



UNIVERSITÄT ZU LÜBECK

From the Institute for Electrical Engineering in Medicine
of the University of Lübeck
Director: Prof. Dr. Philipp Rostalski

Non-invasive Estimation of Respiratory Effort
– **From Surface Electromyography to Gaussian Processes** –

Dissertation

for Fulfillment of Requirements for the Doctoral Degree
of the University of Lübeck

from the Department of Computer Sciences and Technical Engineering

Submitted by
Jan Graßhoff
from Lübeck

Lübeck, 2024

First referee: Prof. Dr.-Ing. Philipp Rostalski

Second referee: Prof. Dr. rer. nat. habil. Jonas Obleser

Date of oral examination: November 21st, 2024

Approved for printing. Lübeck, May 8th, 2025

Abstract

Mechanical ventilation is a pivotal intervention for replacing or supporting spontaneous breathing in critically ill patients. Tailoring mechanical ventilation to the needs of an individual patient is challenging and requires a balance between maintaining adequate gas exchange and preventing lung injury as well as loss of respiratory muscle strength. Recently, many researchers have argued that the patient's own spontaneous breathing effort is an essential factor for improving mechanical ventilation therapy and that it should be monitored more closely. As yet, however, there is no established non-invasive technique to measure this quantity continuously at the bedside. Therefore, the thesis explores respiratory surface electromyography as an emerging non-invasive method for recording the electrical activity of respiratory muscles. This technique is easy to implement in clinical practice and only requires placement of gel electrodes on the thorax of the patient.

Based on these measurements, the thesis adopts a model-based approach to derive a completely non-invasive estimate of the respiratory effort. A model of the respiratory system is used to integrate all available measurements, including electromyography signals and pneumatic data from the ventilator. Latent parameters and states of the model can then be estimated using statistical inference. By combining measurements from different sensors through sensor fusion, noise from individual signals is reduced and a more precise estimate of the respiratory effort can be constructed. The thesis proposes several robust techniques for solving this estimation problem.

An important contribution is an investigation of Gaussian processes as a means to formulate prior physiological knowledge about respiratory signals. By incorporating statistical properties of the data, such as their smoothness or quasi-periodicity, estimates can be regularized towards physiologically plausible solutions. The use of Gaussian processes gives rise to a Bayesian framework in which the prior distribution is conditioned on the measurements to derive a posterior distribution. Due to the high computational cost involved in solving this type of problem, the thesis also investigates scalable Gaussian process algorithms by using various structure exploitation techniques. These considerations also prove to be very useful in other biomedical domains, which is demonstrated in the context of spatio-temporal source separation.

The methods are tested and validated on clinical data, including a large dataset of patients undergoing assisted mechanical ventilation. The results confirm the great potential of surface electromyography and statistical inference for optimizing mechanical ventilation therapy.

Zusammenfassung

Die maschinelle Beatmung ist eine lebenswichtige Intervention, um die Spontanatmung bei Intensivpatienten zu ersetzen oder zu unterstützen. Die individuelle Anpassung der Beatmungstherapie an einzelne Patienten ist anspruchsvoll: sie erfordert eine Abwägung zwischen der Erhaltung des Gasaustausches und der Vermeidung von Lungenschäden sowie eines Verlustes der Atemmuskulatur. Nach Auffassung vieler Wissenschaftler ist die Atemanstrengung von Patienten ein entscheidender Faktor für die Verbesserung der Beatmungstherapie und sollte daher sorgfältig überwacht werden. Derzeit existiert aber noch keine etablierte nicht-invasive Technologie, mit der dieser Wert kontinuierlich am Krankenbett erfasst werden könnte. In dieser Arbeit wird daher die Oberflächen-Elektromyographie als eine neue nicht-invasive Technik zur Aufzeichnung der elektrischen Aktivität der Atemmuskulatur betrachtet. Diese Messung lässt sich in der Klinik leicht umsetzen und bedarf nur der Positionierung weniger Gel-Elektroden auf dem Thorax des Patienten.

In dieser Arbeit wird ein modellbasierter Ansatz verfolgt, um eine vollständig nicht-invasive Schätzung der Atemarbeit zu ermöglichen. Die verfügbaren Messungen, einschließlich elektromyographischer Signale und pneumatischer Daten vom Beatmungsgerät, werden über ein Modell des Atmungssystems kombiniert. Latente Parameter und Zustände des Modells können somit über statistische Inferenz geschätzt werden. Durch die Kombination der verschiedenen Messdaten mittels Sensorfusion wird das Rauschen einzelner Signale reduziert und eine genauere Schätzung der Atemarbeit ermöglicht. Im Rahmen der Arbeit werden mehrere robuste Methoden zur Lösung dieses Schätzproblems vorgeschlagen.

Ein wichtiger Aspekt der Arbeit ist die Untersuchung von Gaußprozessen zur Beschreibung des physiologischen Vorwissens über die Atmungssignale. Durch die Einbeziehung statistischer Eigenschaften der Daten, wie z.B. Glattheit oder Quasiperiodizität, können die Schätzungen zugunsten physiologisch plausibler Lösungen regularisiert werden. Die Verwendung von Gaußprozessen führt dabei zu einem Bayes'schen Verfahren, bei welchem aus der Prior-Verteilung durch Konditionierung auf die Messdaten eine Posterior-Verteilung berechnet wird. Da dies mit hohen Rechenkosten einhergeht, werden skalierbare Gaußprozess-Algorithmen für die Ausnutzung von vorhandener Struktur untersucht. Diese Überlegungen erweisen sich auch in anderen biomedizinischen Bereichen als sehr nützlich, was am Beispiel der räumlich-zeitlichen Quellentrennung demonstriert wird.

Die vorgeschlagenen Methoden werden auf klinischen Daten getestet und validiert, unter anderem auf einem großen Datensatz assistiert beatmeter Patienten. Die Ergebnisse bestätigen das große Potenzial der Oberflächen-Elektromyographie sowie der Verwendung statistischer Inferenz für eine Verbesserung der Beatmungstherapie.

Acknowledgments

The research presented in this thesis was carried out at two institutions: the Institute for Electrical Engineering in Medicine (IME), and the Fraunhofer Research Institution for Individualized and Cell-Based Medical Engineering (IMTE). Writing this thesis has been a rewarding (and long) journey, during which I had the privilege of working with many talented people.

I am deeply grateful to all my collaborators and colleagues for their support. In particular, I wish to thank Philipp Rostalski for his mentorship and supervision. I got to know Philipp during my time as a graduate student, when he was teaching control engineering — one of my favorite courses, and the one that ultimately led me to join his team. The course *Advanced Methods of Control Engineering* had only four students, three of whom would later join the IME. In that course, we would gather around a table while Philipp wrote on a sheet of paper, with us looking over his shoulder. We learned everything there is to know about state-space models and model predictive control. The atmosphere was very engaging and ignited my passion for dynamic systems and state estimation.

At both the IME and IMTE, Philipp created an inspiring atmosphere for research. The early days at the IME were a time when we opened new areas of research with a small team, including factor graphs and Gaussian processes. It was an exciting time. I am grateful to Philipp for his guidance and optimism.

I was very fortunate to collaborate with the Dräger research unit. Dräger did not only provide financial support for my research but also offered a close and fruitful partnership. I have had the great pleasure of working with Marcus Eger and Thomas Handzsj continuously since writing my master's thesis at the Dräger research unit. Since then, we have been thinking about respiratory surface EMG together and are determined to bring it into clinical application. I believe such a trusting and close partnership between industry and academia is rare.

This work would not have been possible without the support from clinical partners. I am especially grateful for the opportunity to collaborate with Stephan Waltersbacher and Niklas Carbon. As physicians, they were willing to engage in technical discussions and answer the countless questions I had. I am also grateful to Eline Oppersma, a technical physician, for her fruitful collaboration and shared vision for surface EMG and respiratory monitoring.

At the IME, I had the privilege of working alongside wonderful colleagues: Christian Herzog, Ayla Nawaz, Julia Sauer, Eike Petersen, Maryam Nezami, Hossam Abbas, Felix Vollmer, and Georg Männel made the time very special. We shared many cherished memories, including retreats, conferences, and late nights at the institute.

Felix and Georg were also among the four students in that one memorable lecture. We started at the same time, and I am grateful that we have walked this path together. With Julia, I worked extensively on both clinical and technical aspects related to surface EMG. I was always impressed by her diligence and patience.

Eike is one of the people I could turn to with any kind of question — not just technical, but also personal — and he always had useful wisdom at hand. It was inspiring to work with him. I could not have asked for a better scientific partner.

Ayla and I shared the same passion for probabilistic graphical models, which we never got tired of discussing. Besides technical conversations, Ayla broadened my perspectives on the world by introducing new angles that challenged my assumptions. Above all, I am grateful for her support when facing roadblocks and difficulties. During the COVID pandemic, Ayla, Eike, and I worked together virtually and kept the feeling of an office alive through (literally) thousands of chat messages.

Though we worked in different research fields, Maryam and I easily connected over science, writing, work, and movies. I truly appreciated our many conversations and exchanges.

At the IMTE, I found a new home in the data science team. I was fortunate to work closely with Lukas Boudnik, Andra Oltmann, Marlin Siebert, and Tobias Barth. Together, we conducted physiological studies, wrote papers, and submitted countless research proposals. It is a fantastic environment, and I enjoy working with them every day.

This journey would not have been possible without my family. My siblings, Hanna and Lars, supported me through the most trying times. I am especially grateful to my parents for laying the early foundations for this work: they fostered my enthusiasm for science and curiosity for open questions. They also taught me the value of working accurately and rigorously.

Finally, and most importantly, I want to thank Chrissie, my partner, for her unwavering support over the years, for always believing in me, and for her patience with the long working hours.

— Jan Graßhoff, May 2025

Contents

Abstract	i
Zusammenfassung	iii
Acknowledgments	v
Contents	ix
1 Introduction	1
1.1 Background and objectives	1
1.2 Thesis outline and contributions	3
2 Fundamentals	7
2.1 Parametric and nonparametric regression	7
2.1.1 Parametric Gaussian regression	7
2.1.2 Robust regression	10
2.1.3 Nonparametric Gaussian regression	12
2.1.4 Linear stochastic differential equations	19
2.2 Respiratory physiology and mechanical ventilation	23
2.2.1 Physiology	23
2.2.2 Respiratory failure and assisted ventilation	25
2.2.3 Lung mechanics modeling	27
2.2.4 Measurement of respiratory muscle effort	29
3 The EMG-Pressure Relationship of Respiratory Muscles	33
3.1 Introduction to respiratory electromyography	33
3.1.1 Physiology	34
3.1.2 Measurement	34
3.1.3 EMG-pressure relationship	36
3.2 Data analysis	38
3.2.1 Study data	38
3.2.2 Signal processing	40
3.2.3 Quantitative analysis	42
3.3 Results	46
3.4 Discussion	50

4	Model-based Inspiratory Effort Estimation using Surface EMG	55
4.1	Introduction and previous work	55
4.2	Estimation framework	58
4.2.1	Occlusion-based effort estimation	58
4.2.2	Model-based effort estimation	60
4.2.3	Combined approach	66
4.3	Validation & performance quantification	66
4.3.1	Performance metrics	67
4.4	Results	67
4.5	Discussion	73
5	A Gaussian Process Model for Respiratory Muscle Pressure: Sensor Fusion, Inference and Learning	79
5.1	Introduction and previous work	80
5.2	Bayesian framework	82
5.2.1	Gaussian process model for P_{mus}	82
5.2.2	Probabilistic sensor fusion model	85
5.2.3	Inference of P_{mus}	87
5.2.4	Equivalent state-space formulation	89
5.2.5	Relation to other models	92
5.3	Data analysis	93
5.3.1	Model learning	95
5.3.2	Performance metrics	97
5.4	Results	99
5.5	Discussion	102
6	Scalable Gaussian Process Regression for Spatio-Temporal Biomedical Data	107
6.1	Introduction and previous work	108
6.1.1	Equivalent state-space representation	108
6.1.2	Toeplitz and Kronecker methods	108
6.2	Scalable separation of GPs with a non-stationary phase	111
6.2.1	Model definition	111
6.2.2	Structure exploitation	113
6.2.3	Spatio-temporal kernels	116
6.2.4	Relation to other methods	117
6.3	Case study: fetal ECG extraction	118
6.4	Case study: reconstruction of periodic processes in electrical impedance tomography (EIT)	121
6.4.1	Temporal EIT solver	122
6.4.2	Results on simulation and clinical data	126
6.5	Discussion	129

7 Conclusion	133
7.1 Summary and key insights	133
7.2 Outlook	135
Appendix	139
A.1 Study data	139
A.2 Kalman filtering and smoothing algorithm	140
A.3 EIT forward equations	141
References	143
List of Figures	168
List of Tables	169
List of Publications	171

Chapter 1

Introduction

Mechanical ventilation is a life-saving intervention allowing to support or fully replace spontaneous breathing of critically ill patients. Recently, this technique has received widespread attention due to its pivotal role in the management of the COVID-19 pandemic. Adjusting ventilation to the individual needs of patients is a challenging task and requires balancing conflicting objectives: adequate gas exchange must be ensured, but ventilation should not worsen lung injury while maintaining respiratory muscle strength. This trade-off is complicated by the fact that many variables related to the respiratory system can only be measured indirectly or via complex, sometimes invasive, techniques [134, 156, 206]. However, even after the acute phase of the pandemic, intensive care staff continue to be heavily burdened by time-consuming procedures and shortage of trained personnel [230], limiting the time available for assessing individual patients. Thus, there is a huge potential to improve quality of care in mechanically ventilated patients by incorporating non-invasive sensors, integrating available measurements, and inferring latent states and parameters of the respiratory system. These improvements can only be achieved through technological innovation, which is the main motivation for the kind of research undertaken in this thesis.

1.1 Background and objectives

In recent decades, many advances have been made in the treatment of critically ill ventilated patients, including new approaches to lung-protective ventilation and lung imaging. One factor that is often overlooked and cannot easily be quantified is the role of the spontaneous breathing activity during assisted mechanical ventilation. A growing body of research has recently revealed the importance of closely monitoring the respiratory muscles and adjusting ventilator settings accordingly [31, 68, 244]. Inadequate respiratory muscle effort often goes unnoticed in clinical practice, yet, it is an important mediator of lung and diaphragm injury, affecting long-term outcomes [53, 66, 103, 244, 245]. Vigorous inspiratory effort increases transpulmonary pressure and tidal volume, which exerts mechanical stress on the lungs and may lead to self-inflicted lung injuries [244]. At the same time, both inactivity and too heavy loading of the diaphragm lead to muscle weakness and injury, ultimately inducing diaphragm

dysfunction [197]. In regard of these mechanisms, researchers have advocated to improve lung and diaphragm health by keeping muscle activity within safe limits [31, 67, 68, 93]. To this end, it will be essential to develop techniques enabling accurate *quantification of respiratory muscle effort* in mechanically ventilated patients at the bedside, which is the main objective of this thesis.

The current standard for estimating respiratory effort relies on measuring esophageal pressure (P_{es}) with an invasive, nasogastric catheter [10, 134]. Despite potential benefits of this measurement, P_{es} is not yet routinely employed in many clinics because it is considered technically cumbersome and difficult to interpret [243]. Establishing a reliable, non-invasive alternative to esophageal pressure could significantly improve patient care and has been the subject of intense research. One potential approach is *electromyography* (EMG), a technique for recording the electrical activity of a muscle. The first main focus of this thesis is on *respiratory surface EMG* (sEMG), a straightforward and completely non-invasive measurement technique, which only requires placement of a few electrodes on the thorax of the ventilated patient [128]. It is easy to implement in clinical practice but usually has a low signal-to-noise ratio and is prone to interference by various types of electrical artifacts. To this date, only a limited number of studies have investigated sEMG as a means of estimating respiratory effort. Previous publications have emphasized its potential but also expressed concerns with regard to the significant level of noise present in the signal [27, 226]. A particularly important challenge in making sEMG practical is to establish a meaningful scale to the signal — that is, to convert the measured electrical activity into the force or pressure exerted by the muscle [29, 161].

The key idea adopted in this thesis is to pose the estimation of respiratory effort as a *model-based inference problem*. To this end, models of the respiratory system will be used, combining all available measurements, namely sEMG signals and pneumatic data from the ventilator. Latent parameters or states of these models can then be inferred completely non-invasively during normal assisted breathing via statistical — in some cases *Bayesian* — estimators. This thesis will investigate robust and efficient techniques for estimating parameters of the respiratory system. By incorporating all available data, i.e., conducting *sensor fusion*, noise from the individual measurements is reduced and a clean estimate of respiratory effort can be constructed.

An important research question is how prior knowledge about the statistics of the data can be incorporated into the model and exploited during inference. One possible approach, and the second main focus area of this thesis, is the use of *Gaussian processes*, a type of *Bayesian nonparametric* model, which allows to formulate key characteristics of the data like smoothness and quasi-periodicity by means of kernel functions. Gaussian processes have been proposed as an explainable approach to solving regression tasks, including, among others, for spatio-temporal statistics and machine learning. The usage of Gaussian processes raises another highly relevant question: How can inference in such models be solved in a scalable way? Two approaches will be investi-

gated: a reformulation in terms of state-space models, enabling to solve the inference sequentially, and the exploitation of existing structure within the kernel matrix. These methods give rise to a general framework for solving regression problems, which will be demonstrated on a number of highly relevant, biomedical datasets.

1.2 Thesis outline and contributions

Figure 1.1 provides an overview of the thesis' main themes and of the methodological chapters 3–6.

Chapter 2 introduces fundamentals that are essential for understanding the later parts of the thesis. It comprises an introduction to different probabilistic inference methods, in particular parametric and nonparametric Gaussian regression, as well as a brief overview of respiratory physiology, mechanical ventilation and esophageal pressure measurement.

Chapter 3 investigates the relation between sEMG and the total pressure generated by respiratory muscles, denoted by P_{mus} , a key measure of respiratory effort. The chapter provides a brief introduction to the relevant literature and then continues with a statistical analysis of a large dataset consisting of 43 patients undergoing bronchoscopy with mechanical ventilation. The main contribution of the chapter is a comprehensive evaluation of the P_{mus} -sEMG relation with regard to correlation, linearity and neuromechanical coupling. A novel technique for automatically selecting the sEMG channel with higher SNR is proposed and analyzed. Finally, the chapter also investigates the validity of a previously proposed occlusion technique, which is a simple maneuver allowing to determine a scaling factor between sEMG and muscle pressure.

Based on these findings, the subsequent chapter 4 investigates a *model-based* approach that incorporates all available data (sEMG and pneumatic signals) to reconstruct the muscle pressure P_{mus} during assisted mechanical ventilation. This chapter contains a quantitative comparison of estimation algorithms on the aforementioned bronchoscopy dataset using a range of different performance measures. The main contribution of the chapter is the proposal of a novel estimation algorithm that is highly robust against typical confounders of ventilatory data such as phase delays, outliers and artifacts. To this end, the chapter analyzes the merits of using *robust regression* instead of *ordinary least squares regression* for estimating parameters of the respiratory system.

Chapter 5 explores *Gaussian processes* as a means of modeling key characteristics of P_{mus} signals, namely smoothness, periodicity, shape of spontaneous efforts, and variance. The usage of Gaussian process priors gives rise to a Bayesian *sensor fusion* framework. The inference with respect to P_{mus} can be solved in an end-to-end fashion by conditioning the prior distribution on the available data. Two numerical solutions to the inference problem are proposed: a closed-form solution that uses the full

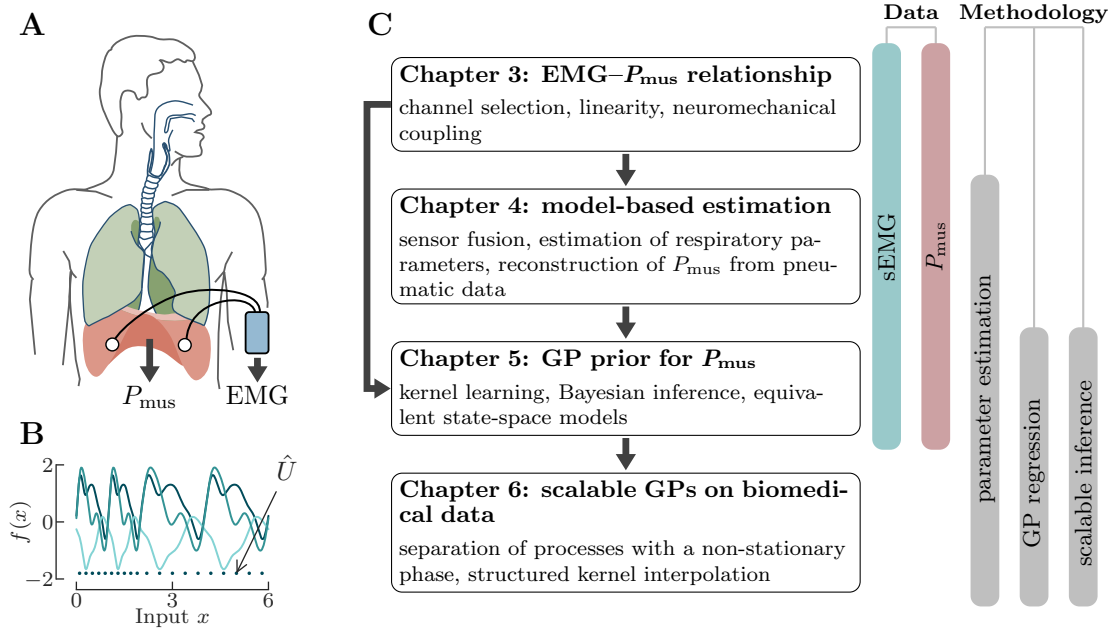


Fig. 1.1: Graphical overview of the thesis’ content and methodological chapters. **A:** An illustration of the respiratory system including the airways, lungs and the diaphragm. The key measures of respiratory effort used in chapters 3–5 are the pressure exerted by the respiratory muscles (P_{mus}) and the respiratory surface electromyogram (EMG or sEMG). **B:** A distribution of functions sampled from a strictly periodic Gaussian process (GP) prior with a non-stationary phase. The set of inducing points denoted by \hat{U} is employed in chapter 6, enabling exploitation of structure in the kernel matrix. **C:** Organization and structure of chapters 3–6 as well as an overview of involved data and methodologies.

kernel matrix, and an approximate, sequential solution in terms of equivalent state-space models. An important contribution of the chapter is a quantitative comparison of different kernel functions and their effect on the posterior distribution. The results demonstrate the merits of incorporating prior physiological knowledge about the respiratory effort in terms of a stochastic process.

Chapter 6 is concerned with questions raised in the previous chapter regarding the scalability of Gaussian processes on (spatio-temporal) biomedical data. A relevant practical problem encountered in many biomedical applications is the separation of additive mixtures involving physiological processes with a non-stationary phase. The non-stationarity may stem from natural variations of, e.g., heart/respiratory rates or gaiting velocity. The approach taken in this chapter is to exploit existing structure within the kernel matrix by means of *structured kernel interpolation*, enabling highly efficient matrix operations. The method is demonstrated on two case studies (both unrelated to the previous problem of estimating respiratory effort). One of these case

studies deals with the separation of non-stationary perfusion and ventilation signals in electrical impedance tomography, a highly relevant problem in itself, to which a first end-to-end solution is proposed.

The final chapter 7 concludes the thesis by summarizing the key insights and providing recommendations for future research directions.

The results of this thesis have been published in a series of conference papers and journals. The author of the thesis has first-authored two peer-reviewed journal publications [76, 78] and six peer-reviewed conference publications [71–75, 79] on the topic of the thesis. Moreover, he has co-authored four patent applications [55, 56, 107, 117], eight peer-reviewed journal publications [90, 105, 112, 116, 148, 162, 194, 203] and 13 peer-reviewed conference publications [1, 35, 36, 39, 82, 91, 115, 118, 130, 161, 165, 218, 231] on topics closely related to the thesis. Parts of this thesis have not yet been published elsewhere, including the Gaussian process framework of chapter 5 and the end-to-end perfusion/ventilation separation on electrical impedance tomography data proposed in chapter 6.

Chapter 2

Fundamentals

The purpose of this chapter is to provide the background necessary to follow along with the methodological contributions of the thesis. The first part of the chapter, section 2.1, introduces relevant methods from the field of probabilistic inference with a focus on parametric and nonparametric regression. These methods form the basis for estimators developed in later chapters. The second part of the chapter, section 2.2, comprises a brief introduction to respiratory physiology and mechanical ventilation. The aim is to provide the fundamentals of respiratory system functions, particularly respiratory mechanics, and challenges related to patient-ventilator interaction such as the invasive measurement of respiratory effort.

2.1 Parametric and nonparametric regression

The question at the heart of this thesis is how latent parameters and states, e.g., of the respiratory system, can be inferred from available, noisy measurements. A suitable candidate for solving such problems is Bayesian inference as it provides a very natural way of dealing with uncertainty in measured data and estimates. Most of the estimation problems treated in this thesis can be posed as regression analyses. The following sections will therefore focus on inference of regression functions, i.e., on learning mappings between training input-output pairs denoted by x_i and y_i , respectively. In the Bayesian setting, this involves placing a prior on a function (or on parameters of the function) and then calculating the function posterior by conditioning on the available data. The following sections will first introduce the standard parametric Gaussian regression model, followed by several extensions, including robust regression models, Gaussian processes and stochastic differential equations.

2.1.1 Parametric Gaussian regression

First, consider the problem of learning functions defined by the linear model $f(x) = \phi(x)^\top \beta$, where $\phi(x) : \mathbb{R}^D \rightarrow \mathbb{R}^F$ maps from the D -dimensional input space to an F -

dimensional feature space. The features are the model's *regressors*, and $\beta \in \mathbb{R}^F$ is the corresponding parameter vector. The goal is to infer β from noisy measurements

$$y_k = f(x_k) + \epsilon_k = \phi(x_k)^\top \beta + \epsilon_k, \quad (2.1)$$

with $k = 1, \dots, N$, and predict the value of the function $f(x_*)$ at a given test point x_* . The variable ϵ_k is intended to account for deviations from the linear relationship. The capacity of the model can be adjusted by changing the feature vector: to give an example, polynomial regression on a one-dimensional domain ($x \in \mathbb{R}$) corresponds to using $\phi(x) = [1, x, x^2, \dots]^\top$. In the following, a Bayesian paradigm is adopted, thus, prior knowledge about the parameters is encoded via $\beta \sim \mathcal{N}(\mu_\beta, \Sigma_\beta)$. For sake of brevity, only zero-mean parameter priors $\mu_\beta = 0$ will be considered, but the extension to $\mu_\beta \neq 0$ is straightforward. Here, it will be assumed that the measurement noise is an identical and independently distributed zero-mean Gaussian variable $\epsilon_k \sim \mathcal{N}(0, \sigma_n^2)$ with variance σ_n^2 . A more general case of non-Gaussian noise will be treated in section 2.1.2. The likelihood function factorizes over the data points and can thus be written as

$$\mathcal{L}(\beta|y) := p(y|\beta) = \prod_{k=1}^N \mathcal{N}(y_k; f(x_k), \sigma_n^2) = \mathcal{N}(y; \Phi^\top \beta, \sigma_n^2 I) \quad (2.2)$$

where $y = [y_1, \dots, y_N]^\top$ denotes the vector of measurements and $\Phi = [\phi(x_1), \dots, \phi(x_N)]$ is the design matrix that contains feature vectors of all data points in its columns. As both the prior and the likelihood have Gaussian distributions (i.e., the prior is conjugate), the posterior

$$p(\beta|y) \propto p(y|\beta)p(\beta) \quad (2.3)$$

can be calculated in closed form and is also Gaussian

$$p(\beta|y) = \mathcal{N}(\beta; \mu_{\beta,\text{post}}, \Sigma_{\beta,\text{post}}) \quad (2.4)$$

with

$$\mu_{\beta,\text{post}} = \sigma_n^{-2} \Sigma_{\beta,\text{post}} \Phi y \quad (2.5)$$

$$\Sigma_{\beta,\text{post}} = \left(\Sigma_\beta^{-1} + \sigma_n^{-2} \Phi \Phi^\top \right)^{-1}, \quad (2.6)$$

which can be derived simply using Gaussian conditioning rules, refer to [170, section 2.1.1], [89, section 4.1]. The predictive distribution over the function value $f_* := f(x_*)$ can then be calculated by propagating the parameter uncertainty forward over the linear model, which yields

$$p(f_*|y) = \mathcal{N}(f_*; \mathbf{m}_{f,\text{post}}(x_*), \sigma_{f,\text{post}}^2(x_*)) \quad (2.7)$$

with

$$\mathbf{m}_{f,\text{post}}(x_*) = \phi(x_*)^\top \mu_{\beta,\text{post}} \quad (2.8)$$

$$\sigma_{f,\text{post}}^2(x_*) = \phi(x_*)^\top \Sigma_{\beta,\text{post}} \phi(x_*) \quad (2.9)$$

and is again a Gaussian distribution. Equation (2.7) implements what Rasmussen and Williams [170] have called the *function space view* as it provides an estimate of the posterior function for all input points x_* . A reformulation of the solution in equations (2.8) and (2.9) can be obtained via the matrix inversion lemma, leading to

$$\mathbf{m}_{f,\text{post}}(x_*) = \phi(x_*)^\top \Sigma_\beta \Phi \left(\Phi^\top \Sigma_\beta \Phi + \sigma_n^2 I \right)^{-1} \mathbf{y} \quad (2.10)$$

$$\sigma_{f,\text{post}}^2(x_*) = \phi(x_*)^\top \Sigma_\beta \phi(x_*) - \phi(x_*)^\top \Sigma_\beta \Phi \left(\Phi^\top \Sigma_\beta \Phi + \sigma_n^2 I \right)^{-1} \Phi^\top \Sigma_\beta \phi(x_*), \quad (2.11)$$

refer to [170, section 2.1.2] for a detailed derivation. In cases where the number of data points is larger than the number of features ($N > F$), it is preferable to use equations (2.8) and (2.9), which involves inversion of an $F \times F$ matrix. In other cases with a larger number of features ($F > N$), which is common in inverse problems, equations (2.10) and (2.11) are more advantageous, requiring inversion of an $N \times N$ matrix.

