In-silico investigation of the neonatal brain physiology using a systems biology approach: modelling birth asphyxia and neuroprotective strategies.

Joshua Russell-Buckland

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Medical Physics and Biomedical Engineering

University College London

December 16, 2020

I, Joshua Russell-Buckland, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

Hypoxic ischaemic encephelopathy (HIE), often resulting from intrapartum hypoxic-ischemic injury, is a significant cause of death and morbidity before, during and after birth. In order to identify and monitor HIE, clinicians use non-invasive techniques including magnetic resonance spectroscopy (MRS) and near-infrared spectroscopy (NIRS). However, interpretation of these signals, particularly to determine the effectiveness of treatment and the severity of injury, is a challenging and difficult task.

This thesis describes an attempt to use a systems biology approach to better understand the mechanisms behind HIE and its outcomes, using mathematical and computational techniques to analyse multimodal data, including broadband near-infrared spectroscopy (bNIRS). These models incorporate submodels of cerebral blood flow, oxygen transport and metabolism into a single cohesive model that attempts to simulate the observed measurements of tissue oxygenation and metabolism. The scope of this work is to both develop a set of computational tools that can be used to better understand existing systems biology models of the brain and to develop a new model which is able to incorporate the effects of therapeutic hypothermia, a common treatment for HIE, on the underlying physiology and its dynamics.

The work begins by redeveloping the existing framework used for running and analysing systems biology models as used previously, before going on to develop a Bayesian framework which allows a better and more comprehensive interpretation of the results. This framework is then used to analyse three new models that incorporate the impact of therapeutic hypothermia on the piglet brain. The model

determined to be most effective is then applied to clinical data from neonates that experience spontaneous desaturations in blood oxygen whilst undergoing hypothermic treatment. In all cases data from subjects with both mild and severe injuries are compared to determine if separate parameter spaces (and therefore physiological mechanisms) can be identified for each.

IMPACT Statement

The work presented in this thesis has the potential to have a broad impact both within and outside academia. Within academia and the study of neonatal physiology, it lays the foundations for new research into methods of analysis that can be used to understand multimodal data collected from clinical machines. In particular, it identifies ways to incorporate treatment via external methods into models. In addition, it sets a strong basis for how to properly distinguish between competing models, allowing for a better and more rigorous comparison of contrasting approaches.

In this case, it has been applied to specific models developed at UCL, but the software and tool developed is agnostic with regards to the models it can analyse The software could be used to analyse any systems biology model where a Bayesian approach could be relevant, for example in drug development. This also highlights a potential avenue for the research to be applicable outside of academia; a great number of industry applications require robust and comprehensive analysis of competing system models, and being able to easily analyse these through a ready made tool, that includes a user friendly interface, increases the likelihood of uptake and use.

Finally, whilst the focus of this work is on biological systems, the approach taken can be applied to any model of a system. These could be models of resource distribution within economics and finance, or within engineering looking at large systems that have multiple components that interact. As long as the model can be written in the BCMD framework language, the new approach can be used to analyse it from a Bayesian perspective and to compare competing variations.

Acknowledgements

Firstly I would like to thank Dr. Ilias Tachtsidis for his continuous support throughout my Ph.D study and the related research. His guidance has helped me to achieve everything I wanted to. I would also like to thank Dr. Chris Barnes for his immense help in understanding Bayesian analysis and how it can be applied to systems biology models.

I would also like to thank the entire multimodal spectroscopy group for helping keep me sane. Whether it was teaching me Italian, playing DnD or going on many long coffee walks, their patience, friendship and support was key to getting through the last four years.

Finally, I would like to thank my beautiful partner, Mia, for her constant love, support, friendship and partnership. You've been with me from the very start of this and without you I would not be who or where I am now.

Publications, Talks and Presentations

Publications

- Joshua Russell-Buckland, Matthew Caldwell, and Ilias Tachtsidis. We-BCMD: A cross-platform interface for the BCMD modelling framework. *Wellcome Open Research*, 2:56, 7 2017b. ISSN 2398-502X. doi:10.12688/wellcomeopenres.12201.1. URL `https://wellcomeopenresearch.org/articles/2-56/v1` (as per **Chapter 6**)

- J. Russell-Buckland, G. Bale, I. de Roever, and I. Tachtsidis. *ABroAD: A machine learning based approach to detect broadband NIRS artefacts*, volume 1072. 2018a. doi:10.1007/978-3-319-91287-5_51 (as per **Chapter 3**)

- Joshua Russell-Buckland, Christopher P. Barnes, and Ilias Tachtsidis. A Bayesian framework for the analysis of systems biology models of the brain. *PLOS Computational Biology*, 15(4):e1006631, 4 2019b. ISSN 1553-7358. doi:10.1371/journal.pcbi.1006631. URL `http://dx.plos.org/10.1371/journal.pcbi.1006631` (as per **Chapter 7**)

- Joshua Russell-Buckland and Ilias Tachtsidis. Developing a model to simulate the effect of hypothermia on cerebral blood flow and metabolism. In *Oxygen Transport to Tissue XLI*, pages 299–306. Springer, 2020 (as per **Chapter 8**)

Conferences

- Joshua Russell-Buckland, Gemma Bale, I. DeRoever, and Ilias Tachtsidis. Abroad: A machine learning based approach to detect broadband nirs artefacts. International Society of Oxygen Transport to Tissue, MLU, Halle, Germany, 2017a

- Joshua Russell-Buckland, Cornelius Bauer, and Ilias Tachtsidis. Investigating the relationship between hypothermia and cerebral metabolism following hypoxic-ischaemic injury. International Society of Oxygen Transport to Tissue, Seoul University, Seoul, South Korea, 2018b

Presentations

- Joshua Russell-Buckland, Matthew Caldwell, and Ilias Tachtsidis. Webcmd: A modern, cross-platform interface to the bcmd modelling framework. International Society of Oxygen Transport to Tissue, MLU, Halle, Germany, 2017c

- Joshua Russell-Buckland, Christopher P. Barnes, and Ilias Tachtsidis. A bayesian framework for the analysis of systems biology models of the brain. eResearch Symposium, UCL, London, United Kingdom, 2019c

- Joshua Russell-Buckland, Matthew Caldwell, and Ilias Tachtsidis. Webcmd:

A modern, cross-platform interface to the bcmd modelling framework. CARnet, Oxford University, Oxford, United Kingdom, 2018c

- J. Russell-Buckland, P. Kaynezhad, S. Mitra, G. Bale, C. Bauer, I. Lingam, C. Meehan, A. Avdic-Belltheus, K. Martinello, A. Brainbridge, N.J. Robertson, and I. Tachtsidis. Systems biology model of cerebral oxygen delivery and metabolism during therapeutic hypothermia:application to the piglet model. International Society of Oxygen Transport to Tissue, Albuquerque, New Mexico, USA, 2019a

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Systems biology models are used to understand complex biological and physiological systems comprised of large numbers of individual elements that give rise to emergent behaviours. These complex systems are dependent on both the properties of the whole network and on the individual elements [Kitano, 2002]. UCL has a history of taking this approach to the modelling of cerebral haemodynamics and metabolism, making it possible to better understand physiology and to reproduce measured data. In doing so, it is possible to identify potential mechanisms for the components that act together to produce the observed, often pathological, system behaviour.

Hypoxic-ischaemic encephelopathy (HIE) is a significant cause of death around the time of birth. Somewhere between 421,000 and 1.05 million deaths amongst neonates can be attributed to birth asphyxia and 45-55% of cases given the primary treatment of therapeutic hypothermia still lead to death or severe neurodevelopmental disability [Lee et al., 2008]. The exact mechanisms behind this difference in outcome are not fully understood, but it has been shown that severe birth asphyxia is accompanied by a 'secondary' cerebral energy failure and impairment of cerebral metabolism. This can occur hours after the initial impairment of brain function during injury [Thoresen et al., 1995]. Treatment given during this period, typically via therapeutic hypothermia, is crucial to improving outcome and reducing the severity of damage done. Understanding how this treatment affects the physiology of the brain is key to proper application and gauging of its efficacy.

In order to identify and monitor HIE, clinicians use non-invasive techniques including magnetic resonance spectroscopy (MRS) and near-infrared spectroscopy (NIRS). However, interpretation of these signals, particularly to determine the effectiveness of treatment and the severity of injury, is a challenging and difficult task. In the Neonatal ICU at UCLH neonates are monitored extensively, and data are collected using a unique broadband near-infrared spectroscopy (bNIRS) instrument that offers the potential to identify early-stage biomarkers of injury, as well as continuous real time data collection of the brain's activity and state, with quantification of oxygenation and metabolism from day 1, providing unique prognostic markers [Bale et al., 2018]. Correct interpretation of these data is key to unlocking its full potential as a clinical aid. Without this there is a great chance of 'information overload', removing the gains obtained by collecting this data and making diagnosis and prognosis more difficult.

This thesis describes an attempt to use a systems biology approach to better understand the mechanisms behind HIE and its outcomes, using mathematical and computational techniques to analyse multimodal data, including bNIRS. These models incorporate submodels of cerebral blood flow, oxygen transport and metabolism into a single cohesive model that attempts to simulate the observed measurements of tissue oxygenation and metabolism. It is the ability of bNIRS to measure both haemodynamics and metabolic activity that makes analysis via systems biology such an appealing approach. By combining the three submodels with the measured data, the interplay between the systemic measurements, the haemodynamics and the metabolic processes, there is the potential to better understand the impact these separate components can have on each other.

Aims and Objectives

The general scope of this PhD is to develop a tool, a computational framework that will allow us to extend and enhanced the previous UCL models by developing a new web based, applying novel statistical frameworks for parameter and variable estimation and finally allowing their application with the existing data collected at

UCL in the Neonatal ICU. In particular the objectives are:

- Redevelop the existing framework used for running and analysing systems biology models.

- Extend this framework to utilise a Bayesian approach to model analysis, providing a more comprehensive interpretation of results than the previous maximum likelihood approach.

- Develop a BrainPiglet-with-hypothermia model as an extension of the existing BrainPiglet model [Moroz et al., 2012a] to incorporate the effects of therapeutic hypothermia in order to better understand the impact of treatment on brain physiology.

- Apply this BrainPiglet-with-hypothermia model to data collected from neonates that have experienced spontaneous desaturation events during hypothermic treatment with differing outcomes.

Thesis Summary

The initial 4 chapters provide background to the work that follows. Chapter 2 provides an overview of cerebral physiology, looking at anatomy, haemodynamics and cellular metabolism before considering HIE, its treatment and some of the differences in the neonatal, piglet and adult brain. Chapter 3 looks at the different clinical techniques that are used to measure physiological data in the brain, including electroencephalography (EEG), magnetic resonance imaging (MRI), MRS and NIRS, before a brief consideration of the ways in which this data must be handled before use. Mathematical modelling is considered in Chapter 4, with an initial look at how mathematical models can be used in biology before a detailed consideration of systems biology models used to understand the brain and its dynamics. We then move on to Chapter 5 in which we look at the different ways these models can be analysed and used to understand the systems they simulate. This starts by looking at the methods that have been used previously to analyse the UCL family of models

before moving on to an overview of how a Bayesian approach can be applied to systems biology models, with a particular focus on approximate Bayesian computation (ABC).

Chapter 6 looks at the original Brain/Circulation Model Developer (BCMD) framework and how it functioned before moving on to the work done to redevelop the system. This leads into Chapter 7 which outlines the work done to rework the previous BCMD framework into something that can be used to perform Bayesian analysis. This is adapted from a previously published paper [Russell-Buckland et al., 2019b] and uses the new framework to analyse both simulated and measured data and compare to the previous. Chapter 8 defines the work done to extend the existing BrainPiglet model to incorporate temperature. Three separate model variants are developed and then compared to each other and the original, temperature-free model before selecting the best variant. This model variant is then applied to neonatal data in Chapter 9. The neonatal data are taken from two neonates with differing injury severities that both experienced spontaneous desaturation events. This data was originally analysed in Bale et al. [2018] and the model is here used to elaborate on this analysis and consider the hypotheses of that work. Finally, Chapter 10 concludes on this work by summarising its findings and outlining future work and directions.

1.1 Personal Statement

There are a number of research collaborations that have made this work possible with this section aiming to make these clearer.

The Brain/Circulation Model Developer (BCMD) framework used to compile and implement the models was developed by Matthew Caldwell and was itself based on the BRAINCIRC software developed by Banaji et al. [2005, 2008] The BCMD software is re-developed in Chapter 6 and then extended upon in Chapter 7 to include the ability to perform Bayesian Analysis.

The BrainPiglet model was originally developed by Moroz et al. [2012a] and is then extended in this thesis to incorporate the effects of external temperature changes, mainly via therapeutic hypothermia. This model is then applied to data collected from piglets that have undergone brain injury followed by hypothermic treatment. This data was collected and initially analysed by Kaynezhad et al. [2019].

The model was then applied to neonatal data collected by Bale et al. [2018] at the Neonatal Intensive Care Unit, University College London Hospital with written, informed consent obtained from parents before each study. All neonates were identified as suffering from hypoxic-ischaemic encephalopathy.

Chapter 2

Brain Physiology

This work as a whole involves a variety of diverse but related subjects relating to systems biology, including distribution of models, model analysis and application of models, but all work has the end goal of better understanding the human brain, particularly that of the human neonate. The brain is arguably the most complex organ in the human body and the human brain in particular may be one of the most complex biological systems, bar that of the body itself.

The brain controls the majority of the functions of the body - from interpreting sensory information to controlling movement, from memory to hormones - and as a result any dysfunction of the brain has a serious and cascading effect on the body as a whole. For example, a stroke occurs when blood supply to the brain is cut off. This can lead to a wide range of effects including slurring of speech, weakness and paralysis of one side of the body, dysphagia and memory problems. The impact of damaging this single organ has a cascade effect that leads to problems with and damage to many parts of the body.

It is important that this complexity is considered when developing a mathematical model of the brain. Any attempt to incorporate all behaviours and effects of the brain is likely to fail due to the necessary complexity of such a model. The resulting model would likely be intractable and were some form of solution possible, interpretations of the output may prove to be non-useful. Instead we must focus the model down to a limited scope that is both tractable and interpretable. George Box is oft-quoted as saying “All models are wrong, but some are useful." [Box, 1979]

Whilst I will go into this in more detail in Chapter 4, it is important to note here. The main goal of this work and the work that has preceded it is to produce a *useful* model, rather than a complete model that can perfectly reproduce our data without providing usable or useful insights.

With the above in mind, this chapter explores the functions and physiology of the brain that relates directly to the BrainSignals models - cellular metabolism and cerebral oxygenation and haemodynamics - before looking at HIE, the injury this work focuses on understanding. We then consider the differences between the adult, neonate and piglet brain, which is important when developing models for each sub-group.

2.1 Physiology

2.1.1 Anatomy

The brain along with the spinal cord make up the central nervous system. It consists of three main components: the cerebrum, the brain stem and the cerebellum. These themselves also have internal structure and organisation.

Cerebellum The cerebellum plays an important part in controlling motor function within humans, and may also be involved in some non-motor functions including personality and mood [Wolf et al., 2014], and damage to the cerebellum can produce dysfunction in fine motor control, motor learning and equilibrium [Fine et al., 2002]. It is connected to the rest of the brain via the brain stem.

Brain stem The brain stem consists of the mid-brain (mesencephalon), pons (metencephalon), and medulla oblongata (myelencephalon) [Sharma and Majsak, 2014]. It has three main functions:

a: Conductance: all information relayed from the body to the cerebrum and cerebellum (and vice versa) is done via the brain stem.

b: Cranial nerves: ten of the twelve pairs of cranial nerves emerge from the brain stem.

c: Critical function: the brain stem is involved in cardiovascular and respiratory control, pain sensitivity and consciousness.

Cerebrum The cerebrum is a large part of the brain containing the cerebral cortex as well as a number of sub-cortical structures including the hippocampus, basal ganglia and olfactory bulb. It has a longitudinal fissure that allows the overall structure to be divided into left and right hemispheres and, whilst most functionality is shared between the two halves, perceptual information is laterally partitioned such that information from the right side of the body is processed by the left side of the brain and vice-versa. Additionally, the cerebral cortex is generally classified into four lobes: frontal, occipital, parietal and temporal. The cerebrum has many functions across its various components, controlling all voluntary actions as well as sensory processing, language and communication, conscious movement and learning and memory.

The central nervous system is surrounded by a series of membranes known as the meninges which are, in order of furthest from the skull to closest, a) the dura mater, b) the arachnoid mater, and c) the pia mater. The cerebrospinal fluid (CSF) is located between the latter two in the subarachnoid space along with the major arteries.

2.1.2 Blood Flow

Four major arteries supply the brain - the right and left internal carotid arteries and the right and left vertebral arteries. The vertebral arteries join to form the basilar artery and this then joins with the internal carotid arteries at the base of the brain to form the *circle of Willis*. From this there are six major vessels that then supply blood to the brain: the left and right middle cerebral arteries (MCAs), the left and right anterior cerebral arteries (ACAs), and the left and right posterior cerebral arteries (PCAs). Branching off from these are many other smaller supply vessels which then supply different regions of the brain. This is illustrated in Figure 2.1. An important feature of the circle of Willis is that perfusion can be maintained across the brain even if some of the vessels are blocked or absent. This redundancy makes the system as a whole much more robust [Payne, 2017, Chap. 1].

At the opposite end of the length scale to the major arteries is another important component of the cerebral vasculatory system: the capillaries. The distribution of these across the brain is heterogeneous due to their role in supplying nutrients and oxygen to local tissue, with the structure and number in a given area being tightly linked to the local functional requirements [Edvinsson et al., 2002].

2.1.3 Oxygen Transport

We can consider oxygen transport to consist of two main phases: dissolved in the plasma and bound to haemoglobin in red blood cells [Pittman, 2011], with the primary method being the binding of oxygen to haemoglobin. The binding of oxygen to one of the haemoglobin molecule's four binding sites, to form oxyhaemoglobin (HbO_2), increases the oxygen transport capability by a factor of around 70 as compared to dissolution of oxygen in plasma alone [Effros, 2012, Chap. 13].
The quantity of dissolved oxygen is calculated using Henry's Law [Henry, 1803]

$$[O_2]^{(plasma)} = a_{O_2} PO_2 \tag{2.1}$$

where a_{O_2} is the solubility co-efficient of oxygen in blood and PO_2 is the partial pressure of oxygen in the blood. At normal conditions, this would account for approximately 3% of the oxygen carried by the blood [Hall and Guyton, 2011], with oxygen reversibly bound to haemoglobin making up the other 97%.
We most commonly model the binding process via the Hill equation [Hill, 1910], which allows us to obtain fractional saturation of haemoglobin by oxygen. Haemoglobin has four binding sites, and so a full description therefore requires "four experimental constants and is given by the Adair equation" [Goldman, 2008]. The Hill equation is given by

$$S(P) = \frac{P^{n_h}}{P^{n_h} + P_{50}^{n_h}} \tag{2.2}$$

where P here gives the local partial pressure of oxygen, P_{50} is the partial pressure giving half maximal saturation and n_h is the hill exponent. The equation is only highly accurate within the 20-80% range of saturation, but this is often deemed

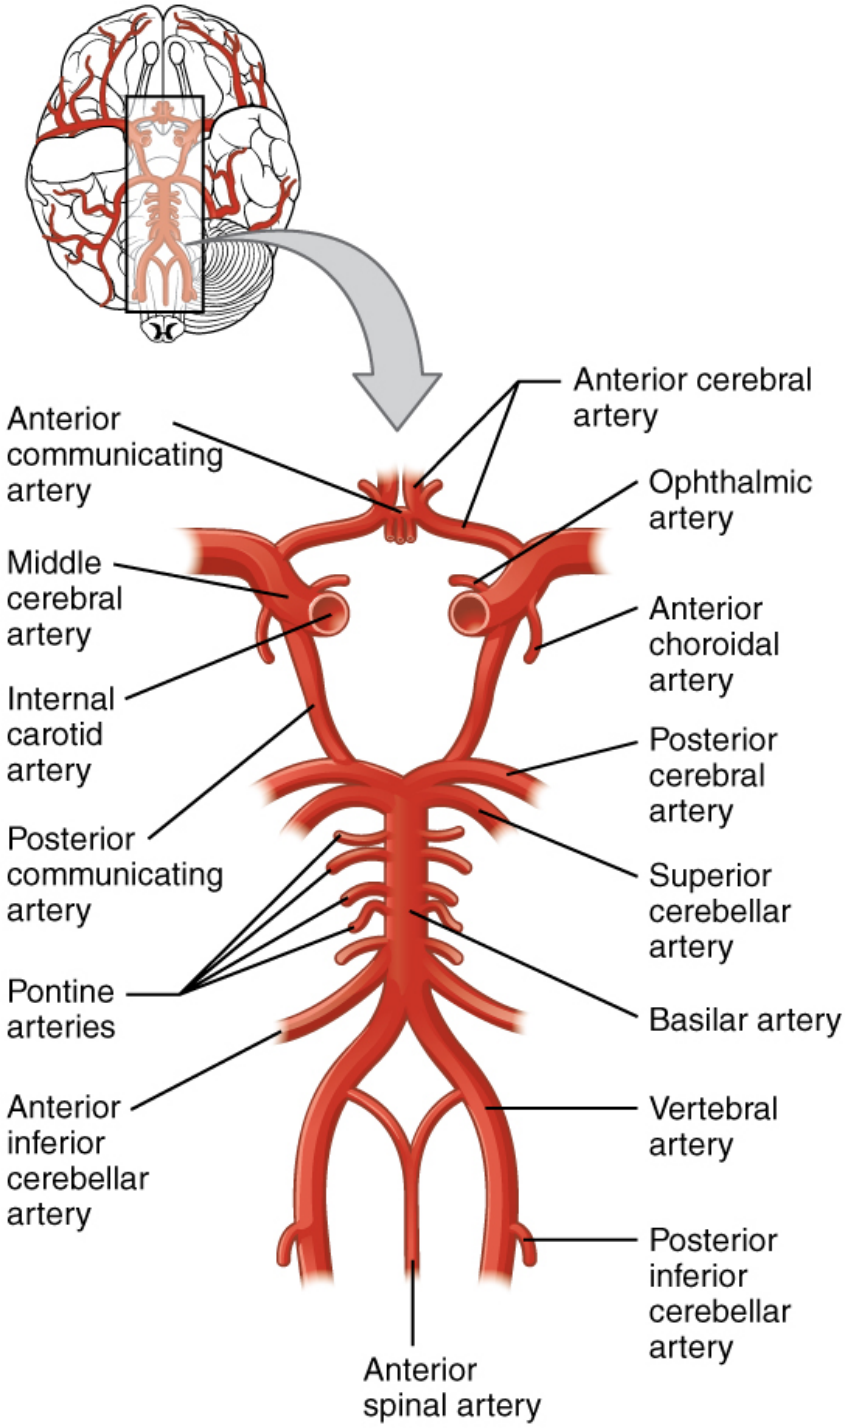

Figure 2.1: Annotated Illustration of the circle of Willis. Image taken from Open Stax College and is reproduced without changes under Creative Commons license `http://creativecommons.org/licenses/by/3.0/`. A schematic of the larger cerebral vessels in the brain, focussing on the Circle of Willis.

sufficient for use in most oxygen transport models given the number of other approximations often made [Goldman, 2008].

Oxygen that dissociates from haemoglobin enters tissue via diffusion, as does the small amount of oxygen dissolved in the plasma. This can be modelled using Fick's Law of Diffusion [Pittman, 2011]. The partial pressure of oxygen (pO_2) in the blood is much higher than that in the interstitial fluid surrounding the cells; approximately 95 mmHg vs 40 mmHg respectively. This creates a pressure gradient. This diffusive process is also a *carrier-facilitated* diffusive process [Schultz et al., 1974], with the facilitative carriers in this case being haemoglobin in the blood and myoglobin in the tissue.

The dissociation of oxygen from haemoglobin depends on a number of factors that control the 'haemoglobin affinity for oxygen'. These factors affect the dissociation curve shown in Figure 2.2. The affinity to oxygen of haemoglobin before the first oxygen molecule is bound is lower than for the subsequent molecules, with the first O_2 molecule inducing a change in the shape of the molecule that increases its affinity for following O_2 molecules. As blood enters the peripheral tissue, the pO_2 decreases leading to unbinding of O_2 molecules. This facilitates further unbinding as with the removal of O_2 the affinity decreases. This is seen in the steep curve below around 50 mmHg where a small decrease in pO_2 leads to a large decrease in oxygen saturation.

Leftward and rightward shifts in the dissociation curve can result from changes in other local environmental factors. A rightwards shift favours unloading oxygen compared to prior to the shift, whilst a leftward shift favours oxygen loading. Table 2.1 summarises some of these changes. Increasing temperature, concentration of 2,3-diphosphoglycerate (2,3-DPG), partial pressure of carbon dioxide (pCO_2) and acidity/concentration of H^+ all lead to a decrease in oxygen affinity, whilst a decrease in any of these leads to an increase in oxygen affinity. The *Bohr effect* describes the relationship between CO_2, acidity and O_2 affinity. An increase in CO_2 leads to an increase in acidity and therefore a decrease in O_2 affinity, leading to increased oxygen delivery. This is clearly helpful in areas where CO_2 is being

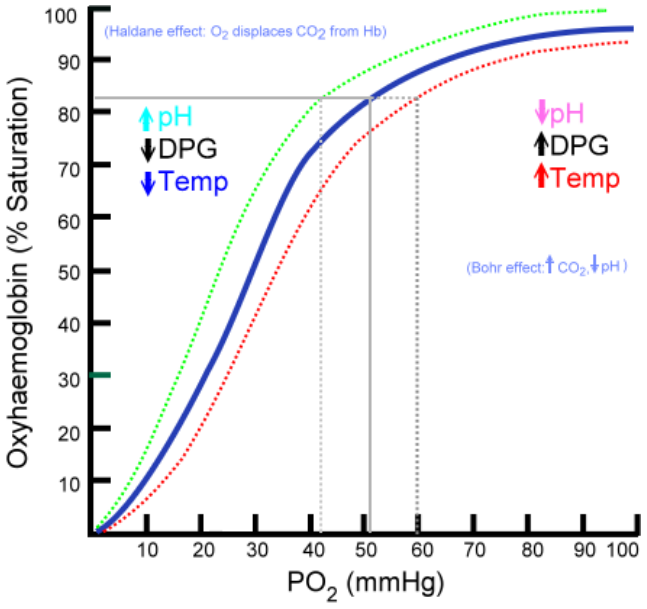

Figure 2.2: The dissociation of oxygen from haemoglobin molecules depends on a number of factors that determine the 'haemoglobin affinity for oxygen' which can be shown, as it is here, as a sigmoidal curve. The *x*-axis gives the pO_2 whilst the *y*-axis gives the oxygen saturation. Image taken from `https://commons.wikimedia.org/wiki/File:Oxyhaemoglobin_dissociation_curve.png` and is used without changes under Creative Commons license `http://creativecommons.org/licenses/by/3.0/`.

Environmental Factor	Change	Curve Shift
Temperature	↑	→
	↓	←
2-3-DPG	↑	→
	↓	←
pCO_2	↑	→
	↓	←
Acidity ($[H^+]$)	↑	→
	↓	←

Table 2.1: Impact of changes in environment on dissociation curve. The table gives the direction of the dissociation curve shift due to changes in various environmental factors. A rightwards shift indicates favouring unloading of O_2 whilst a leftwards shift indicates favouring loading of O_2.

produced due to cellular respiration.

In addition to the above, there are a number of other important factors that affect the dissociation of oxygen. Carbon monoxide CO binds to haemoglobin 240 times more readily than oxygen and will displace it from haemoglobin, forming carboxy-

haemoglobin. This shifts the dissociation curve for the remaining HbO_2 to the left, leading to reduced oxygen delivery in the peripheral tissue.

Finally, of particular relevance to this research, is the difference in the dissociation curve for fetal haemoglobin (HbF) as compared to normal haemoglobin. This is explored in further details in Section 2.3.

2.1.4 Cellular Metabolism

The energy demands of neurons are extremely high and so it's important to have an understanding of how cellular metabolism works. The energy is supplied by adenosine triphosphate (ATP) which must be produced in sufficient quantities in each part of the brain so as to meet the energy needs and maintain healthy function. Glucose is required in order to produce this ATP, with 75% of glucose going towards formation. Glucose is transported from the blood stream into surrounding tissue and is then trapped in cells by the enzyme hexokinase via the addition of a phosphate to the glucose molecule, forming glucose-6-phosphate [Payne, 2017, Chap. 1].

There are two types of metabolism, aerobic and anaerobic, and it is important to briefly consider both, especially given the focus of this work on hypoxic injury. Both share an initial pathway of glycolysis, with aerobic respiration finishing with the Krebs cycle, or tricarboxylic acid (TCA) cycle, and oxidative phosphrylation.

Glycolysis is a metabolic pathway which converts glucose into pyruvate ($C_3H_4O_3$) and a hydrogen ion H^+, occurring in the cytoplasm. The overall reaction is shown in section 2.1.4 giving a net yield of 2 ATP.

$$\begin{aligned} &C_6H_{12}O_6 + 2\,ATP + 2\,NAD^+ + 2\,ADP + 2\,P_i \\ &\longrightarrow 2\,C_3H_4O_3 + 4\,ATP + 2\,NADH + 2\,H^+ \end{aligned} \tag{2.3}$$

Under aerobic conditions, the pyruvate is transported to the mitochondria and converted to acetyl CoA before entering the TCA cycle. During the TCA cycle, additional ATP and NADH are produced. NADH is an electron donor meaning that it can be used to convert oxygen into water, releasing energy in the process [Bale et al., 2016]. This process is mediated via the mitochondrial electron transfer chain

(ETC), which in effect converts adenosine diphosphate (ADP) to ATP in the process via oxidative phosphorylation.

This process harnesses energy through a series of protein complexes (known as complexes I to V) and can be broken down into two parts: the oxidation of NADH and other electron donors in the ETC via complexes I to IV, and the phosphorylation of ADP to ATP at complex V. Of interest to this work is the terminal electron acceptor cytochrome-*c*-oxidase (CCO), also known as complex IV, which reduces the electron transport protein cytochrome *c*, with the electrons being used to convert oxygen to water [Bale et al., 2016].

CCO is of interest because one of its four redox sites, copper A (Cu_A), dominates the absorption spectrum in the near-infrared region, with a strong peak in its oxidised form, and is therefore a useful chromophore. This is examined in more detail in Chapter 3. Figure 2.3 outlines this aerobic metabolic process diagramatically, highlighting the different steps outlined above.

This aerobic process results in the overall equation

$$\begin{aligned} &6\,O_2 + \text{Glucose} + 36\,\text{ADP} + 36\,P_i \\ &\longrightarrow 6\,CO_2 + 6\,H_2O + 36\,\text{ATP} \end{aligned} \tag{2.4}$$

where it is clear to see that a single mole of glucose will produce 36 moles of ATP. If sufficient oxygen is not available however, anaerobic respiration occurs. In this, the pyruvate produced during glycolysis is reduced into lactic acid, with ATP as a by-product and the NADH produced able to be reoxidised to continue the process if oxygen is later available. This yields 2 moles of ATP for every mole of glucose used. Overall this can be defined as

$$\begin{aligned} &\text{Glucose} + 2\,P_i + 2\,\text{ADP} + 2\,NAD^+ \\ &\longrightarrow 2\,\text{Lactate} + 2\,\text{ATP} + 2\,\text{NADH} \end{aligned} \tag{2.5}$$

It is easy to see that aerobic respiration is more efficient than anaerobic respiration, yielding around 36 moles of ATP per mole of glucose compared to 2 moles of ATP per mole of glucose for anaerobic.

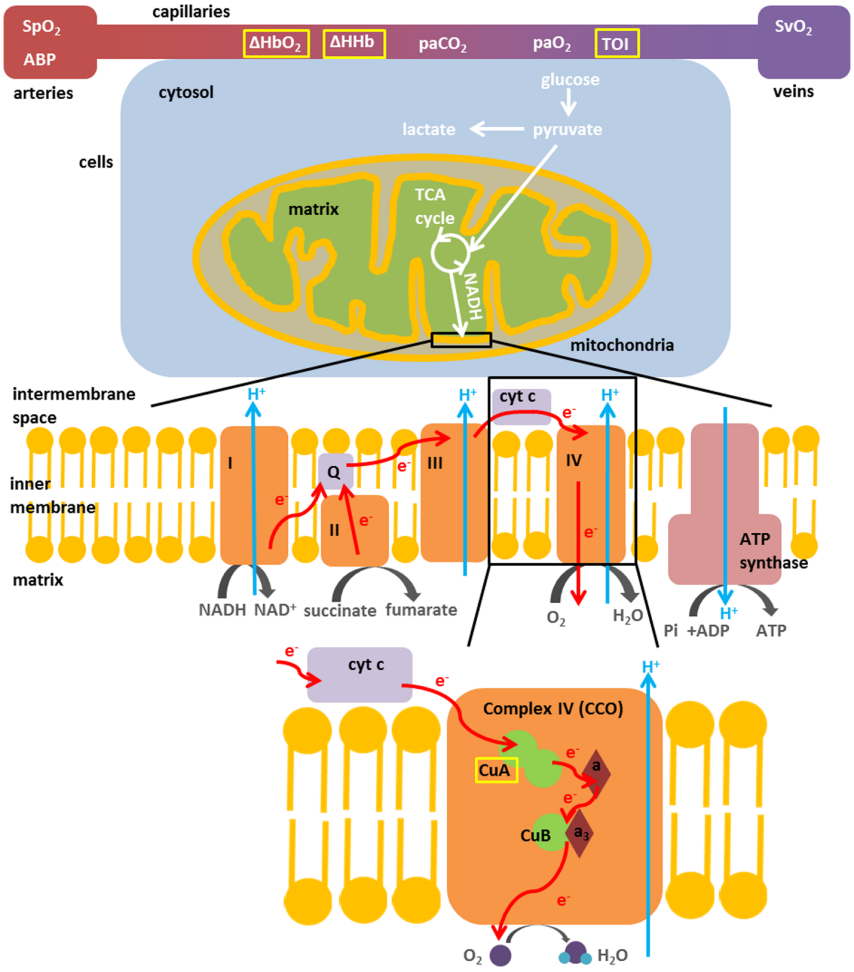

Figure 2.3: Schematic of the cellular metabolism process, taken from Bale et al. [2016]. The ETC and CCO are both expanded, with CCO represented as complex IV. Limited to carbohydrate metabolism, with fat metabolism ignored. Chromophores measured by NIRS

2.2 Hypoxic-Ischemic Encephalopathy

HIE is a condition where a lack of oxygen, due to reduced blood flow, causes damage to the brain [Vannucci, 2000], often resulting in death or long term disability [Huang and Castillo, 2008]. In particular it is an important cause of cerebral palsy and other related disabilities in children [Fatemi et al., 2009]. There are a number of methods available for detection, including EEG, MRI, MRS andNIRS. It is this last one that will be focussed on throughout, particularly bNIRS, but it is important to note alternative methods and their relative merits and drawbacks. These methods

will be explored in more detail in Chapter 3.

Murray et al. [2009] found that normal and mildly abnormal EEG results obtained within 6 hours from birth were associated with normal neurodevelopmental outcomes at 24 months, whilst Shellhaas et al. [2013] evaluated the use of amplitude-integrated EEG (aEEG) and regional oxygen saturation (rSO_2) measured using NIRS to determine short term outcomes for neonates suffering from HIE and treated with therapeutic hypothermia, a common treatment for hypoxic brain injury. Cerebral rSO_2 was found not to have any predictive value, whilst lower systemic rSO_2 variability and invariant, discontinuous aEEG patterns were strong indicators of a poor outcome.

Cowan et al. [1994] used diffusion weighted MRI and found it to indicate the extent and conspicuity of early abnormalities better than in standard imaging, whilst Forbes et al. [2000] found it to have a lower correlation with clinical markers than standard MRI when performed with the technical parameters of that study. Concentrations of brain metabolites, found through MRS, can prove useful as biomarkers of injuries, like HIE. ^{31}P-MRS can be used to measure the quantity of ^{31}P, which can be found in high energy metabolites like phosphocreatine. Moorcraft et al. [1991] found that a global phosphocreatine/inorganic phosphate (PCr/P_i) ratio below the range of values typically expected predicted adverse outcome following asphyxia, whilst Roth et al. [1992] undertook studies of cerebral oxidative metabolism using ^{31}P-MRS, and found that minimum recorded values of PCr/P_i were related to outcome. Despite this, current clinical systems are unable to perform localised ^{31}P-MRS, and so only large areas of the brain can be examined. As a result, ^{31}P-MRS is not routinely used as a clinical tool to assess asphyxiated full term neonates [de Vries and Groenendaal, 2010].

Another key biomarker, is the lactate/N-acetyl-aspartate (Lac/NAA) ratio which is measured using proton magnetic resonance spectroscopy. As has previously been stated, lactate is a byproduct of anaerobic respiration, where oxygen demand exceeds supply and any build up due to poor oxygen delivery can be observed using proton MRS. Changes in other metabolites such as NAA, which decreases with in-

creased oxygen demand, can also be observed [Xu and Vigneron, 2010]. The ratio of these two molecules is therefore directly linked to a hypoxic injury, and has been shown to be a good indicator of outcome. These can be measured through the use of proton MRS, and it has been used in neonates since the early 1990s [Peden et al., 1990, van der Knaap et al., 1990].

NIRS can be used to measure changes in oxygenation within the brain, as well as changes in metabolism, making it possible to investigate physiological changes and their potential causes and effects. This can be important to investigate how the brain responds to stimuli (functional activation)[de Roever et al., 2017], or the impact of injuries such as HIE [Bale et al., 2014]. The bedside nature of NIRS means that it is possible to obtain data at a much earlier stage than with other measurement techniques, and so there is the potential to identify the severity of a hypoxic injury at a much earlier point. Bale et al. [2018] shows that during spontaneous desaturation events following HIE, neonates with a more severe injury had a strong correlation between CCO and haemoglobin difference (HbD). The exact mechanism behind this relationship was hypothesised to be due to a decrease in the cerebral metabolic rate and will be looked at in greater detail in Chapter 9.

The large variety of techniques available and the lack of consensus on their individual reliability means that more work is required to develop their use for diagnosis, especially with a view to improving treatment.

2.2.1 Treatment

As has previously been mentioned, treatment of HIE is often done through the use of therapeutic hypothermia and this generally improves neurodevelopmental outcome [Jacobs et al., 2007]. A number of studies have shown that this improves the outcome of neonates who have undergone a hypoxic injury when compared to receiving only critical care.

Gluckman et al. [2005] found that in a randomised controlled trial, where the outcomes of those treated with and without head cooling were compared, 55% of babies who received treatment had either died or had severe disability at 18 months, compared to 66% amongst those who hadn't received treatment. Further subgroup

analysis revealed that 'head cooling had no effect in infants with the most severe aEEG changes, but was beneficial in infants with less severe aEEG changes', suggesting some underlying conditions or mechanisms may have an impact on treatment outcome. Azzopardi et al. [2009] looked at the outcomes of neonates that suffered asphyxial encephalopathy, comparing those that received intensive care and cooling of the body to those that received intensive care alone. 75 of the 163 ($\sim$45%) that received cooling either died or suffered severe neurodevelopmental disability, compared with 86 out of the 162 ($\sim$53%) that received intensive care alone. They then concluded that treatment with hypothermia 'did not significantly reduce the combined rate of death or severe disability but resulted in improved neurologic outcomes in survivors'. Finally, Shankaran et al. [2005] found that in a comparative trial between treatment of perinatal complications using full body cooling against no treatment, 44% of the group that received hypothermic treatment died or suffered from moderate or severe disability, compared with 62% in the control group. Additionally, the incidence of cerebral palsy was 19% in the hypothermia group as compared with 30% in the control group.

Whilst it is clear that hypothermic treatment does improve outcome, there are still a large number of cases in which treatment does not appear to be effective. One of the key aims of this research is to use data collected through the previously outlined methods, particularly NIRS data, alongside mathematical modelling to try and understand potential reasons for why hypothermic treatment is not as effective for some patients, particularly those with a severe injury. Chapter 8 looks at modelling the effect of hypothermic treatment on the brain's physiology.

2.3 Differences in Adult, Neonate and Piglet Brains

When developing models of the brain it is important to consider the ways in which the physiology of different brains differ from each other. For example, much of the work done to develop the UCL family of models was based originally on the adult brain. Without considering how the adult and neonate brain differ, initial parameter estimates may be incorrect leading to false conclusions about the ways

in which values obtained through fitting vary from the default value or poor initial prior distributions.

The preterm neonatal brain begins smooth and becomes increasingly folded during gestation and into term, increasing surface area. This increase in folding can be seen using a variety of imaging techniques, including MRI, right through to the post natal period [Rutherford, 2002]. In addition to this change in cerebral folding, the neonatal cerebrovascular structure continues to develop in complexity through the post natal period [Coelho-Santos and Shih, 2020]. Studies of preterm neonatal brains have shown that they contain variations in the development of the Circle of Willis [van Kooij et al., 2010] for equivalent term dates. The newborn piglet brain has an intact Circle of Willis and a blood vessel plexus at the base of the brain and so, as per Cady et al. [2008], this make the piglet brain unsuitable for investigating complete cerebral ischaemia, but still suitable for hypoxia-ischaemia.

Another difference can be seen in the dissociation curve for HbF as compared to normal haemoglobin. HbF is shifted leftwards favouring O_2 binding at lower partial pressures as compared to normal haemoglobin. This is important in the womb as it allows the fetus to pull oxygen from the mother's blood. Figure 2.4 illustrates this difference.

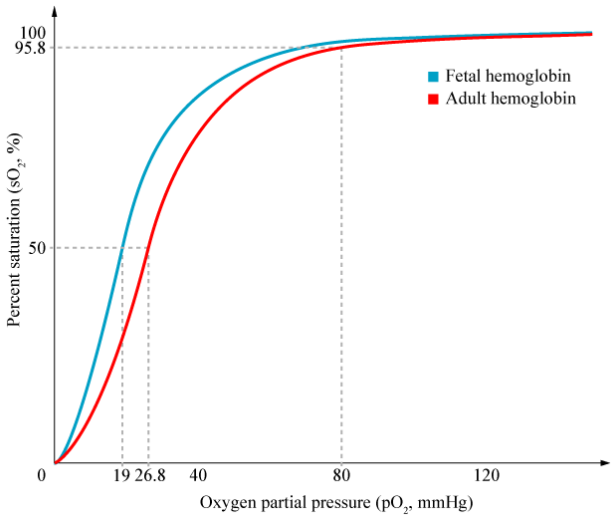

Figure 2.4: Fetal dissociation curve Dissociation of oxygen from fetal haemoglobin differs from that of adult haemoglobin, with a notable leftwards shift in the dissociation curve as seen here.
Image taken from `https://commons.wikimedia.org/wiki/File:HbA_vs_HbF_saturation_curve.png` and is used without changes under Creative Commons license `http://creativecommons.org/licenses/by/3.0/`

2.4 Conclusion

In this chapter I have briefly summarised the physiology of the brain, focussing in particular on those mechanisms and structures relevant to HIE. It is important to understand the biological reality that the modelling is required to emulate. Without this background information it is too easy to separate out the computational and mathematical abstractions from the physical, biological reality we are studying. I have also focussed here on briefly outlining the methods of determining the severity of HIE injury and its treatment via therapeutic hypothermia. Finally, I established some of the physiological differences between the adult, neonatal and piglet brains, which is important when developing the separate models.

Chapter 3

Clinical Monitoring Techniques

3.1 Measurement Techniques

There are a variety of methods available to image and monitor the brain, utilising different phenomena or aspects in order to collect information. In this chapter I will outline a subsection of available measurement techniques, focussing on those that are typically used by the research group at UCL.

3.1.1 Electroencephalography

EEG measures the electrical activity of the brain, typically through the use of electrodes placed along the scalp. Neurons encode information through changes in electrical energy, and has been used within the UCL team as part of research investigating metabolism following neonatal stroke [Mitra et al., 2016]. A neuron possesses a resting potential of -60 to $-70\,\mathrm{mV}$, with there being relatively more sodium ions outside the neuron and more potassium ions inside. Upon receiving stimulation, sodium channels open creating a flux of positive ions into the cell, increasing the potential and potassium channels open slightly later, leading to an out flux of potassium ions, decreasing the potential, shown in figure 3.1. The action potential propagates as a wave along the axon.

At a simple level, EEG is able to measure these changes in voltage due to the ionic current and through this it is possible to get a picture of the electrical activity occurring within the brain. Diagnostic applications generally focus either on specific events, or on the spectral content of the EEG signal. The former investigates po-

Removed due to copyright

Figure 3.1: Simplified diagram of a typical action potential, with the opening and closing of ion channels labelled. Image from *Neuroscience for kids*, University of Washington

tential fluctuations linked directly to an event, such as during functional activation, whilst the latter analyses the type of neural oscillations that can be observed in EEG signals in the frequency domain [Schomer and Lopes da Silva, 2017].

3.1.2 Magnetic Resonance Imaging and Spectroscopy

MRI is a medical imaging technique that uses the phenomena of nuclear magnetic resonance (NMR) to generate detailed images. When placed in a strong external magnetic field, certain atomic nuclei (those that contain an odd number of protons and/or neutrons) are able to absorb and emit radio frequency energy. This energy is at a specific resonant frequency dependent on the strength of the applied magnetic field. In a clinical setting, the most commonly used atomic nuclei are those of the hydrogen atom, which possesses just one proton ion in its most abundant isotope, and is found in water and fat throughout the body. Different methods can be used in order to generate different kinds of images. One of the key benefits that MRI has over other commonly used imaging techniques, such as CT scans, is that it does not require the use of x-rays and is there fore non-ionising.

An alternative use of NMR is magnetic resonance spectroscopy (MRS), which

Removed due to copyright

Figure 3.2: Example MRS spectra, with choline (Cho), creatine (Cr), N-acetyl-aspartate (NAA) and myo-Inositol (mI) labelled. Image from Blüml [2013]

serves as a good complement to MRI. Where MRI uses the signal from hydrogen protons to form anatomical images, MRS instead uses the information to determine, for our purposes here, concentrations of brain metabolites, such as N-acetyl-aspartate, choline, creatine and lactate in the tissue examined [Gujar et al., 2005]. As has previously been mentioned, the resonant frequency of protons is, to a first approximation, a function of the magnetic field strength. However, the electronic

Removed due to copyright

Figure 3.3: Absorption spectra for various chromophores found in human tissue. Note: CtOx is cytochrome-*c*-oxidase (CCO) and O_2Hb is oxyhaemoglobin (HbO_2). Figure taken from Scholkmann et al. [2014].

environment of molecules has an effect on the magnetic field '*seen*' by the proton. If electrons are close to the proton, there is a *shielding* effect, reducing the size of the magnetic field seen by the proton, causing a shift in the resonant frequency. These shifts are typically calculated relative to a reference molecule, tetramethylsilane, which is defined as 0 ppm [Blüml, 2013]. Figure 3.2 shows a MRS spectra, with choline, creatine, N-acetyl-aspartate and myo-Inositol labelled. This information can be used to understand further what is occurring within the brain. For example, lactate is a byproduct of anaerobic respiration, where oxygen demand exceeds supply [Barkovich et al., 1999], thus, in the case of reduced oxygen delivery it builds up within the tissue. This build up can be observed through the use of MRS.

3.1.3 Near-infrared Spectroscopy

Compounds in tissue known as chromophores absorb light as it passes through the body. Visible light has high absorption and scattering in most tissue and is therefore not useful for clinical use. Near-infrared light (650 nm to 1000 nm) however has a lower absorption in water and is therefore able to penetrate further into tissue, whilst also showing a useful difference in absorption for deoxyhaemoglobin (HHb) and HbO_2, allowing the two to be distinguished from each other. Figure 3.3 shows the absorption spectra for various chromophores in human tissue, with the near-infrared optical window marked. By transmitting a specific wavelength or spectra of light into tissue and then measuring what is scattered back to the detector it is possible, using the modified Beer-Lambert law [Delpy et al., 1988], to determine the change in concentration of these chromophores. It is this principle which underlies the use of NIRS.

One notable use is in the monitoring of cerebral haemodynamics and, in the case of bNIRS, metabolism. This is due to two key properties: bone is relatively transparent to near-infrared light [Jöbsis, 1977], and the only compounds within the mam-

malian brain that respond to changes in oxygen concentration are haemoglobin and cytochrome aa3 [Wray et al., 1988]. Because of these two facts, NIRS is able to detect and measure changes in the oxygenation of haemoglobin and cytochrome inside the brain without requiring any invasive techniques.

Typically, haemodynamics are measured by detecting changes in the concentration of HbO_2 and HHb, which can be done with a small number of discrete wavelengths. It is then possible to use these signals to determine the change in total haemoglobin ($\Delta HbT = \Delta HbO_2 + \Delta HHb$), a proxy for cerebral blood volume, and the change in haemoglobin difference ($\Delta HbD = \Delta HbO_2 - \Delta HHb$), a proxy for cerebral blood oxygenation. There is additional benefit in using ΔHbD as the subtraction removes some of the signal noise. Whilst haemoglobin can be measured using a small number of discrete wavelengths, measurement of CCO requires a broader selection of wavelengths to distinguish between the oxidised and reduced state. The enzyme contains four redox active metal centres, with CuA having a broad absorption peak in the near-infrared region [Bale et al., 2014]. However, the concentration of this is much lower than haemoglobin (approximately 10%) which can make detection difficult [Cooper and Springett, 1997]. As a result, bNIRS is required in order to resolve the CCO signal, with the downside that there is a higher SNR than when using the two optimal wavelengths required for haemoglobin alone [Scholkmann et al., 2014].

The main use of measuring CCO is that it is a key indicator of metabolic activity and therefore oxygen use. When considering something like brain injury it is just as important to measure whether supplied oxygen is being used as it is to measure if it is being supplied. As described in Chapter 2, CCO is a key enzyme in the metabolism of oxygen and, because the total quantity remains constant over the measurement time period, the change in its redox state is a good proxy for metabolic activity.

3.2 Systemic Measurements

In addition to measuring the cerebral specific measurements described above, it is important to measure systemic quantities such as blood pressure and arterial oxygen

saturation. Without these, it is difficult to establish information about phenomena such as cerebral autoregulation.

Arterial Oxygen Saturation Arterial oxygen saturation is typically measured using a pulse oximetry device which uses two wavelengths of light to measure HbO_2 concentration in pulsing blood (SpO_2) [Aoyagi, 2003]. The pulsatile volume is assumed to be due to changes in arterial volume and so SpO_2 is used as a measure of arterial arterial oxygen saturation (SaO_2).

Blood Pressure A sphygmomanometer is the most common device in clinical use for measuring blood pressure, with a finometer often used within a research context. It consists of an inflatable cuff, usually placed on the upper arm. Measurements are usually given as diastolic (maximum) and systolic(minimum) arterial pressures, with typical values in a healthy adult being 120 mmHg and 80 mmHg respectively.

Partial pressure of CO_2 The arterial partial pressure of CO_2 ($PaCO_2$) is a good measure of the efficiency of ventilation. It is normally determined by measuring the end tidal CO_2 ($EtCO_2$), the partial pressure of CO_2 breathed out during respiration. This CO_2 diffuses out of the blood in the capillaries into the lungs and, under normal conditions, there is an equilibrium between the blood and the air meaning that $EtCO_2$ is equivalent (with an offset) to $PaCO_2$ [Ward and Yealy, 1998]. $EtCO_2$ can be measured using an infrared capnometer with a normal measurement of 35 mmHg to 40 mmHg[Kodali, 2013].

3.3 Interpretation of Measurements

As with any data collection process, the measuring of physiological signals using any of the above techniques will undoubtedly result in erroneous measurements, redundant data and signal noise. The cause and effect of each of these will vary depending on the measurement technique and the situation in which it was collected. For example, collecting longitudinal data from neonates in intensive care will result

in far more artefacts than an equivalent period of time measuring a piglet in a controlled lab environment. As is discussed below in Section 3.4, these artefacts can come from ambient light, movement of the baby and clinical check ups.

In terms of redundant data and signal noise, depending on what physiological phenomena are under consideration, aspects of the data may prove to be redundant and have the potential to make it more difficult to analyse. An example of this could be the approximately 60 Hz signal often found in haemodynamic data that is due to the resting heart rate. Where we are attempting to analyse data over longer periods, removing this noise can make it easier to isolate specific behaviours relevant to the work. One way of removing this sort of high frequency noise is through applying a low pass filter to the data. In this work, this is often a 5th-order Butterworth filter implemented in Python using SciPy [Virtanen et al., 2020].

The term "artefacts" here refers to erroneous measurements within the data that occur due to outside factors which can lead to data being less reliable and vary in cause and impact size. It is ideal if artefacts are minimised or otherwise outright removed before modelling, thus avoiding fitting a model to data that are not physiology.

In term of removing artefacts, the typical method has been to utilise a combination of automated methods, such as MARA which utilises moving standard deviation and spline interpolation [Scholkmann et al., 2010, Metz et al., 2015], and manual detection and removal. However, a small amount of work was undertaken in the process of this work to try and develop an automated detection technique that could utilise machine learning and the broadband spectra collected from bNIRS.

3.4 ABroAD

The following section has been adapted from conference proceedings published for the ISOTT 2017 conference.

As has been previously mentioned, bNIRS instruments use hundreds of wavelengths of light to observe changes in tissue oxygenation and metabolism, collecting absorption spectra at each time point. Figure 3.4 shows three different spectra collected during the course of an experiment to simulate some typical NIRS artefacts. It can

be seen here that the spectra of light at the three time points are clearly different depending on whether an artefact is present and what kind of artefact it is.

With this in mind, an attempt was made at using machine learning techniques to

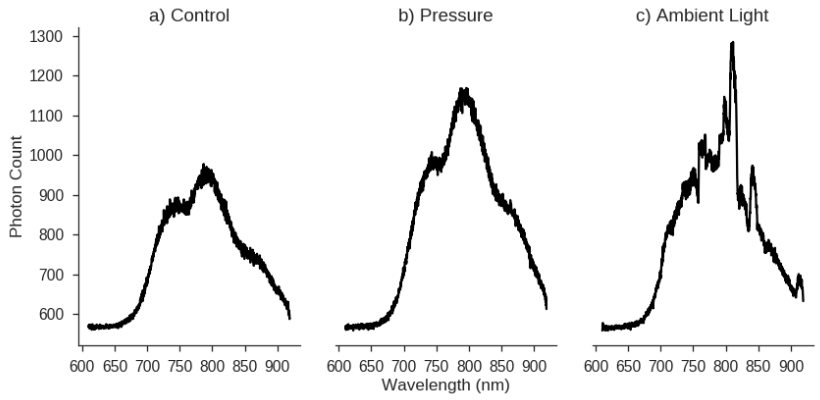

Figure 3.4: Three spectra collected during the simulation of NIRS artefacts. Figure a) shows the spectra collected when no artefact occurred. Figure b) shows the spectrum collected when pressure was applied to the sensor. It has a similar shape to that of the control spectrum, but the photon counts are generally higher across the spectrum. Finally, figure c) shows the spectrum collected when an ambient light source was turned on in the room. Not only are the photon counts generally higher than for the control spectrum, but a number of spikes at specific frequencies are also visible. This is due to the ambient light source being fluorescent, thus producing light at specific wavelengths.

develop a new approach to identifying artefacts in collected data. If such a method were developed, even without the ability to replace the artefacts with reasonable data, the model could be driven by and compared against data that was artefact free. Machine learning refers to the process of identifying 'patterns' within data to try and understand it, preferably in a way that will allow this understanding to be used with new data [Domingos, 2012]. Data is normally represented within machine learning as a set of 'features'. For example, a document may be presented in terms of the words it contains or other information like the length of the sentences. Information about an animal may be done as a mixture of categorical data, like breed or colour, and quantitative data, like size and weight. Features may already be present in the data or new ones may be engineered from the data available.

These data are then used with an algorithm that processes the data and produces output such as a classification or a predicted value. For example it may predict the

breed of dog (a classification), or it may try to predict the weight of the dog (a forecast or regression). The quality of this prediction is then evaluated using a metric. In the case of a predicted value, that may be the error in the value, or in the case of a classification it may be the classification accuracy or some other suitable metric.

For this investigation, bNIRS data was collected from eight different, healthy, physiologically normal subjects in a block test design using the CYRIL NIRS system, previously described by Bale et al. [2014], at a sample rate of 5Hz. Two sensors were used: a short separation sensor at 10 mm, and a long-separation sensor at 30 mm. Six different artefacts were simulated - horizontal motion (shaking head), vertical motion (nodding head), frowning, pressure on sensor, ambient room light and directed torch light - in ten second blocks repeated twice, with ten seconds of rest between each, leading to roughly 50 data points for each block. Figure 3.5 shows the oxyhaemoglobin trace generated during the artefact simulation exercise. A number of significant artefacts can be seen, particularly towards the end of the time series.

The start and end point of each artefact, as well as the start and end points of the

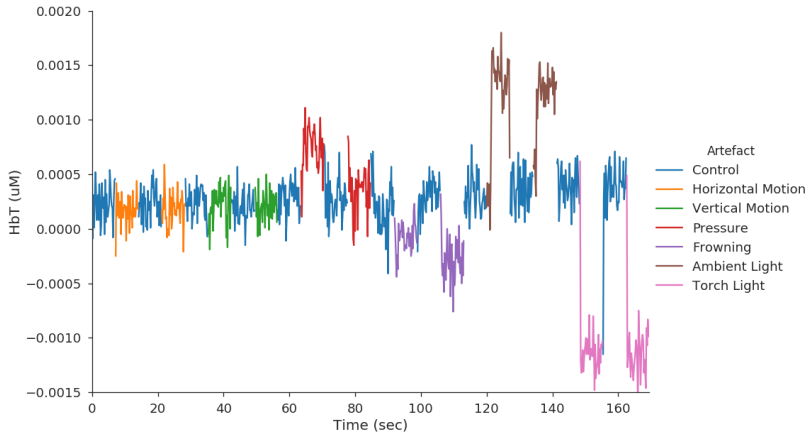

Figure 3.5: Example of artefacts in the total haemoglobin trace of NIRS data. Each artefact type simulated is marked in a different colour, with the signal where no trace occurs in blue. It is clear that some artefacts, such as ambient light, have a bigger impact on the signal than others, such as vertical movement.

experiment, were marked in the output data as an event using the CYRIL software. All artefacts were simulated in the order listed above for all subjects.

At each time point a spectrum of light was collected for 1340 wavelengths between 610 nm and 920 nm. When applying a machine learning algorithm to this data it is important to note that not only does this data have an extremely high dimensionality, but it can also be deemed to be functional i.e each wavelength is functionally related to its neighbouring wavelengths. Many machine learning approaches assume data points to be independent. In order to reduce the dimensionality and generate features that are not functionally related, feature engineering was undertaken.
Feature engineering describes the process by which new features are generated from existing data. In the case of the spectral data used here, this was done to reduce dimensionality and remove any obvious functional relationship between features.
Four features were generated for each spectra, with the distributions of the values for each sensor and for both sensors combined plotted. These four were selected based on heuristics and empirical observations of the spectra from expert users.

Area Under Curve (AUC) Many artefacts, particularly those due to changes in light, showed an increase in intensity for many wavelengths. With the spectra representing the photon count at specific wavelengths, an increase in the area under the curve (spectrum) represents an increase in intensity. The area was calculated in Numpy [van der Walt et al., 2011] by integrating using the trapezoidal rule.

Fractional Power Density The ambient light artefact is one of the most noticeable, as the fluorescent lights used in the room led to spikes in intensity at specific wavelengths, as seen in 3.5c). It was found that the fraction of the integrated power density spectrum occupied by the top 99% of frequencies (referred to as the **fractional power density**) was generally lower in the spectra containing ambient artefacts as compared to a control.

Autocorrelation Autocorrelation is the correlation of a signal with a time delayed copy of itself as a function of delay. Because autocorrelation is often used to identify information that is otherwise hidden by noise, a difference between the control spectra and spectra due to artefacts was expected, as well as a difference between spectra from each artefact.

Sample Entropy Sample entropy is a modification of approximate entropy, chosen due to it being more computationally efficient, and is a measure of the complexity level within a signal [Richman and Moorman, 2000]. It was expected that by applying it to the spectra, a difference in complexity between artefacts and control signals would be observed.

Figure 3.6 outlines the machine learning process. For each subject, i, the spectrum at each time point, t_j was converted into a 4-dimensional feature vector, $\mathbf{x}_{i,j}$ of the form $\langle x_{i,j,1}, x_{i,j,2}, x_{i,j,3}, x_{i,j,4} \rangle$ and assigned a classification $y_{i,j}$ according to the artefact simulated at that time point. This data set was then split into test and training sets, based on the subject number, i, using a k-fold method. This ensured that the algorithm was tested on data from an unseen subject, preventing information about that data leaking into the training process.

The training data was then fed into a machine learning pipeline consisting of two main steps: *scaling* and *estimation*. *Scaling* was done using the 'RobustScaler' from the Scikit-learn library [Pedregosa et al., 2011] to ensure all features were roughly equatable in terms of magnitude. This removes the median and scales the data according to the interquartile range. Without this step, features that were of significantly different magnitudes to others e.g AUC, may receive undue weighting in the estimation process. The *estimation* method chosen was a Random Forest Classifier [Breiman, 2001], which can natively handle multiclass classification problems natively, unlike a binary classifier which requires an additional step to convert from multiclass to binary, and has been shown to be a robust method with minimal requirements for extensive tuning.

The classifier was fit using a grid search, cross validation method [Hsu et al., 2008]. The training data was split by subject into training and test sets M times, thus allowing retesting of the method on different permutations of the overall training set. The classifier was run for different parameter combinations, and the set that provided the best final score was chosen as the best estimator. This was then trained on the total training set and tested on the initial test set. This final score allows the effectiveness of the chosen method to be evaluated.

The metric chosen here was the 'weighted F1-score' [Chinchor, 1992], which takes into account both precision (p) and recall (r) and was defined as $F1 = 2\frac{p \cdot r}{p+r}$. A perfect classification has a score of 1, whilst no correct classifications would have a score of 0. Because it takes into account both precision and recall, it was able to deal with the class imbalance inherent in our data. Precision is the fraction of correct classifications for a class j out of the total number of predictions of that class, whilst recall is the fraction of correct classifications for a class j out of the total number of actual occurrences of that class. By using both of these measures, we ensure that an estimator was developed which finds all occurrences of an artefact j without making too many false classifications.

Computational Tools All of ABroAD was implemented in Python, using SciKit-Learn [Pedregosa et al., 2011], Numpy [van der Walt et al., 2011] and Pandas [Mckinney, 2010]. Data visualisation was done using Matplotlib [Hunter, 2007] and Seaborn [Waskom et al., 2017].

3.4.1 Results

Training and test sets were selected randomly, splitting by subject. The same test and training splits were used for all results. The algorithm was trained using data from subject 1, 2, 3, 4, 6 and 8 and tested against data from subjects 5 and 7.

A green horizontal line is marked on the bar charts to show F1-scores greater than 0.7. The choice of 0.7, whilst having no particular significance, does provide a reasonable benchmark to compare classification between artefacts against. Three combinations of artefacts were considered:

All All artefact types were used for training and testing.

Motion Only control and motion artefacts were used: *Horizontal, Vertical Pressure, Frowning*.

Light Only control and light artefacts were used: *Ambient, Torch*.

Figure 3.7 shows the f1-scores for the short and long separation sensors, with different bars showing each run type. Scores for the long-separation sensor are gener-

ally better than those for the short-separation sensor, with scores for all light-only artefacts being above the chosen threshold of 0.7 for all three metrics. Scores for light-only runs also generally outperform runs for all artefacts and for motion-only artefacts.

In both short and long separation sensors, the control spectra can be identified relatively well when only looking at light artefacts. This suggests that even where the type of artefact cannot be determined, the presence/non-presence of an artefact can still be identified.

The average F1-score for the light-only dataset is significantly higher than the average score for the datasets containing all artefacts. This high level of accuracy is clearly visible in Figure 3.8, which shows the oxyhaemoglobin trace for the control and light-based artefacts. The artefact labels shown in 3.8b were predicted using spectra from the long-separation sensor and without any motion based artefacts in the dataset.

In general, the motion artefacts scores are much lower than the light artefact scores, with the horizontal and vertical artefacts rarely being correctly identified. This may be due to one of two reasons: *a*) the choice of features didn't capture sufficient information for these artefacts, or *b*) during data collection, the simulation process was not sufficient in generating an artefact within the data. The method shown here has great promise. At least one subset of artefacts (light artefacts on the long-separation sensor) were correctly identified to a good level.

In order to develop the process further, two steps are required. A dataset needs to be collated that has time series and spectra collected during 'normal/typical' usage, with artefacts then correctly labelled. The use of "synthetic" artefacts is suitable for exploratory work, but the artefacts simulated may not be representative of those found in real data. In particular, some of the spectra labelled as containing an artefact show very little difference as compared to the control spectra, suggesting no artefact was actually simulated during data collection. Defining what qualifies as an artefact will allow datasets to be correctly labelled and improve the usability of these methods.

Secondly, in this exploratory work, only four features have been engineered from the data and a number of artefacts are unable to be identified using from these. If other features can be engineered from the data that allow better separation of one or more of the poorly performing artefacts, the classification power of the machine learning process could be improved greatly. Thus, future work should consider developing new features from the data and using this to improve the machine learning pipeline. However, in order for this to be successful, it is also a necessary prerequisite that more data are collected to allow more and more rigorous training of the model. It is also important that any further data also include genuine physiological events that could be mistaken for artefacts by an alternative method. This would be both the key test of the ABroAD method, as well as being the most important component of model development. Without including data that contain the events of interest, it would be impossible to train the algorithm to recognise them and to validate that the method will not incorrectly identify these as artefacts.

Once a trained algorithm has been developed and tested sufficiently, there is the potential for real time identification. Incorporating this method alongside existing systems could allow for collected data to have potential artefacts marked in the time series, improving the efficiency of data curation.

Finally, the ability of the model to identify light artefacts to a good degree could be extended to non-broadband systems by collecting light data through a small external sensor. This externally collected light data can then be fed into a similar pipeline and used to identify light-based artefacts.

3.5 Conclusion

In this chapter I have outlined a number of measurement techniques (EEG, MRI, MRS and NIRS)and how they can be used to monitor the brain. I then outlined systemic measurement techniques, of particular interest here as it is systemic data that are used to drive the BrainSignals models and reproduce cerebral physiology measurements. I then discussed a little bit about the interpretation of these measurements before outlining work done to try and improve on this. This work focussed

on the development and use of a novel machine learning based artefact detection technique called ABroAD which was develoepd to determine physical and light-based artefacts in bNIRS data. Whilst this technique was not ultimately successful enough to be used widely, it was able to be used sufficiently well in order to identify light-based artefacts. Further research into its use would require more data that contained labelled, accurate signal collected during verified artefact events.

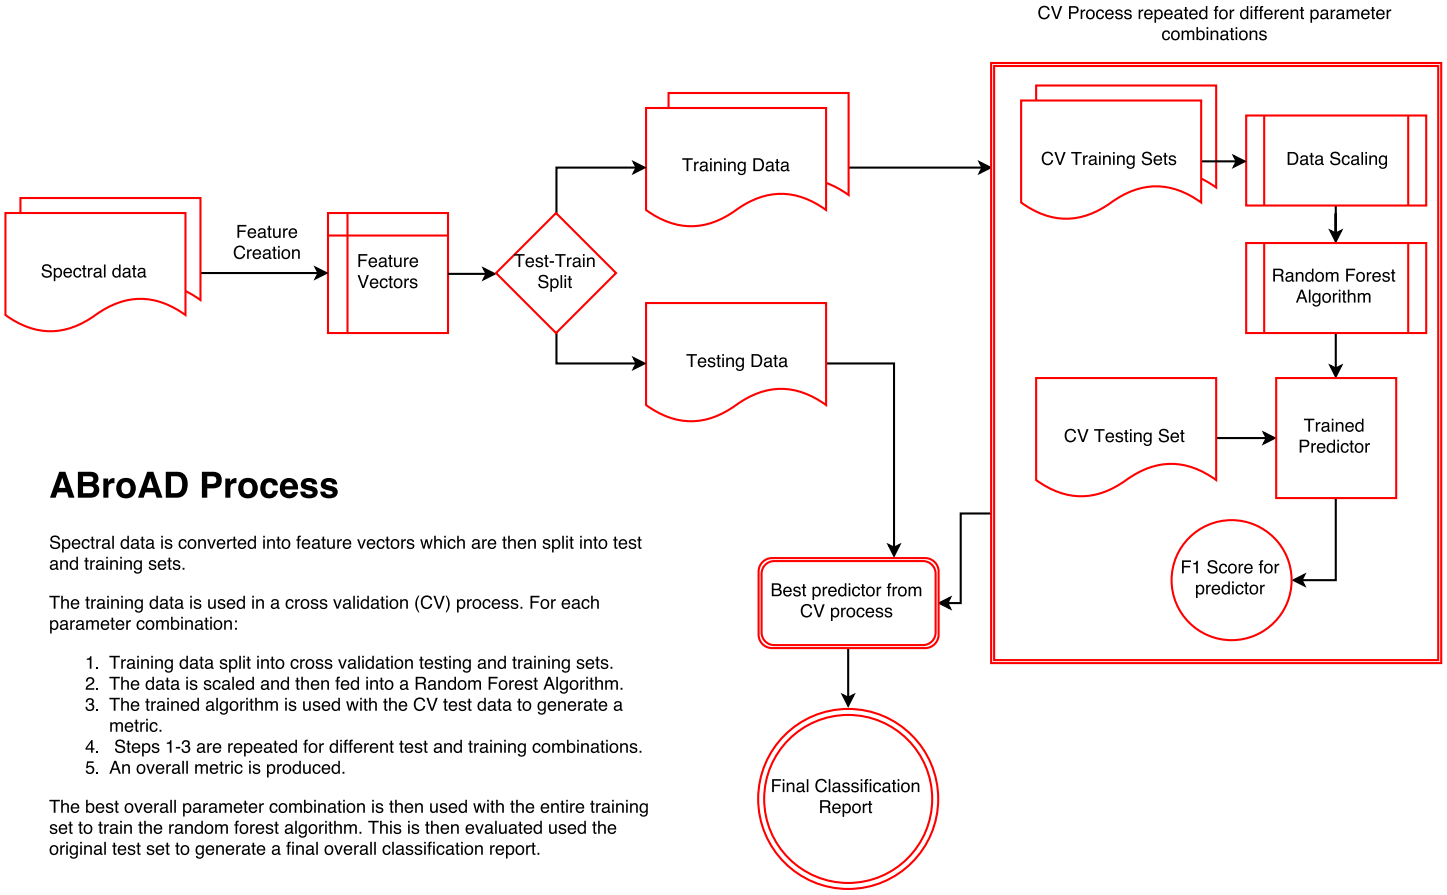

Figure 3.6: The machine learning process.

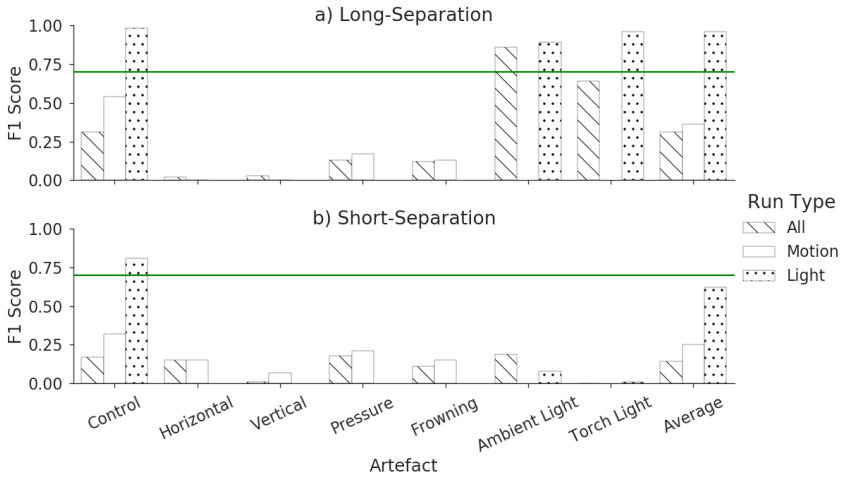

Figure 3.7: F1-scores for each artefact on each sensor, with separate bars for each run type.

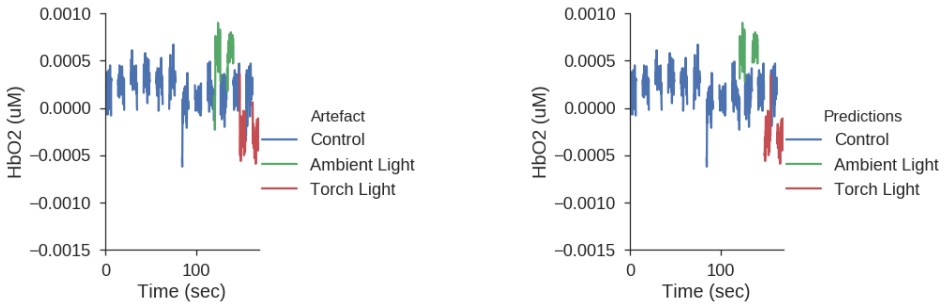

(a) Actual class for each artefact.

(b) Predicted class for each artefact.

Figure 3.8: Trace of the HbO2 signal showing just control and light artefacts. Figure 3.8a shows the actual artefacts whilst figure 3.8b shows the predicted when using spectra from sensor 7. A high level of accuracy is clearly visible when distinguishing between control spectra and light artefact spectra.

Chapter 4

Mathematical Modelling

Mathematical modelling has been a key component of science and research for centuries, ranging from astrophysical models of the stars through to more modern systems biology models of cell-level phenomena. In this chapter we start by providing a brief overview of mathematical modelling and its uses before moving on to look at various model of cerebral physiology and dynamics. Finally we look at the history of modelling within the UCL group.

4.1 Mathematical Modelling

When trying to understand data and physical phenomena, it is often useful to develop a model that can be used to explain behaviour and make predictions. As has been mentioned previously, George Box is oft-quoted as saying that “All models are wrong, but some are useful". He first referred to this concept in Box [1976] outlining the following

> **2.3 Parsimony** Since all models are wrong the scientist cannot obtain a "correct" one by excessive elaboration. On the contrary following William of Occam he should seek an economical description of natural phenomena. Just as the ability to devise simple but evocative models is the signature of the great scientist so overelaboration and overparameterization is often the mark of mediocrity.
>
> **2.4 Worrying Selectively** Since all models are wrong the scientist must be alert to what is importantly wrong. It is inappropriate to be con-

cerned about mice when there are tigers abroad.

This principle must inform all work we do with regards to mathematical modelling. The purpose of the models is not to simply replicate as exactly as possible the phenomena or system in question, but it is to do so in a manner that provides insight and knowledge. In fact, if you want a model that is accurate but uninformative you are best to use the original system.

A useful model should use appropriate simplifications that reproduce the behaviour to a good degree of accuracy whilst also developing our knowledge of the original system. There are a variety of ways in which this can be achieved, two of which are outlined below.

4.1.1 Mechanistic vs Statistical Models

There are two distinct modelling approaches that need to be considered in this work, statistical (or empirical) and mechanistic, focussing on their use in analysing time-series data. When considering these approaches for an open system, where the system has external interactions, we could potentially also term these 'black box' and 'white box' models, with black box models accounting only for the behaviour of the stimulus and response and the inner workings being 'hidden'.

These two approaches have a reasonable amount of overlap but Kendall et al. [1999] highlight four distinct differences *a*) the goals of the analysis, *b*) the way in which the time series is treated, *c*) the kinds of models that are developed and *d*) the methods used for judging the models' explanatory power. In particular, the goals of the analysis provides the simplest description for difference in the approaches. A statistical model seeks to 'describe the data and extrapolate it into the future', whilst a mechanistic model seeks 'typically to understand the causes of a generic phenomenon' [Kendall et al., 1999].

The BrainSignals models considered within this work are mechanistic models. They take a 'bottom-up', systems biology approach towards understanding the mechanisms and behaviours of the brain's physiology. It is through the emergent behaviour of this system that we attempt to understand how the brain works. This is explored in further detail below.

Whilst there are many way to classify models into different groups, and inevitably there will always be 'shades of grey' between those groups, the focus in this section is to look at what we term *systems biology models* and *black box models*. Before beginning it is important to clarify what we mean by both of these terms and also how they potentially coincide with the above description of mechanistic and statistical models.

4.2 Brain Physiology Models

Attempting to model the physiology of the brain and its resulting complex behaviour is a challenge that is both interesting and incredibly difficult. Many different approaches have been taken, focussing on different phenomena, functions or physiology. Attempting to comprehensively review and discuss all of these is beyond the scope of this work but a number of models are listed in Appendix A covering circulation, metabolism or both. Some models deal with whole brain physiology but where that is not the case their work can be considered important or relevant to the development of such a model. The *Focus* column attempts to group the model or paper into one or more BrainSignals sub-compartments and as such it may involve generalising a model that is more specific than this column may imply.

A number of models relevant to this work are explored in further detail below. By detailing these models and some of the results that they have led to it should become clear the benefit of approaching some of these problem areas via mathematical modelling. The models show how it is possible to identify extra information from measured data, as well as how models can be used to propose and examine hypotheses around expected and observed behaviour.

4.2.1 Ursino-Lodi Models

The first Ursino-Lodi model, Ursino-Lodi (1997) was published in 1997 [Ursino and Lodi, 1997] and was based on simplifications of an earlier model published in 1988 [Ursino, 1988a, Ursino et al., 1998] combined with cerebral autoregulation (AR) mechanisms. These AR mechanisms were simulated via time constant and

a sigmoidal static characteristic. The model was use to simulate interactions between intra-cranial pressure (ICP), cerebral blood flow (CBF) and AR. The model was capable of reproducing clinical phenomena concerning ICP changes but could not be used for studying ICP pulsations that were synchronous with cardiac beat or respiration.

The Ursino-Lodi(1997) model was used as the basis for a an extended model published in Ursino and Lodi [1998], the Ursino-Lodi model (1998). This model is of particular note as it served as an integral part of the BRAINCIRC model [Banaji et al., 2008], which is outlined in further detail in Section 4.3. The Ursino-Lodi (1998) model extended both Ursino and Lodi [1997] and Ursino et al. [1998] to produce a combined physiological model of hydrodynamics and CO_2 reactivity in the brain. It investigates the relationship CBF, cerebral blood volume (CBV), ICP and the regulatory mechanisms of CO_2 reactivity and AR. Key aspects of this model over previous versions were the aforementioned inclusion of CO_2 reactivity , its nonlinear interaction with ICP and cerebral AR and the description of the transcranial doppler velocity signal, which can be linked to measurement of V_{mca}. An electrical analogue of the intracranial dynamics, taken from Ursino and Lodi [1998], is shown in Figure 4.1. Simulation results from the model were encouraging and supported data on CBF and CBV reported in the literature concerning both the separate effects of CO_2 and AR, as well as their nonlinear interaction. The model was used for further study in Ursino et al. [1998].

This model was further simplified in Ursino et al. [2000], mainly through not distinguishing between large and small pial arteries and the use of a windkessel model to reproduce the biomechanics of the arterial-arteriolar vasculature. This was validated against ICP and middle cerebral artery velocity (Vmca) data collected from 13 patients, with the model able to reproduce time patterns seen in the monitoring systems using parameter values within range reported in the literature. The Ursino-Lodi (1998) model was also used in Ursino and Giulioni [2003] to investigate the relationship between Vmca pulsatility and cerebral AR, coming to a number of conclusions. One notable finding was that they identified a linear relationship between

Vmca and cranial perfusion pressure (CPP) for a wide range of CPP values and that the slope of this linear relationship may provide information about the AR status.

