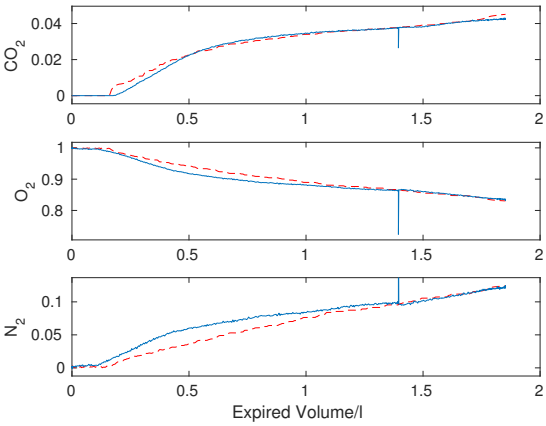

(d) Mid washout

Figure B.12: Example plots of the fit to subject C01 from their first visit.

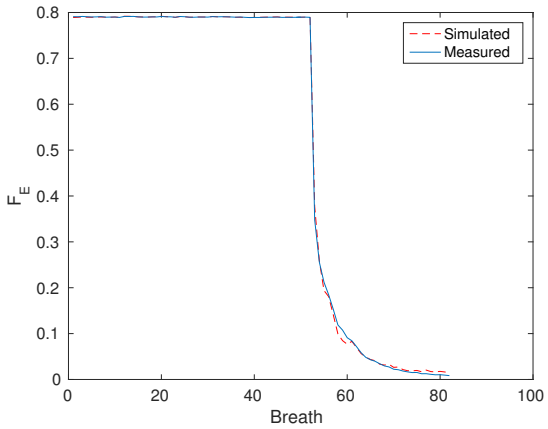

(a) Mixed expired nitrogen fraction

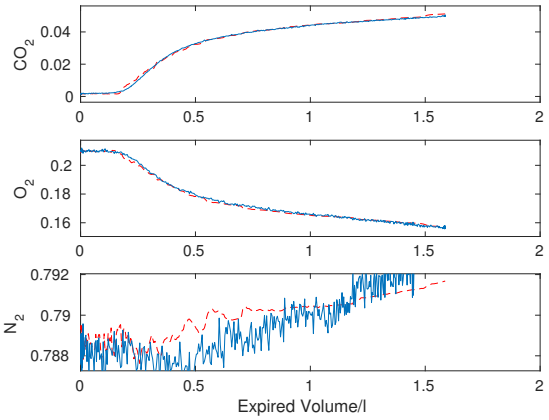

(b) Normal breathing

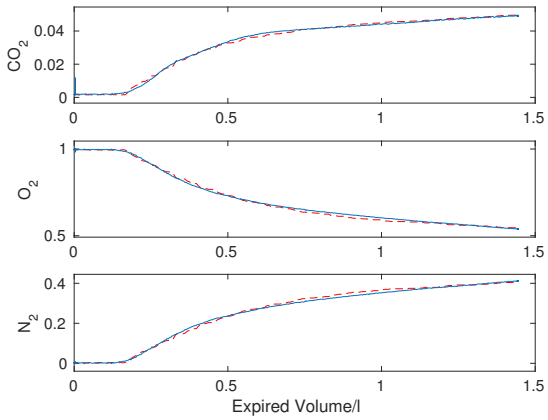

(c) Washout start

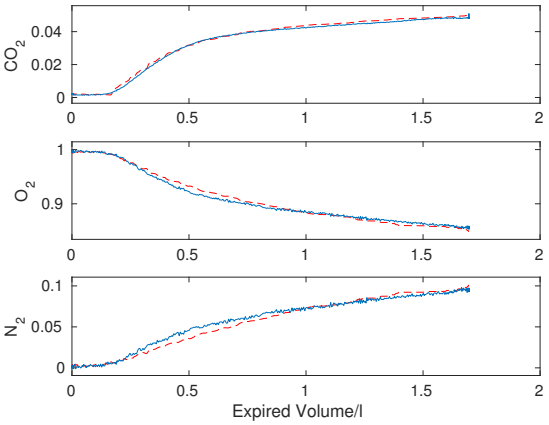

(d) Mid washout

Figure B.13: Example plots of the fit to subject C01 from their second visit.