The *maximum a posteriori* (MAP) parameter, i.e., the mode of the posterior, is often used in statistical inference as a point estimate and serves as a surrogate for the full posterior distribution. In the parametric Gaussian regression introduced above, the MAP estimate coincides with the posterior mean, that is, $\hat{\beta}_{\text{MAP}} = \mu_{\beta,\text{post}}$. MAP estimation on the model in equation (2.3) corresponds to solving the optimization problem

$$\hat{\beta}_{\text{MAP}} = \arg \max_{\beta} p(\mathbf{y}|\beta)p(\beta) \quad (2.12)$$

$$= \arg \min_{\beta} -\log \mathcal{L}(\beta|\mathbf{y}) - \log p(\beta) \quad (2.13)$$

$$= \arg \min_{\beta} \sum_{k=1}^N \frac{1}{2\sigma_n^2} (y_k - \phi(x_k)^\top \beta)^2 + \frac{1}{2} \beta^\top \Sigma_\beta^{-1} \beta + \text{const} \quad (2.14)$$

$$= \arg \min_{\beta} \frac{1}{\sigma_n^2} \|\mathbf{y} - \Phi^\top \beta\|_2^2 + \|\beta\|_{\Sigma_\beta^{-1}}^2, \quad (2.15)$$

where the negative log likelihood $-\log \mathcal{L}(\beta|\mathbf{y})$ leads to a quadratic loss term and $-\log p(\beta)$ gives rise to the *Tikhonov regularization* $\|\beta\|_{\Sigma_\beta^{-1}}^2 := \beta^\top \Sigma_\beta^{-1} \beta$, penalizing the norm of the parameter vector. Under the regularization paradigm, prior knowledge about the parameters is imposed via a penalty on the model complexity and measurements are incorporated via a loss function, which is to be minimized. Variants of equation (2.15) have been widely used for solving regression problems in different

domains, in particular for preventing overfitting as the added regularization term favors simpler models [62]. In equation (2.15) the trade-off between fitting the data and constraining the parameters is controlled by the noise parameter σ_n^2 and the covariance Σ_β . A special case is obtained by using a diagonal covariance matrix $\Sigma_\beta = \sigma_\beta^2 I$, which is known as *ridge regression*. It can be easily verified that the MAP estimate in this case only depends on the ratio $\lambda = \frac{\sigma_n}{\sigma_\beta}$, and thus depends on a single parameter for adjusting the strength of the regularization. Another important special case of the parametric Gaussian MAP regression is obtained by *maximum likelihood* (ML) estimation, which corresponds to using an uninformative prior for β (i.e., by setting $\lambda = 0$) and recovers the famous *ordinary least squares* (OLS) [219] regression result

$$\hat{\beta}_{\text{OLS}} = (\Phi\Phi^\top)^{-1}\Phi y \quad (2.16)$$

which is a special case of equation (2.5).

The linear model with a Gaussian conjugate prior constitutes the simplest instance of the family of regression models used in this thesis. The following sections consider two important generalizations. Firstly, section 2.1.2 discusses robust regression by means of M-estimation, which is a generalization of the maximum likelihood method. Secondly, section 2.1.3 examines distributions over functions given as a Gaussian process (GP), which can be understood as the infinite-dimensional equivalent of the finite-dimensional Gaussian distribution.

2.1.2 Robust regression

The classic Gaussian regression model and its variants (e.g., ridge regression and OLS), have become popular mainly because they have simple closed-form solutions [137, 251]. However, on many real-world datasets the underlying assumptions are unrealistic, for example, biomedical datasets are often strongly perturbed by artifacts [227]. Under the simple Gaussian noise model, the posterior mean is known to be extremely sensitive to outliers: in some cases even a single bad datum may completely distort the parameter estimates. This can be understood as a direct consequence of the squared loss function, which puts an excessive weight on data points with large residuals. Led by the seminal work of Tukey and Huber [97, 223], these considerations have prompted the development of robust statistical estimators that are less affected by perturbations than classic least squares approaches. In the Bayesian setting, robustness can be achieved by using a likelihood function with “fat tails”, i.e., the distribution should have larger kurtosis than the Gaussian distribution to better account for the outliers. A common approach is to use a Laplace noise model (plotted in figure 2.1), for which the maximum likelihood estimator is obtained by minimization of the mean absolute error. In this section, a generalization of maximum likelihood estimation called *M-estimation*, proposed by Huber [97], will be discussed. The method replaces the negative log likelihood $\sum_k -\log \mathcal{L}(\beta|y_k)$ with a more general loss function $\sum_k \rho(r_k)$, which must not

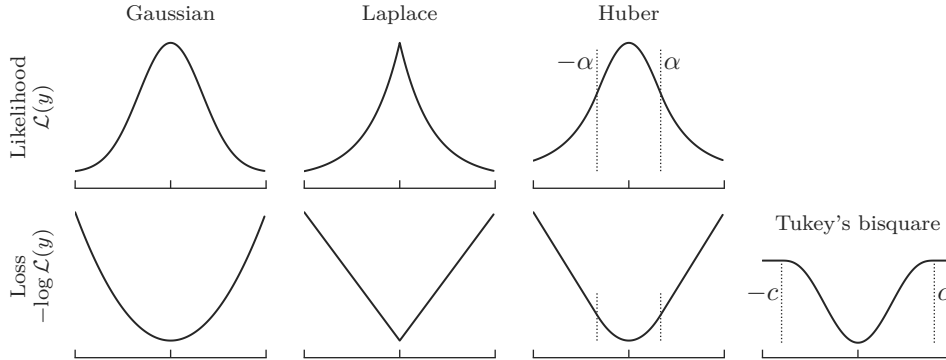


Fig. 2.1: An illustration of likelihood and loss functions used in this thesis. The squared and absolute loss functions are commonly used for solving regression problems and correspond to a Gaussian and Laplace likelihood, respectively. The Huber loss and the Tukey's bisquare loss have been proposed under the M-estimation framework as robust alternatives to the squared loss function. The Huber function [97] is quadratic for absolute values $< \alpha$ and linear elsewhere. Tukey's bisquare [22] has quadratic behavior close to the origin and is constant for absolute values $> c$.

necessarily relate to a proper probability distribution, and here r_k denotes the residuals of the model. In order to reach robustness, ρ is typically constructed to increase less rapidly for larger values than the classic quadratic loss. Two examples are plotted in figure 2.1. A very commonly used robust regression function is *Huber's loss* [97, p. 75], which has a quadratic shape for residuals close to zero, but increases linearly for larger values. Robustness can be adjusted via the parameter α , which determines the transition point between the quadratic and linear regime, refer to figure 2.1. Huber's loss function has been recommended as a reliable choice in most regression problems perturbed by outliers [251, p. 37]. Even higher robustness can be achieved via *Tukey's bisquare function* [22, section 7], also plotted in figure 2.1. Similar to Huber's loss, it has a quadratic shape close to the origin, but puts a constant loss on all larger residuals, and the cutoff is determined by the parameter c shown in figure 2.1. Notably, due to the constant part of the function, it cannot be represented by maximum likelihood estimation. M-estimation is one of the most prevalent techniques for robustifying regression models, and has been employed in a wide range of domains, refer to [227] for an overview. It appears to be particularly useful for estimating parameters in biomedical data, where model mismatch and artifacts are omnipresent. In general, M-estimators do not have closed-form solutions and therefore require iterative solvers. In this thesis, iteratively reweighted least squares (IRLS) is used, which weights all data points according to their influence on the solution of a standard least squares problem. The weights can be determined as $\frac{1}{r} \frac{d\rho(r)}{dr}$, where $\psi(r) = \frac{d\rho(r)}{dr}$ is called the *influence function* and r are *residuals*, refer to [250, section 9.4] for further details.

In many cases, IRLS only requires very few iterations for convergence—of course, gradient descent can also be used for solving M-estimators, but would often require more iterations than IRLS.

2.1.3 Nonparametric Gaussian regression

The parametric Gaussian model, introduced in section 2.1.1, implicitly imposes a prior on the values of the function $f(x)$ [170]. It can be easily shown that its mean is

$$m_f(x) := \phi(x)^\top \mu_\beta \quad (2.17)$$

and the covariance between function values at any two input points x and x' is given by the inner product

$$k_f(x, x') := \phi(x)^\top \Sigma_\beta \phi(x'). \quad (2.18)$$

However, a parametric prior might not be the most natural way of formulating knowledge about the characteristics of the function space because it might be difficult to relate the parameter prior to the desired property, e.g., smoothness, periodicity or stationarity of functions [183]. Thus, the following question arises: is there a more direct way to describe distributions over functions?

An extension of the feature-based Gaussian regression can be obtained by exploiting the fact that in equation (2.10) and (2.11) features only appear as inner products $\phi(x)^\top \Sigma_\beta \phi(x')$ and thus could be rewritten in terms of the functions' prior covariance defined above. Interestingly, for certain sets of features it is possible to analytically derive the solution to the inner product $\phi(x)^\top \Sigma_\beta \phi(x')$ without ever having to explicitly evaluate the feature map $\phi(x)$. This approach is referred to as *kernel trick* and allows to work with high-dimensional, possibly infinite-dimensional, feature sets by defining a *positive definite* kernel function $k(x, x')$ [199]. For instance, the *squared exponential* (SE) kernel

$$k_{\text{SE}}(x, x') = \sigma_{\text{SE}}^2 \exp\left(-\frac{(x - x')^2}{2\ell_{\text{SE}}^2}\right) \quad (2.19)$$

with univariate inputs x and x' , length-scale ℓ_{SE} and variance σ_{SE}^2 can be derived as the inner product of an infinite number of equispaced Gaussian basis functions [129, section 3.1]. In that sense, kernels provide a *nonparametric* approach to the Gaussian regression stated in equations (2.10) and (2.11), allowing to represent an infinite number of features and weights, without having to compute them explicitly.

Gaussian process regression

In the following, the nonparametric regression idea will be explored in terms of *Gaussian processes* (GPs). GPs are stochastic processes that are used to describe distributions over functions on a continuous domain by assigning covariance to finite

subsets of function values. GPs have been employed in a wide range of problems, including, among others, in spatio-temporal statistics and machine learning. The interested reader is referred to the book by Rasmussen and Williams [170] for a comprehensive introduction. Formally, the function $f(x)$ is a Gaussian process with mean function $m_f(x)$ and kernel function $k_f(x, x')$ if any finite subset of function values $f(x_{1:N}) := [f(x_1), \dots, f(x_N)]^\top$ has a joint multivariate Gaussian distribution

$$f(x_{1:N}) \sim \mathcal{N}(m_{f,X}, K_{f,XX}) \quad (2.20)$$

with mean vector

$$m_{f,X} = m_f(x_{1:N}) := [m_f(x_1), \dots, m_f(x_N)]^\top \quad (2.21)$$

and covariance matrix

$$K_{f,XX} = k_f(x_{1:N}, x_{1:N}), \quad (2.22)$$

where $k_f(x_{1:N}, x_{1:N})$ denotes the evaluation of the kernel for all pairs of input points (i.e., $K_{f,XX}^{(i,j)} = k_f(x_i, x_j)$). The GP distributed function is denoted by

$$f(x) \sim \mathcal{GP}(m_f(x), k_f(x, x')). \quad (2.23)$$

Based on the definition above, a GP can be understood as the limit of a multivariate Gaussian distribution, allowing to assign a probability measure over an infinite-dimensional function space [89, section 4.1]. Its finite-dimensional marginal is obtained by extracting the appropriate subvectors and submatrices from the mean and kernel function, respectively. The mean function specifies the average value of the distribution, and the kernel defines how different points co-vary against each other:

$$m_f(x) = \mathbb{E}[f(x)] \quad (2.24)$$

$$k_f(x, x') = \mathbb{E}[(f(x) - m_f(x))(f(x') - m_f(x')))]. \quad (2.25)$$

The mean is usually, without loss of generality, chosen to be constant or zero. Thus, the main tool for modeling the characteristics of the function and encoding prior knowledge is the kernel. In machine learning, a very common approach is to exploit local correspondence between data points: when inputs points are close together, their covariance is modeled to be high and to decrease with larger distance. This is commonly implemented via the squared exponential kernel, refer to equation (2.19). Several other kernels may be considered for biomedical data, some of which will be introduced in section 2.1.3.

GPs can be used as priors over latent functions, allowing to infer hidden structure from measured data and predict new inputs. In *GP regression*, measurements




Fig. 2.2: An illustrative example of Gaussian process regression. The left plot depicts samples and the confidence interval of a zero-mean GP prior $f(x)$ with squared exponential kernel and hyperparameters set to unit values. Conditioning on the data using a small measurement noise $\sigma_n = 0.01$ yields the posterior GP $f_{\text{post}}(x)$, for which samples, mean and confidence interval are depicted in the two remaining plots. The mean is plotted as a black line and the 95% confidence interval as a shaded area.

$y = [y_1, \dots, y_N]^\top$ are modeled using the same Gaussian likelihood function as in section 2.1.1. Thus the GP regression model can be stated as

$$f(x) \sim \mathcal{GP}(m_f(x), k_f(x, x')) \quad (2.26)$$

$$p(y|f(x_{1:N})) = \prod_{k=1}^N \mathcal{N}(y_k; f(x_k), \sigma_n^2), \quad (2.27)$$

where σ_n^2 is the variance of the measurement noise. Given the vector of measurements y , the goal is to predict the value of the function at a given test point x_* . In the case of Gaussian measurement noise, the posterior

$$f_{\text{post}}(x_*) \sim \mathcal{GP}(m_{f,\text{post}}(x_*), k_{f,\text{post}}(x_*, x'_*)) \quad (2.28)$$

has a closed-form solution and is also a Gaussian process defined by

$$m_{f,\text{post}}(x_*) = k_f(x_*, x_{1:N}) \left(K_{f,XX} + \sigma_n^2 I \right)^{-1} y \quad (2.29)$$

$$k_{f,\text{post}}(x_*, x'_*) = k_f(x_*, x'_*) - k_f(x_*, x_{1:N}) \left(K_{f,XX} + \sigma_n^2 I \right)^{-1} k_f(x_{1:N}, x'_*), \quad (2.30)$$

where $K_{f,XX}$ is the $N \times N$ kernel matrix, $k_f(x, x_{1:N}) := [k_f(x, x_1), \dots, k_f(x, x_N)]$ is defined as an N -dimensional vector and $k_f(x_{1:N}, x) := k_f(x, x_{1:N})^\top$ as its transpose, refer to [170, section 2.2] for a detailed derivation. Equations (2.29) and (2.30) can be recognized to be the nonparametric equivalent of the parametric Gaussian regression solution in equations (2.10) and (2.11). Figure 2.2 shows a simple example of a GP prior with squared exponential kernel and its posterior on an illustrative dataset. A key advantage of GPs compared to simpler parametric regression models is their ability to directly incorporate prior assumptions about the data, such as smoothness, noisiness or

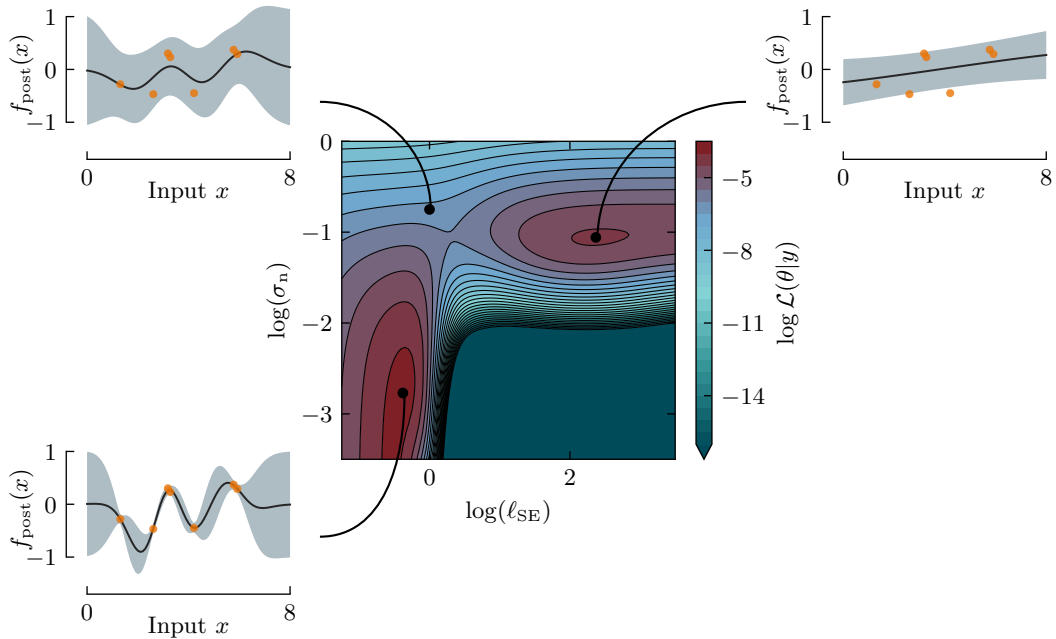


Fig. 2.3: The marginal log likelihood surface for two hyperparameters of a zero-mean GP (with SE kernel) on a toy dataset. The likelihood $\mathcal{L}(\theta|y)$ has two optima, corresponding to different interpretations of the data: a short length-scale ℓ_{SE} with small level of noise σ_n or a large length-scale ℓ_{SE} with high level of noise σ_n . The resulting GP posterior is plotted for three different sets of hyperparameters. The posterior mean is plotted as a black line and here the shaded area represents one positive and one negative standard deviation from the mean (which is different from figure 2.2, but better visualizes the relation between the GP and data points).

periodicity. The main computational challenge of GP regression comes with the need to solve linear systems involving the kernel matrix. In practice, to improve numerical stability, this is done using Cholesky decomposition, which requires $\mathcal{O}(N^3)$ runtime and $\mathcal{O}(N^2)$ storage. These requirements pose significant practical limitations as they render GP regression intractable on all but the smallest datasets [124, 127, 240]. Thus, a recurring theme of this thesis will be the reduction of computational complexity on temporal and spatio-temporal data, which will be discussed in detail in chapters 5–6 and represents a central contribution of this thesis.

An important advantage of GP regression is that — as a Bayesian framework — it provides means to learn hyperparameters θ of the model, such as the measurement noise variance σ_n^2 or parameters of the kernel. Calculating the full posterior distribution over the hyperparameters proves to be analytically intractable (and would require Markov chain Monte Carlo sampling or approximate inference), which is why most researchers resort to maximum marginal likelihood estimation (sometimes called type-II maximum

likelihood estimation). The marginal log likelihood of the hyperparameters has the closed-form solution

$$\log \mathcal{L}(\theta|y) = -\frac{1}{2}y^\top \left(K_{f,XX} + \sigma_n^2 I\right)^{-1} y - \frac{1}{2} \log |K_{f,XX} + \sigma_n^2 I| - \frac{N}{2} \log 2\pi, \quad (2.31)$$

and can be maximized using gradient ascent [170]. An illustrative example of a hyperparameter likelihood surface is shown in figure 2.3. Equation (2.31) requires solving a linear system and calculating the log determinant of the kernel matrix, which is again solved using Cholesky decomposition, leading to an $\mathcal{O}(n^3)$ runtime for hyperparameter learning.

Kernel functions

Not every bivariate function is automatically a proper kernel. It must meet the requirements of a *positive semi-definite function*, which means that for any set of inputs $x_{1:N}$ the corresponding kernel matrix $k(x_{1:N}, x_{1:N})$ must be a positive semi-definite matrix [170, section 4.1]. The following section briefly introduces commonly used kernel functions, focusing on those that are relevant to temporal biomedical data (for example for modeling smoothness or quasi-periodicity). It will also discuss a few techniques for constructing new kernels from existing ones. When a kernel function $k(x, x')$ only depends on the difference $\tau = x - x'$, it is referred to as a *stationary* kernel. In some cases the shortened notation $k(\delta) := k(x, x')$ will be used when a stationary kernel can be written in terms of a distance δ between x and x' . Stationarity can be exploited for easing the computational complexity of GP regression, which will be discussed later on in section 2.1.4 and chapter 6.

The squared exponential kernel k_{SE} has already been introduced in equation (2.19). It is one of the most widely used stationary kernel functions, because it gives rise to a distribution of very smooth (infinitely mean-square differentiable) functions [170, section 4.2.1]. Though, it has been argued that functions generated by the squared exponential kernel may be overly smooth and unrealistic for describing physical processes [170, 217]. A viable alternative is the Matérn kernel family [131], which allows to adjust smoothness of functions via the parameter ν . Functions sampled from a GP with Matérn covariance are k times mean-square differentiable if and only if $\nu > k$ [170, section 4.2.1]. For half-integer values $\nu = k + \frac{1}{2}$ with $k = 0, 1, 2, \dots$ Matérn

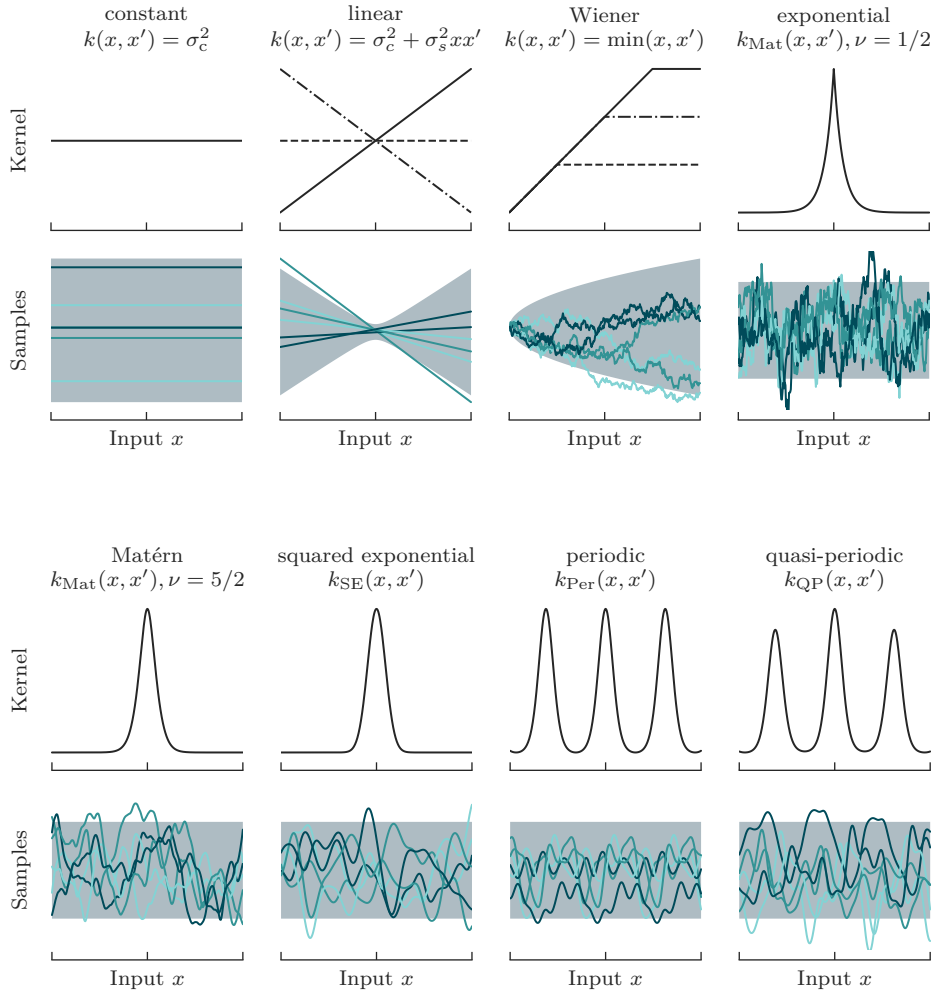


Fig. 2.4: An illustration of commonly used kernel functions. The top rows depict kernels plotted by their first inputs as $k(x, 0)$ for stationary kernels, and as $k(x, x')$ with three selected values of x' for non-stationary kernels (linear and Wiener). The bottom rows depict samples from corresponding zero-mean GPs with 95% confidence intervals visualized by the shaded area. Definitions of kernels are provided in this plot or in the main text.

kernels can be written as simple expressions, and the most commonly used cases are $\nu = \frac{1}{2}$, $\nu = \frac{3}{2}$ and $\nu = \frac{5}{2}$, which are given by

$$k_{\text{Mat},\nu=\frac{1}{2}}(\delta) = \sigma_{\text{Mat}}^2 \exp\left(-\frac{\delta}{\ell_{\text{Mat}}}\right) \quad (2.32)$$

$$k_{\text{Mat},\nu=\frac{3}{2}}(\delta) = \sigma_{\text{Mat}}^2 \left(1 + \frac{\sqrt{3}\delta}{\ell_{\text{Mat}}}\right) \exp\left(-\frac{\sqrt{3}\delta}{\ell_{\text{Mat}}}\right) \quad (2.33)$$

$$k_{\text{Mat},\nu=\frac{5}{2}}(\delta) = \sigma_{\text{Mat}}^2 \left(1 + \frac{\sqrt{5}\delta}{\ell_{\text{Mat}}} + \frac{5\delta^2}{3\ell_{\text{Mat}}^2}\right) \exp\left(-\frac{\sqrt{5}\delta}{\ell_{\text{Mat}}}\right), \quad (2.34)$$

where σ_{Mat}^2 is the variance and ℓ_{Mat} the length-scale parameter. The case $\nu = \frac{1}{2}$ corresponds to the exponential kernel, which on a one-dimensional domain is the covariance function of the Ornstein-Uhlenbeck process [225]. For $\nu \rightarrow \infty$ the Matérn covariance converges to the squared exponential [170, section 4.2.1].

Given positive definite kernels $k_1(x, x')$ and $k_2(x, x')$, the sum $k_1(x, x') + k_2(x, x')$ and product $k_1(x, x') \cdot k_2(x, x')$ are also positive definite kernels. These two simple rules already allow to construct highly expressive models, e.g., by combining existing kernel functions with different characteristic length-scales. Another useful technique for adapting existing models is vertical rescaling with a deterministic function $\sigma : \mathcal{D}_{\text{in}} \rightarrow \mathbb{R}$, which leads to the non-stationary covariance function $\sigma(x)k(x, x')\sigma(x')$. If $k(x, x')$ has unit variance, $\sigma(x)^2$ directly determines the variance of the rescaled kernel at each input x . Another technique worth mentioning is to transform points in the input space via a mapping (or warping) function $\phi : \mathcal{D}_{\text{in}} \rightarrow \mathcal{D}_{\text{out}}$ which gives rise to the kernel $k(\phi(x), \phi(x'))$ and allows to construct kernels with non-stationary length-scales. The output space \mathcal{D}_{out} of the mapping must not necessarily have the same dimensionality as the input space \mathcal{D}_{in} .

Using the warping technique, MacKay [129, section 5.4.3] constructed a particularly useful family of kernels: he proposed to map one-dimensional inputs $x \in \mathbb{R}$ to the two-dimensional space $u(x) = [\sin(2\pi x/p), \cos(2\pi x/p)]^\top$, where p is the period length. When using a (locally) smooth kernel on the transformed points $u(x)$ this gives rise to the periodic kernel $k(u(x), u(x'))$. For instance, by using an isotropic squared exponential kernel, one obtains

$$k_{\text{Per}}(\delta) = \sigma_{\text{Per}}^2 \exp\left(-\frac{2 \sin^2(\pi\delta/p)}{\ell_{\text{Per}}^2}\right), \quad (2.35)$$

where σ_{Per}^2 denotes the variance and ℓ_{Per}^2 the local length-scale. This kernel generates a distribution of strictly periodic functions with period length p and local behavior (on length-scales $< p$) is similar to a squared exponential kernel. The same construction could be done with other local kernels, e.g., to change local smoothness. The periodic kernel proves to be highly useful for modeling biomedical problems, which are often

governed by periodic patterns. To allow for deviations from a fixed waveform over periods, *quasi-periodic* kernels can be constructed by multiplying the strictly periodic function k_{Per} with a smooth kernel function that has length-scale $> p$. In this thesis, quasi-periodicity will be modeled by multiplication with the squared exponential, leading to

$$k_{\text{QP}}(\delta) = k_{\text{SE}}(\delta)k_{\text{Per}}(\delta). \quad (2.36)$$

2.1.4 Linear stochastic differential equations

The regression models introduced in this chapter are usually computed via batch methods, which means that all observations $y = [y_1, \dots, y_N]^\top$ are required to solve the inference. A batch solution has important practical limitations: for instance, the runtime of GP regression scales cubically with the number of data points, which is intractable on long (possibly unbounded) time-series data. This is even more problematic in real-time applications, where the inference would have to be updated with every new data point [72, 211]. The following section therefore investigates reformulations of temporal GP priors in terms of state-space models. In modern machine learning research, the representation of GPs as kernels is prevalent. However, the state-space view of GPs has been known and used for a very long time, for example in statistical physics in the form of Brownian motion and the Ornstein-Uhlenbeck process [225]. This was also precisely the subject of Kalman’s seminal 1960 paper [108] in which he presented a state-space reformulation of the Wiener filtering problem [85]. The reformulation in state-space form gives rise to highly scalable recursive algorithms based on Kalman filters/smoothers, which process one data point at a time, leading to quasi-linear complexity.

Note that for sake of simplicity, in this section, the time axis t is considered as the input domain, but this can of course be generalized to other one-dimensional domains. Here, the aim is to represent the temporal stochastic process $f(t)$ as a linear *stochastic differential equation* (SDE) of the form

$$dx(t) = Fx(t) dt + Ld\beta(t), \quad (2.37)$$

where the state vector $x(t)$ contains the derivatives of $f(t)$, $\beta(t)$ is a Brownian motion (mathematically described by the Wiener process) with diffusion constant q_c , the process is defined by constant system matrices F and L with appropriate dimension, and the expression $d\beta(t)$ corresponds to the increment of an Itô integral, refer to [191, chapter 4] for further details. By defining a measurement matrix H , $f(t)$ can be extracted from the system’s states via $f(t) = Hx(t)$.

The solution trajectory of the SDE is a Gaussian process, because it is given by a linear transformation of a Brownian motion (which is a Gaussian process) and Gaussianity is preserved under linear operations [191, section 6.1]. The stationary

state covariance of $f(t)$ (i.e., after letting the process run for an infinite time) can be calculated in closed-form in terms of the state-space matrices and is given by

$$k_f(\tau) = \begin{cases} HP_\infty\Phi(\tau)^\top H^\top, & \text{if } \tau \geq 0 \\ H\Phi(-\tau)P_\infty H^\top, & \text{if } \tau < 0 \end{cases}, \quad (2.38)$$

where $\Phi(\tau) = \exp(F\tau)$ is a matrix exponential with $\tau = t - t'$, and P_∞ is computed from the continuous-time algebraic Lyapunov equation

$$\frac{dP_\infty}{dt} = FP_\infty + P_\infty F^\top + Lq_c L^\top = 0, \quad (2.39)$$

which has a solution when the model described by the state matrix F corresponds to a stable system [188, 191].

Converting covariance functions to state-space models

The close relation between covariance functions and SDE models provides a range of new techniques for working with Gaussian processes, including Kalman filtering/smoothing algorithms as well as methods from control theory. In a series of papers, Simo Särkkä and colleagues have investigated this very fruitful connection in detail [85, 188–190, 192, 212]. Notably, they have presented practical methods for converting kernel functions to equivalent state-space models.

In the following, the general technique for constructing a state-space model described by F , L and q_c for a given covariance function $k(\tau)$ of a process $f(t)$ will be discussed. The framework presented by Särkkä and Hartikainen [85] relies on the *spectral density*, which is the expectation of the squared Fourier transform $F(i\omega)$ of the process:

$$S(\omega) = \mathbb{E} \left[|F(i\omega)|^2 \right] = \mathbb{E} [F(i\omega)F(-i\omega)] \quad (2.40)$$

The Wiener-Khinchin theorem states that the spectral density forms a Fourier transform pair with the covariance function, that is, $S(\omega)$ is immediately given by the Fourier transform of $k(\tau)$:

$$S(\omega) = \int_{-\infty}^{\infty} k(\tau)e^{-j\omega\tau} d\tau. \quad (2.41)$$

When the spectral density can be written as a rational function, *spectral factorization* can be applied, leading to

$$S(\omega) = G(i\omega)q_c G(-i\omega), \quad (2.42)$$

where all poles of $G(i\omega)$ have positive imaginary parts and all poles of $G(-i\omega)$ have negative imaginary parts. Then, to obtain a state-space model with the same spectral density, one can construct a linear system with transfer function $G(i\omega)$ driven by white noise with spectral density q_c . The construction of a state-space model with a given

rational transfer function $G(i\omega)$ is straightforward, e.g., by using companion form (or canonical controllable/observable form). The reformulation proposed by Särkkä et al. is exact for *Markovian* kernel functions like the Matérn kernel — other kernels like the squared exponential or periodic kernel would theoretically give rise to an infinite-dimensional state dimension, which is why finite-dimensional approximate forms have been proposed [85, 190, 212].