Removed due to copyright

Figure 4.1: Electrical analogue of the intracranial dynamics found in the Ursino-Lodi (1998) model. Electric analog of intracranial dynamics. G_1 and C_1, hydraulic conductance and compliance, respectively, of proximal cerebral arteries; G_2 and C_2, hydraulic conductance and compliance, respectively, of distal cerebral arteries; P_a, systemic arterial pressure (SAP); P_1 and P_2, intravascular pressure of large pial arteries and medium and small arteries, respectively; q, CBF; P_c and P_v, capillary and cerebral venous pressure, respectively; P_{vs} and P_{cv}, sinus venous and central venous pressure, respectively; P_{ic}, ICP; C_{ic}, intracranial compliance; G_{pv} and C_{vi}, hydraulic conductance and compliance of large cerebral veins; G_{vs}, hydraulic conductance of terminal intracranial veins (bridge veins and lateral lacunae or lakes); G_{ve} and C_{ve}, hydraulic conductance and compliance, respectively, of extracranial venous pathways; G_f and G_o, conductances to cerebrospinal fluid (CSF) formation and CSF outflow; q_f and q_o, rates of CSF formation and CSF outflow; and I_i, artificial CSF injection rate. *Image and caption taken from Ursino and Lodi [1998]*

4.2.2 Aubert-Costalat and Cloutier Models

The Aubert (2001) model [Aubert et al., 2001] attempted to group various aspects of brain functional imaging, such as functional magnetic resonance imaging (fMRI), MRS, electroencephelography, magnetoencephelography, within a coherent framework. It implemented a system of differential equations modelling: 1) sodium membrane transport, 2) Na/K ATPase, 3) neuronal energy metabolism (glycolysis, buffering effect of phosphocreatine (PCr), mitochondrial respiration) 4) blood-brain barrier exchange, and 5) brain haemodynamics. They assumed that the correlation between brain activation and metabolism could be due to either changes in ATP and ADP following activation of Na/K ATPase that result from changes in ion concentration, or the involvement of a second messenger, such as calcium, in different phases of metabolism. Following comparison with measured data, it is this second hypothesis that they posit to be the case, however the nature of this second messenger was not identified.

The Aubert (2001) model was then used as the basis for the Aubert-Costalat (2002) model [Aubert and Costalat, 2002]. This model developed the Aubert (2001) model further and included: 1) electrophysiology via the inclusion of membrane sodium

currents, 2) energy metabolism (ATP regeneration via PCr buffer effect, glycolysis, aerobic metabolism), 3) glucose, O_2 and lactate blood-brain barrier exchanges, 4) haemodynamics aspects including the impact of CBF on these exhanges, and 5) linking model results to BOLD data by including venous dilation processes that occur during stimulation (via the use of the Balloon model [Buxton et al., 1998]). They tested four hypotheses on the behaviour of the cerebral metabolic rate of O_2 ($CMRO_2$), especially during sustained activation:

H1) $CMRO_2$ remains at its baseline level during stimulation.

H2) Assume that $CMRO_2$ increases according to a trapezoidal function

H3) $CMRO_2$ depends on intracellular O_2 and pyruvate concentrations and is regulated by the ATP/ADP ratio.

H4) In addition to hypothesis H3, $CMRO_2$ progressively increases due to the presence of a second messenger.

Aubert and Costalat were able to obtain good agreement between model simulations and experimental data under hypothesis H3 and H4 as opposed to H1 and H2. In addition, by looking at the effects of varying a number of physiologically significant parameters on the time course of the simulated BOLD signal, they were able to use the model to formulate hypotheses about the physiological and/or biochemical significance of features in the fMRI data, especially the poststimulus undershoot and the baseline drift.

Aubert and Costalat [2005] introduces a compartmentalised model (Aubert-Costalat (2005) model) of energy metabolism in neurons and astrocytes. Neuronal and astrocytic metabolisms are described based on the Aubert-Costalat (2002) model. They use the model to investigate the Astrocyte-Neuron lactate shuttle (ANLS) hypothesis of astrocyte provided lactate being an energy fuel for neurons, adopting hypotheses highly unfavourable to ANLS. Simulation results always displayed ANLS for at least some time period and could be divided into two groups. At lower neuronal stimulation, ANLS can occur during a prolonged activation as well as during the

poststimulus period, whilst at higher neuronal stimulation ANLS occurs preferably at the start of stimulation and in the poststimulus period. They determined that experimental data available at the time on lactate kinetics was compatible with the ANLS hypothesis.

Cloutier et al. [2009] takes a systems approach towards modelling brain energy metabolism and is based on the Aubert-Costalat (2005) model but with the inclusion of glutamate cycling. The co-ordination of neuronal and astrocytic responses to stimuli is controlled by glutamate cycling between the two cells, activating the sodium pump in astrocytes to maintain the sodium gradient.

4.2.3 Orlowski Models

The Orlowski models are based on the Cloutier (2009) model. The first of these is the Orlowski (2011) model [Orlowski et al., 2011]. It is composed of four compartments: astrocytes, neurons, capillary vessels, and extracellular volume. One of its key aspects is the inclusion of pH dynamics. The model structure is shown in Figure 4.2 with modifications from the Cloutier (2009) model shown in red. The

Removed due to copyright

Figure 4.2: Diagram representing the structure of the Orlowski (2011) model. Modifications in red indicate the changes made to the Cloutier (2009) model to allow the computation of pH dynamics in brain cells.
Image taken from Orlowski et al. [2011].

authors assume a linear relationship between the concentration of H^+ ions and the concentration of ATP, PCr and the production of lactate. the authors focus on modelling pH dynamics following ischaemic stroke. The model is in good agreement with previously published data for pH values under total ischaemia. The results of the model for 80% ischaemia are shown in Figure 4.3. Of note is the reduction in pH for both astrocytes and neurons. This reduction is stated by the authors to be a linear function of the decrease in CBF. The Orlowski model was developed further in Orlowski et al. [2013] to include extracellular metabolite concentrations. This was done to investigate tissue damage caused by oedema following ischaemic stroke. Results from this model are shown in Figure 4.4 under two experimental

Removed due to copyright

Figure 4.3: Model simulation results from a CBF reduction of 80%. Variations of pH, intracellular LAC, CO2t, intracellular energy stores and sodium concentration as a function of time after a CBF reduction of 80% of initial value. Where not mentioned plots are given for neurons.
Image and caption taken from Orlowski et al. [2011].

conditions: with sodium diffusion and with no sodium diffusion. A cubic element of side 1 cm was used to represent tissue under ischaemia to compute the flux owing to diffusion and the space surrounding the cube was assumed to have a constant concentration of sodium ions irrespective of the flux through the cube. Figure 4.4a shows extracellular space fraction under each condition, Figure 4.4b shows intracellular space under each condition and finally Figure 4.4c shows the neuronal sodium ion concentration under each condition. The authors note that with diffusion extracellular volume decreases less than without diffusion whilst there is also an increase in sodium and chloride into the cell. should these additional ions come from capillaries rather than the extracellualr space then there would be an increase in brain volume. The model was then used to observe the diffusion process at the tissue level. This was done for a cubic mesh of 1000 cubic cells for 1000 cm^3 volume.

Removed due to copyright

Figure 4.4: Model simulation results from a CBF reduction of 89% using the Orlowski (2013) model. (a) Size evolution of the extracellular and (b) cellular space; (c) and of neuronal sodium concentration after a reduction of blood flow by 90%. Two cases are taken into account: with or without sodium diffusion-enabled.
Image and caption taken from Orlowski et al. [2013].

4.3 UCL Models

The BCMD framework was developed to run a number of different mathematical models of brain haemodynamics and metabolism. These models are often extremely complex and do require a reasonable amount of domain knowledge to understand. It is recommended that if you are not familiar with these models that you read their relevant papers in order to better understand them. Figure 4.5 provides an outline

of the history of the models and the general relationships between each.

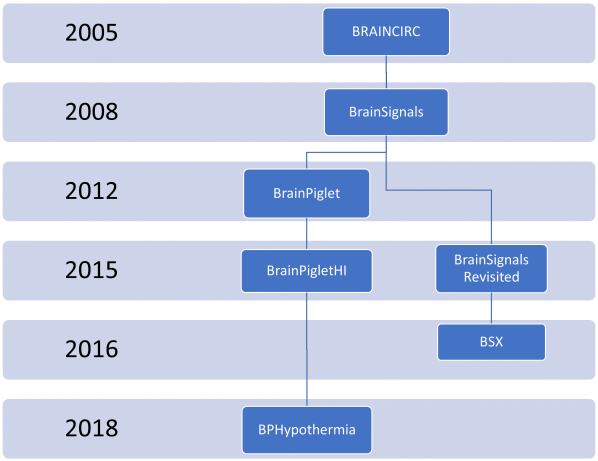

Figure 4.5: Hierarchy of existing BrainSignals models.

4.3.1 BrainCirc

The **BrainCirc** model was first introduced in Banaji [2004] and consists of three 'sites': blood vessels, brain tissue, and vascular smooth muscle (VSM). The blood vessels model is based on an earlier model by Ursino and Lodi [1998]. It contains many more parameters than its successors and focussed on modelling the biophysical and biochemical processes in detail, as opposed to the interpretation of clinically-measurable signals. It was run using the now retired BRAINCIRC interface.

4.3.2 BrainSignals

The **BrainSignals** model [Banaji et al., 2008] is a simplification of the earlier BrainCirc model, adding the ability to model metabolism. All of the BrainSignals derived models retain the same general structure, with many of the inputs and outputs remaining the same, with small variations to allow for model specific additions such as scalp blood flow. A simplification of this structure is shown in figure 4.6. There are four constituent submodels - blood flow, oxygen transport from blood to tis-

sue, oxidative metabolism within the tissue and measurement - with a number of state variables passing information between them. In Banaji et al. [2008] a range of model simulations were presented and simulated data were compared to published data obtained from both *in vivo* and *in vitro* settings. It was shown that the model was able to reproduce observed behaviour in response to stimuli and the authors outline their hope that the model could be used to understands NIRS signals, particularly the oxCCO signal.

The BrainSignals model was also used in Jelfs et al. [2012], which noted that the interpretation of physiological data are a highly non-trivial precondition to using non-invasive measurement methods in a clinical situation. The authors explored the use of BrainSignals to explain and predict physiological signals. Five signals were all non-invasively-measured during hypoxemia in healthy volunteers: the difference in haemoglobin ($\Delta HbD = \Delta HbO_2 - \Delta HHb$), the total haemoglobin ($\Delta HbT = \Delta HbO_2 + \Delta HHb$), tissue oxygen saturation (TOS) and the change in cytochrome *c* oxidase (ΔoxCCO) measured using near-infrared spectroscopy and the middle cerebral artery blood flow measured using transcranial Doppler flowmetry. They found that optimising the model using partial data improved its predictive power, but some discrepancies between model and data persisted despite model optimisation. It was suggested that this could either be used to flag up important questions concerning the underlying physiology, and the reliability and physiological meaning of the signals, or that the model is missing certain physiological mechanisms which would need to be included.

A number of derivative models have been built on top of BrainSignals, such as the simplified BSRV [Caldwell et al., 2015a] which reduced the number of variables and parameters within the model by making various simplifications. This was then developed further to include a scalp submodel and used to investigate the potentially confounding effect of systemic physiological factors on NIRS measurements [Caldwell et al., 2016].

BrainSignals Revisited (BSRV) [Caldwell et al., 2015a] is a simplified version of BrainSignals. The original model was refactored before considering a number of

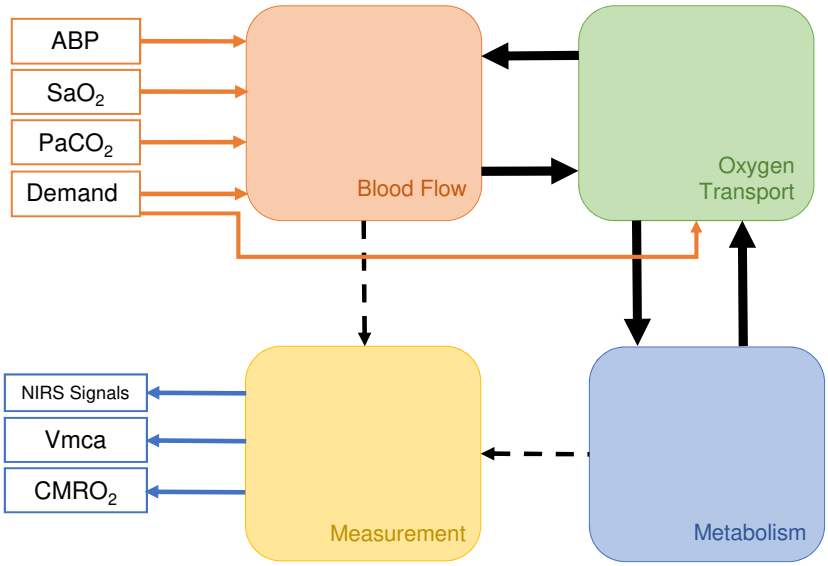

Figure 4.6: Simplified representation of the BrainSignals structure. The model consists of four subcompartments - Blood Flow, Oxygen Transport, Metabolism and Measurement - with Measurement having no input back into the other compartments due to it being solely for reproducing measurement data. Model inputs are shown in red and model outputs are shown in blue.

linearity assumptions in some relationships to reduce the number of parameters and the overall model complexity. This would serve to reduce the number of parameters that need to be optimised as well as providing an increase in the run speed. A number of model variants were developed that were then tested against simulated and experimental data from healthy adults undergoing a hypercapnia challenge. To develop the model variants, the haemodynamics and metabolic submodels were substituted with multiple simplified alternatives. The two best performing variants of each submodel were then combined and these combinations were compared to the full BrainSignals model. It was found that the best performing simplified models were able to reproduce the behaviour of the full BrainSignals model. The authors suggest that the ability to obtain similar behaviour from a simplified variant points to superfluous complexity within the full model. They do note, however, that the context in which the model is being used also dictates what features are superfluous, with some behaviours perhaps only important in a small subset of all use cases. Additionally, the simplification process inevitably leads to the loss of some infor-

mation through the omission of numerous details of internal processes. However, the details were removed in part due to the difficulty in interpreting them. The relevance of the simplified model to different physiological and pathophysiological conditions will require specific testing.

The simplified BrainSignals model was later extended to model extracerebral contamination in a version called **BrainSignals eXtended** (BSX) [Caldwell et al., 2016]. Contamination of NIRS signals can come from blood flow in extracerebral tissue layers due to systemic changes such as blood pressure and oxygen saturation. The inclusion of a scalp submodel within the BrainSignals model allowed the effects of extracerebral blood flow on the haemoglobin NIRS signals to be investigated, as per figure 4.7. The results showed that confounding effects from systemic physiological factors can produce misleading haemodynamic responses in both positive and negative directions, leading to both false positive and false negative signal outputs. This suggests that it is important to record potential confounders in the course of fNIRS experiments. The model may then be able to attribute the observed behaviour to the correct cause.

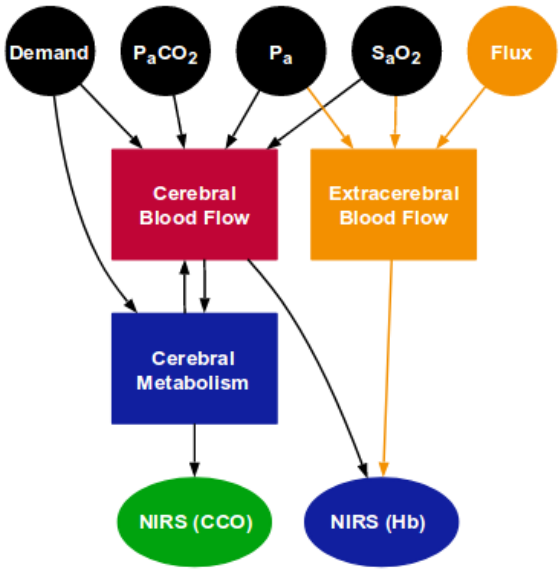

Figure 4.7: A simplified representation of the BrainSignals eXtended structure, as per Caldwell et al. [2016].

4.3.3 BrainPiglet

The BrainSignals model was also adapted for use in interpreting data from studies using a piglet model, a common surrogate for the neonatal brain. The models based around this data are often called **BrainPiglet** models [Moroz et al., 2012a]. Amongst other things, they have been used to investigate the effects of hypoxic-ischemia on the brain [Caldwell et al., 2015b], including functionality to simulate cell death due to oxygen deprivation.

The initial BrainPiglet model was an extension of the BrainSignals models that was used to predict NIRS measurements of the human adult brain.The original model was altered to apply to the anaesthetised piglet brain. It includes metabolites measured by ^{31}P-MRS, namely PCr, inorganic phosphate, and ATP. In Moroz et al. [2012a], the model simulations are compared to measurements from piglets with anoxia. The NIRS and MRS measurements were predicted well, although this required a reduction in blood pressure AR. In addition, the cerebral rate of oxygen consumption and the lactate concentrations were predicted using the model, but were not measured experimentally due to the difficulty in doing so. The model was then used to investigate hypotheses regarding changes in CCO redox state during anoxia.

The BrainPiglet model was extended in Hapuarachchi et al. [2013] to investigate changes in neonatal intracellular brain pH during hypoxia-ischemia. As in the previous version, the model can simulate NIRS and MRS measurements obtained from experiments in piglets.^{31}P-MRS data were used to estimate intracellular pH and to compare simulated pH and oxygenation with measured values. Hapuarachchi et al. [2014] investigate changes in cerebral haemodynamics during and after cerebral hypoxia-ischaemia (HI) in 15 piglets, using the BrainPiglet model alongside NIRS and ^{31}P-MRS data. The model was used to simulate HI using the measurements, using sensitivity analysis to identify important parameters that were then optimised. They successfully simulated changes in cellular metabolism including shifts in intracellular pH observed in the piglet brain during HI. The model fitted the data reasonably, suggesting a 20% drop in glucose consumption, 65% increase in lactate

concentration and 35% drop in cerebral metabolic rate of oxygen during HI.

BrainPigletHI is the most recent modification to the BrainPiglet model, developed to look specifically at the effects of hypoxia-ischemia on the neonatal brain, using measurements from the neonatal piglet brain. Figure 4.8 shows the general structure of the model, with changes to the original BrainPiglet model shown in red. There are a number of key modifications: *a*) the addition of a compartment to model carotid occlusion, a technique used experimentally to model HI, *b*) the modelling of cytoplasmic pH, which affects the metabolic processes during HI, and is an important contributor to a measurable signal and *c*) a representation of cell death, which is a possible explanation for the non-recovery observed in some experimental data. The model was used to simulate data from two piglets: one showing recovery following injury and the other not, with signals not returning to their initial baseline following the injury. Initially the model was optimised using data up to the nadir of the injury in both piglets. Without the inclusion of cell death, only the recovering piglet is modelled reasonably well, with the simulation returning to its initial baseline for both piglets. With the inclusion of cell death, the model shows a much better simulation of the non-recovering piglet, with the cell death parameter fitted to the post-injury data.

4.4 Conclusion

In this chapter I have outlined a brief history of mathematical modelling of the brain, focussing in particular on those models that could be deemed to fall within the category of systems biology models. Systems biology refers to a mathematical approach within biology to understand the holistic and emergent behaviour of biological systems by bringing together the multitude of individual processes that in combination produce the larger macro behaviour of the system. A number of these models, particularly the Ursino-Lodi models, served as the basis for the BRAIN-CIRC and BrainSignals models developed at UCL. These models and the work they have been used in were also discussed and described. By introducing these

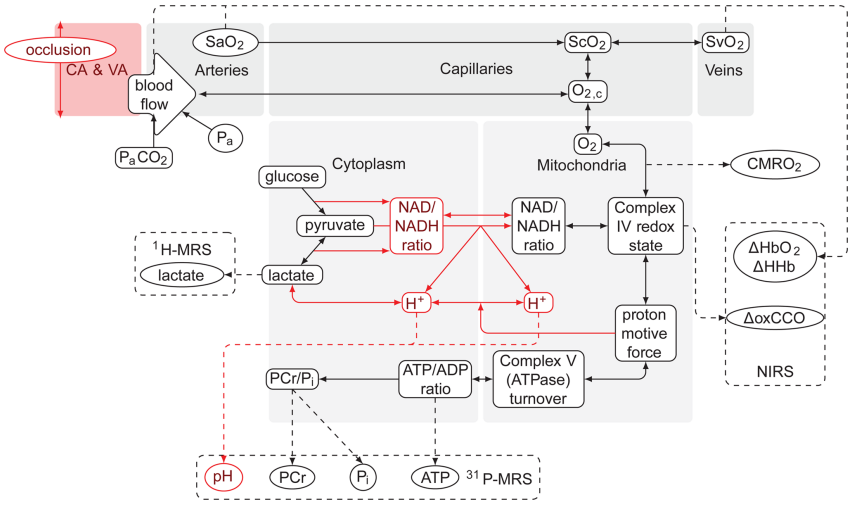

Figure 4.8: Structure of the BrainPigletHI model with major changes shown in red. Key modifications include the ability to model occlusion of the carotid artery, model cytoplasmic pH changes and a representation of cell death. Source: Caldwell et al. [2015b]

models here we serve to establish the foundation upon which future research within this thesis will be built. We highlight the strengths and weaknesses of these models, allowing the research presented herein to be compared and evaluated in context.

Chapter 5

Model Analysis

5.1 Analysis Techniques

Prior to the work undertaken in this PhD to include Bayesian analysis techniques, there were a number of model analysis method already built into BCMD. These can broadly be split into *sensitivity analysis* techniques and *model optimisation* techniques. The uses and specific implementation of each of these is explained below.

5.1.1 Sensitivity Analysis

Sensitivity analysis is principally used in the BCMD framework as a form of model reduction. With the BrainSignals models being inherently very complex, with an incredibly high number of parameters, variables and equations (see Table 5.1), it is important that in order to successfully optimise the model to a given dataset or situation that we use a reduced model that is reliable and reflective of the phenomena or dataset at hand. If the model is not reduced, optimisation methods are likely to be either unable to find an optimal solution or to take an excessively long time.

Frøysa et al. [2018] talks about model reduction in biochemical network dynamics, where the high number of ODEs in the system makes it difficult to analyse the dynamics, much as is the case with the BrainSignals models. They state that model reduction is important in such a situation in order to "identify the main components governing the dynamics of the system" and that the "reduced model should be simpler to analyse, but retain the dynamical behaviour of the original model." They discuss three potential methods - lumping of parameters, time-scale separation to

Model	BRAIN CIRC	Brain Signals	Brain Piglet	Brain Signals Revisited	Brain PigletHI
Year	2005	2008	2013	2015	2015
Focus	Adult	Adult	Piglet	Adult	Piglet
Reactions	81	5	17	5	17
Differential Equations	5	9	21	9	21
Algebraic Relations	72	3	3	3	3
Variables	168	40	128	22	128
Parameters	697	139	226	73	227

Table 5.1: Table highlighting the complexity of the BrainSignals family of models, detailing the number of reactions, differential equations, algebraic relations, variables and parameters.

separate slow and fast dynamics on to separate time scales and finally sensitivity analysis which they describe as being where “the parameters with the least effect on the system output are neglected”.

Over the course of its development, the BrainSignals models have had all of these approaches applied to them in an attempt to reduce their complexity. The previously mentioned BrainSignals Revisited model is a perfect example of lumping multiple parameters into one equivalent parameter, and the model contains time-scale separation in multiple places, specifically when considering autoregulation where different physiological inputs are allowed to operate over different time scales. Finally, sensitivity analysis is always performed before optimising the model to a new phenomena or dataset. This allows us to focus the optimisation only on those parameters that have the most impact on the output of the model.

In the BCMD framework there are multiple sensitivity analysis methods available to the user. Before exploring these in detail there is an important note that should be made. All of these methods use a single summary measure in order to determine their final score. This summary measure is typically a distance or error metric between the “true", measured data and the model output. However, as will be made clear in Chapters 7 to 9, this is not always the best measure to use. Where we want

to identify a specific behaviour the measure needs to be changed to accurately reflect this behaviour e.g. if we were to only want to reproduce a gradient using the model, the best summary measure may be one that looks at the difference in gradient between the measured and simulated data.
BCMD contains the following sensitivity analysis methods by default:

1. `single`: Each selected parameter is varied across its range individually, while all other selected parameters are kept at their default value. Scales $\mathbf{O}(np)$, for p parameters and n levels.

2. `pairwise`: Each pair of selected parameters is varied jointly across their ranges while all other selected parameters are kept at their default. Scales $\mathbf{O}(n^2p^2)$, for p parameters and n levels.

3. `cartesian`: All selected parameters are varied jointly. That is, every possible combination of the available levels of all parameters is tried. Scales $\mathbf{O}(n^p)$, for p parameters and n levels.

4. `morris`: All parameters are varied by single steps in random sequence. This is an implementation of the Morris 'elementary effects' sampling scheme for sensitivity analysis [Morris, 1991, Campolongo et al., 2007]. A single 'trajectory' (of $p+1$ steps) samples sequential changes to all parameters. Scales $\mathbf{O}(kp)$, for p parameters and k trajectories.

5. `fast`: Parameters are varied jointly along oscillatory trajectories with disjoint frequencies in each dimension. This is used for FAST sensitivity analysis [Saltelli et al., 1999]. Scales $\mathbf{O}(kp)$, for p parameters and k samples.

Method descriptions are taken from Caldwell [2015]. In general, the Morris method is the one most likely to be used due to its suitability for models with a large number of parameters. It is this method that is described here and that was used in the research undertaken for this PhD.

Morris Method

The Morris method is one of the most popular sensitivity analysis methods for models with a large number of parameters. It varies one factor (parameter) at a time and is therefore known as an OAT method. It aims to determine parameters with effects that are 1. negligible, 2. linear and additive, and 3. non-linear or involved in interactions with other parameters. Two sensitivity measures were proposed by Morris for each parameter: μ, which gives an estimate of the overall effect of the parameter on the output, and σ, which estimates the ensemble of the second- and higher-order effects in which the parameter is involved, including interaction effects. Furthermore, work by Campolongo et al. [2007] introduced a third , μ_* which "is the estimate of the mean of the distribution of the absolute values of the elementary effects" and avoids the problem of opposite sign effects in a non-monotonic model cancelling each other out. This measure is the one used in this research as it is more effective at identifying parameter importance [Saltelli et al., 2004, Morris, 1991].

The method suggested by Morris is based on what is called an *elementary effect*. For a k-dimensional vector $\mathbf{X}$ of parameters with components X_i, each of which can assume integer values in the set $\{0, 1/(p-1), 2/(p-1), ..., 1\}$. The parameter space is then a k-dimensional, p-level grid. Let Δ be a predetermined multiple of $1/(p-1)$ and so for a given value $\mathbf{x}$ of $\mathbf{X}$ the elementary effect of the ith input factor is

$$d_i(\mathbf{x}) = \frac{[y(x_1, ..., x_{i-1}, x_i + \Delta, x_{i+1}, ..., x_k) - y(\mathbf{x})]}{\Delta} \tag{5.1}$$

where $\mathbf{x} = (x_1, x_2, ..., x_k)$ is any selected value in the parameter space such that $(\mathbf{x} + \mathbf{e}_i\Delta)$, where $\mathbf{e}_i$ is a vector of zeros but with a unit for its ith component, is still inside the parameter space for all $i = 1, ..., k$ [Saltelli et al., 2004]. The function $y(\mathbf{x})$ maps to a single value and so here effectively represents the aforementioned summary metric or measure that we apply to the timeseries data produced by the BrainSignals models.

A distribution of elementary effects F_i is calculated for each ith input factor by randomly sampled different $\mathbf{x}$ from the parameter space. The number of elements

for each F_i is $p^{k-1}[p-\Delta(p-1)]$. This is then used to determine each of the three sensitivity measure. The mean μ of the distribution F_i, the mean of the absolutes μ_* and the standard deviation of F_i is calculated for each parameter. Parameters with a greater influence have a larger μ_* whilst a greater standard deviation σ means the parameter's influence on the output is non-linear and depends on its interaction with other parameters.

5.1.2 Parameter Optimisation

Given that a large portion of the research undertaken for this work was towards developing a better method of performing parameter optimisation using Bayesian techniques, it is important to briefly outline the previous methods used. A more in depth comparison of the two approaches can be found in Chapter 7. As with the sensitivity analysis, there are a number of methods 'pre-baked' into the BCMD framework, namely a particle swarm optimisation method, a genetic algorithm, and a two-array differential evolution algorithm. These are implemented using the `OpenOpt` and `PSwarm` Python libraries [Caldwell, 2015].

PSwarm

The particle swarm or "PSwarm" method is a global optimisation method for variables with specific upper and lower bounds [Vaz and Vicente, 2007, 2009]. The algorithm combines two steps at each iteration: a pattern search (poll) step and a particle swarm (search) step. In the poll step it applies a directional direct search (coordinate search in the pure simple bounds case) and in the search step a particle swarm is used to generate points in the feasible region. By generating a population of points in the search step, the algorithm can then poll around the best particle, improving the robustness of the algorithm [Vaz and Vicente, 2009]. This method is the one that has generally been used previously for optimising the BrainSignals models [Hapuarachchi et al., 2014, Hapuarachchi, 2015, Moroz, 2014, Caldwell et al., 2015b].

5.2 Bayesian Analysis

One of the key developments of this work over previous iterations of BrainSignals will be to develop a Bayesian analysis of the model that will allow model parameters to be described by posterior distributions rather than as point estimates.

5.2.1 Bayesian vs Frequentist

Understanding a statistical analysis of a problem also involves understanding the theories of probability that underlie model estimation. There are two theories regarding probability: the *frequentist paradigm* and the *Bayesian paradigm* [Van de Schoot et al., 2014].

The frequentist paradigm associates probability with a series of outcomes. A typical example would be that of a coin toss. When a frequentist says that the probability of a fair coin toss coming up heads is $\frac{1}{2}$, they mean that the over a long run of coin tosses, approximately half the time the result is heads. The Bayesian approach however would be that the a probability of getting heads being $\frac{1}{2}$ is instead describing the degree of belief, a belief based on prior information such as considering the geometry of the coin. This definition of probability is sometimes called *subjective probability* [Glickman and Van Dyk, 2007].

These two differing approaches to probability affect the way in which we approach statistics by changing how we draw conclusions from the data. For example, by estimating a statistic, such as the mean of a variable or a regression coefficient, that we want to use to make inferences about a dataset, we are trying to determine an unknown parameter.

The key difference between the Bayesian and frequentist approach to statistical inference is in the nature of this unknown parameter. For a frequentist, the parameter has a fixed but unknown value. This value is the *true* value of that parameter, such as the 'true' mean or the 'true' regression coefficient. The Bayesian approach, however, treats all unknown parameters as uncertain and therefore describes them by a distribution.

5.2.2 Bayes' Theorem

In order to understand Bayesian inference a number of background ideas must first be established. Throughout this section the following notation will be used. $p(a)$ is used to define the marginal distribution of the variable a and $p(a|b)$ defines the conditional probability distribution *'the probability of* a *given* b'. Both a and b are dummy variables in this context.

Bayes' theorem defines that a joint probability density function($p(\theta, y)$) can be written as the product of the *prior distribution* ($p(\theta)$) and the *sampling distribution* ($p(y|\theta)$), where θ is a parameter and y is observed data. Bayes' theorem can therefore be used to determine the *posterior density function* where the factor $p(y)$ can be omitted when considered as a constant in the case of fixed y.

$$p(\theta|y) \propto p(\theta)p(y|\theta) \tag{5.2}$$

The posterior density function provides information about the model parameter, θ, given that we have observed the data y, thus allowing the prior distribution to be updated based on that data.

It is also possible to make inferences about an unseen observable. Before any data have been considered, the distribution of the unknown observable y is

$$p(y) = \int p(y, \theta)d\theta = \int p(y|\theta)p(\theta)d\theta \tag{5.3}$$

which is known as the marginal distribution of y or the *prior predictive distribution*. After the data y have been observed we can predict an unknown observable $\tilde{y}$, and the distribution for this quantity is known as the *posterior predictive distribution*:

$$p(\tilde{y}|y) = \int p(\tilde{y}|\theta)p(\theta|y)d\theta \tag{5.4}$$

If we regard this as a function θ for fixed y, it is known as the *likelihood function*. Whilst it is possible to derive an analytical form for the likelihood function in some cases, where the problem is particularly complex another approach is required. In-

stead, it is often possible to sample from the posterior distribution for a parameter θ directly. A method of sampling directly without the use of an explicit likelihood function and without approximation is described by Beaumont [2010], as per Rubin [1984]:

1. Draw $\theta_i \sim \pi(\theta)$.
2. Simulate $x_i \sim p(x|\theta_i)$.
3. Reject θ_i if $x_i \neq y$

where x is the data, θ are parameter values and y is an observation. The accepted samples of θ are sampled directly from $p(\theta|y)$, thus allowing us to obtain a posterior distribution for θ, given y, without having first obtained an analytically complete form for the likelihood. This process of *Approximate Bayesian Computation* is explored further in section 5.3.

5.2.3 Model Checking

Upon constructing a probability model and computing a posterior distribution it is important to then 'check' this model, by assessing its fit to the data and to additional knowledge. In Gelman et al. [2014, chap. 6], Gelman et al. state that rather than asking whether a model is 'true or false', it is more important to ask whether the model's deficiencies have a noticeable effect on inference. This is because the probability models in most cases will not be perfectly true. Failures in the model can lead to problems by creating false inferences about the estimands of interest. Gelman et al. [1996] describe three ways in which a model can be checked: (1) examining sensitivity of inferences to reasonable changes in the prior distribution and the likelihood; (2) checking that the posterior inferences are reasonable, given the substantive context of the model; and (3) checking that the model fits the data. Firstly, the sensitivity of the inferences can be examined by undertaking sensitivity analysis, determining the impact of other reasonable probability models. The sensitivity analysis can be incorporated into the prior-to-posterior analysis by setting up a joint distribution that any data that may be observed are a plausible outcome of

the distribution. However, setting up an exhaustive super-model that considers all possible outcomes is computationally infeasible for all but the simplest of models.

For the second check, we can determine if the model inferences make sense by utilising additional information not formally included in either the prior distribution or the likelihood. For any applied problem it is extremely likely that such information exists.

Finally, checking the fit of the data can be done by evaluating the model through *external validation*, in which the model is used to make predictions about future data that are then collected and compared against. However, it is often the case that the model needs to be checked before new data can be obtained.

Posterior predictive checking is a method of approximating external validation using data already obtained. This involves drawing simulated values from the joint posterior predictive distribution,

$$p(y^{rep}|y) = \int p(y^{rep}|\theta)p(\theta|y)d\theta \tag{5.5}$$

where y is observed data, θ is the vector of parameters and y^{rep} is simulated (replicated) data, and comparing these samples to the observed data. It should be noted that the use of y^{rep} here in contrast to $\tilde{y}$ in equation 5.4 is intentional, with $\tilde{y}$ defining a future value whilst y^{rep} is specifically a replication like y. In the case of the model using explanatory variables, x, they will be the same for y and y^{rep} but $\tilde{y}$ may instead have its own explanatory variables $\tilde{x}$. In the case of the dynamical systems under consideration in this research, posterior predictive checking involves sampling parameters directly from the posterior distributions, running the model with these parameters and the systemic information from the actual data (y) and then comparing the simulated time series data (y^{rep}) to the true values either graphically or via a more formal method of measuring the 'statistical significance' of the lack of fit [Gelman et al., 2014, chap. 6].

The discrepancy between the model and the data can be measured via the use of *test quantities*. Gelman et al. [2014, chap. 6] define a test quantity, or *discrepancy measure*, $T(y,\theta)$ as a 'scalar summary of parameters and data that are used as a

standard when comparing data to predictive simulations'. A test quantity plays the same role in Bayesian model checking that a test statistic, $T(y)$, plays in classical testing, where the test statistic is dependent only on the data.
Given the complex nature of the BrainSignals models and the field in which they are applied (neurophysiology), the possibility of finding a 'true' model is close to zero and so instead a 'good enough' model is sought.

5.3 Approximate Bayesian Computation

Approximate Bayesian Computation (ABC) is defined by Sunnåker et al. [2013] as 'a class of computational methods rooted in Bayesian statistics'. They are used where the analytic form of the likelihood function is intractable. The first example of ABC-related methods is the rejection algorithm posited by Rubin in 1984 [Rubin, 1984], with Sunnåker et al. [2013] noting that the first ABC algorithm for posterior inference was proposed by Tavare et al. in 1997 [Tavare et al., 1997]. This paper looked at methods to estimate the time to the most recent common ancestor for a sample of intraspecies DNA sequences. They note that for most cases in which inferences are drawn about coalescence times from sequencing data, the theory and computational implementation for an exact solution are complex, and instead they proposed an approximate simulation method to obtain the posterior distribution of coalescence times. A sample from the posterior was obtained by accepting/rejecting proposals based on the number of segregating sites in the simulated and real data. The actual term 'Approximate Bayesian Computation' was first coined by Beaumont et al. in Beaumont et al. [2002], where the application of ABC methods to population genetics was considered.

5.3.1 ABC Methods

There are a number of different ABC methods that can be used, but all share a similar generic form as detailed by Toni et al. [2009]:

1. Sample a candidate parameter vector θ^* from the proposal distribution $p(\theta)$.

2. Simulate a dataset y^{rep} from the model described by a conditional probability distribution $p(y|\theta^*)$.

3. Compare the simulated dataset, y^{rep}, to the experimental dataset, y, using a distance function, d, and tolerance, ε. If $d(y, y^{rep}) \leq \varepsilon$, accept θ^*. The tolerance $\varepsilon \geq 0$ is the desired level of agreement between y and y^{rep}.

The output of the ABC algorithm used will be a sample from the distribution $p(\theta|d(y, y^{rep}) \leq \varepsilon)$. If ε is sufficiently small, then $p(\theta|d(y, y^{rep}) \leq \varepsilon)$ will be a good approximation for the posterior $p(\theta|y)$.
The simplest algorithm is the *ABC rejection sampler* [Pritchard et al., 1999], outlined by Toni et al. [2009] as:

R.1 Sample θ^* from $p(\theta)$.

R.2 Simulate a dataset y^{rep} from $p(y|\theta^*)$.

R.3 If $d(y, y^{rep}) \leq \varepsilon$, accept θ^*.

R.4 Return to R.1

The main disadvantage of this approach is that if the prior distribution is very different from the posterior, the acceptance rate of the algorithm can be very low. This can massively increase the time taken to produce a reasonable posterior. To avoid this problem, a Markov chain Monte Carlo approach was introduced [Marjoram et al., 2003]. The *ABC MCM algorithm*, again as outlined by Toni et al. [2009], is:

M.1 Initialise $\theta_i, i = 0$.

M.2 Propose θ^* according to a proposal distribution $q(\theta|\theta_i)$.

M.3 Simulate a dataset y^{rep} from $p(y|\theta^*)$.

M.4 If $d(y, y^{rep}) \leq \varepsilon$, go to M.5, otherwise set $\theta_{i+1} = \theta_i$ and go to M.6.

M.5 Set $\theta_{i+1} = \theta^*$ with probability

$$\alpha = \min\left(1, \frac{p(\theta^*)q(\theta_i|\theta^*)}{p(\theta_i)q(\theta^*|\theta_i)}\right)$$

and $\theta_{i+1} = \theta_i$ with probability $1-\alpha$.

M.6 Set $i = i+1$ and go to M.2

The ABC MCMC algorithm produces a Markov chain with the stationary distribution $p(\theta|d(y,y^{rep}) \leq \varepsilon)$, which as discussed previously is a good approximation of the posterior for sufficiently small ε. It has the potential disadvantages, however, of becoming stuck in regions of low probability for long periods of time and of producing very long chains. To avoid these disadvantages, sequential Monte Carlo (SMC) based methods can instead be used.

Once such method is described in Toni et al. [2009]. In ABC SMC, a number of sampled parameter values, known as *particles*, $\{\theta^{(1)}, \ldots, \theta^{(N)}\}$ are sampled from the prior distribution, $p(\theta)$, and propagated through a sequence of intermediate distributions, $p(\theta|d(y,y^{rep}) \leq \varepsilon_i), i = 1, \ldots, T-1$, until it represents a sample from the target distribution $p(\theta|d(y,y^{rep}) \leq \varepsilon_T)$. The tolerances $\varepsilon_1, \ldots, \varepsilon_T$ are chosen such that $\varepsilon_1 > \ldots \geq \varepsilon_T \geq 0$, thus gradually evolving the distributions towards the target posterior. The large number of particles also, in principle, helps avoid the problem found in ABC MCMC of chains becoming stuck in regions of low probability. The algorithm for this process is:

S.1 Initialise $\varepsilon_1, \ldots, \varepsilon_T$. Set the population indicator $t = 0$.

S.2 Set the population indicator $i = 0$.

S.3 If t=0. sample θ^{**} from $p(\theta)$.
Else, sample θ^* from the previous population $\{\theta_{t-1}^{(i)}\}$ with weights w_{t-1} and perturb the particle to obtain $\theta^{**} \sim K_t(\theta|\theta^*)$, where K_t is a perturbation kernel.
If $p(\theta^{**}) = 0$, return to start of S.3.

S.4 Simulate a candidate dataset $y^{rep} \sim p(y|\theta^{**})$. If $d(y, y^{rep}) \geq \varepsilon_t$, return to S.3.

S.5 Set $\theta_t^{(i)} = \theta^{**}$ and calculate the weight of particle $\theta_t^{(i)}$,

$$w_t^{(i)} = \begin{cases} 1, & \text{if } t = 0, \\ \frac{p(\theta_t^{(i)})}{\sum_N^{j=1} w_{t-1}^{(j)} K_t(\theta_{t-1}^{(j)}, \theta_t^{(j)})}, & \text{if } t > 0. \end{cases}$$

If $i < N$, set $i = i + 1$, go to S.3.

S.6 Normalise the weights.

If $t < T$, set $t = t + 1$ and go to S.2.

Particles that are sampled from the previous distribution are denoted by a single asterisk, whilst after perturbation these particles are denoted by a double asterisk. For the special case that $T = 1$, the ABC SMC algorithm corresponds to the ABC rejection algorithm.

5.4 Conclusion

In this chapter I have outlined some of the existing methods for analysing systems biology models, particularly those that have been used to analyse the BrainSignals models. I started by considering sensitivity analysis and its role in both analysing models and producing a reduced model, something that is necessary given the complexity of the models under analysis. I then considered the currently used method for parameter optimisation, particularly looking at the particle swarm method that has been used to analyse the BrainSignals models previously. I then provided an overview of Bayesian method of analysis, identifying how they may prove more useful for analysing our models and what their strengths are compared to the current maximum-likelihood approach. As part of this I outlined approximate Bayesian computation and how it can be useful where an analytical form of the likelihood function cannot be obtained, something that is the case for the BrainSignals models and which will be explored in more detail in Chapter 7.

Chapter 6

Modelling Framework

The following chapter has been adapted from a paper published via Wellcome Open Research. It can be found at doi:10.12688/wellcomeopenres.12201.1.

6.1 Introduction

Understanding data collected when measuring a biological system can be a complex process. One approach to aid understanding is the use of mathematical models. These can be used alongside the measured data to try and understand the underlying processes that contribute to the observed dynamics. Whilst a lot of time is spent in developing these models, the accessibility and usability of them is often forgotten. UCL's Biomedical Optics Research Laboratory has previously developed two software interfaces for defining and solving complex physiological models, BrainCirc [Banaji and Moroz, Banaji, 2004] and BCMD[Caldwell et al., 2015a]. While the latter provides simplified modelling and improved stability and portability, it remains challenging to install and use, with some facilities such as parameter optimisation only easily available on Linux. In addition, the requirement to create detailed input specification files to define simulation parameters constitutes a significant barrier to use by non-technical specialists.

In this section I will outline the previous framework, before outlining the WeBCMD platform that was created to facilitate ease of access to the models, discussing the ideas and motivation behind this new interface.

6.2 Previous Framework

Before discussing how the software has been modified, it is important to provide a general outline of how the original software worked. The BCMD framework is a system for implementing differential-algebraic equations (DAE). It uses Python 2.7 to generate and run C code that encapsulates the equations and is generated using a template. The C code is solved using the RADAU5 solver [Hairer et al., 1996] written in Fortran. One of the main jobs of the C template code is to mediate between the model equations and the RADAU5 solver via number of wrappers, also written in C. The solver remains untouched in the current version.

6.2.1 Defining and running a model

BCMD could be accessed through two main methods: a command line interface and a basic GUI which has been constructed using TKinter. A model is written as a `*.modeldef` file, which allows the model to be written in something approaching a human readable format. This is then compiled into an executable file that takes a number of arguments. The process of converting the `*.modeldef` file to a C file and then compiling into an executable is handled by a Makefile for the command line interface and is a built-in function of the GUI. Figure 6.1 shows a typical model definition file, this one being for a simple resistor-capacitor circuit. Line 1 is a comment providing simple information about the file and then lines 3 and 4 define model outputs and inputs respectively. Below this, the differential equation for the model is defined on line 8, initial values are defined on lines 11 and 12 and then model parameters are defined on lines 15 and 16. Initial values are set using '`:=`' whilst intermediate variables or expressions are set using '`=`'. It is also possible to define algebraic equations and chemical reactions in the model definition. For more information, please see the BCMD manual [Caldwell, 2015].

The compiled model program can then be run by passing it an input file, which is a text file of specific formatting that contains information about that model run. This includes how many time steps there are, what output should be produced, non-default parameter values, the time points themselves and also any input values at

```
1  # super-simple model of a series RC circuit
2
3  @output Vc V
4  @input V
5
6  # root variable Vc is the voltage across the capacitor
7  # input V is the voltage across the circuit
8  Vc' = (V - Vc) / (R * C)
9
10 # initial values
11 V := 0
12 Vc := 0
13
14 # parameters
15 R := 10
16 C := 0.1
```

Figure 6.1: Typical modeldef file for a BCMD model. This one defines a simple resistor-capacitor (RC) circuit.

those time points. For example, a typical input file may resemble that in figure 6.2, which is an input file for the model defined by the modeldef file in figure 6.1. Line 1 is commented out and used to provide user information, line 2 defines how many time steps the model will run for, line 3 states the value of the parameter `V` will be set at each time steps and all following lines define each of these. We can see that the first time step runs from 0 to 5 seconds and each step after this occurs every 5 seconds (as denoted by the '`+ 5`' at the start of each line). As with the model definition file, for more information, please see the BCMD documentation [Caldwell, 2015].

6.2.2 Batch Processes

It is often the case that rather than simply running the model once for one set of parameters, model behaviour as a whole will need to be explored and in order to do this there needs to be the option to run the model many times with a slightly varied configuration each time. Writing out each input file manually with each new

```
# simple voltage input to our RC circuit model
@5
: 1 V
= 0 5 1
+ 5 0
+ 5 -1
+ 5 0
+ 5 1
```

Figure 6.2: Typical input file to a BCMD model. This one works with the RC model defined in 6.1

configuration would be a fool's errand and so an automated method of doing this was created.

There are two main batch processes that the BCMD software is able to undertake, as well as at least one unmaintained process. Within each process there are various levels of functionality, but not all of these have been fully maintained since their original creation and there may be unseen bugs resulting from later changes to the software. The first is the ability to run 'deterministic or quasi-deterministic batches of simulations', primarily so as to undertake sensitivity analysis of a model [Caldwell, 2015] and the second is to run model optimisation, finding the parameter combination that best fits the model to the data. All batch processes share a number of common arguments and features, such as defining the distribution to draw parameter values from or the distance measure to use when comparing to true data.

Sensitivity Analysis There are a number of different sensitivity analysis modes that can be used, with the parameters being varied drawn from a set of predefined levels. For more information on this see Section 5.1.

Optimisation As with sensitivity analysis, a number of different methods for optimising model parameters are provided, all of which are implemented in the OpenOpt Python library. There are three different kinds of optimisation types,

- *GLP*: A global optimisation, which is generally the type of optimisation that will be required.

- *NLP*: A non-linear local optimisation.
- *NSP*: A non-smooth local optimisation.

and three different solver methods,

- *pswarm*: Particle swarm optimisation method, using the PSwarm library [Vaz and Vicente, 2007].
- *galileo*: Genetic algorithm based method, built into OpenOpt.
- *de*: A two-array differential evolution algorithm, built into OpenOpt.

A distance can be declared which is then minimised to optimise the parameters. Some basic post-processing options are also available, such as zeroing the output, such that the simulation data are of the same form as the experimental data. Reasons for doing this include the fact that experimental NIRS data are often a 'relative' signal rather than absolute, whilst simulated data are given as an absolute value.
Any information about the optimisation process is limited to that available in Caldwell [2015], as the OpenOpt documentation is no longer available and the software does not seem to be maintained any longer. This makes extending the existing optimisation capabilities difficult and is one of the many reasons for updating the software and considering new methods of fitting the models to data.

6.3 Framework Redevelopment

6.3.1 General Modifications

Both the Bayesian framework (BayesCMD), outlined in Chapter 7, and the We-BCMD interface outlined below are based on a single overhaul of the BCMD software, with the Python wrapper for the BCMD models rebuilt. This had two purposes:

a) The new framework is built in Python 3. This was necessary to ensure the model remains open for use and expansion past 2020 after Python 2 is retired.

b) The new structure is better suited to the Bayesian and web frameworks that were developed over the course of this PhD.

This transition required some small changes to the files responsible for parsing modeldef files, such as modifying syntax slightly to allow for the change in how the Python `print` statement works. The C template has been modified slightly to allow for a new command line argument, `-I`, which allows the input file to be passed in directly as a string rather than having to be written to file first.

The new Python wrapper has a number of new classes that are used in conjunction with the existing C templates and command line interface to run and configure models, as well as to run the Bayesian analysis. Models are still defined using modeldef files, but the input files can be written by passing information into Python. It is this change that allows the models to be configured and run through a web interface.

The two key classes for running the models are the `InputCreator` class, which allows input files to be created through Python, and the `ModelBCMD` class which takes in model information, in the form of Python data structures, to then configure and run a model. Both of these are found in the `bcmdModel` subpackage. `InputCreator` is passed the information it requires by `ModelBCMD` before using it to create an input file with the same form as that used in BCMD. This allows the creation of input files to be done separately to running a model, something which may prove useful if they need to be accessed separately to running a model. By default, the input file is written to buffer before being passed directly to the executable as a string. This avoids unnecessary writing to disk. The model output is a Python dictionary, allowing it to be passed easily to the web interface of WeBCMD and making any postprocessing easier to handle. The option of writing this output to file is available but is turned off by default. A simplified diagram of this process can be seen in figure 6.3a).

Figure 6.3b) shows a simplified description of the batch process required to perform a simple rejection-based ABC process. The `Batch` class takes in input data, model configuration data and the prior distributions for each parameter to be analysed. This is parsed into a form that can be handled by `ModelBCMD` before passing it through. A single model run is undertaken and the distance from the true data

calculated. The model run data are stored, as is the parameter combination and distances for that run. Every N runs, the model runs are written to a file, whilst parameters and distances are appended to a single file. The final result is a number of files containing model run outputs and a single `parameters.txt` file containing all parameter combinations and distances.

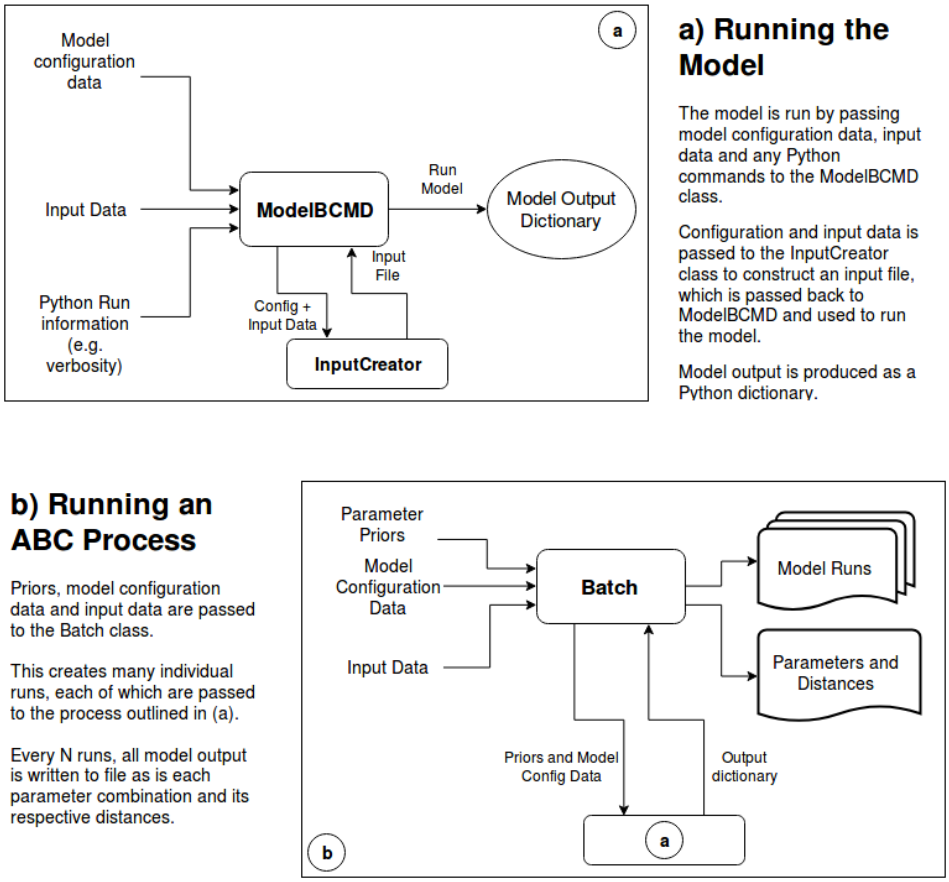

Figure 6.3: Rough outline of how the different components of BayesCMD. Flow chart a) shows how a single model run is implemented using the `ModelBCMD` and `InputCreator` classes. Flow chart b) shows how the `Batch` class uses the process in a) to perform a simple Rejection-based ABC process.

6.3.2 WeBCMD

As has previously been mentioned, installation of the software is difficult and time consuming, especially where the software is only being used for initial exploratory work. Additionally, the learning curve required to use the software led to non-specialists struggling to use the software without a lot of guidance. Because of this

the framework redesign focussed on two main components: making installation and dependency packaging simpler, and improving the interface and usability of the framework.

The initial stages of the framework redesign focussed on making the BCMD software easily accessible across all operating systems. The software includes a number of external dependencies that were either difficult to satisfy or entirely unavailable on Windows and Mac computers. A relatively simple solution to this was to wrap the entire process inside a Docker container, with all necessary installation steps being handled by the Dockerfile when building the image. This however removed access to the original graphical interface. As part of redesigning the BCMD framework the decision was made to improve the user interface and accessibility.

Providing native cross-platform GUI support within a Docker application is impractical. Instead, a web-based interface was developed called WeBCMD, using a Python web server to mediate the mathematical modelling, which allowed the Docker-packaged framework to be accessed via any modern web browser. As well as supporting cross-platform use locally, it was also possible to set up the framework on a remote server to make it accessible as a “cloud service”, which is run remotely and managed by the research team. This allows for non-specialists to simply log on to a website and use the software without having to download anything.

WeBCMD is built primarily around the idea of a "representational state transfer application programming interface" or REST API [Massé, 2012]. A request is sent from the client to the server, which handles the request before sending back a response, as seen in fig. 6.4.

REST systems of varying complexity are ubiquitous across the internet. The form used in WeBCMD is quite basic, serving simply to separate the interface from the mathematical modelling. Doing so makes the interface easily extensible through the addition of new API routes while ensuring that the command line interface remains usable.

The client side component of the interface is written entirely in HTML, CSS and javascript and is compatible with any modern web browser. Frameworks such

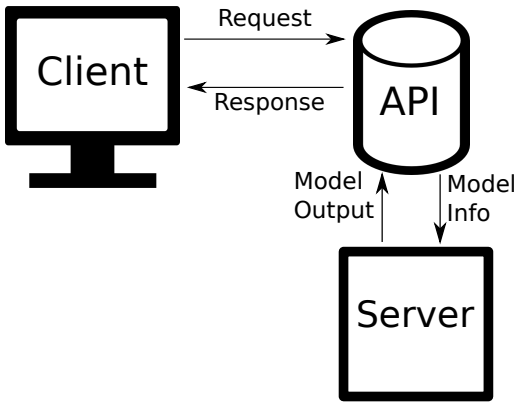

Figure 6.4: A simplified representation of a web API.

as Bootstrap and AngularJS, which provide a modern and user friendly design, are easily integrated. The server side component uses the Flask micro webframework, written in Python. It uses API routes to interface with the BCMD modelling software, processing JSON requests sent from the browser, then returning model information and outputs.

Both client-side and server-side components are distributed and installed using Docker, with the files required to build and run the container separate to the WeBCMD code itself. Because the container will run a Linux-based OS, pulling directly to this rather than to the host computer first can avoid problems with different line endings on Windows and Unix-based computers.

Operation

WeBCMD has been designed to be as simple to operate as possible with specific operation instructions found in Appendix B.

6.3.3 Use Cases

There are two distinct use cases for the WeBCMD software at present:

1. Running a model

2. Comparing steady state simulations of autoregulation against a default BrainSignals run.

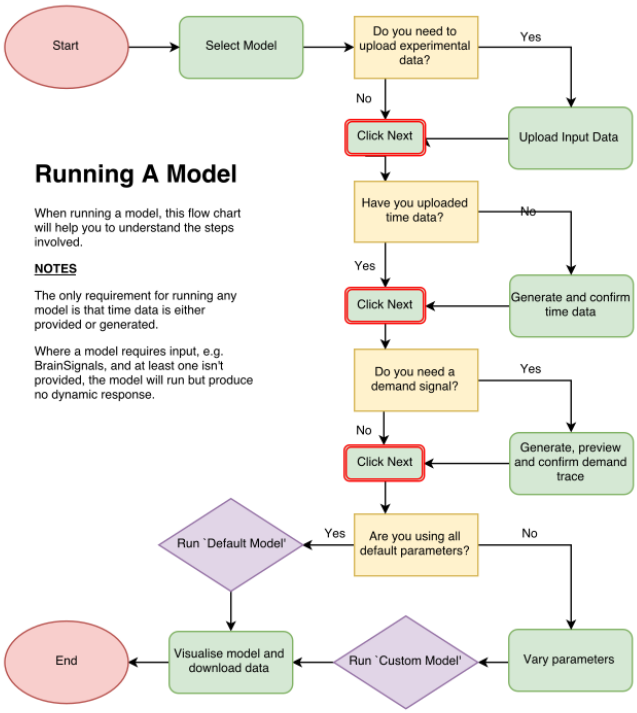

Figure 6.5: Model run process

Running a Model

There are 7 main steps to running a model in WeBCMD, once it has been compiled and its data uploaded. Figure 6.5 shows the general process of running a model in the WeBCMD interface. We will outline this in three specific use cases:

1. A Lotka-Volterra model that requires no input data, run with the default parameters and model inputs and outputs.

2. Driving a BrainSignals model with systemic data, with non-default parameter values.

3. Simulating a number of functional activation tests in BrainSignals, using the demand function generator to create this demand data.

Lotka-Volterra

The Lotka-Volterra model outlined here is available as part of the WeBCMD distribution. Solving models such as this is not the intended use of WeBCMD however, and there are likely more efficient methods of running and investigating such

a model.

The Lotka-Volterra model is a commonly used model for Predator-Prey dynamics in ecology [Freedman, 1980]. The version shown here consists of two equations

$$\frac{\mathrm{d}x}{\mathrm{d}t} = ax - bxy \tag{6.1}$$

$$\frac{\mathrm{d}y}{\mathrm{d}t} = cxy - dy \tag{6.2}$$

where x represents the number of prey in the system, y represents the number of predators in the system, t is time and a, b, c, d are positive, real parameters representing the interactions between the two species.

This model will be simulated using a 'Default Run', where all parameters are set to their default values and any changes made to them by the user will not be included in the simulation.

The seven main steps to running this model are:

1. Select 'Lotka-Volterra' in the model select screen.

2. No data are uploaded, so simply click next.

3. Generate time data here by choosing start and end times, as well as the desired time step, before clicking 'Generate Time'. Once the time data have been generated 't' will appear in the inputs list at the top of the screen.

4. Demand is BrainSignals specific and so is not required. De-select the checkbox here before clicking next.

5. Because we are using default parameters, simply advance to the next page without altering any values.

6. On the 'Default' tab, choose our time point data, which in this case is `t`. All inputs and outputs are the defaults found in the original model definition. Clicking 'Run Model' will send the information to the backend of the interface. When the model output has been returned it will then be possible to click 'Next'.

7. On this final screen it is possible to visualise and download the returned model data as a comma-separated values (.csv) file. This can then be compared to the `model-output-LV.csv` file.

BrainSignals

As has previously been mentioned, the BrainSignals models are a collection of models that represent the physiological dynamics of the brain. The version referenced here, is the original BrainSignals model, prior to the refactoring by Caldwell et al. [2015a]. For this model, input data are used to drive various inputs, such as blood pressure, partial pressure of CO_2 and arterial oxygen saturation.
There are two different run types demonstrated here. The first uses systemic data to drive the model whilst the second uses the built in 'demand creator' to simulate a number of demand increases.

Systemic Data The seven steps to running this simulation are:

1. Select 'BrainSignals' in the model select screen.

2. Click "Browse..." and select the supplied data, which has a default name of `synthetic_input_data.csv`. Select all of the options as inputs, and confirm. You can leave outputs blank.

3. Click 'Next' on the time creation screen as this information is provided in the input file.

4. As we are not simulating functional activation here, we don't need to create a demand trace so just click 'Next' here.

5. Alter parameters to whatever values are desired - in this case the supplied output data were generated by increasing the parameter `R_autp` to 6.

6. On the 'Custom' tab, select the time data to be our `t` data. Then select the appropriate input data for each input parameter: `Pa` for 'P_a', `PaCO2` for 'Pa_CO2' and `SaO2` for 'SaO2sup'. Leave all output values as their default

settings and then set the 'Burn in' time to '250' seconds. Click 'Run Model' and when the model has finished running click to visualise on the next screen.

7. On this final screen it is possible to visualise and download the returned model data as a comma-separated values (.csv) file. This can then be compared to the `model-output-BS-systemic.csv` file.

Changing Demand The second run type doesn't use any input data and instead generates a synthetic input demand. The demand created is not physiologically accurate, but is instead used as a way of showing the possiblities of this functionality. It is expected that more demand input types will be added over time.

1. Select 'BrainSignals' in the model select screen.

2. Click 'Next' on this screen as no input data are needed.

3. Set the final time to 1000 seconds and click 'Generate Time'. Once the time data have been created, the inputs section at the top will now contain the letter `t`. Click 'Next'.

4. The Demand Creation page will have detected the information from the previous screen and already filled in the Start and End times, as well as the time interval.

   To create a demand peak, click 'Add Another Peak'. For this example, set the start time for the peak to 120, the end time to 180, the peak height to 5 and select a peak type of 'Top Hat'. This will create a single square wave with a magnitude of 5 that lasts 60 seconds, starting at 120 seconds. We will repeat this 3 times by clicking repeat peak, typing in 3 for the number of repeats and setting an interval of 60 seconds between each. You can preview this by clicking 'Generate Demand'.

   We will then add 3 more demand peaks of a different type. Click 'Add Another Peak' and enter a start time of 480, end time of 540, peak hight of 5 and select a peak type of 'Wavelet'. Select 'Repeat peak' and repeat it 3 times

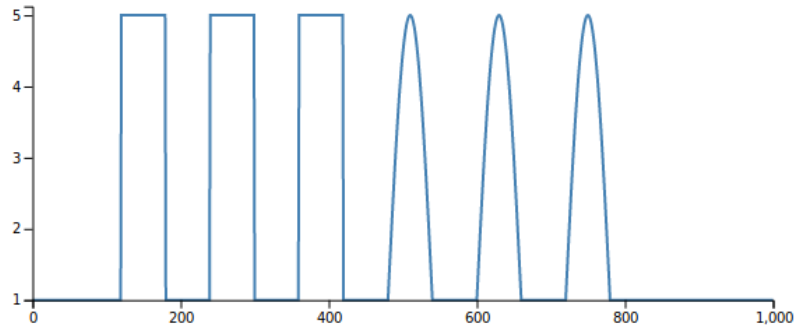

Figure 6.6: Example demand trace

with an interval of 60 seconds. Click 'Generate Demand' and you should get a figure like that in 6.6. Click 'Next'.

5. We will leave all parameters at their default values.

6. On the 'Default' tab, select the time data, `t`, to be our 't' data and `u` to be'u'. Click 'Run Model' and when the model has finished running click to visualise on the next screen.

7. On this final screen it is possible to visualise and download the returned model data as a comma-separated values (.csv) file. This can then be compared to the `model-output-BS-demand.csv` file.

Steady state simulations of autoregulation Steady state simulations of autoregulation are a way of considering the validity of a haemodynamic model like BrainSignals. By altering a single input value, running the model for a period of time and then altering it again, it is possible to compare the autoregulation response of the model to a default dataset created using the 'Default BrainSignals Model'. The input values changed are the arterial blood pressure (P_a), oxygen saturation (SaO_2) and partial pressure of carbon dioxide ($PaCO_2$).

To run a steady state simulation, simply go to the Steady State tab and choose a model. After doing this, you have the option of changing any model parameters, allowing you to visualise their effect on the steady state autoregulation curve. You can then choose whether top run the steady state simulation varying the parameters in one of three ways:

- **up**: Vary from min to max
- **down**: Vary from max to min
- **both**: Vary from min to max to min, allowing for the detection of hysteresis.

Once the model has finished running going to the next page allows you to visualise the autoregulation curve for this model run alongside the Default BrainSignals curve.

Figure 6.7 shows the steady state response curves, when varying from min to max (up), for the BrainSignals model with the `R_auto` parameter, which is responsible for the autoregulatory reactivity to oxygen, set to 6.

Summary

When dealing with mathematical models that are themselves complex, it is import to simplify the process of running and analysing these models as much as is possible. If not, there is a risk that the use of the models by non-technical experts will be limited, preventing the potential for insights from the wider scientific community.

WeBCMD has taken the BCMD framework for brain circulation models and made it significantly easier to install and run. The process of installing and managing dependencies can be handled by the framework developers, leaving the process of running and analysing the models to the user. Additionally, the web interface has made input file creation significantly easier by breaking it into easy to follow steps. The final, and arguably most significant, benefit of the new interface is the cloud based web app. Users who are unable to install the software or simply wish to test the model out before installing are now able to do so by simply accessing a website. By removing the obvious barrier of installation, the framework and its associated models are now accessible to anyone with access to the internet. This also means that the most recent and stable models will be easily accessible.

There are a number of key features we wish to include in WeBCMD going forwards and a number of small improvements that need to be made. The inclusion of sensitivity analysis in the interface would be very useful as this remains an important component of all model analysis techniques Due to the significantly high number of

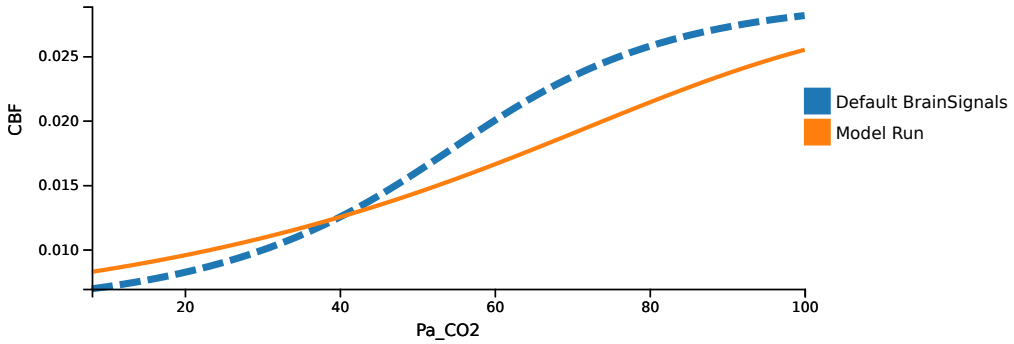

(a) Steady state response for changing partial pressure of CO_2 (Pa_CO2)

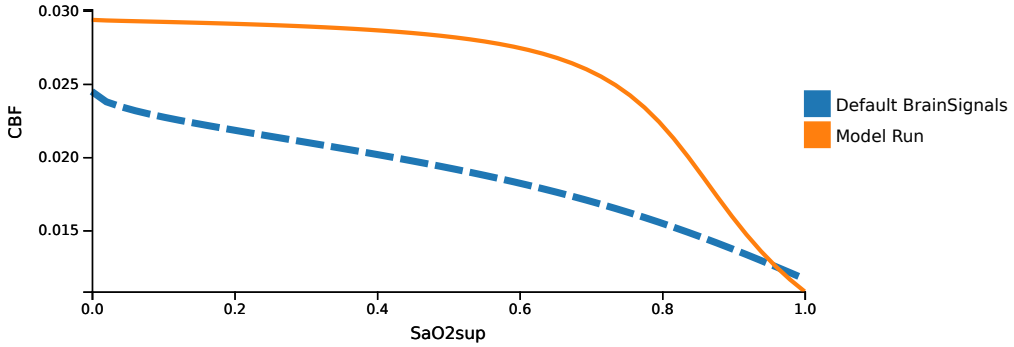

(b) Steady state response for changing arterial O_2 saturation (SaO2sup)

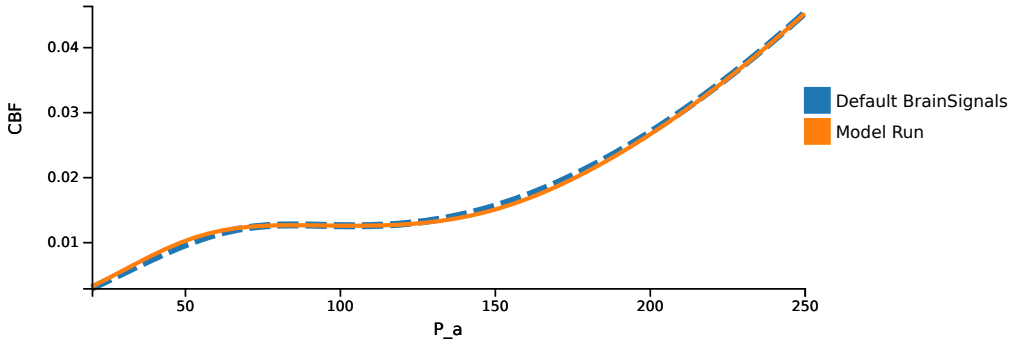

(c) Steady state response for changing blood pressure (P_a)

Figure 6.7: Example steady state output. In blue is the response curve for the default BrainSignals model and in orange is the response curve for this model run's settings (`R_auto`=6).

parameters within the models, identifying what parameters are likely to be the most important for the analysis being undertaken is extremely important. This is considered in greater detail in Model Analysis, when considering existing techniques used in the framework, and Development of a Bayesian framework for model analysis, when looking at the new Bayesian analysis framework. It is hoped that in future work, sensitivity analysis can be added to both the online and offline versions of the interface. However, the main barrier to implementation is the length of time required to perform the analysis. In order to properly implement this, further work would have to be done in terms of preserving the web interface's state whilst the analysis is performed. Without this, closing the web browser may make results inaccessible once complete.

Smaller improvements include the inclusion of units for parameters and variables in the model information section. It is also hoped that some form of model optimisation could be included, but the same problem of long run times, as in sensitivity analysis, would need to be handled.

Chapter 7

Development of a Bayesian framework for model analysis

7.1 Introduction

Systems biology models are used to understand complex biological and physiological systems comprised of large numbers of individual elements that give rise to emergent behaviours. These complex systems are dependent on both the properties of the whole network and on the individual elements [Kitano, 2002]. This inherent complexity within the models can lead to difficulties in determining how best to interpret information obtained through their use.

At University College London, the family of BrainSignals models (and the BRAIN-CIRC model on which they are based) are used to understand the brain's dynamics via a systems biology approach. They bring together a number of mathematical models relating to different aspects of blood circulation, oxygen transport and oxygen metabolism within the brain in order to develop a more complete model that can be used alongside experimental data to simulate physiological phenomena of the brain, such as autoregulation and neural activation. This allows us to understand how our measurements are linked to specific brain physiological and metabolic mechanisms.

All of the models were developed to reproduce broadband near-infrared spectroscopy (NIRS) measurements of brain tissue concentration changes of

haemoglobin (oxygenation and haemodynamics) and cytochrome-*c*-oxidase (mitochondrial metabolism) and vary in their complexity and scope. The first model developed was the 'BRAINCIRC' model in 2005 [Banaji et al., 2005], followed by the 'BrainSignals' model [Banaji et al., 2008] in 2008. A number of additional versions were then developed from this, such as the 'BrainPiglet' model [Moroz et al., 2012a] which was developed to to simulate the physiological and metabolic processes of the piglet brain often used as the neonatal preclinical model. This was extended in BrainPiglet v2.0 to incorporate the effects of cell death during injury [Caldwell et al., 2015b]. In 2015, Caldwell et al. modified and simplified the BrainSignals model to both reduce model complexity and improve model run time, producing the 'BrainSignals Revisited' model [Caldwell et al., 2015a]. All of these models are run using the Brain/Circulation Model Developer environment (BCMD) and are defined in a simple text language. The data collected and analysed with the models primarily consist of broadband NIRS data, providing information about tissue oxygenation, through monitoring of oxy- and deoxy-haemoglobin levels, and cellular metabolism, through the concentration of cytochrome-c-oxidase. This data are then supplemented by systemic information such as blood pressure, arterial oxygen saturation and/or partial pressure of CO_2.

One of the main uses of the models is to fit the model simulations to clinical and experimental data and investigate how model parameters are affected. In the case where data are collected from an injured or sick patient, these changes may illuminate what the underlying causes/mechanisms are behind the illness or injury is.

The models are currently fit using a maximum likelihood based method, with a single value obtained for each parameter. Sensitivity analysis is performed on the models to determine which parameters are most important in influencing each model output for any particular dataset. These parameters are then optimised using the PSwarm method [Vaz and Vicente, 2009] to minimise a given error metric, such as the Euclidean distance, between the modelled and measured signals. Through this each output has a set of optimised parameter values. Parameter values are limited to the same ranges used in the sensitivity analysis [Caldwell et al., 2015b].

This approach has a number of drawbacks. The models are mechanistic and, if fitted to single value parameter estimates, will produce the same output for the same input. Physiology and biology, however, is unlikely to operate in such a constrained manner. Additionally, this set of best-fit parameters for the model may not be representative of the full parameter space [Pullen and Morris, 2014]. In an attempt to try to compensate for this potential drawback, Caldwell et al. [Caldwell et al., 2015b] fit the BrainPiglet model multiple times for two different piglets and found that, whilst parameter values can vary within the same data, separate parameter spaces for each piglet did seem to exist based on the brain physiological status of the piglet following a hypoxic-ischaemic insult.

One of the key ways in which these models are used to extract information from data are through the use of parameter estimation and fitting. However, this step remains a difficult mathematical and computational problem, potentially originating in the lack of identifiability [Chis et al., 2011]. In addition, there has been discussion of 'universal sloppiness' within dynamic systems biology models. Gutenkunst et al. [Gutenkunst et al., 2007] proposed that *sloppiness*, where the parameters of a dynamic model can vary by orders of magnitude without affecting model output, is a universal property of systems biology models. Due to this sloppiness, it may not be possible to make parameter estimations that can be used to make inferences about the system [Brown and Sethna, 2003, Gutenkunst et al., 2007]. Chis et al. have stated however that sloppiness is not equivalent to a lack of identifiability and that a sloppy model can still be identifiable [Chis et al., 2016]. Apgar et al. note that experimental design can be used to constrain a sloppy parameter space by choosing a set of complementary experiments [Apgar et al., 2010].

The use of a Bayesian methodology, by avoiding point estimates, can allow the full uncertainty of the problem to be captured [Pullen and Morris, 2014]. In fact, the use of an Approximate Bayesian Computation (ABC) approach, discussed below, is particularly well suited to these kinds of problems [Liu and Niranjan, 2017]. There are many examples of Bayesian methods being used to analyse bioinformatics data and systems biology models [Wilkinson, 2006], including in sequence analysis [Liu

and Logvinenko, 2003], gene microarray data [Do et al., 2006] and in models of genetic oscillators [Woods et al., 2016] and DNA network dynamics [Woods and Barnes, 2016]. There are a number of models that take a systems biology approach towards understanding physiology, particularly oxygen transport and blood flow, including the previously mentioned BrainSignals [Banaji et al., 2005, 2008, Caldwell et al., 2015b] and BrainPiglet [Moroz et al., 2012a, Caldwell et al., 2015b] models, the Aubert-Costalat model [Aubert and Costalat, 2002], and work by Fantini [Fantini, 2002, 2014, Pierro et al., 2014, Kainerstorfer et al., 2014] and Orlowski and Payne [Orlowski et al., 2013, 2014] where Bayesian parameter estimation could also be applied but has yet to be.

Where a likelihood function can be defined there are a number of these methods that can be used to infer a posterior distribution, $p(\theta|y)$. Although the BrainSignals models are deterministic, the model noise is a combination of process noise and experimental error which is expected to depend on the state in a non-trivial manner. This makes formulating an analytical expression for the likelihood difficult. In this case where a likelihood expression is unobtainable a likelihood-free approach using ABC is required instead. There are a number of different methods available with the simplest being the ABC rejection algorithm (ABC REJ) approach. This has the additional benefit of allowing us to consider different summary statistics that would not be valid in a likelihood-based approach. It may be that these summary statistics allow us to optimise for specific behaviours that have physiological relevance.

The aim of this paper is to introduce the new *BayesCMD* modelling platform that can be used in systems biology models of physiology such as the BrainSignals models, but that can be replicated beyond these. For this work, we have chosen to use ABC REJ as whilst it is less efficient than the other methods mentioned here, the simplicity with which it can be implemented is a significant factor. The models and modelling environment used are already complex and so this initial work focuses on the use of the simplest method as proof of utility. We will demonstrate the effectiveness of this approach by using it to analyse two simulated datasets chosen to represent healthy and impaired brain states, before then using it on experimental

data from a healthy subject undergoing a hypoxia challenge. We will show that the Bayesian approach allows us to extract more information from our data than the previous maximum likelihood approach, with a more complete picture of the parameter space being obtained.

7.2 Materials and methods

Figure 7.1 shows a generalised outline of the final Bayesian analysis process. It can be split into three main sections: sensitivity analysis, Bayesian analysis and model checking. However, before applying the process, data must be generated or collected and an appropriate model chosen.