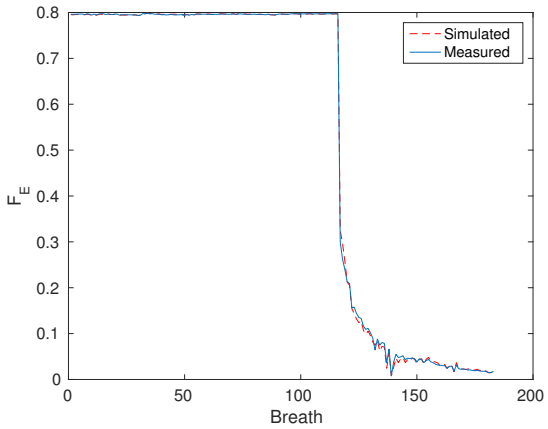

(a) Mixed expired nitrogen fraction

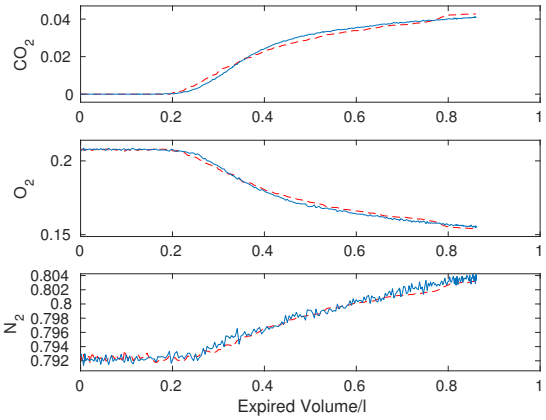

(b) Normal breathing

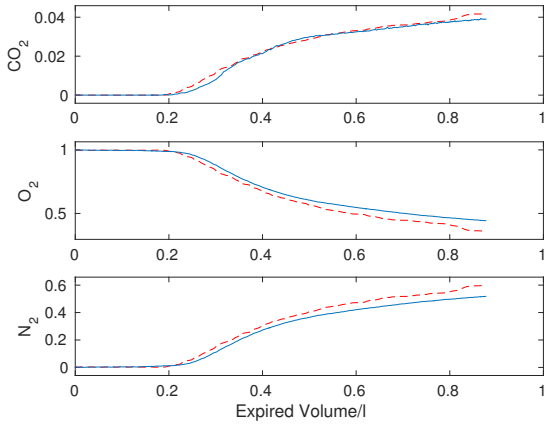

(c) Washout start

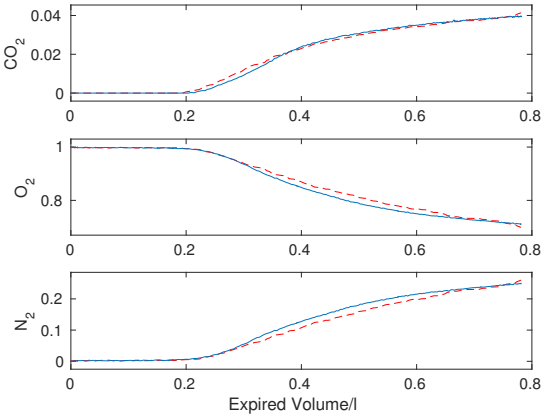

(d) Mid washout

Figure B.14: Example plots of the fit to subject C02.

Bibliography

[1] O. Aguilar, M. Allmaras, W. Bangerth, and L. Tenorio. "Statistics of Parameter Estimates: A Concrete Example". In: *SIAM Review* 57.1 (2015), pp. 131–149.

[2] W. A. Altemeier, H. T. Robertson, and R. W. Glenny. "Pulmonary gas-exchange analysis by using simultaneous deposition of aerosolized and injected microspheres". In: *Journal of Applied Physiology* 85.6 (1998), pp. 2344–2351.

[3] T. Altes, J. Mugler, I. Dregely, S. Ketel, I. Ruset, E. deLange, F. Hersman, and K. Ruppert. "Hyperpolarized ^{129}Xe ventilation MRI: preliminary results in normal subjects and patients with lung disease". In: *Proceedings of the 18th Annual Meeting of ISMRM, Stockholm, Sweden*. 2010, p. 2529.

[4] S. J. Ament, S. Maus, H. Reber, H. G. Buchholz, N. Bausbacher, C. Brochhausen, F. Graf, M. Miederer, and M. Schreckenberger. "PET Lung Ventilation/Perfusion Imaging Using ^{68}Ga Aerosol (Galligas) and ^{68}Ga-Labeled Macroaggregated Albumin". In: *Theranostics, Gallium-68, and Other Radionuclides*. Springer Berlin Heidelberg, 2013, pp. 395–423.

[5] M. Arborelius, H. E. Rosberg, and R. Wiberg. "Multiple breath nitrogen dead space". In: *Clinical Physiology* 8.6 (1988), pp. 561–576.

[6] P. Aurora, A. Bush, P. Gustafsson, C. Oliver, C. Wallis, J. Price, J. Stroobant, S. Carr, and J. Stocks. "Multiple-Breath Washout as a Marker of Lung Disease in Preschool Children with Cystic Fibrosis". In: *American Journal of Respiratory and Critical Care Medicine* 171.3 (2005), pp. 249–256.

[7] A. B. Baker and A. D. Farmery. "Inert Gas Transport in Blood and Tissues". In: *Comprehensive Physiology*. John Wiley & Sons, Inc., 2011.

[8] J. H. Bates, G. K. Prisk, T. E. Tanner, and A. E. McKinnon. "Correcting for the dynamic response of a respiratory mass spectrometer". In: *Journal of Applied Physiology* 55.3 (1983), pp. 1015–1022.

[9] J. H. Bates, G. Schmalisch, D. Filbrun, and J. Stocks. "Tidal breath analysis for infant pulmonary function testing. ERS/ATS Task Force on Standards for Infant Respiratory Function Testing. European Respiratory Society/American Thoracic Society". In: *The European Respiratory Journal* 16.6 (2000), pp. 1180–1192.

[10] K. C. Beck, B. D. Johnson, T. P. Olson, and T. A. Wilson. "Ventilation-perfusion distribution in normal subjects". In: *Journal of Applied Physiology* 113.6 (2012), pp. 872–877.

[11] M. R. Becklake. "A new index of the intrapulmonary mixture of inspired air." In: *Thorax* 7.1 (1952), pp. 111–6.

[12] M. R. Becklake. "Concepts of normality applied to the measurement of lung function." In: *The American Journal of Medicine* 80.6 (1986), pp. 1158–64.

[13] A. Bell, J. Dean, D. Singh, and A. R. Horsley. "LCI as an endpoint in COPD: Comparison with healthy controls, using $SF_6$and nitrogen washout". In: *European Respiratory Journal* 44 (2014).

[14] A. Ben-Tal. "Simplified models for gas exchange in the human lungs". In: *Journal of Theoretical Biology* 238.2 (2006), pp. 474–495.

[15] R. Caflisch. *Random walks with gaussian steps*. URL: `http://www.math.ucla.edu/~caflisch/181.1.03f/Lect4-5.pdf` (visited on 03/03/2016).