Example. To illustrate the approach described above, the equivalent state-space representation for the Matérn kernel is treated below, as detailed in [192, 210]. The spectral density of the Matérn kernel for half-integer values $\nu = p + 1/2$ is simply given by

$$S_{\text{Mat}}(\omega) = q_c(\lambda^2 + \omega^2)^{-(p+1)}, \quad (2.43)$$

which is a rational function, and can be factorized as

$$S(\omega) = (\lambda + i\omega)^{-(p+1)}q_c(\lambda - i\omega)^{-(p+1)}, \quad (2.44)$$

where $\lambda = \sqrt{2\nu}/\ell_{\text{Mat}}$. Thus, a state-space model can be constructed from the transfer function

$$G(i\omega) = (\lambda + i\omega)^{-(p+1)}, \quad (2.45)$$

corresponding to the stable part of $S(\omega)$. For half-integer values of ν , equation (2.45) can be expanded to get the coefficients of the denominator polynomial. To give an example, the transfer function $G(i\omega)$ for $\nu = \frac{3}{2}$ has a denominator polynomial of degree two with coefficients $1, 2\lambda, \lambda^2$. The corresponding SDE with F and L in controllable canonical form is given by

$$dx(t) = \underbrace{\begin{bmatrix} 0 & 1 \\ -\lambda^2 & -2\lambda \end{bmatrix}}_F x(t)dt + \underbrace{\begin{bmatrix} 0 \\ 1 \end{bmatrix}}_L d\beta(t), \quad (2.46)$$

where the diffusion constant of the Brownian motion is $q_c\lambda^3\sigma^2$ [210]. The first component of the state vector is equivalent to the GP with Matérn covariance and can be extracted via the measurement matrix $H = \begin{bmatrix} 1 & 0 \end{bmatrix}$.

When the spectral density can not be written as a rational function, like for the squared exponential or periodic kernel, approximate forms must be used. For the squared exponential kernel, truncated Taylor series or Padé approximants have been proposed, refer to [85, 190]. The spectral density of the periodic kernel consists of a series of Dirac delta functions at the fundamental frequency and higher harmonics — it is therefore not suited for a polynomial approximation [175]. Instead, Solin and Särkkä proposed to use a set of parallel harmonic oscillators, whose initial states are tuned to emulate the behavior of the periodic GP [212]. For these types of kernels, the state-space model may in some cases require a relatively large number of states, e.g., due to a high-degree rational function or a large number of harmonic oscillators, to achieve a good approximation of the original kernel.




Fig. 2.5: A comparison of GP batch regression, Kalman filtering and smoothing solutions for a squared exponential kernel on an illustrative dataset. Hyperparameters of the kernel were set to unit values and $\sigma_n = 0.01$. The state-space model was derived from a sixth-order Taylor expansion of the spectral density. The posterior mean is plotted as a black line and the 95% confidence interval as a shaded area.

Sequential inference

The reformulation in state-space form allows to solve the GP regression efficiently in a sequential fashion. To this end, the continuous model needs to be discretized with respect to the noisy measurements y_1, \dots, y_N at time points t_1, \dots, t_N , which yields the discrete-time model

$$\begin{aligned} x(t_{k+1}) &= A_k x(t_k) + q_k \\ y_k &= Hx(t_k) + \epsilon, \end{aligned} \quad (2.47)$$

where $q_k \sim \mathcal{N}(0, Q_k)$ is process noise and $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$ is measurement noise. The discrete-time matrices can be determined using the analytical expressions

$$A_k = \Phi(\Delta t_k) \quad (2.48)$$

$$Q_k = \int_0^{\Delta t_k} \Phi(\Delta t_k - \tau) L q_c L^\top \Phi(\Delta t_k - \tau)^\top d\tau, \quad (2.49)$$

where $\Phi(\tau) = \exp(F\tau)$, and the integral in (2.49) can be solved numerically by using matrix fraction decomposition, refer to [185, remark 2.2]. The model in (2.47) is now in a form that can be used in standard Bayesian state estimation, e.g., via linear *Kalman filtering/smoothing*, which processes data points sequentially. The batch GP regression solution is recovered by the Rauch-Tung-Striebel smoother [172]: it first executes the standard linear Kalman filter, which goes over data points forward in time, and then employs a backward smoothing pass to improve these estimates, refer to appendix A.2 for details. The forward pass must be initialized with the steady state covariance matrix, i.e., $x(t_0) \sim \mathcal{N}(0, P_\infty)$. The state can be predicted at any given test point t_* by discretizing the model for that point and doing forward/backward prediction

from the previous and subsequent data point. The GP batch solution and Kalman filter/smoothing solutions are visualized in figure 2.5. Computational complexity of GP regression is reduced to $\mathcal{O}(NM^3)$, where N is the number of timesteps and M is the dimension of the state vector. In practice, due to $N \gg M$, the complexity is usually dominated by the linear term. Finally, it is also worth mentioning that the marginal model likelihood can be calculated alongside a single forward Kalman filter pass, which enables highly efficient optimization of hyperparameters, refer to appendix A.2 and to [187, section 12.3.2] for further details.

2.2 Respiratory physiology and mechanical ventilation

Having explained the basics of parametric and nonparametric regression in the previous section, this next section will provide the *physiological and clinical* fundamentals necessary for understanding the methodological parts of this thesis. In the following sections, an introduction to the physiology of the human *respiratory system* is given, ranging from respiratory anatomy to respiratory control and lung mechanics. Means of providing *respiratory support* to patients who cannot sustain sufficient respiration due to failure of parts of the respiratory system are also discussed briefly. Among the many directions currently being explored for improving mechanical ventilation therapy, the main focus of this section (and of the whole thesis) is on the measurement of *respiratory muscle effort*, which is a key prerequisite for implementing a *diaphragm-protective ventilation* strategy. This treatment approach has been proposed as a highly promising technique to improve the long-term outcome and well-being of critically ill patients undergoing mechanical ventilation, and will be treated in detail in section 2.2.2.

2.2.1 Physiology

Most multi-cellular organisms, including humans, rely on aerobic metabolism of high-energy molecules and thus a continuous supply of oxygen to cells to sustain life. The process of taking up oxygen from the environment and removing carbon dioxide, one of the main waste products of metabolism, is referred to as *respiration*. The respiratory system comprises a series of organs and structures responsible for gas exchange. Anatomically, it consists of the airways through which air is conducted, the two lungs, the chest wall, and the respiratory controller [202, chapter 1]. Air enters the body via the upper airways, where it is filtered from pollutants and moistened, and then passes through the trachea into the two primary bronchi, leading to the lungs. The bronchi branch into a tree-like structure of increasingly narrower tubes, finally reaching small air-filled cavities, the *alveoli*. The alveolar membrane is permeable to gases and surrounded by a network of capillaries, thus, oxygen and carbon dioxide can be exchanged between air and blood via diffusion [235].

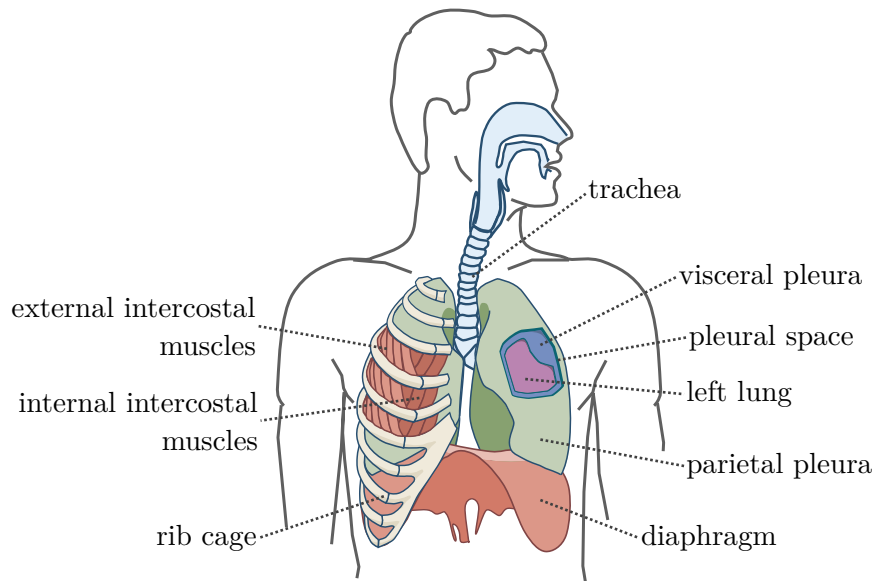


Fig. 2.6: An illustration of the human chest including the lungs, rib cage and the main respiratory muscles located in the chest wall.

In order for air to flow in and out of the lungs, a pressure gradient between alveolar and airway opening pressure must be generated. In healthy, spontaneously breathing humans, this is accomplished by expansion and contraction of the thoracic cavity by means of skeletal muscles located in the chest wall (also referred to as thoracic wall). On its inside, the chest wall is separated from the lungs by the *pleural cavity*, which is a thin fluid-filled gap, enabling the lungs to move freely against it. This cavity is enclosed by two opposing membranes, the visceral and parietal pleura, which stay in close contact to each other through adhesive forces. Refer to figure 2.6 for an illustration of the lungs' anatomy. During relaxed breathing, the pressure inside the pleural cavity, denoted by P_{pl} , is slightly negative due to the inward elastic recoil of the lung tissue which is balanced by the rib cage and diaphragm [202, 235].

The most important inspiratory muscle in humans is the *diaphragm*, see figure 2.6. It is a dome-shaped muscle located at the base of the lungs between the thoracic and the abdominal cavity. Its muscle fibres attach anteriorly/laterally to the costal margin and posteriorly to the 11th and 12th ribs as well as the vertebral column [52]. The fibres converge into a non-contractile plate at the apex of the diaphragm called the central tendon [52]. A contraction flattens the diaphragm, causing the thoracic cavity to expand caudally. In addition to the diaphragm, the *external intercostal muscles* also contribute to inspiratory efforts. They are located in the intercostal spaces shown in figure 2.6 and their contraction causes motion of the rib cage outward, which also increases the thoracic volume [47, 202]. Under relaxed breathing, expiration is often

done passively by means of the elastic recoil forces of the lungs and chest wall. Under exercise or when respiratory load is increased, expiratory muscles can be recruited to actively contract the thoracic cavity. The main muscles involved are the *internal intercostal muscles* shown in figure 2.6, which help to lower the rib cage, and the *abdominal wall muscles*, which increase abdominal pressure, pushing the diaphragm into a more superior position [47].

The rate and depth of spontaneous breathing is controlled by the respiratory center located in the medulla oblongata and pons within the brain stem [202, chapter 2]. The respiratory center receives afferent stimuli from central and peripheral chemoreceptor cells which provide feedback with respect to partial pressures of oxygen and carbon dioxide as well as acidity (pH) of arterial blood. Other types of receptors also give relevant feedback such as stretch receptors and mechanoreceptors located in the lung and diaphragm. The respiratory center aims to keep the partial pressures, and in particular CO₂, which is the primary driver of breathing effort, within tight bounds by generating a breathing pattern [220]. It excites the diaphragm via the phrenic nerve, and the intercostal muscles via the intercostal nerves. The abdominal muscles are innervated by lower intercostal, subcostal and lumbar nerves [221]. Breathing patterns can also be generated consciously, for instance, automatic control can be overwritten or modulated by volitional control via cortical or limbic pathways [92].

2.2.2 Respiratory failure and assisted ventilation

When patients are unable to perform the work of breathing required for respiration—a condition referred to as *respiratory failure*—mechanical ventilation is used as a life-saving intervention. There is a wide range of conditions that can lead to respiratory failure, concerning different components of the respiratory system. This includes failure of the respiratory pump (e.g., due to muscular or neurodegenerative diseases), lung conditions that affect alveolar gas exchange (e.g., pneumonia or interstitial lung disease), and diseases that obstruct airflow (e.g., chronic obstructive pulmonary disease). One condition that has received tremendous attention from researchers due to its severity is the *acute respiratory distress syndrome* (ARDS), which has been described as “the sum of most of the problems encountered in intensive care” [64]. It is a rapid onset of systemic inflammation characterized by pulmonary fluid accumulation and damage to the alveolar barriers [132]. Despite decades of research, its mortality remains very high at around 35% to 40% [28].

Mechanical ventilation is used to support or fully replace spontaneous breathing—per se, it is not intended as therapy for the underlying disease, but rather provides additional time for healing and treatment [64]. A wide range of ventilation modes exist today, most of which generate a pressure gradient against the alveoli by applying positive pressure at the airway opening or in the trachea, denoted by P_{aw} . In *controlled mandatory ventilation* modes, breaths are determined mostly by the ventilator and

have a preset amplitude and shape, without requiring any spontaneous effort from the patient. In contrast, in *assisted ventilation* modes, the work of breathing is shared between patient and machine, which provides relief for the patient while promoting spontaneous efforts [78]. For instance, in *pressure support ventilation* (PSV), which is the prevalent assisted ventilation mode [57], the patient triggers every breath (via a fixed flow threshold) and then the ventilator delivers a constant support pressure. Usage of assisted ventilation is beneficial to many patients — in particular for improving oxygenation [166] and for *weaning* — but it also necessitates very careful monitoring of patient effort and patient-ventilator interaction. Injury to the lungs and diaphragm from excessive or asynchronous spontaneous efforts must be prevented, but at the same time, the level of respiratory activity must be high enough to not exacerbate respiratory muscle atrophy. A brief overview of selected mechanisms involved in lung and diaphragm injury during assisted ventilation is provided below, refer to [31, 68, 244] for more details.

Lung injury during assisted spontaneous breathing

There is accumulating experimental and clinical evidence for mechanisms of lung injury via spontaneous breathing [31]. In severe ARDS, clinical trials have shown that early usage of neuromuscular blockade, preventing spontaneous efforts, reduces both barotrauma and mortality [154]. Experimental data by Yoshida et al. [245] in ARDS animal models confirmed that strong spontaneous breathing causes transpulmonary pressure P_{tp} , i.e., total pressure acting on the lung, to reach injurious levels that exceed the plateau pressure visible on the ventilator. In addition, some types of patient-ventilator asynchrony (e.g., double triggering and reverse triggering) have been linked to an increase of tidal volume and transpulmonary pressure, which also contributes to lung stress [244]. Finally, several other harmful patterns have been associated with strong breathing effort such as pendelluft and increased perfusion causing edema, refer to [31, 244] for further details.

Diaphragm injury during assisted spontaneous breathing

The harmful effects of mechanical ventilation on diaphragm health have also been widely recognized as an important factor for patient outcome [197]. Ventilator-induced diaphragm dysfunction (VIDD) is very common and predisposes to prolonged ventilator dependence [65]. Notably, Dres et al. [53] found that 63% of studied patients had diaphragm dysfunction when undergoing their first spontaneous breathing trial and, in these patients, weaning failure was more likely to occur. Several mechanisms of diaphragm myotrauma were identified and described [31]. One of the most relevant mechanism is muscle atrophy due to overassistance and diaphragm disuse. Goligher et al. [66] reported a decrease in diaphragm thickness in 44% of ventilated patients

(measured over one week) and revealed its correlation with high driving pressure as well as low diaphragm contractile activity. Another important mechanism is load-induced diaphragm injury due to ventilator underassistance. Repeated concentric diaphragm contractions against excessive loads induce acute diaphragm injury and inflammation [103]. The effect of eccentric diaphragmatic loading is believed to be even more damaging [197] and occurs when the diaphragm contracts during the expiration.

Lung- and diaphragm-protective ventilation

Considering the mechanisms described above, the following can be concluded: to avoid injury to the lungs and diaphragm during assisted ventilation, spontaneous breathing must be monitored very closely and kept within a safe range. This paradigm has come to be known as *lung- and diaphragm-protective ventilation*. It has been formulated across a series of publications, see [31, 67, 68, 197], whose authors advocate to tailor an adequate workload share between patient and ventilator and to optimize patient-ventilator interaction by reducing asynchronous breathing. These problems (optimizing workloads and asynchrony) are directly tied to the challenge of accurately recognizing the amplitude and timing of the patient's own respiratory effort, which will be the focus of section 2.2.4 and following chapters.

2.2.3 Lung mechanics modeling

Mathematical models of the lung have long been used for improving understanding of physiology, simulating data, and estimating respiratory parameters. In the following, the most commonly used model of lung mechanics, the linear *equation of motion* (EqM), is derived. It has been known at least since the early 1950s, where it was formulated by Otis et al. [153], and goes back to the pioneering work by Rohrer [178, 179], who discovered many of the essential principles of lung mechanics. Note that this model uses a lumped-element approach, that is, physical quantities will be assumed to be homogeneous and the system will be represented by idealized pneumatic components, which are visualized by the single-compartment model in figure 2.7. Under this model, the total pressure acting on the respiratory system is the sum of the pressure at the airways P_{aw} generated by the ventilator and the pressure P_{mus} (with $P_{mus} > 0$ representing inspiratory effort) exerted by the respiratory muscles:

$$P_{rs}(t) = P_{aw}(t) + P_{mus}(t). \quad (2.50)$$

The applied respiratory system pressure $P_{rs}(t)$ forms an equilibrium with the opposing pressures arising from air, lung and thorax motion. Firstly, the motion of air leads to a pressure drop

$$P_{res}(t) = R\dot{V}(t) \quad (2.51)$$

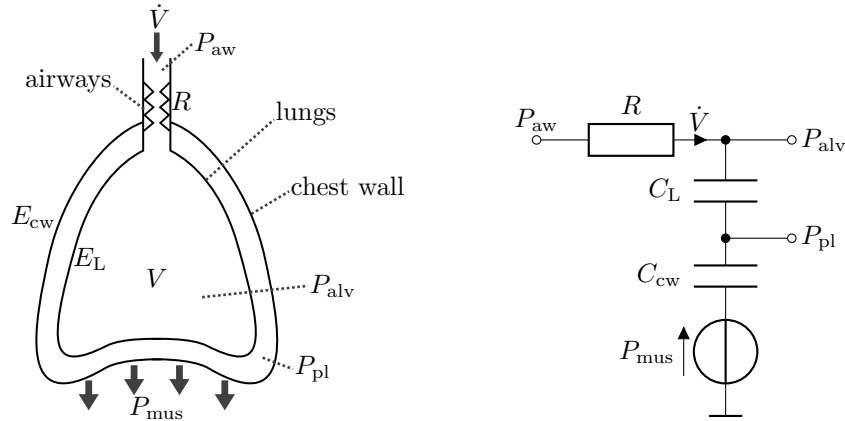


Fig. 2.7: The single-compartment model of lung mechanics. The diagram on the left provides an overview of the model's main components and parameters. The alveolar space is represented by a single cavity with volume V enclosed by two elastic containers, the lung and chest wall, with elastances E_L and E_{cw} , respectively. The airways are represented by a narrow tube, which has a pneumatic resistance R against the flow of air \dot{V} . The equivalent circuit diagram on the right uses electrical components to represent this model. Airflow is represented by current, volume is represented by electrical charge, and pressure is represented by voltage. The respiratory cavity is implemented using two capacitors, whose capacitances are $C_L = 1/E_L$ and $C_{cw} = 1/E_{cw}$, respectively. The resistor models the pressure drop across the airways, and its electrical resistance corresponds to the lung resistance R .

across the airways, which is assumed to be proportional to flow $\dot{V}(t)$. The proportionality constant R is the respiratory *resistance*. Equation (2.51) is a variant of the model proposed by Rohrer [178], who also included a second pressure term proportional to $\dot{V}(t)|\dot{V}(t)|$, which will here be omitted for sake of simplicity. Secondly, the extension of lung tissue and thoracic wall creates an elastic recoil pressure

$$P_{el}(t) = (E_L + E_{cw})V(t), \quad (2.52)$$

which is assumed to be proportional to the lung volume $V(t)$. The proportionality is determined by the parameters E_L and E_{cw} , corresponding to the *elastance* of the lung and chest wall, respectively. In clinical practice, people also often use *compliance* to describe the respiratory system, which is the inverse of the elastance, i.e., $C_L = \frac{1}{E_L}$ and $C_{cw} = \frac{1}{E_{cw}}$. Given the expressions above, and using the equilibrium between applied and opposing pressures, one gets

$$P_{rs}(t) = P_{aw}(t) + P_{mus}(t) = P_{res}(t) + P_{el}(t). \quad (2.53)$$

Thus, the classic formulation of the EqM

$$P_{aw}(t) + P_{mus}(t) = R\dot{V}(t) + (E_L + E_{cw})V(t) \quad (2.54)$$

can be constructed, which is a first-order differential equation in $V(t)$. Another very simple EqM can be derived using an equilibrium of all forces acting on the chest wall, which leads to

$$P_{\text{pl}}(t) = E_{\text{cw}}V(t) - P_{\text{mus}}(t), \quad (2.55)$$

where $P_{\text{pl}}(t)$ is the pleural pressure. This equation has been widely used in clinical research for estimating P_{mus} , which will be discussed in detail in the following section. Note that in this model P_{pl} is zero at end-expiration. A negative end-expiratory P_{pl} could be modeled simply by changing the operating point of the chest wall characteristic. *Transpulmonary pressure* can be expressed as the difference between airway pressure and pleural pressure, i.e.,

$$P_{\text{tp}}(t) = P_{\text{aw}}(t) - P_{\text{pl}}(t) = R\dot{V}(t) + E_{\text{L}}V(t). \quad (2.56)$$

It is important to keep in mind that the herein presented model omits many effects that have been observed in practice such as inertance, nonlinearity, viscoelasticity, chest wall resistance, flow-limitation, and more. Thus, it is not intended as a comprehensive lung mechanics model, but rather as a basic description of the most important characteristics. Much more complex models have been presented in the literature, including nonlinear [16], multi-compartment [20], alveolar-based [95] and finite element method (FEM) [18] models. However, it is precisely because of its simplicity that the equation of motion still has great significance in clinical practice and in research.

2.2.4 Measurement of respiratory muscle effort

The measurement of respiratory muscle effort in mechanically ventilated patients at the bedside is a challenging task. Observing the curves on the ventilator (of airway pressure, flow and volume) is not sufficient because neither the timing nor the amplitude of spontaneous breaths can be inferred with high accuracy due to the underdetermined nature of the problem (i.e., there is linear dependency between the predictors) [228]. A measurement of pleural pressure would be desirable, as it directly reflects the influence of the respiratory muscles on the lungs. This was recognized early on by respiratory physiologists who demonstrated measurement of pleural pressure first in animals [50] and later in humans [15] by puncture of the pleural space, which, however, involves considerable risk of lung injury. The current clinical standard for assessing the respiratory muscle pressure P_{mus} is based on the measurement of pressure in the esophagus with a nasogastric catheter. Esophageal pressure (P_{es}) provides a surrogate for pleural pressure and is now used extensively in critical care research. Further development and research of this approach is being pursued by the Pleural Pressure Working Group (PLUG). Two recent papers by the PLUG group have reviewed measurement techniques and applications of P_{es} in mechanically ventilated patients, refer to [10, 134]. The measurement of P_{es} is usually made using an air-filled balloon, which is positioned in the lower third of the esophagus just above the




Fig. 2.8: An illustrative simulation of pneumatic data under assisted spontaneous breathing. Esophageal pressure P_{es} can be understood as a superposition of respiratory muscle pressure P_{mus} and the recoil pressure $E_{cw}V$ by the chest wall (plotted as an orange line). The plot on the right is a Campbell diagram. It depicts the P_{es} - V loop and the chest wall characteristic $E_{cw}V$ is plotted as a red line. The offset $P_{es,0}$ is assumed to be zero.

diaphragm. After proper inflation of the balloon, intra-thoracic pressure variations are transmitted through the esophageal wall and measured at the proximal end of the catheter using a pressure transducer. This technique is considered to be minimally invasive and does not involve a great risk of injury, but can be uncomfortable for awake patients. The exact relationship between P_{pl} and P_{es} is still being debated and is complicated by the fact that both are spatially inhomogeneous quantities [10]. A recent study showed that P_{es} best represents P_{pl} in the dependent to middle lung region [243]. The relation also strongly depends on the balloon filling: at higher filling volumes, the absolute value of P_{es} usually has a considerable baseline offset due to the elasticity of the esophagus itself. Thus, many researchers have argued that only relative changes of P_{es} should be interpreted and that its absolute value must be treated with caution. The interested reader is referred to [141] for further details on dealing with these baseline artifacts. Moreover, due to the close proximity of the heart, P_{es} is often strongly disturbed by cardiac pressure artifacts, which significantly complicate interpretation of the underlying respiratory activity and were shown to affect clinically important quantities derived from it [201]. Several algorithms for reducing cardiac interference have been proposed, including a template subtraction technique by Graßhoff et al. [75].

Given measurements of P_{es} , respiratory muscle pressure P_{mus} can be estimated under the model stated in section 2.2.3. It can be assumed that

$$P_{es}(t) \approx P_{pl}(t) + P_{es,0}, \quad (2.57)$$

where $P_{es,0}$ denotes the baseline offset pressure generated by the esophageal wall. Thus, using equation (2.55) P_{mus} is estimated by correcting P_{es} for the influence of the chest wall inflation via

$$P_{mus}(t) \approx E_{cw}V(t) - P_{es}(t) + P_{es,0}. \quad (2.58)$$

To this end, an estimate of the chest wall elastance E_{cw} is needed, which can be obtained under passive conditions [134] or during spontaneous breathing when the expiration is passive as discussed by Graßhoff et al. [74]. The offset $P_{es,0}$ can simply be estimated from P_{mus} by identifying a baseline of the signal in between muscle efforts. There are two commonly used metrics for respiratory effort that can be derived from P_{mus} . The first one, work of breathing (WOB), is a measure of pneumatic energy expenditure, calculated as the integral of P_{mus} over volume, i.e.,

$$\text{WOB} = \int P_{mus}(V(t)) dV(t), \quad (2.59)$$

which corresponds to an area enclosed by the P_{es} -volume diagram, also called *Campbell diagram*, refer to figure 2.8. Another common metric is the pressure–time product (PTP), calculated as the integral of P_{mus} over time, i.e.,

$$\text{PTP}_{mus} = \int P_{mus}(t) dt, \quad (2.60)$$

which is simply the area under the P_{mus} curve. Both metrics are typically calculated in a breath-wise fashion as a measure of inspiratory effort. The pressure–time product is considered to be advantageous when little volume is generated, for instance due to muscle weakness or airway obstruction.

Despite its usefulness, P_{es} is not yet routinely employed in many clinics, because it is hampered by technical difficulties [243]. The measurement is very dependent on catheter positioning and therefore requires time-consuming preparation and calibration [21]. The P_{es} signal is considered quite difficult to interpret because, as discussed above, it is influenced by the esophageal wall [141], cardiogenic pressure artifacts [75, 201] and elastic recoil of the chest wall [134]. As a consequence of these challenges, P_{es} is still mainly considered a research tool and only used in very few specialized intensive care units. To support adoption of diaphragm-protective ventilation in intensive care units, a simpler non-invasive alternative to P_{es} is needed. The development of such a technique will be the main focus of the following chapters.

Chapter 3

The EMG-Pressure Relationship of Respiratory Muscles

Both excessive and insufficient respiratory effort often go unnoticed in clinical practice, but were found to be associated with mechanisms of lung and diaphragm injury and to affect long-term outcome in patients under assisted mechanical ventilation. Esophageal pressure, the current clinical gold standard for measuring respiratory effort, requires placement and calibration of a nasogastric catheter, thus, it is only used in very few intensive care units. *Surface electromyography* (surface EMG or sEMG) represents a promising alternative to esophageal pressure for monitoring spontaneous breathing effort. It enables a bedside measurement of the *electrical* activity of respiratory muscles by means of electrodes placed on the thorax of the patient. This chapter contributes to a broader understanding of respiratory sEMG and its relation to muscle pressure P_{mus} measured with an esophageal balloon catheter. The chapter comprises a quantitative analysis of EMG-pressure coupling, including correlation, linearity and channel selection, on a large dataset of patients undergoing bronchoscopy with mechanical ventilation.

Section 3.1 introduces the fundamentals of respiratory electromyography touching both physiological and technical aspects of the measurement. The subsequent section 3.2 introduces the clinical dataset, signal processing chain and statistical evaluation methodology. Finally, in section 3.3, results are reported and then discussed in section 3.4. Some of the results were presented in 2021 [77] at an online meeting of the Pleural Pressure Working Group (PLUG), an international group dedicated to the development of advanced respiratory monitoring techniques. Parts of this chapter were published in two peer-reviewed journal articles [76, 78] first-authored by the author of the thesis. For these publications, he conceived and performed the signal processing/statistical analysis and wrote most of the manuscripts.

3.1 Introduction to respiratory electromyography

The first measurements of electrical respiratory muscle activity were conducted as early as 1951 by Tokizane et al. [222], who used needle electrodes to show that breathing

activity can be monitored by means of individual motor unit's electrical discharges. Since then, many studies have been published on respiratory electromyography ranging from its physiology [24, 25, 54] over technical aspects such as signal acquisition and processing [160, 162, 206, 207] to its clinical application [29, 48, 101, 113, 206]. The field has gained major traction with the development of invasive esophageal measurement techniques in the late 90s and with the implementation of a ventilation mode (NAVA, currently available only in the Servo-i ventilator, Maquet Critical Care, Solna, Sweden) in the 2000s which employs the invasively measured electromyogram for controlling pressure support [140, 205]. *Surface* electromyography of respiratory muscles, which is the main focus of this thesis, is still a relatively unexplored line of research compared to its invasive alternatives. This section provides the necessary background on respiratory electromyography by briefly surveying the key methodologies in the field.

3.1.1 Physiology

Like all skeletal muscles, the diaphragm and the intercostal muscles consist of individual motor units, which are defined as a collection of muscle fibers innervated by one motor neuron [87]. Upon activation by the central nervous system, *action potentials* are propagated along the motor neuron and reach muscle fibres via the neuromuscular junction. When a certain threshold potential at the neuromuscular junction is reached, depolarization of the muscle fibre membrane is triggered and propagated in both directions along the fibre. As a consequence, the muscle fibre contracts. All fibers innervated by one motor neuron are activated and contracted simultaneously. The spatial and temporal summation of all membrane potentials within a motor unit is the so called *motor unit action potential* (MUAP) whose waveform depends on the spatial arrangement of muscle fibres and their specific properties with respect to the conduction of action potentials [155]. A measurement of the electric fields generated by these action potentials is called *electromyogram* (EMG) and is the main focus of this and the following chapters. Depending on the measurement technique, an EMG will usually measure a superposition of many MUAPs within a certain *pickup volume*.

3.1.2 Measurement

The EMG of respiratory muscles, and in particular that of the diaphragm, has been investigated for many decades and has already found application in intensive care medicine [140, 205]. Due to the anatomy of the diaphragm, as a deep and thin muscle, the measurement of its EMG is considered to be notoriously difficult. It requires special measurement techniques and equipment, which will be discussed in the following section.