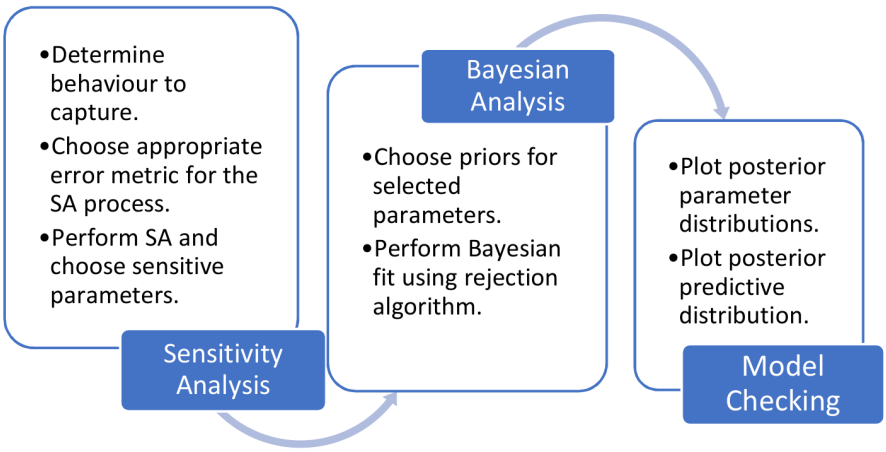

Figure 7.1: Generalised analysis process. A simplified representation of the Bayesian analysis process .

7.2.1 Choice of Model

Whilst a brief overview of the history of the BrainSignals models was given in the introduction, in this section we provide more information about the specifics of the different models. Table 7.1 compares the number of reactions, equations, relations, reactions, variables and parameters in three different models. The BRAINCIRC model from 2005 built on an earlier circulatory model by Ursino and Lodi [Ursino and Lodi, 1998] and combined models for the biophysics of the circulatory system, the brain metabolic biochemistry and the function of vascular smooth muscle. The

BrainSignals model which succeeded it simplified the 'BRAINCIRC' model and added a submodel of mitochondrial metabolism. As previously mentioned, in order to better simulate the physiological and metabolic processes of the piglet brain, which is often used as the neonatal preclinical model, the 'BrainPiglet' model [Moroz et al., 2012a] was developed from the BrainSignals model. It involved modifying the default values for 11 of the 107 parameters used and was extended to include simulated measurements for magnetic resonance spectroscopy values that included brain tissue lactate and ATP production, measurements of which are available in piglet studies. Its extension, BrainPiglet v2, incorporated the effects of cell death during injury in order to to investigate why two piglets showed different recoveries following hypoxia-ischaemia, finding that the differences could be explained by including cell death within the model [Caldwell et al., 2015b].

The 'BrainSignals Revisited' model was produced by making various simplifications to the BrainSignals model by identifying various functions that could be replaced by linear approximators without reducing model applicability. This reduced complexity and decreased the time taken to run a simulation, whilst being able to reproduce the same results and behaviour of the original model. This reduced model of the adult brain was later extended to simulate extracerebral haemodynamics to investigate confounding factors with brain near-infrared spectroscopy measurements, the 'BSX' model [Caldwell et al., 2016].

The models are driven with input signals, such as the blood pressure and/or oxy-

Table 7.1: Comparison of the number of reactions, equations, relations, reactions, variables and parameters in the BRAINCIRC, BrainSignals Revisited and BrainPiglet v2.0 models.

	BRAINCIRC	BrainSignals Revisited	BrainPiglet v2.0
Reactions	81	5	17
Differential Equations	5	9	21
Algebraic Relations	72	3	3
Variables	168	40	128
Parameters	697	139	227

gen saturation, and simulate brain tissue measurments of oxygenation, blood vol-

ume and metabolism, as well as the middle cerebral artery velocity (Vmca) and the cerebral metabolic rate of oxygen ($CMRO_2$). The model can be split into roughly 3 compartments - blood flow, oxygen transport and metabolism - with boundaries chosen to minimise interdependence. Fig 7.2 outlines this in more detail.

All of these models are solved using the BCMD framework and are written in

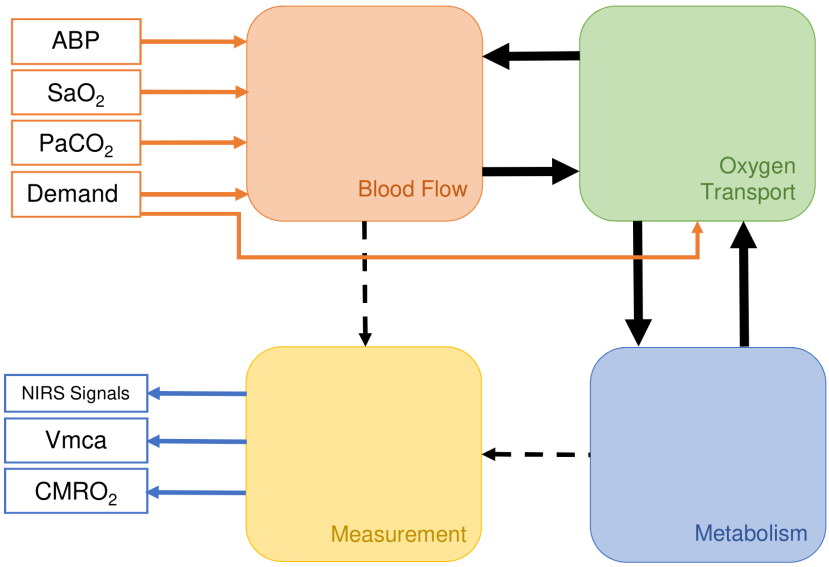

Figure 7.2: Simplified structure of a typical BrainSignals model A typical BrainSignals model can be split into four compartments or submodels. The *blood flow* submodel represents blood flow from arteries to veins via the capillary bed and the *oxygen transport* submodel estimates diffusion of dissolved O_2 from the capillary blood to the brain tissue. Delivered oxygen is then utilised by the *metabolism* submodel. Finally, the *measurement* submodel translates the internal states of the blood flow and metabolism submodels into observable outputs. Model inputs are shown in red and consist of arterial blood pressure (ABP), arterial oxygen saturation (SaO2), partial pressure of CO_2 (PaCO2) and a parameter specifying relative demand, whilst measurable outputs are shown in blue, including NIRS signals as well as middle cerebral artery velocity (Vmca) and cerebral metabolic rate of oxygen ($CMRO_2$).

a simple text format that can be translated to executable C code and solved using the RADAU5 solver [Hairer et al., 1996]. The models take a standard differential-algebraic equation representation, of the form:

$$\mathbf{M}\frac{d\mathbf{y}}{dt} = \mathbf{f}(\mathbf{y}, \boldsymbol{\theta}, t) \tag{7.1}$$

where $\mathbf{y}$ is a vector of variables of interest, $\mathbf{M}$ is a constant, possibly-singular, mass matrix specifying relations among the differential terms, and $\mathbf{f}$ is some vector-valued function, possibly having additional parameters $\boldsymbol{\theta}$. If a row of $\mathbf{M}$ is zero, the corresponding equation in $\mathbf{f}$ is algebraic rather than differential.

In this work we have chosen to use the refactored BrainSignals model [Caldwell et al., 2015a], with a minor modification to include the haemoglobin difference ($\Delta HbO_2 - \Delta HHb = \Delta HbD$) as a model output alongside the normal outputs of oxyhaemoglobin (ΔHbO_2), deoxyhaemoglobin (ΔHHb), total haemoglobin ($\Delta HbO_2 + \Delta HHb = \Delta HbT$), tissue oxygenation index (TOI), and cytochrome-*c*-oxidase (ΔCCO). Both ΔHbD and ΔHbT are included in the experimental dataset due to them being good indicators of brain oxygenation changes and brain blood volume changes respectively, with both being easily measured using broadband NIRS. All NIRS outputs, except TOI, are measured as changes relative to an initial value and therefore both data and model outputs are normalised to an initial value of 0.

7.2.2 Data

Three datasets were used to test the new Bayesian model analysis process. Firstly, 'healthy' data were simulated using the BrainSignals model with the default parameter settings, as per [Banaji et al., 2005, 2008]. Next, the same inputs were used but with the model modified to represent an 'impaired' brain. To do this, a single parameter was changed to reflect a potential pathology or injury, to generate an 'impaired' simulated dataset. Finally, we used experimental data from a healthy adult undergoing a hypoxia challenge, selected due to it being well understood and to minimise complications that could arise from using pathological data.

Simulated Data

Partial pressure of CO_2 ($PaCO_2$) and arterial blood pressure (ABP) were kept at their baseline values of 40 mmHg and 100 mmHg respectively, whilst arterial oxygen saturation (SaO_2) was varied to simulate hypoxia through a decrease in arterial oxygen saturation from 97% to 65%. Initially, all model parameters were kept at

their default values in order to simulate a healthy brain's response to this challenge. Figure 7.3 shows the arterial saturation data and the model response across all considered model outputs.

After simulating the healthy brain response and determining its posterior parameter distribution, the model was altered to include a pathological or impaired brain state. This was done by changing a single parameter to be outside of the healthy parameter space. `r_t`, which affects the shape of the muscular tension relationship, was found to be sensitive in both the sensitivity analysis process (see Simulated Data in the sensitivity analysis results) and the Bayesian analysis. This is clearly seen in its comparatively narrow marginal posterior for the healthy data. Stiffening of blood vessels in the brain has also been noted as a potentially important factor in a number of different pathologies, including Alzheimer's disease [Hughes et al., 2015], and in autoregulation, as seen in Figure 7.4 .

The muscular tension relationship is defined as

$$T_m = T_{max} \exp\left(- \left| \frac{r - r_m}{r_t - r_m} \right|^{n_m} \right), \tag{7.2}$$

where T_m is the muscular tension within the vessel wall and has a bell-shaped dependence on the vessel radius, taking value T_{max} at some optimum radius r_m. r_t and n_m are parameters determining the shape of the curve. Figure 7.4a illustrates the effect of changing r_t on the shape of the curve and shows that decreasing r_t leads to increased muscular tension for the same vessel radius due to a widening of the bell-shaped curve. This can be seen to represent a stiffening of vessels within the brain.

Changing r_t has a significant effect on the brain's ability to autoregulate within the model as seen in figures 7.4b, 7.4c and 7.4d. Figure 7.4b shows that higher blood pressure causes a decrease in cerebral blood flow (CBF) for lower r_t values, as opposed to an increase at the normal value of $r_t = 0.018\,\text{cm}$. Figure 7.4c shows that CBF is lower and decreases quicker for lower r_t values as $PaCO_2$ is decreased and figure 7.4d shows that across all considered oxygen saturations, lower r_t gives a lower CBF.

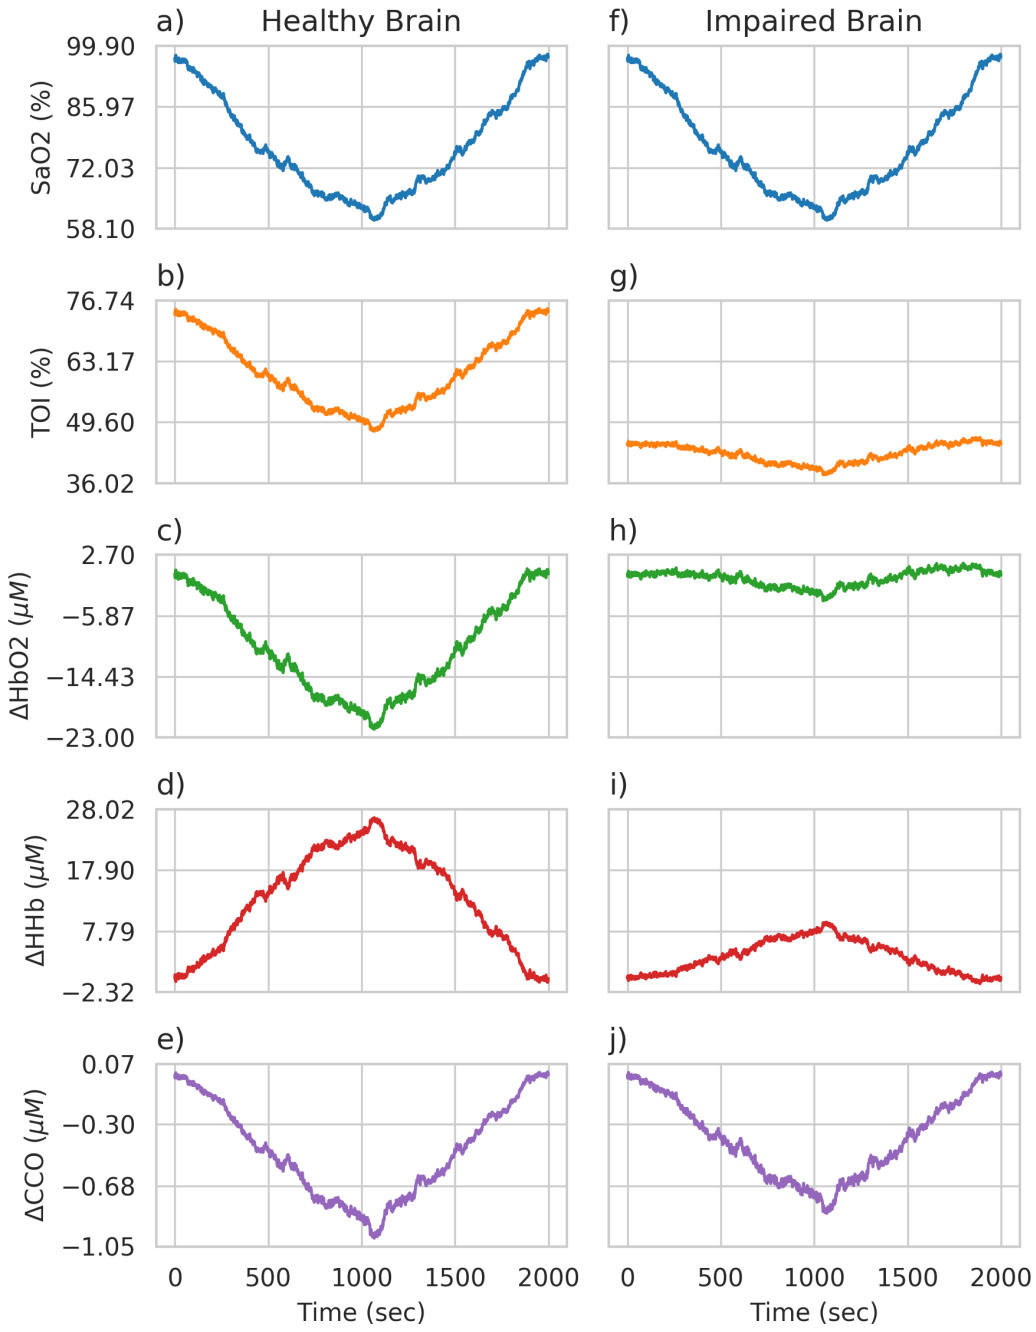

Figure 7.3: Healthy and impaired brain simulations Figures a)-e) show simulations of a healthy brain's response to hypoxia, whilst f)-j) show the impaired brain's response. The input variable of arterial oxygen saturation is shown in blue and is the same for both simulations, whilst the outputs of TOI, ΔHbO_2, ΔHHb and ΔCCO clearly differ between the two brain states.

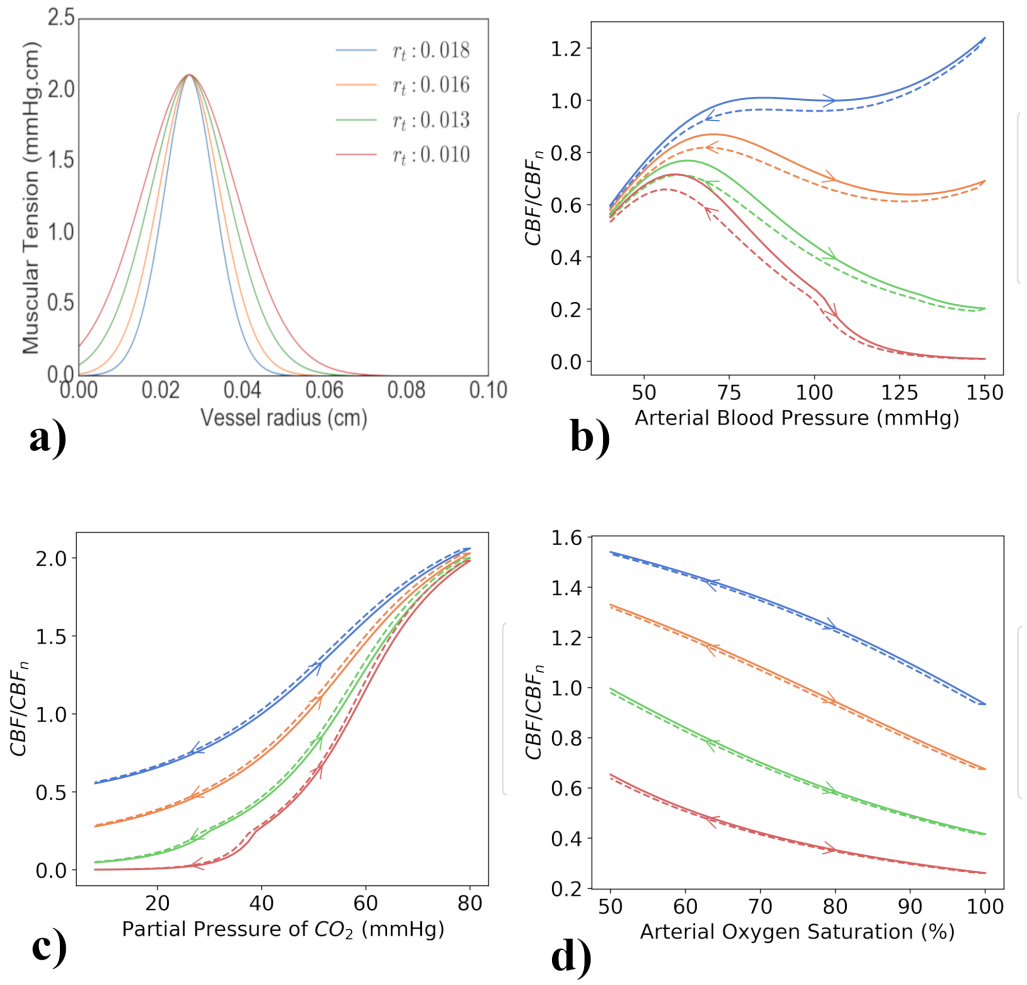

Figure 7.4: Figure 7.4a shows the effect of different r_t values on the shape of the muscular tension curve for a range of vessel radii. It can be seen that reducing r_t widens the curve, leading to increased muscular tension for the same vessel radius. Figures 7.4b, 7.4c and 7.4d show the effect of both increasing and decreasing model inputs on cerebral blood flow for different values of r_t. Cerebral blood flow (CBF) is given as a proportion of the normal CBF ($40\,\mathrm{ml}\,100\mathrm{g}^{-1}\,\mathrm{min}^{-1}$). Changing r_t has a significant effect on the brain's ability to autoregulate within the model. Figure 7.4b shows that higher blood pressures causes a decrease in cerebral blood flow for lower r_t, as opposed to an increase at the normal value of $r_t = 0.018\,\mathrm{cm}$. Figure 7.4c shows that for lower r_t values, CBF decreases quicker as $PaCO_2$ is decreased. Figure 7.4d shows that across all considered oxygen saturations, lower r_t gives a lower CBF.

Figures 7.3f)-j) shows the model response across all considered model outputs for this impaired brain state. The response of the model outputs to the same change

in arterial saturation is much smaller than in the healthy simulation, with the TOI having a lower baseline value of around 45% as compared to around 75%.

Whilst we would expect impairment of a real biological system to stem from multiple parameter changes the intention here was to make the simplest modification possible whilst still representing a potentially real physiological change in order test the method under the simplest conditions. Additionally, it should be noted that a single parameter change will have effects on various physiological variables. As outlined below, we also apply the method to experimental data which are inherently more complex than this simple example and where we expect multiple parameters to differ from baseline.

Experimental Data

Experimental data will inherently contain more uncertainty for parameter fitting than data generated by the model itself. This makes it important to test the Bayesian analysis process on experimental data as well as that simulated from the model. The data used were originally collected by Tisdall et al. [Tisdall et al., 2007] and is shown in figure 7.5. Healthy adult humans had their arterial oxygen saturation reduced from baseline to 80%, whilst minimising changes in end tidal carbon dioxide tension ($EtCO_2$). Data here are from a single subject.

The dataset contains three model inputs: arterial oxygen saturation, end tidal CO_2 and arterial blood pressure, with $EtCO_2$ converted to partial pressure of CO_2 via unit conversion from kPa to mmHg. Blood pressure data were filtered using a low pass 5th order Butterworth filter, with a cut off of 0.05 Hz, to remove noise. The heavily quantised nature of the partial pressure of CO_2 data are not an issue here as the model contains first order filters to smooth input signals over a given time period.

In terms of model outputs, only NIRS signals were used: ΔHbD, ΔHbT, ΔCCO and TOI. All data were resampled to 1 Hz.

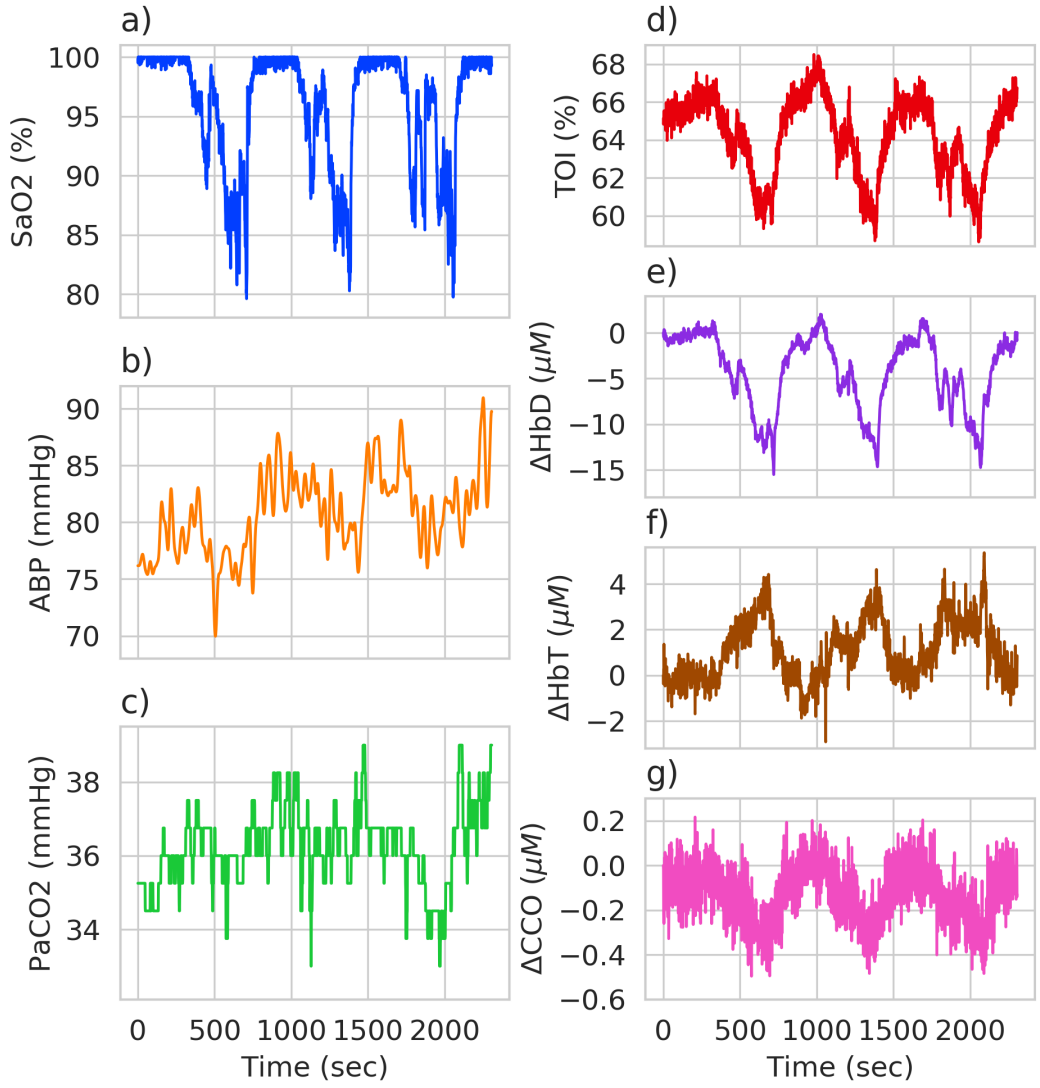

Figure 7.5: Experimental hypoxia data. Data collected from a healthy adult during a hypoxia challenge. Systemic data used as model inputs are shown in figures a), b) and c), with broadband NIRS measurements shown in figures d), e), f) and g).

7.2.3 Sensitivity Analysis

When fitting a model as complex as BrainSignals, it is important to reduce the number of parameters that are required to be fit. We expect that not all parameters will have a significant impact on the model output for a given set of input data. Instead, we can attempt to reduce the number of considered parameters through sensitivity analysis. We used the Morris method [Morris, 1991, Campolongo et al., 2007], which is known to work well with a large number of parameters. The method requires the time series to be reduced to a single number and identifies the parameters that have produce the most variance in this summary value. Previously, we have used the Euclidean distance over the whole time series as our summary value but this has a number of significant drawbacks.

If the summary measure is the distance across the whole time series, we're failing to capture specific changes that we know to be physiologically important. In the case of our hypoxia simulation, for example, we want to select parameters that are important in controlling the overall change from baseline. Taking the Euclidean distance over the time series as a whole however does not prioritise this behaviour. Figure 7.6a shows three sets of data generated from the same toy model function

$$y_i = a\,x\,\sin(x) + b + \varepsilon, \tag{7.3}$$

where a, b are both model parameters and ε is random Gaussian noise.

Assume that without modification, our model produces data y_0, with the default parameters $\Theta_0 : a = 0, b = 0$, and that the behaviour we want to reproduce is sinusoidal but, for some reason, we don't know which parameter is most important in producing this specific behaviour. We decide to undertake sensitivity analysis, using a distance measure of some kind as our summary statistic in order to identify the parameter most important in producing sinusoidal behaviour. If when altering a parameter that distance measure increases, then the behaviour summarised by that distance is sensitive to changes in that parameter. In this case, to produce sinusoidal behaviour, we would want parameter a to be identified as important rather than parameter b.

To generate our data x was varied from 0 to 2π, producing datasets y_1 and y_2 for the parameter sets $\Theta_1 : a = 1, b = 0$, where only a is changed from baseline, and $\Theta_2 : a = 0, b = 0.707$, where only b is changed from baseline, respectively. y_0 and parameter set Θ_0 provide our baseline data. This is seen in figure 7.6a. It is clear from the figure that the two outputs y_1 and y_2 show very different behaviour, the behaviour we want to optimise for is seen in y_1.

Despite both y_1 and y_2 being qualitatively very different they are very similar when

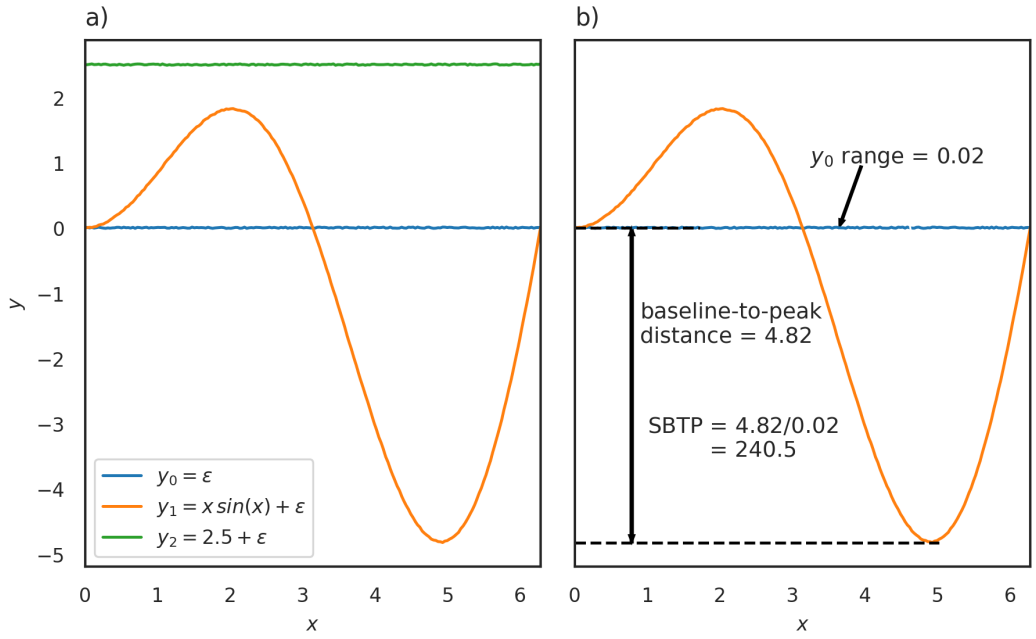

Figure 7.6: Figure 7.6a shows data generated from the same test function $y_i = a\,x\,\sin(x) + b + \varepsilon$, where a, b are both model parameters and ε is random Gaussian noise. x was varied from 0 to 2π, producing data y_0, y_1 and y_2 for the parameter sets $\Theta_0 : a = 0, b = 0$, $\Theta_1 : a = 1, b = 0$ and $\Theta_2 : a = 0, b = 2.5$ respectively. Despite both y_1 and y_2 being qualitatively very different they are very similar when summarised using only the Euclidean distance, with y_1 having a Euclidean distance $\varepsilon_{\text{euc},1} = 35.58$ and y_2 having a Euclidean distance $\varepsilon_{\text{euc},2} = 35.44$. If we instead look at the scaled baseline-to-peak (SBTP) distance we find that y_1 has a SBTP distance $SBTP(y_1) = 240.5$ and y_2 has a SBTP distance $SBTP(y_2) = 0.27$, giving $\varepsilon_{\text{SBTP},1} = 240.2$ and $\varepsilon_{\text{SBTP},2} = 0.11$.
Figure 7.6b illustrates how the scaled baseline-to-peak distance is defined using $x\,\sin(x) + \varepsilon$ as the example signal. The baseline-to-peak distance is the absolute distance from the baseline to $\max(\{|y_{max}|, |y_{min}|\})$. This is then divided by the range of the 'default' data, y_0, to get the distance as a proportion of the total change seen within the data. In this example, baseline-to-peak distance is 4.82 and the range is 0.02, giving the previously mentioned SBTP distance of 240.5.

summarised using only the Euclidean distance, with y_1 having a Euclidean distance $\varepsilon_{\text{euc},1} = 10.01$ and y_2 having a Euclidean distance $\varepsilon_{\text{euc},2} = 10.03$. This means that we would fail to clearly identify parameter a as being important than parameter b in producing sinusoidal behaviour.

Instead we can define a new summary measure, which we will call the "scaled baseline-to-peak" (SBTP) distance. We know that we want to find the parameter that determines how sinusoidal our model is. One way to emphasise this behaviour is to find the distance from our baseline to the maximum or minimum (whichever has the largest absolute value) of our data, as illustrated in Figure 7.6b. We then scale this by the range of our 'default' signal, y_0, to normalise it and avoid issues comparing data of different magnitudes. This gives us

$$SBTP(y_i) = \frac{\max\left(\{|\max(y_i) - y_i(t=0)|, |\min(y_i) - y_i(t=0)|\}\right)}{\max(y_0) - \min(y_0)}) \tag{7.4}$$

We then find the Euclidean distance between the SBTP value for our 'default' data, $SBTP(y_0)$, and $SBTP(y_1)$ and $SBTP(y_2)$

$$\varepsilon_{SBTP,i} = \sqrt{(SBTP(y_0) - SBTP(y_i))^2}, \tag{7.5}$$

where here $i \in \{1,2\}$.

If we use ε_{SBTP} as our summary measure, we find that y_1 has a distance $\varepsilon_{\text{SBTP},1} = 240.2$ and y_2 has a distance $\varepsilon_{\text{SBTP},2} = 0.11$. This would mean that parameter a could be clearly identified as being more important in producing sinusoidal behaviour than parameter b.

We scale our baseline-to-peak distance because a number of model outputs significantly vary over different scales. For example, cerebral oxygenation can be measured through TOI which is a percentage and, as seen in figure 7.3 can vary over 10-20%. Cytochrome-c-oxidase however, varies over a much smaller range, with a change of less than 1 μM being typical. Failing to account for these different scales will lead to parameters that affect larger magnitude outputs being identified as more sensitive than those that affect smaller magnitude outputs, even if the rela-

tive change is significant.

For example, if changing a parameter θ_1 causes the CCO change seen in figure 7.3e) to double to a minimum of $-2\,\mu M$, whilst a change in a parameter θ_2 causes TOI to decrease to 55%, without scaling the model seems more sensitive to θ_2 because the magnitude of the change is much more, even though the relative change is smaller. If we consider this change proportional to the range of our data however, we account for their relative size.

It should also be noted that this choice of metric is specific to the behaviour being optimised for. For example, in the case of a signal that is non-oscillatory, a different summary method would be required based around the behaviour to be replicated within that particular signal. We also acknowledge that there are a variety of different methods for identifying a sinusoidal signal from a linear signal and that our choice of metric here is one of many. We have chosen it as in the case of our hypercapnia data, we expect to see our signal to change from baseline to maxima or minima, depending on the signal, before then returning to baseline. The SBTP distance emphasises this behaviour in a single number whilst also being easily comparable to previous work where the Euclidean distance was used.

We used the Morris elementary effect method [Morris, 1991] variant devised by Saltelli et al. [Saltelli et al., 2002]. This provides us with two notable statistics: the mean of the absolute values of the changes, μ_*, and their standard deviation, σ. The larger the value of μ_*, the more influential parameter is on the output, whilst the larger the standard deviation, the more non-linear the influence of the parameter is. The top ten most sensitive parameters, as per μ_* were chosen to fit the model. σ was not used to determine which parameters to fit as, whilst knowing the non-linearity of a parameter is useful, in previous work [Caldwell et al., 2015b,a] we have opted to use simply μ_* as this gives a good summary of the sensitivity of a single parameter and feel it is pertinent to continue to do so here. The parameter range considered for sensitivity is the default value $\pm 50\%$. Sensitivities are calculated for each output as well as across all outputs jointly. This joint sensitivity is calculated by summing the SBTP value for each output and then determining variability in this total.

7.2.4 Approximate Bayesian computation

After selecting the most important parameters, the model was fit using the rejection algorithm as out lined previously in Section 5.3.1.

The output of the ABC algorithm used will be a sample from the distribution $p(\theta|d(y,y^{rep}) \leq \varepsilon)$. If ε is sufficiently small, then $p(\theta|d(y,y^{rep}) \leq \varepsilon)$ will be a good approximation for the posterior $p(\theta|y)$.

The choice of $d(\cdot,\cdot)$ is important, just as with the sensitivity analysis. We've chosen different metrics to the SBTP distance used in the sensitivity analysis as one metric is chosen for moedul reduction and to isolate specific parameters, whilst the metric used for fitting should instead optimise for simulation accuracy. Previously the Euclidean distance has been used to fit the model but, as in the case of the sensitivity analysis, this fails to account for outputs that vary over different magnitudes. Instead, we have chosen to include a number of other distance metrics including the root-mean-square error (RMSE) and the normalised root-mean-square error (NRMSE). These are defined as

$$RMSE = \sqrt{\frac{\sum_{t=1}^{T}(x_{1,t} - x_{2,t})^2}{T}} \tag{7.6}$$

$$NRMSE = \frac{RMSE(\mathbf{x}_1, \mathbf{x}_2)}{x_{1,max} - x_{1,min}} \tag{7.7}$$

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are the two time series being compared, running over $t = 1$ to $t = T$, with T being the total number of time points.

By dividing the RMSE by the range of the data, the errors for time series that vary over different magnitudes are comparable. Without doing this, parameters that mainly affect outputs that vary over larger magnitudes are preferentially optimised. Normalisation prevents overfitting of one output at the expense of others, providing a more reliable joint posterior distribution after fitting.

After an initial exploratory fitting of the different datasets, it was found that setting an absolute tolerance value was not a suitable selection criteria. This was due to massively differing distance values between datasets, with all parameter combinations in the simulated healthy dataset producing NRMSE values smaller than almost

all parameter combinations on the impaired dataset.

In general, the number of accepted samples that gives an adequate approximation of the posterior distribution is problem dependent; dispersed posterior distributions will ultimately require more samples. Poor estimation of the posterior can in most cases result in a wide posterior predictive distribution which appears to give a poor quality fit because outlier posterior samples cause biases. To address this issue in a pragmatic way, a fixed acceptance rate of 0.01% was set. This meant the 0.01% parameter combinations with the lowest $d(y, y^{rep})$ were used as the posterior. The posterior was visualised through kernel density estimation on a pairplot using the Seaborn plotting package [Waskom et al., 2018]. The posterior predictive density is then generated by sampling directly from the posterior 25 times and the model simulated for each sample. The results are aggregated and plotted, with the median and 95% credible interval marked on the plot.

The model was run in batches of 10,000,000 and the parameter combinations within the acceptance rate were used as a posterior. This batch size was chosen as a compromise between sufficient sampling of the parameter space and the computational time required to run the batch. The quality of the fit obtained from this posterior determined if the model had been run a sufficient number of times to sample the posterior adequately. If the posterior predictive distribution failed to capture the behaviour seen in the “true" data, then the process was repeated until a more adequate fit was obtained.

7.3 Results

Supporting information is found in Appendix C but includes a number of additional model checking figures. These are referred to in the text where relevant.

7.3.1 Sensitivity Analysis

Simulated Data

Sensitivity analysis was performed for the simulated healthy dataset for the CCO, HbO_2, HHb and TOI outputs. Figure 7.7 shows the sensitivity analysis results across all four outputs individually and for the outputs considered jointly. The re-

sults are plotted as bar charts, with sensitivity, as per the μ_* value, on the x-axis. The corresponding σ values for each parameter can be seen in Figure C.1 in S1 Sensitivity Analysis σ values - simulated data.

Table 7.2 shows the selected parameters, their respective μ_* values and their defini-

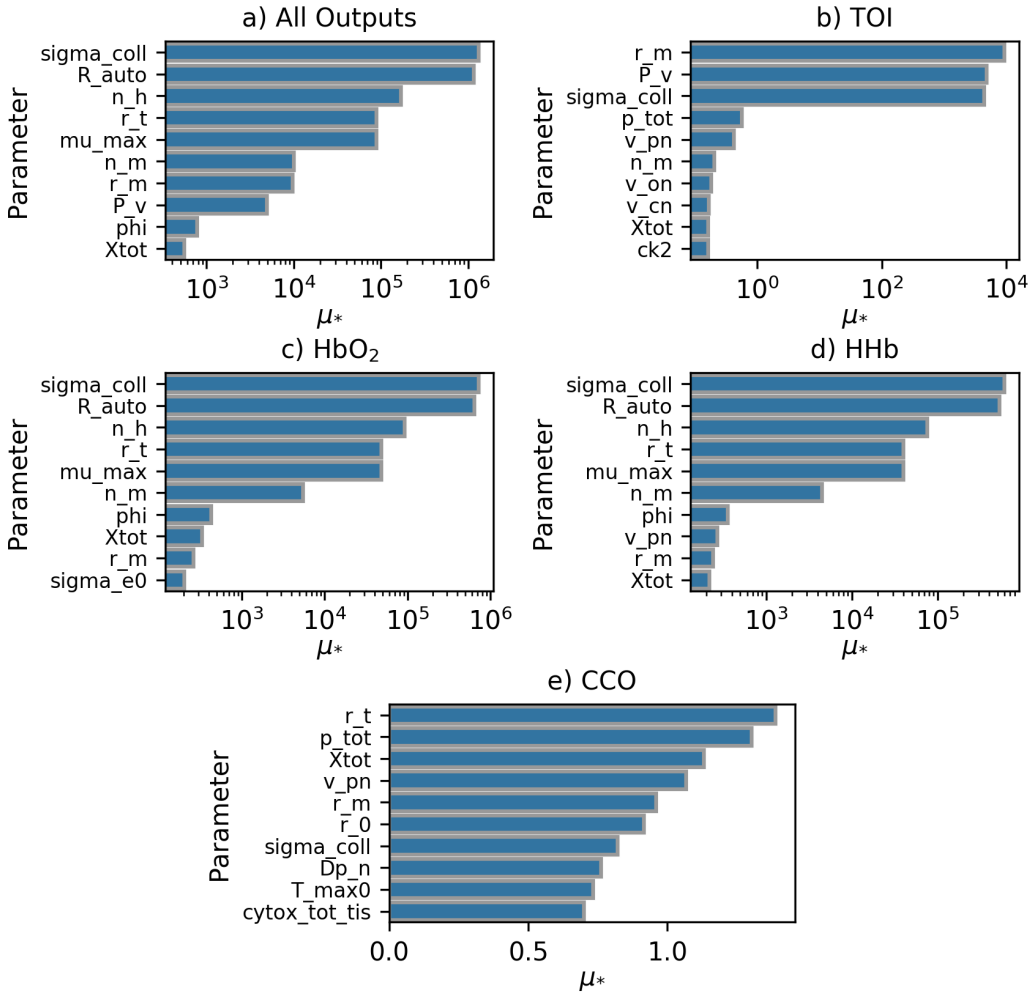

Figure 7.7: Sensitivity analysis across all outputs for simulated dataset. Bar charts showing μ_* for the 10 most sensitive parameters across all model outputs, with values plotted on a log scale where appropriate. Distance used for calculation is the sum of ε_{SBTP} across all model outputs. All outputs except cytochrome-*c*-oxidase alone have μ_* values that vary on a logarithmic scale. Figure 7.7a shows results for all outputs combined, figure 7.7b for TOI, figure 7.7c for HbO_2, figure 7.7d for HHb and figure 7.7e for CCO.

tions and default values. The total sensitivity analysis results, shown in figure 7.7a, produced 10 parameters to be used in fitting the model. Sensitivity analysis based

Table 7.2: Sensitivity analysis results for simulated data, including each selected parameter's definition and default value. * See Caldwell et al. [2015a] and Banaji et al. [2008] for a full explanation of this parameter and the stimulus μ.

Parameter	μ_*	Definition	Default value
sigma_coll	1.32×10^6	Pressure at which blood vessels collapse.	62.79 mmHg
R_auto	1.16×10^6	Autoregulatory reactivity to oxygen.	1.5
n_h	1.68×10^5	Hill coefficient for oxygen dissociation from haemoglobin.	2.5
r_t	8.82×10^4	Radius in the muscular tension relationship.	0.018 cm
mu_max	8.82×10^4	Upper bound for the transformed stimulus μ^*.	1
n_m	9.97×10^3	Exponent in the muscular tension relationship.	1.83
r_m	9.77×10^3	Vessel radius at which muscular tension is maximal.	0.027 cm
P_v	4.88×10^3	Venous blood pressure.	4 mmHg
phi	7.85×10^2	Oxygen concentration at half-maximal saturation.	0.036 mM
Xtot	5.54×10^2	Total concentration of haemoglobin O_2 binding sites in blood.	9.1 mM

on individual outputs showed that different parameters were important for different outputs, with TOI, in figure 7.7b being dominated by `r_m`, `P_v` and `sigma_coll` and oxyhaemoglobin, in figure 7.7c, and deoxyhaemoglobin, in figure 7.7d, dominated by `sigma_coll` and `R_auto`. Cytochrome-*c*-oxidase however showed levels of dependence that were similar across many parameters, as seen in figure 7.7e, with μ_* values falling within a range of 0.7. For all individual outputs and the combined output, only `Xtot`, `r_m` and `sigma_coll` were within the 10 most sensitive in all cases.

Experimental Data

Sensitivity analysis was undertaken on the experimental dataset to determine the parameters to be fit. Table 7.3 shows the selected parameters, their respective μ_* values and their definitions and default values. Figure 7.8 shows the results across all outputs. The corresponding σ values for each parameter can be seen in Fig-

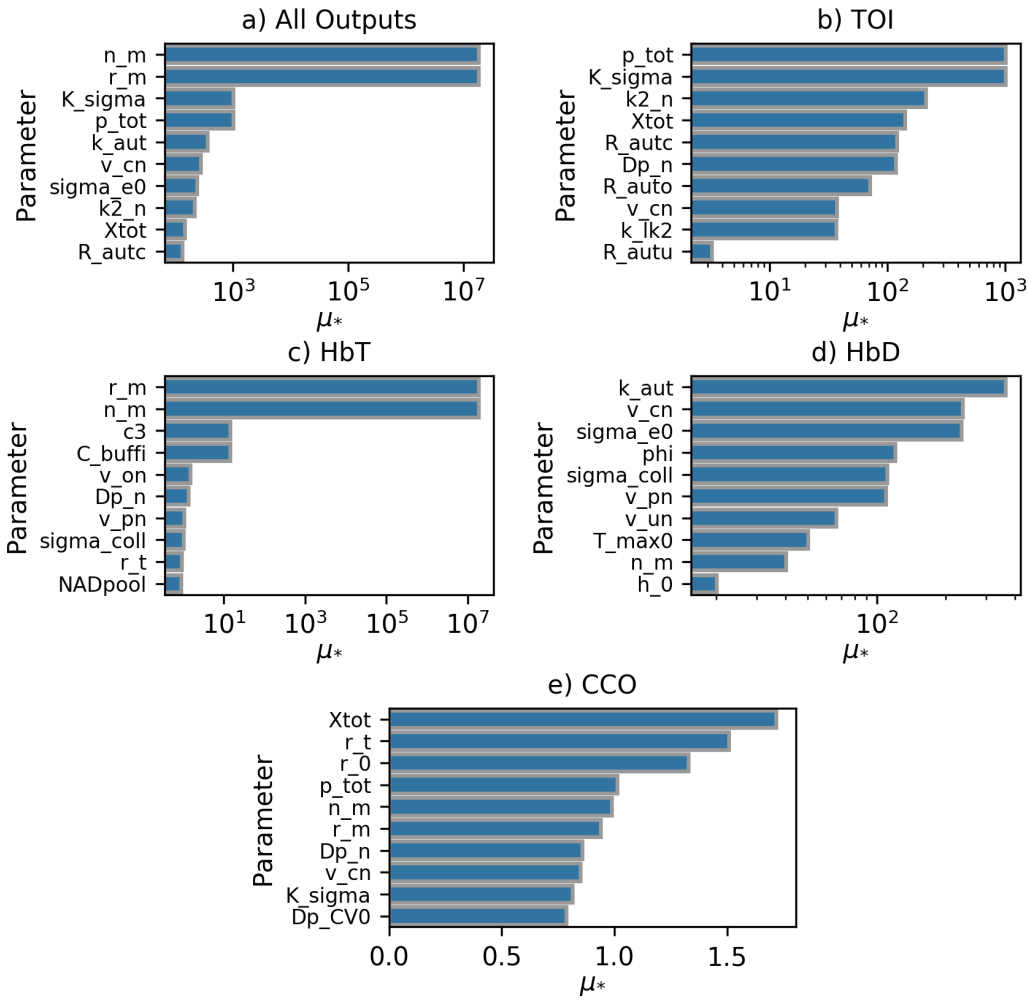

Figure 7.8: Sensitivity analysis across all outputs for experimental dataset. Barplots showing μ_* values for the 10 most sensitive parameters across all model outputs, with the x-axis plotted using a log scale where appropriate. Distance used for calculation is the sum of ε_{SBTP} across all model outputs. Figure 7.8a shows results for all outputs combined, figure 7.8b for TOI, figure 7.8c for HbT, figure 7.8d for HbD and figure 7.8e for CCO.

ure C.2 in S2 Sensitivity Analysis σ values - simulated data. When considering all outputs jointly, the effect of `n_m` and `r_m` is significantly larger than all other parameters, but when looking at the individual outputs it's clear that the other parameters are still important, but the magnitude of the impact `n_m` and `r_m` have on the overall variability is drastically larger.

Unlike the simulated dataset, 9 of the top 10 most sensitive parameters have μ_* values between approximately 10 and 1000 which is significantly smaller than the range of the μ_* values for TOI in the simulated data.

Similarly, the most sensitive parameters for HbD fall within a very small range with no one parameter obviously determining the majority of the output's behaviour. In contrast, the two most sensitive parameters for HbT, `r_m` and `n_m`, are approximately 10^6 times larger than the third highest. As with the simulated data, μ_* values for CCO have much smaller values than all other outputs and fall within a range of 1.0. Unlike the simulated data, no parameters were sensitive across all individual and joint outputs.

Parameters

Whilst a full exploration of the parameters within the BrainSignals model is outside the scope of this paper, we advise the reader to look at the original publications Banaji et al. [2008], Caldwell et al. [2015a], and provide a brief overview of some of those identified as important here.

A number of the parameters identified as being important for the above datasets, such as `R_auto` and `mu_max`, are dimensionless parameters. They are often model specific parameters that cannot be directly measured and instead need to be considered in the context of their meaning within the model. For example, an increase in `R_auto` would mean that the autoregulatory response would become more sensitive to changes in oxygen concentration. In contrast, other parameters such as `Xtot`, which is four times the concentration of haemoglobin, are more easily measured in an experimental or clinical setting.

Some of the parameters identified as important are linked closely to the shape of

Table 7.3: Sensitivity analysis results for experimental data, including each selected parameter's definition and default value.
*See Caldwell et al. [2015a] and Banaji et al. [2008] for a full explanation of this parameter and the stimulus μ.
†This is the arterial $PaCO_2$ input put through a first order filter to simulate varying time response and is typically the same as arterial $PaCO_2$. For more information see Banaji et al. [2008]

Parameter	μ_*	Definition	Default Value
n_m	1.87×10^7	Exponent in the muscular tension relationship.	1.83
r_m	1.87×10^7	Vessel radius at which muscular tension is maximal.	0.027 cm
K_sigma	1.01×10^3	Parameter controlling the sensitivity of σ_e to vessel radius *.	10
p_tot	6.29×10^3	Total protons removed from the mitochondrial matrix by the three modelled electron transport reactions.	20
k_aut	3.65×10^2	Parameter controlling overall functioning of autoregulatory response, with a value of 1 meaning intact autoregulation.	1
v_cn	2.76×10^2	Normal filtered $PaCO_2$†.	40 mmHg
sigma_e0	2.37×10^2	Parameter in the elastic tension relationship.	0.1425 mmHg
k2_n	2.17×10^2	Normal forward reaction rate for the reduction of a_3.	3915.68 s^{-1}
Xtot	1.48×10^2	Total concentration of haemoglobin O_2 binding sites in blood (4 times the haemoglobin concentration).	9.1 mM
R_autc	1.31×10^2	Autoregulatory reactivity to carbon dioxide.	2.2

the autoregulatory response of the model and its sensitivity to changes in model inputs. These are `R_auto` and `mu_max` in the simulated dataset, and `v_on`, `v_un`, `R_autc` and `v_cn` in the experimental dataset. As we are driving the model with a changing input, the identification of these parameters as important seems physiologically sensible. It should also be noted that, despite other parameters not directly controlling the autoregulation response, the interconnected and complex nature of the BrainSignals model means that other parameters may still have an impact on it indirectly, for example the parameter `r_t` controls the stiffness of blood vessel walls, which is important in controlling blood flow during autoregulation.

More detailed information on the exact nature of these parameters and how they function within the BrainSignals model can be found in [Caldwell et al., 2015a].

7.3.2 Bayesian Analysis

Simulated Data

The BrainSignals model was fit to the simulated "healthy" dataset initially. The model was run for 10,000,000 different parameter combinations before determining that the posterior had been estimated sufficiently well, based on the quality of the posterior predictive distribution. The samples in the posterior were found to have $0.019170 \leq \varepsilon_{\text{NRMSE}} \leq 0.098098$. Figure 7.9 shows this posterior distribution in blue. `Xtot`, `phi` and `r_t` show narrow marginal distributions whilst the others are much wider. Median values for all parameters are close to the model value, with `R_auto` showing a skew towards lower values in its marginal distribution that also leads to a median slightly lower than the model value. Figure 7.10a shows the posterior predictive distribution produced by sampling 25 times directly from the posterior, and shows a very good fit.

This healthy posterior was then used to define an impaired brain, as mentioned above. `r_t` was set to 0.013 cm and the model driven with the same inputs as the healthy simulation. This "impaired" dataset was then fit using the same approach as above, using the sensitivity analysis results. The model was run 30,000,000 different parameter combinations, with the increased run number required in order to sufficiently estimate the posterior. With an acceptance rate of 0.01%, a posterior

was produced based on 3000 samples having $0.019170 \leq \varepsilon_{\text{NRMSE}} \leq 0.267152$. Despite the higher error values as compared to the healthy data, the resulting fit was still deemed very good. Figure 7.9 shows this posterior in orange and figure 7.10b shows the time series generated by sampling 25 times directly from this posterior. `Xtot`, `phi` and `r_t` show marginal distributions that are narrower than the others, but wider than those seen in the healthy posterior. All parameters have median values close to the value set in the model. A separation between the healthy `r_t` and impaired `r_t` marginal distributions is clearly visible.

Figures 7.10c and 7.10d show a zoomed in view of each output in order to show the 95% credible interval of the posterior predictive distribution. This is not clearly visible on the full trace as it is reasonably small.

Sections S6 Statistical Analysis - Simulated Healthy Data and S7 Statistical Analysis - Simulated Impaired Data of Appendix C show a number of different statistical analyses of the results for the healthy and impaired datasets respectively. These are all posterior predictive checks, where the posterior predictive distribution is used to produce a number of statistical results that can be used to assess the quality of the model fit. Figure C.3 shows the autocorrelation of both the posterior predictive and observed data as a function of lag for each signal in the healthy data. Figure C.4 shows the distribution of the residuals for each signal in the healthy dataset, with the mean and standard deviation drawn on. Q-Q plots for these distributions are shown in Figure C.5 and are used to assess the normality of the residuals. Figure C.6 shows the prior and posterior distributions for each parameter along with the Kullback-Leibler Divergence for each of these, giving us a sense of how much information was gained when moving from prior to posterior distribution. Figures C.7, C.8, C.9 and C.10 show these same posterior predictive checks but for the impaired distribution.

The autocorrelation plots in figures C.3 and C.7 show that the autocorrelation of the posterior predictive distribution and the observed data match closely across all lag values for both healthy and impaired datasets. Looking at the distributions of the residuals for each signal in figure C.4 and C.8 we can see that the residuals for

TOI appear to be normally distributed, whilst the other three signals all show generally symmetric but leptokurtic distributions. The Q-Q plots in figures C.5 and C.9 confirm this, suggesting that the residuals across all signals are generally normally distributed, with some slight differences at the highest and lowest quantiles. Residuals for HHb and CCO both appear to be generally more leptokurtic than both TOI and HbO_2.

Figure 7.9: Comparison of posterior distributions for healthy and impaired simulated data. Fig 7.9 shows the posteriors for healthy and impaired data based on an acceptance rate of 0.01%. Posterior are shown over the full prior range as defined in S3 Table and S4 Table.

Experimental Data

When approaching the experimental data, the criteria for a good fit were different to those in the simulated dataset. With the simulated dataset, any parameters not chosen for fitting would have the same value during the fitting process as during

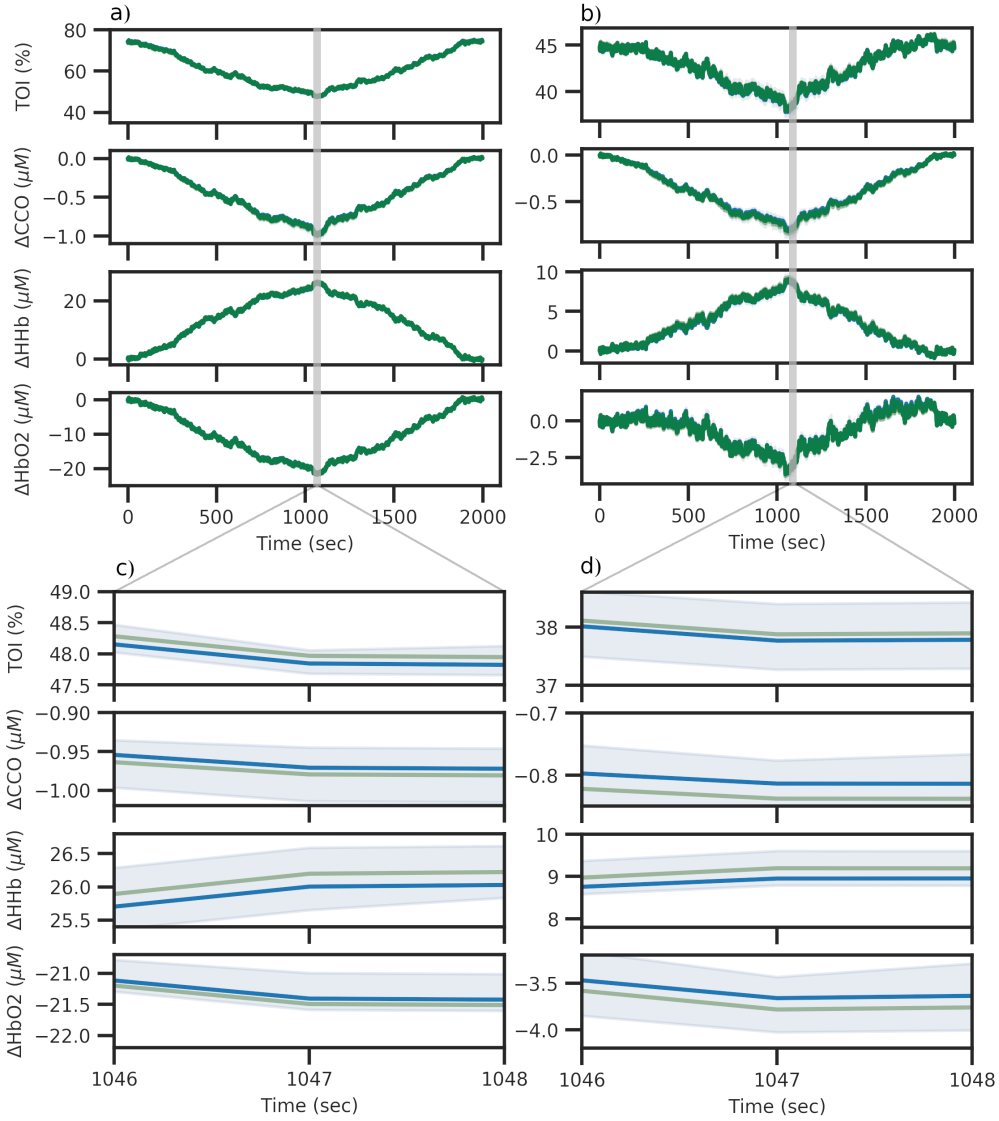

Figure 7.10: Comparison of predictions for healthy and impaired simulated data. Figures 7.10a and 7.10b show the predicted time series data from the healthy and impaired posteriors respectively. Each posterior was sampled 25 times and the resulting runs aggregated, with the median and 95% credible intervals plotted in dark blue and light blue respectively. Figures 7.10c and 7.10d show a zoomed in view of each output in order to show the credible interval of the posterior predictive distribution.

the generation of the simulated dataset. In the experimental data however, it is almost certain that the default values of any parameters not chosen for fitting would not have the exact same value as their biological, real-world analogue. As a result, instead of looking for a perfect fit, we instead look for qualitative behaviours to be reproduced, such as the periodic increase and decrease in output values due to the repeated hypoxia challenges.

The fitting process required 20,000,000 parameter combinations before a satisfactory fit was obtained, and with an acceptance rate of 0.01% the posterior in figure 7.11 consisted of 2000 samples with $0.778492 \leq \varepsilon_{\text{NRMSE}} \leq 0.802900$. The model was also fit using the previous OpenOpt method, and the values obtained from that are also shown for comparison. We can see that for parameters with reasonably well defined posterior, the OpenOpt values and the posterior median are reasonably close, but for those showing a wider distribution, the OpenOpt value can vary massively from the posterior median. For `sigma_e0` and `k2_n` the OpenOpt value is at one extreme end of the prior range, whilst the median remains central due to the distribution being uniform. Figure 7.12 shows the predicted time series for all outputs based on the posterior shown in figure 7.11. The posterior was sampled 25 times with the resulting time series aggregated, with the median and 95% credible intervals plotted. Overall behaviour is reflected in the predicted trace, with 3 distinct periods of hypoxia visible as periodic behaviour within all signals. Shown in green is the fit obtained using the OpenOpt method, which has an error $\varepsilon_{\text{NRMSE}} = 0.77518$. It is clear that both methods are able to achieve similar fits, but the Bayesian method provides more information about the space of possible parameter combinations and the resultant uncertainty in fitted model output.

As with the simulated datasets, we have produced posterior predictive checks to assess the fit of the model to the data. These are shown in Section S8 Statistical Analysis - Experimental Data of Appendix C. Figures C.11, C.13 and C.14 show the autocorrelation comparisons, Q-Q plots and prior-posterior comparisons plots respectively.

The autocorrelation plot in Figure C.11 show that the autocorrelation of the pos-

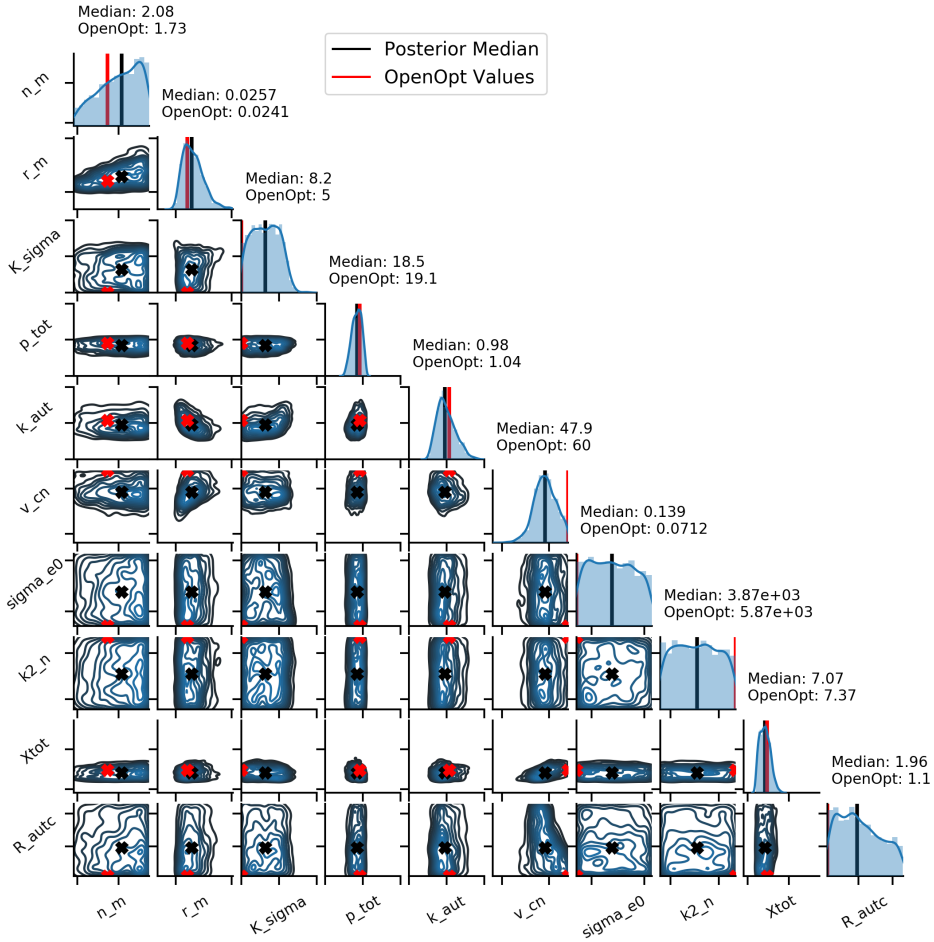

Figure 7.11: Posterior distributions for the experimental dataset. Fig 7.11 shows the posterior distribution for the experimental dataset, based on an acceptance rate of 0.01%. The posterior median is shown in black and the OpenOpt predicted value is shown in red. Posterior are shown over the full prior range as defined in S5 Table

terior predictive distribution and the observed data generally match across all lag values, with the same shape seen in both plots, but with a slight difference in the magnitude of the autocorrelation. The distributions of residuals in figure C.12 all appear to be relatively normally distributed, with TOI, CCO and HbD showing a mean close to zero. HbT however has a mean noticeably less than zero, which is due to the simulation predicting generally lower values than in the observed data, as seen in figure 7.12. These results are more clearly seen in figure C.13, which suggests that the residuals for all signals are generally normal, with some bimodality in the HbT distribution and a significant amount of positive skew in the HbD

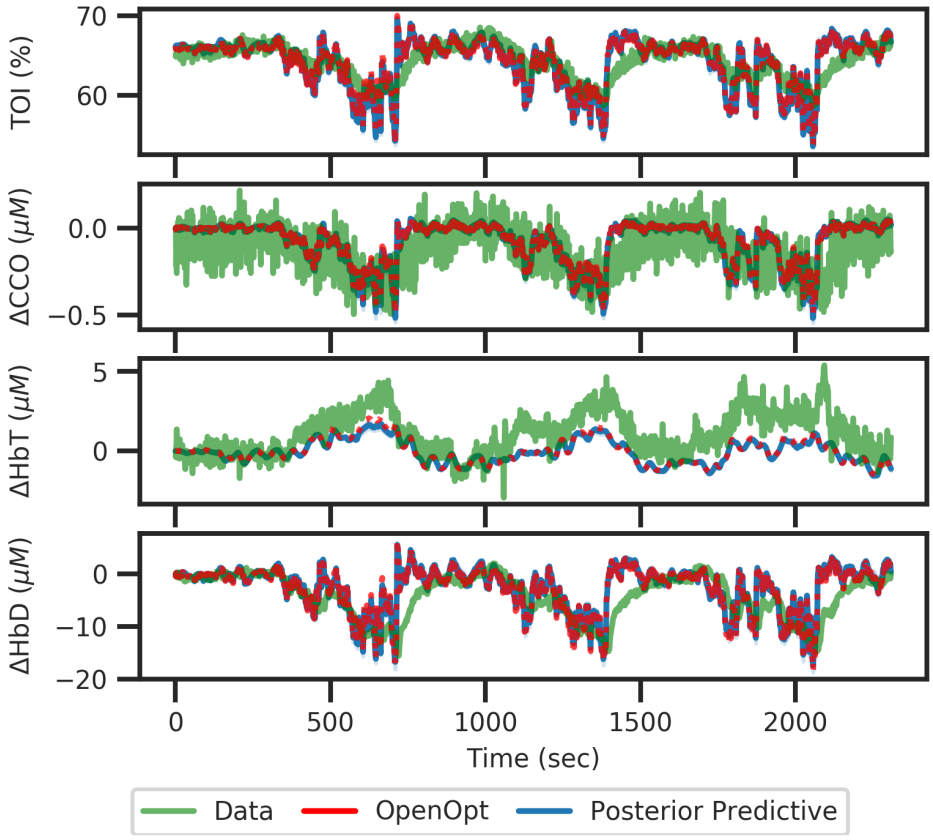

Figure 7.12: Predicted fits for the experimental dataset for all outputs, based on the posterior shown in Fig 7.11. The posterior was sampled 25 times with the resulting time series aggregated, with the median and 95% credible intervals plotted in dark and light blue respectively. Overall behaviour is reflected in the predicted trace, with 3 distinct periods of hypoxia visible as periodic behaviour within all signals. The fit obtained using OpenOpt is shown in red.

distribution.

7.4 Discussion

In this work we have introduced a new Bayesian analysis for interpretation of the BrainSignals models. The process was tested and used to analyse two simulated datasets and one experimental dataset. The Bayesian approach provides us with complete information about the parameter space and takes into account the prior information we have about physiological parameters via the proposal distribution, $p(\theta)$, which allows us to simulate distributions of input parameters. Both of these factors are extremely important when drawing physiological conclusions from any

parameter estimates.

Using the posterior predictive checks in sections S6 Statistical Analysis - Simulated Healthy Data, S7 Statistical Analysis - Simulated Impaired Data and S8 Statistical Analysis - Experimental Data of Appendix C we have shown that the Bayesian method is able to produce good fits for a range of different datasets, including both overly simplistic test cases, with just a single parameter change, and real world measured data which contain inherently more complexity than simulated data. It can also be seen from direct observation of the posterior predictive distribution in Fig 7.12 that the Bayesian method is able to provide a fit equally as good as the previous OpenOpt method.

We have shown how the method can be used to define healthy and impaired parameter spaces, as shown with the simulated datasets, and how for some parameters these spaces may overlap. We have also shown how the new Bayesian approach provides more information about the parameter space than the previous OpenOpt maximum likelihood method. Looking at only the healthy dataset, the parameters `sigma_coll`, `P_v` and `mu_max` all have marginal posteriors with a median at the default value set in the model, but with distributions that cover the entirety of the prior distribution initially set. Determining that a parameter's posterior distribution is not tightly constrained is important when drawing physiological conclusions from the model fitting process. This is seen even more clearly in the direct comparison between prior and posterior distributions for these parameters in Fig C.6 and in their K-L divergence values of 0.009 71 nats, 0.0109 nats and 0.014 nats respectively. If we compare these to the plots and K-L divergence values for `phi` and `Xtot`, which have values of 1.38 nats and 1.04 nats respectively, it is clear that the Bayesian process provides significantly more information than the previous OpenOpt method, both in terms of producing posteriors that have significant information gains compared to prior distributions, and in identifying parameters where a good fit is produced despite minimal information gain. Using only sensitivity analysis and OpenOpt would not provide this extra insight.

This is seen more clearly when looking at the experimental data. Many of the

parameters show relatively narrow marginal posteriors, but `sigma_e0` and `k2_n`, which were both identified as important by the original sensitivity analysis, are both shown to have wide distributions, suggesting insensitivity within the prior range. The previous OpenOpt method produces an almost identical fit as the Bayesian approach but provides significantly less information about the parameter space. For `sigma_e0`, `k2_n`, `v_cn`, `R_autc` and `k_sigma` the OpenOpt values fall outside of the interquartile range of the posterior distribution, yet produce equivalent model simulations. If considering the OpenOpt estimate alone, it would be simple to draw the conclusion that these parameters have shifted away from the default 'healthy' value, showing some sort of physiological change during the hypoxia challenge. However, when we look at the posterior obtained through the Bayesian method, the median value is close to the default value and in fact parameter values across the entire prior range produce similar results. As a result we can instead say that for these data the model is insensitive to these parameters, with a median value that would be considered 'healthy'. Again, if we look at the comparison between prior and posterior distributions and consider the K-L divergence for each parameter in Fig C.14, it is clear to see where we have and have not gained information through the use of the Bayesian process.

There are a number of other methods that can also be used to explore and define the parameter space. The previously used maximum likelihood based method, for example, can provide estimates and confidence limits of parameter values, but under the assumption that the maximum likelihood estimator is normally distributed around the maximum. It may also be possible to use a profile likelihood [Venzon and Moolgavkar, 1988], but whilst this will provide information about the distribution of the parameter space without assuming normality, it is computationally expensive and does not take into account prior information about the parameters.

It is acknowledged that the Bayesian approach is not without its own limitations. Historically, non-trivial problems were not solvable analytically due to the high dimensional integrals required. However, with the relatively recent availability of more computational power, a number of algorithms and approaches are now avail-

able that allow these problems to be approximated. This has seen increased uptake of Bayesian approaches within the fields of systems biology and genetics, where the inherently complex models and noisy data that these fields involve are particularly well suited to being analysed through the Bayesian approach. As long as a statistical model can be used to relate the relevant quantities, Bayesian inference can be used to give full probabilistic information on all unobserved model variables.

One of the main drawbacks to this method is that the number of model runs required to have sufficient samples in the posterior may be prohibitively high, especially where the tolerance is low or the prior distribution is very different from the posterior distribution.

This requirement for a large number of simulations for a reliable posterior is seen in all of the datasets used here. For the simulated 'healthy' data, the model sampled 10,000,000 parameter combinations in order to achieve the obtained fit. In contrast, to fit the 'impaired' simulated data the model required 30,000,000 parameter combinations to be sampled and for the same acceptance rate the accepted samples had generally higher $\varepsilon_{\text{NRMSE}}$ values. Finally, the experimental data were only able to obtain a good posterior after being sampling 20,000,000 parameter combinations and all $\varepsilon_{\text{NRMSE}}$ values were significantly above those seen in the simulated datasets. This is clearly visible in figure 7.13, where the distribution of $\varepsilon_{\text{NRMSE}}$ values for each posterior are clearly very different. This figure clearly highlights the variance in both the error values that define a 'good' fit and the number of samples required for a reliable posterior distribution for different datasets.

It should be noted that all of the obtained posterior distributions produce what are considered good fits, with those obtained for the simulated datasets far more accurate than we would ever expect to achieve when fitting experimental data. When looking at the experimental data in particular, despite the $\varepsilon_{\text{NRMSE}}$ values being much higher than in the simulated data, the obtained fit captures all important behaviour and phenomena, with three clear hypoxia events visible in the inferred data trace.

More efficient methods of ABC, may alleviate the problem of requiring so many model runs to obtain good posteriors. An approach based on MCMC is more effi-

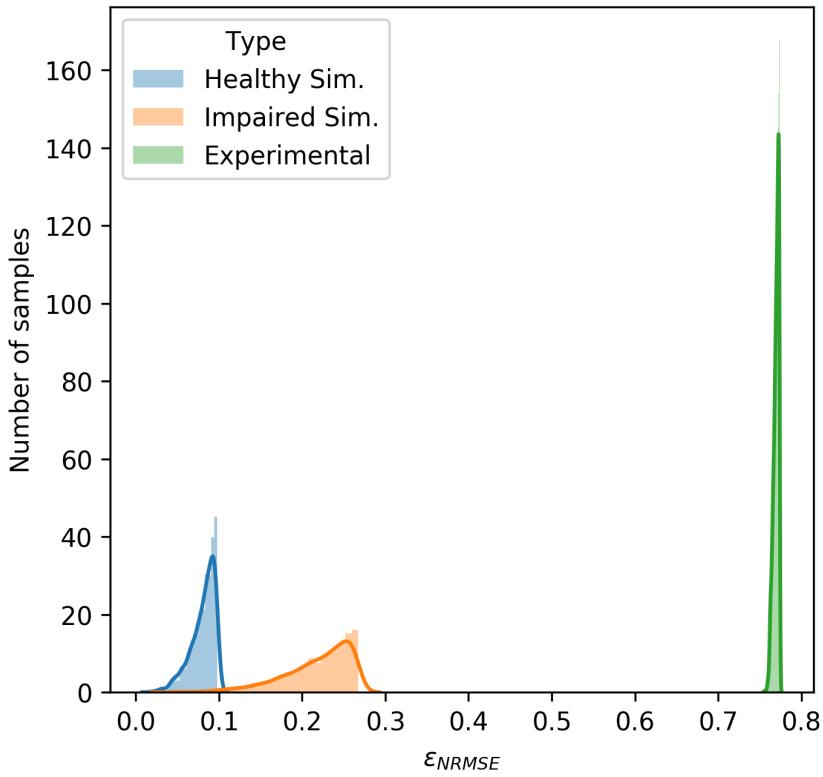

Figure 7.13: Distribution of $\varepsilon_{\text{NRMSE}}$ values for the posteriors of each dataset. It can be seen here that the three datasets had very different distributions $\varepsilon_{\text{NRMSE}}$ values for the samples that made up their respective posteriors. Despite this, the posterior predictive distributions for all datasets were good fits.

cient than ABC REJ but the chain may become stuck in regions of low probability for long periods of time [Marjoram et al., 2003]. In order to deal with this problem and also the disadvantages of the rejection algorithm, an approach based on sequential Monte Carlo (ABC SMC) [Toni et al., 2009] was first proposed by Sisson et al. [Sisson et al., 2007], as well as Beaumont et al. [Beaumont et al., 2009] and Cappé et al. [Cappé et al., 2004]. In this approach, a number of sampled parameter values, known as *particles*, are sampled from the prior distribution and then propagated through a number of intermediate distributions before reaching a final target distribution. The tolerance for each successive distribution is smaller than the previous, allowing them to evolve towards the target posterior. Additionally, for a sufficiently large number of particles, the problem in MCMC of getting stuck in areas of low probability can be avoided. Developing the BayesCMD framework to use an ABC SMC approach is a key focus for future work.

7.4.1 Conclusion

We have outlined how this new Bayesian framework for model analysis can be used with models of brain haemodynamics to extract information from physiological data. A more comprehensive picture of the parameter space is obtained, allowing physiological conclusions to be based on a broader picture. This is most clearly seen in the experimental data, where point estimates suggested that the values for a number of parameters had changed significantly during fitting, whilst the Bayesian method showed that the parameters were defined by a broad, roughly uniform distribution. We have also shown, through the use of data simulated from the BrainSignals model in healthy and impaired states, how the Bayesian approach allows us to better distinguish different parameter spaces. Finally, whilst we have focussed on using the BrainSignals model here, any model that can be written in a format compatible with BCMD can use this method to estimate model parameters. A major interest within our research group is to use these models and approaches to understand and investigate further our novel measures of brain tissue physiology and metabolism and how they are linked to brain injury [Mitra et al., 2017, Bale et al., 2018]. In particular, we are interested in neonatal hypoxic ischaemic injury. The Bayesian approach provides a better representation of the parameter space and can inform a better distinction between different brain states, such as between a mild and severe injury. The method will also be adapted to use more efficient methods of parameter estimation, such as ABC SMC, reducing the number of model runs required to obtain a given tolerance.

Chapter 8

Modelling Therapeutic Hypothermia

8.1 Piglet Data

Data were collected as per Kaynezhad et al. [2019]. Piglets aged less than 36 hours were sedated and their common carotid arteries isolated and occluded. Arterial oxygen saturation was monitored continuously as well as mean arterial blood pressure and heart rate. bNIRS was used to collected haemoglobin and cytochrome-*c*-oxidase signals. Prior to the hypoxic-ischaemic insult piglets were monitored for 1 hour to ensure stability and the insult lasted approximately 20-22 minutes. All animals received TH at 33.5 °C 1 hour after injury with some receiving additional treatments and additional pharmacological interventions throughout the 48-hour experiment.

Data were considered for four piglets that had previously been found to exhibit different behaviour and have different outcomes following the same clinical intervention to emulate HIE. These were LWP475, LWP479, LWP481 and LWP484. Of these, LWP475 and LWP481 were identified as having had a mild injury whilst LWP479 and LWP484 were identified as having had a severe injury. Table 8.1 indicates the HI duration, 10 minute CCO recovery fraction (RF) and Lac/NAA at 24 hours for each piglet. All data needed cleaning in order to remove artefacts and were chosen to encompass the period of cooling only, with data before and after this period removed. Cleaning involved passing data through a 5th-order, low-pass, 0.05 Hz Butterworth filter, as per Caldwell et al. [2015a], to remove high

Piglet	HI Duration (mins)	10 min CCO RF	24hr Lac/NAA	Injury Severity
LWP475	24	122%	0.21	Mild
LWP479	29	69%	1.03	Severe
LWP481	20	55%	0.26	Mild
LWP484	25	53%	8.32	Severe

Table 8.1: Table detailing each piglet's injury severity.

frequency noise before interpolating missing data using the `interpolate` method on the `pandas.DataFrame`[Mckinney, 2010]. During the cleaning process it was decided that only data from LWP477 and LWP479 would be used due to problems with data from the other two piglets. Figure 8.1 shows the data for both of these piglets. Piglet LWP481 showed signs of a large external intervention in the NIRS signals with a sudden increase coinciding with a drop in the measured arterial oxygen saturation and a sharp increase in ABP. Similarly, piglet LWP484 showed a similar sharp change coinciding with a dramatic change in blood pressure and oxygen saturation. Given that these changes are likely to be due to external stimuli, the model would be unable to properly model the observed behaviour and any conclusions and derived model parameters would be unreliable. Using the data up to this point was not a valid approach as it would have prevented us from modelling equivalent periods of cooling.