[16] J. Carpenter and J. Bithell. "Bootstrap confidence intervals: when, which, what? A practical guide for medical statisticians". In: *Statistics in Medicine* 19.9 (2000), pp. 1141–1164.

[17] L. Ciaffoni, D. P. O'Neill, J. H. Couper, G. A. D. Ritchie, G. Hancock, and P. A. Robbins. "In-airway molecular flow sensing: A new technology for continuous, non invasive monitoring of oxygen consumption in critical care". In: *Submitted* (2016).

[18] A. R. Clark, M. H. Tawhai, E. A. Hoffman, and K. S. Burrowes. "The interdependent contributions of gravitational and structural features to perfusion distribution in a multiscale model of the pulmonary circulation". In: *Journal of Applied Physiology* 110.4 (2011), pp. 943–955.

[19] A. L. Coates, R. Peslin, D. Rodenstein, and J. Stocks. "Measurement of lung volumes by plethysmography." In: *The European Respiratory Journal : Official Journal of the European Society for Clinical Respiratory Physiology* 10.6 (1997), pp. 1415–27.

[20] D. Colquhoun, C. J. Hatton, and A. G. Hawkes. "The quality of maximum likelihood estimates of ion channel rate constants". In: *The Journal of Physiology* 547.3 (2003), pp. 699–728.

[21] R. O. Crapo. "Pulmonary-function testing." In: *The New England Journal of Medicine* 331.1 (1994), pp. 25–30.

[22] R. C. Darling, A. Cournand, D. W. Richards, and B. Domanski. "Studies on the intrapulmonary mixtures of gases. V. Forms of inadequate ventilation in normal and emphysematous lungs, Analyzed by means of breathing pure oxygen". In: *Journal of Clinical Investigation* 23.1 (1944), pp. 55–67.

[23] H. Davy. *Researches, chemical and philosophical, chiefly concerning nitrous oxide*. Butterworth-Heinemann, 1800, pp. 472–473.

[24] C. B. Do. *The multivariate Gaussian distribution*. URL: `http://cs229.stanford.edu/section/gaussians.pdf` (visited on 02/01/2016).

[25] A. B. Dubois. "Alveolar CO_2 and O_2 During Breath Holding, Expiration, and Inspiration". In: *Journal of Applied Physiology* 5.1 (1952), pp. 1–12.

[26] L. A. Engel. "Gas mixing within the acinus of the lung". In: *Journal of Applied Physiology* 54.3 (1983), pp. 609–618.

[27] S. B. Fain, F. R. Korosec, J. H. Holmes, R. O'Halloran, R. L. Sorkness, and T. M. Grist. "Functional lung imaging using hyperpolarized gas MRI." In: *Journal of Magnetic Resonance Imaging : JMRI* 25.5 (2007), pp. 910–23.

[28] A. D. Farmery and C. E. W. Hahn. "Response-time enhancement of a clinical gas analyzer facilitates measurement of breath-by-breath gas exchange". In: *Journal of Applied Physiology* 89.2 (2000), pp. 581–589.

[29] W. O. Fenn and P. Dejours. "Composition of Alveolar Air During Breath Holding With and Without Prior Inhalation of Oxygen and Carbon Dioxide". In: *Journal of Applied Physiology* 7.3 (1954), pp. 313–319.

[30] W. S. Fowler. "Lung function studies; the respiratory dead space." In: *The American Journal of Physiology* 154.3 (1948), pp. 405–16.

[31] W. S. Fowler. "Lung function studies; uneven pulmonary ventilation in normal subjects and in patients with pulmonary disease." In: *Journal of Applied Physiology* 2.6 (1949), pp. 283–99.

[32] W. S. Fowler, E. R. Cornish, and S. S. Kety. "Lung function studies. VIII. Analysis of alveolar ventilation by pulmonary N_2 clearance curves." In: *The Journal of Clinical Investigation* 31.1 (1952), pp. 40–50.

[33] F. Garcia-Rio, A. Dorgham, J. M. Pino, C. Villasante, C. Garcia-Quero, and R. Alvarez-Sala. "Lung Volume Reference Values for Women and Men 65 to 85 Years of Age". In: *American Journal of Respiratory and Critical Care Medicine* 180.11 (2009), pp. 1083–1091.

[34] D. J. Gavaghan and C. E. W. Hahn. "A mathematical evaluation of the alveolar amplitude response technique." In: *Respiration Physiology* 102.1 (1995), pp. 105–20.

[35] D. J. Gavaghan and C. E. W. Hahn. "A tidal breathing model of the forced inspired inert gas sinewave technique." In: *Respiration Physiology* 106.2 (1996), pp. 209–21.

[36] H. W. Glindmeyer, J. J. Lefante, C. McColloster, R. N. Jones, and H. Weill. "Blue-collar normative spirometric values for Caucasian and African-American men and women aged 18 to 65." In: *American Journal of Respiratory and Critical Care Medicine* 151.2 Pt 1 (1995), pp. 412–22.

[37] M. Grönkvist, E. Bergsten, and P. M. Gustafsson. "Effects of body posture and tidal volume on inter- and intraregional ventilation distribution in healthy men". In: *Journal of Applied Physiology* 92.2 (2002), pp. 634–642.

[38] P. M. Gustafsson, P. Aurora, and A. Lindblad. "Evaluation of ventilation maldistribution as an early indicator of lung disease in children with cystic fibrosis". In: *European Respiratory Journal* 22.6 (2003), pp. 972–979.