Esophageal measurement (EAdi)

The first widely used method for measuring the electrical activity of the diaphragm (EAdi) employs an esophageal catheter (with an array of ring electrodes) positioned at the level of the gastroesophageal junction. The esophageal measurement technique was first proposed by Agostoni et al. [9] as well as Petit et al. [163] and allows to acquire the EMG of a sample of motor units in the crural diaphragm. The raw EMG signals are amplified, filtered, and then usually a combined EMG signal is calculated from all individual electrodes, which compensates for diaphragm movement relative to the catheter, refer to [207] for details on signal processing of esophageal EMG data. Finally, an amplitude signal is calculated by applying some form of average filter to the rectified EMG and this amplitude signal is usually referred to as the “EAdi signal”. It was postulated to be a robust measure for phrenic nerve activation, and thus for the *neural drive* of a patient’s respiratory center [24]. EAdi has been used in many studies to assess both the timing and amplitude of respiratory activity [29, 101, 133, 206]. The technique was commercialized through the Servo-i mechanical ventilator (Maquet Critical Care, Solna, Sweden), which can display the EAdi signal in real-time and also allows to use it for ventilation control via the ventilation mode *neurally adjusted ventilatory assist* (NAVA) [140, 205]. In NAVA, patients are given pressure support in proportion to EAdi by a predefined proportionality factor, the so called NAVA level, measured in $\text{cmH}_2\text{O}/\mu\text{V}$. Most of the recent studies on respiratory EMG have employed the Servo-i ventilator and its proprietary filtering algorithms for recording and evaluating EAdi, see for example [29, 48, 101, 206]. The invasively measured Eadi shares the potential disadvantage with P_{es} of relying on the correct positioning of an invasive catheter, which is time-consuming in clinical practice, prone to errors and uncomfortable to patients that are awake.

Skin surface measurement (sEMG)

A completely non-invasive alternative for quantifying respiratory muscle activity is the *surface electromyogram* measured from electrodes on the thorax of the patient, which is the method under consideration in this thesis. The application of this signal modality is relatively straightforward in the clinical arena and comfortable for the patient. The sEMG signal usually constitutes a summation of many MUAPs from different parts of the diaphragm muscle, depending on the electrode positions [120]. Also, in contrast to the invasively measured EAdi signal, it allows to acquire the activity of accessory respiratory muscles (e.g., of the intercostal muscles) [2]. There is no consensus yet on the best number and position of electrodes [120, 123] and most likely, for different applications, different electrode setups will have to be used. Recently, researchers have begun to investigate respiratory sEMG for assessing inspiratory effort [27, 76], respiratory muscle training [122, 169], and patient-ventilator interaction in both intensive care [40, 113] and home mechanical ventilation [54, 168]. Surface EMG usually comes

with a smaller signal-to-noise ratio than the invasively measured EAdi and is more prone to crosstalk from other muscles, including to the ECG artifact from the heart, thus, it requires more sophisticated amplifiers and signal processing [160, 195]. The removal of ECG artifacts is possibly the most crucial preprocessing step and strongly influences the utility of the measured data. The separation of ECG and EMG signals cannot be achieved by linear filtering due to the substantial overlap of the two spectra. Several more advanced filtering techniques have been co-developed by the author of this thesis, including wavelet-based approaches, but shall not be the focus of this thesis. The interested reader is referred to Petersen, Sauer, Graßhoff et al. [162] for a quantitative comparison of filtering algorithms.

3.1.3 EMG-pressure relationship

Extensive work has been done to understand the relationship between EMG signals and the corresponding force/pressure exerted by muscles. The interpretation of muscle force based on the EMG is complicated by the physiological conversion between electrical and mechanical activity in the muscle, which depends on many factors such as muscle length, contraction velocity, muscle fatigue, and the composition of recruited fibres [177]. Beyond that, the measurement of the (surface) EMG signal itself is confounded by many unknowns such as electrode skin contact, volume conduction of the electric field through the tissue and measurement noise [46, 155]. Also, in many cases, the EMG will not be representative of the whole motor unit population but rather of a smaller subset of motor units. Despite these confounding factors, studies have shown for several muscles that the amplitude of an appropriately measured and filtered EMG is approximately proportional to the generated force/pressure [59], that is, it follows

$$\text{EMG}(t) \propto P(t) = \frac{F(t)}{A}, \quad (3.1)$$

where $\text{EMG}(t)$ denotes a rectified and smoothed EMG signal, $P(t)$ denotes the generated muscle pressure, $F(t)$ denotes the muscle force acting on an area A . The assumption of linearity is certainly only a rough heuristic of the underlying physiological processes, but has proven to be very useful in many use-cases [13, 29, 96]. There are, however, important caveats to this assumption: curvilinear EMG-force relations have been observed for certain muscles even under isometric conditions [46], and such an assumption will generally not hold under high activation levels or highly dynamical movements.

Dependency on volume

It has been recognized quite early that the conversion of neural activation to pressure by the diaphragm is mediated by lung volume. To briefly summarize, the diaphragm exerts less pressure at high lung volumes when it is shortened. This effect goes back

to a fundamental property of muscle fibres, namely, their *force-length relationship*: At a certain optimal length, muscle fibres develop their maximum force. When the fiber is shortened or lengthened, this force declines (at constant activation) [171]. The relationship is remarkably constant in vertebrate muscles and was confirmed in vitro on bundles of diaphragm muscle fibres in mammals by McCully and Faulkner [135]. They found a symmetric, bell-shaped curve ranging from roughly 50 % to 150 % of the optimal fibre length. When the diaphragm is analyzed not on a fibre level but as a whole, the force-length relation becomes a pressure-volume relation, which was also investigated in vivo by several researchers. It is worth noting that there is also a nonlinear relationship between diaphragm length and lung volume, refer to Braun et al. [38], and during breathing the diaphragm is exposed to both longitudinal and transversal loading, thus, a direct translation from in vitro to in vivo results is difficult [204]. In [83, 208] phrenic nerve stimulation was used to investigate *transdiaphragmatic pressure* (P_{di}) in vivo at different lung volumes. The results show an approximately linear decline of P_{di} with increasing volume. Smith et al. [208] observed a reduction of pressure generation between functional residual capacity (FRC) and total lung capacity (TLC) by more than 50 % (at the same level of activation). The authors also found that the producible pressure levels off towards the residual volume (RV) and hypothesized that the optimal muscle length, where maximum force can be exerted, lies close to RV. Similar results have been obtained by other authors who studied maximum voluntary efforts at different lung volumes [38, 80]. Based on these findings, a comprehensive model of diaphragm muscle mechanics, incorporating the volume-dependency of pressure generation, was introduced by Younes and Riddle [246].

Respiratory neuromechanical coupling

Several studies have characterized the EAdi-pressure relationship of respiratory muscles in spontaneously breathing patients under assisted mechanical ventilation: surprisingly, in patients under pressure support ventilation EAdi was shown to be very tightly correlated to both P_{mus} and *transdiaphragmatic pressure* P_{di} [24, 29] even without correcting for any confounding factors (such as lung volume, diaphragm length or intrinsic PEEP). Thus, despite the earlier experimental data [38, 135], it was suggested that in practice a linear, patient-specific factor termed *neuromechanical coupling* or *neuromechanical efficiency* would suffice for relating EAdi to the generated muscle pressure (over a limited time-frame). This factor, denoted as K_{EMG} , indicates how much force/pressure the diaphragm can produce on average for a certain amount of neural activation [101]. Beck et al. [24] found that the neuromechanical coupling between EAdi and P_{di} was stable across a large range of pressure support (PS) levels in mechanically ventilated patients. Very importantly, the authors did not find lung volume to be a significant confounder, because the diaphragm was most active in the early phase of the inspiration (when lung volume is low) and intrinsic PEEP (iPEEP)

was relatively constant across PS levels. These results were later confirmed for the EAdi- P_{mus} relationship by Bellani et al. [29] who also evaluated the linearity of the relationship across multiple PS levels by fitting different models concluding that the EAdi- P_{mus} characteristic could be well approximated with a linear model. In regard of these studies, EAdi might represent a viable alternative to P_{es} for assessing patients' respiratory efforts under pressure support ventilation. One important caveat of using a linear neuromechanical coupling is that nonlinear effects have been reported at high activation levels [25], and mild nonlinearities in the EAdi- P_{di} relation have also been observed under isometric conditions in [80]. It is also not yet clear if and how long the neuromechanical coupling index remains stable over time.

To this date, very little work is available with respect to the respiratory sEMG-pressure relationship. A first study by Bellani et al. [27] found a good correlation between processed surface EMG signals and P_{mus} as well as EAdi. However, to reduce noise, their analysis relied heavily on averaging across multiple breaths. The study also proposed that —similarly to EAdi— a patient-specific neuromechanical coupling factor could be calculated, which relates the surface EMG to the generated muscle pressure and would allow to approximate P_{mus} non-invasively. In the following sections a more comprehensive analysis of the sEMG-pressure relation on a larger dataset is provided. As opposed to previous studies on respiratory sEMG, the linearity of neuromechanical coupling and the quantification of efforts are investigated without relying on aggregation of breaths.

3.2 Data analysis

This chapter, building upon the previous sEMG-related study by Bellani et al. [27], presents a quantitative analysis of the relationship between respiratory effort derived from sEMG signals and P_{mus} (as well as the pressure P_{di} generated by the diaphragm) under assisted mechanical ventilation on a cohort of 43 patients scheduled for bronchoscopy. As a measure for the sEMG-derived inspiratory effort, the EMG-time product (ETP) is calculated as the integral of the EMG curve against an adaptive baseline. Moreover, a novel channel selection method to exploit the availability of multiple sEMG measurement channels is proposed and tested.

3.2.1 Study data

In this chapter, data from a study conducted at the department of pneumology, cardiology and intensive care of the Klinikum Konstanz (Konstanz, Germany) by Dr. Waltersbacher and colleagues were analyzed. The protocol was approved by the ethics committee of Witten/Herdecke University (Witten, Germany) on July 12, 2018 (protocol number 137/2017). The study was registered in the German Clinical

Trials Register (DRKS00021524). The data under consideration have previously been published by Graßhoff et al. [76, 78].

Adult patients scheduled for elective bronchoscopy under mechanical ventilation using flexible endotracheal tubes were enrolled for the study, please refer to [76] for detailed enrollment criteria. A total of 62 patients were enrolled, 43 patients were included in the data analysis. The first nine patients were excluded due to technical issues. Additionally, patients were excluded from the analysis if they met one of the following criteria (respective number of patients given in brackets): failure to employ the nasogastric catheter (2), technical recording issues (1), corrupted P_{es} signal, e.g., due to balloon positioning issues or scaling factor > 2 against P_{aw} during airway occlusions (5), and failure of the sEMG cardiac artifact gating algorithm (2). Table A.1 summarizes the characteristics of the included patients.

Measurement setup

After enrollment, a double-balloon catheter (Bösch, Gottenheim, Germany) was filled according to the recommendations in [234] and esophageal/gastric pressures (P_{es} , P_{ga}) were measured by means of pressure transducers at the proximal end of the catheter. The correct positioning of the esophageal balloon was confirmed via the airway occlusion technique [21]. Surface EMG was measured using two pairs of pre-gelled Ag/AgCl electrodes at the following positions: bilaterally at the lower costal margin on the midclavicular line and bilaterally in the second intercostal space on the parasternal line. A common electrode was placed above the sternum. Both channels were measured differentially between electrode pairs and at a sampling rate of 1000 Hz. This electrode configuration allows to sample MUAPs from both the costal diaphragm and the intercostal muscles—similar setups have already been employed in earlier studies [27, 49, 113, 128]. The EMG signals were recorded using a dedicated amplifier and acquisition software provided by Dräger (Drägerwerk AG & Co. KGaA, Lübeck, Germany). The same acquisition system was also used to digitize and record the analog signals from the pressure transducers (P_{es} and P_{ga}) at a sampling rate of 200 Hz (sEMG Base, Drägerwerk AG & Co. KGaA, Lübeck, Germany). The airway flow (\dot{V}) and pressure (P_{aw}) tracings from the Dräger V500 ventilator (Drägerwerk AG & Co. KGaA, Lübeck, Germany) were acquired through the ventilator’s RS232 interface at 100 Hz and then synchronized with the remaining signals.

Study protocol

After patients were enrolled in the study, they were intubated and put on assisted spontaneous ventilation using a sedation protocol with propofol. All patients were sedated to a level of moderate/deep sedation corresponding to level -3 to -4 of the Richmond agitation sedation scale for the study measurement period. At the

beginning of the study protocol, the correct positioning of the esophageal balloon was validated via the Baydur maneuver [21]. Following the initial positioning of the esophageal balloon, a series of spontaneous inspiratory efforts against occluded airways was recorded. Initially, patients were ventilated with continuous positive airway pressure (CPAP) of 5 cmH₂O on a Dräger V500 ventilator (Drägerwerk AG & Co. KGaA, Lübeck, Germany). Patients were then switched to pressure support ventilation (PSV) and underwent a trial of three pressure support levels (5, 10 and 15 cmH₂O) in random order. Throughout the protocol a positive end-expiratory pressure (PEEP) of 5 cmH₂O was used.

3.2.2 Signal processing

Volume trend removal

A continuous lung volume signal was generated by running integration over the airflow measurement \dot{V} from the ventilator. Due to small flow sensor inaccuracies or leakages, this integration usually results in long-term drifts in the estimated volume signal. To mitigate volume drift, a trend removal algorithm was applied to the end-expiratory volumes (assuming that on average over multiple breaths the end-expiratory volume is zero). To this end, the end-expiratory volumes were filtered via a moving median, a continuous baseline was calculated by interpolating between these filtered points, and then this baseline was subtracted from the volume signal. The baseline-corrected volume signal was denoted as V .

Esophageal pressure processing

Esophageal pressure signals are often strongly disturbed by cardiogenic pressure artifacts, which complicate interpretation of the underlying respiratory activity. Cardiogenic artifacts were removed from P_{es} signals using a template subtraction method proposed by Graßhoff et al. [75]. Similar artifacts were also visible in the gastric pressure (P_{ga}) signals, which is why the filtering technique was applied there as well. This technique calculates a template for the artifact waveform, which is then subtracted from each cardiac cycle. The method successfully recovered even strongly disturbed recordings and substantially increased the signal quality of pressure signals. Any remaining artifacts were then removed from P_{es} and P_{ga} using a low-pass filter (5th order Butterworth filter with 7 Hz cutoff). The transdiaphragmatic pressure signal was then calculated as the difference between the filtered P_{ga} and P_{es} signals. As described in section 2.2.4, the time course of the respiratory muscle pressure P_{mus} was determined as the difference between filtered esophageal pressure P_{es} and the elastic recoil of the chest wall $P_{cw}(t) = E_{cw} \cdot V(t)$, where V is the trend-corrected volume signal. The determination of the chest wall elastance E_{cw} on spontaneously breathing patients is challenging, refer to Graßhoff et al. [72]. On the data under consideration,

as there were no phases without residual patient activity, E_{cw} was calculated under the highest PS level, where spontaneous breathing effort was lowest. To this end, breaths were annotated, where both the end-inspiratory and the end-expiratory point were not affected by spontaneous efforts. The validity of the selected points was verified by careful manual inspection of the P_{es} and P_{ga} waveforms. The chest wall elastance E_{cw} was then calculated as the slope between end-expiratory and end-inspiratory points in the Campbell diagram and averaged over multiple semi-passive breaths. Finally, as the absolute values of P_{es} and P_{ga} strongly depend on the filling volume and positioning of the pressure balloons, only relative pressure swings were evaluated against baselines

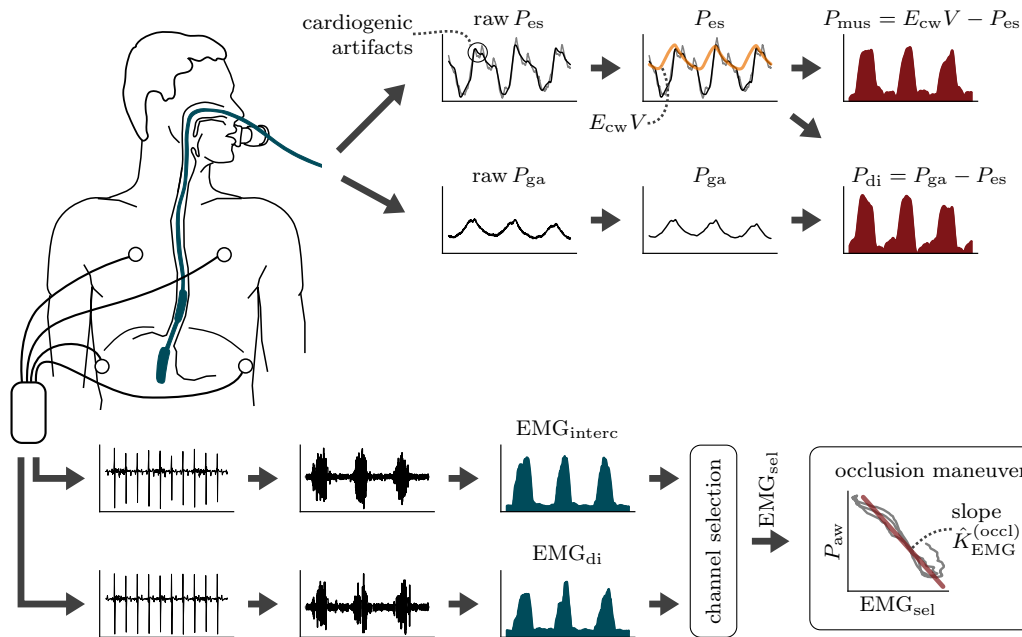


Fig. 3.1: An overview of the signal processing pipeline used in this chapter. Esophageal and gastric pressure were measured by means of a double balloon catheter. Cardiogenic artifacts were removed from both pressure signals using a template subtraction method. Respiratory muscle pressure P_{mus} was calculated as the difference between the elastic recoil pressure of the chest $E_{cw}V$ (plotted as an orange line) and P_{es} . Transdiaphragmatic pressure P_{di} was calculated as the difference between P_{ga} and P_{es} . The sEMG of intercostal muscles and the diaphragm was measured using two pairs of bilateral electrodes. ECG artifacts were removed by means of a simple gating technique, allowing to calculate the root-mean-square envelopes EMG_{di} and EMG_{interc} . The more informative of the two channels, denoted by EMG_{sel} was automatically selected and fitted to the airway pressure P_{aw} over the course of multiple occlusion maneuvers to derive the neuromechanical coupling factor $\hat{K}_{EMG}^{(occl)}$. Figure modified and extended from Graßhoff et al. [76].

$P_{es,0}(t)$ calculated during passive expirations. Figure 3.1 provides a graphical overview of the main processing steps conducted to estimate P_{mus} and P_{di} based on P_{es} .

Next, the available end-expiratory occlusions were used to check and correct the scaling of both P_{es} and P_{mus} . During occlusions, as flow and volume are zero, the pressure drop in P_{aw} can be assumed to be equal to P_{mus} (and to the relative drop in P_{es}), which allows assessing possible scaling errors in P_{es} , e.g., due to catheter positioning errors [21]. Thus, following the balloon positioning procedure, a correction factor $\hat{K}_{P_{es}}^{(occl)}$ was determined by fitting the esophageal pressure waveform to the airway pressure waveform over the course of multiple subsequent occlusions by means of linear regression. In practice, small scaling errors remained even after proper positioning, i.e., the factor $\hat{K}_{P_{es}}^{(occl)}$ was often slightly larger than one. To correct for these remaining errors, the pressure waveform was scaled with the factor $\hat{K}_{P_{es}}^{(occl)}$ determined during occlusions, that is, in this and the following chapters $P_{mus}(t) = \hat{K}_{P_{es}}^{(occl)} (P_{cw}(t) - P_{es}(t) + P_{es,0}(t))$ is used as a measure for the respiratory muscle pressure, which is a slightly modified variant of equation (2.58).

Surface EMG processing

As discussed in section 3.1.2, thoracic sEMG signals are highly susceptible to crosstalk from other muscles. One of its main contaminants is the cardiac artifact, which was removed using a simple gating technique [162]. R peaks within the sEMG signals were detected and then segments affected by the QRS complex were discarded. The amplitude of both EMG channels was then calculated using a moving 250 ms root-mean-square filter. These surface EMG *envelopes* often have significant offsets in the order of several μV due to measurement noise [41]. The offsets can be observed between efforts, when the patient is (mostly) passive. To account for possible changes of the offset (e.g., due to fluctuations of the noise) a time-varying baseline was calculated and subtracted from the envelope. Empirically, the first tercile of envelope values within a moving 5 s window was found to provide a robust estimate for the signal's baseline and was subsequently used for both channels—the resulting corrected envelope signals of the diaphragm and intercostal EMG were denoted as EMG_{di} and EMG_{interc} , respectively. After subtraction of baselines, the EMG envelopes were indeed roughly zero when patients were inactive.

3.2.3 Quantitative analysis

Segments affected by artifacts due to irregular patient behavior (e.g., ventilator fighting or coughing) were manually excluded. Likewise, measurement errors in P_{es} , such as artifacts due to peristalsis etc., were manually marked as invalid and excluded from analyses involving metrics derived from P_{es} .




Fig. 3.2: An illustration of different metrics used to quantify inspiratory muscle effort. The pressure–time products PTP_{mus} and PTP_{di} are calculated as the breath-wise area under the P_{mus} and P_{di} curve, respectively. Similarly, the EMG–time products ETP_{di} and $\text{ETP}_{\text{interc}}$ correspond to the area under the EMG_{di} and $\text{EMG}_{\text{interc}}$ envelope, respectively. Availability of both channels is leveraged by automatically selecting one of the two EMG-derived time–products in each patient, leading to the ETP_{sel} metric.

Inspiratory effort metrics

As a metric for inspiratory effort, the pressure–time product (PTP) was used because it has been suggested to better represent patient activity than WOB when little or no volume is generated. As a first step towards calculating PTP, the data were segmented by detecting inspiratory efforts. To this end, an automatic threshold-based detector similar to the one proposed by Sinderby et al. [206] was applied directly to the P_{mus} signal. Note that a segmentation based on P_{mus} is generally preferable to a flow-based detection because the former better accounts for missed efforts and late triggers. PTP metrics were calculated in a breath-wise fashion as the area under the curve for both P_{mus} and P_{di} on the detected segments and they were denoted by PTP_{mus} and PTP_{di} , respectively. Any efforts exceeding $\text{PTP}_{\text{mus}} = 20 \text{ cmH}_2\text{O s}$ were excluded from further analysis because such strong efforts do not fall into the regular range of breathing activity that was targeted in this study. Finally, as a metric for breathing effort derived from sEMG the area under the curve was calculated for the (baseline-adjusted) EMG_{di} and $\text{EMG}_{\text{interc}}$ envelopes, leading to EMG–time product (ETP) values measured in $\mu\text{V s}$ and denoted by ETP_{di} and $\text{ETP}_{\text{interc}}$, respectively, refer to figure 3.2 for an overview.

Channel selection

The two EMG channels often do not carry the same amount of information, because they measure MUAPs from different muscle groups and therefore capture different parts of the patient’s respiratory activity. The EMG channels also often contain different levels of measurement noise. It was already proposed by Bellani et al. [27] that the availability of multiple EMG channels could be leveraged by summing the corresponding envelopes, which however did not improve results in their study. In this work, another direction was pursued by investigating a simple, automatic method

for selecting the EMG channel that contains the best signal quality at the lowest level of noise. To this end, the signal-to-noise ratio (SNR) of the two channels was approximated by comparing the maximum amplitude reached during inspirations with the level of baseline noise. Thus, the distribution of (not baseline-corrected) envelope values was analyzed over a moving 10s window using the ratio

$$\text{SNR}_{\text{approx}} = \frac{Q_3}{Q_1} \quad (3.2)$$

where Q_3 denotes the third quartile and Q_1 denotes the first quartile of values. For each patient, the envelope with higher median $\text{SNR}_{\text{approx}}$ was selected and used as the measure for the electrical breathing activity. The selected envelope with higher (approximate) SNR was denoted as EMG_{sel} and the corresponding EMG–time product as ETP_{sel} .

Neuromechanical coupling analysis

The next part of the analysis was concerned with characterizing the sEMG- P_{mus} relationship. As discussed above, the work by Beck et al. [24] and Bellani et al. [29] suggests that a linear coupling is viable during normal assisted mechanical ventilation, which was modeled via

$$P_{\text{mus}}(t) = K_{\text{EMG}} \cdot \text{EMG}(t) + P_{\text{bias}}, \quad (3.3)$$

where K_{EMG} is the neuromechanical coupling index and P_{bias} is a constant bias term. Thus, P_{bias} represents systemic offsets that the EMG envelope might have against the muscular pressure curve P_{mus} .

To estimate K_{EMG} , both sides of equation (3.3) were integrated, leading to

$$\text{PTP}_{\text{mus}}(i) = K_{\text{EMG}} \cdot \text{ETP}(i) + P_{\text{bias}} \cdot T_{\text{insp}}(i), \quad (3.4)$$

which represents the EMG-pressure relation in terms of time product values and can be solved for the neuromechanical coupling K_{EMG} and the constant bias term P_{bias} using ordinary least squares regression on a set of breaths indicated by i . The value T_{insp} is the length of the detected effort, which accounts for the integration of bias over the duration of each effort. The main advantage of using the integrated form (3.4), besides a lower computation time, is that all included breaths have an equal weight, independent of their length. An important focus of this study was to analyze the linearity of the EMG- P_{mus} relationship, which was done by also fitting the model

$$\text{PTP}_{\text{mus}}(i) = K_{\text{EMG},1} \cdot f(\text{ETP}(i)) + K_{\text{EMG},2} \cdot \text{ETP}(i) + P_{\text{bias}} \cdot T_{\text{insp}}(i), \quad (3.5)$$

where the function f represents a nonlinear characteristic. As a nonlinear characteristic a quadratic and a square root term were tested, corresponding to an attenuation or amplification of the coupling at higher activation levels.




Fig. 3.3: Example of an occlusion maneuver. The left panel shows the time-course of airway pressure (P_{aw}) and the EMG envelope. The flow of air was interrupted in between the dashed lines—thus, the muscle pressure P_{mus} is directly reflected in P_{aw} . The right panel shows a plot of EMG against P_{aw} over the course of the occlusion. The slope of the regression line corresponds to the neuromechanical coupling index $\hat{K}_{EMG}^{(occl)}$.

Occlusions

Several studies have demonstrated that the neuromechanical coupling can be identified without P_{es} using end-expiratory occlusion maneuvers with both EAdi [29, 101] and sEMG [27]. When the flow and volume are zero, any spontaneous efforts by the patient will be directly visible as a pressure drop in P_{aw} , allowing to analyze its relation to EMG and approximate K_{EMG} completely non-invasively. Thus, the selected EMG channel was fitted to P_{aw} across multiple occlusion maneuvers using linear regression and the slope, denoted as $\hat{K}_{EMG}^{(occl)}$, was used as a surrogate for the ‘true’ neuromechanical coupling K_{EMG} . Figure 3.3 shows an illustrative excerpt of P_{aw} and EMG signals during an occlusion maneuver. Bellani et al. [29] found that the parameter $\hat{K}_{EMG}^{(occl)}$ was well correlated with but higher than K_{EMG} . The authors proposed that $\hat{K}_{EMG}^{(occl)}$ has to be multiplied with a constant correction factor α to account for the deviation:

$$K_{EMG} \approx \alpha \cdot \hat{K}_{EMG}^{(occl)}. \quad (3.6)$$

Bellani et al. [27, 29] gave numerical values for α and argued that the systematic overestimation can be prescribed to the more favorable configuration of the diaphragm at end-expiration, which is a consequence of the force-length relationship discussed in section 3.1.3. Values for α were determined on the herein analyzed study data and compared to the values reported in the earlier studies. It is worth mentioning that the occlusion-based method for analyzing neuromechanical coupling is still a subject of an ongoing scientific debate: a recent EAdi-based study by Jansen et al. [101] found its repeatability to be low as there was an unacceptably high variation between coupling indices derived from subsequent occlusion maneuvers. The authors attributed these problems (at least in part) to shortcomings of the proprietary EAdi filtering algorithms.




Fig. 3.4: An excerpt of the data used in this thesis and previously published by Graßhoff et al. [76]. The plot shows measurements of airway pressure (P_{aw}) and flow (\dot{V}) acquired from the ventilator, and the volume signal (V). The raw esophageal pressure (P_{es}) curve before removal of artifacts is plotted in grey and the elastic recoil of the chest wall ($E_{cw}V$) in orange. PTP values for total respiratory muscle pressure (P_{mus}) and transdiaphragmatic pressure (P_{di}) are represented by shaded areas. The EMG channel selected as the more informative one is denoted by EMG_{ssel} and plotted as a green line. ETP corresponds to the shaded area and was calculated against an adaptive baseline plotted in black.

Statistics

Results are expressed as mean \pm standard deviation. Correlation between variables was quantified by means of Pearson's correlation coefficient r . A two-tailed Wilcoxon signed-rank test was used to test for differences between the channel with high SNR, which was selected as the more informative channel, and the other channel with low SNR.

3.3 Results

As already mentioned, a total of 43 patients were analyzed. In one patient, P_{ga} was not available, which is why this patient was excluded from any analysis relying on P_{di} . In two patients, no sufficiently long occlusions (> 0.35 s) were available. These data were only used in the correlation analysis and excluded from any analysis relying

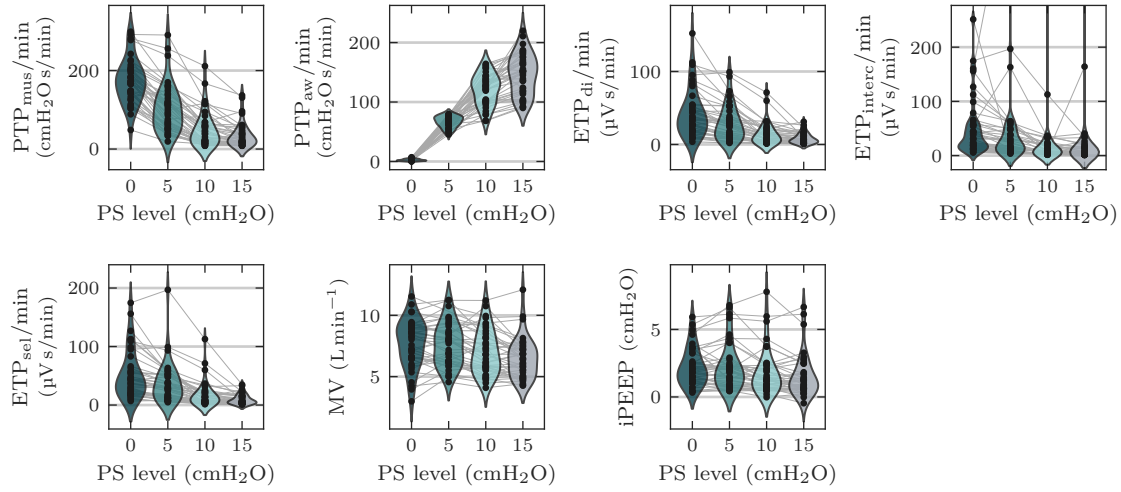


Fig. 3.5: Distribution of muscle and airway pressure–time products (PTP_{mus} and PTP_{aw}), EMG–time products of the diaphragm, intercostal and selected channel (ETP_{di} , ETP_{interc} and ETP_{sel}), minute ventilation (MV) and intrinsic positive end-expiratory pressure (iPEEP) by pressure support (PS) level. PTP/min and ETP/min values were calculated by summing all breaths within each PS level and dividing by the length of the segment. MV is the total volume of air inhaled in one minute, iPEEP is the positive alveolar pressure remaining at the end of an expiration due to incomplete lung deflation. Each point represents one patient and pressure support level. In the ETP_{interc} distribution plot, three outliers $> 400 \mu\text{V s}/\text{min}$ are not shown within the plotting range.

on the absolute value of P_{es} such as the calculation of neuromechanical coupling. An illustrative excerpt of the analyzed data is shown in figure 3.4.