Figure 8.2 shows the final data for piglets LWP475 and LWP479. Though there is still some noise within the data, particularly in LWP475, the data were deemed sufficiently clean for modelling with a balance between removing any obvious artefacts or measurement noise and ensuring all relevant physiological information was retained. A clear difference is seen between the two piglets in terms of their response to cooling. LWP479 shows an initial increase in both ΔHbD and ΔCCO before then decreasing again whilst LWP475 shows a general decrease from the start of cooling. Conversely, LWP479 shows an initial decrease in ΔHbO_2 followed by a sharp increase, whilst LWP475 shows a steady increase from the start of cooling. It is this opposite behaviour that we hope to be able to identify mechanisms for in the hypothermia model.

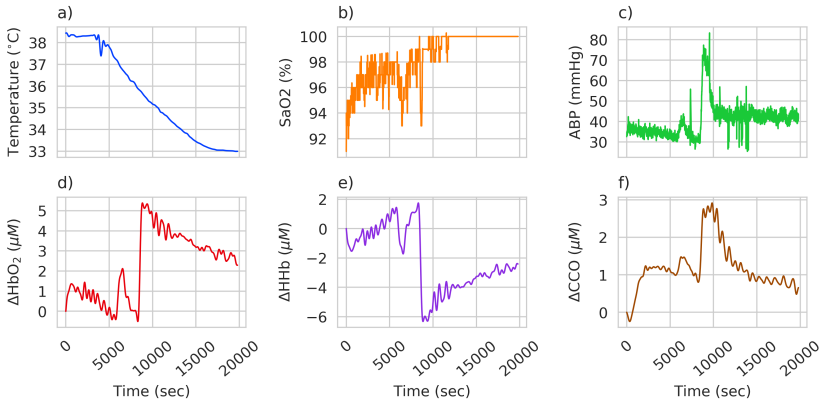

(a) Measured data for piglet LWP481

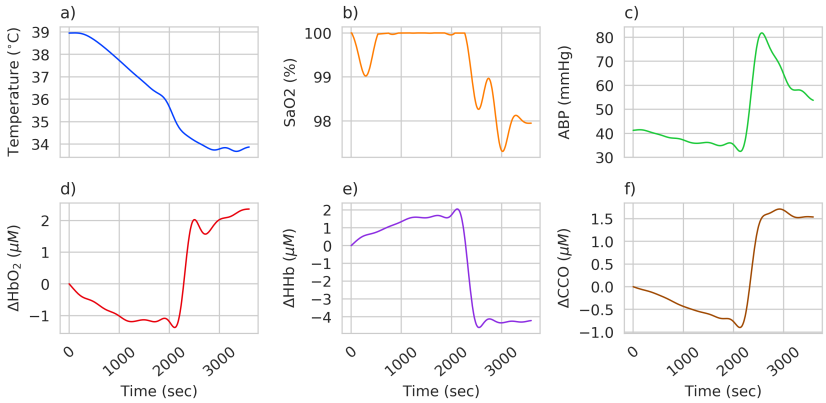

(b) Measured data for piglet LWP484

Figure 8.1: Measured data for piglets LWP481 and LWP484. Data were filtered using a low-pass, 5th-order Butterworth filter at 0.05 Hz. External intervention is visible in both piglets.

.

8.2 Hypothermia models

A total of four BrainPiglet model variants were analysed and compared for this work. BrainPiglet 2.1 is a control model in which temperature is fixed and hypothermia has no direct effect. All three hypothermia model variants, BP Hypothermia 1, 2 and 3, are affected by hypothermia but have slightly varying structures,as outlined below. Figure 8.3 provides a summary of these models as well as showing the general hierarchy of the models. Models are abbreviated as BPH*, where * refers to the model variant. BrainPiglet2.1 is referred to as BPH0 as it is our *control model*.

The BrainPiglet 2.1 model already contains temperature as a constant value in the calculation of the model parameter Z, which is a proportionality constant found

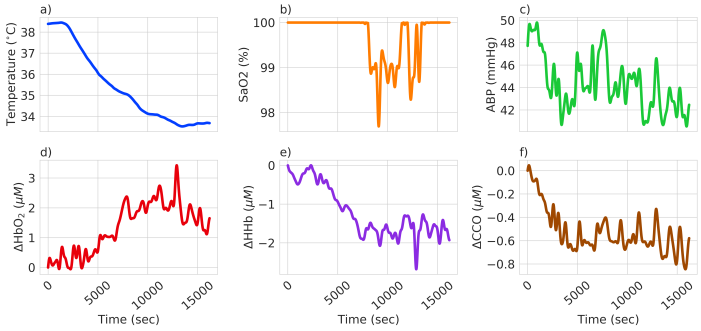

(a) Measured data for piglet LWP475

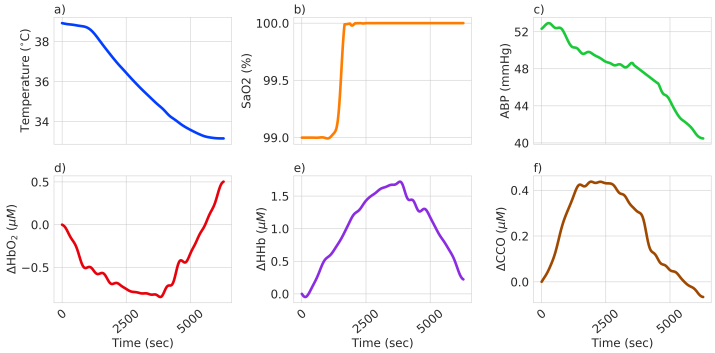

(b) Measured data for piglet LWP479

Figure 8.2: Measured data for piglets LWP475 and LWP479. Data were filtered using a low-pass, 5th-order Butterworth filter at 0.05 Hz with anomalous data removed. Any missing data were interpolated using the `interpolate` method on the `pandas.DataFrame`[Mckinney, 2010].

in equations within the metabolic sub-model that deal with the electron transport chain, specifically in the calculation of driving forces due to concentration differences.

This proportionality constant is defined as:

$$Z = \frac{RT}{F}$$

where R is the ideal gas constant, F is the Faraday constant and T is absolute body temperature as a *constant* value. We make this proportionality 'constant' a function of temperature, $T(t)$, which we set as a model input that can vary over time.

$$Z(T) = \frac{R}{F} \cdot T(t)$$

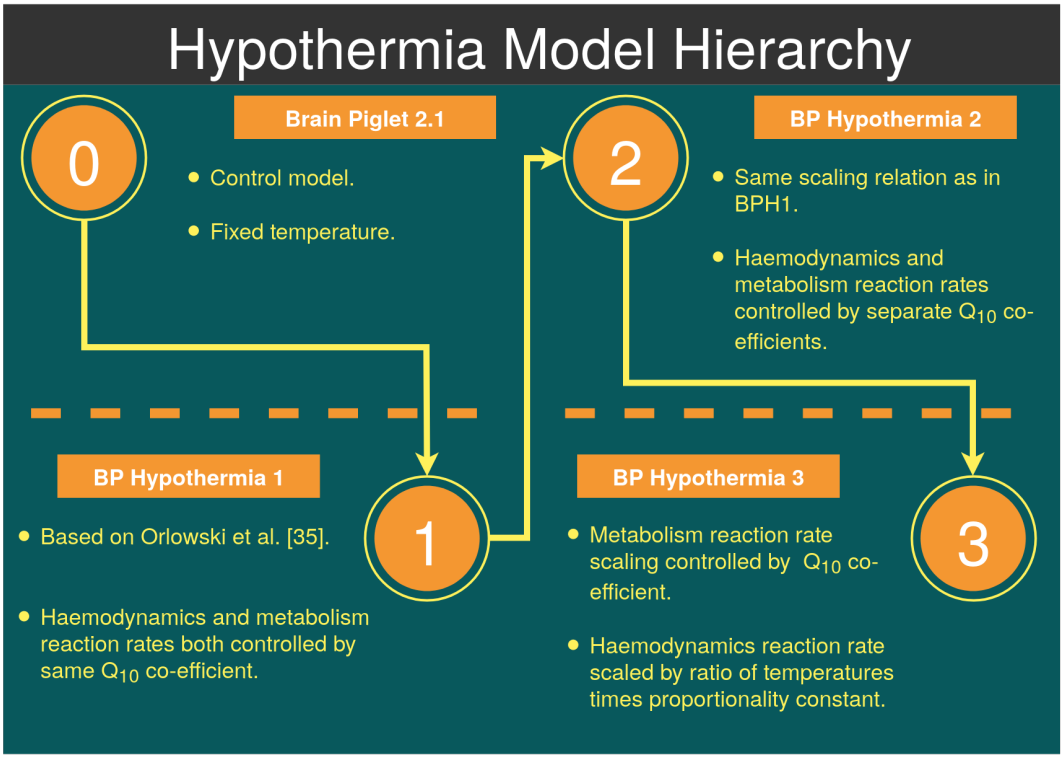

Figure 8.3: Hierarchy of the hypothermia models. Brain Piglet 2.1 is included here as our control model in which temperature is fixed and hypothermia has no effect. Models BP Hypothermia 1, 2 and 3 all have slightly varying structures as outlined below, but all are affected by hypothermia.

This adds temperature into the calculation of driving forces due to concentration differences, such as the proton motive force.

By making this proportionality constant a function of temperature it is possible to incorporate temperature within the calculation of driving forces due to concentration differences, without adding excessive complexity to the model. However, there are still other areas within the model where we expect temperature to have an impact, such as in diffusion of oxygen within the oxygen delivery sub-model or in the Michaelis-Menten type behaviour of many reactions that take place in the metabolic sub-model.

8.2.1 BP Hypothermia 1

The first hypothermia model variant, BP Hypothermia 1 (BPH1), is the simplest of the model variants with only a single major addition to the model. This is based on

work by Orlowski et al. [2014], in which they expanded an existing model of cellular metabolism by incorporating temperature dependence into the rates of reaction and ion diffusion. This was done to model the impact of therapeutic hypothermia on stroke prevention. Their model was based on previous work modelling pH dynamics in brain cells after stroke that included capillaries, neurons, astrocytes, and extracellular space with one compartment per tissue type [Orlowski et al., 2011]. Reaction rates, Michaelis-Menten rate constants and diffusion rates are modified such that

$$k_{i,new} = k_{i,previous} \cdot Q_{10}^{\frac{T_{new}-T_{previous}}{10}}, \tag{8.1}$$

where $k_{i,new}$ is the new rate constant for reaction i at the new temperature T_{new}, $k_{i,previous}$ is the rate constant for reaction i at temperature $T_{previous}$, and Q_{10} is the temperature coefficient, defined as the ratio of reaction rates measured for the same reaction at two temperatures 10 °C apart.

Figure 8.4 illustrates how different Q_{10} values have differing effects on reaction rates. For $Q_{10} = 1$ temperature has no effect on the ratio of reaction rates at the two temperatures, whilst $Q_{10} > 1$ has increasing temperature increase the reaction rate ratio and decreasing temperature decrease the ratio and $Q_{10} < 1$ does the inverse, with increasing temperature decreasing the reaction rate ratio and decreasing temperature increasing it. We expect that Q_{10} will be between 2 and 3 in healthy brains [Reyes et al., 2008], as this is the typical value for organic reactions, but with some of the piglets showing the opposite behaviour to what is expected during cooling and all having suffered a severe injury to the brain, an initial prior range $1 < Q_{10} < 5$ has been set, with the potential to consider $Q_{10} < 1$ should initial explorations suggest this. We add this modification to the model by adding a temporary variable Q_{temp} based on equation 8.2.1 such that

$$Q_{temp} = Q_{10}^{\frac{T_{new}-T_{previous}}{10}} \tag{8.2}$$

and therefore $k_{i,new} = k_{i,previous} \cdot Q_{temp}$.

These additions to the model therefore produce the new model structure seen in

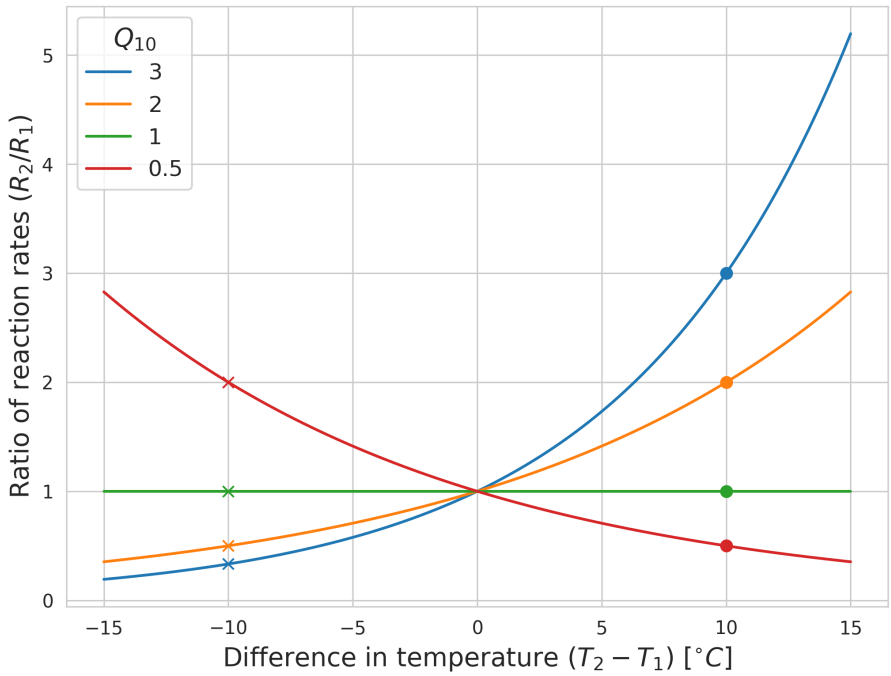

Figure 8.4: Graph showing the impact of Q_10 on reaction rates and how the change in temperature affects the ratio for different values. In particular, it is worth noting that a Q_{10} value between 0 and 1 has the inverse affect of a value greater than 1. For a cooling to 10 °C lower than the starting temperature, the reaction rate doubles for $Q_{10} = 0.5$ whilst for $Q_{10} = 2$, the reaction rate doubles for an increase in temperature to 10 °C *above* the starting temperature. This may prove important when we look at the piglets where their metabolic response is inverted compared to what is expected during the process of cooling.

figure 8.5.

8.2.2 BP Hypothermia 2

BP Hypothermia 2 (BPH2) is based closely on BPH1 but with one small but important change. The decision was made to split the effect of temperature on the reaction rates of diffusion/haemodynamics and metabolism into two separate equations with two separate Q_{10} values. This allows for tighter control over the subtleties of hypothermia and how it affects the two compartments inside the model, without adding the unmanageable levels of complexity that allowing each equation to have its own Q_{10} would entail.

Both equations have the exact same form as in BPH1 but with individual Q_{10} values.

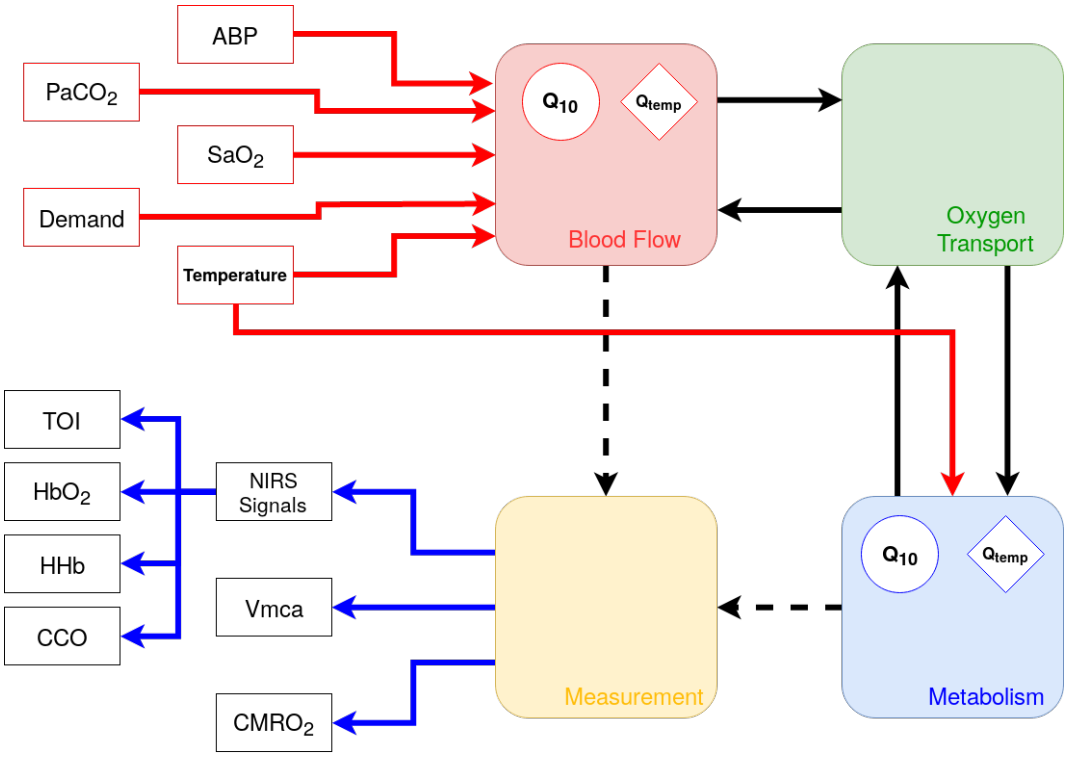

Figure 8.5: Structure of the BPH1 model. Additions are shown in bold, with parameters shown in a circle and temporary variables shown in a diamond. In this case, Q_{10} and Q_{temp} in the Metabolism and Blood Flow compartments refer to the same parameter and variable.

$$k_{i,new} = k_{i,previous} \cdot Q_{10,haemo}^{\frac{T_{new}-T_{prev}}{10}} \tag{8.3}$$

$$k_{j,new} = k_{j,previous} \cdot Q_{10,met}^{\frac{T_{new}-T_{prev}}{10}} \tag{8.4}$$

where all model terms are the same as in equation 8.2.1, but with separation of reactions such that $Q_{10,haemo}$ is the temperature co-efficient for the reaction i within the oxygen transport and blood flow compartments and $Q_{10,met}$ is the temperature co-efficient for reaction j within the metabolic submodel.

These additions to the model therefore produce the new model structure seen in figure 8.6.

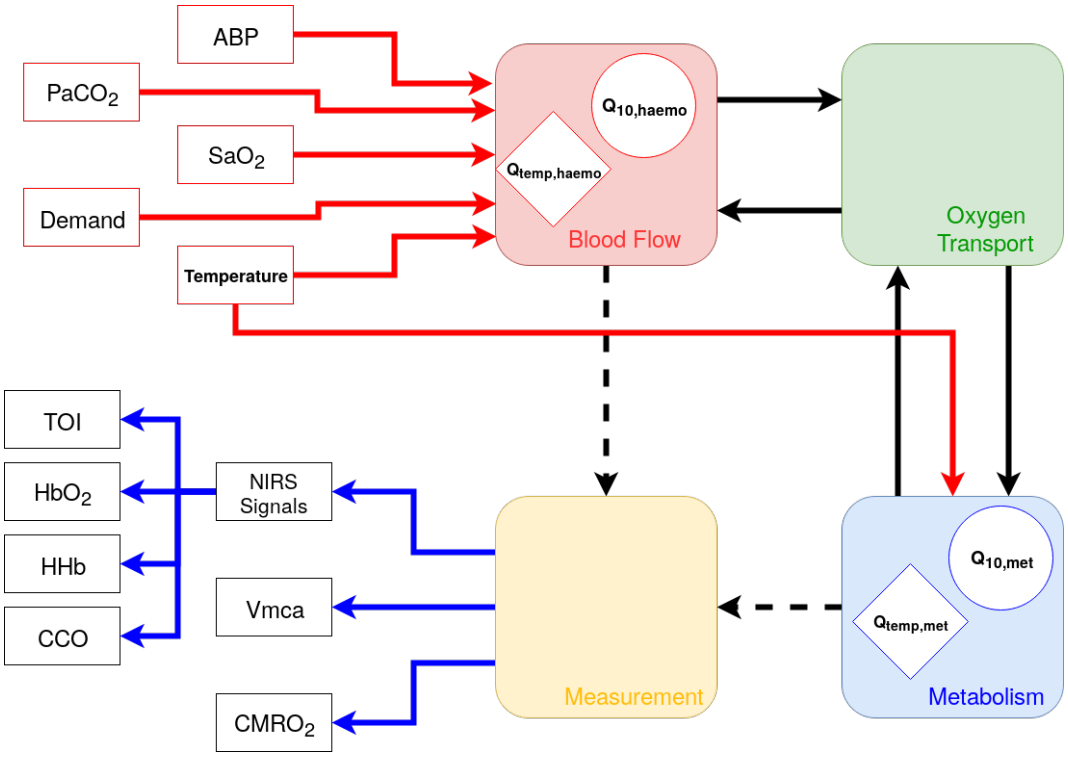

Figure 8.6: Structure of the BPH2 model. Additions are shown in bold, with parameters shown in a circle and temporary variables shown in a diamond. The separation of haemodynamic and metabolic reaction rate modification is clearly seen here.

8.2.3 BP Hypothermia 3

The final model variant, BP Hypothermia 3 (BPH3) is similar to BPH2 in that haemodynamics and metabolic reaction rates are controlled separately. However, since the haemodynamic compartment has only a single reaction that is controlled directly by the reaction rate modifications in BPH1 and BPH2, it was decided that it would be best to try a different approach here to see how this affected the model outcome and results. With this in mind, the metabolic effect is the same as in equation 8.2.2 of BPH2, but the diffusion reaction is instead modified such that

$$D_{O_2,new} = k_{diff} \cdot D_{O_2,prev} \tag{8.5}$$

$$k_{diff} = q_{diff} \cdot \frac{T_{new}}{T_{prev}} \tag{8.6}$$

where q_{diff} is a dimensionless scaling parameter and $k_{diff} > 0$. In doing this the diffusion rate is scaled directly by the change in temperature. This approach is not

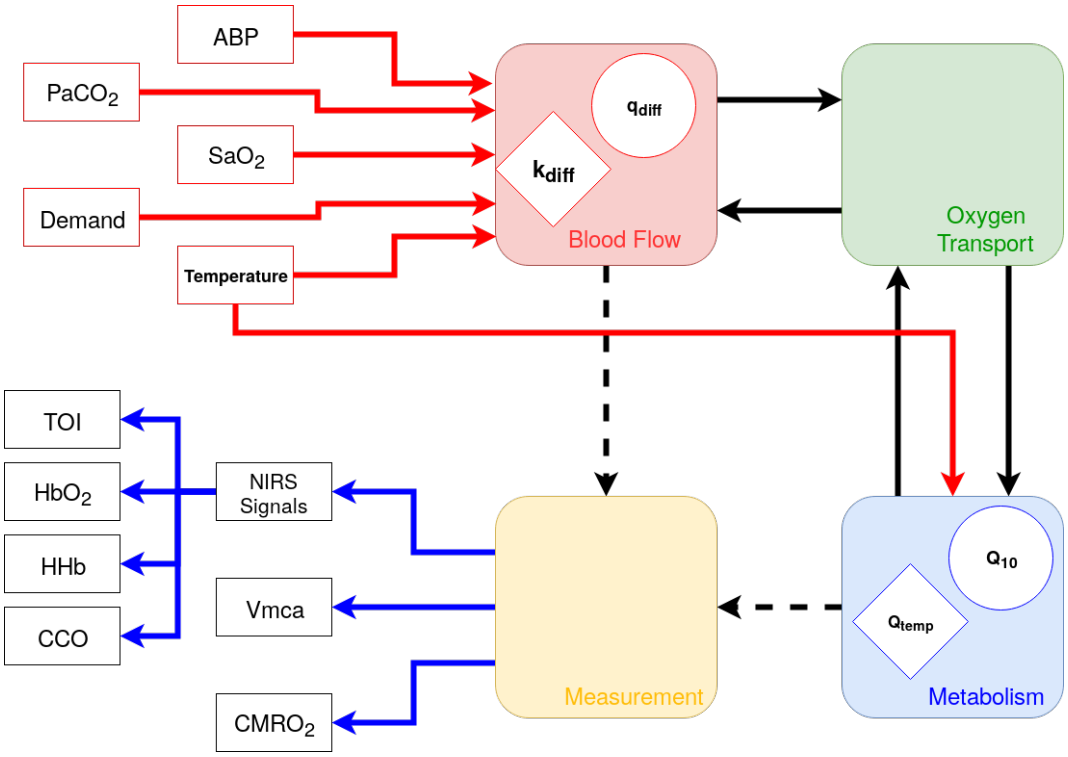

Figure 8.7: Structure of the BPH3 model. Additions are shown in bold, with parameters shown in a circle and temporary variables shown in a diamond.

based on any particular previous approach, but is instead simply a case of applying a constant scaling factor to determine the impact this has. This is arguably the simplest approach possible but is likely to lose much of the nuance captured by the previous two approaches.

These additions to the model therefore produce the new model structure seen in figure 8.7.

8.3 Initial Model Analysis

In ordered to assess the potential viability of each model we can perform some simple checks, looking at steady state behaviour of the models under different conditions. In particular we will focus on autoregulatory behaviour during changes to system input variables and model output behaviour during cooling.

There is some work within the literature which will allow us to qualitatively assess the models based on their responses during these checks. Lee et al. [2011] investigated the effect of cardiac arrest and hypothermia on autoregulation of cerebral

blood flow in neonatal piglets. In this work they found that hypothermia did not change autoregulation for both hyper- and hypotensive changes. As a result, if a model is to be considered a reasonable candidate, we should expect to find minimal to no difference in the autoregulation curves for increasing and decreasing blood pressure when looking at hypothermic temperatures. In terms of considering output response to temperature changes, Ehrlich et al. [2002] looked at the effect of temperature on cerebral blood flow and metabolism in the pig model. In this work, they cooled pigs down to 8 °C, with cerebral blood flow, cerebral oxygen consumption and cerebral vascular resistance determined at 37 °C, 28 °C, 18 °C and 8 °C. They found that both CBF and $CMRO_2$ decrease with decreasing temperature for a $Q10$ value of 2.46.

It should be noted here that the BrainSignals and BrainPiglet models are only expected to operate down to typical therapeutic hypothermia temperatures of 33.5 °C and so cooling will not be investigated as low as the temperatures in the work by Ehrlich et al.. Additionally, quantitative agreement by the models with collected data typically requires fitting of parameters which is not possible here. Regardless, we should expect any good candidate model to be able to qualitatively reproduce their findings and so we will expect to find that both CBF and $CMRO_2$ decrease during cooling within the model.

8.3.1 Autoregulation

In order to assess autoregulatory behaviour, each model was driven by one of the three main systemic input quantities - arterial blood pressure, arterial oxygen saturation and partial pressure of CO_2 - whilst the other two are kept constant. Fifty equally sized steps were defined and the model was run for 100 seconds at each step. The cerebral blood flow value at the end of each 100 second interval was recorded and used for analysis. This process was done for temperatures of 37 °C, 35 °C and 33.5 °C in order to assess behaviour across temperatures.

Figures 8.8, 8.9 and 8.10 shows the autoregulation curves of models BPH1, BPH2 and BPH3 respectively. In these we see the response of cerebral blood flow to changes in each of the three systemic input quantities at three different tempera-

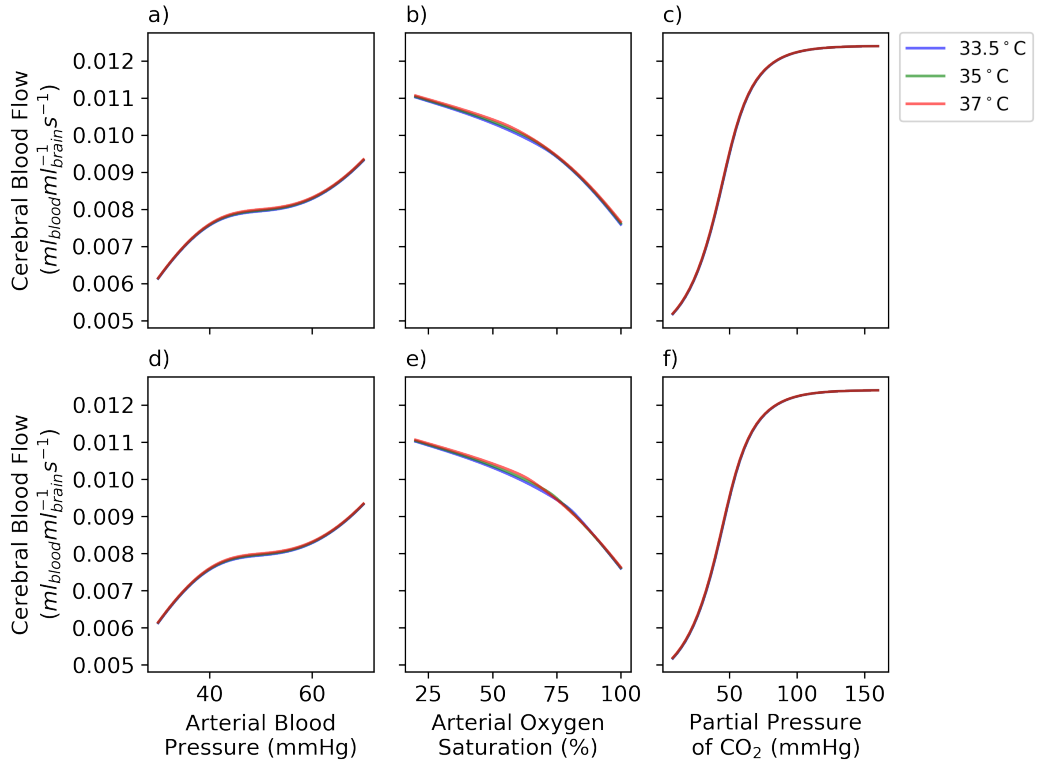

Figure 8.8: Autoregulation curves for each measurable input quantity to model BPH1 - arterial blood pressure, arterial oxygen saturation and partial pressure of CO_2 - at three distinct temperatures. Figures a)-c) show the autoregulation response with decreasing input quantity and figures d)-f) show the autoregulation response for increasing input.

tures. Subfigures a)-c) show the response curves for model inputs decreasing from a maximum value to a minimum value, whilst subfigures d)-f) show the response curves for the opposite. This is to determine the presence of hysteresis within the model. Model parameters were kept at their normal, 'healthy' values in all models. In all models, we see that across all three temperatures, which includes a hypothermic temperature of 33.5 °C, autoregulatory behaviour is maintained. In particular we see that autoregulation is maintained at hypothermic temperatures for both hyper- and hypo-tensive changes in blood pressure, as found by Lee et al. [2011]. Autoregulation curves for models BPH1 and BPH2 are identical for the default, 'healthy' parameter combinations. This is due to model BPH2 being a submodel of BPH1 such that when $Q_{10,met} \equiv Q_{10,haemo}$, the model structure is identical. Autoregulation in model BPH3 also appears to be very similar to that in models BPH1

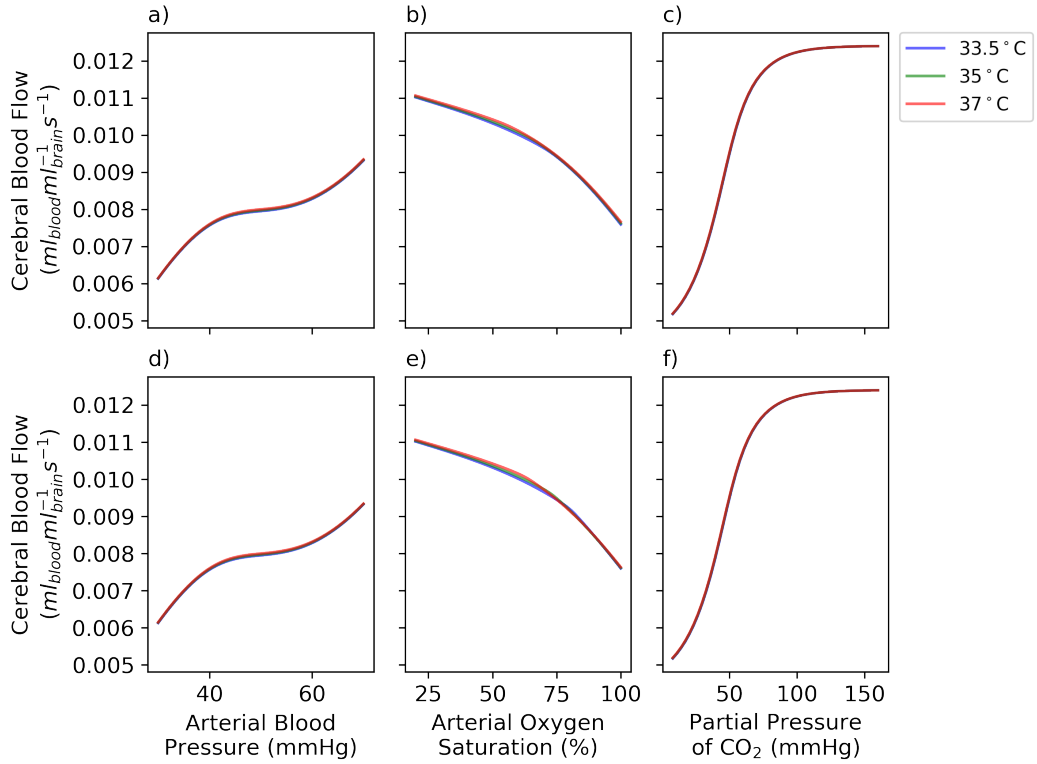

Figure 8.9: Autoregulation curves for each measurable input quantity to model BPH2- arterial blood pressure, arterial oxygen saturation and partial pressure of CO_2 - at three distinct temperatures. Figures a)-c) show the autoregulation response with decreasing input quantity and figures d)-f) show the autoregulation response for increasing input.

and BPH2, with the most notable difference being in the response of arterial oxygen saturation at each temperature. For model BPH1 and BPH2 there is a slight but noticeable difference to the response curve at an oxygen saturation of around 50-75% at each temperature, but for model BPH3 this difference is negligibly small.

In conclusion, the autoregulatory response of each model agrees with the available literature in so far as this is possible, particularly with regards to the maintaining of autoregulation at hypothermia temperatures during both hyper-and hypo-tensive changes. As a result we can say that all models can be considered good candidates based solely on this initial exploration. Further investigation requires analysis and observation of the response of different outputs to changes in temperature.

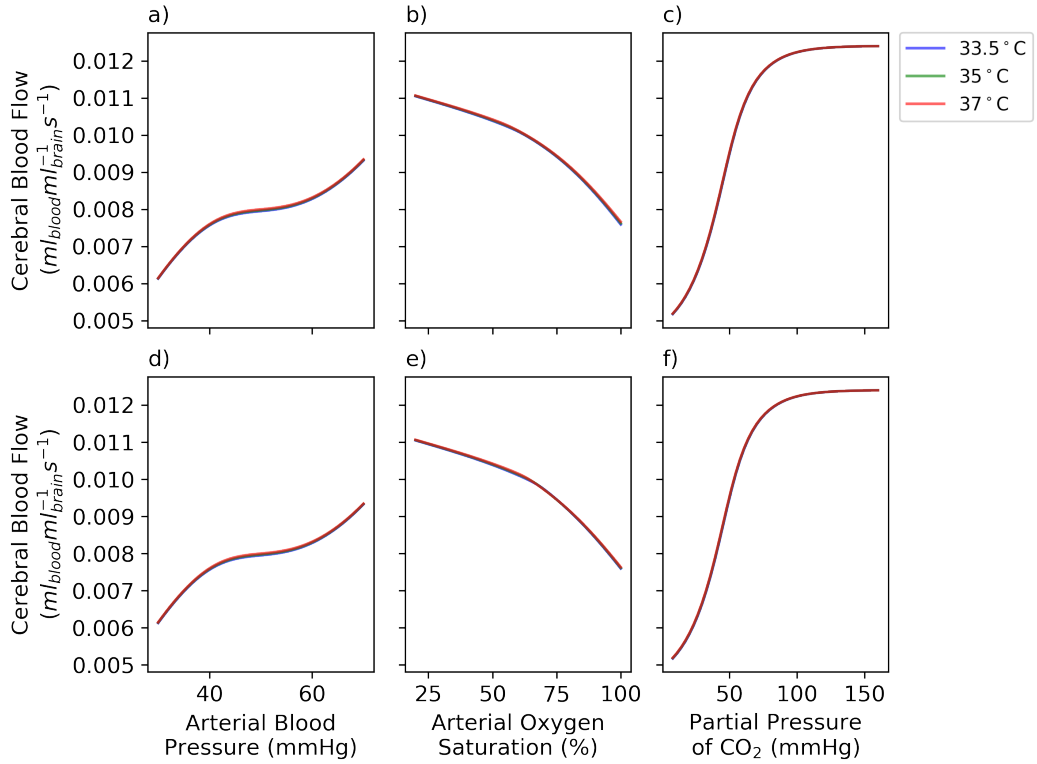

Figure 8.10: Autoregulation curves for each measurable input quantity to model BPH3 - arterial blood pressure, arterial oxygen saturation and partial pressure of CO_2 - at three distinct temperatures. Figures a)-c) show the autoregulation response with decreasing input quantity and figures d)-f) show the autoregulation response for increasing input.

8.3.2 Output response to temperature

To investigate the effect of cooling on model output, the model was run with the three main systemic variables kept at their default values and with temperature reduced gradually from 37 °C to 33.5 °C. To consider the full effect of temperature on model outputs, this process was repeated for a range of parameter values. For BPH1, with only one new parameter added to the model, Q_{10} was varied between 0.1 and 5.0 in steps of 0.1 leading to a total of 50 runs. For model BPH2, both $Q_{10,met}$ and $Q_{10,haemo}$ were varied jointly between 0.1 and 5.0 in steps of 0.1, leading to a total of 2500 runs. Finally, for model BPH3 Q_{10} was varied between 0.1 and 5.0 in steps of 0.1 and q_{diff} was varied between 0.01 and 1.0 in steps of 0.02, leading to a total of 2250 runs.

To visualise this information two approaches are taken. For model BPH1, where

only one model parameter is varied, we can visualise the response of an output across the full range of temperatures for each parameter value. For models BPH2 and BPH3 however, which each have two new parameters to consider, we instead visualise the output value at 33.5 °C only. This is plotted on a heatmap for each parameter combination. Where there are results requiring further investigation, we can visualise output in detail across the full range of temperatures by fixing one parameter at a constant value and varying the other.

Figure 8.11 shows the responses of CBF, $CMRO_2$, ΔoxCCO, ΔHbD and TOI to cooling for a range of Q_{10} values. The color map has been chosen to aid in distinguishing between general ranges of Q_{10} values.

There are a number of important features to note. First of all if we consider $Q_{10} > 1$, which is the range we would expect to find in a healthy subject, we can see a number of important behavioural trends. Both CBF and $CMRO_2$ decrease approximately linearly during cooling to as low as 98.5% and 93.5% respectively. This qualitatively matches up with the experimental observations of Ehrlich et al. [2002] even though the exact amounts don't match. This could be due to a combination of the lack of overall parameter fitting and the maintaining of constant systemic inputs which we would expect to vary with temperature also.

In addition to this confirmation of behaviour, we also gain some useful insights into the behaviour of other measurable quantities. Firstly, both TOI and ΔHbD increase during cooling, suggesting an increase in oxygenation *if systemic signals stay constant*. Secondly, in line with the $CMRO_2$ response, we see a decrease in ΔoxCCO during cooling, indicating a decrease in metabolic activity. These results are in line with the behaviour seen in Piglet LWP475 which suffered a mild injury, as outlined in Piglet Data.

For extremely low Q_{10} values, particularly those at the bottom end of the considered scale at about $Q_{10} = 0.1$, behaviour between 37 °C and 35 °C is opposite to behaviour between 35 °C and 33.5 °C. For example, cerebral blood flow increases when cooling from 37 °C to 35 °C for $Q_{10} = 0.1$ but then decreases when cooling further temperatures less than 35 °C. We wouldn't expect to find these parameter

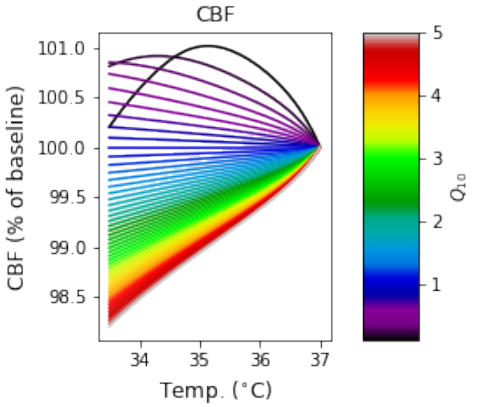

(a) Cerebral blood flow output

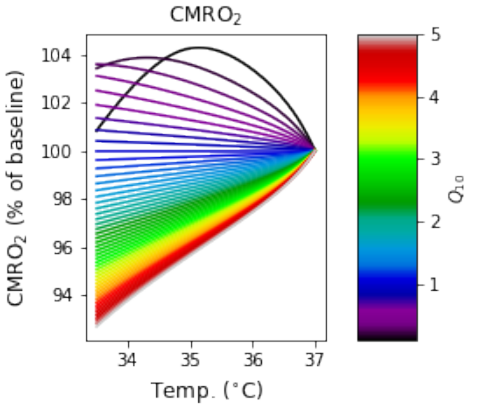

(b) CMRO$_2$ output

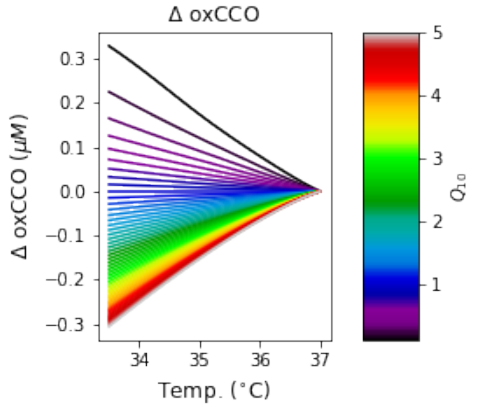

(c) Cytochrome-*c*-oxidase output

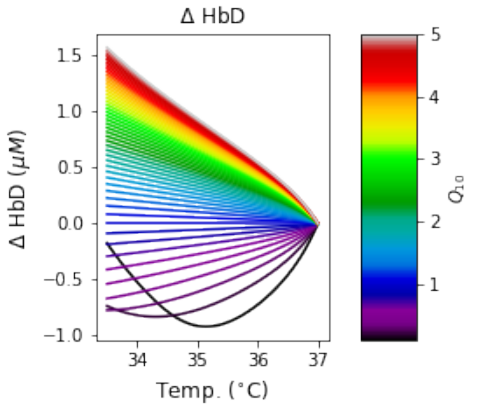

(d) Haemoglobin difference output

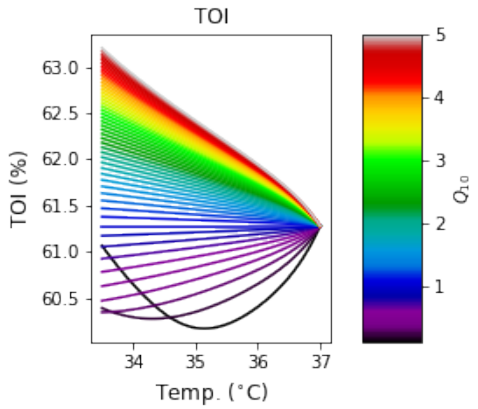

(e) TOI output

Figure 8.11: Output response curves for important measurable outputs.

values in a healthy subject, but it is important to consider them and their behaviour as they may be present in severely injured neonates. In particular, the general inversion of the expected response at these lower values seems in line with the data collected from Piglet LWP479 which suffered a severe injury. This observation will prove useful when determining the prior ranges to be used for Bayesian analysis in this piglet.

Figure 8.12 show the responses of model BPH2 outputs to cooling, with output values at 33.5 °C given for each parameter combination. General behaviour of the model is similar to that seen in BPH1 but with some noticable differences. Figures 8.12a and 8.12b show that cooling for the healthy parameter ranges of $Q_{10,haemo} > 1$ and $Q_{10,met} > 1$ results in reductions in both outputs. Additionally, the majority of the impact seems to come from changes in $Q_{10,met}$, with the colour change due to changes in $Q_{10,haemo}$ being much less drastic and noticeable.

Looking at the NIRS outputs of ΔoxCCO, ΔHbD and TOI we see a similar pattern, with most of the change seeming to come from changing $Q_{10,met}$. However, for ΔoxCCO we also see that for values of $Q_{10,met} < 1$, changing $Q_{10,haemo}$ has a significant effect, with $Q_{10,haemo} >\sim 3.5$ causing the signal to decrease rather than increase. This is seen even more clearly in figure 8.13, where ΔoxCCO has been plotted against $Q_{10,haemo}$ for specific values of $Q_{10,met}$. For $Q_{10,met} = 0.1$ a clear decrease in the value of ΔoxCCO is visible when $Q_{10,haemo} >\sim 3.5$. This behaviour suggests that careful fitting of both parameters will be required, as the effect of one can easily cancel out that of the other.

Figure 8.14 show the responses of model BPH3 outputs to cooling, with output values at 33.5 °C given for each parameter combination. General behaviour of the model is noticeably different to that of models BPH1 and BPH2, with distinctly non-monotonic behaviour observed when increasing q_{diff} in the range 0.1 to 1.0. Figures 8.14a and 8.14b show that for $Q_{diff} <\sim 0.52$ the desired behaviour of cooling causing a decrease in each signal is observed. However, for $Q_{diff} >\sim 0.52$ we instead observe an increase, with this most noticeable at $Q_{10} < 1$, as was the case in Figure 8.13.

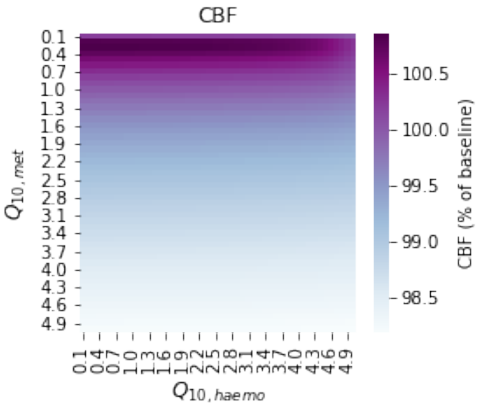

(a) Cerebral blood flow output

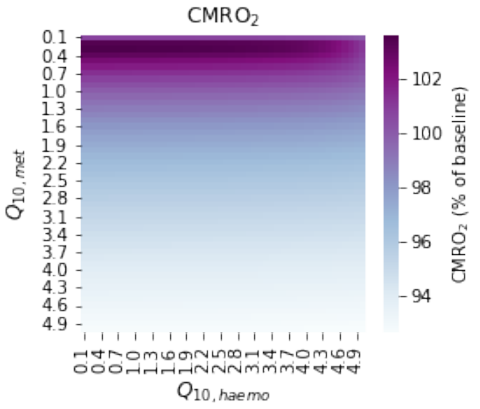

(b) CMRO$_2$ output

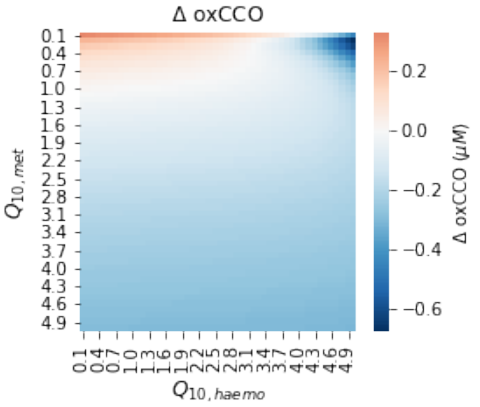

(c) Cytochrome-*c*-oxidase output

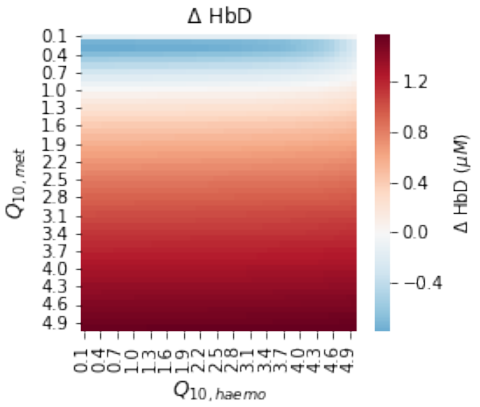

(d) Haemoglobin difference output

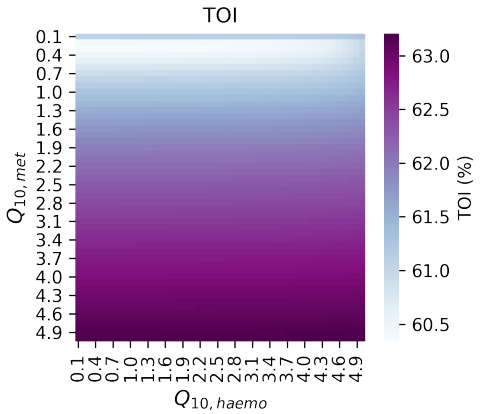

(e) Tissue oxygen saturation output

Figure 8.12: Output response to cooling for different parameter combinations in model BPH2. Figures a)-e) show output values at 33.5 °C for different parameter combinations.

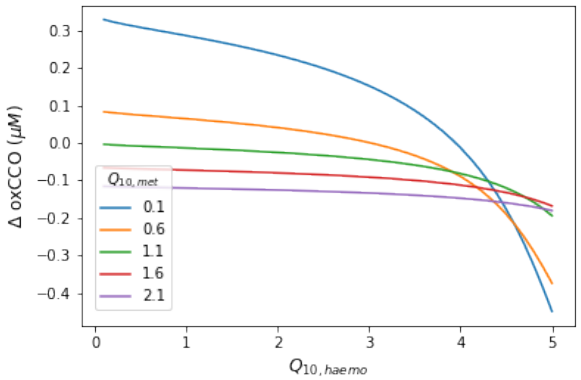

Figure 8.13: Cytochrome-*c*-oxidase output for changing $Q_{10,haemo}$. The figure shows how the value ΔoxCCO at 33.5 °C varies with $Q_{10,haemo}$ for specific $Q_{10,met}$ values. We can see clearly here that increasing $Q_{10,haemo}$ quickly causes a drop in ΔoxCCO.

Looking at the NIRS signals, we see some particularly interesting behaviour. For ΔoxCCO, in figure 8.14c, we see that for the cases where we have either $Q_{10} >\sim 2.5$ and $q_{diff} < 0.52$ or $Q_{10} < 1$ and $q_{diff} > 0.7$, ΔoxCCO increases with cooling, but for either low Q_{10} and low q_{diff} or high Q_{10} and high q_{diff} cooling causes a decrease in q_{diff}. In addition, for a specific area of parameter space, a drastic change in ΔoxCCO can be clearly seen. For $0.52 < q_{diff} < 0.7$ and $Q_{10} < 2.5$, cooling causes a drop in ΔoxCCO much larger than seen outside this region. This is clearly visible in Figure 8.15a, where for $Q_{10} = 0.1$ increasing q_{diff} first causes a gradual and then sudden decrease in ΔoxCCO until $q_{diff} \sim 0.6$, after which ΔoxCCO increases above values seen for larger Q_{10}. In contrast, for $Q_{10} = 4.1$ ΔoxCCO gradually increases as q_{diff} is increased from 0.1 to ~ 0.5 after which there is a decrease between $q_{diff} = 0.5$ and $q_{diff} = 0.6$, after which ΔoxCCO is approximately constant. Q_{10} values between these two extremes show a combination of these behaviours, with the sudden decrease and increase being more prominent as Q_{10} decreases.

For ΔHbD, in figure 8.14d, we see some similarly interesting behaviour. Notably we see that cooling generally causes an increase in ΔHbD apart from in the region of parameter space bounded by approximately $Q_{10} \in [0.1, 1.0]$ and $q_{diff} \in [0.55, 1.0]$, where cooling causes a decrease. Figure 8.15b shows how the value ΔHbD at

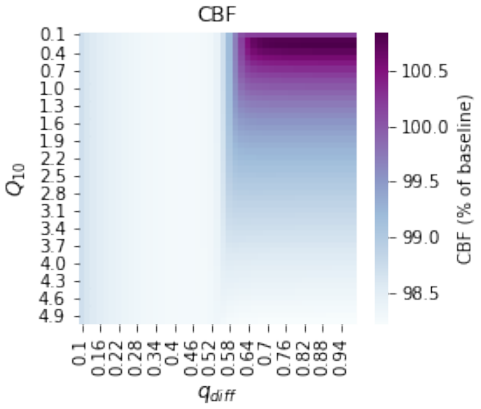

(a) Cerebral blood flow output

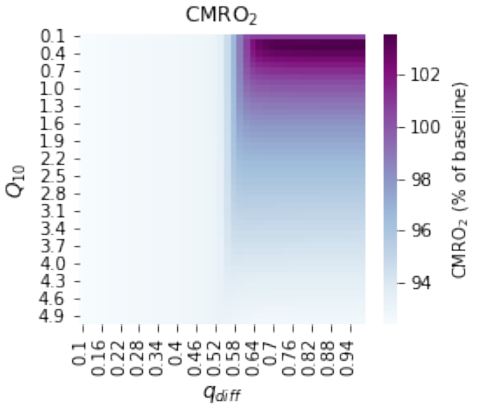

(b) $CMRO_2$ output

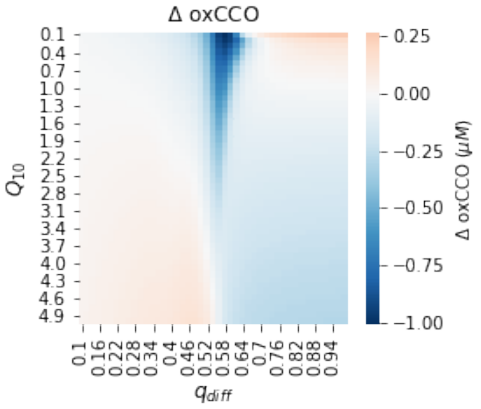

(c) Cytochrome-*c*-oxidase output

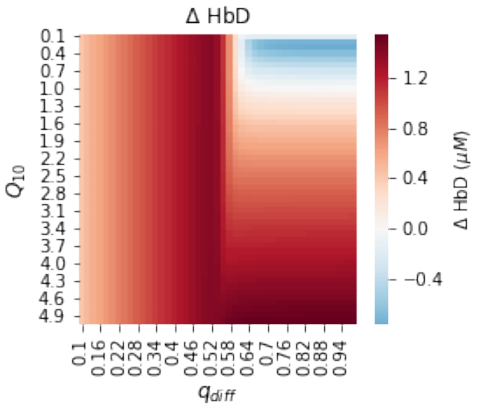

(d) Haemoglobin difference output

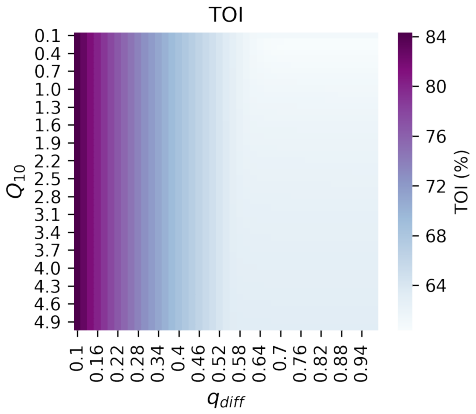

(e) Tissue oxygen saturation output

Figure 8.14: Output response curves for important measurable outputs.

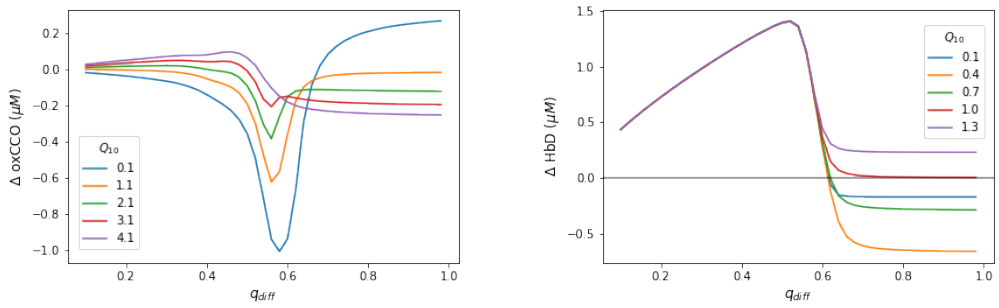

(a) Tissue oxygen saturation output **(b)** Haemoglobin difference output

Figure 8.15: ΔHbD and ΔoxCCO output for changing q_{diff}. Figure 8.15a shows how the value ΔoxCCO at 33.5 °C varies with q_diff for specific Q_{10} values. We can see clearly here that increasing q_{diff} has a varying effect on ΔoxCCO depending on the Q_{10} value. For lower values of Q_{10}, increasing q_{diff} causes first a decrease and then an increase in ΔoxCCO, whilst for higher values of Q_{10} increasing q_{diff} causes a gradual increase followed by a sharp decrease and than an approximately constant value. Figure 8.15b shows how the value ΔHbD at 33.5 °C varies with q_diff for specific Q_{10} values.
Figure 8.15b shows how the value ΔHbD at 33.5 °C varies with q_diff for specific Q_{10} values. Here we can see clearly that behaviour between different values of Q_{10} diverges for $q_{diff} >\sim 0.6$.

33.5 °C varies with q_{diff} for specific Q_{10} values. Here we can see clearly that behaviour between different values of Q_{10} diverges for $q_{diff} >\sim 0.6$, with $Q_{10} > 1$ having ΔHbD > 0 and $Q_{10} < 1$ having ΔHbD < 0.

Finally, the results for TOI are much simpler than those of the other considered outputs. Varying Q_{10} appears to have minimal impact on TOI value, whist increasing Q_{diff} causes TOI to decrease fairly rapidly at first before levelling off to being approximately constant for $q_{diff} > 0.5$.

It is clear from this initial analysis of model BPH3 that the choice of priors for these parameters is particularly important, with the model output being notably non-monotonic across parameter space and sharp changes in behaviour occurring at specific boundary values. We might expect that we will find posterior distributions with small variance, especially if parameters fall within one of the tightly bounded regions, such as that seen in the top right corner of figure 8.14d.

Having now considered our model behaviour and determined that for all three models we are able to qualitatively agree with experimentally determined behaviour, al-

beit within certain parameter boundaries, we can now move on towards using these models to analyse our data. This process has been outlined in Development of a Bayesian framework for model analysis in Figure 6.3, although a number of small changes have been made to the sensitivity analysis process as outlined below.

8.4 Bayesian Modelling

Including the original BrainPiglet 2.1 model, a total of four BrainPiglet model variants were analysed and compared for this work. Additionally, small variations on these models were also considered for dataset LWP479 with different hyperparameters (parameters describing the distribution, e.g. the mean or standard deviation) placed on the prior distributions. This is because initial hyperparameter settings were determined based on the assumption of the expected physiological response to hypothermia. However, initial analysis of dataset LWP479 highlighted a flaw in this assumption: the physiological response of the more severely injured piglet is inverted as compared to the healthy piglet, therefore assumptions that the temperature co-efficient must be greater than 1 are not valid.

As a result it was decided to perform further fitting of the model for this dataset using different hyperparameters for the prior distributions, allowing $Q_{10,\cdot} < 1$ parameters, where $Q_{10,\cdot}$ are the temperature co-efficients in each model. These submodel variants are labelled as BP Hypothermia $x.1$, where x denotes the major model version as outlined below. Information about the changed hyperparameters are mentioned within the description of the main model version. These $x.1$ variants are used below for posterior and posterior predictive distributions for piglet LWP479. The $x.0$ variants are not analysed directly for piglet LWP479 as they failed to produce meaningful results for reasons that will become clear. They are included however in model selection in order to validate the decision to exclude them from analysis.

8.4.1 Sensitivity Analysis

As explained in section 5.1, in order to successfully use the BCMD framework and model it is necessary to perform model reduction via sensitivity analysis. In this case we needed to reduce a model such that the selected parameters would be ap-

plicable to multiple datasets based on brains in distinctly different states. To ensure that this reduced model would be applicable to all piglets and their outcomes, rather than being specific to a single piglet and its brain state, sensitivity analysis results were considered jointly across all piglets . It is fair to assume that a model applied to a piglet having suffered severe HIE is likely to have different parameter sensitivities to a model applied to a piglet with a very mild HIE injury.

Sensitivity analysis was performed twice per dataset and per model - once where failed runs that result in `NaN` values for the summary metric were replaced first by 10,000,000 and then again with them replaced by 0. This was done to try and avoid the problem of parameters being selected that would overwhelmingly produce many failed runs.

When using only one of the above substitutions we can run into two separate issues. When replacing `NaN` values by an extremely large value, we bias the sensitivity analysis method towards parameters that produce failed runs. This is still important however, as the case where varying a parameter can cause the model run to fail is arguably a good sign of that parameter being sensitive. It also makes it very clear where a parameter has failed many times as μ_* values tend to be extremely higher where this is the case.

When replacing `NaN` values by zero however we are able to observe different information. By doing this we can select parameters that are sensitive to variation but are less likely to cause the model runs to fail. This biases the method towards parameters that are potentially less tightly bound than in the case of substitution by a large value, but are still identified as sensitive. The downside to this method however is that parameters that are potentially important but have a very small viable parameter region tend to be ignored.

To try and combine these two approaches the following criteria were used for each, before selecting final parameters as those found to be sensitive in both. The term results here refers to the sensitivity results for a given piglet and output. The top 20 most sensitive parameters for each combination of dataset and output were selected and the union of these is the total set of results considered.

Figures 8.16 to 8.19 show heatmaps of the sensitivity analysis results. The parameters shown are the union of the twenty most sensitive parameters for each output and dataset combination. As a result, the number of parameters considered in each heatmap varies.

These heatmaps provide a general overview of what parameters were shown to be sensitive within each output and dataset. It also allows us to clearly see where specific parameters are found to be sensitive in all situations e.g. NADpool or r_t in Figure 8.16a.

There is an interesting result within Figure 8.16a here we can see that parameter Q_10_haemo is identified as being sensitive in CCO-LWP484 despite not being present within the model or its definition. Because all models call their parameter distributions from the same file, Q_10_haemo appears to have been varied by the SA method and has potentially been altered at the same time as another parameter that was important. As a result the method has erroneously identified it as important despite it having *zero effect by definition*. This is further justification for applying prior knowledge and common sense when performing model reduction. As we saw in chapter 7 and will see below in section 8.4.2, the Bayesian framework provides much more information about the sensitivity of the model to a given parameter, both in term of its joint posterior and its marginal posterior.

What is also clear within these figures is how a parameter may be important within a single output but not all, especially when looking at the metabolic subcompartment and the CCO output. For example, in every BPH model variant where substitution by zero was performed, cytox_tot_tis is identified as important in the CCO output for *all* piglets, but is rarely important in any other output. As a result, if we were to only select parameters that were important in 75% of all Output-Dataset combinations, we would inevitably bias the selection against parameters that were important in only metabolism due to only the CCO output being specific to that compartment.

After having performed this process of selecting parameters that were in 75% of all results or occurred in 3 or more datasets for a single output, the number of pa-

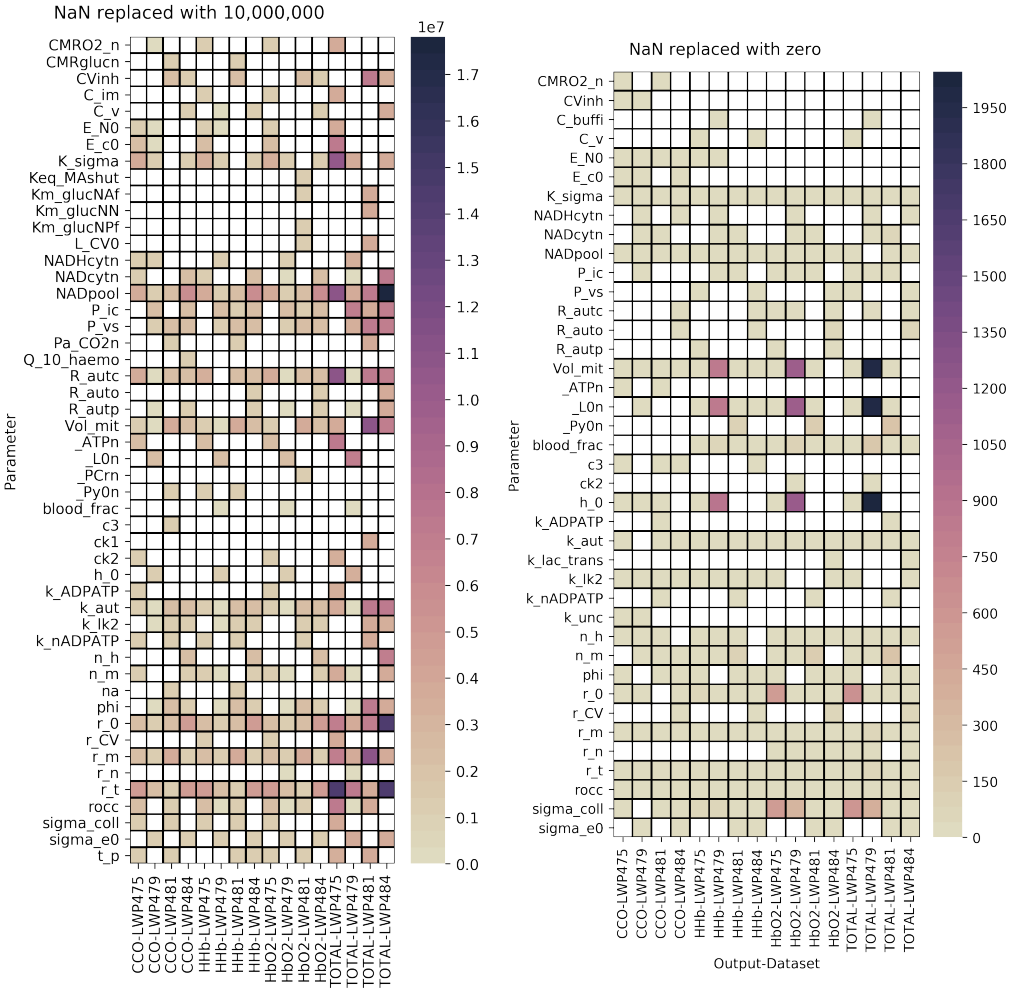

(a) SA Results for BrainPiglet2.1 with substitution by large number.

(b) SA Results for BrainPiglet2.1 with substitution by zero.

Figure 8.16: Sensitivity Analysis results for model BrainPiglet 2.1, with NaN substitution by 10,000,000 and zero. Each cell is coloured based on the μ_* value for that parameter and Output-Run combination.

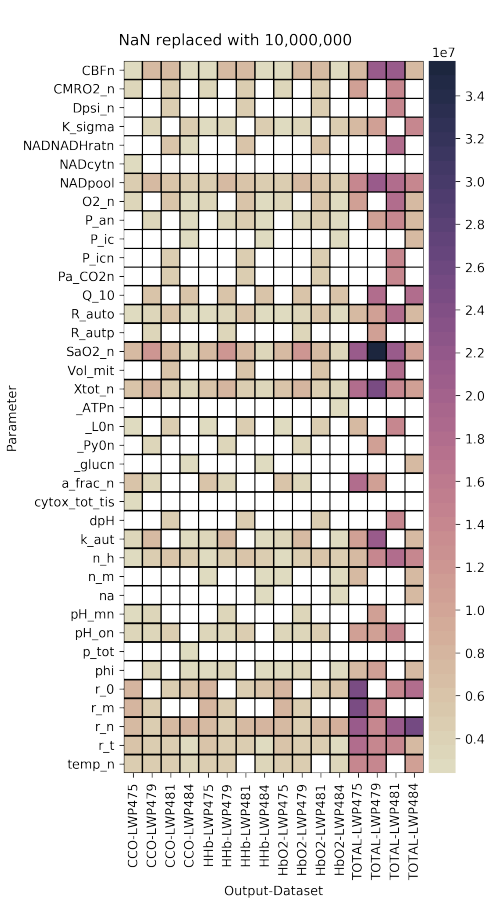

(a) SA Results for BPH1 with substitution by large number.

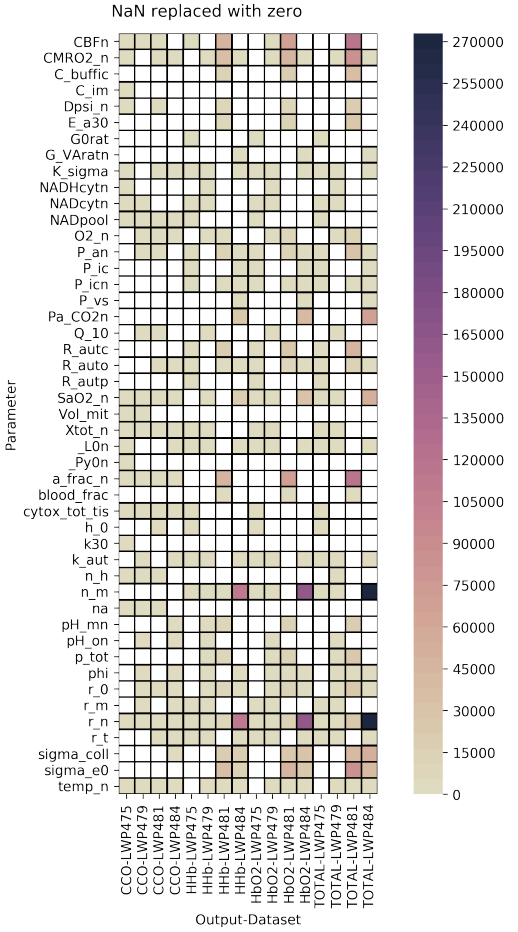

(b) SA Results for BPH1 with substitution by zero.

Figure 8.17: Sensitivity Analysis results for model BPH1, with NaN substitution by 10,000,000 and zero. Each cell is coloured based on the μ_* value for that parameter and Output-Run combination.

rameters was still relatively high. Whilst it was important to make sure we provided adequate coverage of parameter space in order to sufficiently fit the model, we could potentially reduce the number of parameters that required fitting by removing those which were unlikely to vary physiologically.

For example, changing ‘temp_n‘ clearly has a large impact on the model, but it is extremely unlikely that the normal body temperature of a piglet was significantly different from the default value of 37 °C. Similarly, the normal SaO_2 value given by SaO2_n is unlikely to be very different from the default value given that we have recordings of this value at the start of the measurement period and it is similar if not

exactly the same as the default value for most piglets.

Other parameters that may fall into this category include normal blood pressure, P_an, and normal partial pressure of CO2, PaCO2_n. These parameters are important, but when focussing on the impact an injury has had in terms of function, parameters that control what we expect a 'normal' input to be are unlikely to provide useful information when fitting. It is potentially more useful to leave many of the 'normal value' parameters at their default values. The normal value of blood pressure, for example, should probably be kept the same across all piglets. Then, if the impact of blood pressure is different from the expected 'healthy' behaviour, it should be more dynamic parameters that control these changes, rather than modifying what we consider to be a normal blood pressure to try and force a better fit despite the change being unrealistic. This appears to be an inversion of the more typical '*Occam's razor*' approach to solving a problem, in which the simplest answer is usually the correct one, but the complex nature of the models means that many potential solutions may be found to a given problem. With this in mind, a certain amount of common sense must be used when performing model reduction in order to ensure that we focus on producing physiologically reasonable results and conclusions.

Final parameter selection results are outlined below in Tables 8.2 to 8.5, along with each parameter's distribution limits and default value. All parameters are varied over a uniform distribution and parameters are ordered by how many of the sensitivity analysis results they are within the 10 most sensitive parameters, with the most present at the top and the least present at the bottom.

Parameter	Min. Value	Max. Value	Default Value
r_m	0.0216	0.0324	0.0270
rocc	0.0000	1.0000	0.0000
r_0	0.0101	0.0151	0.0126
R_autc	0.0000	10.0000	2.2000
P_ic	3.6000	5.4000	4.5000
K_sigma	8.0000	12.0000	10.0000
k_lk2	0.0304	0.0456	0.0380
r_t	0.0144	0.0216	0.0180
NADpool	1.5000	4.5000	3.0000
phi	0.0288	0.0432	0.0360
Vol_mit	0.0536	0.0804	0.0670
k_aut	0.0000	1.0000	1.0000
n_m	1.4640	2.1960	1.8300

Table 8.2: Table showing final sensitivity analysis parameters for model BrainPiglet 2.1.

Parameter	Min. Value	Max. Value	Default Value
NADpool	1.5000	4.5000	3.0000
r_t	0.0144	0.0216	0.0180
n_h	2.0000	3.0000	2.5000
Q_10	1.0000	6.0000	2.4000
r_0	0.0101	0.0151	0.0126
K_sigma	8.0000	12.0000	10.0000
k_aut	0.0000	1.0000	1.0000
R_auto	0.0000	10.0000	1.5000
r_n	0.0150	0.0224	0.0187
CBFn	0.0064	0.0096	0.0080
O2_n	0.0192	0.0288	0.0240
phi	0.0288	0.0432	0.0360
Xtot_n	3.5000	8.0000	5.4000

Table 8.3: Table showing final sensitivity analysis parameters for model BPH1.

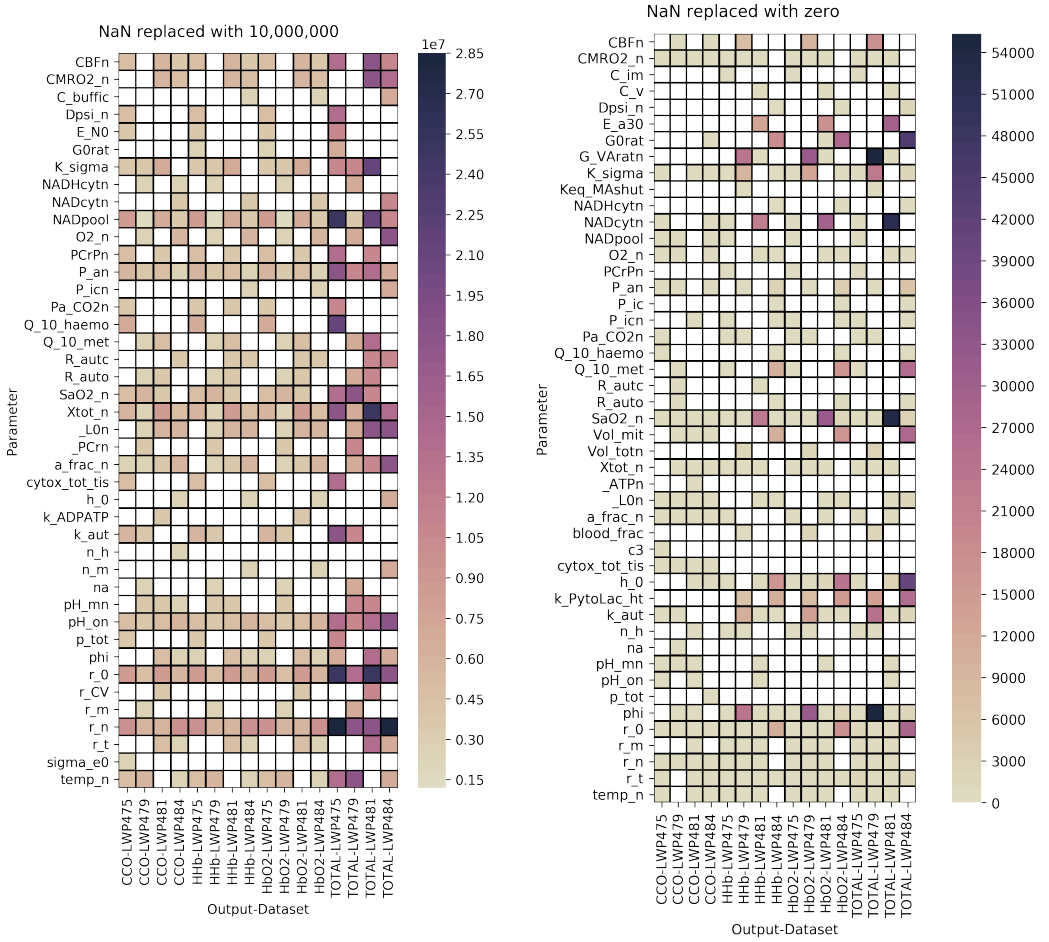

(a) SA Results for BPH2 with substitution by large number.

(b) SA Results for BPH2 with substitution by zero.

Figure 8.18: Sensitivity Analysis results for model BPH2, with NaN substitution by 10,000,000 and zero. Each cell is coloured based on the μ_* value for that parameter and Output-Run combination.

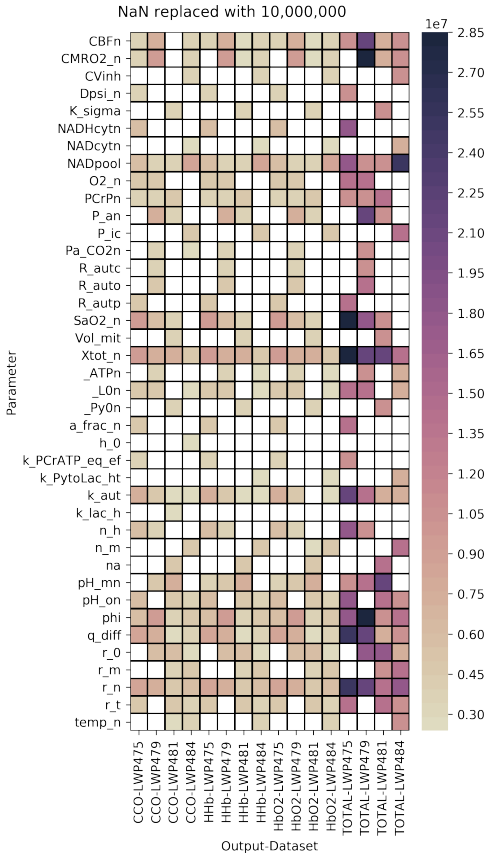

(a) SA Results for BPH3 with substitution by large number.

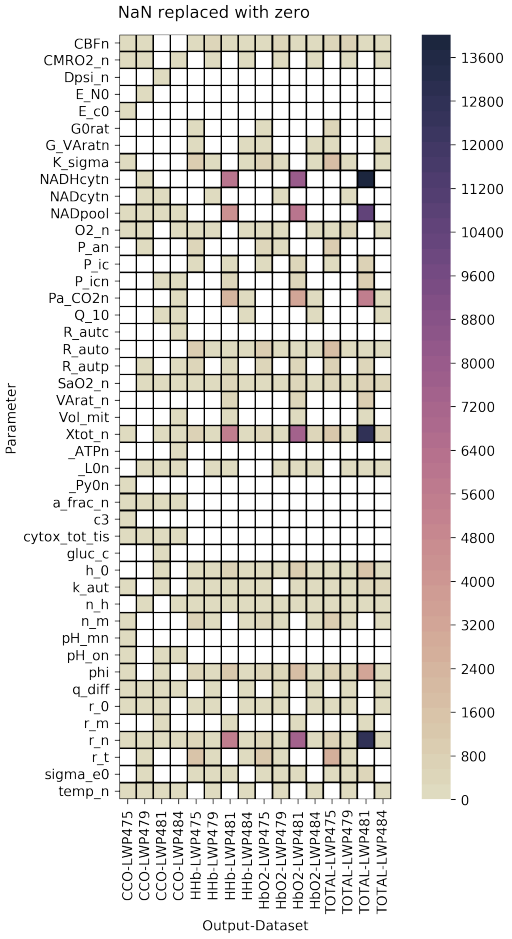

(b) SA Results for BPH3 with substitution by zero.