[39] P. M. Gustafsson, P. D. Robinson, M. Gilljam, A. Lindblad, and B. K. Houltz. "Slow and fast lung compartments in cystic fibrosis measured by nitrogen multiple-breath washout". In: *Journal of Applied Physiology* 117.7 (2014), pp. 720–729.

[40] C. E. W. Hahn. "Oxygen respiratory gas analysis by sine-wave measurement: a theoretical model." In: *Journal of Applied Physiology* 81.2 (1996), pp. 985–97.

[41] C. E. W. Hahn, A. M. Black, S. A. Barton, and I. Scott. "Gas exchange in a three-compartment lung model analyzed by forcing sinusoids of N_2O". In: *Journal of Applied Physiology* 75.4 (1993), pp. 1863–1876.

[42] C. E. W. Hahn and A. Farmery. "Gas exchange modelling: no more gills, please". In: *British Journal of Anaesthesia* 91.1 (2003), pp. 2–15.

[43] J. S. Haldane and J. I. Graham. *Methods of air analysis*. C. Griffin and co., 1935.

[44] R. Hamilton. "Cardiorespiratory measurements using inspired oxygen sinusoids". PhD thesis. Oxford University, 1998.

[45] K. E. Hines. "A Primer on Bayesian Inference for Biophysical Systems". In: *Biophysical Journal* 108.9 (2015), pp. 2103–2113.

[46] K. E. Hines, T. R. Middendorf, and R. W. Aldrich. "Determination of parameter identifiability in nonlinear biophysical models: A Bayesian approach". In: *The Journal of General Physiology* 143.3 (2014), pp. 401–416.

[47] M. P. Hlastala. "Multiple inert gas elimination technique". In: *Journal of Applied Physiology* 56.1 (1984), pp. 1–7.

[48] J. C. Hogg. "Pathophysiology of airflow limitation in chronic obstructive pulmonary disease". In: *Lancet* 364.9435 (2004), pp. 709–721.

[49] F. C. Horn, M. H. Deppe, H. Marshall, J. Parra-Robles, and J. M. Wild. "Quantification of regional fractional ventilation in human subjects by measurement of hyperpolarized 3He washout with 2D and 3D MRI". In: *Journal of Applied Physiology* 116.2 (2014), pp. 129–139.

[50] A. R. Horsley. "Lung clearance index in the assessment of airways disease". In: *Respiratory Medicine* 103.6 (2009), pp. 793–799.

[51] A. R. Horsley, P. M. Gustafsson, K. A. Macleod, C. Saunders, A. P. Greening, D. J. Porteous, J. C. Davies, S. Cunningham, E. W. F. W. Alton, and J. A. Innes. "Lung clearance index is a sensitive, repeatable and practical measure of airways disease in adults with cystic fibrosis". In: *Thorax* 63.2 (2008), pp. 135–140.

[52] A. R. Horsley, K. Macleod, R. Gupta, N. Goddard, and N. Bell. "Enhanced Photoacoustic Gas Analyser Response Time and Impact on Accuracy at Fast Ventilation Rates during Multiple Breath Washout". In: *PLOS ONE* 9.6 (2014), e98487.

[53] J. R. Hurst. "Precision Medicine in Chronic Obstructive Pulmonary Disease". In: *American Journal of Respiratory and Critical Care Medicine* 193.6 (2016), pp. 593–594.

[54] K. Husemann, N. Berg, J. Engel, J. Port, C. Joppek, Z. Tao, F. Singer, H. Schulz, and M. Kohlhäufl. "Double tracer gas single-breath washout: reproducibility in healthy subjects and COPD". In: *European Respiratory Journal* 44.5 (2014), pp. 1210–1222.

[55] K. Ito and P. J. Barnes. "COPD as a disease of accelerated lung aging". In: *Chest* 135.1 (2009), pp. 173–180.

[56] H. U. Kauczor, D. Hofmann, K. F. Kreitner, H. Nilgens, R. Surkau, W. Heil, A. Potthast, M. V. Knopp, E. W. Otten, and M. Thelen. "Normal and abnormal pulmonary ventilation: visualization at hyperpolarized He-3 MR imaging". In: *Radiology* 201.2 (1996), pp. 564–568.

[57] S. S. Kaushik, Z. I. Cleveland, G. P. Cofer, G. Metz, D. Beaver, J. Nouls, M. Kraft, W. Auffermann, J. Wolber, H. P. McAdams, and B. Driehuys. "Diffusion-weighted hyperpolarized ^{129}Xe MRI in healthy volunteers and subjects with chronic obstructive pulmonary disease". In: *Magnetic Resonance in Medicine* 65.4 (2011), pp. 1154–1165.

[58] S. S. Kaushik, M. S. Freeman, Z. I. Cleveland, J. Davies, J. Stiles, R. S. Virgincar, S. H. Robertson, M. He, K. T. Kelly, W. M. Foster, H. P. McAdams, and B. Driehuys. "Probing the regional distribution of pulmonary gas exchange through single-breath gas- and dissolved-phase ^{129}Xe imaging". In: *Journal of Applied Physiology* 115.6 (2013), pp. 850–860.

[59] T. Kawashiro, R. S. Sikand, F. Adaro, H. Takahashi, and J. Piiper. "Study of intrapulmonary gas mixing in man by simultaneous wash-out of helium and sulfur hexafluoride". In: *Respiration Physiology* 28.2 (1976), pp. 261–275.

[60] M. G. Levitzky. *Pulmonary Physiology*. Sixth Edition. McGraw-Hill Medical Publishing Division, 2002.

[61] S. M. Lewis, J. W. Evans, and A. A. Jalowayski. "Continuous distributions of specific ventilation recovered from inert gas washout." In: *Journal of Applied Physiology: Respiratory, Environmental and Exercise Physiology* 44.3 (1978), pp. 416–23.