The recording length per patient was 23 ± 4 min. Across all datasets a total of 19540 inspiratory efforts were detected and analyzed (454 ± 137 breaths per patient). The length of analyzed efforts was 0.9 ± 0.3 s. The number of analyzed occlusions per patient was 4.8 ± 1.5 . The correction factor $\hat{K}_{P_{\text{es}}}^{(\text{occl})}$ determined from occlusions was close to one (1.18 ± 0.18) indicating the validity of the esophageal balloon positioning and filling.

Figure 3.5 provides distributions of relevant parameters against PS levels: all effort metrics, including those derived from the sEMG, increased with lower PS. Intrinsic PEEP was low at $2.1 \pm 1.4 \text{ cmH}_2\text{O}$ (measured from the positive deflection of P_{mus} at the point of zero flow). Only two patients with COPD had $iPEEP > 5 \text{ cmH}_2\text{O}$.

Table 3.1: Breath-wise Pearson correlation coefficient r (mean \pm standard deviation) between effort metrics derived from esophageal/gastric pressure and sEMG, $n = 43$ ($n = 42$ when P_{di} is involved).

	ETP_{interc}	ETP_{sel}	PTP_{mus}	PTP_{di}
ETP_{di}	0.74 ± 0.27	0.89 ± 0.16	0.84 ± 0.16	0.84 ± 0.16
ETP_{interc}	–	0.86 ± 0.28	0.79 ± 0.25	0.77 ± 0.26
ETP_{sel}	★	–	0.87 ± 0.09	0.86 ± 0.10
PTP_{mus}	★	★	–	0.97 ± 0.05

PTP_{mus} and PTP_{di} : pressure–time products of muscular and trans-diaphragmatic pressure; ETP_{di} , ETP_{interc} and ETP_{sel} : EMG–time products of the diaphragmatic, intercostal and selected EMG channel; the star symbol (★) marks metrics given due to symmetry;




Fig. 3.6: Distribution of neuromechanical coupling indices K_{EMG} and biases P_{bias} derived from equation (3.4) for different EMG channels, $n = 41$. Each point represents one patient and EMG channel.

Correlation and channel selection

The breath-wise correlation between effort metrics is provided in table 3.1, revealing a high correlation between the different ETP measures and PTP_{mus} as well as PTP_{di} . To assess the influence of the channel selection technique, differences between the selected channel and the respective other channel with lower (approximate) SNR were tested for, revealing that the ETP of the selected channel had a higher correlation to PTP_{mus} ($p = 0.029$). Overall, the intercostal channel was selected as the more informative of the two channels 25 out of 41 times. The average SNR of the intercostal channel was 1.87 ± 0.95 and that of the diaphragm channel 1.73 ± 0.56 . For the selected channel a high correlation of 0.87 ± 0.09 between ETP_{sel} and PTP_{mus} was observed, refer to table 3.1.




Fig. 3.7: The sEMG- P_{mus} relation for three selected patients. The scatter plot shows diaphragmatic EMG-time product values (ETP_{di}) plotted against muscular pressure-time product values (PTP_{mus}) corrected by the bias term $P_{\text{bias}}T_{\text{insp}}$ from equation (3.4). Each point represents one inspiratory effort. As expected, the slope of the ETP- PTP_{mus} relation, i.e., neuromechanical coupling, varied widely across patients, ranging from $< 1.9 \text{ cmH}_2\text{O}/\mu\text{V}$ for patient 1 to $> 6.1 \text{ cmH}_2\text{O}/\mu\text{V}$ for patient 3. The ETP_{di} - PTP_{mus} correlation was $r = 0.94$ (patient 1), $r = 0.95$ (patient 2) and $r = 0.86$ (patient 3).

Neuromechanical coupling

The neuromechanical coupling K_{EMG} and bias parameter P_{bias} were calculated via equation (3.4) for the three ETP measures (ETP_{di} , $\text{ETP}_{\text{interc}}$, ETP_{sel}). Results are reported in figure 3.6. There was a large variance in neuromechanical coupling K_{EMG} across patients for all sEMG channels, which is in line with earlier results by Bellani et al. [27, 29] and may be explained by inter-patient variability in muscle geometry, volume conduction and skin-electrode impedances. Figure 3.7 provides a scatter plot of the ETP_{sel} - PTP_{mus} relation for three patients, visualizing the large range of neuromechanical coupling. There was weak positive correlation between K_{EMG} of the diaphragm and the intercostal channel ($r = 0.38$, $p = 0.014$). In the channel with higher SNR, both K_{EMG} and P_{bias} were smaller than in the respective other channel with lower SNR ($p = 0.27$ and $p < .001$, respectively). These results confirm the benefits of the proposed channel selection technique, which is why in this and the following chapters EMG_{sel} will be used as the overall measure for electrical patient activity.

Next, the linearity of the sEMG-pressure relationship was analyzed. Models with an additional nonlinear term were fitted to the ETP_{sel} - PTP_{mus} data, refer to equation (3.5), and compared to the linear model in equation (3.4). For the nonlinear characteristic f a quadratic and a square root term were tested. For all three models




Fig. 3.8: Scatter plot of neuromechanical coupling indices determined invasively from esophageal pressure (K_{EMG}) and non-invasively using occlusions ($\hat{K}_{EMG}^{(occl)}$). Each point represents one patient ($n = 41$).

the adjusted coefficient of determination r_{adj}^2 was calculated. There was a significant ($p < .001$) but small difference between the linear model ($r_{adj}^2 = 0.79 \pm 0.13$) and the two nonlinear models ($r_{adj}^2 = 0.80 \pm 0.12$ and $r_{adj}^2 = 0.81 \pm 0.10$), which demonstrates that the sEMG-pressure relationship can be well approximated using a linear model across multiple PS levels.

Finally, the possibility to determine the neuromechanical coupling index completely non-invasively using end-expiratory occlusion maneuvers was investigated. Figure 3.8 provides a plot of the parameter $\hat{K}_{EMG}^{(occl)}$ calculated from occlusions against K_{EMG} calculated from P_{mus} . There was a high correlation between the two parameters ($r = 0.95$, $p < .001$, slope = 0.84) and, similar to the study by Bellani et al [29], on average the ratio $K_{EMG}/\hat{K}_{EMG}^{(occl)}$ was smaller than one (0.81 ± 0.23), indicating a systematic overestimation when determining neuromechanical coupling indices during occlusions. In agreement with Bellani et al. [27, 29], in this thesis, $\alpha = 0.8$ will be adopted as a correction factor to account for the deviation (Bellani et al. used a factor of 1/1.25 in their studies).

3.4 Discussion

The results presented in this chapter contribute to a better understanding of the relationship between respiratory sEMG signals and muscle pressure. Based on the findings from this chapter, respiratory sEMG might be a viable alternative to esophageal pressure for measuring the breath-wise inspiratory effort of patients under assisted mechanical ventilation. In contrast to esophageal pressure and EAdi, sEMG can be measured non-invasively in a straightforward manner, allowing to easily assess the activity of patients. The results of the analysis can be summarized as follows. Firstly, EMG-time

product (ETP) values and pressure–time product (PTP) values of invasively measured P_{mus} and P_{di} were highly correlated on a large cohort of patients scheduled for bronchoscopy. Secondly, the availability of multiple channels could be leveraged via a channel selection method, improving correlation to PTP metrics. Thirdly, the EMG– P_{mus} relationship could be well approximated using a linear neuromechanical coupling index.

The results are consistent with earlier EAdi-based studies by Beck et al. [24] who observed a high correlation ($r = 0.84 \pm 0.12$) between invasively measured EAdi and P_{di} across multiple pressure support levels. Similar results have been reported by Bellani et al. [29] for the correlation between EAdi and P_{mus} . Prior to the herein conducted study, results on sEMG were quite limited except for a first study by Bellani et al. [27], which, however, relied on averaging of multiple breaths. In this study, for the first time, a high correlation ($r = 0.87 \pm 0.09$) has been reported between inspiratory efforts derived from surface EMG and P_{mus} without relying on aggregation of breaths. Thus, this work provides first evidence for the feasibility of quantifying inspiratory efforts on a breath-by-breath basis using surface EMG.

The results also showed that the ETP–PTP relation could be well approximated by a linear model and that a nonlinear model did not provide a substantially better representation of the data. Thus, these results support the simple model $\text{EMG} \propto P_{\text{mus}}$ for relating electrical to pneumatic muscle activity, which will be employed in the following chapters. For the proportionality factor, termed neuromechanical coupling, a large variance across patients was observed, which also coincides with earlier reports by Bellani et al. [29]. Thus, the absolute value of the measured sEMG amplitude can correspond to a wide range of muscle pressures and should really only be interpreted in concert with its neuromechanical coupling. It is not yet clear if the neuromechanical coupling is stable over longer periods of time. The data from this chapter indicate that K_{EMG} can be kept constant over periods of approximately 20 minutes, but during longer measurements changes in skin conductance (e.g., due to sweat) may influence the coupling.

This chapter has also analyzed the occlusion-based approach for measuring neuromechanical coupling when P_{es} is not available. The results corroborate findings by Bellani et al. [29], who found a systematic overestimation of the coupling index determined during occlusions. This effect can be attributed to the force-length relationship of muscle fibres: when the volume is low, a higher force/pressure can be generated than at higher lung volumes when the diaphragm muscle fibres are shortened [38, 83]. Based on this chapter’s findings and those by Bellani et al. [29], there is strong evidence for using a correction factor of $\alpha = 0.8$ on most patients. This gives rise to a very simple approximation of muscle pressure, namely $P_{\text{mus}}(t) \approx \alpha \hat{K}_{\text{EMG}}^{(\text{occl})} \text{EMG}_{\text{sel}}(t)$, which will be further explored in the next chapter.

For the evaluation of respiratory sEMG recordings, as for any other physiological measurement, a high SNR is essential. The SNR of respiratory sEMG depends on a

number of factors: (1) the breathing pattern of the patient, e.g., presence of abdominal or thoracic breathing, (2) the volume conduction of MUAPs through the tissue to recording electrodes, and (3) the amount of physiological crosstalk and measurement noise. A good SNR is obtained through a favorable combination of these conditions, i.e., strong muscle activation and good transmission of the signal at a low level of noise. The previous study on respiratory sEMG by Bellani et al. [27] relied on channel averaging to increase SNR, which however did not seem to improve results. A novel channel selection technique was proposed, which automatically detects the channel with more favorable measurement conditions using a simple SNR heuristic. It was shown that the proposed method attains a substantially higher correlation of ETP to PTP_{mus} (as well as PTP_{di}). These results demonstrate the merits of measuring multiple EMG channels to capture MUAPs from different muscles and then making a manual or automated channel selection based on the measurement conditions.

It is currently an open question whether the comparison of multiple respiratory sEMG channels provides further physiologically relevant information, e.g., whether a precise quantification of the activity in individual muscle groups is feasible. It is well known that under high respiratory stress, intercostal and other accessory muscles become more active [11]. A quantification of this effect could be very valuable for diagnostic purposes. In the analyzed study data, a high correlation between the ETP values of the diaphragm and the intercostal channel was observed ($r = 0.74 \pm 0.27$), indicating that muscle activity did not shift from the diaphragm to accessory muscles throughout the protocol. As the included patients did not have respiratory failure, it is indeed unlikely that recruitment would change substantially even under CPAP. This is also confirmed by the high correlation between PTP_{di} and PTP_{mus} . In many patients, large signal amplitudes in the intercostal channel were found and its estimated SNR was higher than that of the diaphragm channel. Consequently, the intercostal channel was selected more often as the more informative one. This finding indicates that relevant information about the total breathing activity of patients is contained in the intercostal channel and that it is worth to be measured. Recent work suggests that the availability of both intercostal and diaphragmatic channels is highly useful for conducting an algorithmic muscle source separation, refer to Sauer et al. [195] for further details.

When analyzing the EMG-pressure relationship, it must be considered that the measurement of esophageal pressure is itself subject to substantial measurement errors, resulting from cardiogenic artifacts, peristalsis or inaccurate balloon positioning/filling. In this work, special care was taken to mitigate these errors. The balloon filling volumes were adjusted using results from the benchmark study by Waltersbacher et al. [234]. A templates subtraction method proposed by Graßhoff et al. [75] was used to reduce cardiogenic artifacts. Moreover, a correction factor, denoted as $\hat{K}_{P_{\text{es}}}^{(\text{occl})}$, was calculated by fitting P_{es} to the P_{aw} pressure drop during end-expiratory occlusion maneuvers.

The P_{mus} signal was then scaled by $\hat{K}_{P_{\text{es}}}^{(\text{occl})}$ prior to the analysis to account for the deviation. This scaling correction helps to reduce positioning errors, but it should generally be used with caution and only in cases where the scaling factor is already close to one.

The study has several limitations that will be discussed in the following. None of the patients had signs of severe respiratory failure and had been ventilated over a prolonged period of time. The transfer of results to intensive care patients is therefore challenging and requires further research. Moreover, obesity is a very common comorbidity in ventilated patients, but was an exclusion criterion in this study. Overall, more study data is needed in the intensive care setting on severe cases, including severe COPD and ARDS patients. In the data analyzed here, a very wide range of patient activities have been included (from four PS levels), which has to be considered when comparing the reported correlation to other studies. Bellani et al. [29] studied a smaller range of activities (± 4 cmH₂O from baseline PS), but Beck et al. [24] used a similar range in their study, also including CPAP ventilation.

Chapter 4

Model-based Inspiratory Effort Estimation using Surface EMG

Having established the close relationship between respiratory sEMG and generated pressure P_{mus} in the previous chapter, this chapter adopts a model-based approach to estimate respiratory effort. The herein proposed method combines the sEMG signals with pneumatic measurements from the ventilator to determine the lung resistance and elastance as well as the neuromechanical coupling of respiratory muscles. Using the estimated parameters, a P_{mus} signal can be reconstructed solely from the pneumatic ventilator data using the equation of motion. By combining all available measurement data (electrical and pneumatic) the resulting estimate is less affected by noise in the individual signals and a higher estimation quality can be attained—this concept of combining multiple measurements for estimating a latent variable is referred to as *sensor fusion*. The herein proposed method would allow to non-invasively estimate P_{mus} at the bedside, without requiring any occlusions or other disruptive maneuvers, and thus, has high potential to benefit the diagnostics of mechanically ventilated patients.

This chapter begins with an overview of previous work on *model-based estimation* of respiratory effort (section 4.1). The following section 4.2 introduces the novel model-based estimation framework and also very briefly reiterates the occlusion-based estimation from the previous chapter. Next, in section 4.3, the validation procedure is presented and results are reported in section 4.4. Finally, the chapter is concluded by a brief discussion in section 4.5. Parts of this chapter have been published in a peer-reviewed journal article [78], for which the author of this thesis conceived the methodology, conducted the data analysis and wrote the majority of the manuscript. This work has also resulted in the co-authorship of two patent applications [55, 56].

4.1 Introduction and previous work

The measurement of a mechanically ventilated patient’s own inspiratory effort (in terms of pressure or pneumatic energy expenditure) is a challenging task. Establishing a viable alternative to measuring P_{mus} from esophageal pressure is difficult and has

been the topic of many publications [29, 32, 66, 157, 200, 247]. To this date, the problem remains unsolved, though, significant progress has been made.

Bellani et al. [29] were the first to demonstrate that EAdi, when multiplied with a neuromechanical coupling index, could be used to approximate the absolute value of P_{mus} and that a simple airway occlusion would be sufficient as a calibration maneuver. Similar to EAdi, the occlusion-based technique was also proposed for approximating P_{mus} from sEMG [27, 76]. In these recent studies, a linear neuromechanical coupling index was calculated during airway occlusions and P_{mus} was approximated by scaling the respiratory sEMG. Such approaches rely heavily on the proportionality of the EMG-pressure relationship, refer to chapter 3 for a detailed analysis.

Despite its simplicity, the occlusion-based calculation of P_{mus} from EAdi and sEMG has some practical drawbacks. Firstly, occlusion maneuvers necessarily interrupt the patient's natural breathing pattern and, therefore, their use should be minimized. Secondly, after the initial maneuver, the P_{mus} estimate is formed only from EMG data without incorporating the pneumatic measurements from the ventilator in any way. Therefore, the quality of the estimated P_{mus} depends strongly on the validity of the airway occlusion maneuver and on the signal-to-noise ratio of the measured EMG. Any disturbances or artifacts in the EMG envelope will significantly distort the estimation results. A recent study on EAdi by Jansen et al. [101] found the repeatability of the occlusion approach to be low. It appears to be necessary to perform multiple subsequent occlusion maneuvers to mitigate these issues (but this in turn entails a stronger disruption to the patient).

Another much pursued direction and topic of this chapter is the *model-based estimation* of respiratory effort [114, 200, 241]. This approach usually employs the *equation of motion* (EqM)

$$P_{\text{aw}}(t) = E \cdot V(t) + R \cdot \dot{V}(t) - P_{\text{mus}}(t) + P_0 \quad (4.1)$$

introduced in section 2.2.3 as a lung mechanics model. When estimates of respiratory parameters are available, the latent variable $P_{\text{mus}}(t)$ can be calculated simply by rearranging equation (4.1). However, given measurement of $P_{\text{aw}}(t)$, $V(t)$, $\dot{V}(t)$ at different time points t_1, \dots, t_N , it is easily apparent that the problem of solving for the unknown parameters is *underdetermined*, because there are infinitely many solutions for R , E , P_0 and $P_{\text{mus}}(t_{1:N})$ [228]. From a technical standpoint, the problem of simultaneously estimating respiratory parameters and P_{mus} can be understood as a *joint parameter and latent variable estimation* task.

A solution can be obtained when $P_{\text{mus}}(t) = 0$, for example under mandatory ventilation, because in that case equation (4.1) is in linear regression form and can be solved for the parameters with appropriate methods. For passive patients, Wald et al. [233] were the first to propose a computational solution to respiratory parameter estimation using the least-mean squares algorithm. Based on this, many similar studies have been conducted that rely on fitting the linear EqM to passive patients using *ordinary*

least squares (OLS) or *recursive least squares* (RLS) regression [17, 100, 121, 224, 232]. Recently, more complex, nonlinear models have also been proposed for respiratory system identification, refer to [119, 145]. The potential of model-based methods for lung parameter estimation has been widely recognized, yet, some authors have urged caution: respiratory parameter identification using the least squares approach appears to be highly sensitive to small phase delays and noise [215] and the standard EqM does not represent typical pathologies such as flow limitations [91, 126, 232]. Yamada et al. [241] were the first who proposed to determine lung mechanical parameters (E and R) when the patient is passive and then to use these parameters for calculating P_{mus} during assisted ventilation from the available non-invasive measurements of airway pressure and flow. The authors reported a high level of agreement to esophageal pressure albeit on a very small cohort of patients. Similar approaches were also investigated by other authors, see [114, 200]. Because these approaches necessitate to completely suppress respiratory muscle activity throughout the parameter estimation period, they will not be applicable in most cases.

Due to the underdetermined nature of the problem, fewer attempts have been made to estimate lung mechanical parameters and breathing activity simultaneously during assisted mechanical ventilation. Researchers have proposed different ways to ensure identifiability. Two main directions can be discerned: (1) making *additional assumptions* about the shape of P_{mus} and (2) using ventilatory maneuvers to *excite the system* and to reduce the linear dependency of regressors.

Regarding the first direction, some studies have investigated to estimate lung parameters only in selected regions, where patient activity is presumably low and $P_{\text{mus}} \approx 0$ can be assumed, see for example [94, 158, 229]. A special case of this approach is to estimate lung parameters under an increased PS level, where respiratory drive is reduced, which was proposed by Iotti et al. [100] and recently tested by two studies [147, 180], but both concluded that the accuracy of the derived P_{mus} curve was not satisfactory. Other authors have used physiological knowledge about the shape of respiratory efforts to regularize the estimated P_{mus} waveform. This was demonstrated by incorporating P_{mus} template waveforms [174, 182] or physiologically motivated constraints into the parameter estimation [176, 228] and will be discussed more comprehensively in the next chapter 5. All of the approaches making assumptions about P_{mus} to estimate parameters have so far only been demonstrated on small patient cohorts and it is still unclear if these assumptions will hold in real ventilation scenarios.

The second direction (using maneuvers) was most prominently pursued by Younes et al. [248, 249], who proposed to use occlusion and pressure pulses to identify lung parameters and to derive a P_{mus} signal during assisted ventilation. The lung elastance was determined using end-inspiratory occlusions while the lung resistance was calculated from brief airway pressure pulses. The method was tailored to the proportional assist ventilation (PAV) mode and commercialized in the PAV+ ventilation software (Puritan Bennett, Medtronic). The approach is promising but of course the maneu-

vers necessarily interrupt the normal ventilation, potentially disturbing the patient. Later, Younes et al. presented an estimator that uses heuristics rather than disruptive maneuvers [247], but one study [111] reported only a moderate agreement against P_{mus} derived from P_{es} . Another method worth mentioning was introduced by Becher et al. [23] who used random variations of the PS level (also called noisy PSV [214]) to identify lung parameters. As opposed to the method by Younes et al. [248, 249], it does not disrupt the natural breathing rhythm of the patient but also provides sufficient excitation to the system, allowing to approximate mechanical parameters. As with all ventilatory maneuvers, there is an inherent trade-off between their *obtrusiveness* and reaching *identifiability* of parameters. A flexible solution to this trade-off is, of course, complex and requires more in-depth research.

To conclude, many authors have emphasized the great potential of the model-based approach and some encouraging results have been obtained, but it has also been subject to critique regarding its accuracy [111, 147, 180], its robustness against noise and phase delay [215], and its capability to represent certain pathologies [126, 232]. Both technical advances and evidence from clinical studies will be needed to bring model-based estimation procedures into practice.

4.2 Estimation framework

This work pursues a *model-based* approach to gain further information about P_{mus} from respiratory surface EMG signals. The herein proposed method employs *sensor fusion* to combine ventilator signals with respiratory sEMG during assisted breathing and non-invasively estimate the lung parameters and inspiratory effort. One of the main advantages of using model-based methods is that the estimation quality can be significantly improved when all available signals are combined. Thus, a higher accuracy may be achieved than with the simpler occlusion-based approach. Also, P_{mus} can be estimated without using airway occlusions or other disruptive maneuvers. In contrast to previous methods, the estimation procedure was designed to be highly robust against phase delays, outliers and artifacts. To this end, this chapter employs robust regression methods instead of OLS regression, because OLS had been reported to be overly sensitive to phase delays and artifacts [215]. This section will start with a brief recap on the occlusion-based P_{mus} estimation and then the novel estimators will be introduced.

4.2.1 Occlusion-based effort estimation

The simple occlusion-based method serves as a baseline for the model-based approaches that are introduced in the following section. As described in the previous chapter in section 3.2.3, K_{EMG} can be determined over the course of multiple subsequent occlusions by fitting the selected sEMG channel EMG_{sel} to P_{aw} using linear regression.

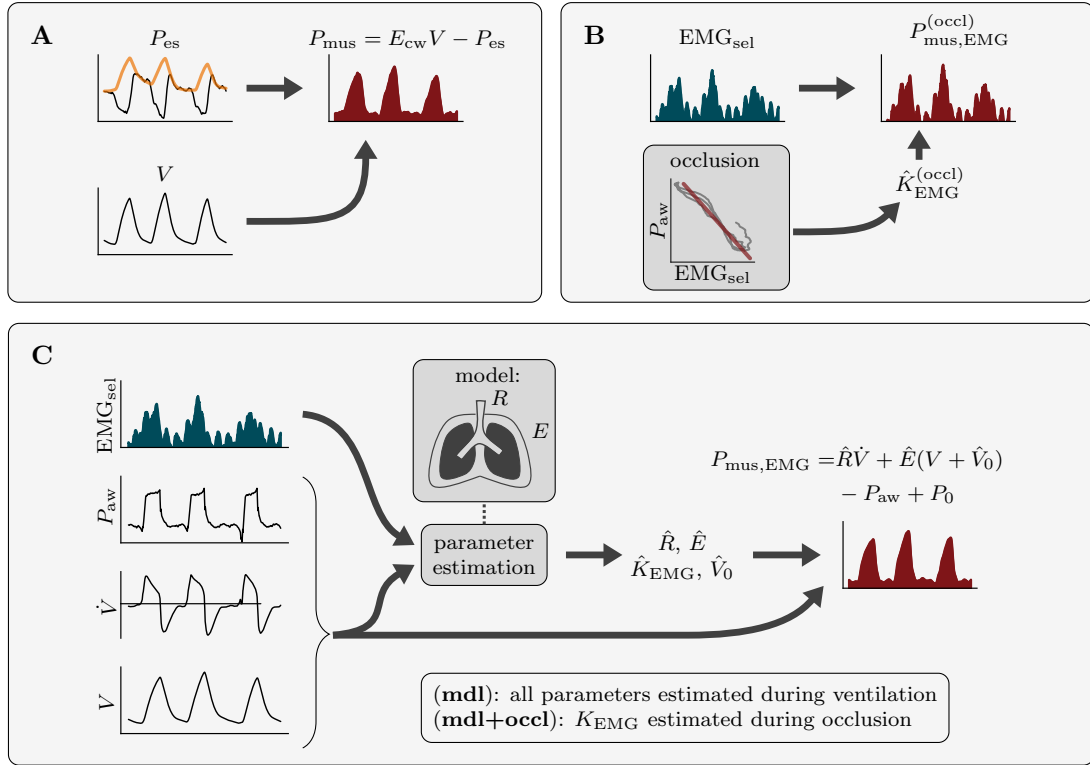


Fig. 4.1: An overview of different methods for estimating respiratory effort. **(A)** Esophageal pressure-based method: P_{es} is recorded with an esophageal balloon, cardiac artifacts are filtered out, and P_{mus} is calculated by correcting for the elastic recoil of the chest wall as described in section 3.2.2. **(B)** Occlusion-based method: The more informative of the two sEMG channels, denoted as EMG_{sel} , is selected automatically based on the SNR. The neuromechanical coupling index $\hat{K}_{EMG}^{(occl)}$ is determined during airway occlusion maneuvers from the slope of the EMG - P_{aw} curve and $P_{mus,EMG}^{(occl)}$ is calculated via equation (4.2). **(C)** Model-based method: A model of the respiratory system is fitted to the EMG and ventilatory data. The estimated parameters are then used to reconstruct P_{mus} purely from the pneumatic signals (P_{aw} , \dot{V} , V). Two variants are investigated: in the first one ($P_{mus,EMG}^{(mdl)}$), all parameters are estimated during normal ventilation at multiple pressure support levels, in the second one ($P_{mus,EMG}^{(mdl+occl)}$), the neuromechanical coupling is determined from occlusion maneuvers and the remaining parameters are estimated during normal ventilation. Figure modified and extended from Graßhoff et al. [78].

Based on the previous results and in accordance with Bellani et al. [27, 29], a constant correction factor $\alpha = 0.8$ was used to account for the systematic overestimation of neuromechanical coupling during occlusions. An sEMG-derived estimate of P_{mus} can then be calculated via

$$P_{mus,EMG}^{(occl)}(t) = \alpha \cdot \hat{K}_{EMG}^{(occl)} \cdot EMG(t), \quad (4.2)$$

where $EMG(t) = EMG_{sel}(t) - EMG_{sel,0}(t)$ corresponds to the envelope of the automatically selected sEMG channel corrected by its baseline $EMG_{sel,0}(t)$. In this work, the performance of this estimate was compared to the following model-based approaches. A visual overview of all methods used in this chapter, including the occlusion-based method described in this section, is given in figure 4.1.

4.2.2 Model-based effort estimation

The occlusion method can already be understood as being a variant of a model-based approach, because a simple model of the respiratory system is used to relate EMG to airway pressure during occlusions and to estimate a parameter of the physiological system (namely K_{EMG}). In this section, all available measurements are to be represented in a single model and its parameters will be estimated to form a highly reliable P_{mus} signal.

Model of respiratory mechanics

As a respiratory mechanics model this chapter employs

$$P_{aw}(t) = R \cdot \dot{V}(t) + E \cdot (V(t) + V_0) - P_{mus}(t) + P_0, \quad (4.3)$$

which is a variant of the standard equation of motion from equation (4.1) with the addition of a volume offset V_0 to represent intrinsic PEEP (iPEEP) and $P_0 = PEEP$ can be assumed to be given. As in the occlusion-based methods, the relationship between P_{mus} and EMG was modeled using a linear neuromechanical coupling index: $P_{mus}(t) = K_{EMG} \cdot EMG(t)$. Inserting this into equation (4.3) gives the combined equation

$$P_{aw}(t) = R \cdot \dot{V}(t) + E \cdot (V(t) + V_0) - K_{EMG} \cdot EMG(t) + P_0, \quad (4.4)$$

which is in the form of a linear regression model. By introducing the EMG as an additional regressor, the problem is no longer underdetermined and can be solved for the parameters using appropriate linear regression methods (such as OLS). A solution can be obtained on data collected during normal assisted ventilation without using maneuvers (given that the sample data matrix has full rank). As before, the channel selection method described in section 3.2.3 was employed and thus only the most informative channel was used, i.e., $EMG(t) = EMG_{sel}(t) - EMG_{sel,0}(t)$. As discussed earlier, automatic channel selection based on the SNR improves correlation to P_{mus} . Finally, given some estimate for the respiratory mechanical parameters (\hat{R} , \hat{E} , \hat{K}_{EMG} and \hat{V}_0), the muscular pressure P_{mus} can be reconstructed using

$$P_{mus,EMG}(t) = \hat{R} \cdot \dot{V}(t) + \hat{E} \cdot (V(t) + \hat{V}_0) - P_{aw}(t) + P_0, \quad (4.5)$$

which only uses pneumatic measurements from the ventilator that are less affected by noise than the EMG signals, thus offering an improvement of the SNR over using the EMG directly. The pneumatic P_{mus} calculation coincides with what has been widely used in the literature [114, 241] and in proportional ventilation modes such as PAV and PPS [247]. The following section will introduce a robust solution to the parameter estimation problem in equation (4.4).

Robust parameter and effort estimation

The estimation of ‘good’ parameters solving equation (4.4) turns out to be challenging. Empirical results have shown that the simple sample-wise OLS solution is brittle: the estimation of lung parameters was reported to be highly susceptible to small phase delays as well as model mismatch [44, 126, 161, 180, 215, 232]. Experience gained from this chapter also indicated that the OLS solution often fails to correctly capture the ‘big picture’ of ventilation data because it is overly sensitive to single outliers and tends to overfit uninformative data such as sections containing many repetitive breaths, refer to Petersen, Graßhoff et al. [161]. In the following, these problems are mitigated by introducing a novel estimator tailored to the characteristics of ventilatory data. Two key novelties are proposed: (1) the respiratory time constant $\tau = R/E$ is estimated separately on expiratory data and (2) the remaining parameters are estimated in a breath-wise integrated form of the equation of motion. In both steps, this work advocates to use robust regression methods instead of OLS regression to reduce the influence of outliers.