Figure 8.19: Sensitivity Analysis results for model BPH3, with NaN substitution by 10,000,000 and zero. Each cell is coloured based on the μ_* value for that parameter and Output-Run combination.

Parameter	Min. Value	Max. Value	Default Value
K_sigma	8.0000	12.0000	10.0000
Q_10_met	1.0000	6.0000	2.4000
a_frac_n	0.0000	1.0000	0.6700
pH_mn	7.2000	7.6000	7.4000
Q_10_haemo	1.0000	6.0000	2.4000
Xtot_n	3.5000	8.0000	5.4000
NADpool	1.5000	4.5000	3.0000
r_0	0.0101	0.0151	0.0126
r_n	0.0150	0.0224	0.0187
_L0n	0.1000	5.0000	3.0000
phi	0.0288	0.0432	0.0360

Table 8.4: Table showing final sensitivity analysis parameters for model BPH2.

Parameter	Min. Value	Max. Value	Default Value
CBFn	0.0064	0.0096	0.0080
phi	0.0288	0.0432	0.0360
NADpool	1.5000	4.5000	3.0000
q_diff	0.0000	5.0000	1.0000
_L0n	0.1000	5.0000	3.0000
CMRO2_n	0.0160	0.0240	0.0200
pH_on	6.8000	7.2000	7.0000
r_0	0.0101	0.0151	0.0126
Q_10	1.0000	6.0000	2.4000
Xtot_n	3.5000	8.0000	5.4000
k_aut	0.0000	1.0000	1.0000
r_t	0.0144	0.0216	0.0180
r_n	0.0150	0.0224	0.0187

Table 8.5: Table showing final sensitivity analysis parameters for model BPH3.

8.4.2 Model Posterior and Posterior Predictive Distributions

For each dataset and model we can generate posterior and posterior predictive distributions. These are generated for an acceptance rate of 0.1%, leading to 50,000 samples for each dataset-model combination. By comparing the posterior predictive distributions we can assess the quality of fit for each model. This is done prior to the model selection process, so the posterior and posterior predictive distributions for each model and dataset are determined independently of each other.

For each model we will first look at the posterior predictive distributions in each piglet for each model. This will allow us to consider the posterior distributions of each model in the context of how well it has performed.

We will then look at the posterior distributions for dataset LWP475, which is known to be the data from a piglet showing the expected *healthy* response. Following this we will look at the posterior distributions for dataset LWP479, which is the piglet showing the unexpected response, for each model $x.1$ variant.

Posterior Predictive Distributions

Figures 8.20 and 8.21 show the posterior predictive distributions for each dataset.

Figure 8.20 shows the posterior predictive distributions for piglet LWP475 for each output and model. BPH0 is clearly unable to properly reflect the CCO, HbO_2 and HHb signals. It does appear to be able to better reflect some of the more small scale changes than the new model variants, closely matching a number of small periodic increases, but is unable to properly reflect broader scale changes that are caused by changes in temperature.

The three new model variants appear to match the longer scale and broader changes caused by decreasing temperature and out of these, model BPH2 seems to be able to more closely match some of these broader changes. It is also generally smoother than models BPH1 and BPH3. This is notable as all three model variants appear to exhibit an inverted behaviour during the small scale periodic increases that occur throughout the signal. This suggests that the reduced form of these models appears to be unable to fully capture these short term changes and that the parameter controlling this is either missing or is damped by the added temperature changes.

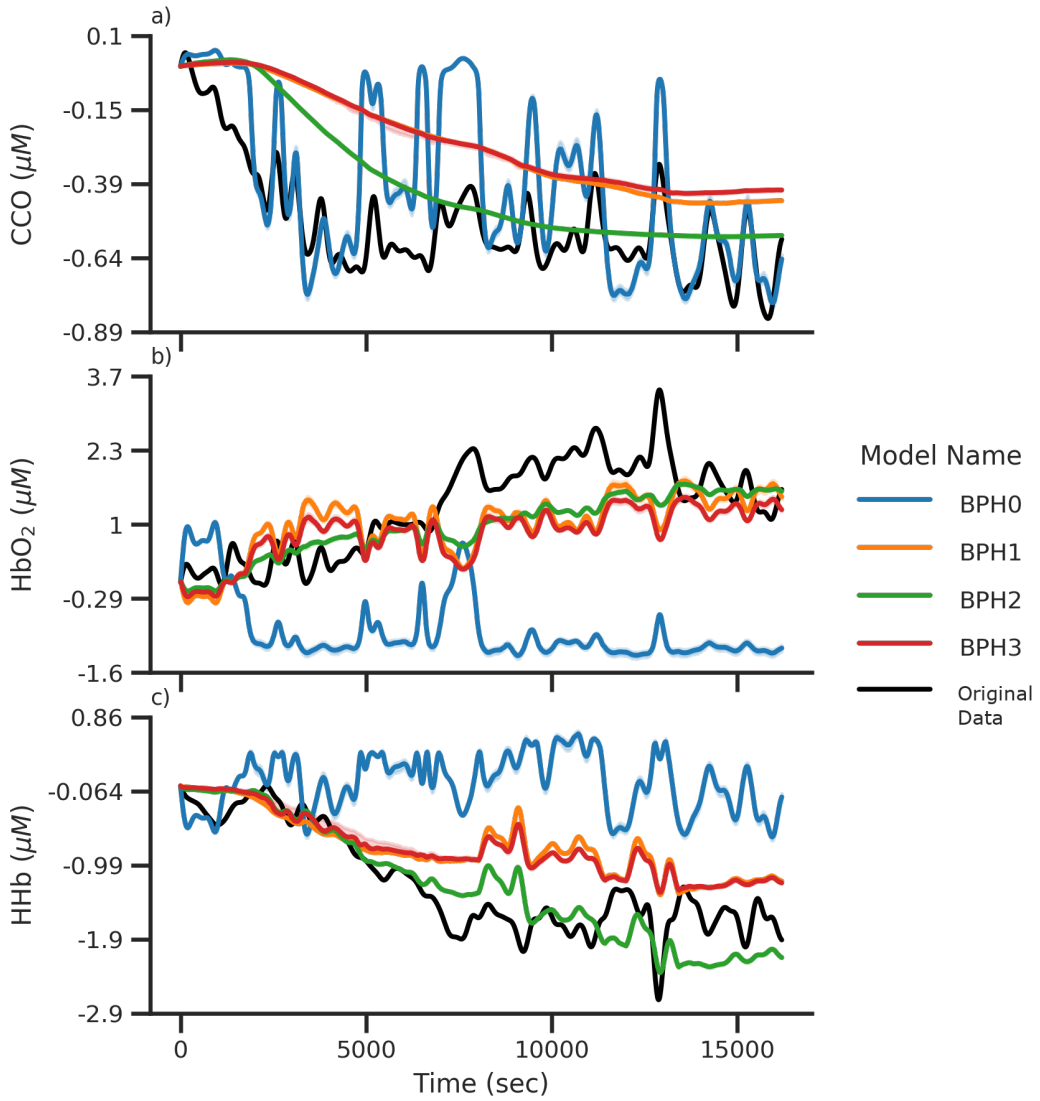

Figure 8.20: Posterior predictive distributions for piglet LWP475. Figures 8.20a, 8.20b and 8.20c show the posterior predictive distributions for CCO, HbO_2 and HHb respectively for each of the four models. 95% credible intervals are included but too small to be visible.

Figure 8.21 shows the posterior predictive distributions for piglet LWP479 for each

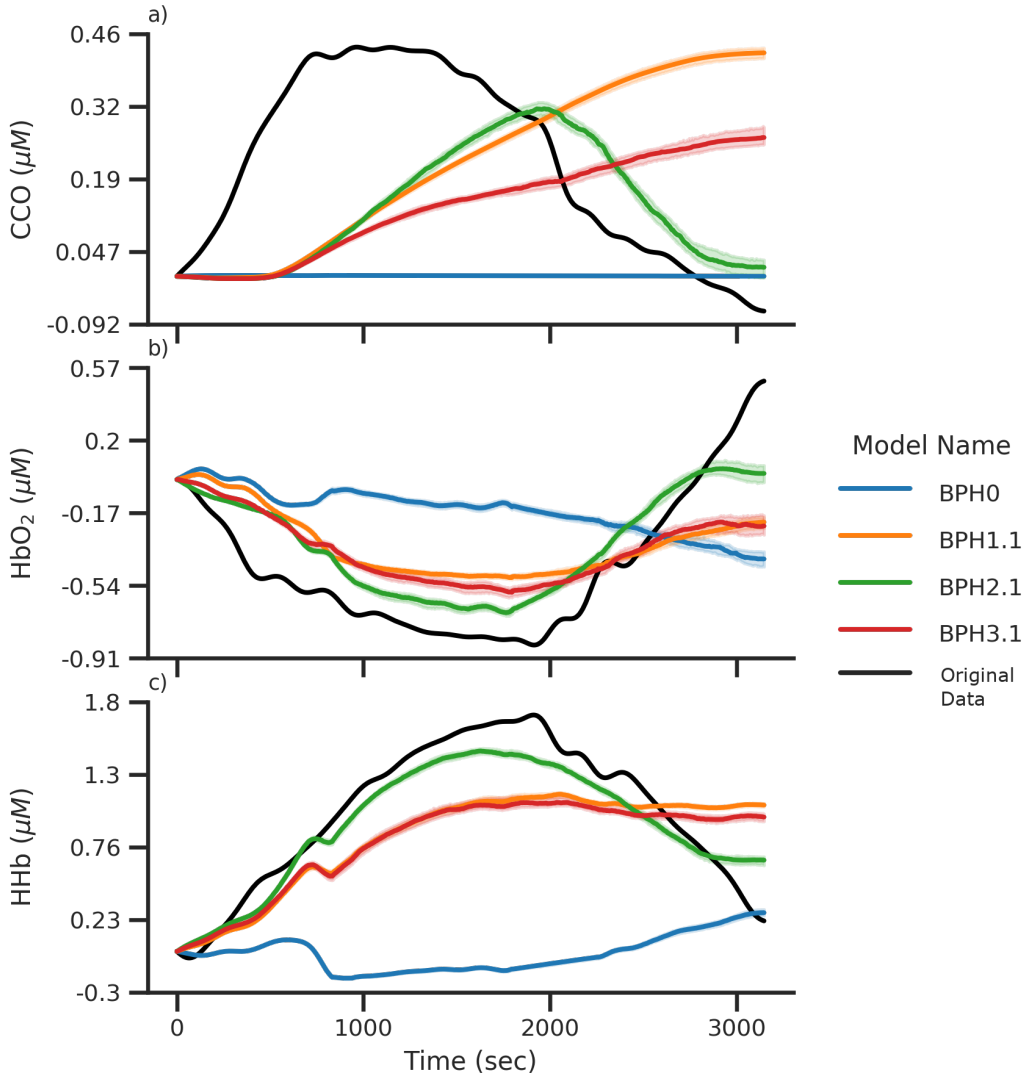

Figure 8.21: Posterior predictive distributions for piglet LWP479. Figures 8.21a,8.21b and 8.21c show the posterior predictive distributions for CCO, HbO_2 and HHb respectively for each of the four models.

output and model. A clear difference can be seen between the three new model variants that include temperature as a variable and the original unmodified model. In particular we can see that BPH0 is unable to respond to reflect the changes in CCO, HbO_2 and HHb that are caused by temperature changes. This is most obvious in the CCO response which is completely flat.

When comparing between the new model variants, in Figure 8.21a we see that by

setting a prior range to include $Q_{10,met} < 1$ in model BPH2.1, the CCO response is able to invert as seen in the data. A small amount of lag is visible in the response, suggesting that other parameters not selected by the sensitivity analysis process may be important here. Models BPH1.1 and BPH3.1, despite the prior range for Q_{10} including values less than 1, are less able to closely match the CCO response, with both increasing over time but failing to return to baseline within the time period measured. In the case of BPH1.1 this may be due to the haemodynamic and metabolic compartments sharing a temperature co-efficient.

All three model variants seem to be able to reasonably reproduce the haemodynamic measurements, with model BPH2.1 generally performing better and both model BPH1.1 and model BPH3.1 producing similar results.

BP Hypothermia 1.0/1.1 Posterior Distributions

Figure 8.22 shows the posterior distributions of models BPH1.0 and BPH1.1 with piglets LWP475 and LWP479 respectively, with piglet LWP475 in blue and piglet LWP479 in orange. Looking at the marginal posteriors we can see that only a small number of parameter posteriors differ between piglets, such as Q_10 of CBFn. Others, such as NADpool or r_n, are almost exactly the same in both piglets. We can consider the distributions in more detail by looking at them in comparison to the prior distribution alongside their K-L divergence values. Figure 8.23 gives this information for model BPH1.*. Here we can see that both k_aut and Q_10 show significant information gain as compared to the prior in both piglets, whilst R_auto only show significant information gain for piglet LWP475. Additionally, we see a number of parameters that differ from the prior by a not insignificant amount and have almost identical posteriors and K-L divergences values in both piglets. These are NADpool, r_n and r_0. These highly similar posteriors suggest a shared parameter space for the piglet models that differs from the uniform priors established here. Finally, k_sigma, phi, n_h, Xtot_n and O2_n all have posteriors that have small K-L divergence values, suggesting that minimal information is gained as compared to the prior. This suggests that the model is insensitive to changes in these parameters within the physiologically reasonable region and highlights the

difficulties faced when using the previously outlined sensitivity analysis method in identifying parameters that are actually important and/or physiologically important.

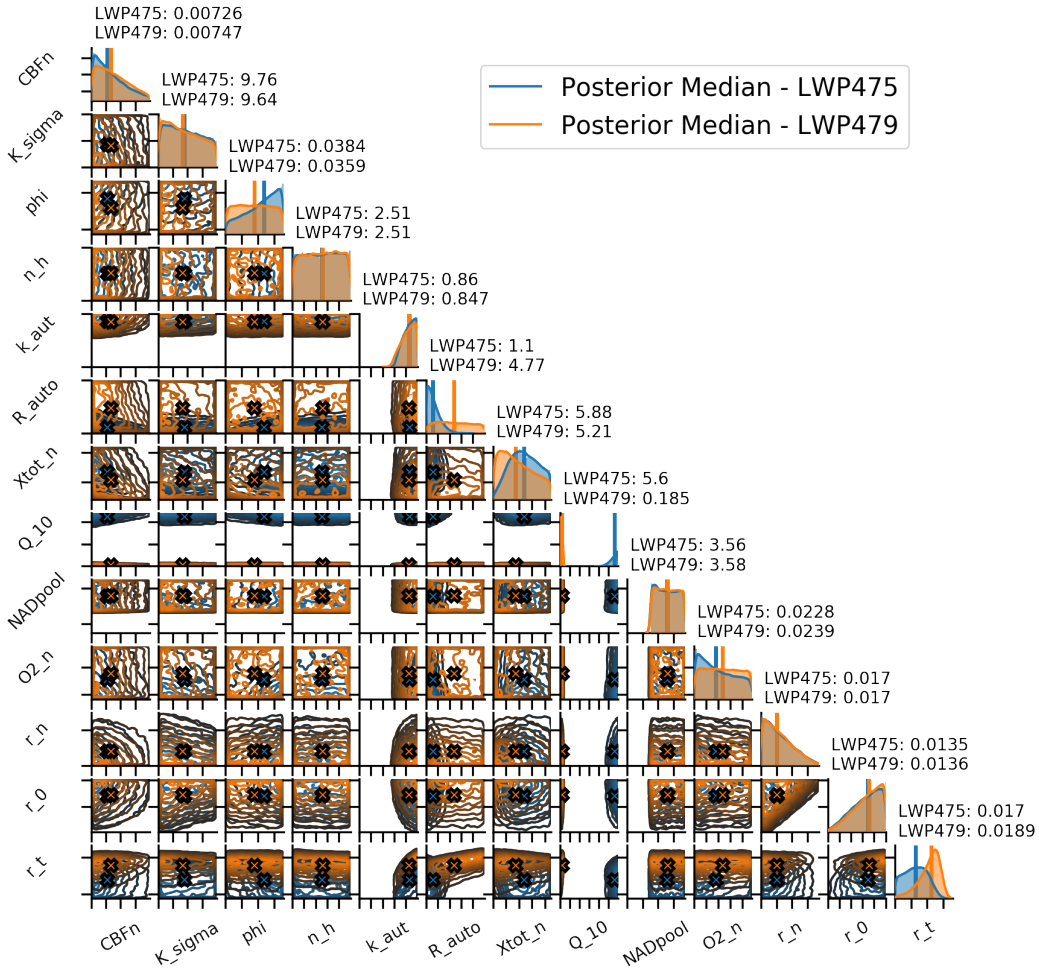

Figure 8.22: Posterior distributions using models BPH1.0/1.1 with piglets LWP475 and LWP479 respectively Posteriors for piglet LWP475 are shown in blue and those for piglet LWP479 are shown in in orange.

BP Hypothermia 2.0/2.1 Posterior Distributions

As above, Figure 8.24 shows the posterior distributions of models BPH2.0 and BPH2.1 with piglets LWP475 and LWP479 respectively, with piglet LWP475 in blue and piglet LWP479 in orange.

Looking at the marginal posteriors we again see that only a small number of parameter posteriors differ between piglets, which in this case are Q_10_met,

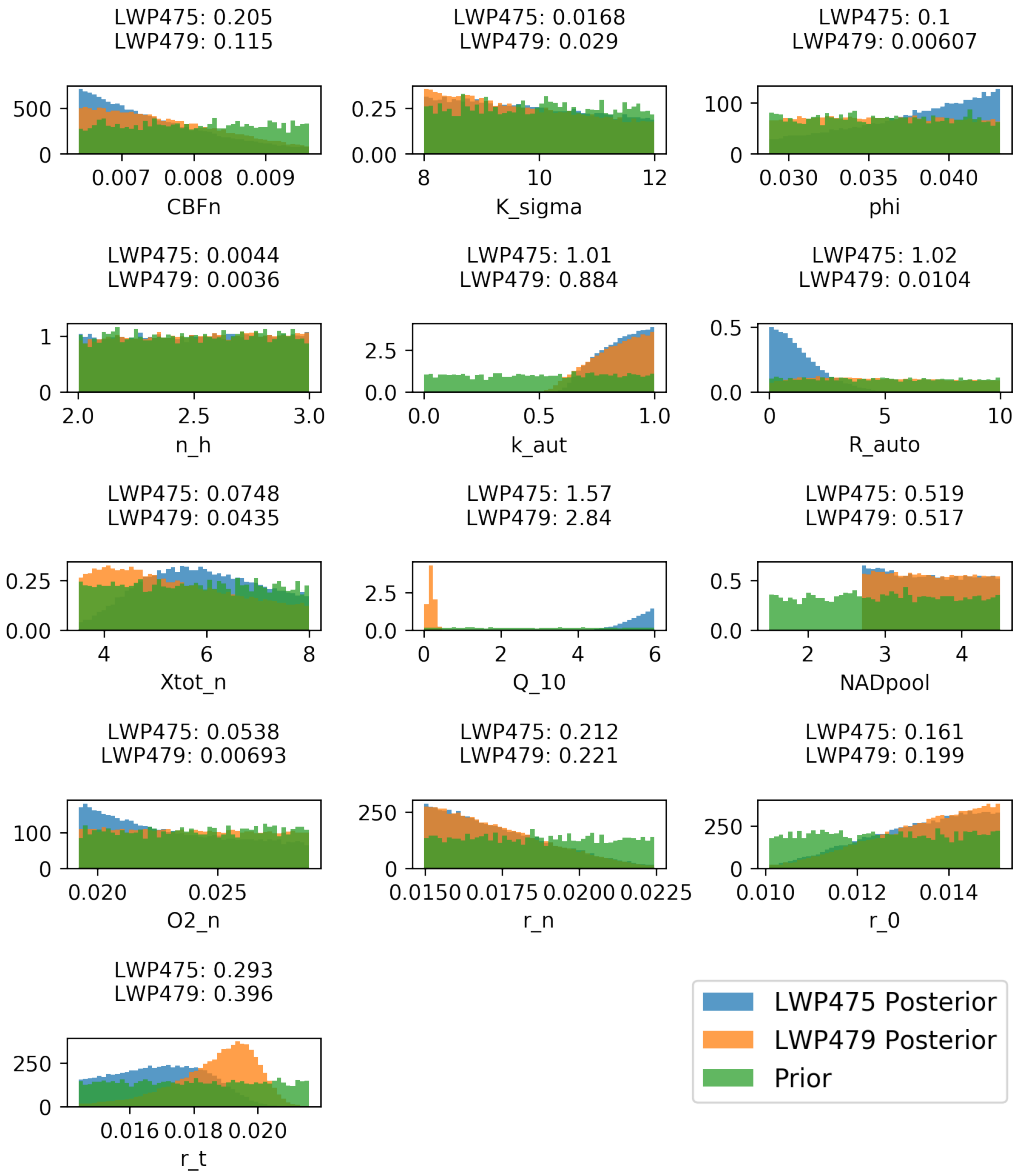

Figure 8.23: K-L Divergence plots for model BPH1.0/1.1 The figure shows marginal distributions for both piglet LWP475 and piglet LWP479, as well as the prior distribution, for model BPH1.*.

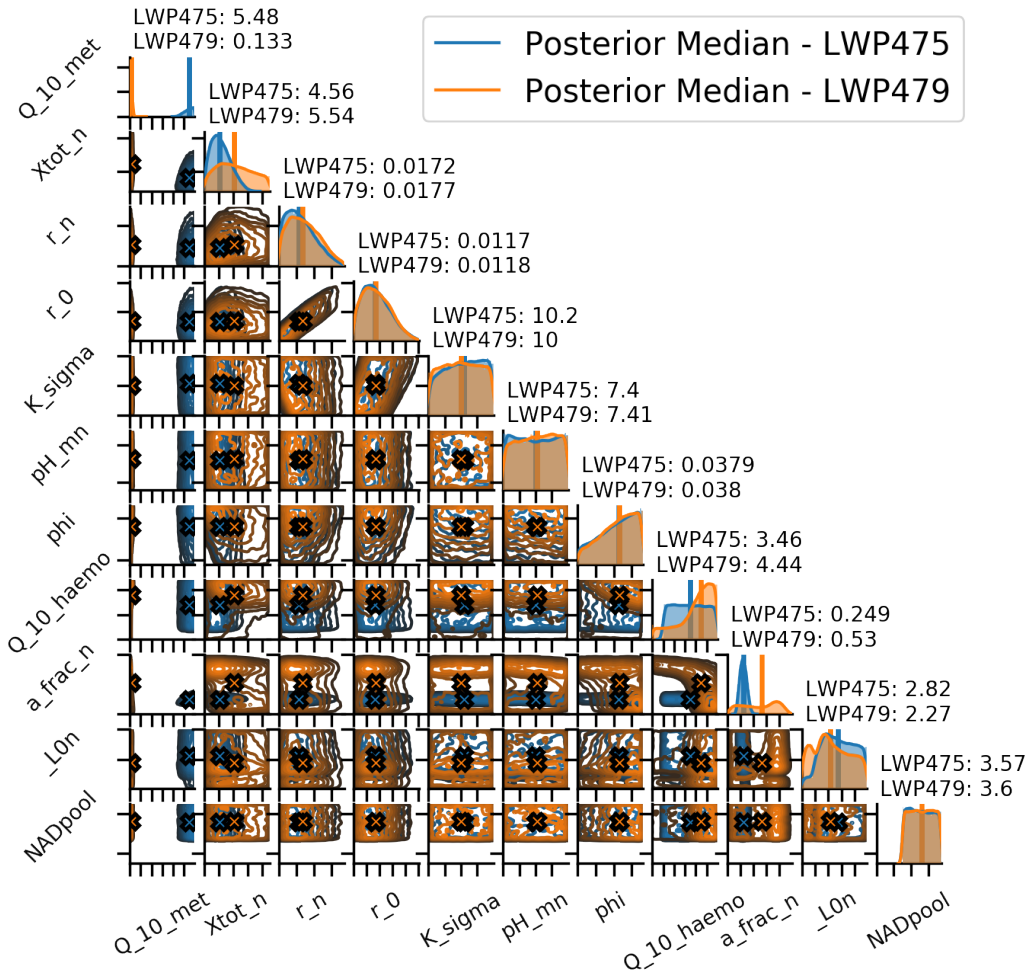

Figure 8.24: Posterior distributions using model BPH2.0/2.1 with piglets LWP475 and LWP479 respectively Posteriors for piglet LWP475 are shown in blue and those for piglet LWP479 are shown in in orange.

Q_10_haemo, a_frac_n and Xtot_n. Figure 8.25 shows the marginal posteriors for each piglet, the prior range and the K-L divergence values. A number of things are clear from this figure. We again see a potentially shared 'piglet parameter space', with the posteriors of NADpool, r_n, r_0 and phi being similar in both piglets, yet dissimilar to the prior.

We also see that Q_10_met is very different in the two piglets with a lot of information gained in both cases as compared to the prior. In particular, we see that Q_10_met for LWP479 is a tightly bound distribution well below 1, suggesting that the temperature response of the metabolic reactions is inverted in this piglet. We also see that q_10_haemo for LWP479 is more heavily skewed towards larger val-

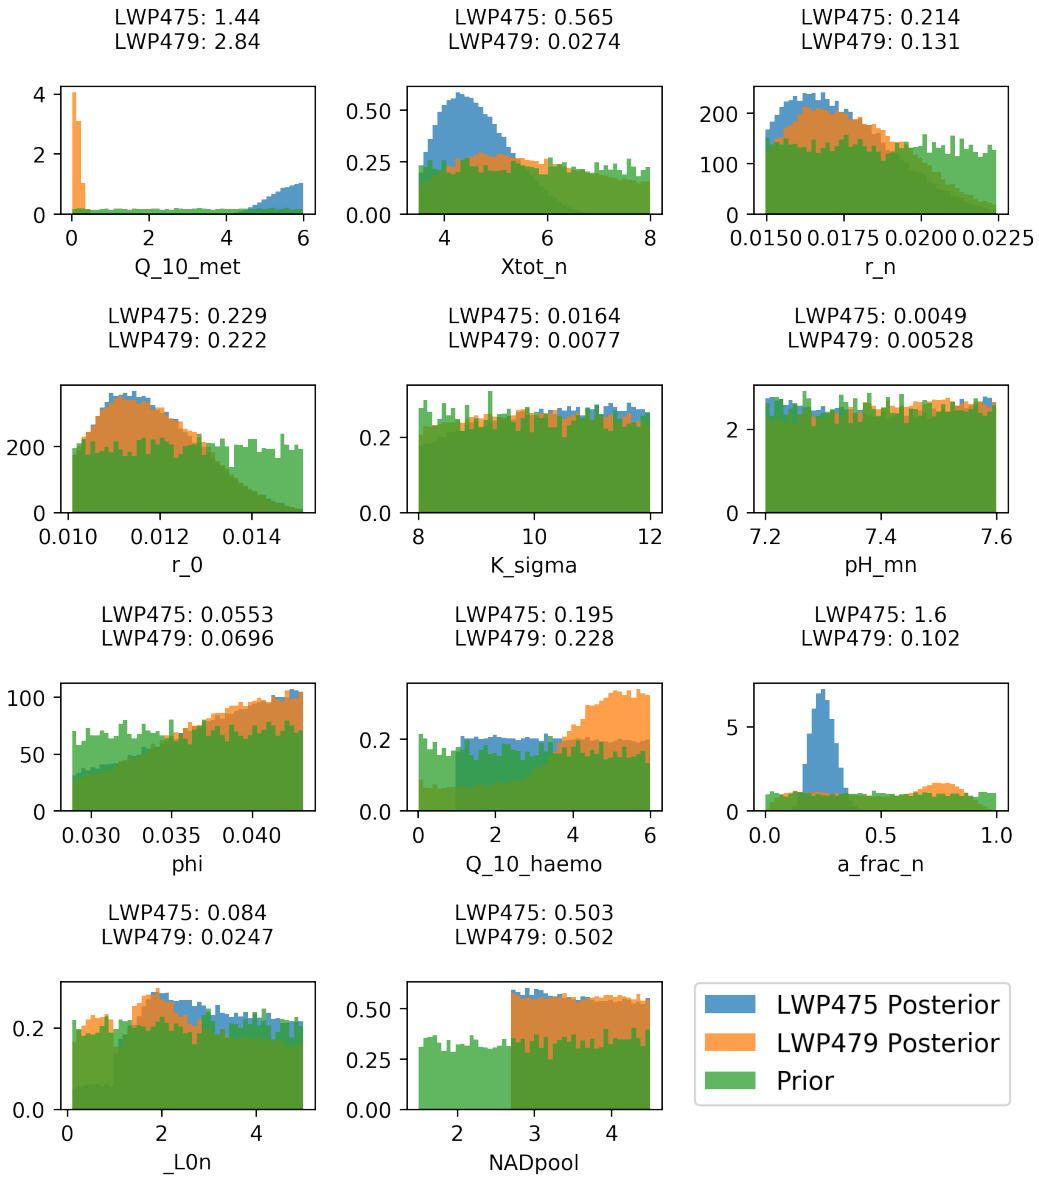

Figure 8.25: K-L Divergence plots for model BPH2.0/2.1 The figure shows marginal distributions for both piglet LWP475 and piglet LWP479, as well as the prior distribution, for model BPH2.*.

ues for LWP479 than LWP475 suggesting that the cooling effects of hypothermia on haemodynamic reactions are more pronounced in the more severely injured piglet. Finally, it is interesting to note that where the two piglets show different parameter spaces, the K-L divergence values for piglet LWP475 are generally higher, such as in a_frac_n (LWP475=1.6, LWP479=0.102) or Xtot_n (LWP475=0.565, LWP479=0.0274). This suggests that the less severely injured piglet has a more tightly bound parameter space.

BP Hypothermia 3.0/3.1 Posterior Distributions

Finally, Figures 8.26 and 8.27 give the posteriors and K-L divergence plots for models BPH3.0 and BPH3.1 with piglets LWP475 and LWP479 respectively, with piglet LWP475 in blue and piglet LWP479 in orange. As with both model BPH1.* and BPH2.*, only a small number of parameter show any noticeable difference between the two piglets, and a number of other parameters suggesting a shared 'piglet parameter space'.

In contrast to the model BPH2.*, where the haemodynamic and metabolic temperature parameters were also split, only the metabolic parameter shows a distinct separation between the piglets. For model BPH3.*, q_diff appears to share an overlapping parameter space in the two piglets.

BrainPiglet 2.1 (BPH0) Posterior Distributions

Figure 8.28 shows the posterior distributions of model BPH0 for both piglet LWP475 and LWP479. When analysing these posteriors we should note the results seen in Figures 8.20 and 8.21. Model BPH0 was unable to produce meaningful results and as a result any conclusions drawn from its parameter space reflect a poorly fitted model. Instead we can use these posteriors to try and understand what parameter changes occur when the model tries and fails to reproduce temperature effects without allowing temperature as an input variable. Where the same parameter occurs in this reduced model and in one of the new model variants, the difference in posterior highlights incorrect conclusions that could have been drawn had we not improved our model.

One of the most noticeable differences between the posterior distributions of model

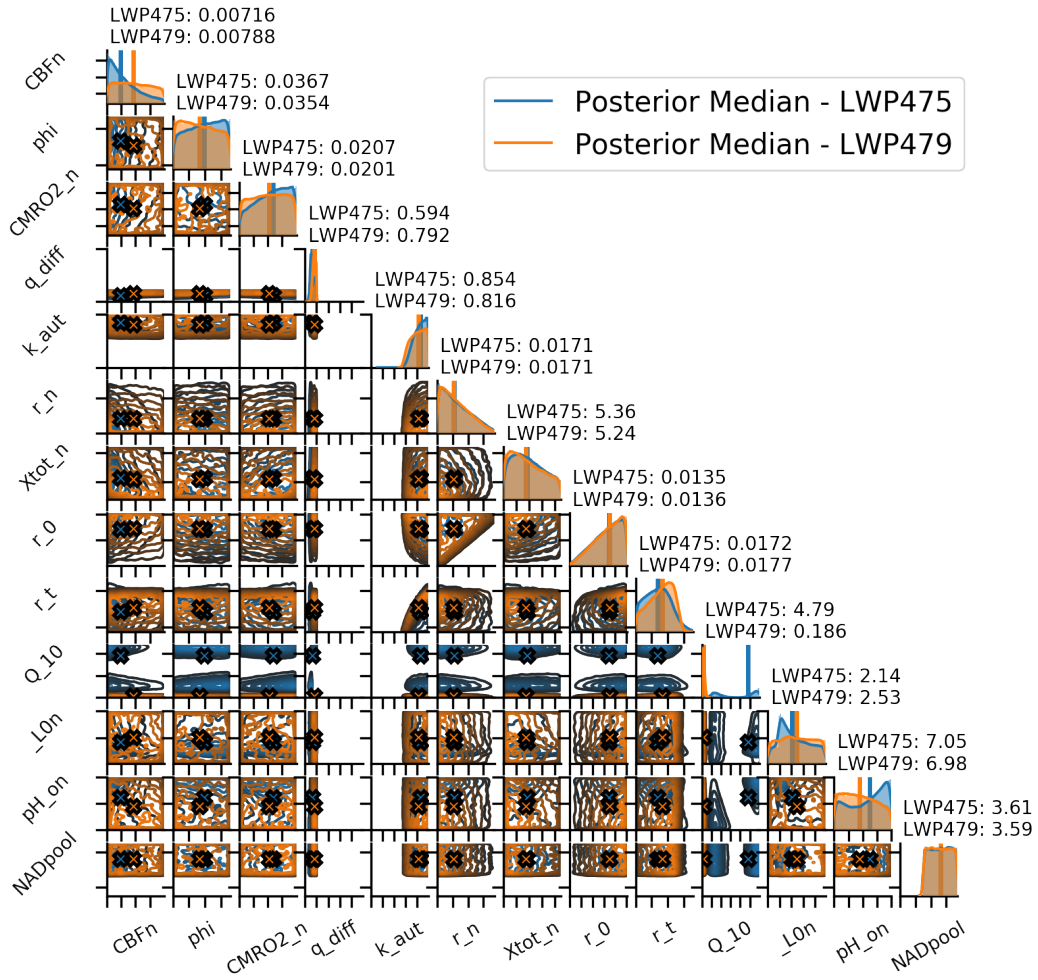

Figure 8.26: Posterior distributions using models BPH3.0/3.1 with piglets LWP475 and LWP479 respectively Posteriors for piglet LWP475 are shown in blue and those for piglet LWP479 are shown in in orange.

BPH0 and the new model variants is how many parameters show posteriors that are clearly different between the two piglets in BPH0 and not in the newer variants. For example, in the newer variants the marginal posteriors for r_0 were generally the same but in BPH0 they are clearly different. This is also the case for NADpool, phi and k_aut. This distinct difference in the posteriors between the new model variants and the old highlights a clear failing in the older models to properly capture behaviour and parameter information during cooling to hypothermic temperatures. Had the older model been used to try and draw conclusions about the differences in the two piglets, it would have identified a number of parameter spaces as being distinct when they instead appear to closely overlap.

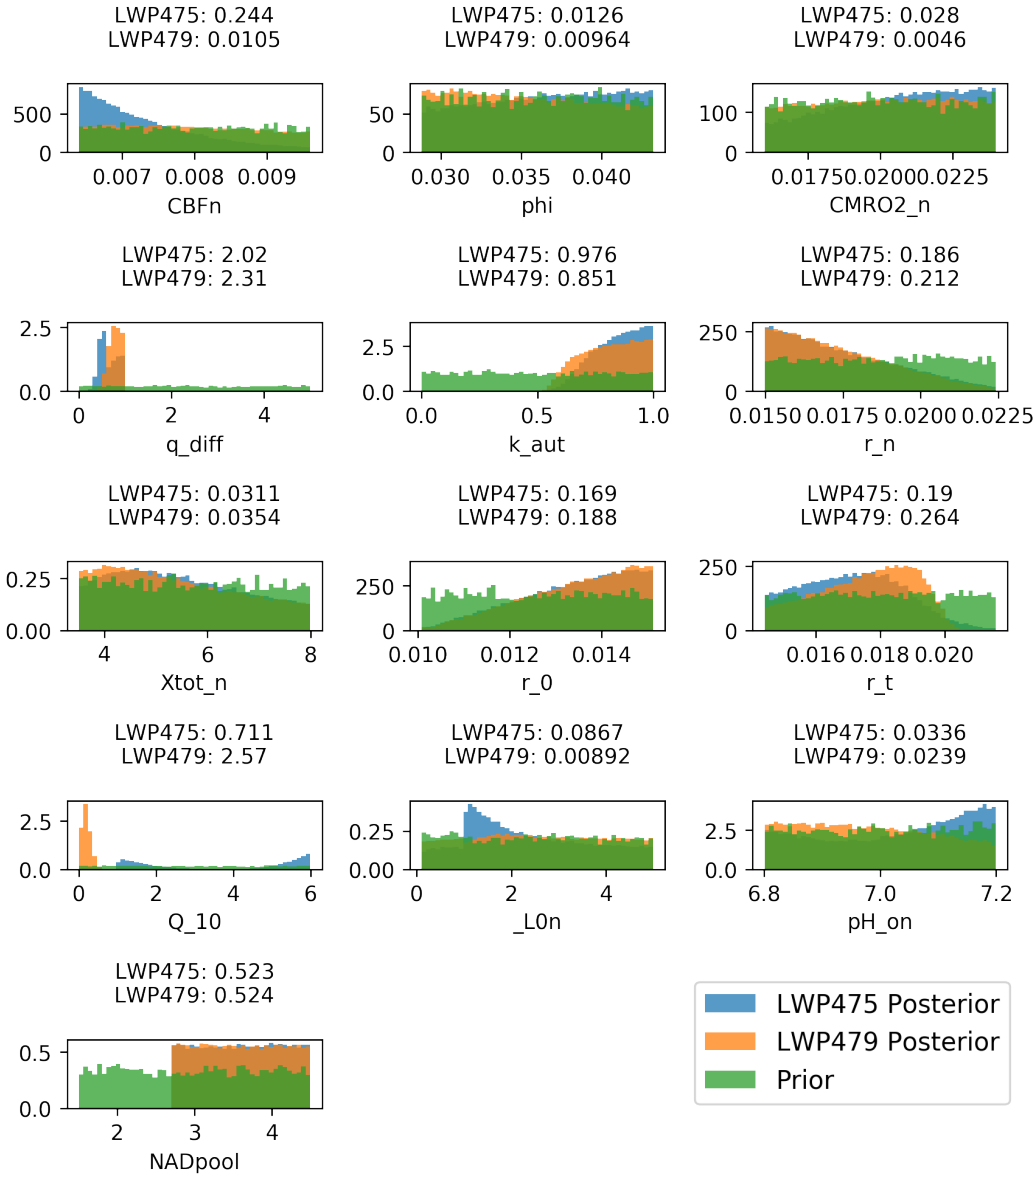

Figure 8.27: K-L Divergence plots for model BPH3.0/3.1 The figure shows marginal distributions for both piglet LWP475 and piglet LWP479, as well as the prior distribution, for model BPH3.*.

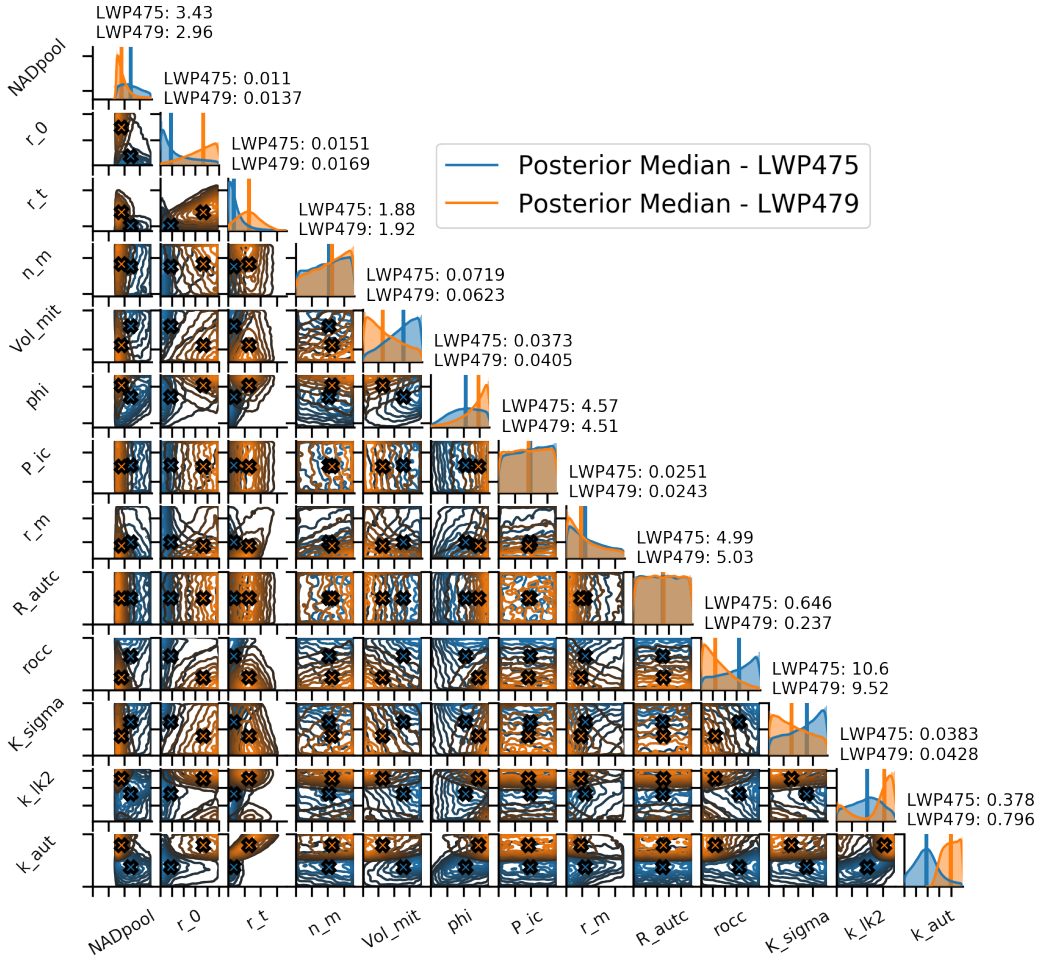

Figure 8.28: Posterior distributions using model BPH0 with piglets LWP475 and LWP479 respectively Posteriors for piglet LWP475 are shown in blue and those for piglet LWP479 are shown in in orange.

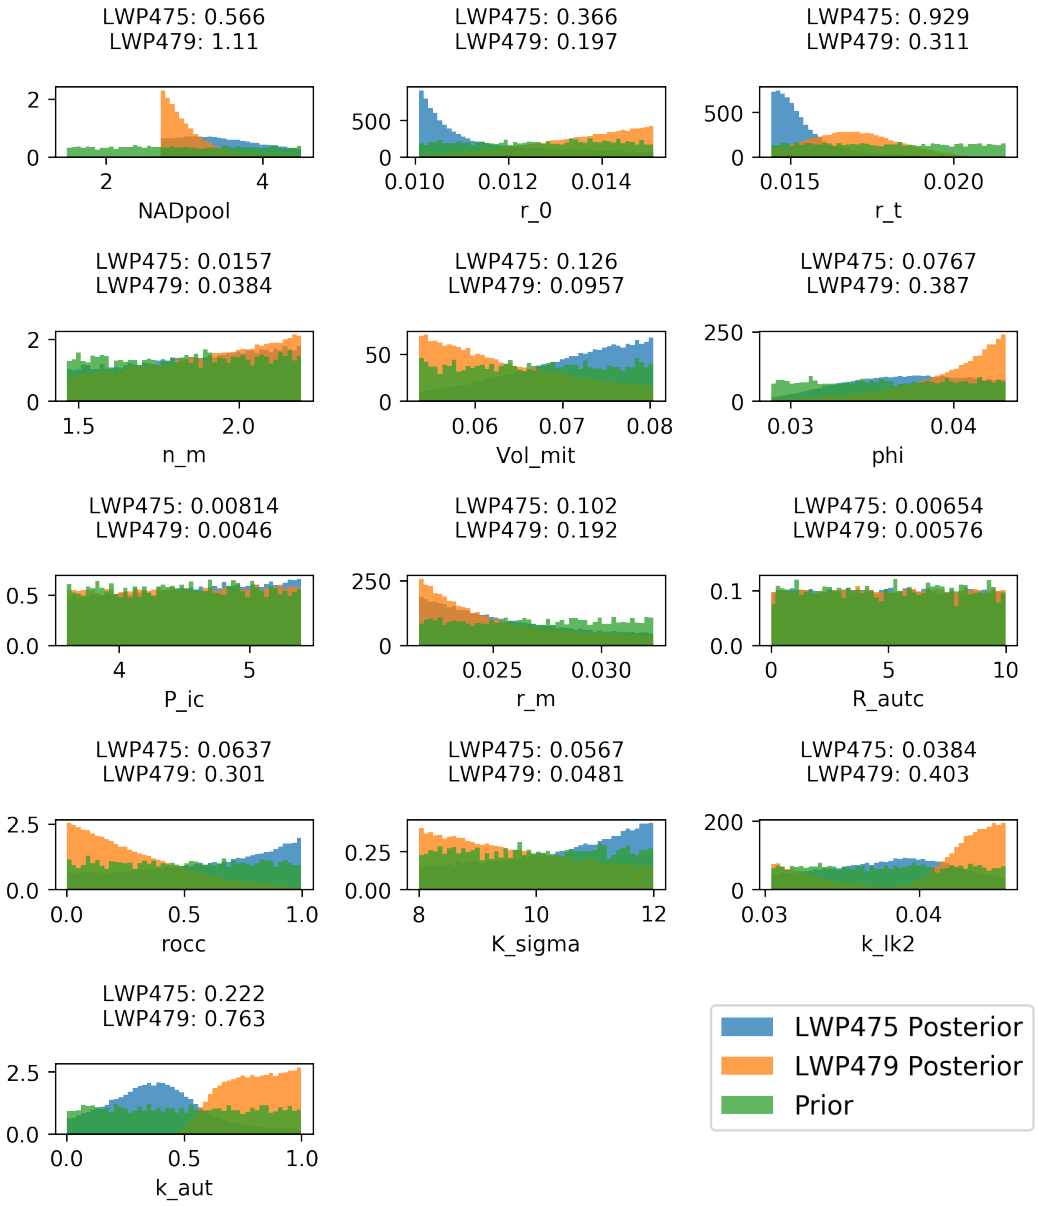

Figure 8.29: K-L Divergence plots for model BPH0 The figure shows marginal distributions for both piglet LWP475 and piglet LWP479, as well as the prior distribution, for model BPH0.

8.4.3 Model Selection

It is reasonable to say that *a priori* all of the models are equally as likely and thus we can apply a uniform prior across model space. Assuming this, we can use the ABC framework to perform model selection on our above models, where model probability is the fraction of the joint posterior model space occupied by a given model. We also calculate the Bayes' factor for every model.

The Bayes' factor is defined as:

$$K = \frac{P(D \mid M_2)}{P(D \mid M_1)} = \frac{P(M_2 \mid D)}{P(M_1 \mid D)} \frac{P(M_1)}{P(M_2)}, \tag{8.7}$$

where K is the Bayes' Factor, $P(D \mid M_i)$ is the probability of seeing data D given model M_i being true, $P(M_i \mid D)$ is the posterior probability of model M_i given data D and $P(M_i)$ is the prior probability of model M_i being true [Gelman et al., 2013]. Where we apply a uniform distribution across model space, $P(M_1) = P(M_2)$, making

$$K = \frac{P(D \mid M_2)}{P(D \mid M_1)} = \frac{P(M_2 \mid D)}{P(M_1 \mid D)}. \tag{8.8}$$

These are presented in a table such that

$$K_{i,j} = \frac{P(M_j \mid D)}{P(M_i \mid D)} \tag{8.9}$$

where i is the row number, j is the column number and $P(M_i \mid D)$ and $P(M_j \mid D)$ are the model probabilities in row i and column j respectively.

Kass and Raftery [1995] provide a table outlining interpretations of different ranges of Bayes' factors, which is itself based on work by Jeffreys [1998]. These classifications are outlined in Table 8.6.

LWP475

Table 8.7 shows the model probabilities for those models with a non-zero probability in piglet LWP475. An acceptance rate of 0.1% was used, leading to a total posterior size of 200,000 samples. Only model BPH2 and BPH3 appear in the joint model posterior for this piglet, and we can see that model BPH2 is overwhelmingly

$\log_{10}$(K)	K	Evidence against M_1
0 to 1/2	1 to 3.2	Not worth more than a bare mention
1/2 to 1	3.2 to 10	Substantial
1 to 2	10 to 100	Strong
> 2	> 100	Decisive

Table 8.6: Table outlining the interpretations of Bayes' Factor values by Kass and Raftery [1995].

Model Name	Model Probability	Median NRMSE
BPH2	99.9825%	0.6025
BPH3	0.0175%	0.6064

Table 8.7: Model selection results for LWP475. Acceptance rate of 0.1% was used, leading to a total posterior size of 200,000 samples. BPH2 is the most likely model here, with a probability of 99.9825%.

the most probable with a probability of 99.9825%. It also has a marginally better median NRMSE value of 0.6025 against 0.6064 for model BPH3.

Table 8.8 show the Bayes' factors for each model in piglet LWP475. We can see that models BPH0 and BPH1 have a Bayes' factor of zero due to their zero probability. Model BPH2 has a Bayes' factor of 5713.3 over BPH3. As per Table 8.6, this is *decisive* evidence to choose model BPH2 over all others in this case.

LWP479

Table 8.9 shows the model probabilities for those models with a non-zero probability in piglet LWP479, considering both only $x.1$ and all model variants. An acceptance rate of 0.1% was used, leading to a total posterior size of 200,000 samples $x.1$ model variations ($Q_{10,\cdot} < 1$) and 350,000 samples for all model variations.

	BPH0	BPH1	BPH2	BPH3
BPH2	0	0	—	0.00018
BPH3	0	0	5713.3	—

Table 8.8: Bayes' Factors for each model using dataset LWP475. Bayes factor is calculated based on the probability of a model in a given column divided by the probability of the model in a given row. A Bayes' Factor of zero means that a model had zero probability.

Model Name	$x.1$ Models Probability	$x.1$ Models Median NRMSE	All Models Probability	All Models Median NRMSE
BPH2.1	72.15%	0.99944	64.54%	1.059
BPH1.1	16.05%	1.054	24.58%	1.117
BPH3.1	11.8%	1.010	10.89%	1.067

Table 8.9: Model selection results for LWP479. Acceptance rate of 0.1% was used, leading to a total posterior size of 200,000 samples $x.1$ model variations ($Q_{10,met} < 1$) and 350,000 samples for all model variations. BP Hypothermia 2.1 is the most likely model in both cases, with a probability of 72.15% for only $x.1$ variants and 64.54% for all model variants. This is followed by BP Hypothermia 1.1 with a probability of 16.05% in $x.1$ variants and 24.58% in all models.

In both cases models BPH1.1, BPH2.1 and BPH3.1 appear in the joint model posterior for this piglet. We can see that model BPH2 is still the most probable, albeit less so than for piglet LWP475, with a probability of 72.15% for only $x.1$ variants, and 64.54% for all variants. Additionally, model BPH1.1 is the second most probable for this piglet, whilst it wasn't at all probable for LWP475. It is also important to note that when considering all model variants, it is clear that the $x.1$ variants, which have a prior allowing $Q_{10,\cdot} < 1$, are more probable than the $x.0$ variants, which have a minimum $Q_{10,\cdot}$ of 1. This implies that allowing $Q_{10,\cdot} < 1$ is important when fitting subjects with a potentially severe injury.

When looking at the median NRMSE values for the three models, model BPH2.1 has the lowest error of all the three models, but BPH1.1 actually has a higher median error than model BPH3.1 despite being a less probable model. We can clearly see this if we plot a histogram of the error values for each model within the joint posterior. When we do this we see that whilst more BPH1.1 samples appear in the posterior, they tend to be at the upper end of NRMSE values.

Looking at Table 8.10 we can see the Bayes' factors for each of the $x.1$ model variants. As only $x.1$ variants have non-zero probability in both of the considered cases, the probabilities of *only* these variants was used to calculate the Bayes' factors. We can see that model BPH0 again has a Bayes' factor of zero due to its zero probability. Model BPH2.1 has a Bayes' factor of 4.95 over model BPH1.1 and 6.114 over model BPH3.1. Using Table 8.6, this is *substantial* evidence to use

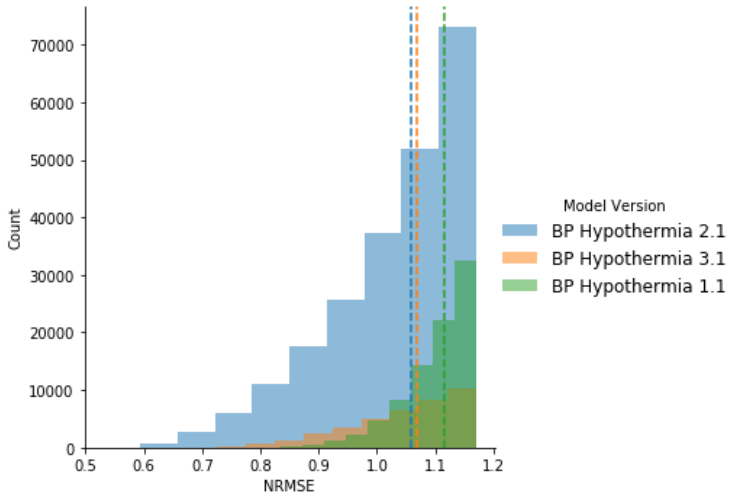

Figure 8.30: Distribution of NRMSE values within the joint posterior. The histogram shows the distribution of NRMSE values within the joint posterior of all models for piglet LWP479. The median for each version is marked by a dotted line. We can see that whilst more samples for model BP Hypothermia 1.1 appear within this posterior, model BP Hypothermia 3.1 has a wider spread of error values, with a lower median value.

	BPH0	BPH1.1	BPH2.1	BPH3.1
BPH1.1	0	—	0.222	1.360
BPH2.1	0	4.95	—	6.114
BPH3.1	0	0.735	0.164	—

Table 8.10: Bayes' Factors for each model using dataset LWP479. Bayes factor is calculated based on the probability of a model in a given column divided by the probability of the model in a given row. A Bayes' Factor of zero means that a model had zero probability. Here we are using the probabilities shown in columns 2 and 3 of Table 8.9, as only the $x.1$ variants appear to be valid for piglet LWP479.

model BPH2.1 over either model BPH1.1 or model BPH3.1. Whilst this evidence is not as strong as in piglet LWP475, it is still enough to suggest that a model BPH2 variant is the most likely model.

In conclusion, as seen in Tables 8.7, 8.8, 8.9 and 8.10 all of the hypothermia model variants are more probable than the original BrainPiglet2.1 model. Of these, the BPH2 variants perform best in both piglets and for the severely injured piglet all $x.1$ variants being more probable in LWP479 than the $x.0$ variants. $x.1$ variants were not tested for piglet LWP475, so conclusions cannot and should not be drawn about the validity of those models against the $x.0$ variants for that piglet. Given that

the $Q_{10,\cdot}$ posteriors for LWP475 were at the higher end of the prior, and the only difference between the $x.0$ and $x.1$ variants were at the lower end of these priors, it is likely that there would be minimal difference. This would however require confirmation in any further work.

8.5 Model Fit Analysis

Having determined the most probable model for each piglet we can look at the quality of the fit in both cases. One way of doing this is by considering the distribution of the residuals for each fit. This distribution can allow us to consider how well the fitting process has performed and what kind of fit it has provided. We do this by looking at both the distribution of the residuals directly and by using a Q-Q plot. Figures 8.31 and 8.32 show the residual distributions and Q-Q plots for piglet LWP475 and LWP479 respectively.

Figure 8.31a shows that BPH2 seems to generally overestimate CCO values, with approximately 75% of the residuals being greater than 0. Additionally there are three distinct peaks visible, most likely due to the simulation failing to capture the periodic increases and decrease in CCO that are seen in the measured data. In contrast BPH2 seems to slightly underestimate the HbO_2 signal, with most of the residuals being below zero. Finally, the HHb signal appears to have been simulated relatively well with a mean of approximately zero. None of the distributions appear to be particularly normally distributed with all showing some degree of skew and kurtosis.

It is easier to identify this in the Q-Q plots in Figure 8.31b. The plots use a standardised diagonal line, where the expected order statistics are scaled by the standard deviation of the sample residuals and have the mean added to them. The theoretical mean and sample means are also shown to allow easier comparison.

All of the Q-Q plots heavier tails than would be expected in a normal distribution, as indicated by the points in the middle generally following the diagonal line but trailing off at the extremities. The skew of each distribution is clear from where the

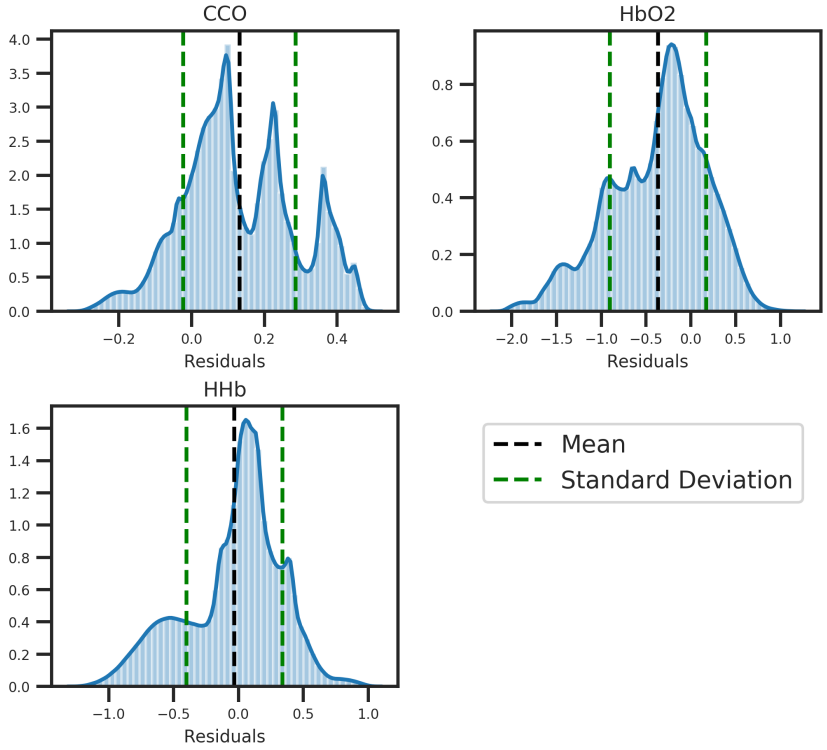

(a) Distribution of residuals for model BPH2 and piglet LWP475.

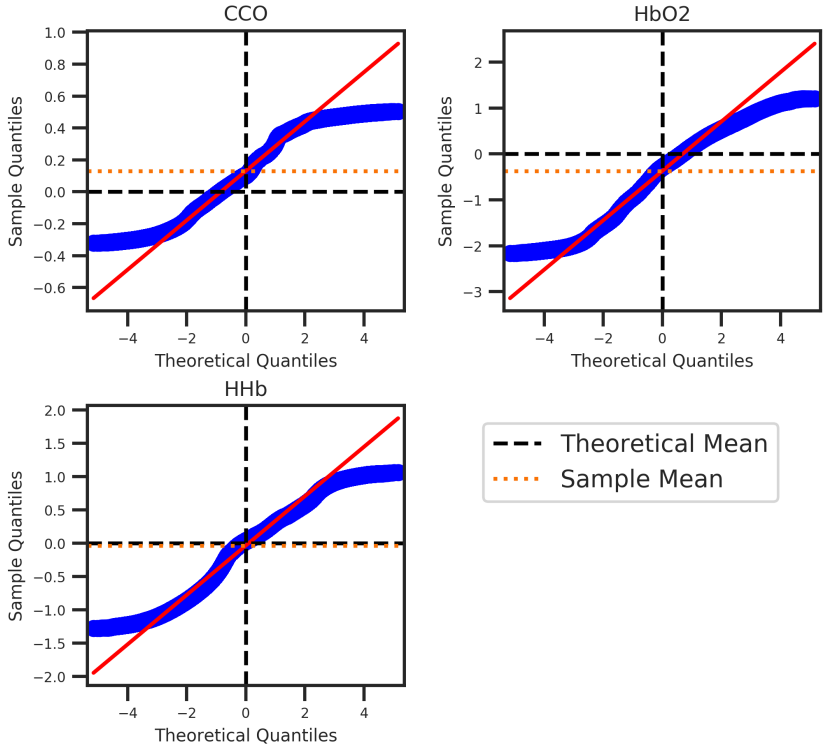

(b) Q-Q plot for residuals from model BPH2 and piglet LWP475.

Figure 8.31: Residual analysis for model BPH2 and piglet LWP475.

sample and theoretical mean both fall, with the sample mean of CCO being above the theoretical mean, the sample mean of HbO_2 being below the theoretical mean and the sample mean of HHb falling almost directly on the theoretical mean.

Figure 8.32a shows the residual distributions for model BPH2.1 and piglet LWP479. Here we see that the model appears to heavily underestimate the CCO signal, with a reasonably large amount of skew towards negative values and two peaks both below zero. The HbO_2 signal appears to be slightly skewed towards positive values, as would be expected based on its posterior predictive in Figure 8.21b, with a long tail into negative values. This long tail is likely to be due to the simulation underestimating the HbO_2 signal for the later times past around 2750 seconds. The HHb signal distribution suggests a slightly better fit and a mean value close to zero but again with some skew, this time towards negative values.

The Q-Q plots support these conclusions, with the CCO signal appearing to be fairly platykurtic with some negative skew and the HbO_2 signal being fairly leptokurtic with positive skew. We can also see that, based on the Q-Q plot, the residuals for HHb appear to be generally normally distributed.

From Figure 8.31 and 8.32 we can see that model BPH2.* variants perform a good job of reproducing the observed behaviour during hypothermia. Considering this alongside the Bayesian model selection process it can be said with a good degree of certainty that the model that is best able to incorporate temperature, and reproduce physiology during TH, is model BPH2. It is important that the prior range include values of $Q_{10,\cdot} < 1$ in order to be able to successfully reproduce behaviour where the typical response of metabolic and haemodynamic reactions to temperature has inverted, as in LWP479.

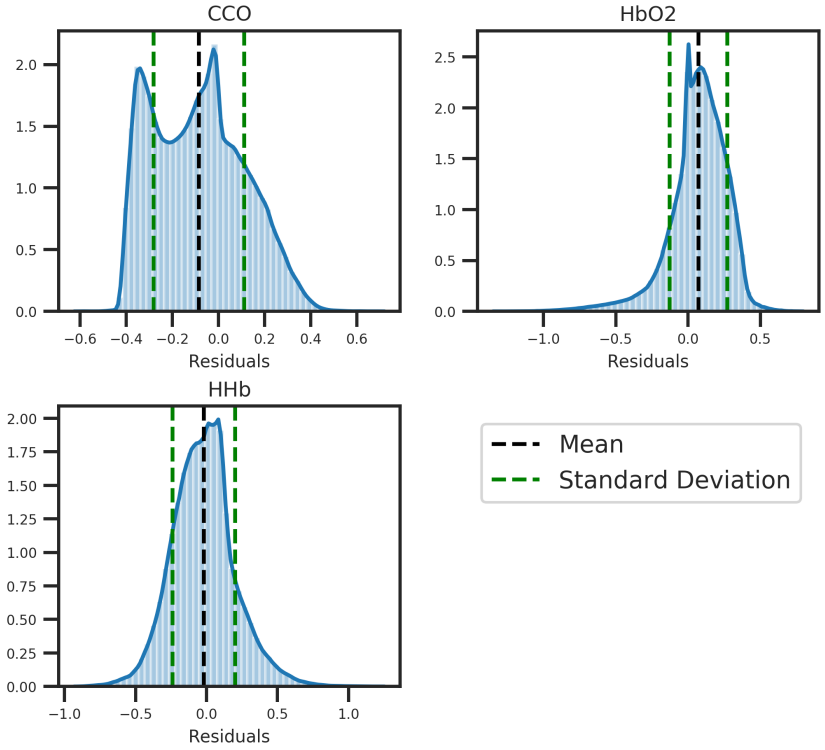

(a) Distribution of residuals for model BPH2.1 and piglet LWP475.

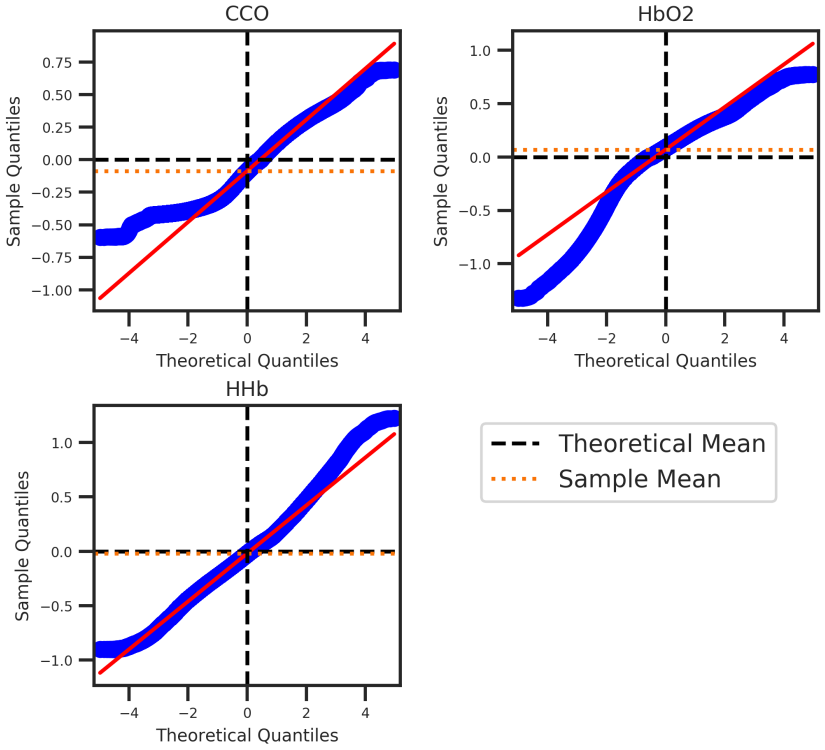

(b) Q-Q plot for residuals from model BPH2.1 and piglet LWP479.

Figure 8.32: Residual analysis for model BPH2.1 and piglet LWP479.

8.6 Conclusion

In this chapter I have outlined a number of model variants(BPH1, BPH2 and BPH3) that are able to incorporate temperature and its effect on physiology. Of these three model variants, model BPH2 was found to be the most effective at reproducing the expected behaviour of the cerebral physiological effects during therapeutic hypothermia. It did this by separating the effect of temperature on metabolic and haemodynamic reactions into two separate temperature co-efficients. We were then able to evaluate that the inverted response seen in piglet LWP479, which had a more severe outcome, could be reproduced by allowing the metabolic temperature co-efficient to be less than 1, effectively inverting the impact of temperature on metabolic reactions. Whilst the sample size of two piglets is understandably very small, this does point to an exciting prospect for understanding what is behind the differing outcomes from therapeutic hypothermia. In the next chapter we will use model BPH2 in order to investigate potential mechanisms behind differences observed in human neonates during spontaneous desaturation, based on work in Bale et al. [2018].

Chapter 9

Neonatal Modelling

9.1 Motivation

Work by Bale et al. [2018] found that a strong correlation between CCO and HbD during spontaneous desaturation events was an indicator of newborn brain injury. They postulated that it is "due to mitochondrial dysfunction and reduction in cerebral metabolic rate as a result of severe encephalopathy". It may be possible to use the BPH2 model developed and outlined in Chapter 8 to investigate these data and their hypothesis on the mechanisms behind the correlation.

Figure 9.1 shows the gradient between ΔHbD and ΔCCO for two neonates, neo021 (Figure 9.1a)) and neo007 (Figure 9.1b)) during desaturation events. In their paper, Bale et al. identify this relationship to be a key indicator of the severity of cerebral injury. The less severely injured neonate, neo021, shows a lesser relationship between the two variables whilst the more severly injured neonate, neo007, shows a stronger relationship as evidenced by the steeper gradient. In this section we look at using our systems biology approach to better understand this and the mechanisms behind it.

9.1.1 In-silico Experimentation

Before looking at these data directly, we will now leverage one of the key benefits of our modelling approach, which is to perform in-silico experimentation. We will take the posteriors fitted to each piglet in the previous chapter for model BPH2 and simulate the response to a spontaneous desaturation event.

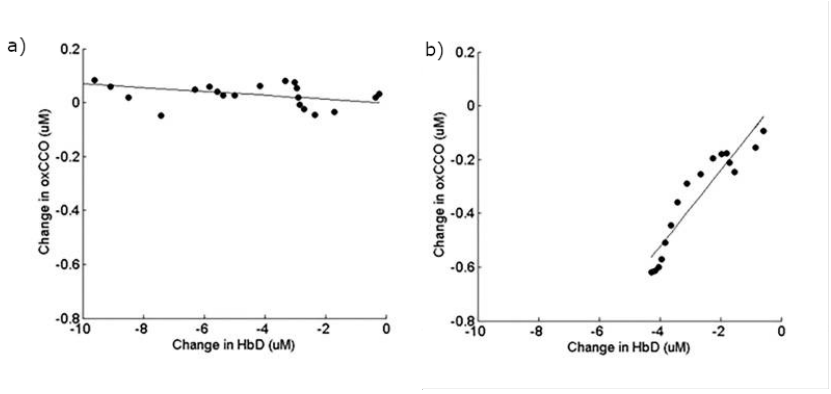

Figure 9.1: Gradients of ΔHbD and ΔCCO for neonates a) neo021 and b) neo007. Neo007 shows a much steeper gradient, hypothesised to be due to the more severe injury leading to a tighter coupling between brain blood oxygenation (ΔHbD) and cerebral metabolism (ΔCCO).
Image: Bale et al. [2018].

The posterior was sampled from 1000 times and the model was driven by two inputs: a constant temperature of 33.5 °C and an SaO_2 signal generated to simulate a potential desaturation event, with SaO_2 dropping to 82% at its lowest over a period of approximately 5 minutes. Figure 9.2 shows the SaO_2 signal and the model outputs for each piglet. Column a) of Figure 9.2 shows the model response for piglet LWP475. We can see that the haemodynamic signals correlate strongly with the change in SaO_2 whilst the CCO signal shows little change. In contrast, column b) shows the response of piglet LWP479 to the same stimuli and we can see a distinct difference. The CCO signal is closely linked to the SaO_2 signal.

9.2 Data selection

Data were selected based on the work outlined above. Of the data available, two neonates were chosen for fitting with each representing a different outcome. These neonates were neo007 and neo021 and were chosen based on the quality of the data and for good comparison with the work of Bale et al. [2018]. Neonate neo007 showed an *unfavourable* outcome with a Lac/NAA ratio of 1.32 whilst neonate neo021 showed an *favourable* outcome with a Lac/NAA ratio of 0.2 [Bale et al., 2018].

The data used are the exact same as that used to draw conclusions in Bale et al.

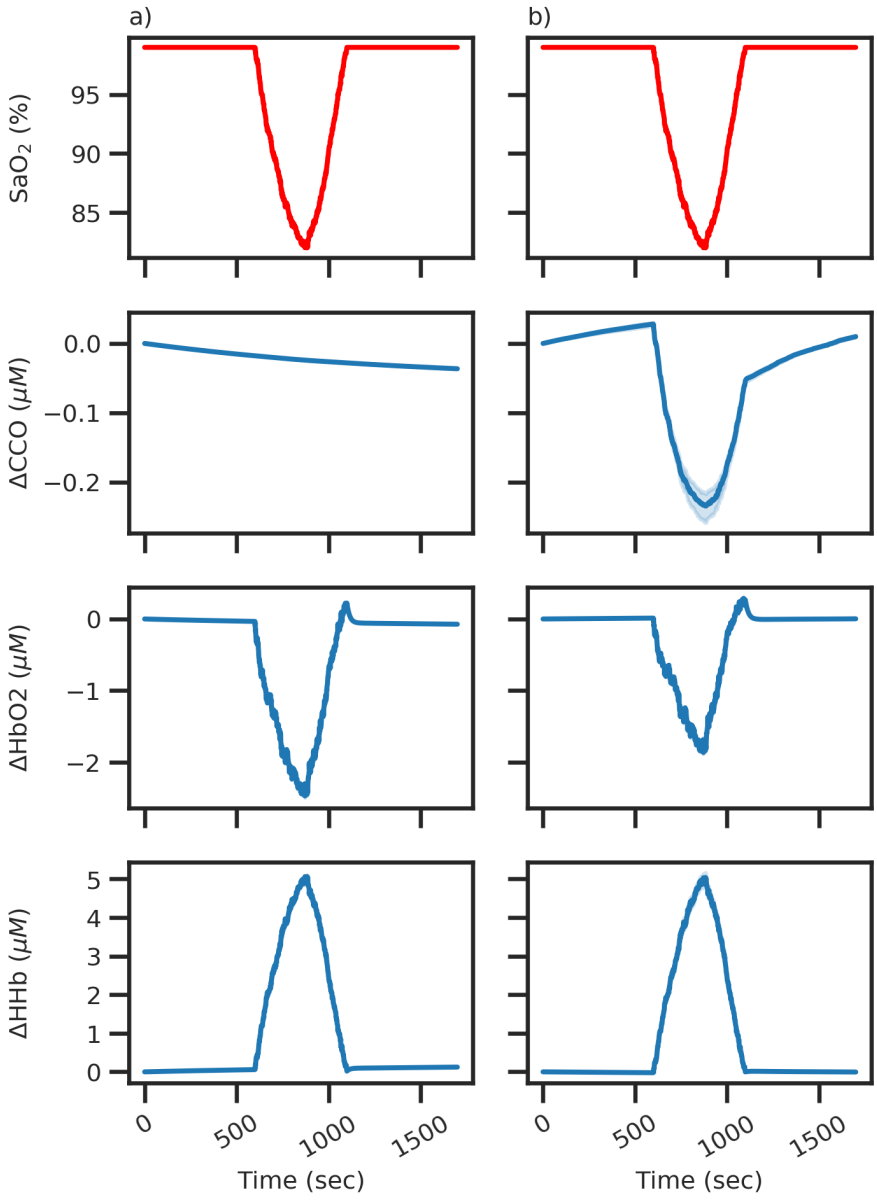

Figure 9.2: In-silico desaturation experimentation. Simulations of a desaturation event using posteriors for piglets LWP475 and LWP479. Column a) show the posterior predictive distribution of the desaturation event using the posterior from piglet LWP475, representing a mild injury response. Column b) show the posterior predictive distribution of the desaturation event using the posterior from piglet LWP479, representing a severe injury response.

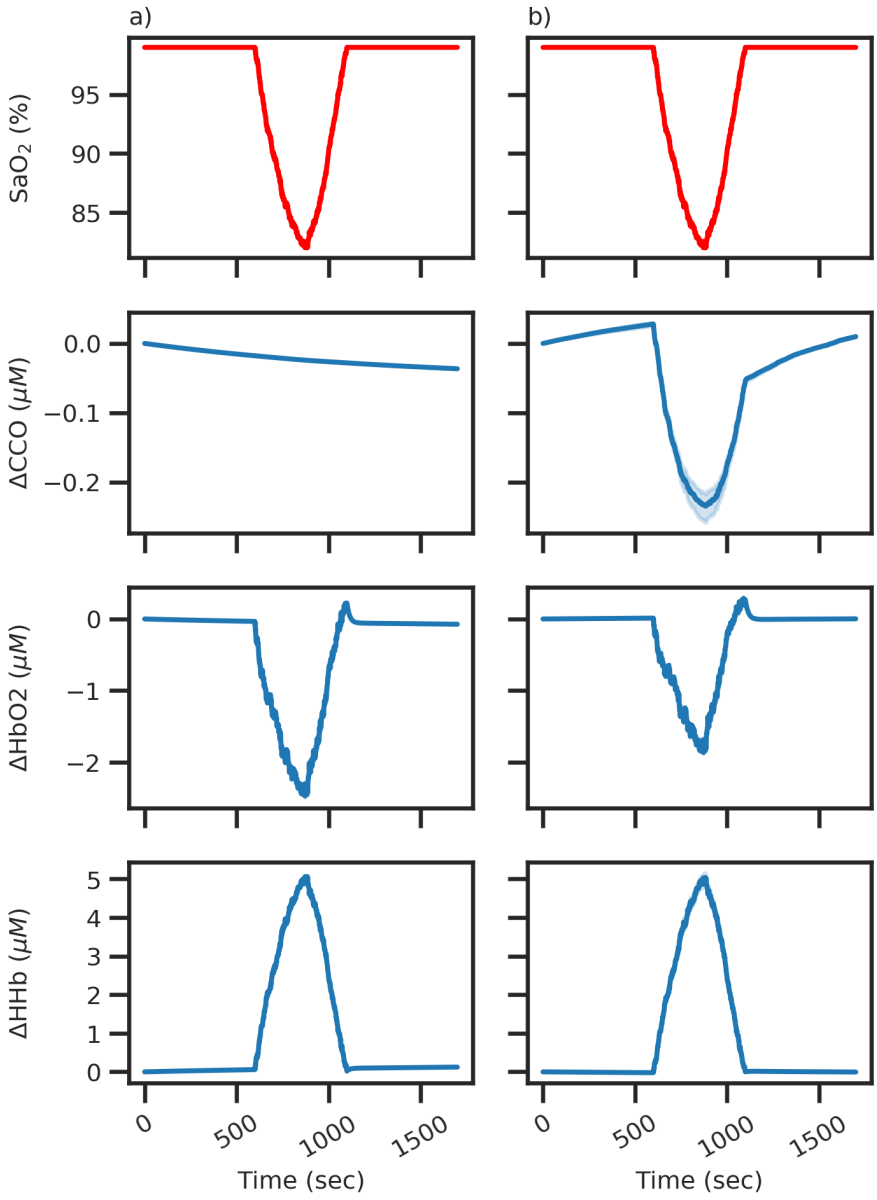

Figure 9.3: In-silico desaturation experimentation. Simulations of a desaturation event using posteriors for piglets LWP475 and LWP479. Column a) show the posterior predictive distribution of the desaturation event using the posterior from piglet LWP475, representing a mild injury response. Column b) show the posterior predictive distribution of the desaturation event using the posterior from piglet LWP479, representing a severe injury response.