[62] J. Liepe, S. Filippi, M. Komorowski, and M. P. H. Stumpf. "Maximizing the Information Content of Experiments in Systems Biology". In: *PLoS Computational Biology* 9.1 (2013), e1002888.

[63] K. Lin, B. Watkins, T. Johnson, J. A. Rodriguez, and M. B. Barton. "Screening for Chronic Obstructive Pulmonary Disease Using Spirometry: Summary of the Evidence for the U.S. Preventive Services Task Force". In: *U.S. Preventive Services Task Force Evidence Syntheses, formerly Systematic Evidence Reviews*. Agency for Healthcare Research and Quality (US), 2008.

[64] N. Macintyre, R. O. Crapo, G. Viegi, D. C. Johnson, C. P. M. van der Grinten, V. Brusasco, F. Burgos, R. Casaburi, A. Coates, P. Enright, P. Gustafsson, J. Hankinson, R. Jensen, R. McKay, M. R. Miller, D. Navajas, O. F. Pedersen, R. Pellegrino, and J. Wanger. "Standardisation of the single-breath determination of carbon monoxide uptake in the lung." In: *The European Respiratory Journal : Official Journal of the European Society for Clinical Respiratory Physiology* 26.4 (2005), pp. 720–35.

[65] D. M. Mannino and B. J. Make. "Is it time to move beyond the "O" in early COPD?" In: *European Respiratory Journal* 46.6 (2015), pp. 1535–1537.

[66] Mathworks. *MATLAB Documentation*. URL: `http://uk.mathworks.com/help/matlab/` (visited on 02/01/2016).

[67] J. E. McDonough, R. Yuan, M. Suzuki, N. Seyednejad, W. M. Elliott, P. G. Sanchez, A. C. Wright, W. B. Gefter, L. Litzky, H. O. Coxson, P. D. Pare, D. D. Sin, R. A. Pierce, J. C. Woods, A. M. McWilliams, J. R. Mayo, S. C. Lam, J. D. Cooper, and J. C. Hogg. "Small-Airway Obstruction and Emphysema in Chronic Obstructive Pulmonary Disease". In: *New England Journal of Medicine* 365.17 (2011), pp. 1567–1575.

[68] M. N. Melsom, J. Kramer-Johansen, T. Flatebø, C. Müller, and G. Nicolaysen. "Distribution of pulmonary ventilation and perfusion measured simultaneously in awake goats". In: *Acta Physiologica Scandinavica* 159.3 (1997), pp. 199–208.

[69] G. R. Meneely and N. L. Kaltreider. "The volume of the lung determined by helium dilution. Description of the method and comparison with other procedures." In: *The Journal of Clinical Investigation* 28.1 (1949), pp. 129–39.

[70] M. Meyer, C. Hook, H. Rieke, and J. Piiper. "Gas mixing in dog lungs studied by single-breath washout of He and SF_6". In: *Journal of Applied Physiology* 55.6 (1983), pp. 1795–1802.

[71] B. E. Miller, R. Tal-Singer, S. I. Rennard, A. Furtwaengler, N. Leidy, M. Lowings, U. J. Martin, T. R. Martin, D. D. Merrill, J. Snyder, J. Walsh, and D. M. Mannino. "Plasma Fibrinogen Qualification as a Drug Development Tool in Chronic Obstructive Pulmonary Disease. Perspective of the Chronic Obstructive Pulmonary Disease Biomarker Qualification Consortium". In: *American Journal of Respiratory and Critical Care Medicine* 193.6 (2016), pp. 607–613.

[72] M. R. Miller, J. Hankinson, V. Brusasco, F. Burgos, R. Casaburi, A. Coates, R. Crapo, P. Enright, C. P. M. van der Grinten, P. Gustafsson, R. Jensen, D. C. Johnson, N. MacIntyre, R. McKay, D. Navajas, O. F. Pedersen, R. Pellegrino, G. Viegi, and J. Wanger. "Standardisation of spirometry." In: *The European Respiratory Journal : Official Journal of the European Society for Clinical Respiratory Physiology* 26.2 (2005), pp. 319–38.

[73] J. H. Mitchell, E. A. Hoffman, and M. H. Tawhai. "Relating indices of inert gas washout to localised bronchoconstriction". In: *Respiratory Physiology & Neurobiology* 183.3 (2012), pp. 224–233.

[74] R. R. Mitchell, R. M. Wilson, and D. Sierra. "ICU monitoring of ventilation distribution." In: *International Journal of Clinical Monitoring and Computing* 2.4 (1986), pp. 199–206.

[75] M. Monjezi, R. Dastanpour, M. S. Saidi, and A. R. Pishevar. "Prediction of particle deposition in the respiratory track using 3D–1D modeling". In: *Scientia Iranica* 19.6 (2012), pp. 1479–1486.

[76] K. W. Morton and D. F. Mayers. *Numerical Solution of Partial Differential Equations: An Introduction*. Cambridge University Press, 2005.

[77] J. P. Mugler and T. A. Altes. "Hyperpolarized 129Xe MRI of the human lung". In: *Journal of Magnetic Resonance Imaging* 37.2 (2013), pp. 313–331.

[78] G. Musch, J. D. H. Layfield, R. S. Harris, M. F. V. Melo, T. Winkler, R. J. Callahan, A. J. Fischman, and J. G. Venegas. "Topographical distribution of pulmonary perfusion and ventilation, assessed by PET in supine and prone humans". In: *Journal of Applied Physiology* 93.5 (2002), pp. 1841–1851.

[79] Y. Oba and N. A. Lone. "Efficacy and safety of roflumilast in patients with chronic obstructive pulmonary disease: a systematic review and meta-analysis". In: *Therapeutic Advances in Respiratory Disease* 7.1 (2013), pp. 13–24.