The first novelty of this work concerns the estimation of the time constant τ , which is to be determined separately and prior to other parameters. To this end, the proposed method exploits the fact that most patients under assisted ventilation are less active during expirations, which can be recognized from the typical exponential shape of the expiratory flow curve. In such cases, the respiratory time constant could be calculated easily by reading off the V/\dot{V} ratio, which has already been used in previous works [173, 229]. However, in most patients, a direct calculation is complicated by residual patient activity, for example by missed efforts and active expirations. Moreover, possible offsets of the volume V_0 have to be taken into account before forming the V/\dot{V} ratio. For these reasons, this chapter proposes a model-based estimation of the time constant. As a model for expirations

$$V(t) = -\tau \cdot \dot{V}(t) - V_0 \quad (4.6)$$

was used, which follows directly from equation (4.3), when $P_{\text{aw}} = P_0 = \text{PEEP}$ and $P_{\text{mus}} = 0$ is assumed. This equation was fitted simultaneously to multiple expirations and robust regression was used to reduce the influence to outliers caused by expiratory patient activity. To account for possible changes of the volume offset V_0 in between expirations (due to iPEEP or remaining drifts of the volume signal), a *parallel slopes regression* model was used, where the slope τ is shared between all expirations, while

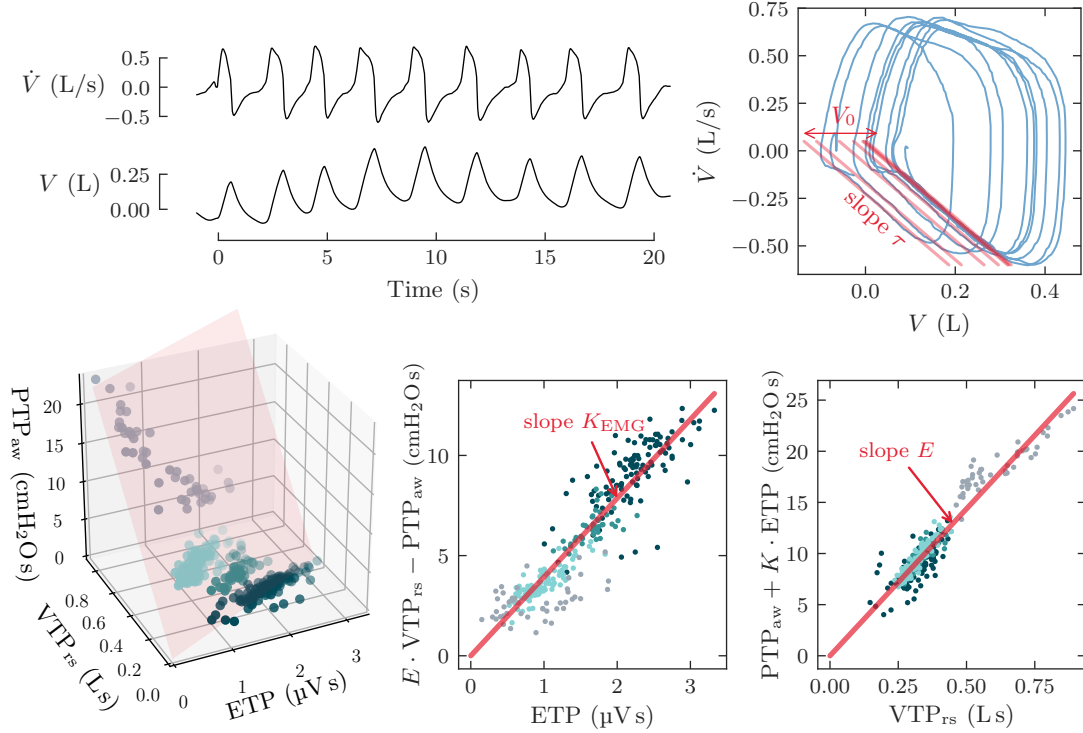


Fig. 4.2: Key steps involved in the estimation procedure. The upper left panel shows the time courses of airway flow \dot{V} and lung volume V of a patient under assisted mechanical ventilation. Passive expirations can be recognized from the exponential shape of the flow and volume curve, allowing to determine the respiratory time constant τ from the slope of the flow–volume loop shown in the upper right panel. Due to remaining drift of the volume signal and intrinsic PEEP, *parallel slopes regression* is used to estimate τ and a breath-wise offset parameter V_0 . Given τ , only two parameters, the neuromechanical coupling K_{EMG} and elastance E , remain to be estimated, which is illustrated in the bottom panels. To increase robustness, these parameters are estimated using a breath-wise *integrated equation of motion* introduced in (4.13). Each breath provides a point within the time product space defined by VTP_{rs} , ETP and PTP_{aw} values. The regression can be visualized as a 2D-plane with slopes K_{EMG} and E along the two independent variables. The color of points corresponds to the four PS levels: CPAP (●), PS at 5 cmH₂O (●), PS at 10 cmH₂O (●), PS at 15 cmH₂O (●).

an individual offset parameter $V_{0,i}$ was used in each expiration; i indicates breaths. Thus, given measurements of expiratory time points t_k , the estimator used

$$V(t_k) = -\tau \cdot \dot{V}(t_k) - \sum_i V_{0,i} \cdot Y(t_k, i) + \epsilon_k, \quad (4.7)$$

where $Y(t_k, i) = 1$ if t_k belongs to expiration of breath i and $Y(t_k, i) = 0$ otherwise, and minimized $\sum_k \rho_{\text{Tukey}}(\epsilon_k)$ with respect to the parameters τ and $V_{0,i}$ using iteratively




Fig. 4.3: A simulation of pneumatic signals under the equation of motion (4.3). When there is a superposition of P_{aw} and P_{mus} , the volume V and flow \dot{V} tend to be difficult to interpret. Given an estimate of τ , the flow and volume can be aggregated via $V_{\text{rs}}(t) := V(t) + \tau \cdot \dot{V}(t)$. The variable V_{rs} is proportional to the total pressure $P_{\text{rs}}(t) = P_{\text{aw}}(t) + P_{\text{mus}}(t) - P_0$ acting on the respiratory system, and thus, gives direct insight into the pneumatic output from both the ventilator and patient.

reweighted least squares. The function ρ_{Tukey} denotes Tukey’s bisquare loss function introduced in section 2.1.2. An example of the regression result on patient data is depicted in figure 4.2. The parameter estimates were denoted as $\hat{\tau}$ and $\hat{V}_{0,i}$. The parameters $\hat{V}_{0,i}$ account for any breath-wise volume offsets for example due to iPEEP. As a next step, a continuous estimate for the volume offset $\hat{V}_0(t)$ was calculated based on the individual (expiratory) offsets $\hat{V}_{0,i}$. This was done mainly by smoothing and interpolating between the baselines of expiratory segments; more details are provided in the appendix of [78]. The resulting signal $\hat{V}_0(t)$ is intended to provide a smooth representation of volume offsets across the full length of recordings.

The proposed method provides a reliable, automated estimate of the respiratory time constant τ . The use of robust regression automatically discards any expirations that are too disturbed by patient activity. A good estimate can be obtained as long as no more than roughly half of the breaths are affected. In contrast, with OLS all outliers directly have a negative impact on the regression result. To further increase robustness, additional selection criteria can be used for selecting ‘good’ sections of the data: in this work, only expirations were included for which $0.2\text{L} < V_{\text{T}} < 1.0\text{L}$ and $0.6\text{s} < T_{\text{exp}} < 3.0\text{s}$, where V_{T} denotes the tidal volume, T_{exp} denotes the length of the expiration. In the future, more sophisticated criteria could be used for selecting suitable expirations, for example based on the shape of the ventilator curves similar to [32], or based directly on the EMG.

The identification of the time constant τ and the volume offset V_0 solves an important part of the overall estimation problem. The estimates were incorporated by inserting $R = E \cdot \tau$ into (4.4), from which

$$P_{\text{aw}}(t) = E \cdot (V(t) + V_0) + E \cdot \tau \cdot \dot{V}(t) - K_{\text{EMG}} \cdot \text{EMG}(t) + P_0 \quad (4.8)$$

$$= E \cdot V_{\text{rs}}(t) - K_{\text{EMG}} \cdot \text{EMG}(t) + P_0, \quad (4.9)$$

can be obtained, where

$$V_{\text{rs}}(t) := V(t) + V_0 + \tau \cdot \dot{V}(t). \quad (4.10)$$

The variable $V_{\text{rs}}(t)$ is fully defined in terms of previous parameters (an estimate $\hat{V}_{\text{rs}}(t)$ can be constructed directly from $\hat{\tau}$ and $\hat{V}_0(t)$). Now only two parameters (K_{EMG} and E) remain to be estimated in equation (4.10), which is a much simpler problem than estimating all parameters at once. It should be pointed out that the variable V_{rs} is quite interesting in itself, which is illustrated in figure 4.3. Using the equation of motion defined in equation (4.3), one can show that $E \cdot V_{\text{rs}}(t) = P_{\text{aw}}(t) + P_{\text{mus}}(t) - P_0 = P_{\text{rs}}(t)$. Thus, $V_{\text{rs}}(t)$ is proportional to the total pressure $P_{\text{rs}}(t)$ acting on the respiratory system and the proportionality constant is the respiratory elastance E . In that sense, $V_{\text{rs}}(t)$ is a surrogate for the pneumatic output that is jointly generated by the ventilator and the patient.

Regarding the second novelty proposed in this work, a new form of the EqM was introduced, based on breath-wise integration over time of equation (4.9), resulting in

$$\int P_{\text{aw}}(t) dt = \int E \cdot (V(t) + V_0) + R \cdot \dot{V}(t) dt - \int K_{\text{EMG}} \cdot \text{EMG}(t) dt + \int P_0 dt, \quad (4.11)$$

where the integrated terms correspond to time product values, which were already used in chapter 3 and also widely employed in the mechanical ventilation literature. Chapter 3 has demonstrated that the EMG–time product is a viable metric for patient activity and that it is highly correlated to the P_{mus} –time product. The breath-wise evaluation of equation (4.11) leads to an *integrated equation of motion*

$$\text{PTP}_{\text{aw}} = E \cdot \text{VTP}_{\text{rs}} - \text{PTP}_{\text{mus}} + \text{PTP}_0 \quad (4.12)$$

$$= E \cdot \text{VTP}_{\text{rs}} - K_{\text{EMG}} \cdot \text{ETP} + \text{PTP}_0, \quad (4.13)$$

where the variables are given in terms of time product values and can be easily calculated: PTP_{aw} is the pressure–time product of the airway pressure, VTP_{rs} is calculated by integrating V_{rs} leading to a volume–time product, ETP is the EMG–time product and PTP_0 is calculated by integrating the PEEP. In this work, one set of time product values were calculated from each inspiration and inspirations were detected directly from the flow signal. Given a set of breaths

$$\text{PTP}_{\text{aw}}(i) = E \cdot \text{VTP}_{\text{rs}}(i) - K_{\text{EMG}} \cdot \text{ETP}(i) + \text{PTP}_0 + \epsilon_i \quad (4.14)$$

was used in the estimation and the robust regression loss $\sum_i \rho_{\text{Huber}}(\epsilon_i)$ was minimized with respect to the parameters K_{EMG} and E by means of iteratively reweighted least squares; the function ρ_{Huber} represents the Huber loss from section 2.1.2 and, as before, i indicates breaths. Using the respiratory elastance estimate \hat{E} , one can calculate the muscle pressure via

$$P_{\text{mus,EMG}}(t) = \hat{E} \cdot \hat{V}_{\text{rs}}(t) - P_{\text{aw}}(t) + P_0, \quad (4.15)$$

which corresponds to the pneumatic reconstruction in equation (4.5). An estimate for the respiratory resistance can be calculated as $\hat{R} = \hat{E} \cdot \hat{\tau}$.

The reconstruction of P_{mus} from pneumatic measurements often causes high frequency artifacts due to model mismatch or noise in the pneumatic signals. This was already noticed by Younes et al., who reported small spikes in the P_{mus} signal when the flow changes rapidly, refer to the supplement of [247]. These artifacts have a much higher frequency than the respiratory waveforms and can be removed using low-pass filtering. To this end, a seventh order Butterworth low-pass filter with a cutoff frequency of 3.5 Hz was used. Empirical tests showed that this cutoff managed to remove most artifacts and did not strongly affect the physiological waveform of the P_{mus} signal. The filtered signal was denoted as $P_{\text{mus,EMG}}^{(\text{mdl})}$ and the corresponding parameter estimates as $\hat{K}_{\text{EMG}}^{(\text{mdl})}$, $\hat{E}^{(\text{mdl})}$ and $\hat{R}^{(\text{mdl})}$.

Remarks on the proposed method

The herein proposed integrated form of the EqM implies a breath-wise interpretation, that is, the model is fitted to explain complete breaths rather than individual samples. In that sense, the model relates the flow and volume generated across full breaths (which is aggregated in VTP_{rs}) to the aggregated pressure outputs from the ventilator (PTP_{aw}) and from the patient (PTP_{mus}). Thus, proportionality between EMG and P_{mus} waveforms is not required, because EMG data are only ever used in their integrated form. Very importantly, $PTP_{\text{mus}}(i) = K_{\text{EMG}} \cdot \text{ETP}(i)$ is a much weaker assumption than $P_{\text{mus}}(t) = K_{\text{EMG}} \cdot \text{EMG}(t)$, and a tight coupling between the inspiratory effort PTP_{mus} and the EMG–time product ETP was previously demonstrated, refer to chapter 3. The regression of time product values is much less susceptible to small phase delays, which were reported to strongly affect the estimation of respiratory parameters [215], and to autocorrelated measurement errors. To give one example, cardiogenic artifacts remaining in the EMG envelope pose a much smaller problem when using ETP values and are therefore less likely to negatively affect the regression result.

It can be easily seen in equation (4.13) that two linearly independent sets of time product values are needed to solve the regression with respect to E and K_{EMG} . Thus, some form of change in workload shares between ventilator and patient will be needed, that is, there must be sufficient variation of VTP_{rs} against ETP to ensure

identifiability. The regression space is visualized in figure 4.2 across four PS levels. A change in workload shares can be achieved by relying on natural variations of the patient's breathing pattern or by introducing external variation for example by changing the ventilator PS level. Thus, the parameter estimation can be solved during tidal ventilation without relying on highly disruptive maneuvers such as airway occlusions or pressure pulses as in the methods proposed by Younes et al. [248, 249]. Notice that the herein proposed methodology is not restricted to ETP and could employ any additional measurement with a high correlation to PTP_{mus} to estimate respiratory parameters. As a last remark, the proposed EqM (4.11) could also be calculated by integrating over volume instead of time, which would then give rise to a formulation in terms of *work of breathing* instead of time product values.

4.2.3 Combined approach

In the model-based approach presented above, respiratory parameters are estimated during normal tidal breathing on multiple pressure support levels and the P_{mus} curve is formed using only pneumatic signals. In contrast, the occlusion-based approach presented in section 4.2.1 provides an immediate estimate for the neuromechanical coupling parameter K_{EMG} , but is highly dependent on the signal quality of the EMG. As a final approach, the merits of both methods were combined. To this end, K_{EMG} was estimated from occlusions maneuvers and then used as a fixed parameter in the model-based estimation described above. Here, the method follows the same parameter identification steps, but in the robust regression equation (4.13), the parameter $\alpha \cdot \hat{K}_{\text{EMG}}^{(\text{occl})}$ was used for K_{EMG} and not estimated. As before, the P_{mus} estimate was then constructed from pneumatic signals via equation (4.15). The P_{mus} estimate was denoted as $P_{\text{mus,EMG}}^{(\text{mdl+occl})}$ and the corresponding parameters as $\hat{K}_{\text{EMG}}^{(\text{mdl+occl})}$, $\hat{E}^{(\text{mdl+occl})}$ and $\hat{R}^{(\text{mdl+occl})}$. Notice that the neuromechanical coupling parameter K_{EMG} is identical to the one used in the occlusion based method, but the two approaches differ in how the P_{mus} signal is calculated from K_{EMG} .

4.3 Validation & performance quantification

The robust model-based estimator was validated against P_{mus} derived from esophageal pressure on the study data presented in chapter 3 comprising 43 patients and compared to the simpler occlusion-based approach. An overview of all considered methods is given in figure 4.1. The data preprocessing followed the same steps as in section 3.2.2. Respiratory parameters were estimated across all PS levels (CPAP, PS at 5, 10 and 15 cmH₂O) using the automatically selected EMG channel and the pneumatic signals from the ventilator, i.e., the airway flow \dot{V} , the trend-corrected volume V , and the airway pressure P_{aw} . The model-based estimates $P_{\text{mus,EMG}}^{(\text{mdl})}$ and $P_{\text{mus,EMG}}^{(\text{mdl+occl})}$ were

validated against P_{mus} while the occlusion-based estimate $P_{\text{mus,EMG}}^{(\text{occl})}$ provides a baseline performance for the model-based approaches.

4.3.1 Performance metrics

The model-based estimates of P_{mus} were evaluated in terms of breath-wise pressure-time product values. Thus, PTP_{mus} was calculated for the reference P_{mus} , $\text{PTP}_{\text{mus,EMG}}^{(\text{occl})}$ for the occlusion-based approach, and $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl})}$ as well as $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl+occl})}$ for the model-based estimators. The start and end of inspiratory efforts were detected directly from the P_{mus} signal as described in section 3.2.3. As a second metric for inspiratory effort PTP/min values, measured in $\text{cmH}_2\text{O s}/\text{min}$, were calculated by summing up the PTP values of all detected efforts in each PS level and dividing it by the total length of the segment. The following performance metrics were used to characterize the relation between estimated $\text{PTP}_{\text{mus,EMG}}$ and measured PTP_{mus} :

Slope and bias. To analyze possible scaling errors and biases between $\text{PTP}_{\text{mus,EMG}}$ and PTP_{mus} , the linear regression model $\text{PTP}_{\text{mus}}(i) = K \cdot \text{PTP}_{\text{mus,EMG}}(i) + P_{\text{bias}} \cdot T_{\text{insp}}(i)$ was used, where K is the slope parameter, P_{bias} is a constant bias term and T_{insp} is the length of detected efforts. The regression was solved for each patient and section 4.4 reports mean \pm standard deviation of the parameters across all patients.

Correlation. In each patient, Pearson's correlation index r was used to assess the relation between estimated and measured PTP values and the mean \pm standard deviation are reported across all patients.

Deviation. The deviation between $\text{PTP}_{\text{mus,EMG}}$ and PTP_{mus} was analyzed using the limits of agreement method by Bland and Altman [33]. As the number of detected efforts varied widely between patients a variant of the classical Bland-Altman method accounting for repeated measurements was used, refer to [34]. As a second metric for deviation, the mean absolute deviation (MAD), weighted by the number of efforts from all patients, was employed.

Differences between parameter estimates were analyzed by means of the two-tailed Wilcoxon signed rank test.

4.4 Results

Below, results are reported for a total of $n = 43$ patients, refer to table A.1 for clinical characteristics. An excerpt of patient data and the proposed model-based estimate $P_{\text{mus,EMG}}^{(\text{mdl})}$ are depicted in figure 4.4.

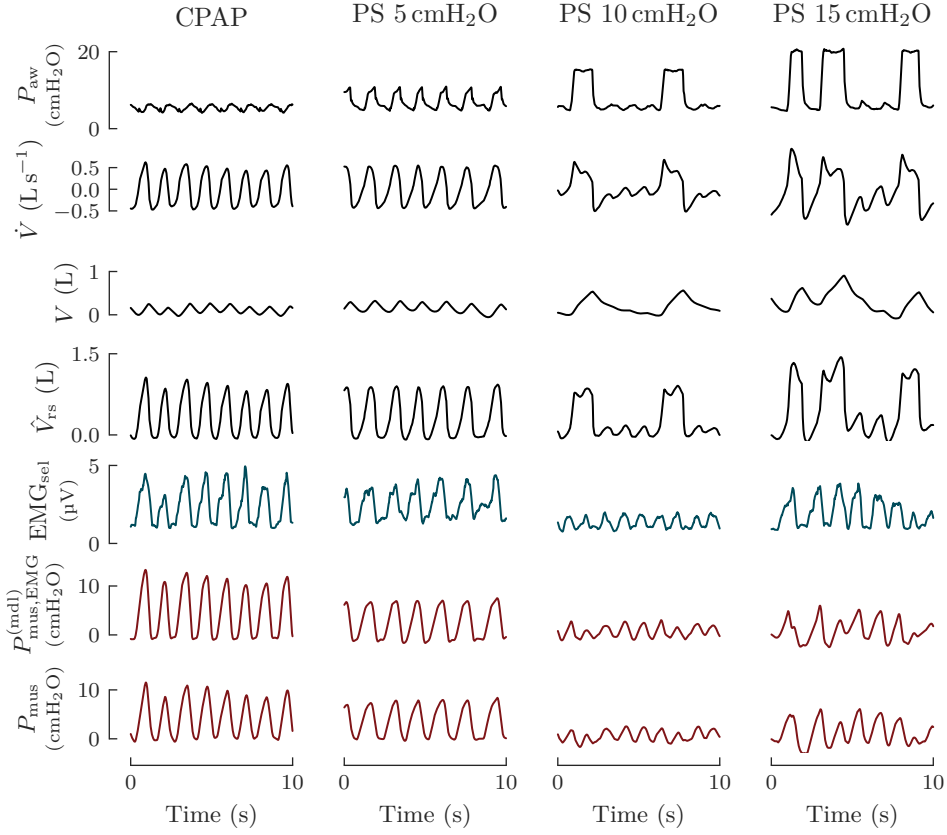


Fig. 4.4: Measurements and model-based estimates of \hat{V}_{rs} and P_{mus} from one patient over pressure support (PS) levels. In all four levels, the non-invasive estimate ($P_{mus,EMG}^{(mdl)}$) is strongly correlated to the reference (P_{mus}) derived from esophageal pressure.

Validation against PTP_{mus}

Estimated $PTP_{mus,EMG}$ values were validated against the reference PTP_{mus} derived from esophageal pressure. As a first step, the merits of using the proposed estimation techniques were investigated in comparison to simpler (OLS-based) estimators. To this end, different combinations of the novelties proposed in this work were tested, i.e., the separate estimation of τ , the integrated equation of motion, and using robust regression instead of OLS. Results are provided in table 4.1. On average the highest breath-wise correlation per patient between estimated and measured PTP_{mus} was attained by combining all aforementioned techniques, see last row of table 4.1, corresponding to the $PTP_{mus,EMG}^{(mdl)}$ estimate. This estimate also had the smallest bias P_{bias} of all methods and the smallest standard deviation of breath-wise deviation against PTP_{mus} . These results demonstrate the benefits of using the proposed robust regression approach as compared to simpler OLS methods for estimating respiratory parameters and P_{mus} .

Table 4.1: Performance metrics (mean \pm standard deviation) for baseline ordinary least squares (OLS) estimators and different combinations of robust regression (RR) methods proposed in this chapter, $n = 43$ ($n = 41$ when occlusions are involved in the calculation of a metric). Definitions of performance metrics are provided in section 4.3.1.

		sep. τ estimator	integr. EqM	correlation	slope K	bias P_{bias} (cmH ₂ O)	PTP dev. (cmH ₂ O s)
baseline methods (OLS) [*]	const. V_0	–	–	0.902 ± 0.067	1.03 ± 0.22	1.84 ± 1.60	1.62 ± 2.14
	non-const. V_0	–	–	0.903 ± 0.076	1.10 ± 0.29	0.91 ± 1.56	0.99 ± 2.31
proposed methods		OLS	–	0.930 ± 0.056	0.93 ± 0.18	1.28 ± 1.72	0.79 ± 2.22
		RR	–	0.931 ± 0.054	0.93 ± 0.18	1.27 ± 1.71	0.79 ± 2.21
		RR	OLS	0.946 ± 0.051	0.82 ± 0.15	0.49 ± 1.54	-0.45 ± 1.89
		RR	RR	0.953 ± 0.037	0.79 ± 0.14	0.29 ± 1.43	-0.83 ± 1.73

integr. EqM: integrated equation of motion; PTP dev.: pressure–time product deviation; *: for the two baseline methods, sample-wise OLS was applied directly to the EqM $P_{\text{aw}}(t) = R \cdot \dot{V}(t) + E \cdot V(t) - K_{\text{EMG}} \cdot \text{EMG}(t) + P_0$ and solved for R , E , K_{EMG} and the offset parameter $P_0 = E \cdot V_0 + \text{PEEP}$; in a first variant (assuming a constant V_0) a single offset P_0 was estimated for each patient, in a second variant (assuming non-constant V_0) a time-varying P_0 was implemented by estimating weights of Gaussian basis functions (which had standard deviation $\sigma = 12$ s and were placed at regular 12 s intervals over the length of the recording);

Next, the performance of the model-based estimates $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl})}$ and $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl+occl})}$ was compared to the simpler occlusion-based estimate $\text{PTP}_{\text{mus,EMG}}^{(\text{occl})}$. Results are reported in table 4.2. The breath-wise correlation between the occlusion-based estimate $\text{PTP}_{\text{mus,EMG}}^{(\text{occl})}$ and PTP_{mus} was $r = 0.87 \pm 0.09$, while both model-based approaches attained a very high correlation to PTP_{mus} ($r = 0.95 \pm 0.04$ for both), which was an improvement over the occlusion-based method ($p < .001$). The improvement in correlation demonstrates that the model-based calculation of P_{mus} substantially reduces the random scatter associated with using the sEMG envelopes directly. The bias P_{bias} of both occlusion- and the model-based estimators in table 4.2 was small. There was a systematic difference in the scale of PTP_{mus} against both model-based approaches, because the slope K was smaller than 1 ($p < .001$ for both), indicating that the estimated values were scaled up compared the reference. The comparison of different methods in table 4.1 reveals a progression of this effect throughout the variants: in fact, the best performing estimators tended to show a stronger overestimation against the esophageal pressure reference.

The breath-wise deviation to PTP_{mus} was analyzed using the limits of agreement approach by Bland and Altman [34] and results are also reported in table 4.2. This analysis was conducted on $n = 41$ patients, where sufficiently long occlusion maneuvers

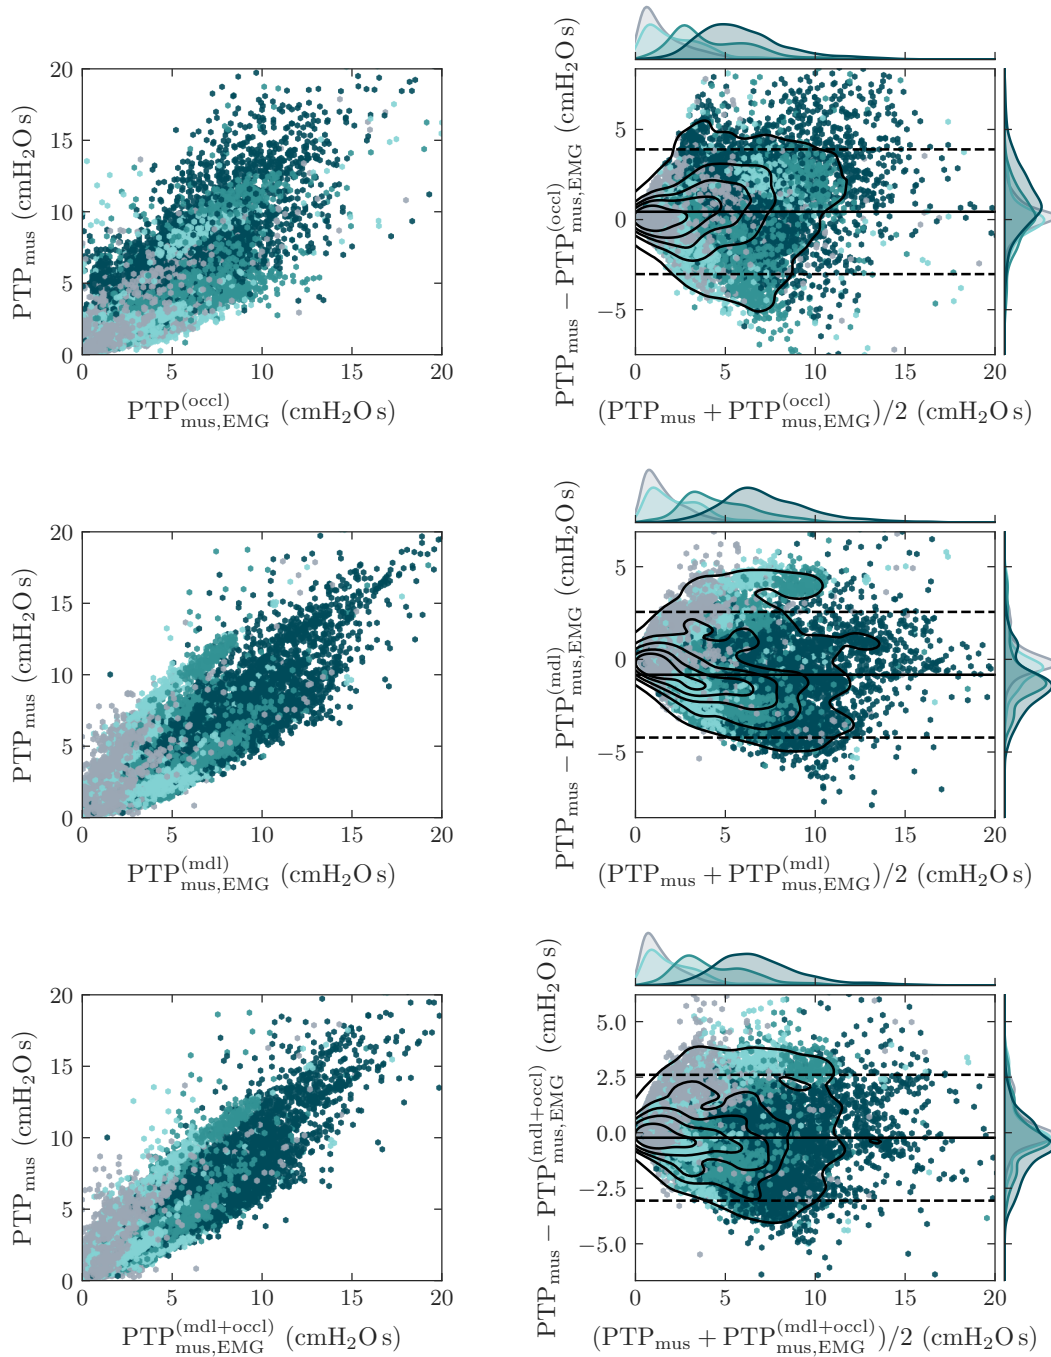


Fig. 4.5: Relation between estimated $PTP_{\text{mus,EMG}}$ and invasively measured PTP_{mus} values. The left column shows breath-wise scatter plots and the right column shows Bland-Altman plots of breath-wise differences. The plots contain $m = 18\,341$ breaths from $n = 41$ patients; each point represents one breath. The color of points corresponds to the four PS levels: CPAP (●), PS5 (●), PS10 (●), PS15 (●). The mean is visualized by the solid black line, 95% limits of agreement are visualized by black dashed lines.