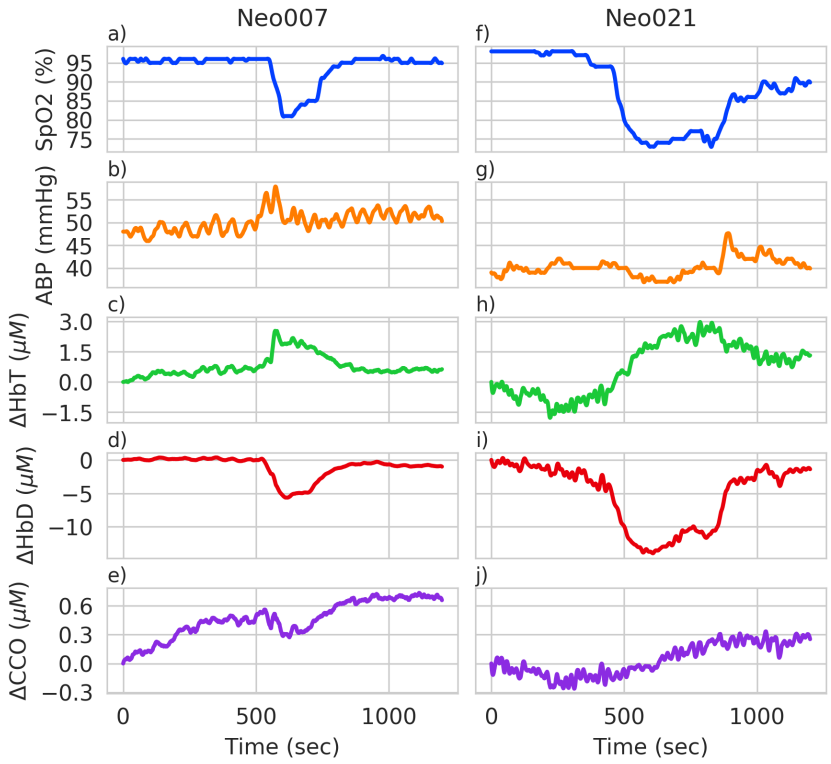

Figure 9.4: Clinical Data for Neonates 007 and 021. Shown is arterial oxygen saturation (SpO2), arterial blood pressure (ABP), change in total haemoglobin (ΔHbT), change in haemoglobin difference (ΔHbD) and change in cytochrome-*c*-oxidase ((ΔCCO).

[2018], thus ensuring that any confirmation of their hypothesis is not a side effect of the data cleaning process. Figure 9.4 shows the arterial oxygen saturation (SpO2), arterial blood pressure (ABP), change in total haemoglobin (ΔHbT), change in haemoglobin difference (ΔHbD) and change in cytochrome-*c*-oxidase (ΔCCO) data for each neonate.

9.3 Model Fitting

9.3.1 Sensitivity Analysis

Sensitivity analysis was performed using the gradient between HbD and CCO as the summary metric. This is in line with the work in Bale et al. [2018] outlined above, and the findings from the end of Chapter 8, where the gradient between HbD and CCO was shown to be a potential biomarker of brain injury. Using this it may

possible to identify parameters that both control the behaviour seen during desaturation events and provide good separation between the brain states of mild and severe injury .

Figure 9.5 shows these sensitivity analysis results. Figure 9.5a shows the μ_* values for the 15 most sensitive parameters in each neonate. We can see clearly here that that the most sensitive parameters in neonate neo007 were more sensitive than those in neonate neo021, as shown by their overall higher μ_* values. Figure 9.5b directly compares the μ_* values in each neonate, with a blank cell showing that a given parameter was not in the most sensitive for that neonate. What is also interesting to note that is neonate neo007, both CMRO2_n and CBFn were identified as being sensitive parameters, but this was not the case in neonate neo021. Given that the physiological reality that these parameters relate to were identified as being important by Bale et al. [2018], this suggests that the sensitivity analysis process has potentially already identified important information: more severely injured systems show greater sensitivity to changes in the normal states of cerebral metabolic rate of oxygen and cerebral blood flow.

Having completed the initial sensitivity analysis run, the model parameters were then reduced further by using prior knowledge to consider if a sensitive parameter is also one that needs fitting. As previously, this involved considering each of the fifteen identified parameters and considered if they could reasonably be expected to actually differ from their default values in terms of real world physiology. This both aids in reducing the model down further, as well as reducing the risk of 'false positive' parameter fits, i.e. those parameters that have large impact on the model outcome but where the fitted value is unlikely to match up with what could be reasonably expected in the real-world. The parameters chosen to be removed were NADpool, P_an, SaO2_n, Xtot_n, a_frac_n, pH_on. The rationale for removing these parameters is to reduce the complexity of the model as much as possible whilst retaining information. Finally, the parameter Q_10_met was added in as this was found to be important in Chapter 8 when fitting data collected during therapeutic hypothermia, as was the case here. Table 9.1 shows the final parameter selection,

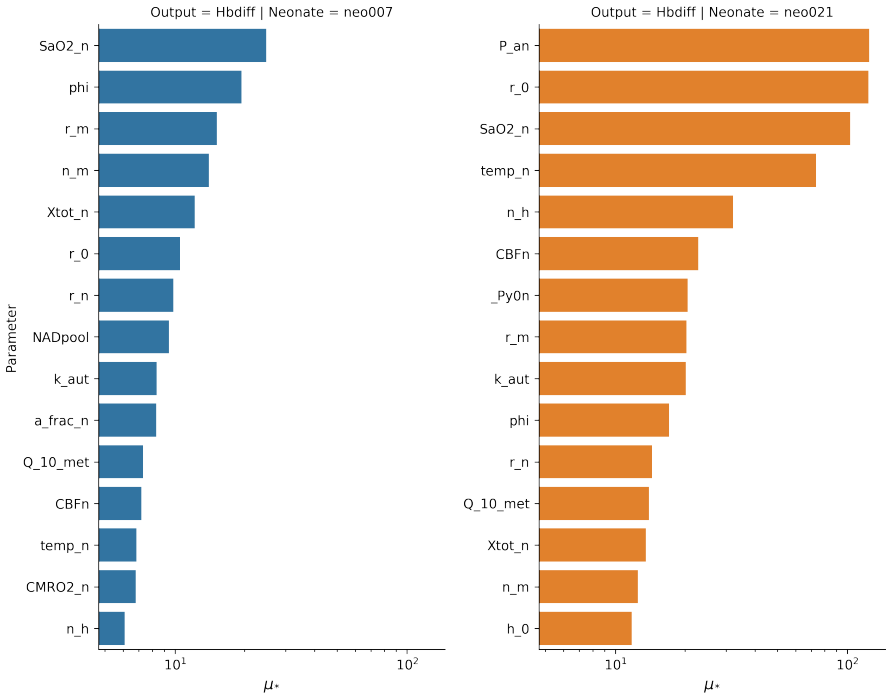

(a) Bar chart of μ_* values for the 10 most sensitive parameters in each neonate. Both charts share an x-axis scale.

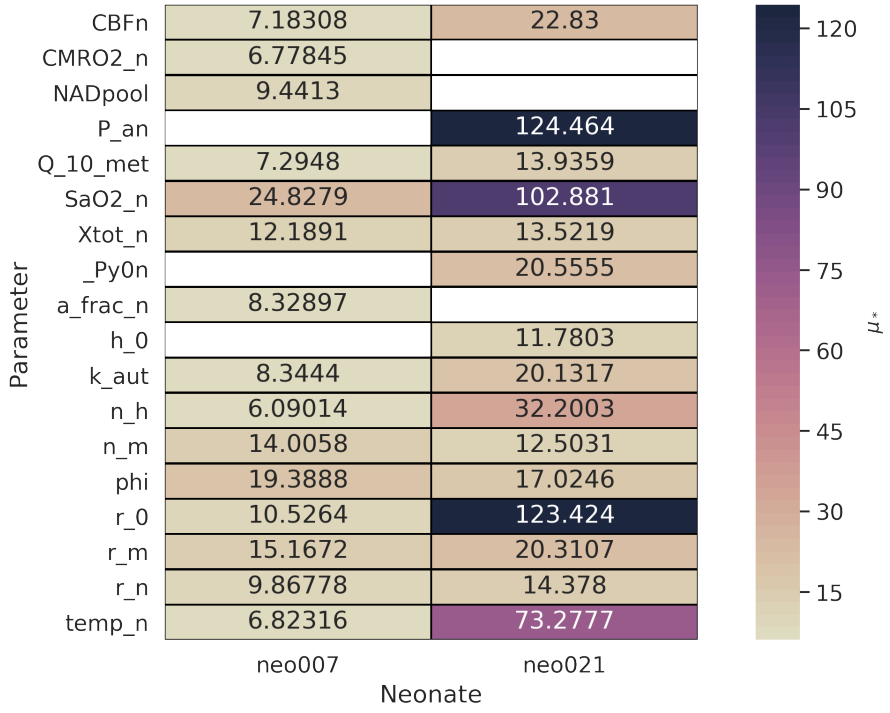

(b) Heatmap of μ_* values for the 10 most sensitive parameters in each neonate. An empty cell marks that a parameter wasn't within the 10 most sensitive parameters for that neonate.

Figure 9.5: Sensitivity analysis results

each parameter's minimum and maximum values and also its default value.

Parameter	Min. Value	Max. Value	Default Value
k_aut	0.0000	1.0000	1.0000
r_0	0.0101	0.0151	0.0126
r_m	0.0216	0.0324	0.0270
r_n	0.0150	0.0224	0.0187
n_m	1.4640	2.1960	1.8300
n_h	2.0000	3.0000	2.5000
phi	0.0288	0.0432	0.0360
_Py0n	0.0800	0.1200	0.1000
CMRO2_n	0.0160	0.0240	0.0200
Q_10_haemo	0.1000	6.0000	2.4000
Q_10_met	0.1000	6.0000	2.4000
h_0	0.0024	0.0036	0.0030
CBFn	0.0064	0.0096	0.0080

Table 9.1: Table showing final sensitivity analysis parameters for noenatal desaturation data.

9.3.2 Bayesian Model Fitting

The model was run 40,000,000 times for each neonate using the prior distributions described in Table 9.1. Not all runs were successful and any parameter samples that produced a failed run were excluded from the final posterior distribution in the same way as a parameter sample with a large NRMSE value is. For neo007 33,912,900 samples produced successful runs and for neo021 32,813,436 samples produced successful runs.

Before looking at the posterior distributions formed from these samples, we should look at the distribution of errors in the two neonates. This can be useful as it helps to identify if the model is performing better for one neonate over the other.

Figure 9.6 shows the cumulative distribution of errors in the lowest 50% of samples in each noenate based on the overall NRMSE value, i.e. the NRMSE value of all signals combined. Two things are clear:

1. The overall NRMSE values for neonate neo021, the neonate with the milder injury, are generally lower than those for neonate neo007. This suggests a

potential difficulty in the model correctly reproducing behaviour for the more severe injury case.

2. The errors for neonate neo007 are less widely spread than those for neonate no021. This suggests that the model showed less variance across its potential outputs for the prior parameter distributions that were defined. This reinforces what is seen in Figure 9.5b, where the $\mu*$ values for neonate neo007 were generally a lot lower than those seen in neonate neo021.

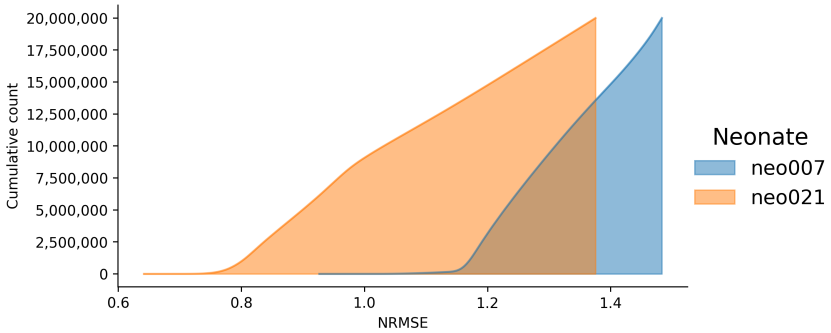

Figure 9.6: Cumulative distribution of errors for the bottom 50% of samples in each neonate, based on the overall NRMSE value.

9.3.3 All Signals

Posterior and Posterior Predictive Distributions

Final posterior distributions were calculated using an acceptance rate of 0.001% based on NRMSE values giving a final posterior size of 4000 samples. Posterior predictive distributions were calculated using a random 50% of the values in the posterior. The decision not to use all samples was due to larger sample sizes being unable to be plotted on available hardware. The type of NRMSE value used depends on the focus. We start here by looking at the NRMSE of all signals combined. As explained below, we will then focus on each individual signal and therefore the posteriors will be determined using the NRMSE values of that particular signal.

We start by looking at the posterior predictive distributions in order to assess the quality of the model fit across all signals. As discussed above, we will expect to see a difference in the qualuity of fit between the two neonates. It is clear to see from

Figure 9.7a and 9.7b that whilst the model is able to reasonably capture the haemodynamic behaviour in both neonates, it is unable to capture the metabolic response seen in the CCO signal. There is a reasonable amount of noise in the posterior predictive simulation of HbT for neo007 which reflects the same behaviour seen in the ABP signal in Figure 9.4b).

Though the metabolic signal for both neonates is clearly a poor fit, it is also apparent that the general behaviour of the simulated CCO signal for neo007 is closer to the measured signal, even if the absolute values differ. This may be due to baseline drift in the signal, something that was only apparent after fitting the model. A clear decrease in CCO is visible at the onset of the desaturation event and the subsequent recovery also occurs at a similar point as in the measured data. In contrast, the posterior predictive distribution of CCO for neo021 bears minimal resemblance to the measured signal. In fact, the measured signal shows an increase during desaturation whilst the posterior predictive shows a small but noticeable decrease. Figure 9.8

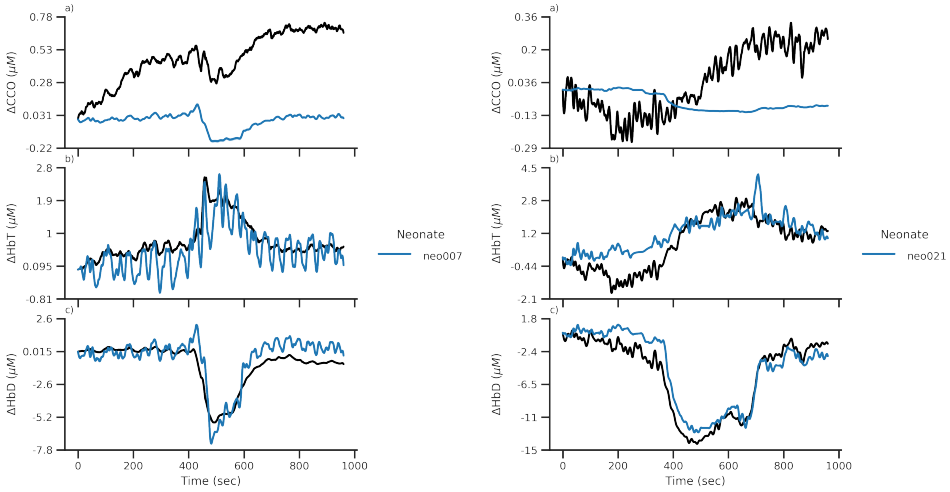

(a) Posterior predictive distribution for neo007 based on NRMSE of all signals.

(b) Posterior predictive distribution for neo021 based on NRMSE of all signals.

Figure 9.7: Posterior predictive distributions for neo007 and neo021 based on NRMSE of all signals.

shows the posteriors for each neonate in comparison to each other. When analysing these posterior distributions we need to remember the quality of the model fit for each neonate. With this in mind, the key thing to notice here is that both the median value of each distribution and the distributions themselves are almost identical

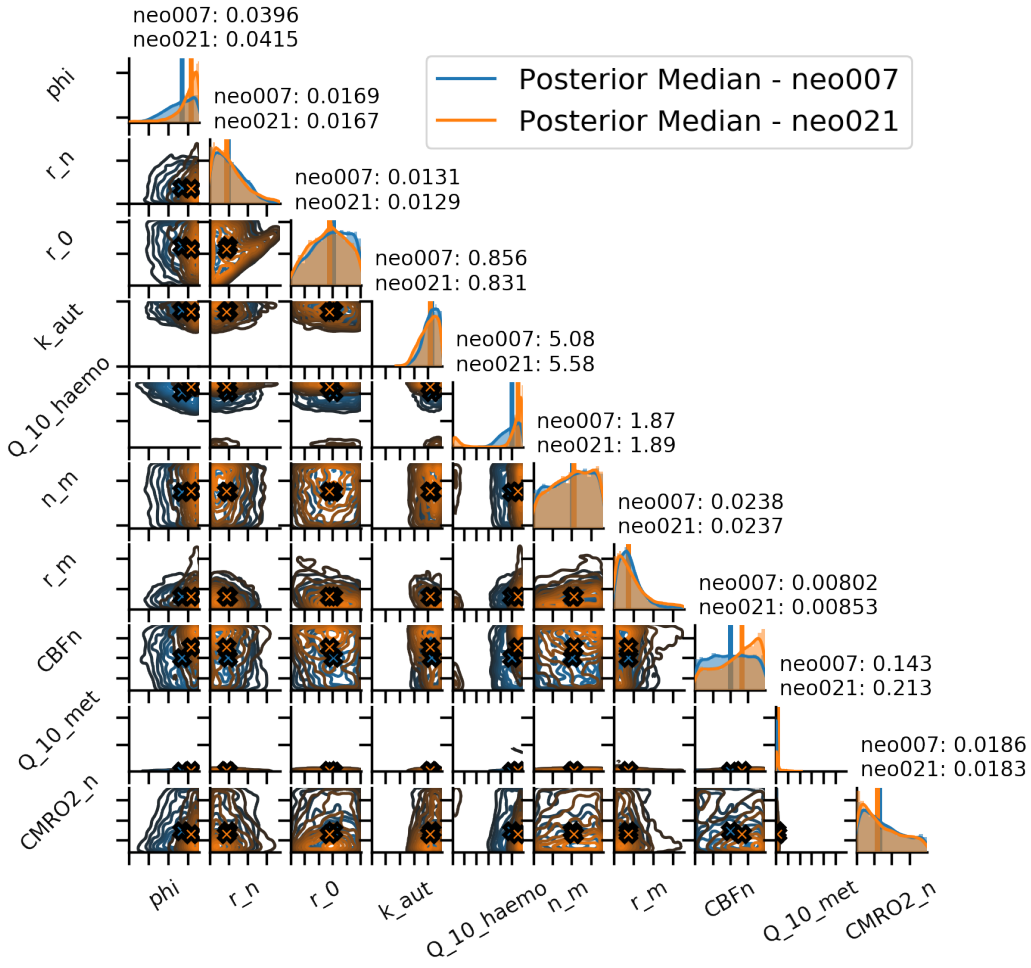

Figure 9.8: Posterior distributions for Neonates 007 and 021 based on the NRMSE of all signals.

across the two neonates. Given that the fit of the metabolic signal for neither neonate is good and that there was such a difference in the quality of the model fit between metabolic and haemodynamic signals, further investigation is needed. This can be achieved by generating posterior and posterior predictive distributions for each signal individually and then performing this same comparison.

Model Analysis

We begin by looking at the distribution of residuals for both neonates in Figure 9.9b. It's clear from this how poor the fit of CCO is in both neonates, with neither showing a distribution close to normal and a mean below zero in both cases. For

neonate Ne021 we see a particularly poor fit with a strongly bimodal distribution reflective of the initial over-prediction followed by under-prediction seen in Figure 9.7b. Both the HbT and HbD signal seen in each neonate have distributions much closer to being normally distributed around 0 as would be expected given the good quality of the fits seen in Figure 9.7. The distribution of HbD signal for neo021 appears to have a negligibly small, but notably present, set of residuals that are far below zero. Ignoring these values, the distribution is otherwise indicative of a good fit.

We can further investigate the residuals by looking at the Q-Q plots in Figure 9.9.

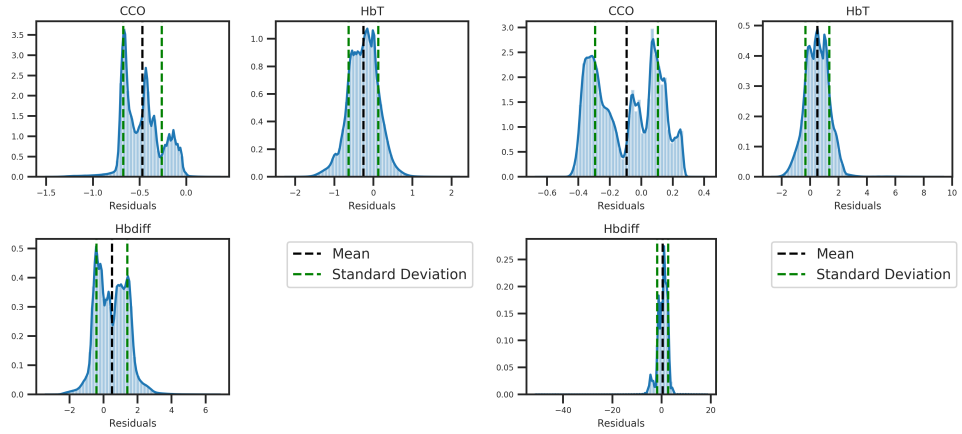

(a) **(b)**

Figure 9.9: Distribution of residuals for neo007 (Figure 9.9a) and neo021 (Figure 9.9b) when fit using NRMSE of all signals.

One of the clearest things to note here is that the Q-Q plots seen in Figure 9.10a suggest a generally better fit than those seen in seen in Figure 9.10b. CCO is generally along a diagonal but with a sample mean far below the theoretical mean, but both HbT and HbD are along the diagonal and with sample means close to the theoretical mean.

Figure 9.10b however shows that none of the fits for neo021 fit the diagonal particularly well ,even if the sample and theoretical means are close to each other. For CCO we see clear evidence of the bimodality also seen in Figure 9.9b. For HbT we can see that up until around quantile 3, the residuals lie generally along the diagonal, but between quantile 3 and 4 there is a sudden move away from the diagonal suggesting

a heavy right tail due to a larger spread in residuals than a theoretical, "good" fit would have. Finally, for HbD, we see a disconnect between quantile -3 and quantile -4. This is due to the small number of residuals that gave a heavy lower tail in the HbD residual distribution. Aside from this we see a slightly heavier upper tail, but with most residuals falling generally along the diagonal. Finally, we look at the K-L

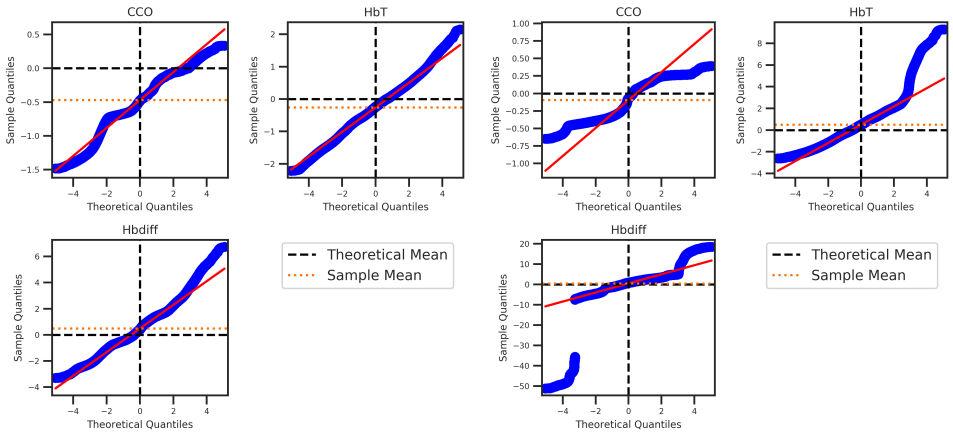

(a) **(b)**

Figure 9.10: Q-Q plots of residuals for neo007 (Figure 9.10a) and neo021 (Figure 9.10b) when fit using NRMSE of all signals.

divergence of the prior and posterior distributions for each neonate in Figure 9.11. The first thing to note is that the posteriors for each neonate are very similar with only a small difference in the K-L divergence values for each set of distributions. For each parameter, both neonates share a general trend, with the difference being by how much. For example, in both neonates we see a very tight posterior distribution for `Q_10_met` values, with both close to zero. The main difference seems to be that the distribution is slightly tighter for neo007 as indicated by the K-L divergence value of 3.09 vs 2.27 for neo021. The only parameter to show some kind of difference is `CBFn`, with neo021 having a K-L divergence approximately 16 times larger than neo007. However, even with this relative discrepancy, both neonates have K-L divergence values for `CBFn` that are fairly small.

Whilst we are not here able to specify a threshold above which a K-L divergence value indicates significant information gain, it is notable that only 4 parameters have values above 0.5 and of these, neo007 has only one greater than 1 (`Q_10_met`) and

neo021 has only two (`Q_10_met` and `Q_10_haemo`).

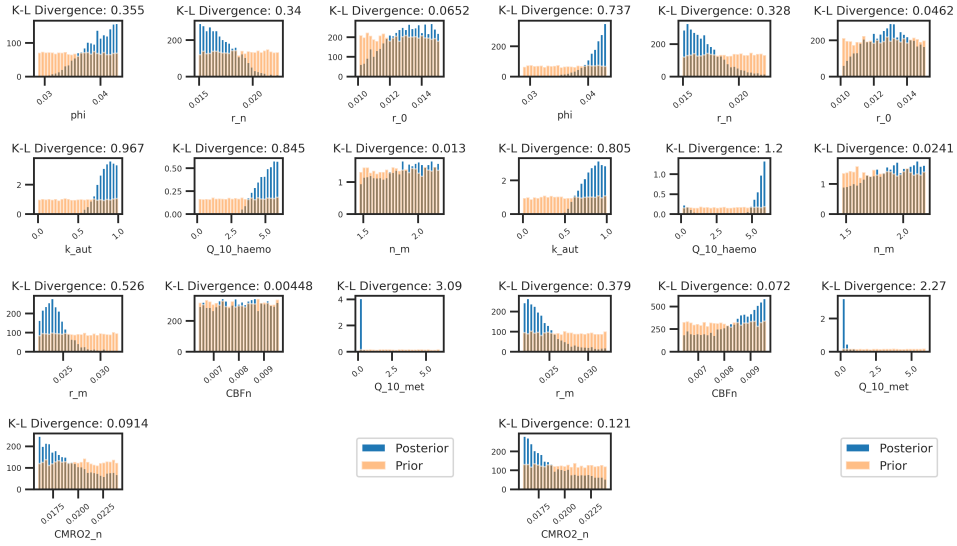

(a) **(b)**

Figure 9.11: K-L divergence plots comparing prior and posterior distributions for neo007 (Figure 9.11a) and neo021 (Figure 9.11b) when fit using NRMSE of all signals. Plots are given as separate sub-figures due to the close similarity of the distributions making it difficult to distinguish between the two neonates if given on the same axes.

Gradient between HbDiff and CCO

A final important check we can do to determine the quality of our simulation is to focus on the key indicator of injury severity, the gradient between ΔHbD and ΔoxCCO, as determined by Bale et al. [2018]. Figure 9.12 shows the distribution of gradients for all simulations in the posterior predictive distribution, with gradient calculated between the start of the desaturation event and the nadir.

It is clear to see that the distribution of gradients for neo021, the less severely injured neonate, is generally lower than that for neo007, with respective medians of 0.00732 and 0.0293. This is in line with the findings of Bale et al. [2018] that a larger gradient between ΔHbD and ΔoxCCO can be observed in the more severely injured neonate.

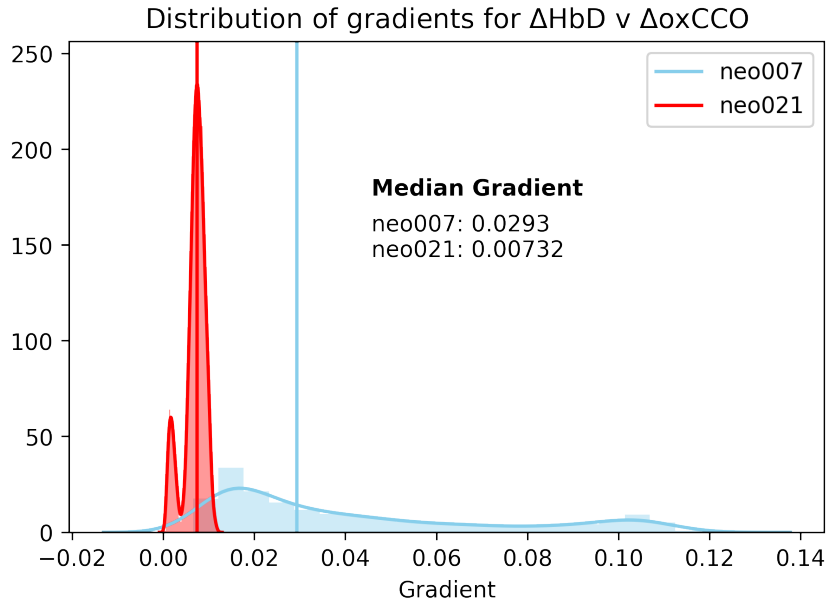

Figure 9.12: Distribution of gradients for simulations within the posterior predictive distribution. Gradient is calculated between ΔHbD and ΔoxCCO. The median values are noted and marked as vertical lines. The median gradient for neo021 is clearly much lower than that of neo007, indicating that the simulation was able to capture the important feature of a ΔHbD-ΔoxCCOgradient nearer to 0 for the less severely injured neonate.

9.3.4 HbD

Posterior and Posterior Predictive Distributions

We focus here on the HbD signal, which is a good proxy for the oxygenation of the tissue and therefore the NRMSE value used was the value for the HbD signal alone. Figure 9.13a and 9.13b show the posterior predictive distribution for this. We can see here that, as would be expected, the HbD signal is a good fit for both neonates. In addition, for neo007 the HbT signal is a reasonably good fit, whilst the same signal for neo021 is a poor fit. In both neonates, the CCO signal is a poor fit, as expected, indicating the previously suggested difference in the parameter state for haemodynamic and metabolic processes. For neonate neo021 the 95% confidence interval around the median is also much more noticeable suggesting that the posterior distributions for the HbD signal constrain the behaviour of the metabolic reactions less tightly for the mild injury than they do for the severe injury. Reasons for this are not fully clear but may suggest a less tight coupling between metabolic

and haemodynamic processes for the mild injury parameter space than in the severe injury parameter space. Figure 9.14 show the posterior distributions for each

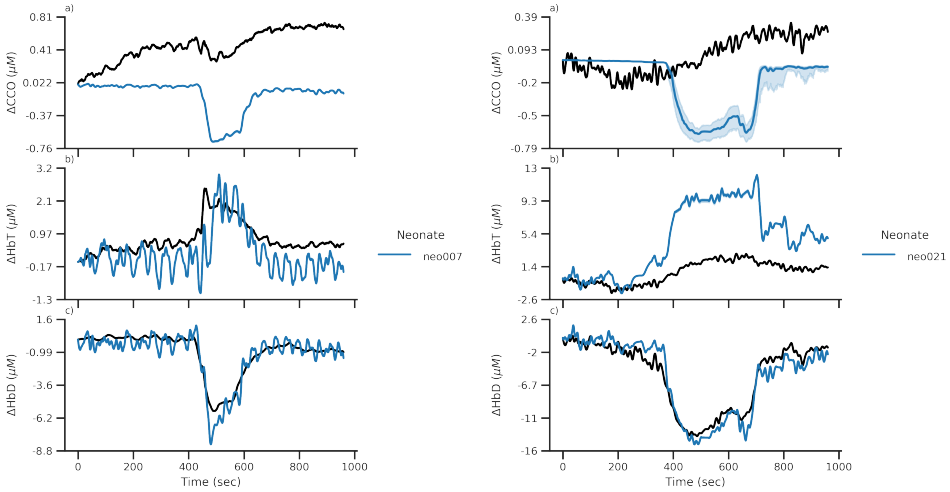

(a) Posterior predictive distribution for neo007 based on NRMSE of only the HbD signal.

(b) Posterior predictive distribution for neo021 based on NRMSE of only the HbD signal.

Figure 9.13: Posterior predictive distributions for neo007 and neo021 based on the NRMSE of the HbD signal.

neonate. In contrast to Figure 9.8, we do see a number of distinctly different parameter distributions for certain parameters. We also see differences in how tightly constrained different parameter spaces are for each injury state. `r_m` shows notably different distributions for each neonate, with neo007 having a much lower median and distribution suggesting a narrowing of the blood vessels in the more severely injured neonate. This is not a clinical conclusion however and there may be other effects not being accounted for that this behaviour is attempting to mimic. `CMRO2_n` is much more tightly constrained and has a lower median value for neonate neo021 than neonate neo007. In terms of the response to hypothermia, there is a small difference in the `Q_10_met` values but neither are particularly well constrained, which is expected given that we are focussing here on a haemodynamic signal. There is some difference in the `Q_10_haemo` distributions, with the neo021 distribution being more tightly constrained and with a higher median value, suggesting that a change in temperature for the mild injury state has a stronger influence on the haemodynamic reactions *when focussing on the HbD signal*. Comparing this observation

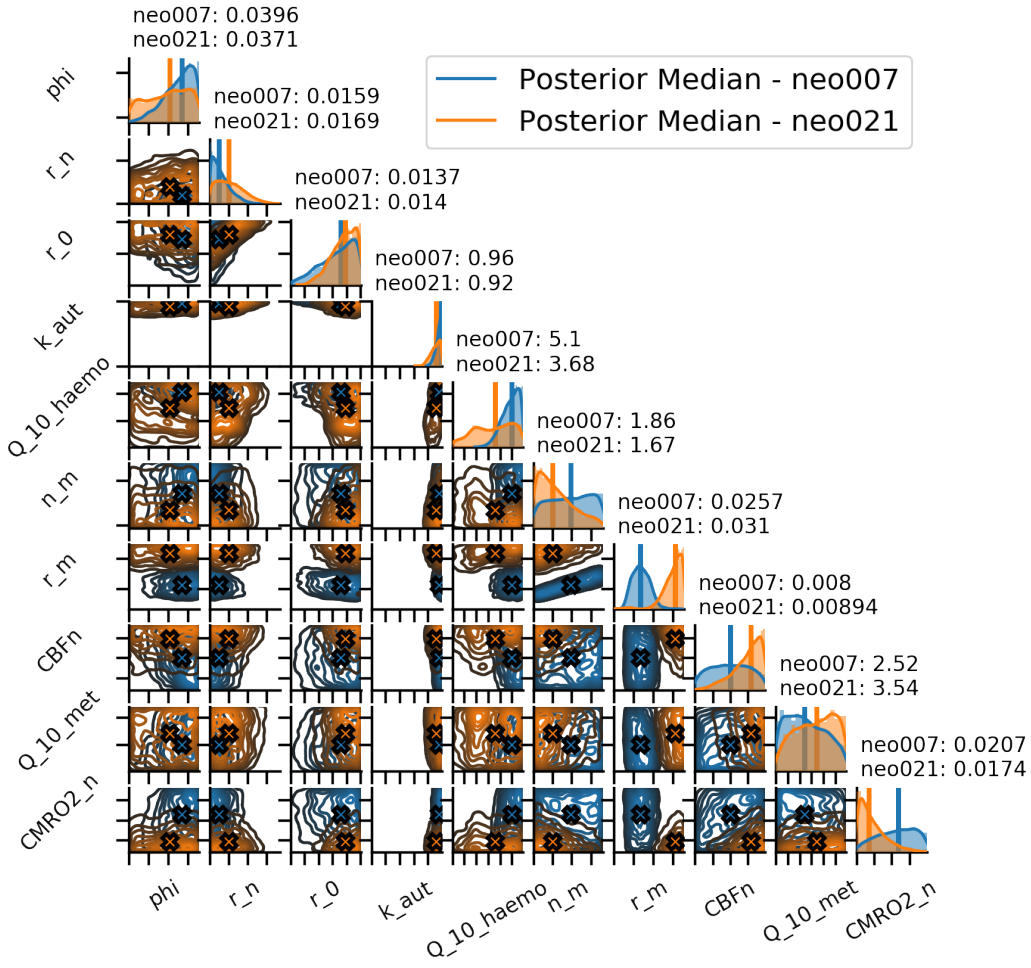

Figure 9.14: Posterior distributions for Neonates 007 and 021 based on the NRMSE of HbD.

against what we find for the HbT signal will give us a more reliable picture of the difference in the effect of temperature here.

Model Analysis

Figure 9.15 shows the distribution of residuals for the posterior predictive distributions of both neo007 and neo021 when fitting with the NRMSE of HbD. As we would expect, the distribution of the HbD signal is better than for either of the other two signals, but we also see that of the other two signals, CCO is clearly a much worse fit than HbT, particularly for neo021 where the distribution of CCO is clearly not at all normally distributed. This is further reinforced when we look at the Q-Q plots in Figure 9.16. Other than the HbD signal being a good fit in general here, we

also see that the plot for HbD in neo021 is missing the disconnect seen previously in Figure 9.10b, suggesting that it may have been an optimisation for one of the other two signals that caused the massive underfit in a small number of runs. We also see that when optimising for HbD, for both neonates the CCO signal is clearly not unimodal, indicated by the S-shaped Q-Q plots. The HbT signal is a reasonable fit for neo007 but for neo021 we see heavy right-skewing in both the distribution and the Q-Q plot reinforcing what was seen earlier in the posterior predictive. This could suggest that there is an underlying difference in the physiological behaviour of the two neonates meaning that the HbD and HbT signals are likely controlled by similar mechanisms in neo007 but not in neo021. This would be indicative of a coupling between blood volume and oxygenation in the more severely injured neonate.

Figure 9.17 shows the K-L divergence values and comparisons between prior and

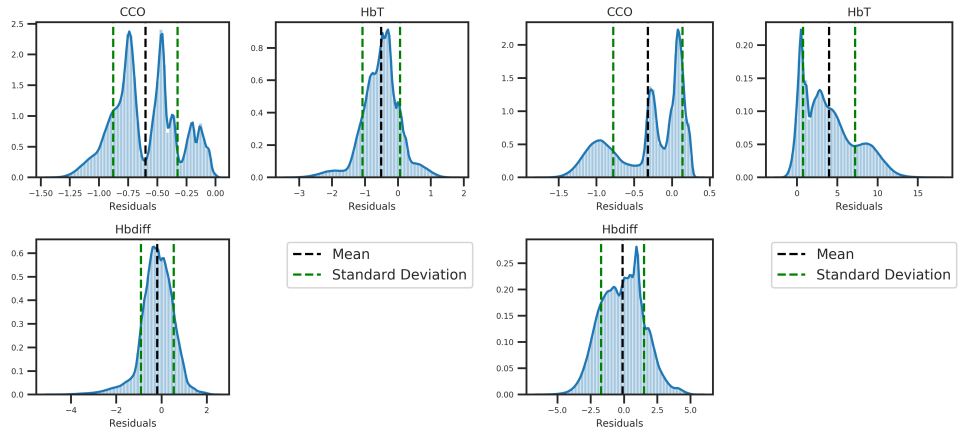

Figure 9.15: Distribution of residuals for neo007 (Figure 9.15a) and neo021 (Figure 9.15b) when fit using NRMSE of the HbD signal.

posterior distributions. In contrast to when fitting using the NRMSE of all signals, there are a number of distributions that vary massively between the two neonates. In neo007, `phi` and `Q_10_haemo` show much more information again than it does in neo021, whilst `n_m`, `CBFn` and `CMRO2_n` all show much higher information gain in neo021 as compared to neo007. This difference may potentially explain some of the difference in the HbT fit between the two neonates, with the difference in sensitivity and distribution in the posteriors for the two affecting the HbT signal

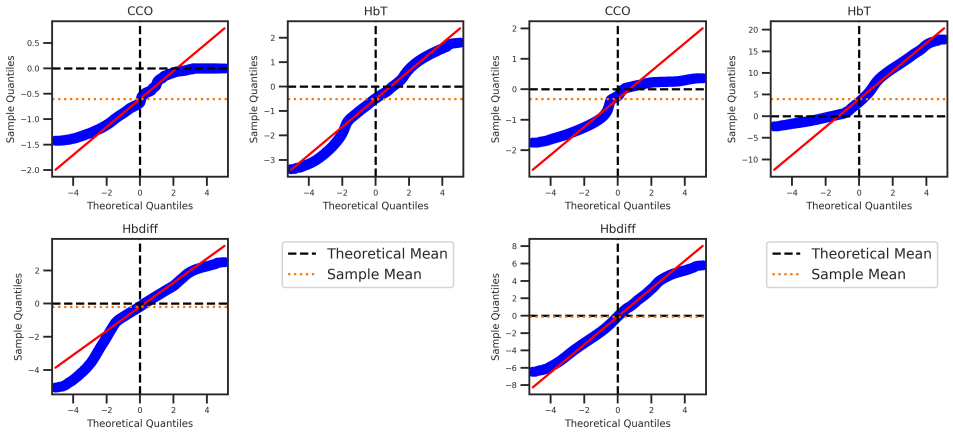

(a) **(b)**

Figure 9.16: Q-Q plots of residuals for neo007 (Figure 9.16a) and neo021 (Figure 9.16b) when fit using NRMSE of the HbD signal.

differently. Aside from this difference in information again, the two neonates show similar trends in terms of distribution skew. Where both neonates show information gain as compared to using the prior, only the `r_m` parameter shows a difference in terms of skew, with neo007's posterior being approximately normally distributed around 0.025, whilst neo021 has a posterior skewed heavily to the right and pushed up against the upper boundary of the prior. Other parameters such as `r_n` or `r_0` have similar distributions across the two neonates.

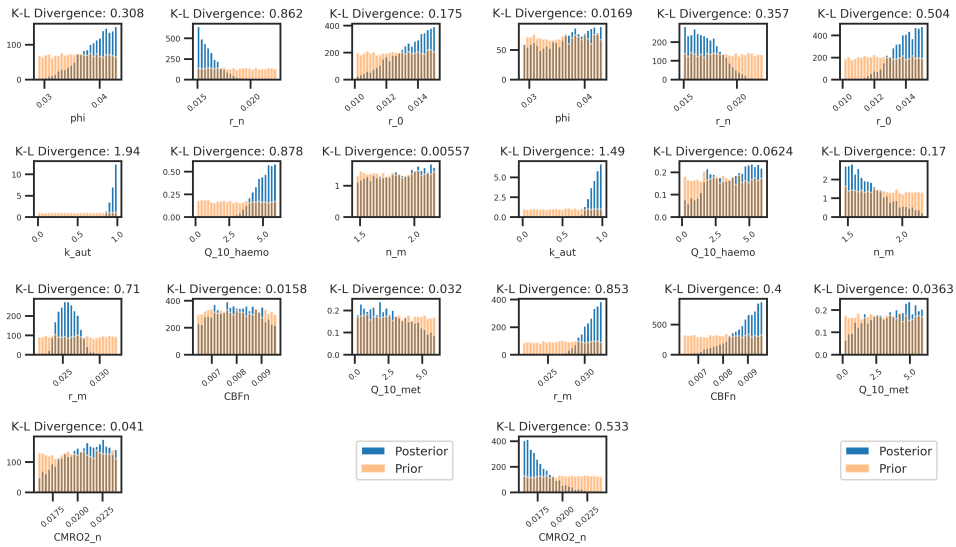

(a) **(b)**

Figure 9.17: K-L divergence plots comparing prior and posterior distributions for neo007 (Figure 9.17a) and neo021 (Figure 9.17b) when fit using NRMSE of HbD.

9.3.5 HbT

Posterior and Posterior Predictive Distributions

Here the NRMSE value used to determine the posterior was the value for the HbT signal alone. Figure 9.18a and 9.18b show the posterior predictive distributions for neonates neo007 and neo021 respectively. We see a very different set of results to those seen in Figure 9.13a and 9.13b. Notably, whilst the HbT signal is a good fit in both, the HbD signal is a better fit for neo021, the mild injury state neonate. This is the inverse of what was seen when fitting for HbD. HbD for neo007 does show some of the time period but it predicts a much larger desturation event than what was actually observed.

The CCO signal shows very different behaviour in each neonate, with neo007 showing almost no response whatsoever whilst neo021 shows a decrease around the time of desaturation that fails to recover. Also, in contrast to when optimising for the HbD signal, this time the CCO signal for neo021 is tightly constrained with a negligibly small 95% confidence interval.

Looking at the posterior distributions in Figure 9.19 we see very different distri-

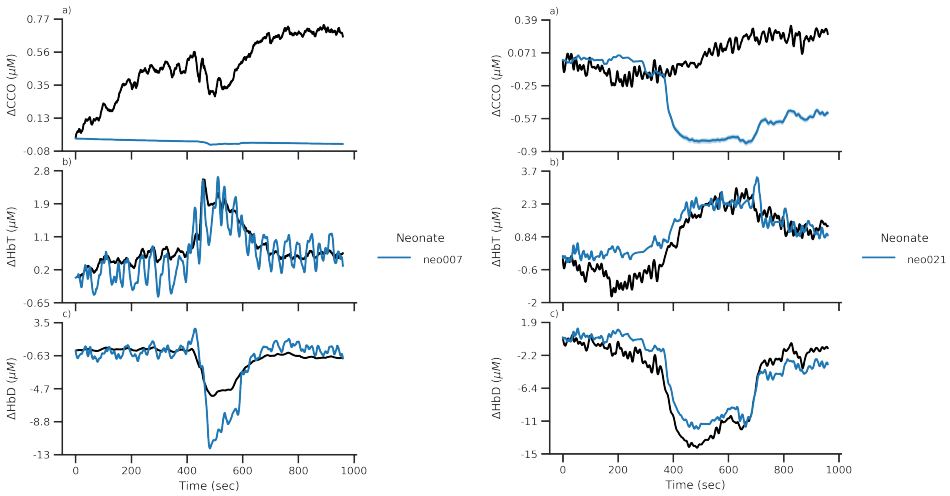

(a) Posterior predictive distribution for neo007 based on NRMSE of only the HbT signal.

(b) Posterior predictive distribution for neo021 based on NRMSE of only the HbT signal.

Figure 9.18: Posterior predictive distributions for neo007 and neo021 based on the NRMSE of the HbT signal.

butions to those seen in Figure 9.14. For example, whilst optimising for HbD, the

CMRO2_n distribution for neonate neo021 was generally grouped towards the lower bound of the prior distribution, with a median of 0.0174 mM s^{-1}, but when optimising for HbT we instead see a distribution grouped towards the upper bound and with a median of 0.0218 mM s^{-1}. We also see that in contrast to the HbD posteriors, both of the Q_{10} parameters show differing distributions between neonates, rather than just the haemodynamics reactions parameter. Q_10_met in particular here shows very different distributions for the two neonates which goes some way towards explaining the very different CCO signals seen. For neonate neo007, the Q_10_met distribution and median is biased heavily towards higher values, whilst for neo021 the distribution is biased heavily towards lower values, with a median of 0.78. We can also see that the posterior parameter distributions for the normal cerebral blood flow, CBFn, are biased towards opposite ends of the prior distribution range for each neonate, with neo007 biased towards the upper bound and neo021 biased towards the lower bound.

Model Analysis

Figure 9.20 shows the distributions of residuals for each neonate when fitting against the NRMSE of HbT. It is clear here that the HbT distribution in each neonate is roughly normally distributed around zero for each neonate, but CCO and HbD are both poorly fit. This is also seen in the Q-Q plots in Figure 9.21. All signals in neo007 show platykurtic behaviour in the Q-Q plots, with the distribution of CCO being clearly non-unimodal in Figure 9.20a. This same behaviour is seen in the CCO signal of neo021. The HbT and HbD signals of neo021 appears to be reasonably well distributed aside from a small number of outlying residuals, as seen previously in the HbD signal in Figure 9.9b. This is seen clearly in the Q-Q plots for both signals in Figure 9.21b where the residuals generally lie on the diagonal until around quantile 3 after which they move sharply up and away. Looking at K-L divergence phi, Q_10_haemo, CBFn, Q_10_met and CMRO2_n all show the opposite behaviour in neo021 as compared to neo007. This would suggest that there may be a distinct difference between the two neonates in terms of how these parameter affect HbT and therefore the cerebral blood volume. Aside from this subset of

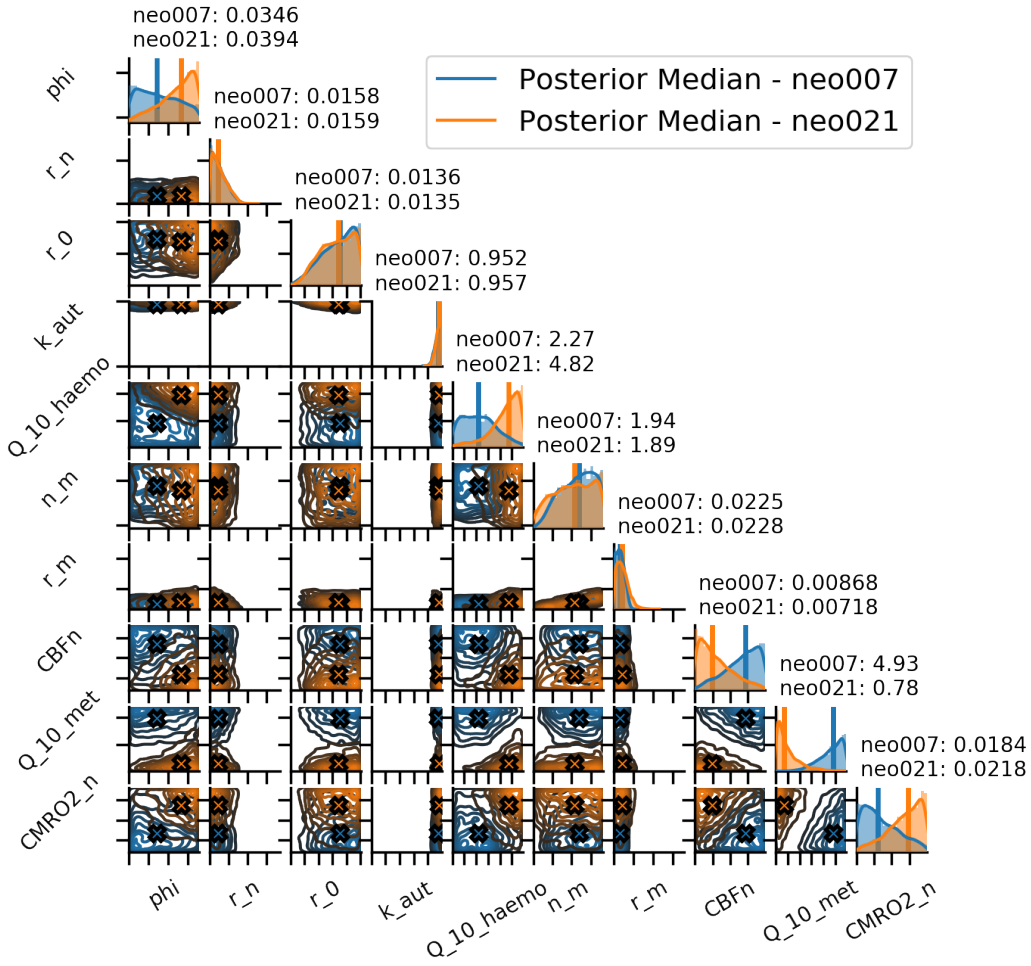

Figure 9.19: Posterior distributions for Neonates 007 and 021 based on the NRMSE of HbT.

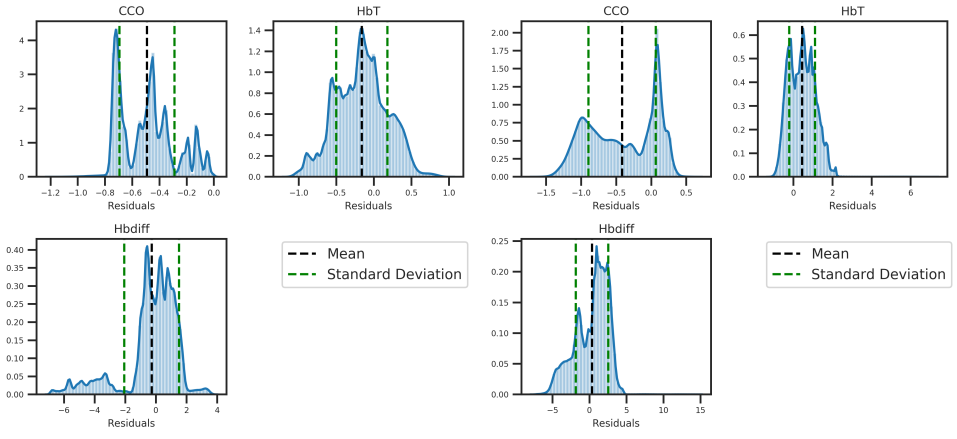

(a) **(b)**

Figure 9.20: Distribution of residuals for neo007 (Figure 9.20a) and neo021 (Figure 9.20b) when fit using NRMSE of the HbT signal

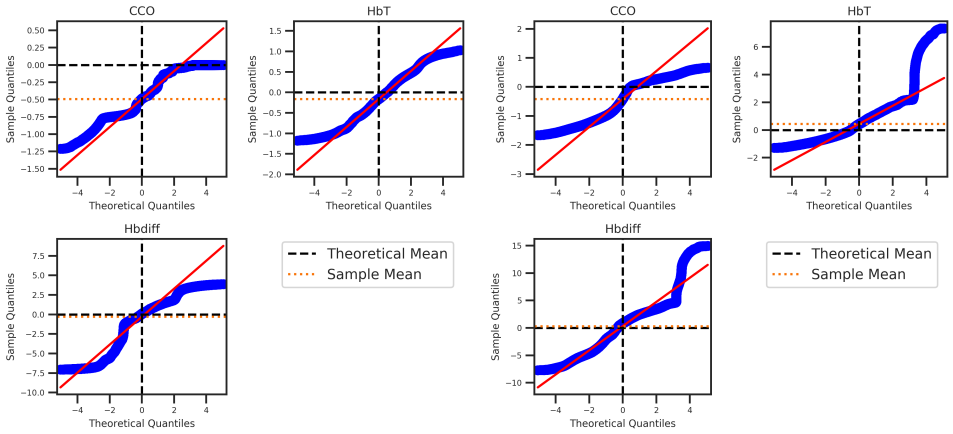

(a) **(b)**

Figure 9.21: Q-Q plots of residuals for neo007 (Figure 9.21a) and neo021 (Figure 9.21b) when fit using NRMSE of the HbT signal.

parameters, the K-L divergence values and the posteriors are similar between the two neonates indicating a potentially shared parameter joint distribution. This may prove useful in a hierarchical Bayesian model where in a shared parameter space can be used to determine parameter distributions at a more general level, allowing the model fitting process to focus on fitting parameters that are more specific to a given strata of the hierarchy e.g. at a species, age or outcome level.

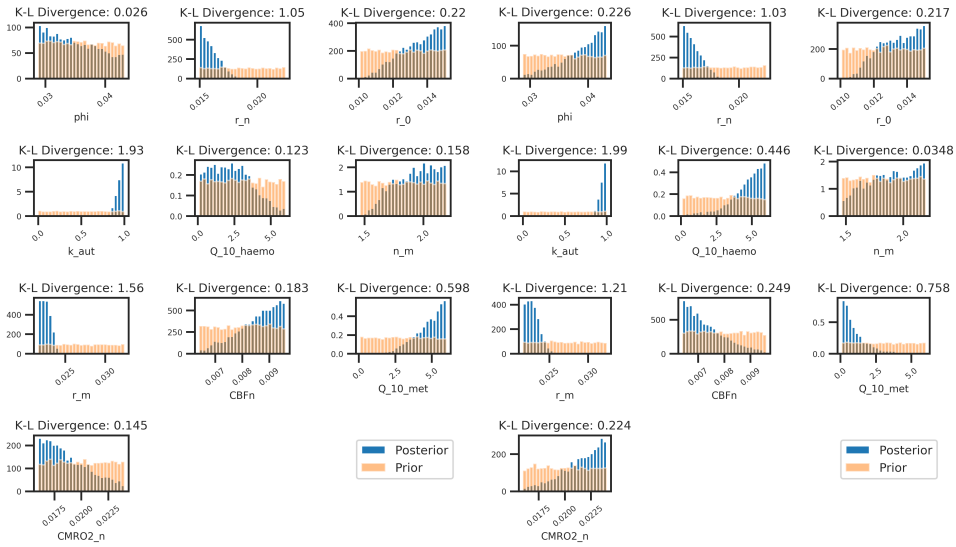

(a) **(b)**

Figure 9.22: K-L divergence plots comparing prior and posterior distributions for neo007 (Figure 9.22a) and neo021 (Figure 9.22b) when fit using NRMSE of HbT.

9.3.6 CCO

Posterior and Posterior Predictive Distributions

Finally, we consider the CCO signal which was reproduced poorly in both neonates when using all signals to generate our posterior. Figure 9.23a and 9.23b show the posterior predictive distributions. Interestingly, when we focus on optimising our parameters to reproduce the CCO signal only, we still fail to produce a good simulation of the CCO signal. Neonate neo021 shows almost no improvement in its simulation despite us focussing solely on that signal, whilst neonate neo007 shows some improvement but still under predicts the values. The main thing that we can conclude here than is that the reason for the poor simulation is not that the posterior distributions in Figure 9.8 were trying to compromise between haemodynamic and metabolic signal and failing to cover both, but that we are not optimising the correct parameters during our optimisation step. This is likely due to a failure by our model reduction step in identifying the parameter most important for simulating our desired behaviour. One thing that is interesting to note here however is the difference in the haemodynamic simulations when using the parameter spaces for the two neonates. Even though the parameter space identified for neo021 failed to accurately reproduce our CCO signal, it has given a much better approximation of the HbT and HbD signals than in neo007. This may be because for the mild injury state and for the parameters specified here, mechanisms that provide the best possible fit for the CCO signal (as poor as it actually is) are similar to those for a good simulation of the haemodynamic processes. For the severe injury state however, the mechanisms that control the CCO signal are notably different to those for the haemodynamic processes, and therefore parameters and how they influence the underlying mechanisms are also very different. Figure 9.24 shows us the posterior distributions for the optimised CCO signal. One very important thing to note when considering these is that the posterior predictive distribution for both neonates was a poor fit. Therefore we cannot analyse these posteriors in terms of what we can reliably predict the underlying physiology to be doing, but instead we can only say what the parameter space was in order to generate the *best possible fit available*

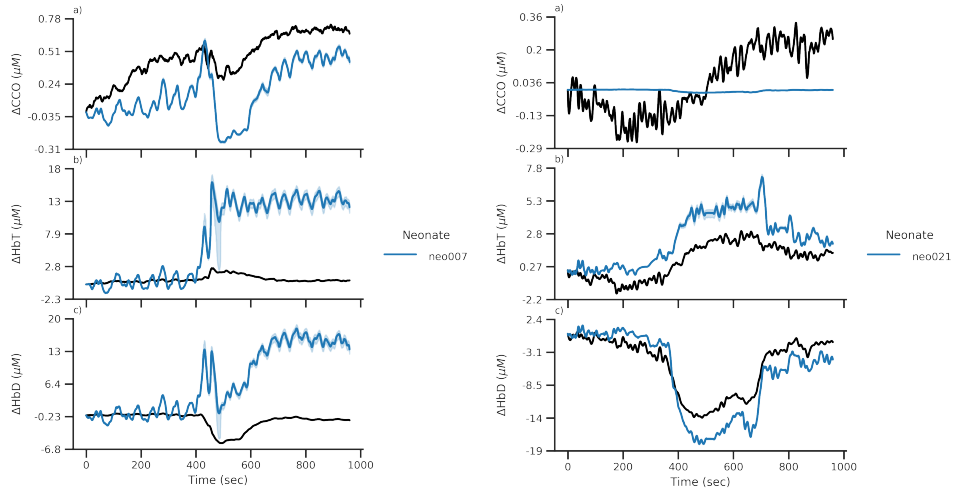

(a) Posterior predictive distribution for neo007 based on NRMSE of only the CCO signal.

(b) Posterior predictive distribution for neo021 based on NRMSE of only the CCO signal.

Figure 9.23: Posterior predictive distributions for neo007 and neo021 based on the NRMSE of the CCO signal.

given the model and parameters used.

One important difference in parameter distributions can be seen in contrast the posterior distributions of the heamodynamics signals: `k_aut` is much lower in neo007 than in neo021, suggesting that the ability to autoregulate needs to be significantly reduced in order to obtain the best CCO simulation possible. It is also notable that both neonates have very similar `Q_10_met` distributions despite having very different response in the CCO signal, which is a proxy for metabolic activity. Instead, we see that the `Q_10_haemo` distributions are different here, with the posterior for neo021 being heavily biased towards the lower bound and with a median of 0.138.

Model Analysis

The first thing to note for the residuals and Q-Q plots when optimising for CCO is that the HbD and HbT plots show some extremely outlying residuals. In Figure 9.25a and Figure 9.25b this is indicated by a near invisible distribution due to the x-axis scale whilst in Figure 9.26a and Figure 9.26b is is seen in a line which veers massively away from the diagonal at the fourth quartile.

Focussing on the CCO signal we can see that the distribution and Q-Q plot for

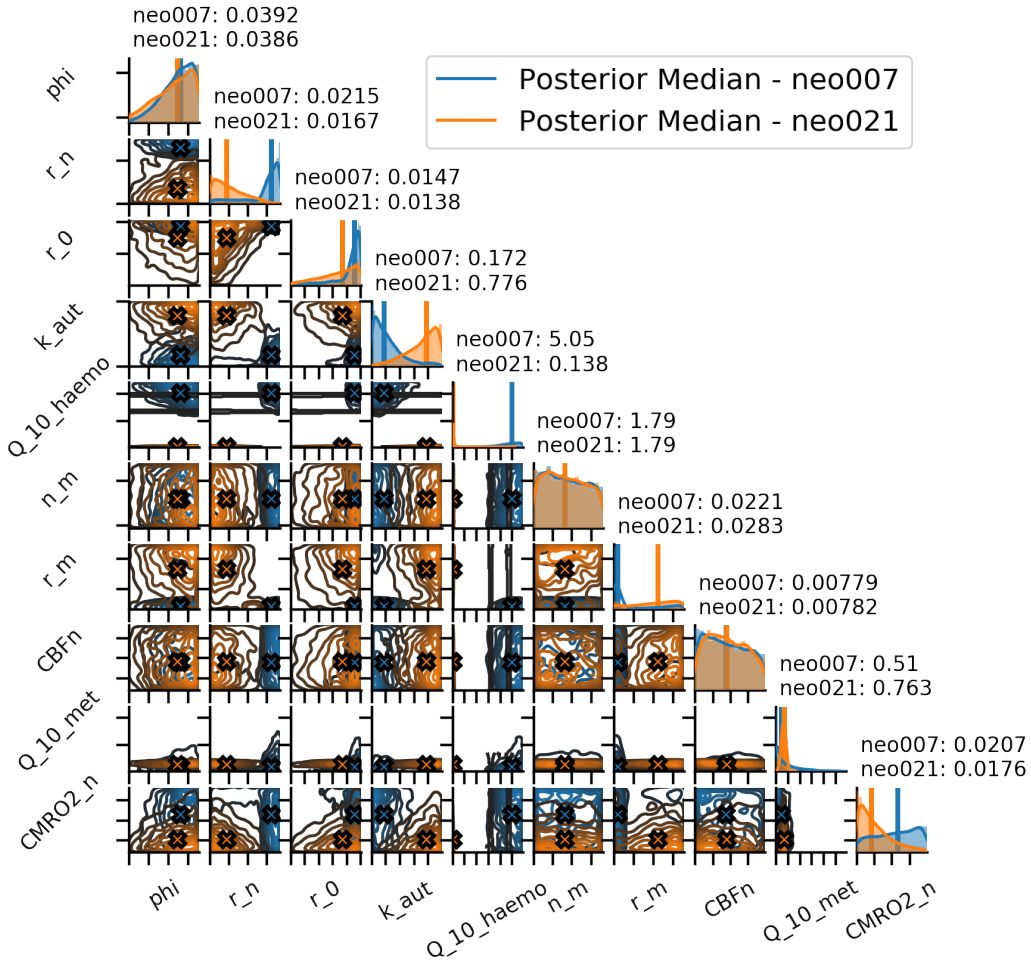

Figure 9.24: Posterior distributions for Neonates 007 and 021 based on the NRMSE of CCO.

neo007 shows a reasonably good distribution with a median slightly below 0. For neo021 however we see a bimodal distribution centred around 0. This is likely due to the fact that we are unable to capture the true dynamics of the measured signal, as seen in Figure 9.23b, and we both overestimate the signal at early time points and underestimate at later time points.

Looking at the K-L divergence plots in Figure 9.27, `Q_10_met` and `Q_10_haemo`) show very narrow distributions in both neonates with reasonably high K-L divergence values suggesting the signal is sensitive to changes in these parameters in all injury states. In contrast, `CBFn` and `n_m` have low K-L divergence values in both neonates and are visibly similar to the uniform prior, suggesting that in terms of

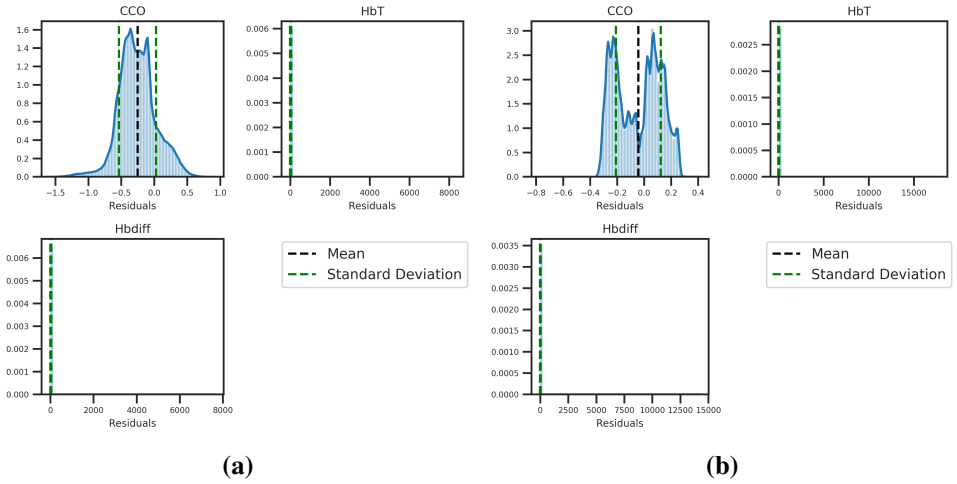

Figure 9.25: Distributions of residuals for neo007 (Figure 9.25a) and neo021 (Figure 9.25b) when fit using NRMSE of the CCO signal.

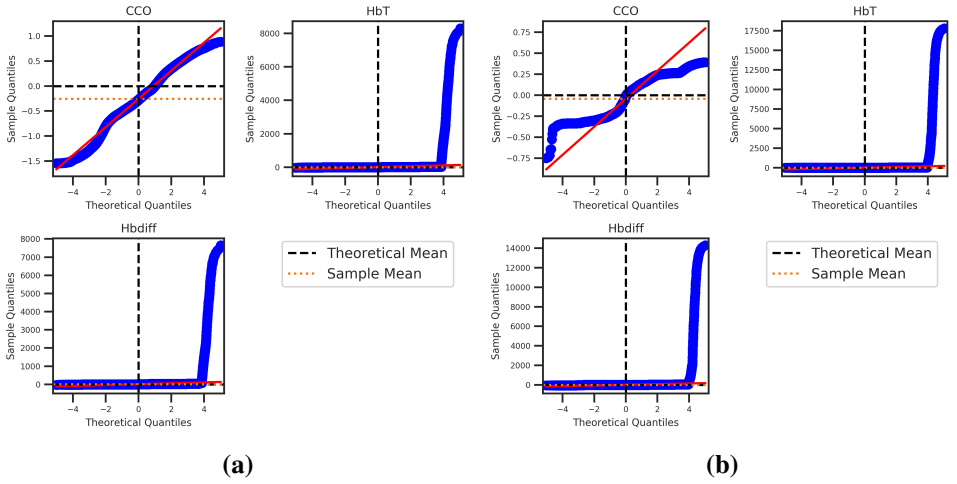

Figure 9.26: Q-Q plots of residuals for neo007 (Figure 9.26a) and neo021 (Figure 9.26b) when fit using NRMSE of the CCO signal.

obtaining a good fit for the CCO signal alone they provide minimal information. Considering differences in the two neonates, a notable difference is seen in `CMRO2_n`, with neo007 showing less sensitivity than neo021 suggesting a possible difference in the importance of this parameter based on injury state. Conversely, `r_m`, the vessel radius at maximal muscular tension, has a high K-L divergence value in the more severe injury state suggesting a higher sensitivity to changes in this parameter. There are also a number of parameters that show opposite skew. `r_n` skews to the upper end of the prior for neo007 and towards the lower end for neo021 indicating a larger *normal* blood vessel radius in the more severely injured

neonate, with values close to the radius of maximal muscular tension. `k_aut` skews lower for the more severely injured neonate indicating a more impaired autoregulatory mechanism. Finally, `Q_10_haemo` is skewed towards higher values in the more severely injured neonate and towards lower values in the less severely injured neonate. This would indicate a higher sensitivity to changes in temperature for haemodynamic reactions. However, given that we are here focussing on the metabolic reactions when fitting this may not be an important parameter to focus on. In terms of the most direct interpretation, given that the neonates were cooled to below normal body temperature, it would suggest that in order to mimic the CCO signal for the severely injured state haemodynamic reactions needed to be slowed down. It may be that fitting the diffusion rate parameter directly would reduce the differences in `Q_10_haemo` distributions for the two neonates. This again highlights the importance of identifying the correct parameters for fitting.

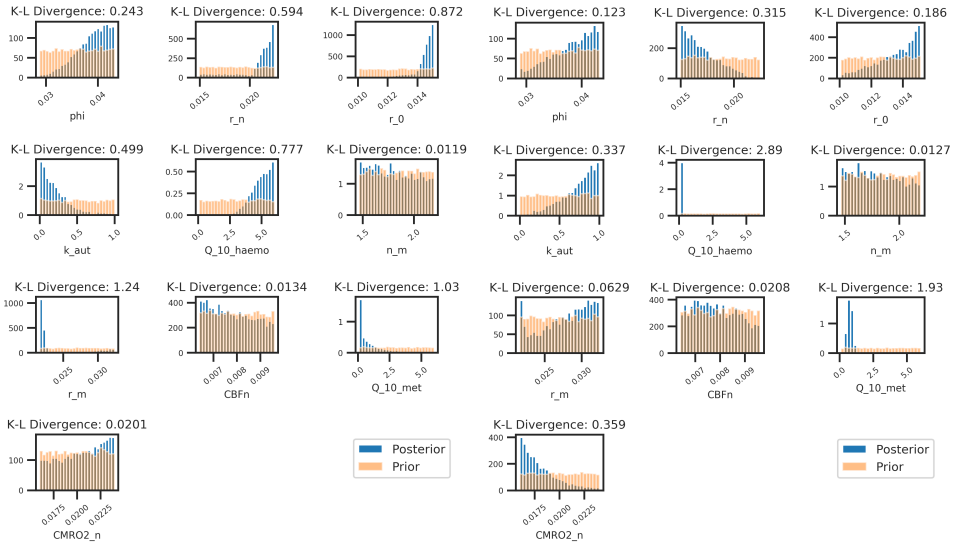

(a) **(b)**

Figure 9.27: K-L divergence plots comparing prior and posterior distributions for neo007 (Figure 9.17a) and neo021 (Figure 9.17b) when fit using NRMSE of CCO.

9.4 Discussion

The purpose of using the new framework with a systems biology model is to understand more about the underlying physiology than a simple black box model could give us. It is the analysis process that provides us with more information

and here this process was able to give us a good amount of information about what is happening. We were able to look at how the behaviour of the model and our obtained posterior distributions varied between signals. Without drawing any final conclusions, one thing that did prove interesting to consider was how a posterior parameter space obtained when optimising for a specific signal affected the other two signals that weren't being optimised for.

Further, our analysis of the residuals and the posterior distributions for each individual signal has allowed us to look at how focussing on different signals affects our posterior parameter distributions and how the different signals may be interconnected. For example, when fitting against the NRMSE of HbD alone we were able to observe that the ability to fit HbT with the same distributions differed in the two neonates, which allows us to conclude a difference in the two neonates in terms of how closely related the mechanisms behind both blood volume and blood oxygenation are. We were also able to see across all the different optimisation metrics that some parameters both differed from the prior and were similar between the two neonates. This suggests the potentials for future work to focus on developing a hierarchical model where in specific parameter distributions are determined to be shared at a certain hierarchy e.g. species (human/pig) or age (adult vs neonate). This would potentially allow for more accurate modelling dependent on the desired outcome. In this particular case, by using but not fitting the the shared distributions, computational effort could instead have been used to optimise parameters neglected here. This may have led to a better fit and therefore allowed us to analyse on a much more quantitative and granular level the differences in physiology between the two neonates.

Finally, by considering the distribution of gradients for the simulations fit using NRMSE, we were able to reproduce the work of Bale et al. [2018] in finding that the system replicating the mild injury state of neo021 didin deed have a smaller gradient between ΔHbD and ΔCCO than neo007. It is by considering these important features that we can assess the usefulness of these models. Where they were unable to reproduce signals exactly they were still able to reproduce important

physiological features. This serves to remind us that a useful model is not always the one that is able to reproduce a system exactly but one that can tell us useful information based on simplifications or approximations.

The purpose of these models is not to simply provide an accurate simulation alone. We were able to determine distinct parameter spaces for each neonate, particularly when isolating a single signal. For example, when isolating HbT, both `CBFn` and `CMRO2_n` are skewed towards opposite ends of the prior for each neonate. Given the significance of the physiological quantities these represent, the difference in their marginal posterior distributions in order to reproduce HbT reinforces the work by Bale et al. [2018] where they hypothesise that a reduction in $CMRO_2$ is likely to cause a larger reduction in CCO for a given change in SpO_2. The use of the systems biology models both reinforces the hypotheses and findings developed from experimental work as well as helping to draw attention to possible new avenues for research. Again looking at the HbT fitting, we were able to identify a separation in the marginal posteriors for both Q_{10} temperature co-efficients. This difference leads to the hypothesis that, when focussing on cerebral blood volume, there is a difference in how the two neonates responded to treatment by therapeutic hypothermia. The more severely injured neonate will have a greater reduction in the rate of metabolic reactions due to a given reduction in temperature than the less severely injured neonate. Conversely, the same reduction in temperature will reduce the rate of haemodynamic reactions in the mildly injured neonate more than in the more severely injured neonate. This could help to inform research around the effectiveness of different treatments.

9.4.1 Limitations

When considering the shortcomings of the model and approach, there are two key areas for improvement. Firstly, the CCO signal data are much noisier then the haemodynamic data. This difference in noise could partially explain the comparative failure in reproducing the measured CCO signal by simulation as compared to the haemodynamic signals. Better pre-processing could help reduce the disparity in

the different signal types and improve the ability to draw conclusion using the systems biology approach. For example, there appears to be a drift in the CCO data for neo007 which could potentially have been removed in a pre-processing step. Figure 9.28 shows the effect of "detrending" the CCO signal to remove drift. There appears to be an improvement in both neonates with the detrended signal closer to the posterior predictive than the original signal. This “improvement” is only clear having already produced the posterior predictive. Without this simulated signal to compare to it is hard to determine if the drift seen is physiologically significant or a result of measurement error. This is itself a strong point for using systems biology models to assess the quality of measured data and to identify measurement artefacts.

Secondly, we can see that the model in its reduced form is unable to properly

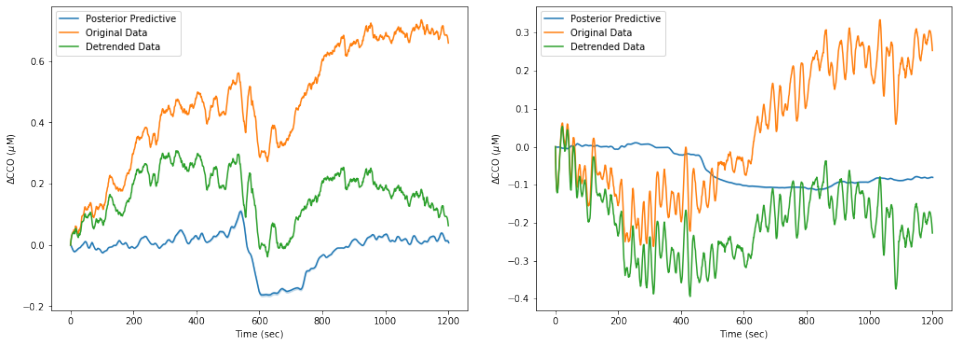

(a) Comparison of the posterior predictive, original and detrended CCO signal for neo007.

(b) Comparison of the posterior predictive, original and detrended CCO signal for neo021.

Figure 9.28: Detrended CCO signals for both neo007 and neo021. An improvement can be seen for both neonates with the detrended data closer to the simulated data. This assumes that the drift seen was measurement error as opposed to something physiologically significant.

capture the metabolic behaviour seen in the CCO signal. This is likely due to the reduced model’s parameter selection not including the parameters required to properly model the behaviour seen. Therefore, one possible place for improvement would be to focus on the model reduction step. The fact that a number of parameter identified as being sensitive are then found to have such small K-L divergence values as compared to the uniform prior suggests that their actual importance in the fitting process is not being correctly identified by the current model reduction step.

This may be a poor choice of summary metric, or it may point to the need for more data from more events in order to determine a sufficiently focussed set of parameters for the reduced model. An improvement here would lead to a model that is better focussed on the behaviour being considered. This improvement could come from better summary statistics in the sensitivity analysis method, the use of something like a Bayesian hierarchical model in which only a subset of all parameters need to be considered, or a combination of the two.

In conclusion, we have demonstrated here that there is indeed a difference in the parameter spaces of the two neonates, indicating a significant difference in the effect of changes in arterial oxygen saturation and temperature depending on injury severity. The hypothesis of Bale et al. [2018] that there is a difference in a reduction in $CMRO_2$ is likely to cause a larger reduction in CCO for a given change in SpO_2 for a more severely injured neonate. The difference in the normal $CMRO_2$ between the two neonates when fitting HbT is a sign of this difference. We are also able to establish from the new analysis process that an improvement in the initial model reduction step would likely provide a better fit, with a number of parameters showing minimal information gain after fitting. Before the introduction of the new process, this would not have been possible to identify. We also see from Figure 9.28 that the modelling framework gives us the ability to identify behaviour in the data that are likely to be a result of measurement error rather than a physiological process. In this case, we were able to identify drift in the data, which is easily remedied once identified.

Chapter 10

Conclusions

10.1 Summary

In this thesis I have developed and presented both a novel software tool for analysis of cerebral measurement data via systems biology and a new model for analysis of data collected from patients that are undergoing treatment via therapeutic hypothermia. The combination of these two developments allows for a new insight into the mechanisms behind differing outcomes following treatment. Additionally, the software allows for a new critical analysis of models and their predictive and analytic capabilities. By taking a Bayesian approach to model analysis, a fuller and more useful interpretation of models is available. By using the posterior predictive distribution, we are able to obtain a better understanding of what the model fitting process can tell us, forging a better critical analysis of the pre-processing step of sensitivity analysis and obtaining a fuller and more revealing picture of the mechanisms behind the observed data.