[80] M. Ochs, J. R. Nyengaard, A. Jung, L. Knudsen, M. Voigt, T. Wahlers, J. Richter, and H. J. G. Gundersen. "The Number of Alveoli in the Human Lung". In: *American Journal of Respiratory and Critical Care Medicine* 169.1 (2004), pp. 120–124.

[81] D. P. O'Neill and P. A. Robbins. "A mechanistic physiochemical model of carbon dioxide transport in blood". In: *Submitted* (2016).

[82] M. Paiva. "Gas transport in the human lung." In: *Journal of Applied Physiology* 35.3 (1973), pp. 401–10.

[83] M. Paiva and L. Engel. "Pulmonary interdependence of gas transport". In: *Journal of Applied Physiology Respiratory Environmental and Exercise Physiology* 47.2 (1979), pp. 296–305.

[84] T. J. Pedley, R. C. Schroter, and M. F. Sudlow. "The prediction of pressure drop and variation of resistance within the human bronchial airways". In: *Respiration Physiology* 9.3 (1970), pp. 387–405.

[85] R. Pellegrino, G. Viegi, V. Brusasco, R. O. Crapo, F. Burgos, R. Casaburi, A. Coates, C. P. M. van der Grinten, P. Gustafsson, J. Hankinson, R. Jensen, D. C. Johnson, N. MacIntyre, R. McKay, M. R. Miller, D. Navajas, O. F. Pedersen, and J. Wanger. "Interpretative strategies for lung function tests." In: *The European Respiratory Journal : Official Journal of the European Society for Clinical Respiratory Physiology* 26.5 (2005), pp. 948–68.

[86] J. L. Plewes, A. J. Olszowka, and L. E. Farhi. "Amount and rates of CO2 storage in lung tissue". In: *Respiration Physiology* 28.3 (1976), pp. 359–369.

[87] G. K. Prisk, A. R. Elliott, H. J. Guy, S. Verbanck, M. Paiva, and J. B. West. "Multiple-breath washin of helium and sulfur hexafluoride in sustained microgravity". In: *Journal of Applied Physiology* 84.1 (1998), pp. 244–252.

[88] G. K. Prisk, H. J. Guy, A. R. Elliott, M. Paiva, and J. B. West. "Ventilatory inhomogeneity determined from multiple-breath washouts during sustained microgravity on Spacelab SLS-1." In: *Journal of Applied Physiology* 78.2 (1995), pp. 597–607.

[89] J. D. Quirk, A. L. Sukstanskii, J. C. Woods, B. A. Lutey, M. S. Conradi, D. S. Gierada, R. D. Yusen, M. Castro, and D. A. Yablonskiy. "Experimental evidence of age-related adaptive changes in human acinar airways". In: *Journal of Applied Physiology* 120.2 (2016), pp. 159–165.

[90] H. Rahn. "A concept of mean alveolar air and the ventilation-blood flow relationships during pulmonary gas exchange". In: *The American Journal of Physiology* 158.1 (1949), pp. 21–30.

[91] J.-C. Richard, M. Janier, F. Lavenne, C. Tourvieille, D. Le Bars, N. Costes, G. Gimenez, and C. Guerin. "Quantitative Assessment of Regional Alveolar Ventilation and Gas Volume Using $^{13}N_2$ Washout and PET". In: *Journal of Nuclear Medicine* 46.8 (2005), pp. 1375–1383.

[92] K. F. Riley, M. P. Hobson, and S. J. Bence. *Mathematical Methods for Physics and Engineering (3rd edition): A Comprehensive Guide*. Cambridge University Press, 2006.

[93] R. L. Riley and A. Cournand. "Ideal alveolar air and the analysis of ventilation-perfusion relationships in the lungs". In: *Journal of Applied Physiology* 1.12 (1949), pp. 825–847.

[94] R. R. Rizi, J. E. Baumgardner, M. Ishii, Z. Z. Spector, J. M. Edvinsson, A. Jalali, J. Yu, M. Itkin, D. A. Lipson, and W. Gefter. "Determination of regional VA/Q by hyperpolarized 3He MRI". In: *Magnetic Resonance in Medicine* 52.1 (2004), pp. 65–72.

[95] P. D. Robinson, M. D. Goldman, and P. M. Gustafsson. "Inert gas washout: theoretical background and clinical utility in respiratory disease". In: *Respiration; international review of thoracic diseases* 78.3 (2009), pp. 339–355.

[96] P. D. Robinson, P. Latzin, S. Verbanck, G. L. Hall, A. R. Horsley, M. Gappa, C. Thamrin, H. G. M. Arets, P. Aurora, S. I. Fuchs, G. G. King, S. Lum, K. Macleod, M. Paiva, J. J. Pillow, S. Ranganathan, F. Ratjen, F. Singer, S. Sonnappa, J. Stocks, P. Subbarao, B. R. Thompson, and P. M. Gustafsson. "Consensus statement for inert gas washout measurement using multiple- and single- breath tests". In: *European Respiratory Journal* 41.3 (2013), pp. 507–522.

[97] R. Rodríguez-Roisin, M. Drakulovic, D. A. Rodríguez, J. Roca, J. A. Barberà, and P. D. Wagner. "Ventilation-perfusion imbalance and chronic obstructive pulmonary disease staging severity". In: *Journal of Applied Physiology* 106.6 (2009), pp. 1902–1908.

[98] A. Ron. *Introduction to numerical analysis*. URL: `http://pages.cs.wisc.edu/~amos/412/lecture-notes/lecture02.pdf` (visited on 02/23/2016).