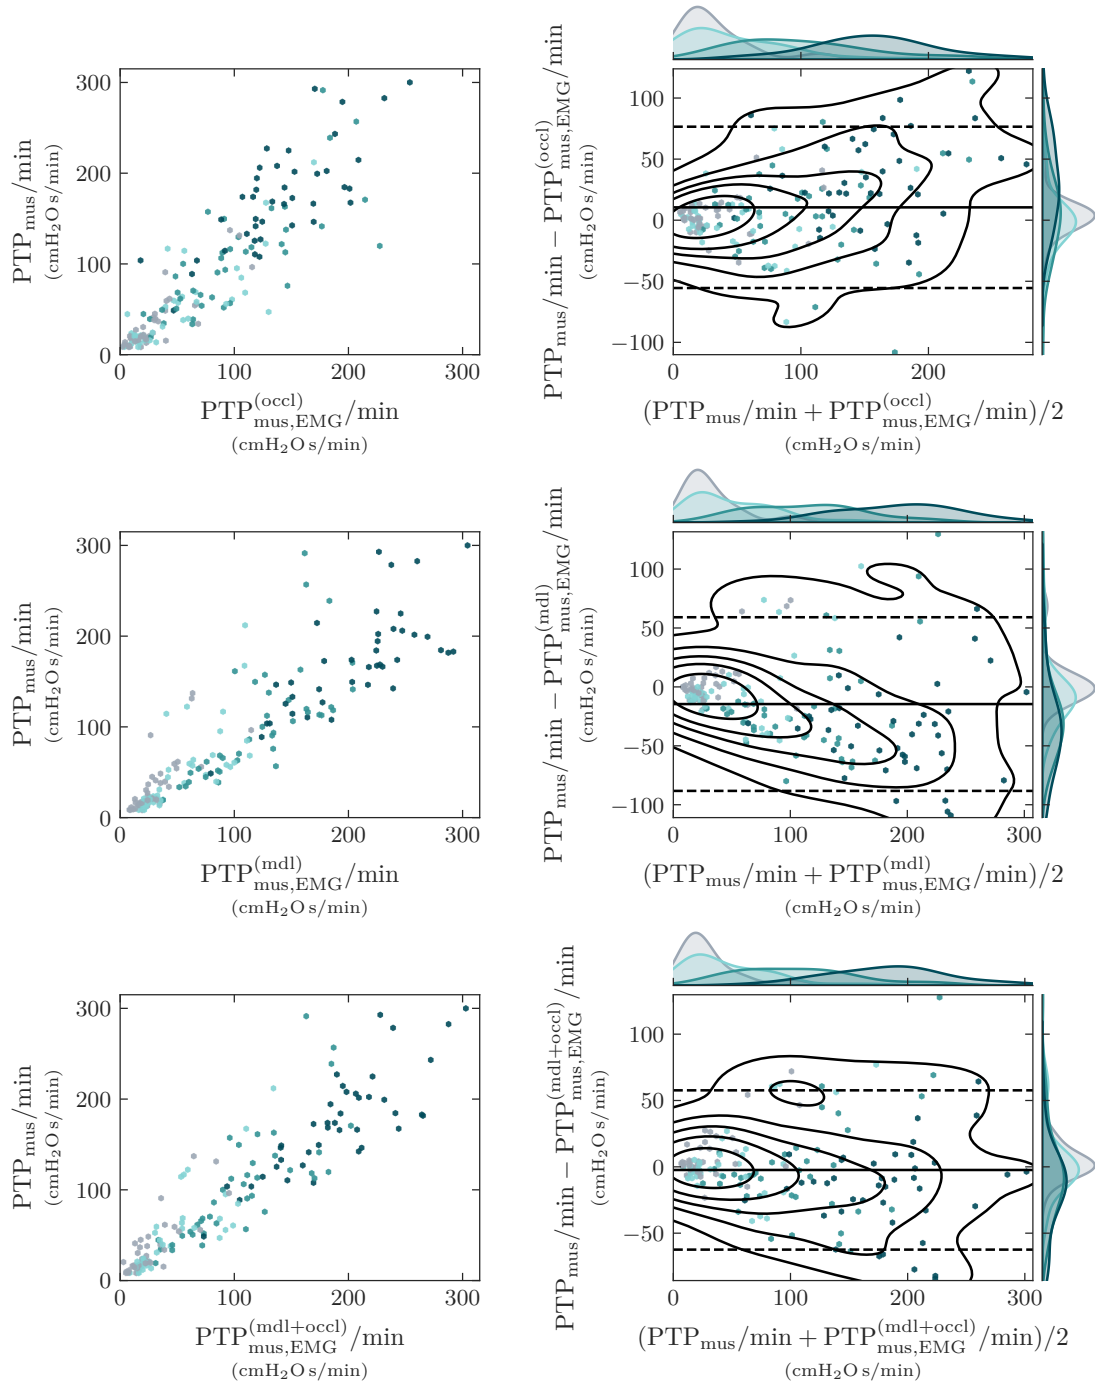


Fig. 4.6: Relation between estimated $PTP_{mus,EMG}/min$ and invasively measured PTP_{mus}/min values. The left column shows scatter plots and the right column shows Bland-Altman plots. The plots contain $k = 161$ from $n = 41$ patients; each point represents one PS level from one patient. The color of points corresponds to the four PS levels: CPAP (●), PS5 (●), PS10 (●), PS15 (●). The mean is visualized by the solid black line, 95 % limits of agreement are visualized by black dashed lines.

were available. The smallest deviation to PTP_{mus} was achieved by the combined estimator $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl}+\text{occl})}$ ($-0.23 \pm 1.44 \text{ cmH}_2\text{O s}$). The corresponding Bland-Altman plots are provided in figure 4.5, which show that both model-based approaches would allow to clinically assess the inspiratory effort of patients on a breath-by-breath basis. Finally, the ability of different estimators to quantify inspiratory efforts over a longer time-scale in terms of PTP/min values was investigated. Again, Bland-Altman plots were used to analyze the limits of agreement with $\text{PTP}_{\text{mus}}/\text{min}$, see figure 4.6. The deviation of the combined estimator (mdl+occl) was $-2.37 \pm 30.60 \text{ cmH}_2\text{O s}/\text{min}$. Taking into account a clinical value range of $0-250 \text{ cmH}_2\text{O s}/\text{min}$ a typical target range of $100-150 \text{ cmH}_2\text{O s}/\text{min}$, the reported agreement of the model-based estimator would allow to accurately assess patient effort.

Table 4.2: Performance metrics (mean \pm standard deviation) for the occlusion-based method from section 4.2.1, the model-based estimator described in section 4.2.2 and the combined approach from section 4.2.3, $n = 43$ ($n = 41$ when occlusions are involved in the calculation of a metric). Definitions of performance metrics are provided in section 4.3.1.

	$\text{PTP}_{\text{mus,EMG}}^{(\text{occl})}$	$\text{PTP}_{\text{mus,EMG}}^{(\text{mdl})}$	$\text{PTP}_{\text{mus,EMG}}^{(\text{mdl}+\text{occl})}$
correlation r	0.87 ± 0.09	0.95 ± 0.04	0.95 ± 0.04
slope K	1.01 ± 0.29	0.79 ± 0.14	0.83 ± 0.14
bias P_{bias} (cmH_2O)	0.71 ± 1.46	0.29 ± 1.43	0.64 ± 1.27
PTP deviation ($\text{cmH}_2\text{O s}$)	0.43 ± 1.77	-0.83 ± 1.73	-0.23 ± 1.44
PTP MAD ($\text{cmH}_2\text{O s}$)	1.26	1.48	1.04
PTP/min deviation ($\text{cmH}_2\text{O s}/\text{min}$)	10.5 ± 33.7	-14.6 ± 37.6	-2.37 ± 30.60
PTP/min MAD ($\text{cmH}_2\text{O s}/\text{min}$)	24.3	29.4	20.5

PTP: pressure–time product; PTP/min: aggregated pressure–time product; MAD: mean absolute deviation; $\text{PTP}_{\text{mus,EMG}}^{(\text{occl})}$: occlusion-based estimate; $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl})}$: model-based P_{mus} estimate; $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl}+\text{occl})}$: estimated via combined method;

Parameter estimates


The estimated parameter values of the two model-based methods are provided in table 4.3. The parameter values were in a physiologically plausible range and coincided with values reported in [14] for healthy and COPD patients. Reference values for lung mechanics parameters were not available in the study, which is why the focus was on the evaluation of neuromechanical coupling. The neuromechanical coupling estimated by the first model-based method ($\hat{K}_{\text{EMG}}^{(\text{mdl})}$) was highly correlated to the reference

Table 4.3: Parameter estimates of the model-based estimator (mdl) introduced in section 4.2.2 and the combined estimation technique (mdl+occl) from section 4.2.3.

	\hat{C} (mL/cmH ₂ O)	\hat{R} (cmH ₂ O s/L)	\hat{K}_{EMG} (cmH ₂ O/ μ V)
(mdl)	56.2 \pm 15.0	13.7 \pm 2.3	5.5 \pm 3.7
(mdl+occl)	62.8 \pm 18.0	12.5 \pm 3.1	4.1 \pm 2.7
reference	–	–	4.4 \pm 3.8

C : respiratory compliance; R : respiratory resistance; K_{EMG} : neuromechanical coupling index;

value K_{EMG} derived from esophageal pressure according to equation (3.4) ($r = 0.91$, $p < .001$, slope = 0.91), see figure 4.7, demonstrating that a completely unobtrusive, model-based estimation of K_{EMG} is feasible. Testing for differences between the estimated parameter $\hat{K}_{\text{EMG}}^{(\text{mdl})}$ and the parameter $\hat{K}_{\text{EMG}}^{(\text{occl})}$ derived from occlusions revealed that the former was larger ($p < .001$), see average values in table 4.3. In turn, and in agreement with the EqM, the resistance and elastance values were larger in the first model-based estimator (mdl) than in the combined approach (mdl+occl) ($p < .001$).

**Fig. 4.7:** Scatter plot of neuromechanical coupling indices determined invasively from esophageal pressure (K_{EMG}) and non-invasively by the fully model-based estimator ($\hat{K}_{\text{EMG}}^{(\text{mdl})}$). Each point represents one patient, $n = 41$.

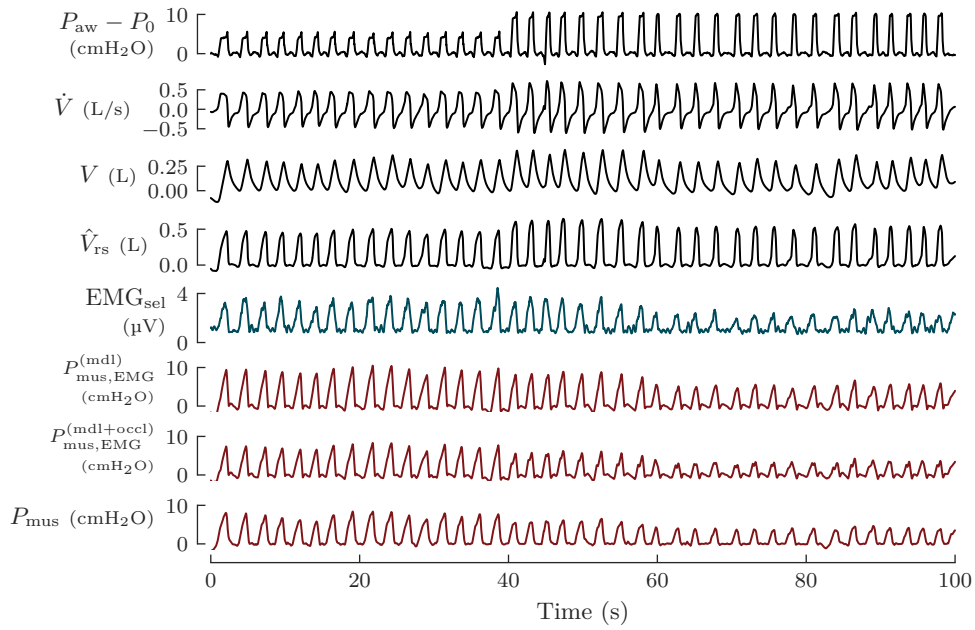
4.5 Discussion

The results demonstrate that the proposed model-based estimator provides a reliable estimate of respiratory muscle pressure P_{mus} . The main findings are: In a cohort of 43 patients under assisted mechanical ventilation, the model-based estimates $\text{PTP}_{\text{mus,EMG}}$ had a very high correlation to the reference PTP_{mus} derived from esophageal pressure ($r = 0.95 \pm 0.04$ for both $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl})}$ and $\text{PTP}_{\text{mus,EMG}}^{(\text{mdl+occl})}$). The breath-wise deviation


between the second model-based estimator ($PTP_{\text{mus,EMG}}^{(\text{mdl}+\text{occl})}$) and the reference (PTP_{mus}) was as low as $-0.23 \pm 1.44 \text{ cmH}_2\text{O s}$.

The occlusion-based method, previously proposed by Bellani et al. [29], uses a scaled version of the EMG curve to approximate P_{mus} . The approach is simple, but its performance strongly relies on a good quality of the EMG signal obtained during the occlusion maneuver, which poses practical difficulties regarding the repeatability of the method [101]. The herein proposed model-based methods use all available signals during normal tidal breathing to form a P_{mus} estimate. They are therefore much less sensitive to noise and artifacts in individual measurements. The results show that the model-based estimators manage to substantially reduce measurement noise and attain a significantly higher correlation to P_{mus} . Figure 4.8b shows an illustrative patient with very low EMG signal-to-noise ratio, for which the model-based estimator still manages to reconstruct a viable P_{mus} curve. The results show that the proposed model-based approach overcomes one of the most prominent drawbacks of surface EMG, namely its low SNR, and thus might help to bring this technique into clinical practice.

Another advantage of model-based methods is their ability to provide insight into lung parameters. The analysis conducted in this chapter has shown, for the first time, that a model-based method enables estimation of the neuromechanical coupling of respiratory muscles without disruptive maneuvers: the estimated parameter $\hat{K}_{\text{EMG}}^{(\text{mdl})}$ had a very high correlation with the reference value derived from esophageal pressure. The respiratory parameters (resistance and elastance) were in the order of magnitude expected for the cohort, but further studies are needed to validate these estimates against passive measurements (which were not available in the analyzed data). Two model-based methods have been investigated: the first one estimates all parameters simultaneously during normal assisted ventilation (mdl), while the second one takes K_{EMG} from occlusions and then estimates the remaining parameters during normal ventilation (mdl+occl). Both methods had, on average, an equally high correlation to PTP_{mus} , but the combined approach attained an overall lower deviation (in terms of MAD and the limits of agreement). The larger deviation of the fully model-based estimator (mdl) is not caused by random scatter of estimated PTP_{mus} values in individual patients, because the correlation per patient was very high, but can be attributed to a systematically higher scaling of P_{mus} . Both model-based methods showed a slight overestimation, but this effect was most pronounced in the first, fully model-based method (mdl). The measured PTP_{mus} values were on average $\sim 21\%$ smaller than the model-based estimates $PTP_{\text{mus,EMG}}^{(\text{mdl})}$, refer to table 4.2. In accordance with this, the estimated neuromechanical coupling ($\hat{K}_{\text{EMG}}^{(\text{mdl})}$) and lung parameters ($\hat{E}^{(\text{mdl})}$ and $\hat{R}^{(\text{mdl})}$) were significantly larger than in the estimator that uses occlusions to estimate K_{EMG} . When analyzing the scaling of the estimated P_{mus} signal, it must be taken into account that esophageal pressure itself is strongly affected by scaling errors due to incorrect balloon positioning/filling. The overestimation of P_{mus} by the



(a)



(b)

Fig. 4.8: Excerpts of measured data and P_{mus} estimates from two patients. In both cases the model-based estimation technique attains a high correlation to the reference derived from esophageal pressure. (a) Response of the patient to a pressure support step. (b) An example with very low sEMG signal-to-noise ratio.

model-based estimator should therefore not be dismissed as an error too quickly, but might be considered as an independent result, complementing a P_{es} -based analysis.

Several important contributions to respiratory model identification were made in this chapter. Most of the previous studies have relied on ordinary least squares (OLS) regression to identify respiratory parameters, but several authors have found potential shortcomings of this approach: the OLS parameter estimates were reported to be highly sensitive to phase delay, noise, residual patient activity and model mismatch [126, 147, 161, 180, 215, 232]. The results in table 4.1 confirm potential shortcomings of directly using OLS regression for the identification of lung parameters. In part due to these problems, it appears there is still a strong sense of skepticism among many authors regarding the application of model-based approaches in mechanically ventilated patients [81]. In this work, in order to increase the robustness of respiratory parameter estimation, several new techniques have been proposed. Firstly, the respiratory time constant was estimated over many expirations using a parallel slopes regression model that is robust against iPEEP and expiratory patient efforts, as long as they are rare enough. Secondly, the remaining parameters were estimated using an integrated form of the equation of motion (EqM), which was fitted to explain complete breaths rather than individual data points. Very importantly, the electrical and pneumatic measurements were only incorporated as time product values, which is why the model is highly robust against signal delays, pneumatic artifacts and the shape of the EMG envelope. In both identification steps, it seems beneficial to use robust regression instead of OLS. The results show that the best performance is achieved through a combination of all aforementioned techniques, refer to table 4.1. One of the strengths of the herein conducted data analysis is the size of the patient cohort ($n = 43$), which was considerably larger than previous studies treating model-based approaches.

Several limitations of the data analysis have to be discussed. Most of the patients were only undergoing ventilation for the duration of the bronchoscopy and had relatively healthy lungs, refer to clinical characteristic in table A.1. Only 16 of the analyzed patients were diagnosed with a COPD and of those patients only 3 showed substantial airflow limitation. Therefore, there is not yet sufficient evidence for the validity of model-based P_{mus} estimation in severe COPD patients. As already suggested in the literature [111], it can be assumed that the lung mechanics models will have to be extended to better represent such patients, e.g., by using two time constants. Recently, this idea has been explored by demonstrating that a switching time constant can be identified when the expiration is passive, refer to the papers by Hennigs et al. [90, 91], which were co-written by the author of this thesis. Overall, further studies will be needed to prove validity of the model-based approach in long-term ventilated patients, including severe ARDS and COPD patients.

Another limitation concerns the composition of the input data, which does not reflect normal assisted ventilation. On the analyzed data, models were fitted to all available PS levels (CPAP and PSV at 5, 10 and 15 cmH₂O), which entailed a

wide range of different patient-ventilator workload shares, and thus, facilitated the identification of lung parameters. Variations in VTP_{rs} tend to increase visibility of the lung elastance E , while variations in ETP improve the identifiability of K_{EMG} . As discussed in section 4.2.2, at least two linear independent set points are needed to solve the regression in equation (4.13), which was clearly satisfied in the data with four different support/activity levels. An important research direction for the future is how identifiability can be achieved with the least obtrusive maneuvers possible. It is reasonable to assume that two PS levels would be sufficient in most patients, and it might even be possible to completely rely on the natural variation of breathing activity. An attractive ventilation mode to be investigated in the future is noisy PSV, which mimics the random variations of natural breathing, and is well tolerated by most patients [214].

Finally, the usefulness of the model-based estimator beyond the herein discussed quantification of patient effort should be emphasized. It has already been suggested by Younes et al. [247] that a calculated P_{mus} signal could be used to non-invasively assess the timing of patient efforts, which would prove highly useful for diagnosing patient-ventilator asynchrony. Moreover, the model-based estimator naturally lends itself to proportional ventilation modes, where pressure support is to be given in proportion to the patient's own inspiratory effort. Thus, the calculated P_{mus} signal, when multiplied with a patient-specific support factor, could be used directly as a control input for the ventilator pressure — this idea was discussed in detail by Petersen, Graßhoff et al. [161]. Interestingly, the herein proposed P_{mus} estimator positions itself between existing proportional ventilation modes such as neurally adjusted ventilatory assist (NAVA), where support is given in proportion to the EMG, and proportional assist ventilation (PAV/PAV+) as well as proportional pressure support (PPS), where support is calculated from the flow and volume via the equation of motion (and parameters are chosen manually or identified via maneuvers).

Chapter 5

A Gaussian Process Model for Respiratory Muscle Pressure: Sensor Fusion, Inference and Learning

After establishing the viability of a linear sEMG-pressure model in chapter 3 and demonstrating the merits of a model-based estimation approach in chapter 4, this chapter will explore another possible source of information, namely, physiological knowledge about the respiratory muscle pressure P_{mus} . To this end, *Gaussian processes* (GPs) are proposed as a tool for modeling key characteristics of P_{mus} signals such as smoothness, periodicity and the average shape of efforts, which can then be exploited for estimating P_{mus} from non-invasive measurements. Previous P_{mus} models from the literature are mostly based on fixed parametric waveforms or simplistic state-space models which do not appear to account very well for the variability of spontaneous breathing curves. In contrast, GPs — as Bayesian *nonparametric* models — provide a much more expressive way of encoding the statistics of the data via dedicated kernel functions. The use of GPs as a modeling tool opens up new perspectives on the P_{mus} estimation problem from chapter 4. The estimation can now be posed as a *Bayesian inference* problem, which means that the prior GP is conditioned on the data to obtain the posterior. Based on this idea, this chapter will develop a comprehensive GP-based framework for non-invasive P_{mus} estimation. It will touch upon (1) Bayesian sensor fusion of ventilatory and EMG data, (2) efficient solutions to the inference via equivalent state-space models and (3) learning GP hyperparameters from esophageal pressure data.

Section 5.1 gives an overview of previously proposed models for respiratory muscle pressure. The following section 5.2 covers the novel GP framework, introducing the different Bayesian models and closed-form solutions to the inference. In section 5.3, the data analysis procedure is described and results are provided in section 5.4. Finally, in chapter 5.5, results will be discussed and a brief outlook on future work will be given. None of the results presented in this chapter have previously been published elsewhere.

5.1 Introduction and previous work

This section provides an overview of previously proposed models for the respiratory muscle pressure P_{mus} — it begins by surveying parametric functions and then proceeds to discuss dynamic state-space models.

One of the most widely adopted models for the shape of respiratory signals is to use sinusoidal functions, which has already been proposed by Otis [153] as early as 1950. One of the first *parametric* formulations for P_{mus} goes back to the 1966 simulation study of Jodat [104], who proposed to use a combination of two sinusoidal functions, corresponding to the inspiration and expiration, respectively. Based upon Jodat’s work, similar sinusoidal functions were later used to represent P_{mus} in several works, refer to [16, 99, 196]. Notably, a sinusoidal half-wave function plotted in figure 5.1 is the preset waveform for spontaneous efforts in the ASL 5000, which is an active mechanical test lung used for training and ventilator performance verification [99]. Also, the numerical simulation studies by Athanasiades et al. [16] employed a sinusoidal function as the driving pressure input. Despite its successful use for simulation purposes, the sinusoidal model is of course highly simplistic and does not reflect very well what is observed in most mechanically ventilated patients. One of the first empirical investigations of the inspiratory effort waveforms was presented in [138] by Milic-Emili and Zin. The authors evaluated the shape of occlusion pressure waveforms in humans and found an approximately quadratic curve for the muscle contraction (i.e., for the rising part of the P_{mus} curve). Based on these results, a quadratic model for the contraction has been employed for both simulation by Yamada et al. [242] and later for estimation by Saatçi and Akan [182]. Further studies concerning the muscular pressure waveform were conducted by Mecklenburgh and Mapleson [136], who fitted quadratic polynomials to both the contraction and relaxation — the resulting curve is shown in figure 5.1. The data presented in [26, 252] suggest a more exponential decay of pressure during muscle relaxation, which has also been implemented in the publications by Yamada et al. [242] as well as Saatçi and Akan [182] (as mentioned above, both combine the exponential relaxation with a quadratic contraction). The idea of an exponential P_{mus} curve was further developed by Fresnel et al. [61] in the form of a capacitor charge and discharge model. Here, the muscle itself was modeled as a dynamic system with an exponential charging behavior and different time constants for inspiration and expiration were used — the resulting pressure curves are shown in figure 5.1. All of the parametric (sinusoidal, quadratic, exponential) P_{mus} functions proposed so far are either based on expert knowledge or on relatively small studies — none of the models have been validated in a larger cohort of patients under assisted mechanical ventilation with esophageal pressure. Thus, the evidence for using fixed, parametric models is very limited. In general, it seems rather unrealistic that the large variability of spontaneous breathing curves in mechanically ventilated patients could be well represented by using a fixed P_{mus} template function.




Fig. 5.1: Examples of parametric P_{mus} functions. The left panel shows the sinusoidal half-wave function with a short hold-period used as a preset in the ASL 5000 simulator [99]. The middle panel depicts the piece-wise quadratic model derived by Mecklenburgh and Mapleson [136]. The right panel depicts the family of exponential models described by Fresnel et al. [61].

The work by Fresnel et al. [61] already conveys the notion that P_{mus} could be described as a *dynamical system* in terms of (deterministic or stochastic) *state-space models*. In the following, the available dynamical system models for P_{mus} are surveyed briefly. The most comprehensive family of models for spontaneous breathing activity has been introduced by Smith, Molkov et al. [142–144, 209] across several publications. Their work is mainly concerned with the control of breathing by the respiratory center and thus has a strong *neurobiological* focus, featuring a network of (non-spiking) central pattern generators. Molkov et al. [144] already demonstrated that the neural output of the model can be used as a driver for phrenic nerve activity to simulate realistic breathing patterns. Despite its wide-ranging capabilities, the application of this model to real-world data would be accompanied by many problems, because it is a high-dimensional, nonlinear model and has a very large number of unknown parameters. Clearly, the model was designed as a *forward model* and was not intended to be used as an *inverse model*.

Other authors have proposed much simpler *stochastic* models for representing P_{mus} , refer to [159, 182, 196]. The model by Saatçi et al. [182] uses the parametric P_{mus} template described above and incorporates a random walk for the dynamics of amplitude changes. Despite enabling adaptation of the amplitude, Saatçi’s model would fail to represent patients whose respiratory activity does not match the fixed template (consisting of a quadratic contraction and exponential relaxation). Other authors have proposed to use *second-order systems* driven by white noise for representing P_{mus} [159, 196]. Scheel et al. [196] have used a harmonic oscillator whose frequency was set to the breathing rate of patients—this model gives rise to a sinusoidal waveform, and variations of amplitude and phase are enabled through a process noise term. Petersen [159] proposed to use a damped harmonic oscillator and determined its parameters on esophageal pressure measured from mechanically ventilated patients. All of

the previously proposed reduced-order models for P_{mus} are rather simplistic and only capture certain parts of its stochastic properties while omitting others. No comprehensive stochastic model has been proposed that accounts for all relevant characteristics of P_{mus} . An important advantage of the stochastic modeling approach is that it can be easily combined with respiratory mechanics models, allowing to use *Bayesian state estimation* on the pneumatic measurements to determine P_{mus} non-invasively. To this end, both linear [159, 196] and nonlinear [182] Kalman filters have been employed.

Another important research direction is the *regularization* approach by Vicario et al. [228] and Reinders et al. [176]. In both of these works, the goal was to estimate P_{mus} non-invasively using only the pneumatic ventilator data. The estimation was posed as a regularized optimization problem and P_{mus} models were incorporated implicitly by means of additional cost-terms and constraints that favor certain physiological solutions. Vicario et al. [228] used a quadratic loss for fitting the pneumatic model and added linear constraints, enforcing monotonicity of P_{mus} during the muscle contraction and relaxation, respectively. By constraining the possible shapes of respiratory efforts, model assumptions about P_{mus} are incorporated into the optimization. Reinders et al. [176] proposed another approach in which the second derivative of P_{mus} is regularized via the ℓ_1 norm, which in turn implies a smooth P_{mus} model. The approaches by Vicario et al. [228] and Reinders et al. [176] have so far not been validated against esophageal pressure in human patients, thus, it is not yet clear what kind of benefits the proposed regularization terms provide in practice.

5.2 Bayesian framework

This work investigates *Gaussian processes* (GPs) for modeling the respiratory muscle pressure P_{mus} . Building upon earlier works [159, 182, 196], the chapter proposes a comprehensive family of stochastic models that implement the key characteristics of P_{mus} signals, namely smoothness, periodicity, average shape of efforts and variance of efforts. The GPs are then incorporated into a *Bayesian sensor fusion* model, allowing to estimate P_{mus} non-invasively from pneumatic ventilator data and from EMG measurements in an *end-to-end* fashion. Both a closed-form solution (using batch GP regression) and a recursive solution (using Kalman smoothers) will be proposed. An important advantage of the herein proposed GP framework is that prior knowledge about the dynamics of P_{mus} can be formulated in a concise way and exploited during inference, thus, leading to more physiological estimates.

5.2.1 Gaussian process model for P_{mus}

The key assumption used in this chapter is

$$P_{\text{mus}}(t) \sim \mathcal{GP}(m_{P_{\text{mus}}}(t), k_{P_{\text{mus}}}(t, t')), \quad (5.1)$$

where $k_{P_{\text{mus}}}(t, t')$ models the covariance of the muscle pressure and $m_{P_{\text{mus}}}(t)$ its mean. One of the main goals in this chapter is to introduce and compare different GP models, i.e., different mean and covariance functions that implement prior knowledge about P_{mus} . The following characteristics of respiratory muscle pressure are to be incorporated into the GP model.

Smoothness. The first characteristic of respiratory muscle pressure to be considered is smoothness. The temporal dynamics of respiratory pressure generation are determined by a multitude of factors including, among others, the force response of motor units, the activation pattern and the biomechanics of the muscle [38, 208]. The impulse response of the diaphragm as a whole can be measured in vivo using phrenic nerve stimulation. The pressure response to a single bilateral twitch was characterized by Bellemare et al. [30]: the maximum pressure was reached after ~ 80 ms and the total duration of the pressure twitch was longer than 200 ms in all subjects. These results give a first impression of the temporal dynamics of P_{mus} : the pressure can be expected to follow a smooth trajectory within a time-frame of *at least* 200 ms. During normal breathing, the time constants will even tend to be longer, because the neural drive itself also takes some time to build up. To approximate the (short-term) temporal dynamics of P_{mus} , two covariance functions were tested, namely the squared exponential and the Matérn kernel. The former provides a model for very smooth functions (infinitely differentiable), while the latter models rougher functions and the roughness/smoothness can be adjusted by the parameter ν , refer to section 2.1.3 for details. Thus, the P_{mus} covariance using either a Matérn or the squared exponential kernel can be stated as

$$k_{P_{\text{mus}}}(t, t') = \sigma_{P_{\text{mus}}}(t) \cdot k_{\{\text{Mat}, \text{SE}\}}(t, t') \cdot \sigma_{P_{\text{mus}}}(t'), \quad (5.2)$$

where it is assumed that both k_{Mat} and k_{SE} have unit variance, thus, the (time-varying) standard deviation of the kernel is fully defined by $\sigma_{P_{\text{mus}}}(t)$. For both kernels, the length-scale parameters (ℓ_{Mat} and ℓ_{SE}) were optimized to fit the dynamics of P_{mus} .