Biology is complex and consists of processes, systems and mechanisms that operate at different levels in terms of both time and scale, which then come together to produce holistic behaviour we can then observe. Mathematical and computational models are powerful tools for starting to understand these different scales and systems and how they interact. However, a model alone is not enough. It is how we develop, use and analyse this model that gives it utility. By developing new and better tools for doing this, we can start to produce more and better models that will

give better and more reliable insight into the data we can collect, and it is from that point that we can use these data and this understanding to provide better treatment. There is a great history in modelling the brain and its processes, as was outlined in Chapter 4, and all provide different perspectives and use cases. The purpose of building a model is also incredibly important. Is the goal to perfectly reproduce the data that are measured or is it to provide a better insight in conjunction with data? I would argue that it is the latter that is the driving force behind modelling, with the former being a consideration when assessing the quality or accuracy of the model. The UCL family of models are a part of this history and also is the ways in which they are analysed. BRAINCIRC was both reduced and extended to produce the BrainSignals model, which was then analysed using the BCMD framework and used to understand the human adult cerebral metabolism. Similarly, the first BrainPiglet model was then developed from the BrainSignals model to simulate the cerebral metabolism of a piglet. Finally, BrainPigletHI has then extended the BrainPiglet model to incorporate temperature in order to understand treatment of injury via thereapeutic hypothermia. It has then been used with the new BayesCMD framework to develop a new and better understanding of how the models can be used to understand the biological mechanisms. From the original BRAINCIRC model through to the BrainpigletHI model developing in this thesis, this interplay between the aims and use cases of the models, the tools used to analyse them and the available measurement techniques has driven the research forwards.

The redevelopment of the BCMD tool in Chapter 6 lays the groundwork for ensuring that this tool remains usable for future work and well as making it more accessible through the WeBCMD interface. This is the first step towards making this tool and the complex models is integrates with more available and usable. It also lays the groundwork for the incorporation of Bayesian techniques. The Bayesian approach allows for a better understanding of the problems space than previous techniques. Some work had been undertaken in Caldwell et al. [2015b] in an attempt to better understand the existence of a parameter space, as opposed to a single point estimate, but the Bayesian approach outlined in this thesis is a significant improvement

on this approach. Two main benefits are clearly visible from the work undertaken here. Firstly, the combination of visual checks of posterior parameter distributions with more quantitative methods such as K-L divergence provides much more insight into both the sensitivity of the model behaviour to a given parameter than sensitivity analysis prior to model fitting. Whilst sensitivity analysis prior to fitting will remain a key step in the process of model reduction, by looking at the shape of the posterior distribution and the information gain from it over a generic uniform prior, we can establish how important fitting a given parameter is for controlling the model's behaviour under a given physiological context.

Secondly, the shape and description of the posterior parameter distributions provide much more context and information than the prior point estimate. A key example is Figure 7.11 in Chapter 7 where a number of parameters are shown to have roughly uniform posterior distributions, but with a point estimate that is at one of the extreme ends of the prior range. Without being aware of the full distribution, it would be easy to say that the parameter had shifted far from its default value when in fact the model's behaviour in this situation is simply insensitive to changes to this parameter within the prior range.

Three variants of a new BrainPiglet model, BrainPigletHI, were then developed and the Bayesian framework used to determine the best. The three model variants were developed to better model the effects of hypothermia on the brain, an important consideration when the most common treatment for HIE is therapeutic hypothermia, and applied to data from two piglets who both suffered hypoxic brain injury followed by treatment via therapeutic hypothermia. The first piglet, LWP475, had suffered a mild severity injury and showed what would could be considered the expected response to cooling, with ΔHbD and ΔCCO decreasing over the period of cooling. The second piglet, LWP479, suffered a severe injury and showed the converse response, with cooling leading to an initial increase in ΔHbD and ΔCCO before then decreasing. This was done by incorporating the Q_{10} temperature co-efficient following work by Orlowski et al. [2014], with the first variant, BPH1, scaling both metabolic and haemodynamic reactions by the same co-efficient. The

second variant, BPH2, separated the effects of temperature out into two separate Q_{10} co-efficients for metabolism and haemodynamics respectively. The final variant, BPH3, had a Q_{10} co-efficient for metabolic reactions but scaled the rate of diffusion, the sole reaction within the haemodynamic compartment, by temperature directly with a proportionality constant set as a model parameter. Initial analysis of these three variants showed that all were able to maintain autoregulatory behaviour under typical physiological conditions.

Sensitivity analysis was then performed for each model and each piglet, with a reduced model then used for analysis. I found that variant 2 performed best, with the effect of temperature best separated out into a haemodynamic and a metabolic component. This was confirmed both visually and through the use of Bayes' Factors, with there being a 99.9825% that variant 2 is the correct model for piglet LWP475 and a 64.54% chance of it being the correct model for piglet LWP479. Finally the model was then used to predict the behaviour of the two piglets should they undergo spontaneous desaturation.

Here we see all of the benefits of the framework to analysing and considering cerebral haemodynamic and metabolism. Firstly, we are able to use the framework to develop and test different possible hypotheses for observed behaviour. We are able to do this with minimal effort when compared to the effort that would be required to develop an experimental paradigm capable of measuring the mechanisms in question. Secondly, we can then use this framework to analyse data and validate the models in question. In particular, the Bayesian approach has given us the ability to quantify the probability of a given model being the correct one from a set of variants. Finally, we are able to take a model and use it to quickly and easily perform different experiments and explore different hypotheses. In this case, it was used to perform in-silico experiments around spontaneous desaturation and consider the behaviour of a specific model under those conditions. In this case, we were able to identify that the effects of desturation during hypothermia on metabolism in a more severely injured piglet are different to those of a mildly injured piglet.

Chapter 9 applies all of the above to investigate data collected by Bale et al. [2018]

which found that a strong correlation between CCO and HbD during spontaneous desaturation events was an indicator of newborn brain injury. They postulated that it is "due to mitochondrial dysfunction and reduction in cerebral metabolic rate as a result of severe encephalopathy". The previously developed BrainPigletHI model was then applied to this in order to investigate their hypothesis around the mechanisms behind the observed correlation. Data from individual desaturation events for two neonates with differing outcomes were analysed using the BPH2 variant. Both neonates were being treated with hypothermia at the time of data collection. The analysis process was able to partially validate that the hypothesis by Bale et al. [2018] regarding a difference in $CMRO_2$ being responsible for the difference in the observed behaviour, with a difference in the normal $CMRO_2$ being observed when fitting HbT. Whilst there is room for improvement with regards to the accuracy of the model simulations, as has been mentioned previously, the usefulness of a model comes not from its accuracy alone, but from the ways in which it can be used to analyse and understand a system. In this respect, the framework was useful a on a number of fronts. It was possible to analyse and consider the differences in how the model parameters change when focussing on different physiological signals, providing greater insight into coupling between these signals and the parameters that control them. It was also possible to use the new framework to identify clear areas for improvement in the model analysis process. By comparing the posterior predictive distributions and the K-L divergence values to the initial sensitivity analysis results, we are able to better analyse and critically consider the existing model reduction step. Without the new framework and the outlook it provides this would not be possible. We also see in Chapter 9 another potential use for the modelling framework: the ability to identify likely measurement error and artefacts. In Figure 9.28 we see drift in the data, likely to be due to measurement error, which is easily remedied through detrending. Without comparison against a modelling simulation it is harder to identify that this is error rather than physiological behaviour.

I would like to take a short moment to draw attention to the accuracy of the measurements from both the animal models and the clinical data from neonates. Where

there is a lot of noise and/or artefacts within the data, there is a greater chance that the model reduction step involving sensitivity analysis will produce a set of parameters that fail to fully identify the behaviour that we are trying to model. This is most noticeable in the neonatal data where the selected parameters failed to reproduce the metabolic signal with any great degree of accuracy. As has been mentioned previously, one way around this would be to perform a more rigorous and thorough pre-processing stage, for example detrending data that appears to have a baseline shift. There is also the potential to include more data during the sensitivity analysis step to ensure that the parameters that are selected are, perhaps paradoxically, the most general and most specific reduced set possible. By this I mean that we obtain parameters that are specific to the observed behaviour across all datasets, but which are generalised across the subjects. This appears to have been the case with the piglet data which was able to produce a reduced set of data that was able to capture the larger trend of the behaviour. In addition it is worth considering that data collected in a lab environment from an animal model will contain far less environmental noise and external factors than data collected from a clinical setting where we are less able to control for other factors and the primary motive is the health of the patient. This also reinforces the benefit of a hierarchical model, as outlined below, where a shared parameter space could be investigated in greater detail using an animal model, as compared to those parameters that may only be captured from clinical data i.e. at a strata specific to humans who are critically ill and in need of treatment.

In summary what I have produced in this thesis is a software tool and a process that can be used in the following ways:

1. to test hypotheses around behaviour under differing physiological and environmental conditions

2. analyse data from neonates, including those undergoing therapeutic hypothermia at the time of data collection),

3. to identify potential measurement errors in data,

4. to run in-silico experiments to understand potential behaviours,

5. to critically assess established existing methods of model analysis.

Next I will consider future applications and improvements to the framework and models.

10.2 Future Work and Directions

10.2.1 Data Availability

Throughout this work, ranging from the machine learning work on artefact detection right through to the hypothermia models and investigations, a key barrier to further and more rigorous development is the availability of varied and reliable data. In all cases, a key requirement to develop modelling approaches is to have data from a variety of sources that has been confirmed to be both reliable and clean of artefacts or data points that could serve to confound or obfuscate the research being undertaken. In the simplest case of the artefacts detection, simply having more data that has confirmed and clearly labelled artefact spectra would make the process of training the algorithm much simpler.

For the more complex case of developing the BrainSignals models and any new modelling approaches, having a variety of data that covers a range of pathologies and subject types will allow model fitting to be more rigorous and the results to be more reliable. For example, developing a complete model with posterior parameter distributions that can be used to create a 'typical behaviour' model would require a range of data from a number of different subjects that share traits such as health, age and gender. From that point there could then be further work to establish a healthy piglet model, a healthy neonatal model and so on. All of these require a range of data for each 'grouping' to make sure that any typical parameter distributions are not in fact simply typical of that individual. These typical models would help inform clinical insights, allowing comparison between any distributions fitted to pathology and those from the healthy, typical model. This would also then help with further model development, identifying areas of weakness or areas that are likely to need more work to understand differences between groups.

10.2.2 Model Reduction

It is clear from the work so far that there is a need for improvement in the model reduction step. Correct identification of the best subset of parameters required to sufficiently model a specific behaviour is needed in order to produce reliable simulations and to draw good analyses. We've seen from the posterior predictive distributions that a number of parameters that were identified as important by the current sensitivity analysis method were then found to be minimally useful in actual model fitting. We see this in both the distributions themselves and the calculated K-L divergence values.

One possible solution to this would be the use of a Bayesian hierarchical model. As described by Gelman et al. [2013], many statistical applications involve parameters that are connected by the problem structure. In the example given, this is the effectiveness of cardiac treatment and the hospital in which treatment is received. For our use case, this could occur at multiple levels. Figure 10.1 outlines a simplified and possible structure for the BrainSignals models. The benefit of such a stratification is that fewer parameters would need to be fit for a given level. For example, once a model has been defined down to the Age level, there is only need to fit the parameters that are specific to the Injury Severity level. As Gelman et al. [2013] state, "hierarchical models can have enough parameters to fit the data well, while... avoiding problems of overfitting."

Should a verified hierarchical model be developed, the models then have the potential to be used to perform initial in-silico experiments to help direct in-vivo experimentation. This has the benefit of reducing costs, reducing the experimental feedback loop and reducing ethical concerns.

10.2.3 Environmental Stimuli

In this thesis I've outlined one extension of the model, the inclusion of temperature, in order to better model the environment in which data are collected. However, there are further examples of this that could also be considered as future directions for this work. There are a number of pharmacological interventions that can occur during treatment. The ability to simulate and model these interventions would provide two

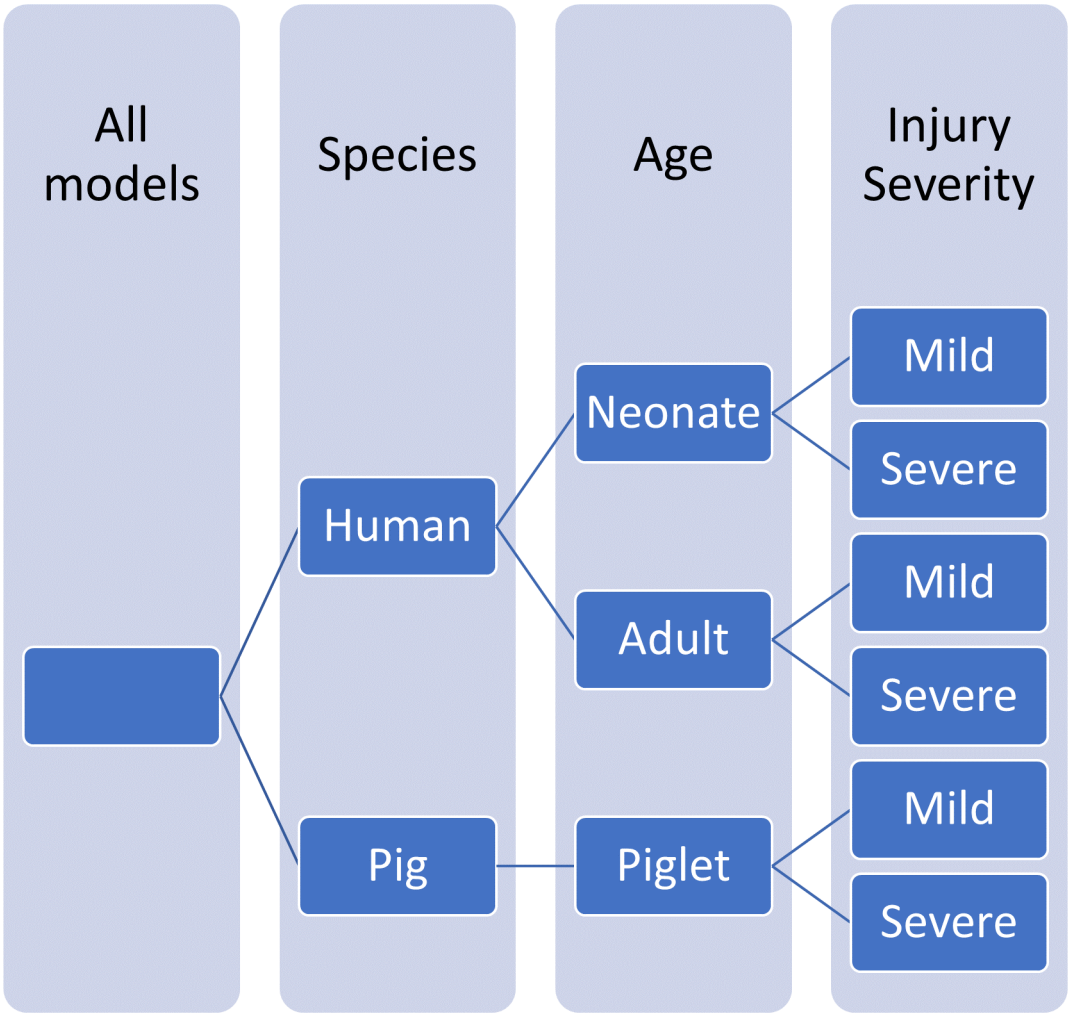

Figure 10.1: Example of a possible Bayesian hierarchical model for the BrainSignals models. Here we have stratified at the levels of Species, Age and Injury Severity.

obvious benefits. Firstly, data collected from subjects that have had a modelled treatment could be interpreted with less need to clean and remove these treatments from the data prior to analysis. Secondly, by introducing the ability to simulate external stimuli, there is the potential to perform in-silico experiments of these interventions under different conditions: hypothermia, severe injury, increased inhaled oxygen concentration etc. There are a number of forms this inclusion could possibly take. All stimuli could be modelled via a small number of parameters that vary in value to reflect the different treatment but which are shared regardless of intervention type, or different interventions could be modelled through individual

parameters, with the inclusion of the parameters being decided at the model reduction step. Regardless of the form it takes, the inclusion of intervention would widen the number of situations to which these models could be applied and improve their validity.

10.3 Impact

The impact of hypoxic injury on a newborn infant and their families cannot be underestimated. Whether it ends in death or severe neurodevelopmental disability, the consequences of an injury that cannot be treated are grave and serious and anything that can be done to improve outcome must be done. At present, therapeutic hypothermia, the main treatment, is only effective around half the time and so a better and more through understanding of the injury, the impact of treatment and the interplay between these two factors needs to be obtained. One of the best ways to do this is through clinical monitoring and so the ability to parse and process these data is paramount to any developments.

The modelling approach outlined in this thesis provides many possible pathways leading to this end goal that analysis of the data alone could not. Furthermore, it is only by developing a robust analysis process and toolset that these models prove their effectiveness and here we have provided both. Whether it's through obtaining a better understanding of how the different levels of the system interact, identifying measurement error in the data, a faster feedback loop for exploratory experimentation, or the verification of experimentally determined hypotheses, development of modelling approaches is a key component in our journey to understand the brain and the possible treatments we have for injury.

Appendix A

Table of Mathematical Models

Table of brain physiology models.

Paper	Title	Focus	Summary
Hudetz et al. [1987]	Pressure distribution in the pial arterial system of rats based on morphometric data and mathematical models.	Blood Flow	Computational model of rat pial arterial system. Vessels were visualised using a corrosion compound and classified in to branching orders. Steady-state pressure distribution in the system was calculated assuming flow at bifurcations was partitioned in proportion to diameter exponent of daughter branches. The models produced a very detailed description of the rat cerebral vasculature, with differences between normotensive and hypertensive rats.
Ursino [1988a]	A mathematical study of human intracranial hydrodynamics. Part 1 - The cerebrospinal fluid pulse pressure	Blood Flow	Compartmentalised model of intracranial vasculature used to simulate production of pressure pulse waves. Model parameters were given normal values calculated using physiology and anatomical data. Model is able to reproduce some typical clinical tests.
Ursino [1988b]	A mathematical study of human intracranial hydrodynamics. Part 2 - Simulation of clinical tests	Blood Flow	Testing of model defined in Ursino [1988a] by simulating results of tests based on a mock CSF injection. Simulation results correlate well with experimental data.

Table of brain physiology models.

Paper	Title	Focus	Summary
Ursino [1991]	A mathematical model of overall cerebral blood flow regulation in the rat.	Blood Flow	Extensive model of cerebral autoregulation in rats including feedback pathways such as the neurogenic, myogenic and chemical regulatory mechanisms. Simulations were performed of experiments with 1) neural fibers cut and artificially stimulated via external means; 2) the neurological mechanism in a physiologically closed loop condition. Comparison of the model with experimental data showed that the model produced 60-70% of the experimental regulatory capacity.
Korzeniewski and Froncisz [1991]	An extended dynamic model of oxidative phosphorylation	Metabolism	Model of concentration changes in time of chemical compounds and thermodynamic forces due to respiration in isolated cells. Model results were in good agreement with experimental data both qualitatively and quantitatively.
Ursino and Di Giammarco [1991]	A mathematical model of the relationship between cerebral blood volume and intracranial pressure changes: The generation of plateau waves.	Blood Flow	Model investigating the relationship between ICP, CBV and CBF based on an electrical analogue of the cerebral circulation. Particular emphasis on the mechanical properties of the proximal cerebral arteries and small pial arterioles. The model is able to reproduce experimental results on cerebral vessel dilation and CBF regulation, following a decrease in CPP.

Table of brain physiology models.

Paper	Title	Focus	Summary
Czosnyka et al. [1993]	CO2 cerebrovascular reactivity as a function of perfusion pressure: a modelling study	Blood Flow	Model of cerebrovascular resistance and compliance as a function of CPP and arterial CO2 pressure. Authors looked at identifying behaviour to differentiate between normal and impaired autoregulation, as well as interpreting interhemispheric asymmetry of CBF reactivity to changes in arterial CO2 in patients with carotid artery stenosis.
Ursino et al. [1995]	Intracranial pressure dynamics in patients with acute brain damage: A critical analysis with the aid of a mathematical model	Blood Flow	Same model as in Ursino [1988a,b]. Time pattern of ICP in response to typical clinical tests was studied in patients with acute brain damage. Simulation results demonstrate that in most patients the ICP time pattern cannot be explained on the basis of CSF dynamics alone, needing additional consideration of cerebral hemodynamics and blood volume changes.

Table of brain physiology models.

Paper	**Title**	**Focus**	**Summary**
Ursino and Lodi [1997]	A simple mathematical model of the interaction between intracranial pressure and cerebral hemodynamics	Blood Flow	Model based on simplifications from Ursino [1988a,b] combined with cerebral AR mechanisms simulated by means of a time constant and a sigmoidal static characteristic. It is used to simulate interactions between ICP, CBF and AR as well as specific phenomena. Simulation results show that the model was capable of reproducing clinical phenomena concerning ICP changes but was not suitable for studying ICP pulsations synchronous with cardiac beat or respiration.
Ursino et al. [1997]	Intracranial pressure dynamics in patients with acute brain damage	Blood Flow	Same model as in Ursino and Lodi [1997]. Used to analyse time pattern of ICP during a PVI test in 20 patients with acute brain injury. In most cases a satisfactory fitting of model response to patient data was achieved by adjusting only four parameters: CSF outflow resistance, intracranial elastance coefficient and the gain and time constant of autoregulation.

Table of brain physiology models.

Paper	Title	Focus	Summary
Czosnyka et al. [1997]	Contribution of mathematical modelling to the interpretation of bedside tests of cerebrovascular autoregulation	Blood Flow	Physiological, compartmentalised model of cerebral circulation. Non-linear differential equations were used to model time dependent interactions between pressure, flow and volume of CSF and cerebral blood. Patterns of blood flow were simulated during either short or long lasting decreases in CPP, which can be considered clinically equivalent to short compressions of the common caroitd artery, systemic hypotension and intracranial hypertension.
Ursino et al. [1998]	Relationships among cerebral perfusion pressure, autoregulation, and transcranial Doppler waveform: a modeling study.	Blood Flow	Same model as in Ursino [1988a,b]. Investigated presence of a relationship between CPP, AR and the transcranial doppler waveform. Concluded that under conditions of altered intracranial dynamics, information contained in the TCD waveforms depends on the superimposition of many complex phenomena.
Ursino and Lodi [1998]	Interaction among autoregulation, CO2 reactivity, and intracranial pressure: a mathematical model	Blood Flow	Extended previous work in Ursino and Lodi [1997], Ursino et al. [1998] to produce a combined physiological model of hydrodynamics a CO2 reactivity of the brain. Simulation results were encouraging and the model was used for further study in Lodi et al. [1998].

Table of brain physiology models.

Paper	Title	Focus	Summary
Lodi et al. [1998]	Modeling cerebral autoregulation and CO2 reactivity in patients with severe head injury	Blood Flow	Same model as in Ursino and Lodi [1998]. Model was used to simulate the time pattern of ICP and Vmca in response to maneuvers simultaneously affecting ABP and CO2. Sensitivity analysis suggested that the dynamic ICP-Vmca relationship is linear within the regulatory range and the slope is affected by the intracranial elastance coefficient. The model was then used to satisfactorily simulate clinical data with paramater values within ranges reported within the literature.
Hyder et al. [1998]	A model for regulation of cerebral oxygen delivery	Oxygen Transport	Model linking CMRO2 to CBF through an effective diffusivity for oxygen in the capillary bed. The model is able to fit observed proportionalities between CBF and CMRO2 for a large range of physiological data. This suggests that O2 diffusivity plays a large role in regulating cerebral O2 delivery in-vivo.

Table of brain physiology models.

Paper	**Title**	**Focus**	**Summary**
Gao et al. [1998]	Mathematical considerations for modeling cerebral blood flow autoregulation to systemic arterial pressure	Blood Flow	Compartmental model developed to synthesise the AR curve. The resistive arterial and arteriolar network was simplified as an autoregulation device consisting of 4 compartments in series controlling CBF, with exach compartment consisting for a group of identical vessels in parallel. The response of each vessel category to changes in CPP was simulated using experimental data. The predicted AR curve was consistent with experimental observations.

Table of brain physiology models.

Paper	Title	Focus	Summary
Panerai et al. [1999]	Linear and nonlinear analysis of human dynamic cerebral autoregulation	Blood Flow	Four linear models and one non-linear model were used to investigate the limitations of the assumption of a linear relationship between ABP and CBF velocity, comparing the performance of the different techniques in using ABP and CBF vecloity data collected from 47 normal subjects at rst in the supine position and during thigh cuff tests. They found that whilst it is possible to model dynamic cerebral AR in human using linear methods, a second-order nonlinear component significantly improves model accuracy from for a segment of data collected in the supine position and used to estimate model parameters. However, this could not be easily extended to represent the nonlinear component of the velocity response for other sgements of data or the transient changes in the thigh cuff test.
Cassot et al. [2000]	Hemodynamic role of the circle of WIllis in stenoses of internal carotid arteries. An analytical solution of a linear model	Blood Flow	Model of blood flow through the circle of Willis developed within a linear framework. Autoregulation was not included. Model simulations with data from patients with carotid artery diseases produced no physiologically signfiicant differences as compared to non-linear models.

Table of brain physiology models.

Paper	Title	Focus	Summary
Ursino et al. [2000]	Cerebral hemodynamics during arterial and CO2 pressure changes: in vivo prediction by a mathematical model	Blood Flow	Simplified version of model seen in Ursino and Lodi [1998], Lodi et al. [1998]. Main simplification is that the model does not distinguish between larg and small pial arteries, and that the biomechanics of the arterial-arteriolar vasculature are reproduced by a windkessel model. Model was validated by comparing imulations with Vmca and ICP data measured in 13 patients. The model was able to reproduce time patterns seen in the monitoring systems by using parameter values within the ranges reported in the literature.
Aubert et al. [2001]	Modelling of the coupling between brain electrical activity and metabolism:	Blood flow; Oxygen Transport; Metabolism	Model attempting to group various aspect of brain functional imaging within a coherent frame implemented via a system of differential equations. It includes 1) sodium membrane transport, 2) Na/K ATPase, 3) neuronal energy metabolism, 4) blood-brain barrier exchanges and 5) brain haemodynamics. They show how the model enables interpretation of published MRS and fMRI data obtained during prolonged neuronal stimulation.

Table of brain physiology models.

Paper	**Title**	**Focus**	**Summary**
Golding and Golding [2001]	Mathematical modelling of responses of cerebral blood vessels to changing intraluminal pressure	Blood Flow	Model is based on the way the total tension within the blood vessel walls varies due to specific ions interacting and affecting the vascular smooth muscle cells and the vascular walls. It uses three parameters: diameter at zero pressure, myogenic response as pressure tends to zero and a term associated with myogenic tone. Model was validated and then used to analyse experimental data on the resistance o cerebral vessels that were isolated from rats.
Kirkham et al. [2001]	A new mathematical model of dynamic cerebral autoregulation based on a flow dependent feedback mechanism	Blood Flow	Model represents dynamic cerebral AR as a flow dependent feedback mechanism. Model is based on a linear transfer function and has only two parameters: rate of restoration and time delay. Model was tested against physiological data from a thigh cuff test and lower body negative pressure. Parameter estimates are consistent in both suggesting they are independent of the pressure change stimulus and depend only on the dynamic cerebral autoregulation process itself.

Table of brain physiology models.

Paper	Title	Focus	Summary
Piechnik et al. [2001]	A model of the cerebral and cerebrospinal fluid circulations to examine asymmetry in cerebrovascular reactivity	Blood Flow	Detailed anatomical model of both brain hemispheres. Arterial part of the model contains both myogenic and metabolic mechanisms responsible for AR. Model is used to investigate the "steal phenomenon" [Symon, 1968] where only one hemisphere is reactive. Findings suggested that vasoreactive effects on ICP had major influence on circulation in both hemispheres and the "steal effect" is unlikely to occur in subjects with nonstenosed carotid arteries.
Olufsen et al. [2002]	Dynamics of cerebral blood flow regulation explained using a lumped parameter model	Blood Flow	Windkessel model with two resistors and a capacitor is used to reproduce beat-to-beat changes in Vmca in reponse to changes in ABP. Despite the model's simplicity is is able to capture dynamic changes due to a shift in posture.
Liu and Allen [2002]	Analysis of dynamic cerebral autoregulation using an ARX model based on arterial blood pressure and middle cerebral artery velocity simulation	Blood Flow	Model of cerebrovascular system using a linear ARX model based on data from model in Ursino et al. [1995] to look at applicability of a linear ARX model to assess cerebral AR under different conditions. Concluded that AR can be assessed using the step response of an appropriate ARX mdoel, even when data is noisy.

Table of brain physiology models.

Paper	Title	Focus	Summary
Gjedde [2002]	Cerebral blood flow change in arterial hypoxemia is consistent with negligible oxygen tension in brain mitochondria	Oxygen Transport	Compartmental model of oxygen delivery to brain tissue used to investigate regulation of blood flow during neuronal activation, namely that minimal oxygen reserves within mitochondria requires blood flow to rise more than oxygen consumption to compensate.
Aubert and Costalat [2002]	A Model of the Coupling between Brain Electrical Activity, Metabolism, and Hemodynamics: Application to the Interpretation of Functional Neuroimaging	Blood flow; Oxygen Transport; Metabolism	Based on Aubert et al. [2001]. Model of cell membrane ionic currents, energy metabolism, blood-brain barrier exchanges and haemodynamics. Model is used to investigate various hypotheses around CMRO2. Model was able to provide good agreement with experimental data from MRI and MRS when under two of the four investigated hypotheses.
Panerai [2003]	The critical closing pressure of the cerebral circulation	Blood Flow	Hydrodynamical model incorporating arterial elasticity and active wall tension. It is used to look at the ABP at which CBF approaches zero. Authors used their simple linear model with models from Ursino and Di Giammarco [1991], Ursino et al. [1995]. Concluded that a more complex anatomical model would give more information about this critical closing pressure.

Table of brain physiology models.

Paper	Title	Focus	Summary
Yang et al. [2003b]	The myogenic response in isolated rat cerebrovascular arteries: vessel model	Blood Flow	Physiological mdoel of the rat cerebrovascular vessel simulating the myogenic response and include elements from Yang et al. [2003a]. It smulates measured data well and provide a link between events at teh subcellular level and macroscopic changes in vessel diameter.
Yang et al. [2003a]	The myogenic response in isolated rat cerebrovascular arteries: smooth muscle cell model	Blood Flow	Physiological model of the cerebrovascular smooth muscle cell incorporating electrochemical and chemomechanical submodels. The model is able to simulate known smooth muscle mechanics and could also predict events that are difficult to measure.
Ursino and Giulioni [2003]	Quantitative assessment of cerebral autoregulation from transcranial Doppler pulsatility	Blood Flow	Same model as Ursino [1988b], Ursino and Lodi [1998]. Looked at relationship between Vcma pulsatility and cerebral AR and found various conclusions. One notable finding is the linear relationship between Vmca and CPP for a wide range of CPP values and that the slope of this relationship maybe provide information about AR status.

Table of brain physiology models.

Paper	**Title**	**Focus**	**Summary**
Lu et al. [2003]	Cerebral autoregulation and gas exchange studied using a human cardiopulmonary model	Blood Flow; Oxygen Transport	Combined model of the cardiopulmonarysystem, central nervous system and cerebral hemo- and hydrodynamics. It is used to simulate various physiological conditions and experiments and could closely mimix experimental findings. It could be used to investigate mechanisms behind the dynamics and interactions simulated.
Aubert and Costalat [2005]	Interaction between astrocytes and neurons studied using a mathematical model of compartmentalized energy metabolism	Oxygen Transport; Metabolism	Compartmentalised model of energy metabolism in neurons and astrocytes based on Aubert et al. [2001]. Used to investigate the Anstrocyte-Neuron lactate shuttle hypothesis of astrocyte-provided lactate being an energy fuel for neurons.

Table of brain physiology models.

Paper	Title	Focus	Summary
Jung et al. [2005]	A mathematical model of cerebral circulation and oxygen supply	Blood Flow	Compartmental model of cerebral circulation and oxygen supply, including autoregulation. Circulatory compartments has 7 sub-compartments of arteries, capillaries, veins, brain tissue, CSF, sagittal sinus and an artificial compartment to simulate brain swelling. Oxygen supply is based on the Krogh model. Model simulations reprorduce experimentally verified behaviour, as well as identifying linear relationship between partial oxygen pressure and MABP where impaired autoregulation is present.
Vatov et al. [2006]	Modeling brain energy metabolism and function: a multiparametric monitoring approach	Metabolism	Model to simulate changes in brain energy metabolism under various pathophysiological conditions. Model paramaters include CBF, extra-cellular potassium, partial oxygen pressure and mitochondrial NADH redox state. Simulations show the model's ability to mimic pathological conditions and higjlight the potential for a properuly tuned model allows monitoring of only one or two variables with others computed in real time.

Table of brain physiology models.

Paper	Title	Focus	Summary
Simpson et al. [2007]	Supply and demand in cerebral energy metabolism: the role of nutrient transporters	Metabolism	Model of glucose and lactate levels in the brain in response to neuronal stimulation, including a description of transporter proteins. Conlude that neurons are responsible for glucose uptake and metabolism as well as the generation of lactate transients on activation.
Banaji et al. [2008]	A model of brain circulation and metabolism: NIRS signal changes during physiological challenges	Blood Flow; Oxygen Transport; Metabolism	
Alastruey et al. [2008]	Reduced modelling of blood flow in the cerebral circulation: Coupling 1-D, 0-D and cerebral auto-regulation models	Blood Flow	Model is used to simulate pulsatile blood flow within cerebral circulation. Model consists of a 1-D model of arterial haemodynamics combined with lumped parameter and flow autoregulation models of perfusion. SImulations look in particular at the effects of anatomical variations of the circle of Willis on blood flow following sudden carotid artery occlusion.
Cloutier et al. [2009]	An integrative dynamic model of brain energy metabolism using in vivo neurochemical measurements	Metabolism	Systems approach to model brain energy metabolism, based on Aubert and Costalat [2005] but with inclusion of glutamate cycling. Model calibrated and verified from in-vivo data collected from rats, with good agreement between model and data.

Table of brain physiology models.

Paper	Title	Focus	Summary
Linninger et al. [2009]	A mathematical model of blood, cerebrospinal fluid and brain dynamics.	Blood Flow	Based on first principles of fluid and solid mechanics to build comprehensive, compartmental model of human intracranial dynamics. Includes blood, CSF, brain parenchyma and the spinal canal. Model provides transition from healthy to diseased states during the onset of communicating hydrocephalus. Model is verified and in good agreement with clinical measurements.
DiNuzzo et al. [2010]	Changes in glucose uptake rather than lactate shuttle take center stage in subserving neuroenergetics: evidence from mathematical modeling	Metabolism	Combination of Aubert and Costalat [2002, 2005], Simpson et al. [2007]. Model is used to look at metabolite trafficking within the brain during functional activation. They conclude that energy substrate usage is mainly controlled via glucose uptake as opposed to the lactate shuttle.
Orlowski et al. [2011]	Modelling of pH dynamics in brain cells after stroke	Metabolism; Oxygen Transport	Model of brain cellular metabolism extended from Cloutier et al. [2009]. It is composed of 4 compartments: astrocytes, neurons, capillary vessel and extracellular volume and incorporates pH dynamics. Model is used to investigate pH dynamics following ischaemic stroke and model simulations of intracellularpH are in good agreement with previously reported measurements.

Table of brain physiology models.

Paper	Title	Focus	Summary
Calvetti and Somersalo [2011]	Dynamic activation model for a glutamatergic neurovascular unit	Oxygen Transport; Metabolism	A dynamic spatially lumped model for brain energy metabolism is developed before using the results of an MCMC based flux balance analysis to estimate the kinetic model parameters. Model consists of 5 compartments and is govern by kinetic mass balance equations, with Michaelis-Menten type expressions for reaction rates and transport between compartments.
Moroz et al. [2012a]	Computational modelling of the piglet brain to simulate near-infrared spectroscopy and magnetic resonance spectroscopy data collected during oxygen deprivation	Blood Flow; Oxygen Transport; Metabolism	
Moroz et al. [2012b]	Development of a model to aid NIRS data interpretation: Results from a hypercapnia study in healthy adults	Blood Flow; Oxygen Transport; Metabolism	

Table of brain physiology models.

Paper	Title	Focus	Summary
Somersalo et al. [2012]	The Metabolism of Neurons and Astrocytes Through Mathematical Models	Metabolism	A review of metabolic models and how they diverge in their findings. Authors then apply a stochastic approach towards obtaining parameter distributions as applied to the model outlined in Calvetti and Somersalo [2011]. This approach reveals how much intrinsic variability a metabolic system may contain and how this may be a valid explanation of variability in model findings.
Hapuarachchi et al. [2013]	Modelling Blood Flow and Metabolism in the Piglet Brain During Hypoxia-Ischaemia: Simulating pH Changes	Blood Flow; Oxygen Transport; Metabolism	
Orlowski et al. [2013]	Modelling of the physiological response of the brain to ischaemic stroke	Metabolism	An extension of Orlowski et al. [2011] to include metabolite concentrations within the extracellular volume. Model is used to investigate the response of the brain to ischaemic stroke, looking at the evolution of the stroke within an individual.
Hapuarachchi et al. [2014]	Simulating NIRS and MRS Measurements During Cerebral Hypoxia-Ischaemia in Piglets Using a Computational Model	Blood Flow; Oxygen Transport; Metabolism	

Table of brain physiology models.

Paper	Title	Focus	Summary
Caldwell et al. [2015a]	Brainsignals revisited: simplifying a computational model of cerebral physiology.	Blood Flow; Oxygen Transport; Metabolism	
Jolivet et al. [2015]	Multi-timescale Modeling of Activity-Dependent Metabolic Coupling in the Neuron-Glia-Vasculature Ensemble	Metabolism	Biophysical model of the brain's metabolic interactions, integrating three modeling approaches. Their findings support the astrocyte-neuron lactate shuttle hypothesis and provide a quantitative mathematical description of the metabolic activation in neurons and glial cells, as well as of the macroscopic measurements obtained during brain imaging.
Caldwell et al. [2015b]	Modelling blood flow and metabolism in the preclinical neonatal brain during and following hypoxic-ischaemia.	Blood Flow; Oxygen Transport; Metabolism	

Table of brain physiology models.

Paper	Title	Focus	Summary
Calvetti et al. [2015]	A spatially distributed computational model of brain cellular metabolism	Oxygen Transport; Metabolism	3-D spatially distributed model of brain metabolism. Investigates how locaus of synaptic activity relative to capillaries and diffusion affects model behaviour. Uses spatial averaging to compare this against a lumped model. Authors find that varying the locus of the synaptic activity relative to capillaries significantly alters the lumped model parameters, indicating that model uncertainty is an inherent feature of lumped models.
Blanchard et al. [2016]	A New Computational Model for Neuro-Glio-Vascular Coupling: Astrocyte Activation Can Explain Cerebral Blood Flow Nonlinear Response to Interictal Events	Blood Flow; Metabolism	Development of a model linking neuronal activity to regional CBF dynamics through neuro-glio-vascular coupling. Model is used to interpret regional multimodal recordings and reproduce experimental data. Results indicate that the relationship between neuronal activity and CBF magnitudes can be nonlinear for isolated events and that this nonlinearity is due to astrocytic activity.
Hapuarachchi et al. [2016]	Simulation of Preterm Neonatal Brain Metabolism During Functional Neuronal Activation Using a Computational Model	Blood Flow; Oxygen Transport; Metabolism	

Table of brain physiology models.

Paper	Title	Focus	Summary
Mathias et al. [2017a]	A model of neurovascular coupling and the BOLD response: PART I	Blood Flow; Oxygen Transport; Metabolism	Developed a model describing a signalling mechanism of neurovascular coupling (NVC) with a model of a pyramidal neuron and its corresponding fMRI BOLD response. The first of two papers, it describes integrating the extended neurovascular coupling model with a model providing a BOLD signal taking its input from CBF and CMRO2.
Mathias et al. [2017b]	A model of neurovascular coupling and the BOLD response: PART II	Blood Flow; Oxygen Transport; Metabolism	Second of two papers describing a model of a pyramidal neuron and its corresponding BOLD fMRI response. In this paper the model is used to predict variations of the BOLD response from variations in neuronal activity. Results indicate that the BOLD signal could be used as an initial biomarker for neuronal dysfunction or variations in the perfusion of blood to cerebral tissue. Model simulation are in good agreement with experimental data.
Mathias et al. [2018]	Integrated models of neurovascular coupling and BOLD signals: Responses for varying neural activations	Blood Flow; Oxygen Transport; Metabolism	The model described in Mathias et al. [2017a,b] has been extended to include a transient sodium ion channel expressed in the neuron and integrated into a complex NVC model. This model is able to simulate BOLD responses associated with continuous spiking, bursting and cortical spreading depression (CSD).

Table of brain physiology models.

Paper	**Title**	**Focus**	**Summary**
Calvetti et al. [2018]	A computational model integrating brain electrophysiology and metabolism highlights the key role of extracellular potassium and oxygen	Blood Flow; Oxygen Transport; Metabolism	Development of a new multi-compartment (4 compartments chosen for a balance between resolution and complexity) model based on the “Reduced Kinetic Model” outlined in Somersalo et al. [2012]. Simulations show that neuron and astrocyte are asymmetric in sharing the cost of activation: numerical simulations show a longer depletion of the phosphorylation state in the astrocyte due to the increased energy cost of extra-cellular potassium cleaning.

Appendix B

Details of WeBCMD Operation

To build the Docker container, Docker and Docker Compose need to be installed. Once this is done, the WeBCMD Docker files can be downloaded in zipped form or cloned from the GitHub repository [Github]. The latter allows for updates to be easily pushed and installed using Git. After downloading, navigate to the directory using a terminal emulator such as CMD or Powershell on Windows, Terminal on Mac or xterm on Linux. Build the Docker container using the command:

```
docker-compose up
```

which will also load all environment variables.

There are two Docker containers used in the local distribution: the WeBCMD container, which is the same as that used for the online version, and a MongoDB container. The latter is used to store information about models, such as parameter values and default inputs and outputs. It is expected that over time other features will be added that utilise this database, such as detailed descriptions of the physiological significance of each model variable.

Other Uses There are some features which have not been fully implemented within the graphical interface. partly due to uncertainty over their continued use, such as sensitivity analysis and OpenOpt model optimisation. It is possible to still access these features inside the Docker container by using the command line.

To build and launch the container to do this a slightly different process is required.

1. **Build** The container needs to be built before running, and the following command needs to be run from inside the directory with the Dockerfile in and ensures that the most recent version of each intermediate container is used:

   ```
   docker build --no-cache --rm -t webcmd:latest .
   ```

2. **Running** When using the command line interface, if you want access and store data on the host computer, you will need to use the 'working' directory with the following command. Any data you wish to use with the models must be stored in the `host_data` directory. These can be individual files or child directories:

   ```
   docker run -it -v
    ↪  /home/user/path/to/host_data:/BCMD/working
     --entrypoint /bin/bash webcmd
   ```

The WeBCMD software is available via two main methods: accessing the software via "the cloud", by visiting the WeBCMD website or by downloading the software and running it from inside the Docker container. It is also possible, though not recommended, to run the interface directly using Python without using Docker. This is most likely to work if the operating system used is a Unix variant, as this is the operating system used within the Docker container.

If accessing the online version of the software, all models and their associated information have already been compiled and stored. The user needs only to follow the on screen steps.

If using the local distribution, the installation steps outlined above will need to be followed. Once this has been done and the containers built and launched, models must be compiled and their information uploaded to the local database. This can be done by accessing the Admin panel in the navigation menu. The local distribution has its own admin user and password which give access to this section. They are

USER: LOCAL_USER

PASSWORD: LOCAL_PASSWORD

Inside the admin section of the interface there are the options to compile a model and to upload model information. Models will need compiling before information can be uploaded, as the model information is generated by running the model itself. Once this has been done, returning to the 'Home' tab allows users to review the models that have been compiled and to run them.

Appendix C

Supporting Information for Chapter 7

S1 Sensitivity Analysis σ values - simulated data. Bar charts of the σ values from the sensitivity analysis of the simulated data.

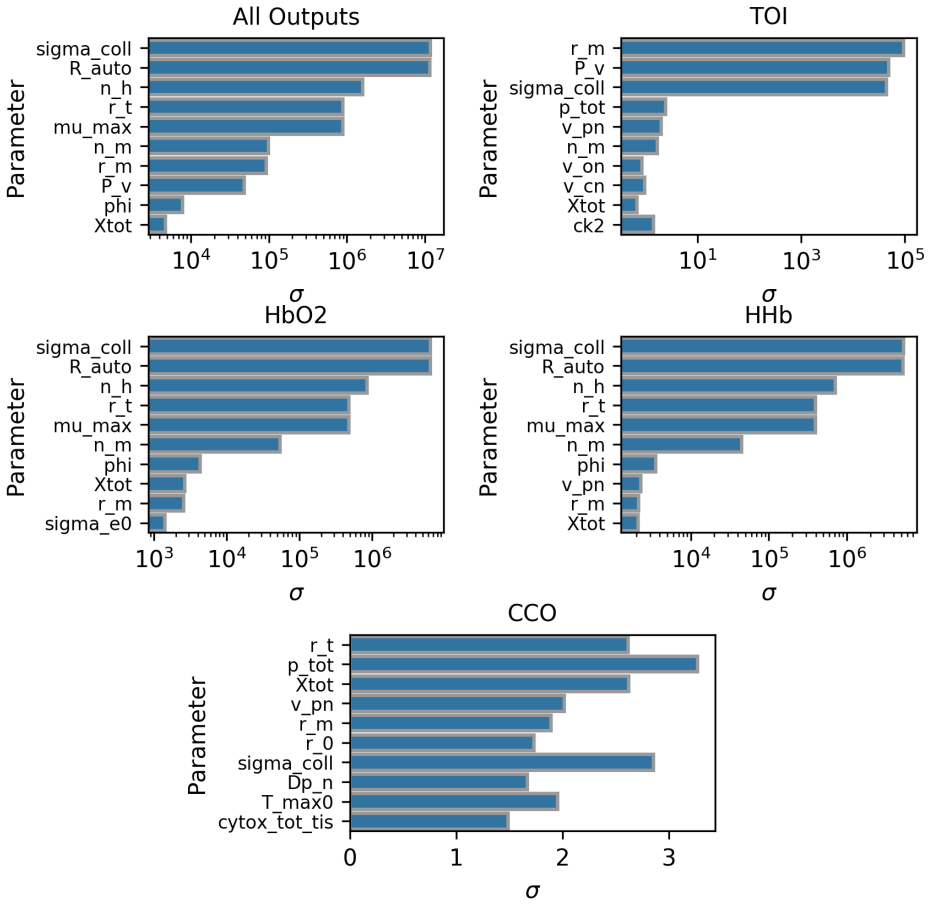

Figure C.1: σ values for each signal as per the sensitivity analysis. Shown are the σ values for each of the 10 most sensitive parameters in each signal, as per μ_*, for the simulated datasets.

S2 Sensitivity Analysis σ values - simulated data. Bar charts of the σ values from the sensitivity analysis of the experimental data.

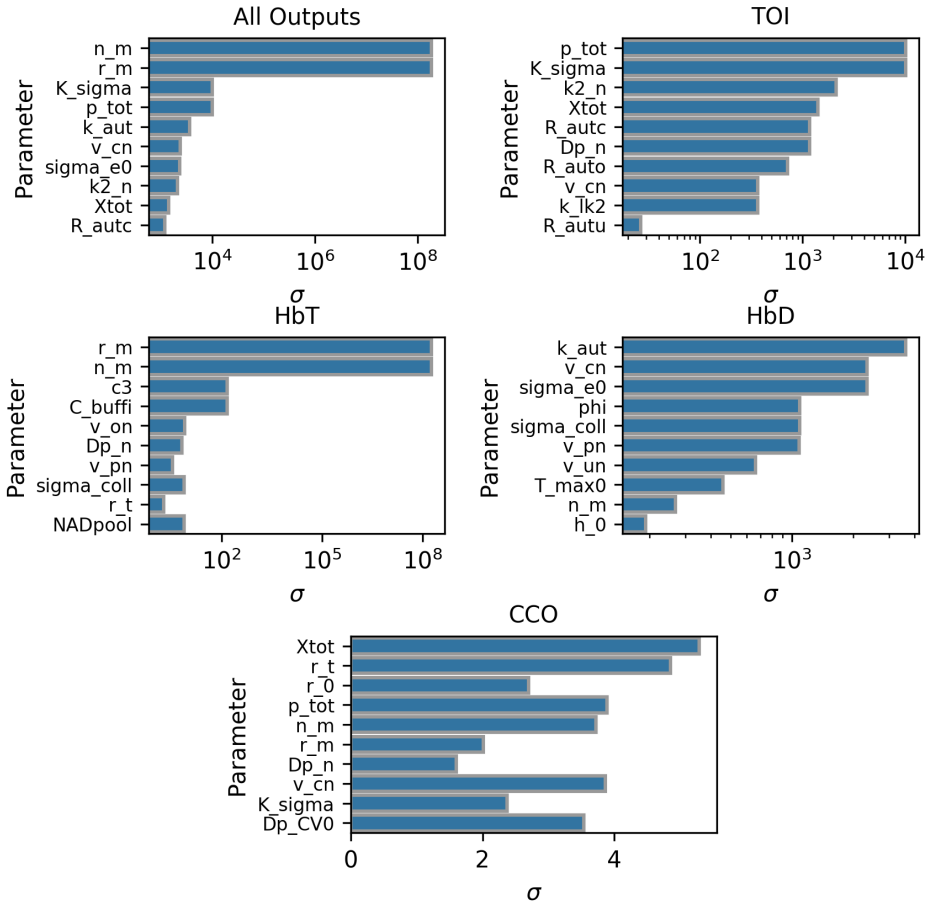

Figure C.2: σ values for each signal as per the sensitivity analysis. Shown are the σ values for each of the 10 most sensitive parameters in each signal, as per μ_*, for the experimental dataset.

S3 Table. Table of posterior and prior distribution information for healthy simulated data.

	Posterior				Prior	
Parameter	Median	Lower Quartile	Upper Quartile	IQR	Prior Min.	Prior Max.
sigma_coll	65.04	49.4	79.87	30.47	31.395	94.185
R_auto	1.353	1.053	1.721	0.6679	0.75	2.25
n_h	2.541	1.988	3.079	1.09	1.25	3.75
r_t	0.01822	0.01692	0.01938	0.002461	0.009	0.027
mu_max	1.014	0.7679	1.26	0.4922	0.5	1.5
n_m	1.992	1.493	2.369	0.876	0.915	2.745
r_m	0.02879	0.02568	0.03273	0.007059	0.0135	0.0405
P_v	4.112	3.095	5.067	1.972	2	6
phi	0.03609	0.03474	0.03775	0.003013	0.018	0.054
Xtot	8.898	8.463	9.406	0.943	4.55	13.65

Table C.1: Posterior and prior distribution information for healthy simulated data. Posterior distribution values are given to 4 significant figures. Prior range values are given as their exact values.

S4 Table. Table of posterior and prior distribution information for impaired simulated data.

	Posterior				Prior	
Parameter	Median	Lower Quartile	Upper Quartile	IQR	Prior Min.	Prior Max.
sigma_coll	64.58	48.86	79.79	30.93	31.395	94.185
R_auto	1.488	1.183	1.831	0.6479	0.75	2.25
n_h	2.337	1.742	3.014	1.272	1.25	3.75
r_t	0.0136	0.01183	0.01508	0.00325	0.009	0.027
mu_max	1.076	0.8283	1.289	0.4609	0.5	1.5
n_m	1.928	1.471	2.344	0.873	0.915	2.745
r_m	0.02738	0.02372	0.03174	0.00802	0.0135	0.0405
P_v	4.022	2.966	4.964	1.998	2	6
phi	0.03722	0.03349	0.04122	0.007732	0.018	0.054
Xtot	8.764	7.534	10.18	2.646	4.55	13.65

Table C.2: Posterior and prior distribution information for impaired simulated data. Posterior distribution values are given to 4 significant figures. Prior range values are given as their exact values.

S5 Table. Table of posterior and prior distribution information for experimental data.

	Posterior				Prior	
Parameter	Median	Lower Quartile	Upper Quartile	IQR	Prior Min.	Prior Max.
n_m	2.078	1.645	2.446	0.8011	0.915	2.745
r_m	0.02574	0.02342	0.02841	0.004988	0.0135	0.0405
K_sigma	8.195	6.701	9.663	2.961	5	15
p_tot	18.48	17.45	19.4	1.943	10	30
k_aut	0.9799	0.9073	1.074	0.167	0.5	1.5
v_cn	47.88	43.93	52.13	8.199	20	60
sigma_e0	0.1386	0.1043	0.1739	0.06966	0.07125	0.2137
k2_n	3870	2903	4822	1919	1957.83	5873.52
Xtot	7.067	6.589	7.555	0.9663	4.55	13.65
R_autc	1.955	1.516	2.517	1.001	1.1	3.3

Table C.3: Posterior and prior distribution information for experimental data. Posterior distribution values are given to 4 significant figures. Prior range values are given as their exact values.

S6 Statistical Analysis - Simulated Healthy Data. Figures of statistical checks used to verify the method and results for the simulated healthy data.

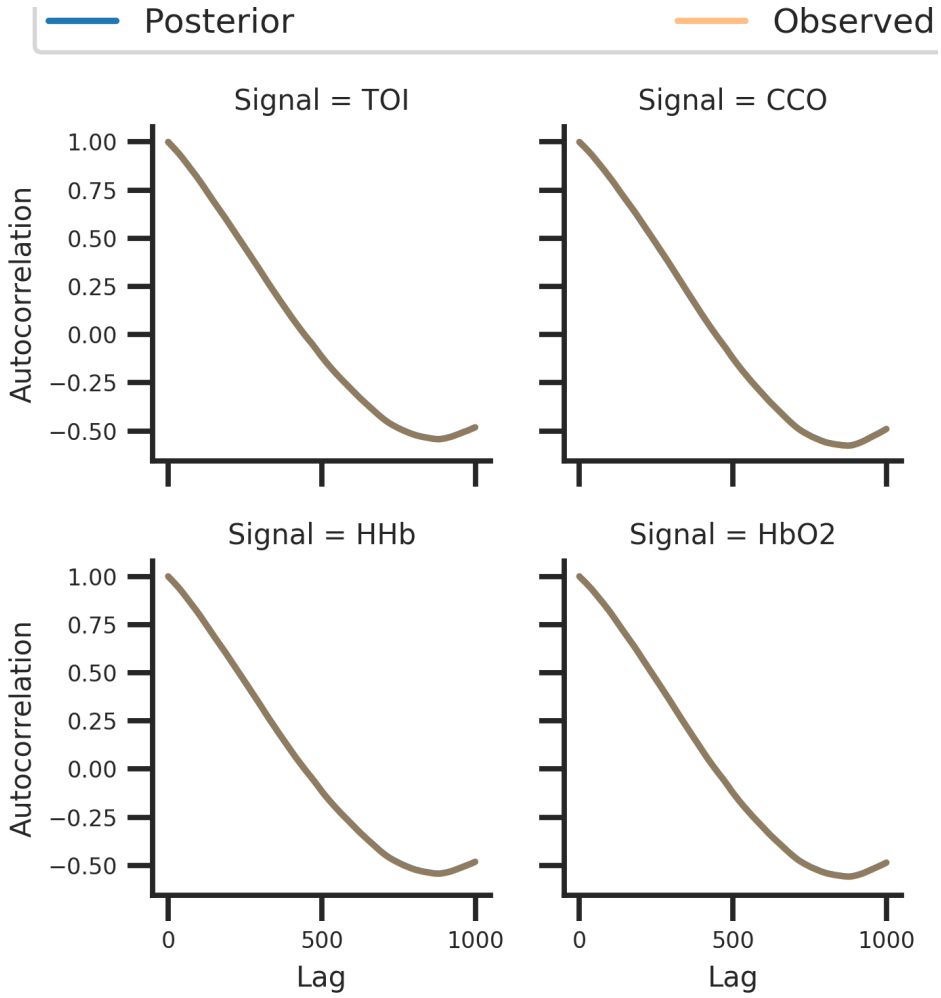

Figure C.3: Autocorrelation of posterior predictive and the observed data. The graph shows a comparison of the autocorrelation value between observed and posterior predictive time series as a function of lag. We see significant overlap between the two suggesting an extremely good fit.

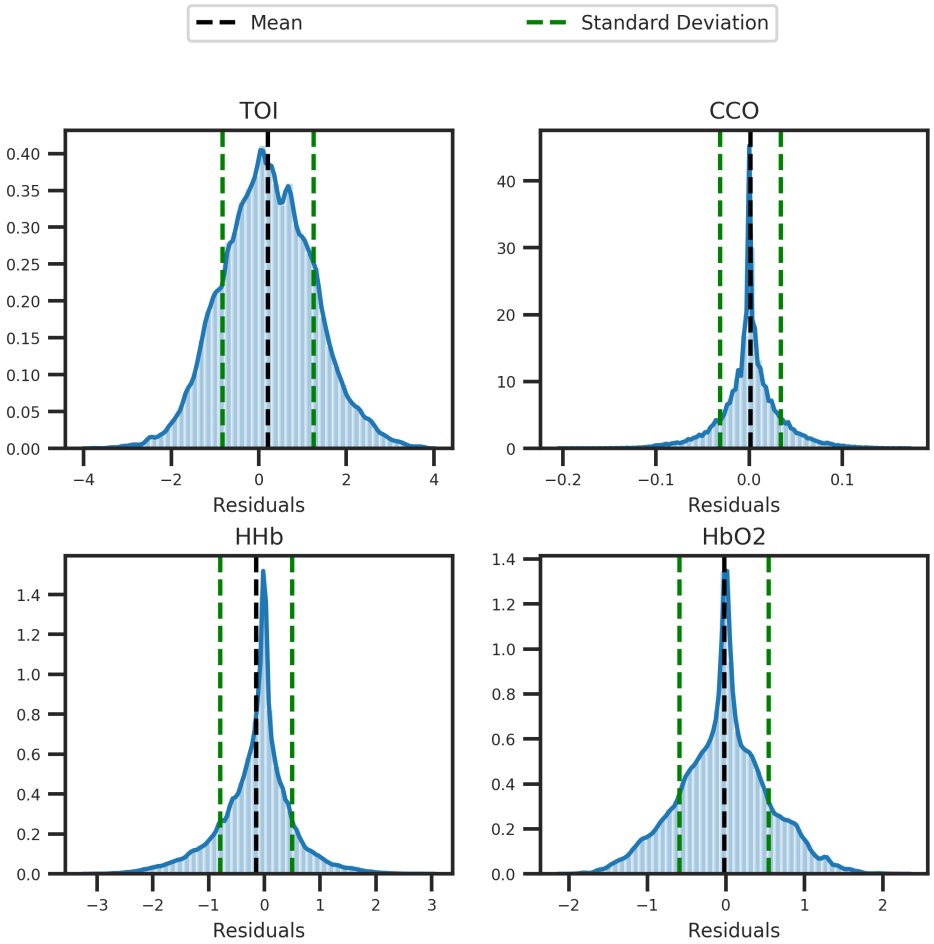

Figure C.4: Distributions of residuals. Distributions of the residuals for each signal. Mean and standard deviation of each plot are indicated by black and green lines respectively.

S7 Statistical Analysis - Simulated Impaired Data. Figures of statistical checks used to verify the method and results for the simulated impaired data.

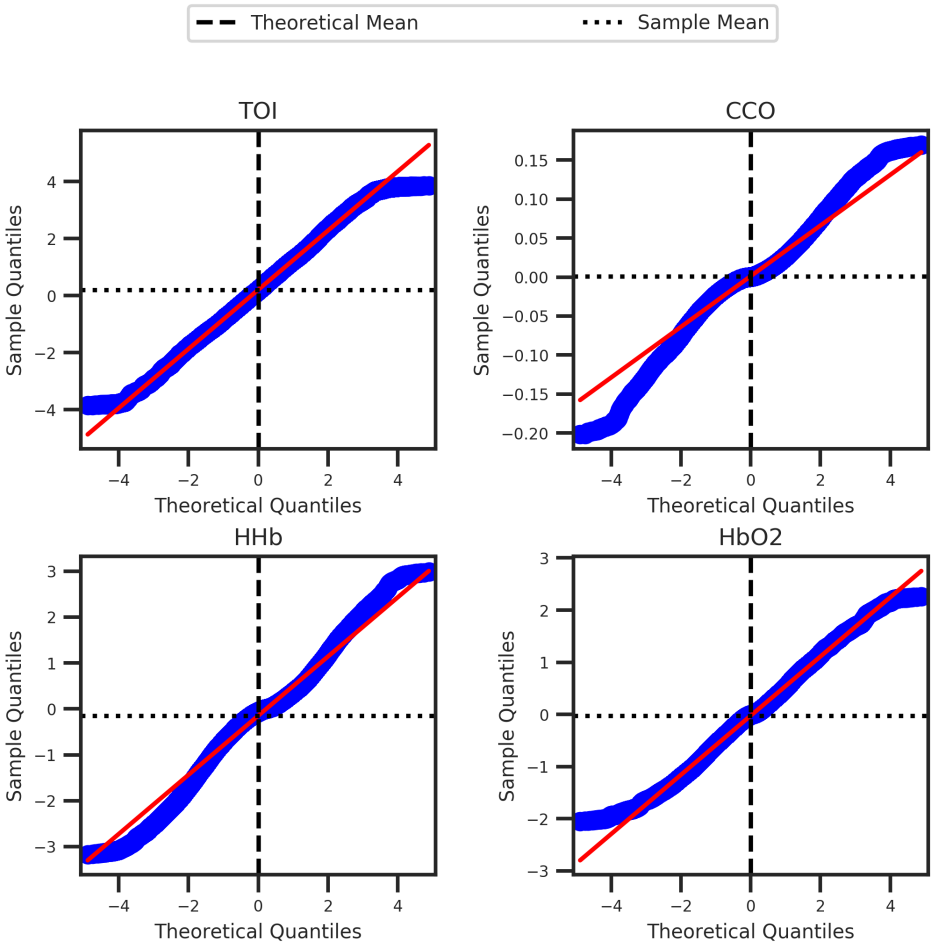

Figure C.5: Q-Q plot of residuals. Q-Q plots looking at the residuals for each signal. The plots use a standardised diagonal line, where the expected order statistics are scaled by the standard deviation of the sample residuals and have the mean added to them.

S8 Statistical Analysis - Experimental Data. Figures of statistical checks used to verify the method and results for the experimental data.

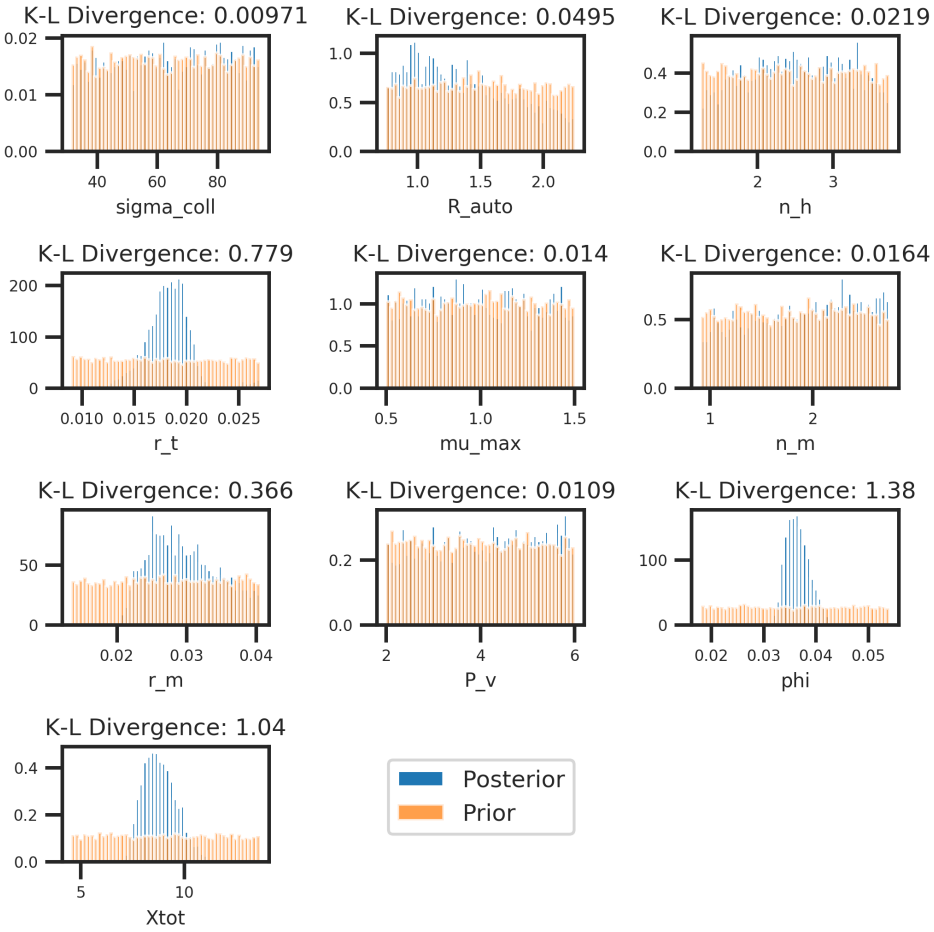

Figure C.6: Comparison of marginal posterior and prior distributions for each parameter. A clearer comparison is made here between the marginal prior and posterior distributions for each fitted parameter. The Kullback–Leibler divergence is shown for each.

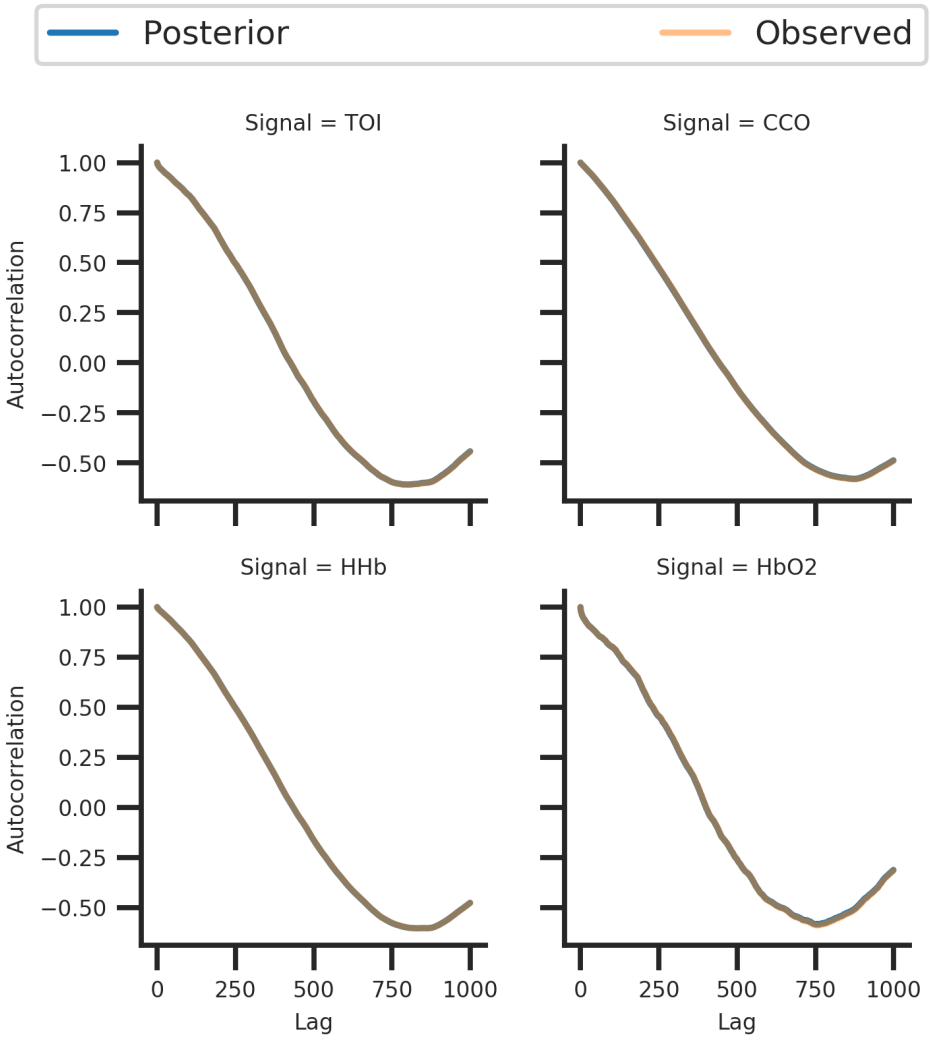

Figure C.7: Autocorrelation of posterior predictive and the observed data. The graph shows a comparison of the autocorrelation value between observed and posterior predictive time series as a function of lag. We see significant overlap between the two suggesting an extremely good fit.

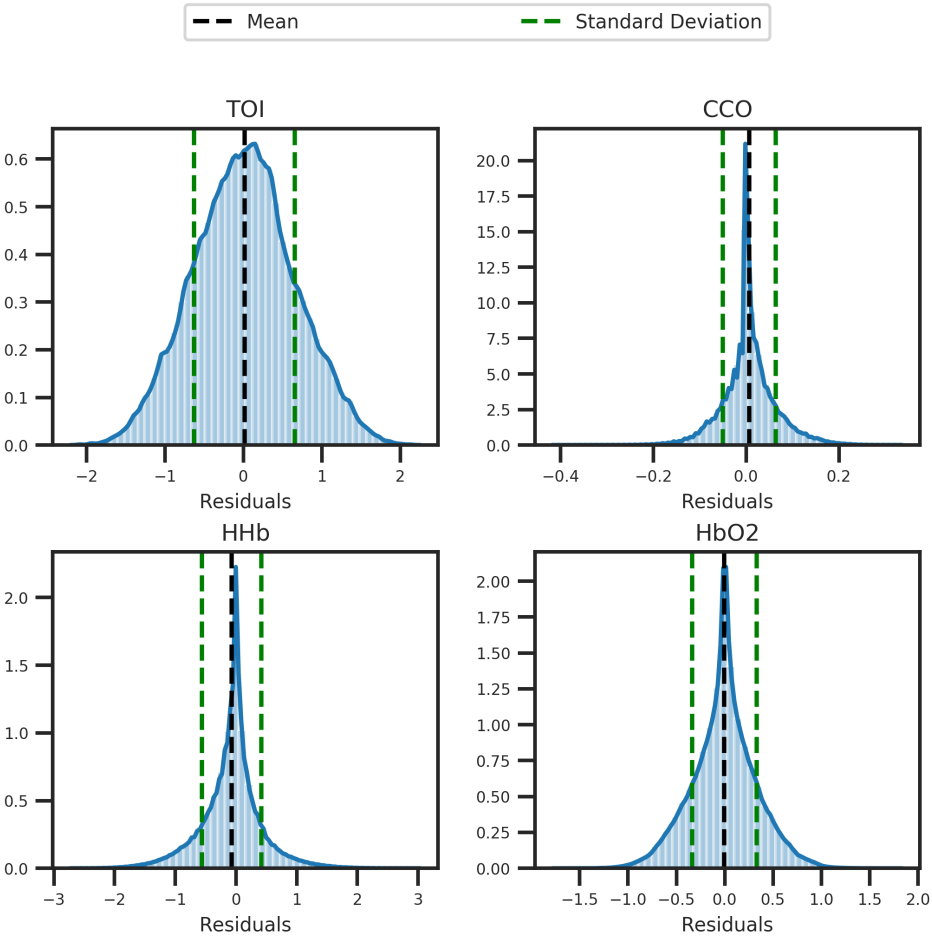

Figure C.8: Distributions of residuals. Distributions of the residuals for each signal. Mean and standard deviation of each plot are indicated by black and green lines respectively.

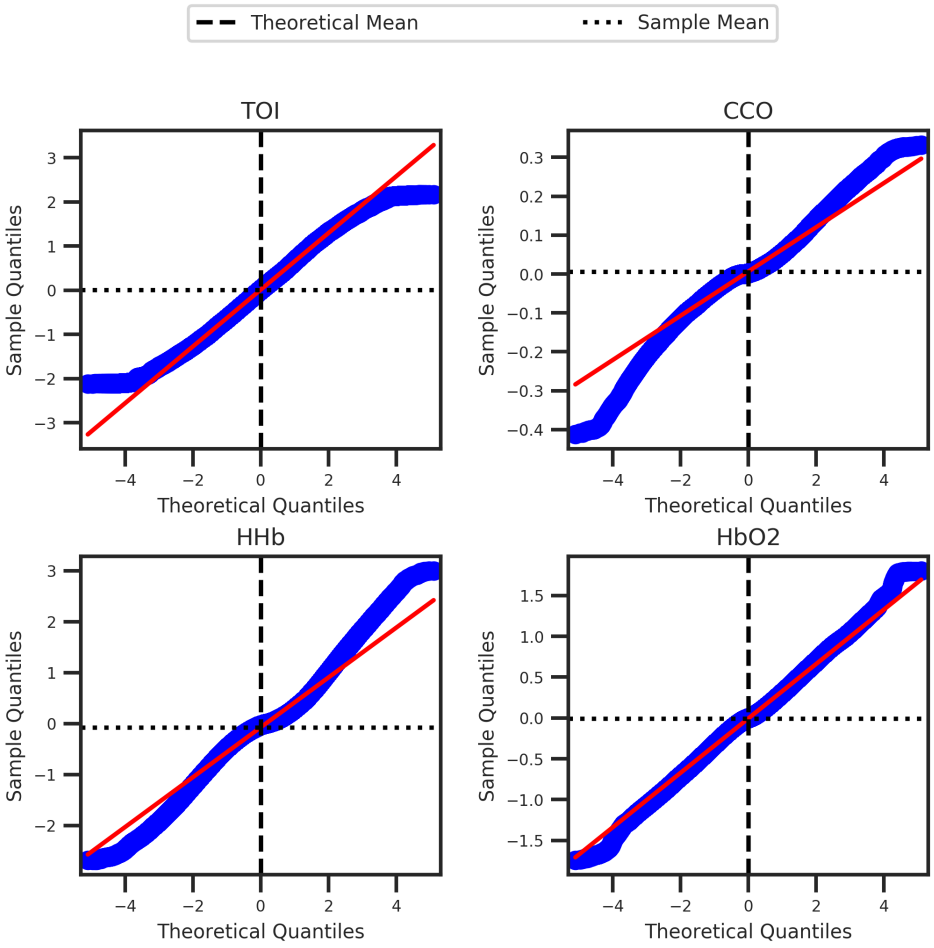

Figure C.9: Q-Q plot of residuals. Q-Q plots looking at the residuals for each signal. The plots use a standardised diagonal line, where the expected order statistics are scaled by the standard deviation of the sample residuals and have the mean added to them.

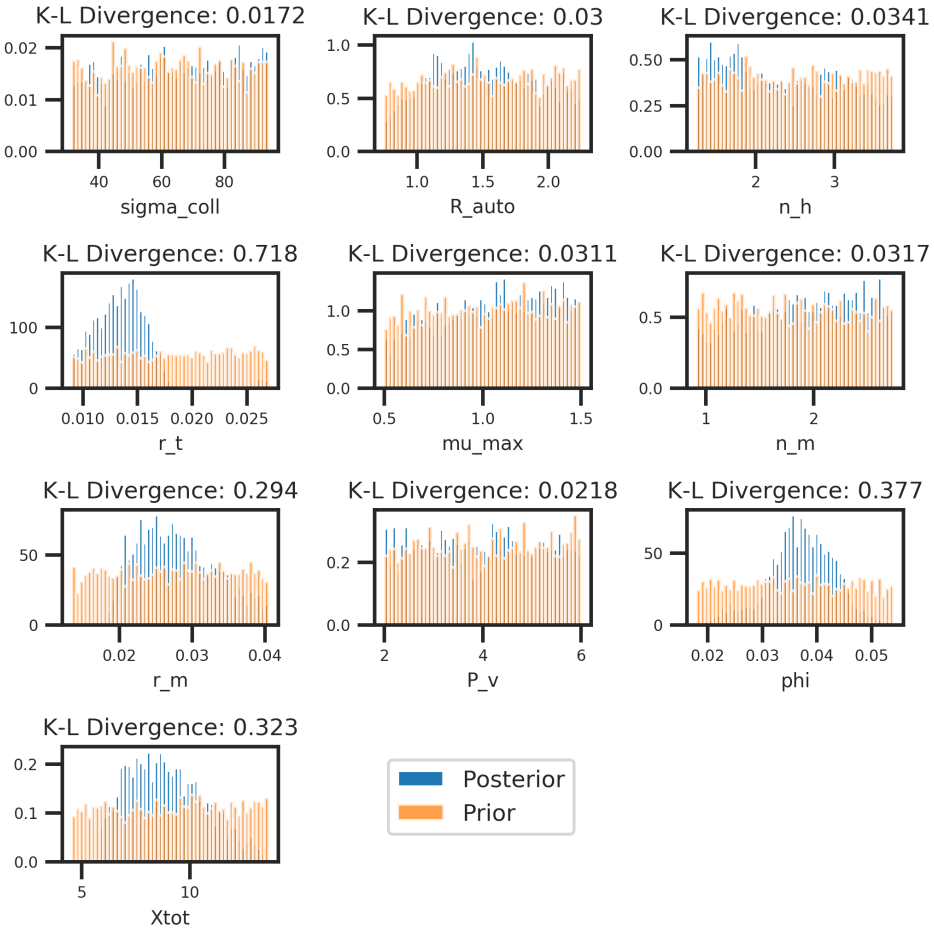

Figure C.10: Comparison of marginal posterior and prior distributions for each parameter. A clearer comparison is made here between the marginal prior and posterior distributions for each fitted parameter. The Kullback–Leibler divergence is shown for each.

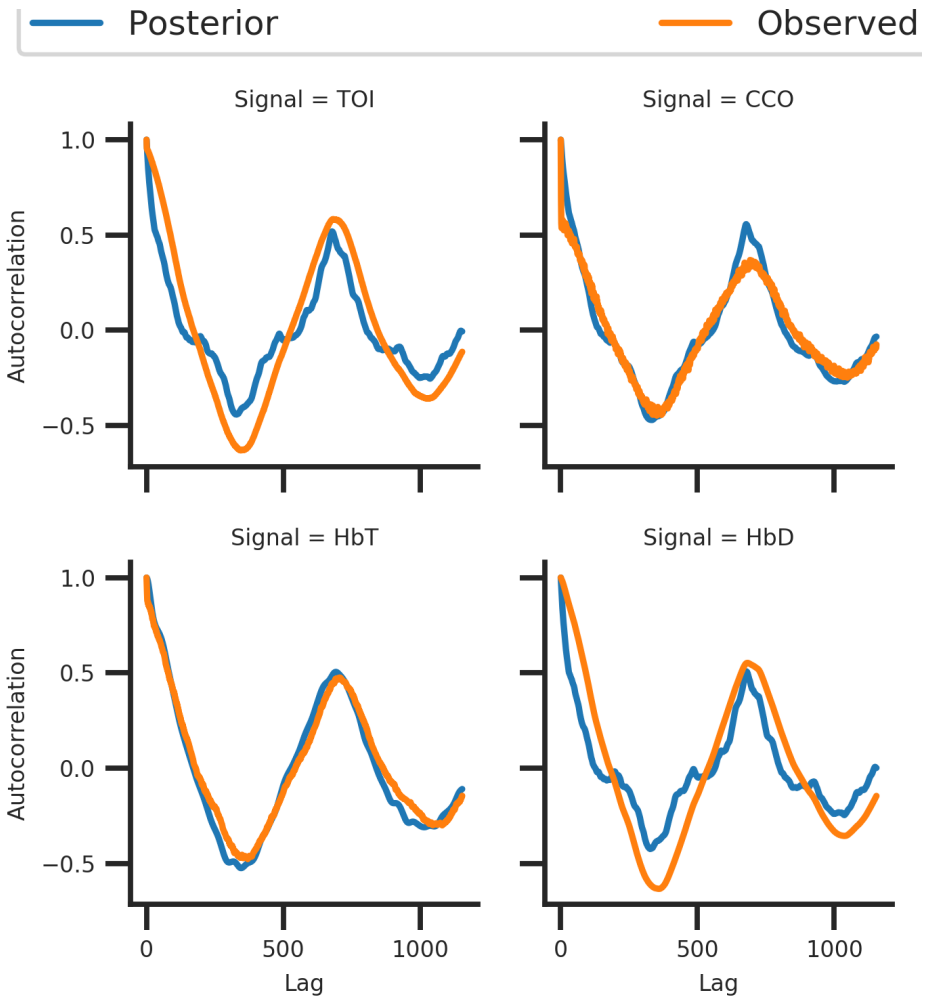

Figure C.11: Autocorrelation of posterior predictive and the observed data. The graph shows a comparison of the autocorrelation value between observed and posterior predictive time series as a function of lag. We see significant overlap between the two suggesting an extremely good fit.