[99] B. T. Saam, D. A. Yablonskiy, V. D. Kodibagkar, J. C. Leawoods, D. S. Gierada, J. D. Cooper, S. S. Lefrak, and M. S. Conradi. "MR imaging of diffusion of 3He gas in healthy and diseased lungs". In: *Magnetic Resonance in Medicine* 44.2 (2000), pp. 174–179.

[100] M. A. Sackner, K. A. Feisal, and A. B. DuBois. "Determination of tissue volume and carbon dioxide dissociation slope of the lungs in man". In: *Journal of Applied Physiology* 19.3 (1964), pp. 374–380.

[101] P. Scheid, M. P. Hlastala, and J. Piiper. "Inert gas elimination from lungs with stratified inhomogeneity: Theory". In: *Respiration Physiology* 44.3 (1981), pp. 299–309.

[102] P. W. Scherer, S. Gobran, S. J. Aukburg, J. E. Baumgardner, R. Bartkowski, and G. R. Neufeld. "Numerical and experimental study of steady-state CO_2 and inert gas washout". In: *Journal of Applied Physiology* 64.3 (1988), pp. 1022–1029.

[103] P. W. Scherer, L. H. Shendalman, and N. M. Greene. "Simultaneous diffusion and convection in single breath lung washout". In: *The Bulletin of Mathematical Biophysics* 34.3 (1972), pp. 393–412.

[104] N. Siafakas, P. Vermeire, N. Pride, P. Paoletti, J. Gibson, P. Howard, J. Yernault, M. Decramer, T. Higenbottam, and D. Postma. "Optimal assessment and management of chronic obstructive pulmonary disease (COPD). The European Respiratory Society Task Force". In: *European Respiratory Journal* 8.8 (1995), pp. 1398–1420.

[105] J. Stocks and P. H. Quanjer. "Reference values for residual volume, functional residual capacity and total lung capacity. ATS Workshop on Lung Volume Measurements. Official Statement of The European Respiratory Society". In: *The European Respiratory Journal* 8.3 (1995), pp. 492–506.

[106] R. Street, G. Watters, and J. Vennard. *Elementary fluid mechanics*. Wiley, 1996.

[107] B. Suki and H. Parameswaran. "Computational modeling helps uncover mechanisms related to the progression of emphysema". In: *Drug Discovery Today: Disease Models* (2014).

[108] A. J. Swan, A. R. Clark, and M. H. Tawhai. "A computational model of the topographic distribution of ventilation in healthy human lungs". In: *Journal of Theoretical Biology* 300 (2012), pp. 222–231.

[109] A. Talag and P. Wilcox. "Clinical physiology of chronic obstructive pulmonary disease". In: *British Columbia Medical Journal* 50.2 (2008), p. 97.

[110] M. H. Tawhai, A. J. Pullan, and P. J. Hunter. "Generation of an Anatomically Based Three-Dimensional Model of the Conducting Airways". In: *Annals of Biomedical Engineering* 28.7 (2000), pp. 793–802.

[111] M. H. Tawhai and K. S. Burrowes. "Developing integrative computational models of pulmonary structure". In: *The Anatomical Record Part B: The New Anatomist* 275B.1 (2003), pp. 207–218.

[112] M. H. Tawhai and P. J. Hunter. "Modeling water vapor and heat transfer in the normal and the intubated airways." In: *Annals of Biomedical Engineering* 32.4 (2004), pp. 609–22.

[113] M. H. Tawhai and P. J. Hunter. "Multibreath washout analysis: modelling the influence of conducting airway asymmetry". In: *Respiration Physiology* 127.2–3 (2001), pp. 249–258.

[114] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M. P. H. Stumpf. "Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems". In: *Journal of The Royal Society Interface* 6.31 (2009), pp. 187–202.

[115] S. Treppo, S. M. Mijailovich, and J. G. Venegas. "Contributions of pulmonary perfusion and ventilation to heterogeneity in V/Q measured by PET". In: *Journal of Applied Physiology* 82.4 (1997), pp. 1163–1176.

[116] S. Verbanck and M. Paiva. "Model simulations of gas mixing and ventilation distribution in the human lung". In: *Journal of Applied Physiology* 69.6 (1990), pp. 2269–2279.

[117] S. Verbanck, D. Schuermans, A. Van Muylem, M. Paiva, M. Noppen, and W. Vincken. "Ventilation distribution during histamine provocation". In: *Journal of Applied Physiology* 83.6 (1997), pp. 1907–1916.

[118] S. Verbanck, D. Shuermans, A. Van Muylem, C. Melot, M. Noppen, W. Vincken, and M. Paiva. "Conductive and acinar lung zone contributions to ventilation inhomogeneity in COPD". In: *American Journal of Respiratory and Critical Care Medicine* 157.5 (1998), pp. 1573–1577.

[119] S. Verbanck and M. Paiva. "Gas Mixing in the Airways and Airspaces". In: *Comprehensive Physiology*. John Wiley & Sons, Inc., 2011.

[120] S. Verbanck, M. Paiva, D. Schuermans, S. Hanon, W. Vincken, and A. Van Muylem. "Relationships between the lung clearance index and conductive and acinar ventilation heterogeneity." In: *Journal of Applied Physiology* 112.5 (2012), pp. 782–90.

[121] E. K. Verbeken, M. Cauberghs, I. Mertens, J. Clement, J. M. Lauweryns, and K. P. Van de Woestijne. "The senile lung. Comparison with normal and emphysematous lungs. 2. Functional aspects". In: *Chest* 101.3 (1992), pp. 800–809.