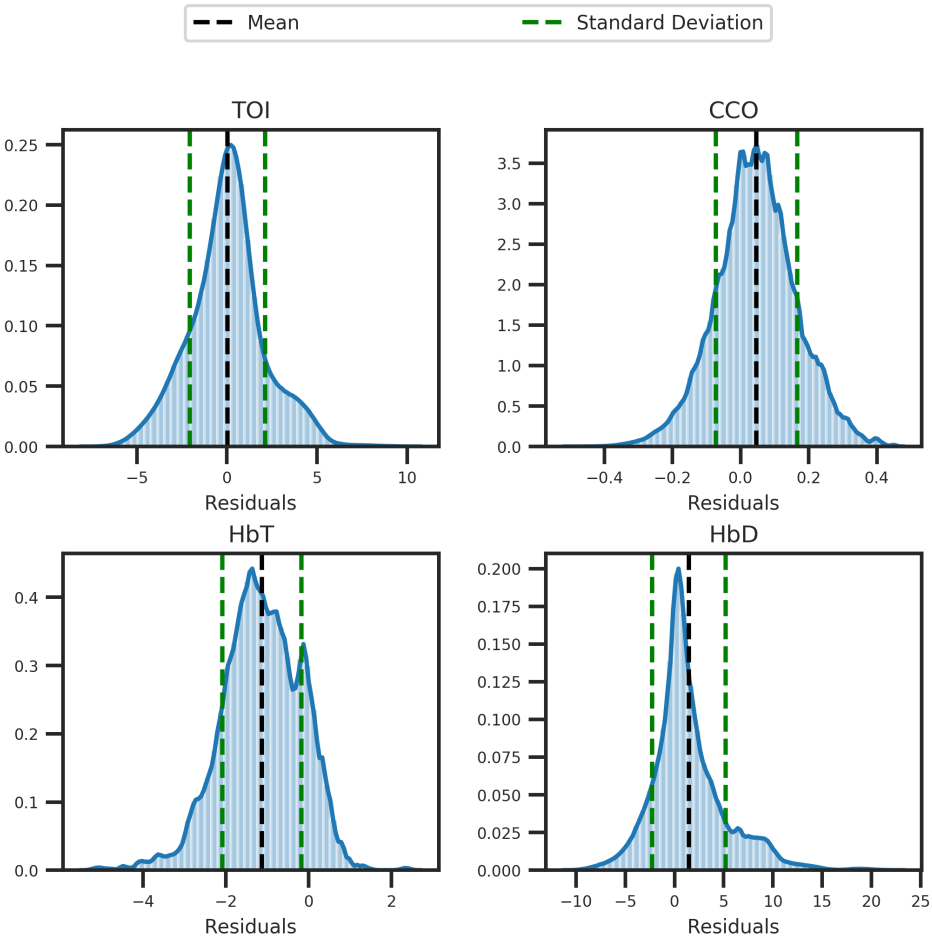

Figure C.12: Distributions of residuals. Distributions of the residuals for each signal. Mean and standard deviation of each plot are indicated by black and green lines respectively.

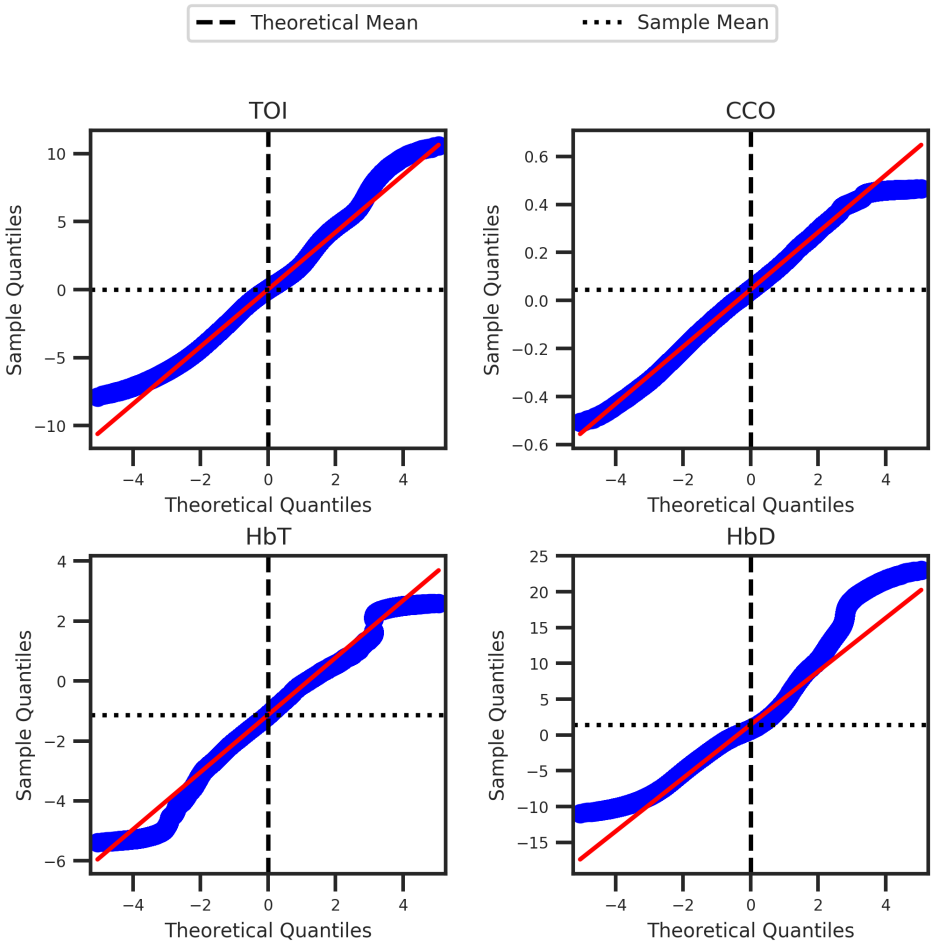

Figure C.13: Q-Q plot of residuals. Q-Q plots looking at the residuals for each signal. The plots use a standardised diagonal line, where the expected order statistics are scaled by the standard deviation of the sample residuals and have the mean added to them.

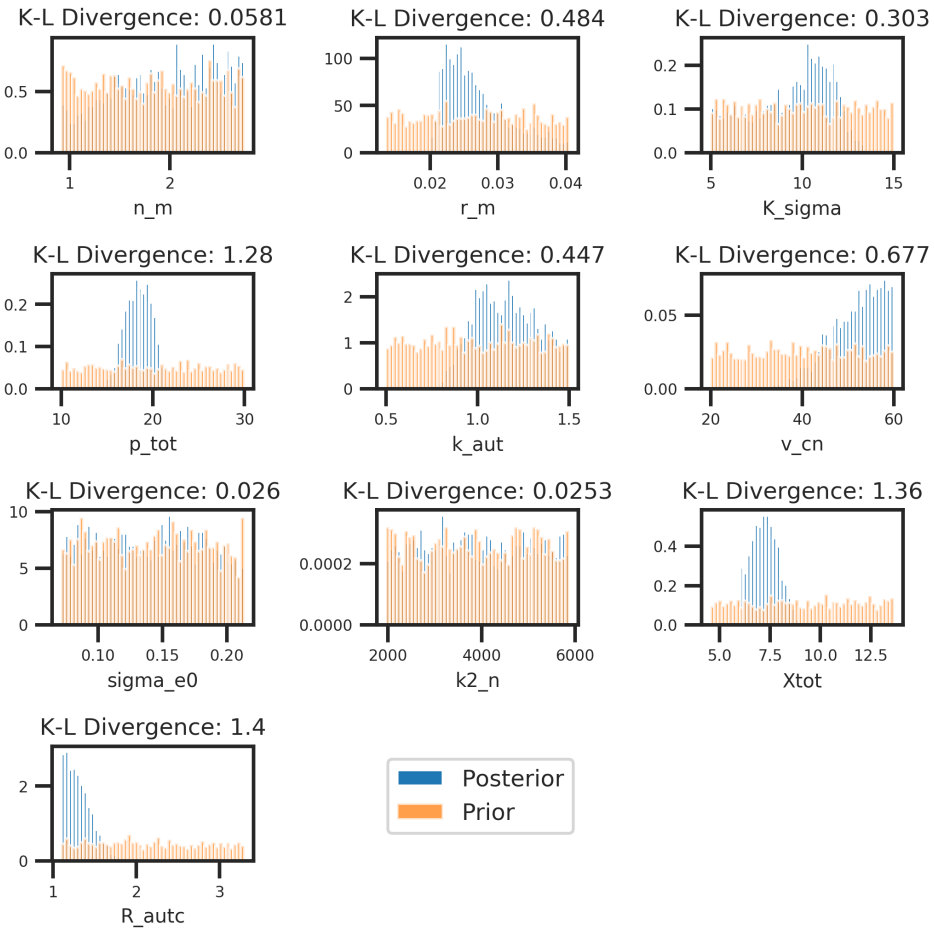

Figure C.14: Comparison of marginal posterior and prior distributions for each parameter. A clearer comparison is made here between the marginal prior and posterior distributions for eacg fitted parameter. The Kullback–Leibler divergence is shown for each.

Bibliography

Jordi Alastruey, S. M. Moore, K. H. Parker, T. David, J. Peiró, and S. J. Sherwin. Reduced modelling of blood flow in the cerebral circulation: Coupling 1-D, 0-D and cerebral auto-regulation models. *International Journal for Numerical Methods in Fluids*, 56(8):1061–1067, 3 2008. ISSN 02712091. doi:10.1002/fld.1606. URL http://doi.wiley.com/10.1002/fld.1606.

Takuo Aoyagi. Pulse oximetry: its invention, theory, and future. *Journal of Anesthesia*, 17(4):259–266, 11 2003. ISSN 0913-8668. doi:10.1007/s00540-003-0192-6. URL http://link.springer.com/10.1007/s00540-003-0192-6.

Joshua F. Apgar, David K. Witmer, Forest M. White, and Bruce Tidor. Sloppy models, parameter uncertainty, and the role of experimental design. *Molecular BioSystems*, 6(10):1890, 2010. ISSN 1742-206X. doi:10.1039/b918098b. URL http://xlink.rsc.org/?DOI=b918098b.

Agnès Aubert and Robert Costalat. A model of the coupling between brain electrical activity, metabolism, and hemodynamics: Application to the interpretation of functional neuroimaging. *NeuroImage*, 17(3):1162–1181, 11 2002. ISSN 10538119. doi:10.1006/nimg.2002.1224. URL https://www.sciencedirect.com/science/article/pii/S1053811902912243.

Agnès Aubert and Robert Costalat. Interaction between Astrocytes and Neurons Studied using a Mathematical Model of Compartmentalized Energy Metabolism. *Journal of Cerebral Blood Flow & Metabolism*, 25(11):1476–1490, 11 2005. ISSN 0271-678X. doi:10.1038/sj.jcbfm.9600144. URL http://journals.sagepub.com/doi/10.1038/sj.jcbfm.9600144.

Agnès Aubert, Robert Costalat, and Romain Valabrègue. Modelling of the coupling between brain electrical activity and metabolism. *Acta Biotheoretica*, 49(4):301–326, 2001. ISSN 00015342. doi:10.1023/A:1014286728421.

Denis V. Azzopardi, Brenda Strohm, A. David Edwards, Leigh Dyet, Henry L. Halliday, Edmund Juszczak, Olga Kapellou, Malcolm Levene, Neil Marlow, Emma Porter, Marianne Thoresen, Andrew Whitelaw, and Peter Brocklehurst. Moderate Hypothermia to Treat Perinatal Asphyxial Encephalopathy. *New England Journal of Medicine*, 361(14):1349–1358, 10 2009. ISSN 0028-4793. doi:10.1056/NEJMoa0900854. URL `http://www.nejm.org/doi/abs/10.1056/NEJMoa0900854`.

Gemma Bale, Subhabrata Mitra, Judith Meek, Nicola Robertson, and Ilias Tachtsidis. A new broadband near-infrared spectroscopy system for in-vivo measurements of cerebral cytochrome-c-oxidase changes in neonatal brain injury. *Biomedical optics express*, 5(10):3450–66, 2014. ISSN 2156-7085. doi:10.1364/BOE.5.003450. URL `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4206316&tool=pmcentrez&rendertype=abstract`.

Gemma Bale, Clare E Elwell, and Ilias Tachtsidis. From J{ö}bsis to the present day: a review of clinical near-infrared spectroscopy measurements of cerebral cytochrome-c-oxidase. *Journal of biomedical optics*, 21(9):91307, 2016.

Gemma Bale, Subhabrata Mitra, Isabel de Roever, Magdalena Sokolska, David Price, Alan Bainbridge, Roxana Gunny, Cristina Uria-Avellanal, Giles S Kendall, Judith Meek, Nicola J Robertson, and Ilias Tachtsidis. Oxygen dependency of mitochondrial metabolism indicates outcome of newborn brain injury. *Journal of Cerebral Blood Flow & Metabolism*, page 0271678X1877792, 5 2018. ISSN 0271-678X. doi:10.1177/0271678X18777928. URL `http://journals.sagepub.com/doi/10.1177/0271678X18777928`.

Murad Banaji. The BRAINCIRC model: Modelling the Brain Circulation, 2004. URL http://www.medphys.ucl.ac.uk/braincirc/.

Murad Banaji and Tracy Moroz. The Braincirc modelling environment. URL https://sourceforge.net/projects/braincirc/.

Murad Banaji, Ilias Tachtsidis, David Delpy, and Stephen Baigent. A physiological model of cerebral blood flow control. *Mathematical Biosciences*, 194(2):125–173, 2005. ISSN 00255564. doi:10.1016/j.mbs.2004.10.005.

Murad Banaji, Alfred Mallet, Clare E. Elwell, Peter Nicholls, and Chris E. Cooper. A model of brain circulation and metabolism: NIRS signal changes during physiological challenges. *PLoS Computational Biology*, 4(11), 2008. ISSN 1553734X. doi:10.1371/journal.pcbi.1000212.

A J Barkovich, K Baranski, D Vigneron, J C Partridge, D K Hallam, B L Hajnal, and D M Ferriero. Proton MR spectroscopy for the evaluation of brain injury in asphyxiated, term neonates. *AJNR. American journal of neuroradiology*, 20(8): 1399–405, 9 1999. ISSN 0195-6108. URL http://www.ncbi.nlm.nih.gov/pubmed/10512219.

M. A. Beaumont, J.-M. Cornuet, J.-M. Marin, and C. P. Robert. Adaptive approximate Bayesian computation. *Biometrika*, 96(4):983–990, 12 2009. ISSN 0006-3444. doi:10.1093/biomet/asp052. URL https://academic.oup.com/biomet/article-lookup/doi/10.1093/biomet/asp052.

Mark A. Beaumont. Approximate Bayesian Computation in Evolution and Ecology. *Annual Review of Ecology, Evolution, and Systematics*, 41 (1):379–406, 12 2010. ISSN 1543-592X. doi:10.1146/annurev-ecolsys-102209-144621. URL http://www.annualreviews.org/doi/10.1146/annurev-ecolsys-102209-144621.

Mark A Beaumont, Wenyang Zhang, and David J Balding. Approximate Bayesian Computation in Population Genetics. *Genetics*, 162(4):2025 LP – 2035, 12 2002. URL http://www.genetics.org/content/162/4/2025.abstract.

Solenna Blanchard, Sandrine Saillet, Anton Ivanov, Pascal Benquet, Christian-George Bénar, Mélanie Pélégrini-Issac, Habib Benali, and Fabrice Wendling. A New Computational Model for Neuro-Glio-Vascular Coupling: Astrocyte Activation Can Explain Cerebral Blood Flow Nonlinear Response to Interictal Events. *PLOS ONE*, 11(2):e0147292, 2 2016. ISSN 1932-6203. doi:10.1371/journal.pone.0147292. URL https://dx.plos.org/10.1371/journal.pone.0147292.

Stefan Blüml. Magnetic Resonance Spectroscopy: Basics. In *MR Spectroscopy of Pediatric Brain Disorders*, pages 11–23. Springer New York, New York, NY, 2013. doi:10.1007/978-1-4419-5864-8_2. URL http://link.springer.com/10.1007/978-1-4419-5864-8_2.

George E. P. Box. Science and Statistics. *Journal of the American Statistical Association*, 71(356):791, 12 1976. ISSN 01621459. doi:10.2307/2286841. URL https://www.jstor.org/stable/2286841?origin=crossref.

George E.P. Box. Robustness in the Strategy of Scientific Model Building. In *Robustness in Statistics*, volume 71, pages 201–236. Elsevier, 12 1979. doi:10.1016/B978-0-12-438150-6.50018-2. URL https://www.jstor.org/stable/2286841?origin=crossrefhttps://linkinghub.elsevier.com/retrieve/pii/B9780124381506500182.

Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, 2001. ISSN 08856125. doi:10.1023/A:1010933404324. URL http://link.springer.com/10.1023/A:1010933404324.

Kevin S. Brown and James P. Sethna. Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E*, 68(2):021904, 8 2003. ISSN 1063-651X. doi:10.1103/PhysRevE.68.021904. URL https://link.aps.org/doi/10.1103/PhysRevE.68.021904.

Richard B. Buxton, Eric C. Wong, and Lawrence R. Frank. Dynamics of blood flow and oxygenation changes during brain activation: The balloon model.

Magnetic Resonance in Medicine, 39(6):855–864, 6 1998. ISSN 07403194. doi:10.1002/mrm.1910390602. URL http://doi.wiley.com/10.1002/mrm.1910390602.

Ernest B. Cady, Osuke Iwata, Alan Bainbridge, John S. Wyatt, and Nicola J. Robertson. Phosphorus magnetic resonance spectroscopy 2 h after perinatal cerebral hypoxia-ischemia prognosticates outcome in the newborn piglet. *Journal of Neurochemistry*, 10 2008. ISSN 00223042. doi:10.1111/j.1471-4159.2008.05662.x. URL http://doi.wiley.com/10.1111/j.1471-4159.2008.05662.x.

Matthew Caldwell. BCMD modelling system 0.6b, 2015. URL https://github.com/bcmd/BCMD/blob/master/doc/manual.pdf.

Matthew Caldwell, Tharindi Hapuarachchi, David Highton, Clare Elwell, Martin Smith, and Ilias Tachtsidis. BrainSignals revisited: Simplifying a computational model of cerebral physiology. *PLoS ONE*, 10(5):1–28, 2015a. ISSN 19326203. doi:10.1371/journal.pone.0126695.

Matthew Caldwell, Tracy Moroz, Tharindi Hapuarachchi, Alan Bainbridge, Nicola J. Robertson, Chris E. Cooper, and Ilias Tachtsidis. Modelling blood flow and metabolism in the preclinical neonatal brain during and following hypoxic-ischaemia. *PLoS ONE*, 10(10):1–25, 2015b. ISSN 19326203. doi:10.1371/journal.pone.0140171.

Matthew Caldwell, Felix Scholkmann, Ursula Wolf, Martin Wolf, Clare Elwell, and Ilias Tachtsidis. Modelling confounding effects from extracerebral contamination and systemic factors on functional near-infrared spectroscopy. *NeuroImage*, 143:91–105, 2016. ISSN 10959572. doi:10.1016/j.neuroimage.2016.08.058. URL http://ac.els-cdn.com/S1053811916304505/1-s2.0-S1053811916304505-main.pdf?_tid=5d36054a-4224-11e7-8614-00000aab0f6b&acdnat=1495811140_472fbdf3de5edbf34d535003f001dc39.

D. Calvetti, G. Capo Rangel, L. Gerardo Giorda, and E. Somersalo. A computational model integrating brain electrophysiology and metabolism highlights the key role of extracellular potassium and oxygen. *Journal of Theoretical Biology*, 446:238–258, 6 2018. ISSN 00225193. doi:10.1016/j.jtbi.2018.02.029. URL `https://linkinghub.elsevier.com/retrieve/pii/S0022519318300900`.

Daniela Calvetti and Erkki Somersalo. Dynamic activation model for a glutamatergic neurovascular unit. *Journal of Theoretical Biology*, 274(1):12–29, 4 2011. ISSN 00225193. doi:10.1016/j.jtbi.2010.12.007. URL `https://linkinghub.elsevier.com/retrieve/pii/S0022519310006624`.

Daniela Calvetti, Yougan Cheng, and Erkki Somersalo. A spatially distributed computational model of brain cellular metabolism. *Journal of Theoretical Biology*, 376:48–65, 7 2015. ISSN 00225193. doi:10.1016/j.jtbi.2015.03.037. URL `https://linkinghub.elsevier.com/retrieve/pii/S0022519315001526`.

Francesca Campolongo, Jessica Cariboni, and Andrea Saltelli. An effective screening design for sensitivity analysis of large models. *Environmental Modelling and Software*, 22(10):1509–1518, 10 2007. ISSN 13648152. doi:10.1016/j.envsoft.2006.10.004.

O. Cappé, A. Guillin, J. M. Marin, and C. P. Robert. Population Monte Carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929, 12 2004. ISSN 10618600. doi:10.1198/106186004X12803. URL `http://www.tandfonline.com/doi/abs/10.1198/106186004X12803`.

F. Cassot, M. Zagzoule, and J. P. Marc-Vergnes. Hemodynamic role of the circle of Willis in stenoses of internal carotid arteries. An analytical solution of a linear model. *Journal of Biomechanics*, 33(4):395–405, 2000. ISSN 00219290. doi:10.1016/S0021-9290(99)00193-1.

Nancy Chinchor. Muc-4 evaluation metrics. In *Proceedings of the 4th Conference on Message Understanding*, MUC4 '92, page 22–29, USA, 1992. Association for

Computational Linguistics. ISBN 1558602739. doi:10.3115/1072064.1072067. URL https://doi.org/10.3115/1072064.1072067.

Oana-Teodora Chis, Julio R. Banga, and Eva Balsa-Canto. Structural Identifiability of Systems Biology Models: A Critical Comparison of Methods. *PLoS ONE*, 6(11):e27755, 11 2011. ISSN 1932-6203. doi:10.1371/journal.pone.0027755. URL http://dx.plos.org/10.1371/journal.pone.0027755.

Oana-Teodora Chis, Alejandro F. Villaverde, Julio R. Banga, and Eva Balsa-Canto. On the relationship between sloppiness and identifiability. *Mathematical Biosciences*, 282:147–161, 12 2016. ISSN 0025-5564. doi:10.1016/J.MBS.2016.10.009. URL https://www.sciencedirect.com/science/article/pii/S0025556416302668.

Mathieu Cloutier, Fiachra B. Bolger, John P. Lowry, and Peter Wellstead. An integrative dynamic model of brain energy metabolism using in vivo neurochemical measurements. *Journal of Computational Neuroscience*, 27(3):391–414, 12 2009. ISSN 0929-5313. doi:10.1007/s10827-009-0152-8. URL http://link.springer.com/10.1007/s10827-009-0152-8.

Vanessa Coelho-Santos and Andy Y. Shih. Postnatal development of cerebrovascular structure and the neurogliovascular unit. *WIREs Developmental Biology*, 9(2), 3 2020. ISSN 1759-7684. doi:10.1002/wdev.363. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/wdev.363.

Chris E. Cooper and Roger Springett. Measurement of cytochrome oxidase and mitochondrial energetics by near–infrared spectroscopy. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 352 (1354):669–676, 6 1997. ISSN 0962-8436. doi:10.1098/rstb.1997.0048. URL https://royalsocietypublishing.org/doi/10.1098/rstb.1997.0048.

F M Cowan, J M Pennock, J D Hanrahan, K P Manji, and A D Edwards. Early detection of cerebral infarction and hypoxic ischemic encephalopathy in neonates using diffusion-weighted magnetic resonance imaging. *Neuropediatrics*, 25(4):

172–5, 8 1994. ISSN 0174-304X. doi:10.1055/s-2008-1073018. URL http://www.ncbi.nlm.nih.gov/pubmed/7824088.

M. Czosnyka, N. G. Harris, J. D. Pickard, and S. Piechnik. CO2 cerebrovascular reactivity as a function of perfusion pressure - a modelling study. *Acta Neurochirurgica*, 121(3-4):159–165, 1993. ISSN 00016268. doi:10.1007/BF01809269.

Marek Czosnyka, Stefan Piechnik, Hugh K. Richards, Peter Kirkpatrick, Piotr Smielewski, and John D. Pickard. Contribution of mathematical modelling to the interpretation of bedside tests of cerebrovascular autoregulation. *Journal of Neurology Neurosurgery and Psychiatry*, 63(6):721–731, 1997. ISSN 00223050. doi:10.1136/jnnp.63.6.721.

Isabel de Roever, Gemma Bale, Robert J Cooper, and Ilias Tachtsidis. Functional NIRS Measurement of Cytochrome-C-Oxidase Demonstrates a More Brain-Specific Marker of Frontal Lobe Activation Compared to the Haemoglobins. In *Oxygen Transport to Tissue XXXIX*, pages 141–147. Springer, 2017.

Linda S de Vries and Floris Groenendaal. Patterns of neonatal hypoxic-ischaemic brain injury. *Neuroradiology*, 52(6):555–66, 6 2010. ISSN 1432-1920. doi:10.1007/s00234-010-0674-9. URL http://www.ncbi.nlm.nih.gov/pubmed/20390260http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2872019.

D T Delpy, M Cope, P van der Zee, S Arridge, S Wray, and J Wyatt. Estimation of optical pathlength through tissue from direct time of flight measurement. *Physics in Medicine and Biology*, 33(12):1433–1442, 12 1988. ISSN 0031-9155. doi:10.1088/0031-9155/33/12/008. URL https://iopscience.iop.org/article/10.1088/0031-9155/33/12/008.

Mauro DiNuzzo, Silvia Mangia, Bruno Maraviglia, and Federico Giove. Changes in Glucose Uptake Rather than Lactate Shuttle Take Center Stage in Subserving Neuroenergetics: Evidence from Mathematical Modeling. *Journal of Cerebral Blood Flow & Metabolism*, 30(3):586–602, 3 2010. ISSN 0271-678X.

doi:10.1038/jcbfm.2009.232. URL http://journals.sagepub.com/doi/10.1038/jcbfm.2009.232.

Kim-Anh Do, Peter Müller, and Marina Vannucci. *Bayesian inference for gene expression and proteomics*. Cambridge University Press, 2006.

Pedro Domingos. A few useful things to know about machine learning. *Communications of the ACM*, 55(10):78–87, 2012.

Lars Edvinsson, Eric T MacKenzie, and James McCulloch. *Cerebral Blood Flow and Metabolism*. Nature Publishing Group, 2002. ISBN 9788578110796. doi:10.1016/B0-12-226870-9/00190-8. URL http://dx.doi.org/10.1016/B0-12-226870-9/00190-8.

Richard Effros. *Microcirculation: current physiologic, medical, and surgical concepts*. Elsevier, 2012.

Marek P Ehrlich, Jock N McCullough, Ning Zhang, Donald J Weisz, Tatu Juvonen, Carol A Bodian, and Randall B Griepp. Effect of hypothermia on cerebral blood flow and metabolism in the pig. *The Annals of Thoracic Surgery*, 73(1):191–197, 1 2002. ISSN 00034975. doi:10.1016/S0003-4975(01)03273-8. URL http://linkinghub.elsevier.com/retrieve/pii/S0003497501032738.

S Fantini. A haemodynamic model for the physiological interpretation of in vivo measurements of the concentration and oxygen saturation of haemoglobin. *Phys Med Biol*, 47(18):249–57, 2002. ISSN 0031-9155. doi:10.1088/0031-9155/47/18/402. URL http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=12375832.

Sergio Fantini. A new hemodynamic model shows that temporal perturbations of cerebral blood flow and metabolic rate of oxygen cannot be measured individually using functional near-infrared spectroscopy. *Physiological Measurement*, 35(1):N1–N9, 1 2014. ISSN 0967-3334. doi:10.1088/0967-3334/35/1/N1. URL http://stacks.iop.org/0967-3334/35/i=1/a=N1?key=crossref.cec58af990d93032143fd91f47fc92da.

Ali Fatemi, Mary Ann Wilson, and Michael V Johnston. Hypoxic-ischemic encephalopathy in the term infant. *Clinics in perinatology*, 36(4): 835–58, 12 2009. ISSN 1557-9840. doi:10.1016/j.clp.2009.07.011. URL http://www.ncbi.nlm.nih.gov/pubmed/19944838http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2849741.

Edward J. Fine, Catalina C. Ionita, and Linda Lohr. The History of the Development of the Cerebellar Examination. *Seminars in Neurology*, 22(4):375–384, 2002. ISSN 0271-8235. doi:10.1055/s-2002-36759. URL http://www.thieme-connect.de/DOI/DOI?10.1055/s-2002-36759.

K P Forbes, J G Pipe, and R Bird. Neonatal hypoxic-ischemic encephalopathy: detection with diffusion-weighted MR imaging. *AJNR. American journal of neuroradiology*, 21(8):1490–6, 9 2000. ISSN 0195-6108. URL http://www.ncbi.nlm.nih.gov/pubmed/11003285.

Herbert I. Freedman. *Deterministic mathematical models in population ecology*. M. Dekker, 1980. ISBN 0824766539. URL https://books.google.is/books/about/Deterministic_mathematical_models_in_pop.html?id=0ggUAQAAIAAJ&redir_esc=y.

Håvard G. Frøysa, Shirin Fallahi, and Nello Blaser. Evaluating model reduction under parameter uncertainty. *BMC Systems Biology*, 12(1):79, 12 2018. ISSN 1752-0509. doi:10.1186/s12918-018-0602-x. URL https://bmcsystbiol.biomedcentral.com/articles/10.1186/s12918-018-0602-x.

Erzhen Gao, William L. Young, John Pile-Spellman, Eugene Ornstein, and Qiyuan Ma. Mathematical considerations for modeling cerebral blood flow autoregulation to systemic arterial pressure. *American Journal of Physiology - Heart and Circulatory Physiology*, 274(3 43-3), 1998. ISSN 03636135.

Andrew Gelman, Xiao-Li Meng, and Hal Stern. POSTERIOR PREDICTIVE ASSESSMENT OF MODEL FITNESS VIA REALIZED DISCREPANCIES. *Sta-*

tistica Sinica, 6:733–807, 1996. URL http://www3.stat.sinica.edu.tw/statistica/oldpdf/A6n41.pdf.

Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 2013.

Andrew Gelman, John B Carlin, Hal S Stern, David B. Dunson, Aki Vehtari, and Donald B Rubin. Model checking. In *Bayeisan Data Analysis*, chapter 6, pages 141–164. Taylor \& Francis Group, third edition, 2014.

Albert Gjedde. Cerebral blood flow change in arterial hypoxemia is consistent with negligible oxygen tension in brain mitochondria. *NeuroImage*, 17(4):1876–1881, 2002. ISSN 10538119. doi:10.1006/nimg.2002.1272.

Mark E. Glickman and David A. Van Dyk. *Basic Bayesian Methods*. Springer, 2007. doi:10.1007/978-1-59745-530-5_16. URL http://link.springer.com/10.1007/978-1-59745-530-5_16.

Peter D. Gluckman, John S. Wyatt, Denis Azzopardi, Roberta Ballard, A. David Edwards, Donna M. Ferriero, Richard A. Polin, Charlene M. Robertson, Marianne Thoresen, Andrew Whitelaw, and Alistair J. Gunn. Selective head cooling with mild systemic hypothermia after neonatal encephalopathy: Multicentre randomised trial. *Lancet*, 365(9460):663–670, 2 2005. ISSN 01406736. doi:10.1016/S0140-6736(05)17946-X.

Elke M. Golding and Raymund M. Golding. Mathematical modelling of responses of cerebral blood vessels to changing intraluminal pressure. *Physiological Measurement*, 22(4):727–743, 2001. ISSN 09673334. doi:10.1088/0967-3334/22/4/307.

Daniel Goldman. Theoretical Models of Microvascular Oxygen Transport to Tissue. *Microcirculation*, 15(8):795–811, 1 2008. ISSN 1073-9688. doi:10.1080/10739680801938289. URL http://doi.wiley.com/10.1080/10739680801938289.

Sachin K Gujar, Sharad Maheshwari, Isabella Björkman-Burtscher, and Pia C Sundgren. Magnetic resonance spectroscopy. *J Neuroophthalmol*, 25(3): 217–226, 9 2005. ISSN 1070-8022. URL `http://www.ncbi.nlm.nih.gov/pubmed/16148633http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=16148633&retmode=ref&cmd=prlinks`.

Ryan N. Gutenkunst, Joshua J. Waterfall, Fergal P. Casey, Kevin S. Brown, Christopher R. Myers, and James P. Sethna. Universally sloppy parameter sensitivities in systems biology models. *PLoS Computational Biology*, 3(10):1871–1878, 2007. ISSN 1553734X. doi:10.1371/journal.pcbi.0030189.

Ernst Hairer, Gerhard Wanner, and others. *Solving ordinary differential equations II: Stiff and differential-algebraic problems*. Springer, second rev edition, 1996.

John E. Hall and Arthur C. Guyton. Guyton and Hall Textbook of Medical Physiology. In *Journal of Chemical Information and Modeling*, volume 53, pages 502–509. Saunders Elsevier, 2011.

T. Hapuarachchi, T. Moroz, A. Bainbridge, S. Faulkner, D. Price, K. D. Broad, D. Thomas, E. Cady, X. Golay, Nicola Robertson, and Ilias Tachtsidis. Simulating NIRS and MRS Measurements During Cerebral Hypoxia-Ischaemia in Piglets Using a Computational Model. In *Advances in Experimental Medicine and Biology*, volume 812, pages 187–194. Springer, New York NY, 2014. ISBN 0065-2598 (Print)\r0065-2598 (Linking). doi:10.1007/978-1-4939-0620-8_25. URL `http://link.springer.com/10.1007/978-1-4939-0620-8_25`.

T. Hapuarachchi, F. Scholkmann, M. Caldwell, C. Hagmann, S. Kleiser, A. J. Metz, M. Pastewski, M. Wolf, and I. Tachtsidis. Simulation of Preterm Neonatal Brain Metabolism During Functional Neuronal Activation Using a Computational Model. In *Advances in Experimental Medicine and Biology*, volume 876, pages 111–120. Springer, New York, NY, 2016. ISBN 9781493930227. doi:10.1007/978-1-4939-3023-

4_14. URL http://dx.doi.org/10.1007/978-1-4939-3023-4_14http://link.springer.com/10.1007/978-1-4939-3023-4_14.

Tharindi Hapuarachchi. Modelling metabolism in the neonatal brain. *Doctoral thesis, UCL (University College London).*, 2015. URL http://discovery.ucl.ac.uk/1471679/.

Tharindi Hapuarachchi, Tracy Moroz, Alan Bainbridge, David Price, Ernest Cady, Esther Baer, Kevin Broad, Mojgan Ezzati, David Thomas, Xavier Golay, Nicola J. Robertson, and Ilias Tachtsidis. Modelling Blood Flow and Metabolism in the Piglet Brain During Hypoxia-Ischaemia: Simulating pH Changes. In *Advances in Experimental Medicine and Biology*, volume 789, pages 331–337. Springer, New York NY, 2013. ISBN 9781461472568. doi:10.1007/978-1-4614-7411-1_44. URL http://link.springer.com/10.1007/978-1-4614-7411-1_44.

W. Henry. Experiments on the Quantity of Gases Absorbed by Water, at Different Temperatures, and under Different Pressures. *Philosophical Transactions of the Royal Society of London*, 93:29–274, 1 1803. ISSN 0261-0523. doi:10.1098/rstl.1803.0004. URL http://rstl.royalsocietypublishing.org/cgi/doi/10.1098/rstl.1803.0004.

AV Hill. The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *Journal of Physiology*, 40, 1910.

Chih-Wei; Hsu, Chih-Chung; Chang, and Chih-Jen; Lin. A Practical Guide to Support Vector Classification. *BJU international*, 101(1):1396–400, 2008. ISSN 1464-410X. doi:10.1177/02632760022050997. URL http://www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf.

Benjamin Y Huang and Mauricio Castillo. Hypoxic-Ischemic brain injury: Imaging findings from birth to adulthood. *RadioGraphics*, 28(2):417–439, 3 2008. ISSN 0271-5333. doi:10.1148/rg.282075066. URL http://pubs.rsna.org/doi/abs/10.1148/rg.282075066.

Antal G. Hudetz, Karl A. Conger, James H. Halsey, Miklos Pal, Orsolya Dohan, and Arisztid G. B. Kovach. Pressure Distribution in the Pial Arterial System of Rats Based on Morphometric Data and Mathematical Models. *Journal of Cerebral Blood Flow & Metabolism*, 7(3):342–355, 6 1987. ISSN 0271-678X. doi:10.1038/jcbfm.1987.72. URL `http://journals.sagepub.com/doi/10.1038/jcbfm.1987.72`.

Timothy M Hughes, Suzanne Craft, and Oscar L Lopez. Review of 'the potential role of arterial stiffness in the pathogenesis of Alzheimer's disease'. *Neurodegenerative Disease Management*, 5(2):121–135, 4 2015. ISSN 1758-2024. doi:10.2217/nmt.14.53. URL `https://www.futuremedicine.com/doi/10.2217/nmt.14.53`.

John D. Hunter. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. ISSN 1521-9615. doi:10.1109/MCSE.2007.55. URL `http://ieeexplore.ieee.org/document/4160265/`.

Fahmeed Hyder, Robert G. Shulman, and Douglas L. Rothman. A model for the regulation of cerebral oxygen delivery. *Journal of Applied Physiology*, 85(2):554–564, 1998. ISSN 87507587.

Susan E Jacobs, Rod Hunt, William O Tarnow-Mordi, Terrie E Inder, and Peter G Davis. Cooling for newborns with hypoxic ischaemic encephalopathy. In Susan E Jacobs, editor, *Cochrane Database of Systematic Reviews*, number 4, page CD003311. John Wiley & Sons, Ltd, Chichester, UK, 10 2007. doi:10.1002/14651858.CD003311.pub2. URL `http://www.ncbi.nlm.nih.gov/pubmed/17943788http://doi.wiley.com/10.1002/14651858.CD003311.pub2`.

Harold Jeffreys. *The theory of probability*. OUP Oxford, 1998.

Beth Jelfs, Murad Banaji, Ilias Tachtsidis, Chris E. Cooper, and Clare E. Elwell. Modelling noninvasively measured cerebral signals during a hypoxemia challenge: Steps towards individualised modelling. *PLoS ONE*, 7(6):e38297,

6 2012. ISSN 19326203. doi:10.1371/journal.pone.0038297. URL http://dx.plos.org/10.1371/journal.pone.0038297.

F F Jöbsis. Noninvasive, infrared monitoring of cerebral and myocardial oxygen sufficiency and circulatory parameters. *Science (New York, N.Y.)*, 198(4323): 1264–7, 12 1977. ISSN 0036-8075. URL http://www.ncbi.nlm.nih.gov/pubmed/929199.

Renaud Jolivet, Jay S. Coggan, Igor Allaman, and Pierre J. Magistretti. Multi-timescale Modeling of Activity-Dependent Metabolic Coupling in the Neuron-Glia-Vasculature Ensemble. *PLoS Computational Biology*, 11(2), 2015. ISSN 15537358. doi:10.1371/journal.pcbi.1004036.

Andreas Jung, Rupert Faltermeier, Ralf Rothoerl, and Alexander Brawanski. A mathematical model of cerebral circulation and oxygen supply. *Journal of Mathematical Biology*, 51(5):491–507, 11 2005. ISSN 0303-6812. doi:10.1007/s00285-005-0343-5. URL http://link.springer.com/10.1007/s00285-005-0343-5.

Jana M. Kainerstorfer, Angelo Sassaroli, Bertan Hallacoglu, Michele L. Pierro, and Sergio Fantini. Practical Steps for Applying a New Dynamic Model to Near-Infrared Spectroscopy Measurements of Hemodynamic Oscillations and Transient Changes. *Academic Radiology*, 21(2):185–196, 2 2014. ISSN 10766332. doi:10.1016/j.acra.2013.10.012. URL https://linkinghub.elsevier.com/retrieve/pii/S107663321300500X.

Robert E Kass and Adrian E Raftery. Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995. doi:10.1080/01621459.1995.10476572. URL https://www.tandfonline.com/doi/abs/10.1080/01621459.1995.10476572.

Pardis Kaynezhad, Subhabrata Mitra, Gemma Bale, Cornelius Bauer, Ingran Lingam, Christopher Meehan, Adnan Avdic-Belltheus, Kathryn A.

Martinello, Alan Bainbridge, Nicola J. Robertson, and Ilias Tachtsidis. Quantification of the severity of hypoxic-ischemic brain injury in a neonatal preclinical model using measurements of cytochrome-c-oxidase from a miniature broadband-near-infrared spectroscopy system. *Neurophotonics*, 6(04):1, 11 2019. ISSN 2329-423X. doi:10.1117/1.NPh.6.4.045009. URL https://www.spiedigitallibrary.org/journals/neurophotonics/volume-6/issue-04/045009/Quantification-of-the-severity-of-hypoxic-ischemic-brain-injury-in/10.1117/1.NPh.6.4.045009.full.

Bruce E Kendall, Cheryl J Briggs, and William W Murdoch. Why do populations cycle? A synthesis of statistical and mechanistic modeling approaches. 1999. doi:10.1890/0012-9658(1999)080[1789:WDPCAS]2.0.CO;2. URL https://escholarship.org/content/qt7sr3f7bk/qt7sr3f7bk.pdf.

S. K. Kirkham, R. E. Craine, and A. A. Birch. A new mathematical model of dynamic cerebral autoregulation based on a flow dependent feedback mechanism. *Physiological Measurement*, 22(3):461–473, 2001. ISSN 09673334. doi:10.1088/0967-3334/22/3/305.

Hiroaki Kitano. Computational systems biology. *Nature*, 420:206, 11 2002. URL http://dx.doi.org/10.1038/nature01254http://10.0.4.14/nature01254.

Bhavani Shankar Kodali. Capnography Outside the Operating Rooms. *Anesthesiology*, 118(1):192–201, 1 2013. ISSN 0003-3022. doi:10.1097/ALN.0b013e318278c8b6. URL http://anesthesiology.pubs.asahq.org/Article.aspx?doi=10.1097/ALN.0b013e318278c8b6.

Bernard Korzeniewski and Wojciech Froncisz. An extended dynamic model of oxidative phosphorylation. *Biochimica et Biophysica Acta - Bioenergetics*, 1060 (2):210–223, 1991. ISSN 00052728. doi:10.1016/S0005-2728(09)91009-X.

A. CC. Lee, L. C. Mullany, J. M. Tielsch, J. Katz, S. K. Khatry, S. C. LeClerq, R. K.

Adhikari, S. R. Shrestha, and G. L. Darmstadt. Risk Factors for Neonatal Mortality Due to Birth Asphyxia in Southern Nepal: A Prospective, Community-Based Cohort Study. *PEDIATRICS*, 121(5):e1381–e1390, may 2008. ISSN 0031-4005. doi:10.1542/peds.2007-1966. URL `http://pediatrics.aappublications.org/cgi/doi/10.1542/peds.2007-1966`.

Jennifer K. Lee, Ken M. Brady, Jennifer O. Mytar, Kathleen K. Kibler, Erin L. Carter, Karen G. Hirsch, Charles W. Hogue, Ronald B. Easley, Lori C. Jordan, Peter Smielewski, Marek Czosnyka, Donald H. Shaffner, and Raymond C. Koehler. Cerebral blood flow and cerebrovascular autoregulation in a swine model of pediatric cardiac arrest and hypothermia*. *Critical Care Medicine*, 39(10):2337–2345, 10 2011. ISSN 0090-3493. doi:10.1097/CCM.0b013e318223b910. URL `https://insights.ovid.com/crossref?an=00003246-201110000-00019`.

Andreas A. Linninger, Michalis Xenos, Brian Sweetman, Sukruti Ponkshe, Xiaodong Guo, and Richard Penn. A mathematical model of blood, cerebrospinal fluid and brain dynamics. *Journal of Mathematical Biology*, 59(6):729–759, 12 2009. ISSN 0303-6812. doi:10.1007/s00285-009-0250-2. URL `http://link.springer.com/10.1007/s00285-009-0250-2`.

Jun S Liu and T Logvinenko. Bayesian methods in biological sequence analysis. *Balding, DJ*, 2003.

Xin Liu and Mahesan Niranjan. Parameter Estimation in Computational Biology by Approximate Bayesian Computation coupled with Sensitivity Analysis. 2017. URL `https://arxiv.org/pdf/1704.09021.pdf`.

Y. Liu and R. Allen. Analysis of dynamic cerebral autoregulation using an ARX model based on arterial blood pressure and middle cerebral artery velocity simulation. *Medical and Biological Engineering and Computing*, 40(5):600–605, 2002. ISSN 01400118. doi:10.1007/BF02345461.

C. A. Lodi, A. Ter Minassian, L. Beydon, M. Ursino, W. I. Rosenblum, J. L. Patterson, and J. R. Little. Modeling cerebral autoregulation and CO2 reactivity in patients with severe head injury. *American Journal of Physiology - Heart and Circulatory Physiology*, 274(Heart Circ. Physiol. 43):H1729–H1741, 5 1998. ISSN 0363-6135. URL http://ajpheart.physiology.org/content/274/5/H1715.long.

K. Lu, J. W. Clark, F. H. Ghorbel, C. S. Robertson, D. L. Ware, J. B. Zwischenberger, and A. Bidani. Cerebral Autoregulation and Gas Exchange Studied Using A Human Cardiopulmonary Model. In *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, volume 1, pages 395–397, 2003.

Paul Marjoram, John Molitor, Vincent Plagnol, and Simon Tavare. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences of the United States of America*, 100(26): 15324–8, 12 2003. ISSN 0027-8424. doi:10.1073/pnas.0306899100. URL http://www.ncbi.nlm.nih.gov/pubmed/14663152http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC307566.

Mark. Massé. *REST API design rulebook*. O'Reilly, 2012. ISBN 1449310508. URL https://books.google.co.uk/books?hl=en&lr=&id=eABpzyTcJNIC&oi=fnd&pg=PR3&dq=REST+API&ots=vyXA_6kdHH&sig=oh9m63mMxXxGW_l8u5Fv6zkvYYI.

E. J. Mathias, M. J. Plank, and T. David. A model of neurovascular coupling and the BOLD response: PART I. *Computer Methods in Biomechanics and Biomedical Engineering*, 20(5):508–518, 4 2017a. ISSN 1025-5842. doi:10.1080/10255842.2016.1255732. URL https://www.tandfonline.com/doi/full/10.1080/10255842.2016.1255732.

E. J. Mathias, M. J. Plank, and T. David. A model of neurovascular coupling and the BOLD response PART II. *Computer Methods in Biomechan-*

ics and Biomedical Engineering, 20(5):519–529, 4 2017b. ISSN 1025-5842. doi:10.1080/10255842.2016.1255733. URL https://www.tandfonline.com/doi/full/10.1080/10255842.2016.1255733.

Elshin J. Mathias, Allanah Kenny, Michael J. Plank, and Tim David. Integrated models of neurovascular coupling and BOLD signals: Responses for varying neural activations. *NeuroImage*, 174:69–86, 7 2018. ISSN 10538119. doi:10.1016/j.neuroimage.2018.03.010. URL https://linkinghub.elsevier.com/retrieve/pii/S1053811918302052.

Wes Mckinney. Data Structures for Statistical Computing in Python. *PROC. OF THE 9th PYTHON IN SCIENCE CONF*, pages 51–56, 2010. URL http://conference.scipy.org/proceedings/scipy2010/pdfs/mckinney.pdf.

Andreas Jaakko Metz, Martin Wolf, Peter Achermann, and Felix Scholkmann. A new approach for automatic removal of movement artifacts in near-infrared spectroscopy time series by means of acceleration data. *Algorithms*, 8(4):1052–1075, 2015. ISSN 19994893. doi:10.3390/a8041052.

Subhabrata Mitra, Gemma Bale, Judith Meek, Sean Mathieson, Cristina Uria, Giles Kendall, Nicola J. Robertson, and Ilias Tachtsidis. In Vivo Measurement of Cerebral Mitochondrial Metabolism Using Broadband Near Infrared Spectroscopy Following Neonatal Stroke. In *Oxygen Transport to Tissue XXXVII*, pages 493–500. Springer, New York, NY, 2016. doi:10.1007/978-1-4939-3023-4_62. URL http://link.springer.com/10.1007/978-1-4939-3023-4_62.

Subhabrata Mitra, Gemma Bale, David Highton, Roxanna Gunny, Cristina Uria-Avellanal, Alan Bainbridge, Magdalena Sokolska, David Price, Angela Huertas-Ceballos, Giles S Kendall, Judith Meek, Ilias Tachtsidis, and Nicola J Robertson. Pressure passivity of cerebral mitochondrial metabolism is associated with poor outcome following perinatal hypoxic ischemic brain injury. *Journal of Cerebral Blood Flow & Metabolism*, 0(00):1–13, 9 2017. ISSN 0271-

678X. doi:10.1177/0271678X17733639. URL http://journals.sagepub.com/doi/10.1177/0271678X17733639.

J Moorcraft, N M Bolas, N K Ives, R Ouwerkerk, J Smyth, B Rajagopalan, P L Hope, and G K Radda. Global and depth resolved phosphorus magnetic resonance spectroscopy to predict outcome after birth asphyxia. *Archives of disease in childhood*, 66(10 Spec No):1119–23, 10 1991. ISSN 1468-2044. URL http://www.ncbi.nlm.nih.gov/pubmed/1750759http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC1590291.

T Moroz. *Computational modelling of brain energy metabolism and circulation in the neonatal animal model*. PhD thesis, 2014.

Tracy Moroz, Murad Banaji, Nicola J. Robertson, Chris E. Cooper, and Ilias Tachtsidis. Computational modelling of the piglet brain to simulate near-infrared spectroscopy and magnetic resonance spectroscopy data collected during oxygen deprivation. *Journal of the Royal Society Interface*, 9(72):1499–1509, 7 2012a. ISSN 17425662. doi:10.1098/rsif.2011.0766. URL http://www.ncbi.nlm.nih.gov/pubmed/22279158http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3367814http://rsif.royalsocietypublishing.org/cgi/doi/10.1098/rsif.2011.0766.

Tracy Moroz, Murad Banaji, Martin Tisdall, Chris E. Cooper, Clare E. Elwell, and Ilias Tachtsidis. Development of a model to aid NIRS data interpretation: Results from a hypercapnia study in healthy adults. In *Advances in Experimental Medicine and Biology*, volume 737, pages 293–300. Springer, 2012b. ISBN 9781461415657. doi:10.1007/978-1-4614-1566-4_43.

Max D Morris. Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2):161–174, 1991.

Deirdre M Murray, Geraldine B Boylan, Cornelius A Ryan, and Sean Connolly. Early EEG findings in hypoxic-ischemic encephalopathy predict out-

comes at 2 years. *Pediatrics*, 124(3):459–67, 9 2009. ISSN 1098-4275. doi:10.1542/peds.2008-2190. URL http://www.ncbi.nlm.nih.gov/pubmed/19706569.

Mette S. Olufsen, Ali Nadim, and Lewis A. Lipsitz. Dynamics of cerebral blood flow regulation explained using a lumped parameter model. *American Journal of Physiology - Regulatory Integrative and Comparative Physiology*, 282(2 51-2), 2002. ISSN 03636119.

P. Orlowski, M. Chappell, C. S. Park, V. Grau, and S. Payne. Modelling of pH dynamics in brain cells after stroke. *Interface Focus*, 1(3):408–416, 6 2011. ISSN 2042-8898. doi:10.1098/rsfs.2010.0025. URL http://rsfs.royalsocietypublishing.org/cgi/doi/10.1098/rsfs.2010.0025.

P. Orlowski, D. O'Neill, V. Grau, Y. Ventikos, and S. Payne. Modelling of the physiological response of the brain to ischaemic stroke. *Interface Focus*, 3(2):20120079–20120079, 2 2013. ISSN 2042-8898. doi:10.1098/rsfs.2012.0079. URL http://rsfs.royalsocietypublishing.org/cgi/doi/10.1098/rsfs.2012.0079.

Piotr Orlowski, Flora Kennedy McConnell, and Stephen Payne. A Mathematical Model of Cellular Metabolism During Ischemic Stroke and Hypothermia. *IEEE Transactions on Biomedical Engineering*, 61(2):484–490, 2 2014. ISSN 0018-9294. doi:10.1109/TBME.2013.2282603. URL http://ieeexplore.ieee.org/document/6603279/.

R.B. Panerai. The critical closing pressure of the cerebral circulation. *Medical Engineering & Physics*, 25(8):621–632, 10 2003. ISSN 13504533. doi:10.1016/S1350-4533(03)00027-4. URL https://linkinghub.elsevier.com/retrieve/pii/S1350453303000274.

Ronney B. Panerai, Suzanne L. Dawson, and John F. Potter. Linear and nonlinear analysis of human dynamic cerebral autoregulation. *American Journal of Physiology-Heart and Circulatory Physiology*, 277(3):H1089–H1099, 9

1999. ISSN 0363-6135. doi:10.1152/ajpheart.1999.277.3.H1089. URL http://www.physiology.org/doi/10.1152/ajpheart.1999.277.3.H1089.

Stephen Payne. *Cerebral Blood Flow and Metabolism: A Quantitative Approach*. World Scientific, 2017.

Carol J. Peden, Frances M. Cowan, David J. Bryant, Janet Sargentoni, I. Jane Cox, David K. Menon, David G. Gadian, Jimmy D. Bell, and Lilly M. Dubowitz. Proton MR Spectroscopy of the Brain in Infants. *Journal of Computer Assisted Tomography*, 14(6):886–894, 11 1990. ISSN 0363-8715. doi:10.1097/00004728-199011000-00004. URL http://content.wkhealth.com/linkback/openurl?sid=WKPTLP:landingpage&an=00004728-199011000-00004.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(Oct):2825–2830, 2011. ISSN ISSN 1533-7928.

S. K. Piechnik, M. Czosnyka, N. G. Harris, P. S. Minhas, and J. D. Pickard. A model of the cerebral and cerebrospinal fluid circulations to examine asymmetry in cerebrovascular reactivity. *Journal of Cerebral Blood Flow and Metabolism*, 21(2):182–192, 2001. ISSN 0271678X. doi:10.1097/00004647-200102000-00010.

Michele L. Pierro, Bertan Hallacoglu, Angelo Sassaroli, Jana M. Kainerstorfer, and Sergio Fantini. Validation of a novel hemodynamic model for coherent hemodynamics spectroscopy (CHS) and functional brain studies with fNIRS and fMRI. *NeuroImage*, 85:222–233, 1 2014. ISSN 10538119. doi:10.1016/j.neuroimage.2013.03.037. URL https://linkinghub.elsevier.com/retrieve/pii/S1053811913002863.

Roland N. Pittman. Oxygen Transport, 2011. URL http://www.ncbi.nlm.nih.gov/books/NBK54103/.

J K Pritchard, M T Seielstad, A Perez-Lezaun, and M W Feldman. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular biology and evolution*, 16(12):1791–8, 12 1999. ISSN 0737-4038. URL http://www.ncbi.nlm.nih.gov/pubmed/10605120.

Nick Pullen and Richard J Morris. Bayesian model comparison and parameter inference in systems biology using nested sampling. *PloS one*, 9 (2):e88419, 2014. ISSN 1932-6203. doi:10.1371/journal.pone.0088419. URL http://www.ncbi.nlm.nih.gov/pubmed/24523891http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3921180.

Bryan A. Reyes, Julie S. Pendergast, and Shin Yamazaki. Mammalian Peripheral Circadian Oscillators Are Temperature Compensated. *Journal of Biological Rhythms*, 23(1):95–98, 2 2008. ISSN 0748-7304. doi:10.1177/0748730407311855. URL http://journals.sagepub.com/doi/10.1177/0748730407311855.

Joshua S Richman and J Randall Moorman. Physiological time-series analysis using approximate entropy and sample entropy. *American Journal of Physiology-Heart and Circulatory Physiology*, 278(6):H2039–H2049, 2000.

S C Roth, A D Edwards, E B Cady, D T Delpy, J S Wyatt, D Azzopardi, J Baudin, J Townsend, A L Stewart, and E O Reynolds. Relation between cerebral oxidative metabolism following birth asphyxia, and neurodevelopmental outcome and brain growth at one year. *Developmental medicine and child neurology*, 34 (4):285–95, 4 1992. ISSN 0012-1622. URL http://www.ncbi.nlm.nih.gov/pubmed/1572514.

Donald B Rubin. Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics*, 12(4):1151–1172, 1984.

J. Russell-Buckland, G. Bale, I. de Roever, and I. Tachtsidis. *ABroAD: A machine learning based approach to detect broadband NIRS artefacts*, volume 1072. 2018a. doi:10.1007/978-3-319-91287-5_51.

J. Russell-Buckland, P. Kaynezhad, S. Mitra, G. Bale, C. Bauer, I. Lingam, C. Meehan, A. Avdic-Belltheus, K. Martinello, A. Brainbridge, N.J. Robertson, and I. Tachtsidis. Systems biology model of cerebral oxygen delivery and metabolism during therapeutic hypothermia:application to the piglet model. International Society of Oxygen Transport to Tissue, Albuquerque, New Mexico, USA, 2019a.

Joshua Russell-Buckland and Ilias Tachtsidis. Developing a model to simulate the effect of hypothermia on cerebral blood flow and metabolism. In *Oxygen Transport to Tissue XLI*, pages 299–306. Springer, 2020.

Joshua Russell-Buckland, Gemma Bale, I. DeRoever, and Ilias Tachtsidis. Abroad: A machine learning based approach to detect broadband nirs artefacts. International Society of Oxygen Transport to Tissue, MLU, Halle, Germany, 2017a.

Joshua Russell-Buckland, Matthew Caldwell, and Ilias Tachtsidis. WeBCMD: A cross-platform interface for the BCMD modelling framework. *Wellcome Open Research*, 2:56, 7 2017b. ISSN 2398-502X. doi:10.12688/wellcomeopenres.12201.1. URL `https://wellcomeopenresearch.org/articles/2-56/v1`.

Joshua Russell-Buckland, Matthew Caldwell, and Ilias Tachtsidis. Webcmd: A modern, cross-platform interface to the bcmd modelling framework. International Society of Oxygen Transport to Tissue, MLU, Halle, Germany, 2017c.

Joshua Russell-Buckland, Cornelius Bauer, and Ilias Tachtsidis. Investigating the relationship between hypothermia and cerebral metabolism following hypoxic-ischaemic injury. International Society of Oxygen Transport to Tissue, Seoul University, Seoul, South Korea, 2018b.

Joshua Russell-Buckland, Matthew Caldwell, and Ilias Tachtsidis. Webcmd: A

modern, cross-platform interface to the bcmd modelling framework. CARnet, Oxford University, Oxford, United Kingdom, 2018c.

Joshua Russell-Buckland, Christopher P. Barnes, and Ilias Tachtsidis. A Bayesian framework for the analysis of systems biology models of the brain. *PLOS Computational Biology*, 15(4):e1006631, 4 2019b. ISSN 1553-7358. doi:10.1371/journal.pcbi.1006631. URL `http://dx.plos.org/10.1371/journal.pcbi.1006631`.

Joshua Russell-Buckland, Christopher P. Barnes, and Ilias Tachtsidis. A bayesian framework for the analysis of systems biology models of the brain. eResearch Symposium, UCL, London, United Kingdom, 2019c.

Mary Rutherford. *MRI of the Neonatal Brain*. 2002. URL `http://www.mrineonatalbrain.com/`.

A. Saltelli, S. Tarantola, and K. P.-S. Chan. A quantitative model-independent method for global sensitivity analysis of model output. *Technometrics*, 41 (1):39–56, 2 1999. ISSN 15372723. doi:10.1080/00401706.1999.10485594. URL `http://www.tandfonline.com/doi/abs/10.1080/00401706.1999.10485594`.

Andrea Saltelli, Stefano Tarantola, Francesca Campolongo, and Marco Ratto. *Sensitivity Analysis in Practice*. John Wiley & Sons, Ltd, Chichester, UK, 2 2002. ISBN 0470870931. doi:10.1002/0470870958.

Andrea Saltelli, Stefano Tarantola, Francesca Campolongo, and Marco Ratto. Sensitivity analysis in practice: a guide to assessing scientific models. *Chichester, England*, 2004.

F Scholkmann, S Spichtig, T Muehlemann, and M Wolf. How to detect and reduce movement artifacts in near-infrared imaging using moving standard deviation and spline interpolation. *Physiological Measurement*, 31(5):649–662, 5 2010. ISSN 0967-3334. doi:10.1088/0967-3334/31/5/004.

Felix Scholkmann, Stefan Kleiser, Andreas Jaakko Metz, Raphael Zimmermann, Juan Mata Pavia, Ursula Wolf, and Martin Wolf. A review on continuous wave functional near-infrared spectroscopy and imaging instrumentation and methodology. *NeuroImage*, 85:6–27, 1 2014. ISSN 10538119. doi:10.1016/j.neuroimage.2013.05.004. URL http://www.sciencedirect.com/science/article/pii/S105381191300494lhttp://linkinghub.elsevier.com/retrieve/pii/S1053811913004941.

Donald L. Schomer and Fernando H. Lopes da Silva, editors. *Niedermeyer's Electroencephalography*, volume 1. Oxford University Press, 11 2017. ISBN 9780190228484. doi:10.1093/med/9780190228484.001.0001. URL http://www.oxfordmedicine.com/view/10.1093/med/9780190228484.001.0001/med-9780190228484.

Jerome S. Schultz, Joe D. Goddard, and Shyam R. Suchdeo. Facilitated transport via carrier-mediated diffusion in membranes: Part I. Mechanistic aspects, experimental systems and characteristic regimes. *AIChE Journal*, 20(3):417–445, 5 1974. ISSN 0001-1541. doi:10.1002/aic.690200302. URL http://doi.wiley.com/10.1002/aic.690200302.

Seetha Shankaran, Abbot R. Laptook, Richard A. Ehrenkranz, Jon E. Tyson, Scott A. McDonald, Edward F. Donovan, Avroy A. Fanaroff, W. Kenneth Poole, Linda L. Wright, Rosemary D. Higgins, Neil N. Finer, Waldemar A. Carlo, Shahnaz Duara, William Oh, C. Michael Cotten, David K. Stevenson, Barbara J. Stoll, James A. Lemons, Ronnie Guillet, and Alan H. Jobe. Whole-Body Hypothermia for Neonates with Hypoxic–Ischemic Encephalopathy. *New England Journal of Medicine*, 353(15):1574–1584, 10 2005. ISSN 0028-4793. doi:10.1056/NEJMcps050929. URL http://www.nejm.org/doi/abs/10.1056/NEJMcps050929.

S. Sharma and M.J. Majsak. Brain Anatomy. In *Encyclopedia of the Neurological Sciences*, volume 1, pages 463–466. Elsevier, 2014. ISBN

9780123851574. doi:10.1016/B978-0-12-385157-4.01121-0. URL http://linkinghub.elsevier.com/retrieve/pii/B9780123851574011210.

Renée A Shellhaas, Brian J Thelen, Jayapalli R Bapuraj, Joseph W Burns, Aaron W Swenson, Mary K Christensen, Stephanie A Wiggins, and John D E Barks. Limited short-term prognostic utility of cerebral NIRS during neonatal therapeutic hypothermia. *Neurology*, 81(3):249–55, 7 2013. ISSN 1526-632X. doi:10.1212/WNL.0b013e31829bfe41. URL http://www.ncbi.nlm.nih.gov/pubmed/23771483http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3770165.

Ian A. Simpson, Anthony Carruthers, and Susan J. Vannucci. Supply and Demand in Cerebral Energy Metabolism: The Role of Nutrient Transporters. *Journal of Cerebral Blood Flow & Metabolism*, 27(11):1766–1791, 11 2007. ISSN 0271-678X. doi:10.1038/sj.jcbfm.9600521. URL http://journals.sagepub.com/doi/10.1038/sj.jcbfm.9600521.

S. A. Sisson, Y. Fan, and M. M. Tanaka. Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765, 2 2007. ISSN 0027-8424. doi:10.1073/pnas.0607208104. URL http://www.pnas.org/cgi/doi/10.1073/pnas.0607208104.

E. Somersalo, Y. Cheng, and D. Calvetti. The Metabolism of Neurons and Astrocytes Through Mathematical Models. *Annals of Biomedical Engineering*, 40(11):2328–2344, 11 2012. ISSN 0090-6964. doi:10.1007/s10439-012-0643-z. URL http://link.springer.com/10.1007/s10439-012-0643-z.

Mikael Sunnåker, Alberto Giovanni Busetto, Elina Numminen, Jukka Corander, Matthieu Foll, and Christophe Dessimoz. Approximate Bayesian Computation. *PLoS Computational Biology*, 9(1):e1002803, 1 2013. ISSN 1553-7358. doi:10.1371/journal.pcbi.1002803. URL http://dx.plos.org/10.1371/journal.pcbi.1002803.

L Symon. Experimental evidence for "intracerebral steal" following CO2 inhalation. *Scandinavian journal of clinical and laboratory investigation. Supplementum*, 102:XIII:A, 1968. ISSN 0085-591X. URL http://www.ncbi.nlm.nih.gov/pubmed/4974242.

Simon Tavare, David J Balding, +j R C Griffiths ', and Peter Donneuyst2. Inferring Coalescence Times From DNA Sequence Data. *Genetics1*, 145(2):505–518, 1997. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1207814/pdf/ge1452505.pdf.

Marianne Thoresen, Juliet Penrice, Ann Lorek, E B Cady, Marzena Wylezinska, V Kirkbride, C E Cooper, G C Brown, A D Edwards, J S Wyatt, and E O R Reynolds. Mild Hypothermia after Severe Transient Hypoxia-Ischemia Ameliorates Delayed Cerebral Energy Failure in the Newborn Piglet. *Pediatric Research*, 37(5):667–670, 5 1995. ISSN 0031-3998. doi:10.1203/00006450-199505000-00019. URL http://www.nature.com/doifinder/10.1203/00006450-199505000-00019.

Martin M. Tisdall, Ilias Tachtsidis, Terence S. Leung, Clare E. Elwell, and Martin Smith. Near-infrared spectroscopic quantification of changes in the concentration of oxidized cytochrome c oxidase in the healthy human brain during hypoxemia. *Journal of Biomedical Optics*, 12(2):024002, 2007. ISSN 10833668. doi:10.1117/1.2718541. URL http://biomedicaloptics.spiedigitallibrary.org/article.aspx?doi=10.1117/1.2718541.

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. *Journal of The Royal Society Interface*, 6(31):187–202, 2 2009. ISSN 1742-5689. doi:10.1098/rsif.2008.0172. URL http://rsif.royalsocietypublishing.org/cgi/doi/10.1098/rsif.2008.0172.

M. Ursino and M. Giulioni. Quantitative assessment of cerebral autoregulation from transcranial Doppler pulsatility: a computer simulation study. *Medical Engineer-*

ing & Physics, 25(8):655–666, 10 2003. ISSN 13504533. doi:10.1016/S1350-4533(02)00251-5. URL https://linkinghub.elsevier.com/retrieve/pii/S1350453302002515.

M. Ursino, C. A. Lodi, S. Rossi, and N. Stocchetti. Intracranial pressure dynamics in patients with acute brain damage. *Journal of Applied Physiology*, 82(4):1270–1282, 1997. ISSN 87507587. doi:10.1152/jappl.1997.82.4.1270.

M. Ursino, A. Ter Minassian, C. A. Lodi, and L. Beydon. Cerebral hemodynamics during arterial and CO2 pressure changes: In vivo prediction by a mathematical model. *American Journal of Physiology - Heart and Circulatory Physiology*, 279 (5 48-5), 2000. ISSN 03636135.

Mauro Ursino. A mathematical study of human intracranial hydrodynamics part 1—The cerebrospinal fluid pulse pressure. *Annals of Biomedical Engineering*, 16(4):379–401, 7 1988a. ISSN 0090-6964. doi:10.1007/BF02364625. URL http://link.springer.com/10.1007/BF02364625.

Mauro Ursino. A mathematical study of human intracranial hydrodynamics part 2—Simulation of clinical tests. *Annals of Biomedical Engineering*, 16(4):403–416, 7 1988b. ISSN 0090-6964. doi:10.1007/BF02364626. URL http://link.springer.com/10.1007/BF02364626.

Mauro Ursino. A Mathematical Model of Overall Cerebral Blood Flow Regulation in the Rat. *IEEE Transactions on Biomedical Engineering*, 38(8):795–807, 1991. ISSN 15582531. doi:10.1109/10.83592.

Mauro Ursino and Patrizia Di Giammarco. A mathematical model of the relationship between cerebral blood volume and intracranial pressure changes: The generation of plateau waves. *Annals of Biomedical Engineering*, 19(1): 15–42, 1 1991. ISSN 0090-6964. doi:10.1007/BF02368459. URL http://link.springer.com/10.1007/BF02368459.

Mauro Ursino and Carlo Alberto Lodi. A simple mathematical model of the interaction between intracranial pressure and cerebral hemodynamics.

Journal of Applied Physiology, 82(4):1256–1269, 1997. ISSN 87507587. doi:10.1152/jappl.1997.82.4.1256.

Mauro Ursino and Carlo Alberto Lodi. Interaction among autoregulation, CO2 reactivity, and intracranial pressure: a mathematical model. *American Journal of Physiology-Heart and Circulatory Physiology*, 274(5):H1715–H1728, 5 1998. ISSN 0363-6135. doi:10.1152/ajpheart.1998.274.5.H1715. URL `http://ajpheart.physiology.org/content/274/5/H1715.abstracthttp://www.physiology.org/doi/10.1152/ajpheart.1998.274.5.H1715`.

Mauro Ursino, Maurizio Iezzi, and Nino Stocchetti. Intracranial Pressure Dynamics in Patients with Acute Brain Damage: A Critical Analysis with the Aid of a Mathematical Model. *IEEE Transactions on Biomedical Engineering*, 42(6): 529–540, 1995. ISSN 15582531. doi:10.1109/10.387192.

Mauro Ursino, Marco Giulioni, and Carlo Alberto Lodi. Relationships among cerebral perfusion pressure, autoregulation, and transcranial Doppler waveform: A modeling study. *Journal of Neurosurgery*, 89(2):255–266, 1998. ISSN 00223085. doi:10.3171/jns.1998.89.2.0255.

Rens Van de Schoot, David Kaplan, Jaap Denissen, Jens B. Asendorpf, Franz J. Neyer, and Marcel A.G. van Aken. A Gentle Introduction to Bayesian Analysis: Applications to Developmental Research. *Child Development*, 85(3):842–860, 5 2014. ISSN 14678624. doi:10.1111/cdev.12169. URL `http://doi.wiley.com/10.1111/cdev.12169`.

M S van der Knaap, J van der Grond, P C van Rijen, J A Faber, J Valk, and K Willemse. Age-dependent changes in localized proton and phosphorus MR spectroscopy of the brain. *Radiology*, 176(2):509–515, 8 1990. ISSN 0033-8419. doi:10.1148/radiology.176.2.2164237. URL `http://pubs.rsna.org/doi/10.1148/radiology.176.2.2164237`.

Stéfan van der Walt, S Chris Colbert, and Gaël Varoquaux. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science & Engi-*

neering, 13(2):22–30, 3 2011. ISSN 1521-9615. doi:10.1109/MCSE.2011.37. URL http://ieeexplore.ieee.org/document/5725236/.

B.J.M. van Kooij, J. Hendrikse, M.J.N.L. Benders, L.S. de Vries, and F. Groenendaal. Anatomy of the Circle of Willis and Blood Flow in the Brain-Feeding Vasculature in Prematurely Born Infants. *Neonatology*, 97(3):235–241, 2010. ISSN 1661-7819. doi:10.1159/000253754. URL https://www.karger.com/Article/FullText/253754.

Robert C. Vannucci. Hpoxic-Ischemic Encephelopathy. *American Journal of Perinatology*, Volume 17(Number 03):113–120, 2000. ISSN 07351631. doi:10.1055/s-2000-9293. URL http://www.thieme-connect.de/DOI/DOI?10.1055/s-2000-9293.

Larisa Vatov, Ziv Kizner, Eytan Ruppin, Sigal Meilin, Tamar Manor, and Avraham Mayevsky. Modeling Brain Energy Metabolism and Function: A Multiparametric Monitoring Approach. *Bulletin of Mathematical Biology*, 68(2):275–291, 2 2006. ISSN 0092-8240. doi:10.1007/s11538-005-9008-1. URL http://link.springer.com/10.1007/s11538-005-9008-1.

A. I.F. Vaz and L. N. Vicente. PSwarm: a hybrid solver for linearly constrained global derivative-free optimization. *Optimization Methods and Software*, 24(4-5):669–685, 10 2009. ISSN 1055-6788. doi:10.1080/10556780902909948. URL http://www.tandfonline.com/doi/abs/10.1080/10556780902909948.

A. Ismael F. Vaz and Luís N. Vicente. A particle swarm pattern search method for bound constrained global optimization. *Journal of Global Optimization*, 39(2):197–219, 9 2007. ISSN 09255001. doi:10.1007/s10898-007-9133-5. URL http://link.springer.com/10.1007/s10898-007-9133-5.

D J Venzon and S H Moolgavkar. A method for computing profile-likelihood-based confidence intervals. *Applied statistics*, pages 87–94, 1988.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser,

Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, and Paul van Mulbregt. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, 3 2020. ISSN 1548-7091. doi:10.1038/s41592-019-0686-2. URL `http://www.nature.com/articles/s41592-019-0686-2`.

Kevin R. Ward and Donald M. Yealy. End-tidal Carbon Dioxide Monitoring in Emergency Medicine, Part 1: Basic Principles. *Academic Emergency Medicine*, 5 (6):628–636, 6 1998. ISSN 10696563. doi:10.1111/j.1553-2712.1998.tb02473.x. URL `http://doi.wiley.com/10.1111/j.1553-2712.1998.tb02473.x`.

Michael Waskom, Olga Botvinnik, Drew O'Kane, Paul Hobson, Saulius Lukauskas, David C Gemperline, Tom Augspurger, Yaroslav Halchenko, John B. Cole, Jordi Warmenhoven, Julian De Ruiter, Stephan Hoyer, Jake Vanderplas, Santi Villalba, Gero Kunter, Eric Quintero, Pete Bachant, Marcel Martin, Kyle Meyer, Alistair Miles, Yoav Ram, Cameron Pye, Tal Yarkoni, Mike Lee Williams, Constantine Evans, Clark Fitzgerald, Brian, Chris Fonnesbeck, Antony Lee, and Adel Qalieh. Mwaskom/Seaborn: V0.8.0 (July 2017). *doi.org*, 2017. doi:10.5281/zenodo.824567. URL `https://zenodo.org/record/824567`.

Michael Waskom, Olga Botvinnik, Drew O'Kane, Paul Hobson, Joel Ostblom, Saulius Lukauskas, David C Gemperline, Tom Augspurger, Yaroslav Halchenko, John B Cole, Jordi Warmenhoven, Julian de Ruiter, Cameron Pye, Stephan Hoyer, Jake Vanderplas, Santi Villalba, Gero Kunter, Eric Quintero, Pete Bachant, Marcel Martin, Kyle Meyer, Alistair Miles, Yoav Ram, Thomas Brunner, Tal Yarkoni, Mike Lee Williams, Constantine Evans, Clark Fitzgerald, Brian, and Adel Qalieh. mwaskom/seaborn: v0.9.0 (July 2018), 7 2018. URL `https://doi.org/10.5281/zenodo.1313201`.

D. J. Wilkinson. Bayesian methods in bioinformatics and computational systems biology. *Briefings in Bioinformatics*, 8(2):109–116, 12 2006. ISSN 1467-5463. doi:10.1093/bib/bbm007. URL https://academic.oup.com/bib/article-lookup/doi/10.1093/bib/bbm007.

Uri Wolf, Mark J. Rapoport, and Tom A. Schweizer. Evaluating the Affective Component of the Cerebellar Cognitive Affective Syndrome. *The Journal of Neuropsychiatry and Clinical Neurosciences*, 21(3):245–253, 7 2014. ISSN 0895-0172. doi:10.1176/jnp.2009.21.3.245. URL http://psychiatryonline.org/doi/abs/10.1176/jnp.2009.21.3.245.

Mae L. Woods and Chris P. Barnes. Mechanistic Modelling and Bayesian Inference Elucidates the Variable Dynamics of Double-Strand Break Repair. *PLOS Computational Biology*, 12(10):e1005131, 10 2016. ISSN 1553-7358. doi:10.1371/journal.pcbi.1005131. URL http://dx.plos.org/10.1371/journal.pcbi.1005131.

Mae L. Woods, Miriam Leon, Ruben Perez-Carrasco, and Chris P. Barnes. A Statistical Approach Reveals Designs for the Most Robust Stochastic Gene Oscillators. *ACS Synthetic Biology*, 5(6):459–470, 6 2016. ISSN 2161-5063. doi:10.1021/acssynbio.5b00179. URL http://pubs.acs.org/doi/10.1021/acssynbio.5b00179.

Susan Wray, Mark Cope, David T. Delpy, John S. Wyatt, and E.Osmund R. Reynolds. Characterization of the near infrared absorption spectra of cytochrome aa3 and haemoglobin for the non-invasive monitoring of cerebral oxygenation. *Biochimica et Biophysica Acta (BBA) - Bioenergetics*, 933(1):184–192, 3 1988. ISSN 00052728. doi:10.1016/0005-2728(88)90069-2. URL https://linkinghub.elsevier.com/retrieve/pii/0005272888900692.

Duan Xu and Daniel Vigneron. Magnetic resonance spectroscopy imaging of the newborn brain–a technical review. *Seminars in perinatology*, 34 (1):20–7, 2 2010. ISSN 1558-075X. doi:10.1053/j.semperi.2009.10.003.

URL http://www.ncbi.nlm.nih.gov/pubmed/20109969http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2842012.

Jin Yang, John W. Clark, Robert M. Bryan, and Claudia Robertson. The myogenic response in isolated rat cerebrovascular arteries: smooth muscle cell model. *Medical Engineering & Physics*, 25(8):691–709, 10 2003a. ISSN 13504533. doi:10.1016/S1350-4533(03)00100-0. URL https://linkinghub.elsevier.com/retrieve/pii/S1350453303001000.

Jin Yang, John W. Clark, Robert M. Bryan, and Claudia S. Robertson. The myogenic response in isolated rat cerebrovascular arteries: vessel model. *Medical Engineering & Physics*, 25(8):711–717, 10 2003b. ISSN 13504533. doi:10.1016/S1350-4533(03)00101-2. URL https://linkinghub.elsevier.com/retrieve/pii/S1350453303001012.