[122] J. Vestbo, S. S. Hurd, A. G. Agustí, P. W. Jones, C. Vogelmeier, A. Anzueto, P. J. Barnes, L. M. Fabbri, F. J. Martinez, M. Nishimura, R. A. Stockley, D. D. Sin, and R. Rodriguez-Roisin. "Global Strategy for the Diagnosis, Management, and Prevention of Chronic Obstructive Pulmonary Disease". In: *American Journal of Respiratory and Critical Care Medicine* 187.4 (2013), pp. 347–365.

[123] P. D. Wagner, D. R. Dantzker, R. Dueck, J. L. Clausen, and J. B. West. "Ventilation-perfusion inequality in chronic obstructive pulmonary disease." In: *Journal of Clinical Investigation* 59.2 (1977), pp. 203–216.

[124] P. D. Wagner, H. A. Saltzman, and J. B. West. "Measurement of continuous distributions of ventilation-perfusion ratios: theory". In: *Journal of Applied Physiology* 36.5 (1974), pp. 588–599.

[125] P. D. Wagner. "The multiple inert gas elimination technique (MIGET)". In: *Intensive Care Medicine* 34.6 (2008), pp. 994–1001.

[126] J. Wanger, J. L. Clausen, A. Coates, O. F. Pedersen, V. Brusasco, F. Burgos, R. Casaburi, R. Crapo, P. Enright, C. P. M. van der Grinten, P. Gustafsson, J. Hankinson, R. Jensen, D. Johnson, N. Macintyre, R. McKay, M. R. Miller, D. Navajas, R. Pellegrino, and G. Viegi. "Standardisation of the measurement of lung volumes." In: *The European Respiratory Journal : Official Journal of the European Society for Clinical Respiratory Physiology* 26.3 (2005), pp. 511–22.

[127] E. R. Weibel. *Morphometry of the human lung*. Springer, 1963.

[128] J. B. West. *Ventilation/blood flow and gas exchange*. Fifth edition. Blackwell Scientific Publications Ltd, 1990.

[129] J. B. West. *Respiratory Physiology: The Essentials*. Lippincott Williams & Wilkins, 2012.

[130] F. M. White. *Fluid Mechanics with Student DVD*. Seventh edition. McGraw-Hill Higher Education, 2010.

[131] J. P. Whiteley and D. J. Gavaghan. "Efficient computations of gas transport in the respiratory tree". In: *Mathematical Medicine and Biology* 20.1 (2003), pp. 91–104.

[132] J. P. Whiteley, D. J. Gavaghan, and C. E. W. Hahn. "A mathematical evaluation of the multiple breath nitrogen washout (MBNW) technique and the multiple inert gas elimination technique (MIGET)." In: *Journal of Theoretical Biology* 194.4 (1998), pp. 517–39.

[133] J. P. Whiteley, D. J. Gavaghan, and C. E. W. Hahn. "A tidal breathing model for the multiple inert gas elimination technique." In: *Journal of Applied Physiology* 87.1 (1999), pp. 161–9.

[134] J. P. Whiteley, D. J. Gavaghan, and C. E. W. Hahn. "A tidal breathing model of the inert gas sinewave technique for inhomogeneous lungs". In: *Respiration Physiology* 124.1 (2000), pp. 65–83.

[135] J. P. Whiteley, M. J. Turner, A. B. Baker, D. J. Gavaghan, and C. E. W. Hahn. "The effects of ventilation pattern on carbon dioxide transfer in three computer models of the airways." In: *Respiratory Physiology & Neurobiology* 131.3 (2002), pp. 269–84.

[136] J. P. Whiteley, D. J. Gavaghan, and C. E. W. Hahn. "Mathematical modelling of pulmonary gas transport". In: *Journal of Mathematical Biology* 47.1 (2003), pp. 79–99.

[137] E. M. Williams, R. M. Hamilton, L. Sutton, J. P. Viale, and C. E. W. Hahn. "Alveolar and dead space volume measured by oscillations of inspired oxygen in awake adults." In: *American Journal of Respiratory and Critical Care Medicine* 156.6 (1997), pp. 1834–9.

[138] E. M. Williams, M. C. Sainsbury, L. Sutton, L. Xiong, A. M. S. Black, J. P. Whiteley, D. J. Gavaghan, and C. E. W. Hahn. "Pulmonary blood flow measured by inspiratory inert gas concentration forcing oscillations". In: *Respiration Physiology* 113.1 (1998), pp. 47–56.

[139] T. A. Wilson and K. C. Beck. "Contributions of ventilation and perfusion inhomogeneities to the VA/Q distribution". In: *Journal of Applied Physiology* 72.6 (1992), pp. 2298–2304.

[140] S. Woods, D. Charney, W. Goodman, and G. Heninger. "Carbon Dioxide—Induced Anxiety: Behavioral, Physiologic, and Biochemical Effects of Carbon Dioxide in Patients With Panic Disorders and Healthy Subjects". In: *Archives of General Psychiatry* 45.1 (1988), pp. 43–52.

[141] J. S. Yem, M. J. Turner, A. B. Baker, I. H. Young, and A. B. H. Crawford. "A tidally breathing model of ventilation, perfusion and volume in normal and diseased lungs". In: *British Journal of Anaesthesia* 97.5 (2006), pp. 718–731.

[142] J. Zinserling, H. Wrigge, D. Varelmann, R. Hering, and C. Putensen. "Measurement of functional residual capacity by nitrogen washout during partial ventilatory support". In: *Intensive Care Medicine* 29.5 (2003), pp. 720–726.

[143] A. Zwart, R. C. Seagrave, and A. Van Dieren. "Ventilation-perfusion ratio obtained by a noninvasive frequency response technique." In: *Journal of Applied Physiology* 41.3 (1976), pp. 419–24.