Periodicity. A very prominent characteristic of the P_{mus} signal is periodicity. A healthy respiratory center produces an approximately periodic breathing rhythm via its central pattern generators. The pattern is modulated by physiological feedback loops, refer to section 2.2.1, and is of course not strictly periodic, i.e., there is a natural variation of the pattern between breaths. Here the quasi-periodic kernel

$$k_{P_{\text{mus}}}(t, t') = \sigma_{P_{\text{mus}}}(t) \cdot k_{\text{QP}}(\phi(t), \phi(t')) \cdot \sigma_{P_{\text{mus}}}(t') \quad (5.3)$$

was used to approximate this behavior, where k_{QP} is defined by the product of a 2π -periodic kernel k_{Per} and a squared exponential kernel k_{SE} as in section 2.1.3. As before, k_{QP} is assumed to have unit variance such that $\sigma_{P_{\text{mus}}}(t)$ determines

the standard deviation. The function $\phi(t)$ was used to map from the time domain to the phase domain, with each breath being mapped to the interval $[0, 2\pi]$. The quasi-periodic kernel models locally periodic functions and the decay of periodicity is determined by the length-scale ℓ_{SE} of the squared exponential kernel. Within periods (on a scale shorter than 2π), the kernel produces smooth functions and short-term behavior is determined by the length-scale parameter ℓ_{Per} . All hyperparameters were optimized using esophageal pressure-derived measurements of P_{mus} in mechanically ventilated patients. The performance of the quasi-periodic kernel was compared to that of the aforementioned smooth kernel functions.

Average effort shape. As discussed above, several parametric models for the average shape of P_{mus} have been proposed [61, 104, 136, 242], which however do not appear to be flexible enough to account for the large variability of respiratory efforts in mechanically ventilated patients. In contrast, the GP framework constitutes a nonparametric model, where prior knowledge about the average shape of $P_{mus}(t)$ can be incorporated quite naturally via the mean function $m_{P_{mus}}(t)$. As a first baseline model the constant mean function

$$m_1(t) = \text{const.} \quad (5.4)$$

was used, which does not induce any further assumptions about the P_{mus} waveform. The second considered mean function was

$$m_2(t) = m_t(\phi(t)), \quad (5.5)$$

which is a periodic function depending on the patient's breathing phase $\phi(t)$ and on a template function m_t . The function m_t represents the average waveform of muscle efforts across one breath comprising its inspiration and expiration. In this work, the template m_t was learned from a large number of patients using esophageal pressure (P_{es}) measurements. Despite using a fixed template m_t for the mean, the herein proposed GP framework will be able to represent many different P_{mus} waveforms, because deviations from the mean are modeled through the kernel. One of the goals of this work is to investigate the capabilities of GP models using $m_1(t)$ and $m_2(t)$, respectively.

Time-varying variance. Another important characteristic of the P_{mus} signal is its non-stationarity, and the most noticeable is that the function distribution changes between inspiration and expiration. In most patients, respiratory activity is highest during inspirations, resulting in a larger variance of P_{mus} values. Thus, similar to the mean function, two different models for the GP's standard deviation $\sigma_{P_{mus}}(t)$ were tested. The simple baseline model

$$\sigma_1(t) = \text{const.} \quad (5.6)$$

implements the assumption that the standard deviation of P_{mus} is constant over time. As a non-stationary model for $\sigma_{P_{\text{mus}}}(t)$, the herein proposed framework uses

$$\sigma_2(t) = \sigma_t(\phi(t)), \quad (5.7)$$

where $\phi(t)$ is the patient's breathing phase and σ_t is a template function modeling the change of standard deviation throughout the breathing cycle. Again, both σ_1 and σ_2 were learned directly from esophageal pressure recordings across a large number of patients.

One of the main goals in this chapter is to investigate the performance of GP models that implement different assumptions (smoothness, periodicity, time-varying mean and variance) about P_{mus} . To this end, a total of six different models were tested, all of which are listed in table 5.1: each of the three kernels (squared exponential, Matérn, quasi-periodic) is combined with the stationary model (using $m_1(t)$ and $\sigma_1(t)$) as well as with the non-stationary model (using $m_2(t)$ and $\sigma_2(t)$).

Table 5.1: Overview of the Gaussian process models that are used to represent $P_{\text{mus}}(t)$. The prior mean is denoted by $m_{P_{\text{mus}}}(t)$, the prior covariance by $k_{P_{\text{mus}}}(t, t')$.

	$m_{P_{\text{mus}}}(t)$	$k_{P_{\text{mus}}}(t, t')$	stationary	locally smooth	periodic	hyper-parameters
GP _{SE1}	$m_1(t)$	$\sigma_1(t) \cdot k_{\text{SE}}(t, t') \cdot \sigma_1(t')$	✓	✓		ℓ_{SE}
GP _{SE2}	$m_2(t)$	$\sigma_2(t) \cdot k_{\text{SE}}(t, t') \cdot \sigma_2(t')$		✓		ℓ_{SE}
GP _{Mat1}	$m_1(t)$	$\sigma_1(t) \cdot k_{\text{Mat}}(t, t') \cdot \sigma_1(t')$	✓	✓		ν, ℓ_{Mat}
GP _{Mat2}	$m_2(t)$	$\sigma_2(t) \cdot k_{\text{Mat}}(t, t') \cdot \sigma_2(t')$		✓		ν, ℓ_{Mat}
GP _{QP1}	$m_1(t)$	$\sigma_1(t) \cdot k_{\text{QP}}(\phi(t), \phi(t')) \cdot \sigma_1(t')$		✓	✓	$\ell_{\text{SE}}, \ell_{\text{Per}}$
GP _{QP2}	$m_2(t)$	$\sigma_2(t) \cdot k_{\text{QP}}(\phi(t), \phi(t')) \cdot \sigma_2(t')$		✓	✓	$\ell_{\text{SE}}, \ell_{\text{Per}}$

SE: squared exponential kernel; Mat: Matérn kernel; Per: periodic kernel; QP: quasi-periodic kernel;

5.2.2 Probabilistic sensor fusion model

After having introduced different GP models for P_{mus} , their application in mechanically ventilated patients will be treated below. The goal is to incorporate the GP into the estimation problem given in chapter 4, i.e., P_{mus} is to be estimated non-invasively from pneumatic ventilator data and respiratory sEMG measurements. Similar to chapter 4, this chapter also employs a sensor fusion approach, where all available measurements are combined in a model to estimate the unknown quantity (P_{mus}). To this end, the problem from section 4.2.2 is reformulated in a *probabilistic* way and posed as a *Bayesian inference* problem. To briefly summarize the main idea, the GP *prior* for P_{mus} is to be conditioned on the available measurements to derive the *posterior* GP distribution.

Similar to Petersen [159], this work proposes a probabilistic model that comprises the following components: (1) a model of respiratory mechanics incorporating the pneumatic measurements ($P_{\text{aw}}, V, \dot{V}$) from the ventilator, (2) a neuromechanical coupling model incorporating the sEMG measurements, and (3) the GP model for P_{mus} described in the previous section. Gaussian priors

$$E \sim \mathcal{N}(\mu_E, \sigma_E^2) \quad (5.8)$$

$$R \sim \mathcal{N}(\mu_R, \sigma_R^2) \quad (5.9)$$

$$K_{\text{EMG}} \sim \mathcal{N}(\mu_{K_{\text{EMG}}}, \sigma_{K_{\text{EMG}}}^2) \quad (5.10)$$

were used for all parameters and the GP prior

$$P_{\text{mus}}(t) \sim \mathcal{GP}(m_{P_{\text{mus}}}(t), k_{P_{\text{mus}}}(t, t')) \quad (5.11)$$

from section 5.2.1 was used for $P_{\text{mus}}(t)$. The measurement equation was given by

$$P_{\text{aw}}(t) = E \cdot V(t) + R \cdot \dot{V}(t) + P_0(t) - P_{\text{mus}}(t) + \nu_1(t) \quad (5.12)$$

$$0 = K_{\text{EMG}} \cdot \text{EMG}(t) - P_{\text{mus}}(t) + \nu_2(t), \quad (5.13)$$

which can be easily recognized as the equation of motion, refer to section 2.2.3, and the linear neuromechanical coupling model discussed in section 3.2.3. The residuals $\nu_1(t)$ and $\nu_2(t)$ of the two equations were modeled as zero-mean GPs

$$\nu_1(t) \sim \mathcal{GP}(0, k_{\nu_1}(t, t')) \quad (5.14)$$

$$\nu_2(t) \sim \mathcal{GP}(0, k_{\nu_2}(t, t')), \quad (5.15)$$

which were assumed to both have a white noise and a (short-term) autocorrelated component. Their covariances were defined as

$$k_{\nu_1}(t, t') = k_{\text{Mat}}(t, t') + \sigma_{\text{n}_1}^2 \delta_{tt'} \quad (5.16)$$

$$k_{\nu_2}(t, t') = k_{\text{Mat}}(t, t') + \sigma_{\text{n}_2}^2 \delta_{tt'}, \quad (5.17)$$

where the Matérn kernel accounts for autocorrelation for example due to phase delays or cardiac artifacts in the EMG, and the white noise kernel accounts for all remaining errors. The length-scale and variance of the Matérn kernel ($\ell_{\text{Mat}}, \sigma_{\text{Mat}}^2$ for both ν_1 and ν_2) as well as the white noise variances $\sigma_{\text{n}_1}^2$ and $\sigma_{\text{n}_2}^2$ were treated as hyperparameters which have to be learned or tuned manually. Finally, the time-varying parameter $P_0(t)$ accounts for pressure offsets caused by PEEP/iPEEP and volume drift. As the offset may change over time, it was described by the GP prior

$$P_0(t) \sim \mathcal{GP}(\text{PEEP}, k_{P_0}(t, t')), \quad (5.18)$$

where a Matérn kernel $k_{P_0}(t, t') = k_{\text{Mat}}(t, t')$ models the slow dynamic of iPEEP changes and volume drift (or of any other offsets within the data). Note that in this chapter, P_0 was defined as a time-varying process, which is different to chapter 4, where P_0 was considered to be constant and offsets were accounted for via V_{rs} .

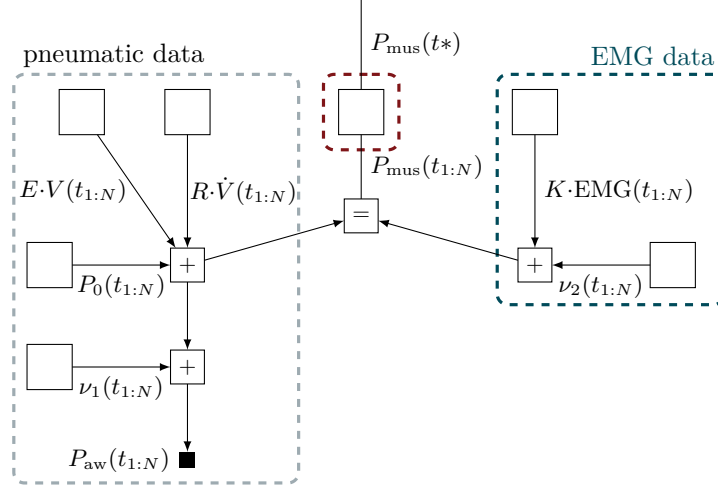


Fig. 5.2: A factor graph of the joint probability density function described by equations (5.8)–(5.18) at given time points t_1, \dots, t_N . The nodes and edges of the graph represent factors of the distribution and multivariate random variables, respectively. Each node is a function of the edges that it is connected with. The ‘+’, ‘■’, and ‘=’ nodes are defined according to Loeliger et al. [125]. A solution to the inference problem can be obtained by calculating the marginal distribution of $P_{\text{mus}}(t_*)$.

5.2.3 Inference of P_{mus}

The inference problem to be solved is the following: given the probabilistic model stated in equations (5.8)–(5.18) and measurements of $P_{\text{aw}}(t)$, $V(t)$, $\dot{V}(t)$ and $\text{EMG}(t)$ at time points t_1, \dots, t_N , calculate the posterior $P_{\text{mus,post}}(t_*)$, where t_* is a given test point. In the following, a column vector of measurements from a variable $y(t)$ is denoted as $y(t_{1:N})$ and a row vector as $y(t_{1:N})^\top$. As a first step towards solving the problem, it can be easily seen that

$$E V(t) \sim \mathcal{GP}(\mu_E V(t), V(t) \sigma_E^2 V(t')) \quad (5.19)$$

$$R \dot{V}(t) \sim \mathcal{GP}(\mu_R \dot{V}(t), \dot{V}(t) \sigma_R^2 \dot{V}(t')) \quad (5.20)$$

$$K_{\text{EMG}} \text{EMG}(t) \sim \mathcal{GP}(\mu_{K_{\text{EMG}}} \text{EMG}(t), \text{EMG}(t) \sigma_{K_{\text{EMG}}}^2 \text{EMG}(t')) \quad (5.21)$$

and thus, all priors of the model are GPs. When the equations are evaluated on a finite set of input points, the joint distribution over all random variables can be represented by the *factor graph* in figure 5.2, refer to [125] for a comprehensive introduction to this type of probabilistic graphical models. The graph is tree-structured and all nodes either have Gaussian probability distributions or correspond to linear building blocks (such as the addition and equality node described in [125]). Thus, the marginalization with respect to $P_{\text{mus}}(t_*)$ is tractable. The closed-form solution is provided in the following. The solution to the inference problem can be achieved by first conditioning

the prior GP on the EMG data, which is implemented with equation (5.22), and then conditioning on the pneumatic data, which is implemented with equation (5.25). Both equations (5.22) and (5.25) are simply variants of the general GP regression equation introduced in section 2.1.3. First calculate

$$\begin{aligned} m_{\text{post}_1}(t) &= -m_{P_{\text{mus}}}(t) + k_{P_{\text{mus}}}(t, t_{1:N})K_{\text{sum}_1}^{-1}m_{\text{sum}_1} \\ k_{\text{post}_1}(t, t') &= k_{P_{\text{mus}}}(t, t') - k_{P_{\text{mus}}}(t, t_{1:N})K_{\text{sum}_1}^{-1}k_{P_{\text{mus}}}(t_{1:N}, t'), \end{aligned} \quad (5.22)$$

where

$$m_{\text{sum}_1} = m_{P_{\text{mus}}}(t_{1:N}) - \mu_{K_{\text{EMG}}} \text{EMG}(t_{1:N}) \quad (5.23)$$

$$K_{\text{sum}_1} = k_{P_{\text{mus}}}(t_{1:N}, t_{1:N}) + \text{EMG}(t_{1:N})\sigma_{K_{\text{EMG}}}^2 \text{EMG}(t_{1:N})^\top + k_{\nu_1}(t_{1:N}, t_{1:N}), \quad (5.24)$$

and $k(t_{1:N}, t)$ denotes the column vector of covariances between $t_{1:N}$ and t , $k(t, t_{1:N})$ denotes its transpose, $k(t_{1:N}, t_{1:N})$ denotes a square matrix of covariances between all time points. Then, using the posterior from equation (5.22), calculate

$$\begin{aligned} m_{\text{post}_2}(t) &= -m_{\text{post}_1}(t) + k_{\text{post}_1}(t, t_{1:N})K_{\text{sum}_2}^{-1}m_{\text{sum}_2} \\ k_{\text{post}_2}(t, t') &= k_{\text{post}_1}(t, t') - k_{\text{post}_1}(t, t_{1:N})K_{\text{sum}_2}^{-1}k_{\text{post}_1}(t_{1:N}, t'), \end{aligned} \quad (5.25)$$

where

$$m_{\text{sum}_2} = P_{\text{aw}}(t_{1:N}) - \mu_E V(t_{1:N}) - \mu_R \dot{V}(t_{1:N}) + m_{\text{post}_1}(t_{1:N}) - \text{PEEP} \quad (5.26)$$

$$\begin{aligned} K_{\text{sum}_2} &= k_{\text{post}_1}(t_{1:N}, t_{1:N}) + V(t_{1:N})\sigma_E^2 V(t_{1:N})^\top + \dot{V}(t_{1:N})\sigma_R^2 \dot{V}(t_{1:N})^\top \\ &\quad + k_{P_0}(t_{1:N}, t_{1:N}) + k_{\nu_2}(t_{1:N}, t_{1:N}). \end{aligned} \quad (5.27)$$

Finally, the posterior GP for P_{mus} is given by

$$P_{\text{mus,post}}(t_*) \sim \mathcal{GP}(m_{\text{post}_2}(t_*), k_{\text{post}_2}(t_*, t_*)). \quad (5.28)$$

To numerically calculate a solution to equation (5.22) and equation (5.25), Cholesky decomposition was used instead of directly calculating the matrix inverse (in order to improve numerical stability).

The herein proposed solution to estimating P_{mus} offers several advantages over simpler non-Bayesian methods. Firstly, the P_{mus} signal is reconstructed using all available sources of information, namely the EMG envelope, pneumatic data and the prior GP. The solution is calculated in an *end-to-end* fashion by means of Bayesian inference. In contrast, the model-based estimators in chapter 4 comprised two separate steps (namely parameter estimation and P_{mus} reconstruction), and P_{mus} was calculated using only pneumatic data. Secondly, the herein proposed estimator incorporates prior knowledge about lung parameters and about the distribution of P_{mus} . In this way, the posterior is regularized towards physiologically plausible solutions. Thirdly, the derived posterior of P_{mus} is itself a GP, which allows to evaluate it at any given test point t_* .




Fig. 5.3: A comparison of exact kernel functions and their state-space approximations. The squared exponential kernel (plotted as a light blue line) was approximated using the n th order Taylor expansion of its spectral density. The Matérn kernel (plotted as a light red line) with $\nu = \frac{5}{2}$ has exact representation with $n = 3$ states. The quasi-periodic kernel (plotted as a light turquoise line) was approximated using m harmonic oscillators and a 5th order Taylor expansion of the squared exponential kernel. Kernel hyperparameters were optimized to approximate P_{mus} measurements, refer to section 5.3.1.

Thus, the framework allows for very easy interpolation and upsampling/resampling of estimates. Fourthly, in contrast to most previous approaches, the proposed method allows to quantify the posterior confidence, which is an important feature for bringing such methods into clinical application.

5.2.4 Equivalent state-space formulation

As discussed in section 2.1.4, the practical application of GP-based inference to large datasets is challenging due to heavy memory and computational requirements. The main bottleneck comes from storing the kernel matrix K and solving $K^{-1}y$, which have $\mathcal{O}(N^2)$ and $\mathcal{O}(N^3)$ complexity, respectively, where N is the number of data points. Thus, the batch solution proposed in the previous section is intractable for datasets exceeding $\sim 10^4$ data points, which however is quickly reached in mechanical ventilation time-series data with recording lengths of >10 min and sampling rates of 50–100 Hz. To address these limitations, equivalent state-space formulations for the GP model can be derived. To this end, the model from section 5.2.2 was reformulated as a stochastic differential equation and then inference was done by means of Kalman filters/smoothers, which have quasi-linear time complexity in the number of data points, refer to section 2.1.4 for more details.

First consider a reformulation of the P_{mus} models from section 5.2.1. For the smooth model $\mathcal{GP}(m_{P_{\text{mus}}}(t), \sigma_{P_{\text{mus}}}(t) \cdot k_{\{\text{Mat}, \text{SE}\}}(t, t') \cdot \sigma_{P_{\text{mus}}}(t'))$ using the squared exponential or Matérn kernel, one can employ the stochastic differential equation

$$\begin{aligned} dx(t) &= Fx(t)dt + Ld\beta(t) \\ P_{\text{mus}}(t) &= \sigma_{P_{\text{mus}}}(t)Hx(t) + m_{P_{\text{mus}}}(t), \end{aligned} \quad (5.29)$$

where $x(t)$ is the state of the model, $\beta(t)$ is a Brownian motion with diffusion constant Q , the functions $\sigma_{P_{\text{mus}}}(t)$ and $m_{P_{\text{mus}}}(t)$ are defined as in section 5.2.1, and the dynamics of the (unit variance) kernel are approximated by the model matrices F , L and H as discussed in section 2.1.4. For the Matérn kernel, the representation in state-space form is exact [85]. For some kernels that can not be written as rational functions such as the squared exponential kernel, the state-space reformulation would require an infinite state dimension. Therefore, as a finite-dimensional approximations to the squared exponential kernel, the Taylor series approach by Hartikainen and Särkkä [85] was used. The exact kernel and its approximation are depicted in figure 5.3. Similarly, for the quasi-periodic model $\mathcal{GP}(m_{P_{\text{mus}}}(t), \sigma_{P_{\text{mus}}}(t) \cdot k_{\text{QP}}(\phi(t), \phi(t')) \cdot \sigma_{P_{\text{mus}}}(t'))$, the stochastic differential equation

$$\begin{aligned} dx(\phi(t)) &= Fx(\phi(t))d\phi(t) + Ld\beta(\phi(t)) \\ P_{\text{mus}}(t) &= \sigma_{P_{\text{mus}}}(t)Hx(\phi(t)) + m_{P_{\text{mus}}}(t) \end{aligned} \quad (5.30)$$

was used, which is defined over the breathing phase $\phi(t)$, where the matrices F , L and H are again constructed to approximate the dynamics of the P_{mus} kernel. As discussed in 2.1.4, the periodic kernel would also give rise to an infinite-dimensional state-space, but can be approximated using a finite number of parallel harmonic oscillators, see [212] for details. The approximation to the quasi-periodic kernel is illustrated in figure 5.3.

The noise processes ν_1 and ν_2 were also converted to state-space form, leading to

$$\begin{aligned} dx_{\nu_1}(t) &= F_{\nu_1}x_{\nu_1}(t)dt + L_{\nu_1}d\beta_{\nu_1}(t) \\ \nu_1(t) &= H_{\nu_1}x_{\nu_1}(t) + \epsilon_{n_1}(t) \end{aligned} \quad (5.31)$$

and

$$\begin{aligned} dx_{\nu_2}(t) &= F_{\nu_2}x_{\nu_2}(t)dt + L_{\nu_2}d\beta_{\nu_2}(t) \\ \nu_2(t) &= H_{\nu_2}x_{\nu_2}(t) + \epsilon_{n_2}(t), \end{aligned} \quad (5.32)$$

where $x_{\nu_1}(t)$ and $x_{\nu_2}(t)$ are the state vectors, $\beta_{\nu_1}(t)$ and $\beta_{\nu_2}(t)$ are Brownian motions with diffusion constants Q_{ν_1} and Q_{ν_2} , the matrices F_{ν_1} , L_{ν_1} , H_{ν_1} and F_{ν_2} , L_{ν_2} , H_{ν_2} are constructed to represent Matérn kernels, ϵ_{n_1} and ϵ_{n_2} are white noise processes with spectral densities $\sigma_{n_1}^2$ and $\sigma_{n_2}^2$ respectively. The state-space models (5.31) and (5.32) allow for exact representations of the two noise processes ν_1 and ν_2 . The offset model P_0 was also reformulated to

$$\begin{aligned} dx_{P_0}(t) &= F_{P_0}x_{P_0}(t)dt + L_{P_0}d\beta_{P_0}(t) \\ P_0(t) &= H_{P_0}x_{P_0}(t) + \text{PEEP}, \end{aligned} \quad (5.33)$$

where $x_{P_0}(t)$ is the state vector, $\beta_{P_0}(t)$ denotes a Brownian motion with diffusion constant Q_{P_0} , and the matrices F_{P_0} , L_{P_0} , H_{P_0} are constructed to represent a GP with Matérn kernel.

Finally, combining all components enables to construct the *augmented model*

$$\begin{aligned} dz(t) &= F_a z(t)dt + L_a d\beta_a(t) \\ y(t) &= H_a(t) z(t) + m_a(t) + \epsilon_a(t), \end{aligned} \quad (5.34)$$

where $z^\top = [x^\top, E, R, K_{\text{EMG}}, x_{P_0}^\top, x_{\nu_1}^\top, x_{\nu_2}^\top]$ is the augmented state vector and β_a is a Brownian motion with diffusion constant Q_a . The model uses the matrices

$$\begin{aligned} F_a &= \text{blkdiag}(F, 0, 0, 0, F_{P_0}, F_{\nu_1}, F_{\nu_2}), \\ L_a &= \text{blkdiag}(L, 0, 0, 0, L_{P_0}, L_{\nu_1}, L_{\nu_2}), \\ Q_a &= \text{blkdiag}(Q, 0, 0, 0, Q_{P_0}, Q_{\nu_1}, Q_{\nu_2}), \end{aligned} \quad (5.35)$$

where $\text{blkdiag}(\cdot)$ denotes an operator that arranges matrices into a block diagonal matrix. Notice that this model implements the assumption that the parameters E , R and K_{EMG} are constant, i.e., they do not have any temporal dynamics and are not subject to noise. The measurement equation in (5.34) is

$$\underbrace{\begin{bmatrix} 0 \\ P_{\text{aw}}(t) \end{bmatrix}}_{y(t)} = \underbrace{\begin{bmatrix} \sigma_{P_{\text{mus}}}(t) H & 0 & 0 & \text{EMG}(t) & 0 & 0 & H_{\nu_2} \\ \sigma_{P_{\text{mus}}}(t) H & V(t) & \dot{V}(t) & 0 & H_{P_0} & H_{\nu_1} & 0 \end{bmatrix}}_{H_a(t)} z(t) + m_a(t) + \epsilon_a(t), \quad (5.36)$$

where

$$m_a(t) = \begin{bmatrix} -m_{P_{\text{mus}}}(t) + \mu_{K_{\text{EMG}}} \text{EMG}(t) \\ -m_{P_{\text{mus}}}(t) + \mu_E V(t) + \mu_R \dot{V} + \text{PEEP} \end{bmatrix}, \quad \epsilon_a(t) = \begin{bmatrix} \epsilon_{n_1}(t) \\ \epsilon_{n_2}(t) \end{bmatrix}, \quad (5.37)$$

which concludes the definition of the equivalent state-space model. Notice that the periodic process given in equation (5.30) is not yet compatible with the augmented model, because it was defined over the breathing phase $\phi(t)$ and not over t . Thus, the state-space model (5.30) was rewritten to have a time derivative, leading to $d\tilde{x}(t) = \phi'(t)F\tilde{x}(t)dt + \phi'(t)Ld\beta(t)$, where the substitution $\tilde{x}(t) = x(\phi(t))$ is used. In this form, it can be directly incorporated into the augmented model (5.34).

Sequential inference of P_{mus}

To solve the inference problem from section 5.2.3, the augmented continuous state-space model was discretized for the time points t_1, \dots, t_N , leading to the discrete-time model

$$\begin{aligned} z(t_{k+1}) &= A_{a,k} z(t_k) + q_{a,k} \\ y(t_k) &= H_{a,k} z(t_k) + m_a(t_k) + \epsilon_{a,k}, \end{aligned} \quad (5.38)$$

where $q_{a,k} \sim \mathcal{N}(0, Q_{a,k})$ represents process noise, $\epsilon_{a,k} = \epsilon_a(t_k) \sim \mathcal{N}(0, R_a)$ represents measurement noise, the matrices $A_{a,k}$ and $Q_{a,k}$ were determined by analytical discretization as discussed in section 2.1.4, and $H_{a,k} = H_a(t_k)$. Based on the discrete model, the inference problem for P_{mus} could now be solved by means of Bayesian state estimation. The linear Kalman smoother recovers the result of the batch GP solution given in section 2.1.4 (up to errors caused by the approximation of the squared exponential and periodic kernels). To this end, the Kalman smoother has to be initialized with $z(t_0) \sim \mathcal{N}(0, \Sigma_{a\infty})$, where

$$\Sigma_{a\infty} = \text{blkdiag}(\Sigma_{\infty}, \sigma_E^2, \sigma_R^2, \sigma_{K_{\text{EMG}}}^2, \Sigma_{P_0\infty}, \Sigma_{\nu_1\infty}, \Sigma_{\nu_2\infty}), \quad (5.39)$$

and the covariances Σ_{∞} , $\Sigma_{\nu_1\infty}$, $\Sigma_{\nu_2\infty}$ and $\Sigma_{P_0\infty}$ are solutions to matrix Riccati equations corresponding to the state-space models (5.29), (5.31), (5.32) and (5.33), respectively. The state can be interpolated for any given test point t_* using forward/backward prediction from adjacent data points.

5.2.5 Relation to other models

The herein proposed model belongs to the class of *latent force models*, which were introduced by Alvarez et al. [12] as a compromise between mechanistic modeling and data driven machine learning. The authors aimed to mitigate the shortcomings of models derived from first principles, which often do not accurately represent all dynamics of a physical system, and of purely data driven approaches, which tend to perform poorly on unseen data. They proposed to augment machine learning to classical state-space models and showed that GPs can be used to learn unknown inputs — the so called *latent forces* — acting on the physical system. Alvarez et al. [12] solved the inference problem by analytically deriving the covariance function of the full augmented system, but later work has focused more on the equivalent state-space approach [188]. The model proposed in this chapter is an instance of the approach by Alvarez et al. [12], and could be described as a *latent pressure model*, because the unknown pressure input P_{mus} acting on the respiratory system is estimated by means of a latent GP. It is also worth mentioning, that the augmented state-space model naturally lends itself to control applications, which was explored in detail by Graßhoff et al. [72] and Männel, Graßhoff et al. [130].

Another technique closely related to the proposed method is *regularization*, which is often employed to solve ill-posed problems. As discussed in section 2.1.1, there is a direct relation between Bayesian inference and optimization, i.e., the maximum of the Bayesian posterior corresponds to the solution of a regularized optimization problem. Under the regularization paradigm, prior knowledge is encoded via a regularization term which — similar to a Bayesian prior — imposes a penalty on the complexity of the solution. In the case of a Gaussian process with Gaussian measurement noise, the regularizer is given by a *reproducing kernel Hilbert space* (RKHS) norm $\|f\|_{\mathcal{H}} =$

$\sqrt{\langle f, f \rangle_{\mathcal{H}}}$ with $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ being an inner product on the RKHS \mathcal{H} and the data-fit term is realized by the ℓ_2 norm [170, section 6.1]. Thus, a very close variant of the inference problem considered in this chapter can be recovered by minimization of the functional

$$J[P_{\text{mus}}(t), \theta] = \lambda \|P_{\text{mus}}(t)\|_{\mathcal{H}}^2 + \|\Gamma\theta\|_2^2 + \|\epsilon_1(t_{1:N})\|_2^2 + \|\epsilon_2(t_{1:N})\|_2^2, \quad (5.40)$$

where \mathcal{H} denotes an RKHS corresponding to the P_{mus} kernel, λ and Γ are regularization parameters, ϵ_1 and ϵ_2 correspond to the residuals in the equation of motion and EMG- P_{mus} model, θ is the vector of model parameters (and for the sake of brevity, the time-varying offset P_0 as well as the autocorrelated errors are not included). Notice that when the discrete problem is considered, i.e., when P_{mus} is only estimated at discrete time instants t_1, \dots, t_N , the cost term can be realized simply by using the *generalized Tikhonov regularization* $\|P_{\text{mus}}(t_{1:N})\|_{K^{-1}}^2$, where K is the covariance matrix. The reformulation in equation (5.40) gives important insight into the relation of the proposed Bayesian method to earlier approaches of estimating P_{mus} . Reinders et al. [176] also used a quadratic data fit term and proposed to regularize the second derivative of P_{mus} using the ℓ_1 norm, which also enforces smoothness. An earlier study by Vicario et al. [228] also used the quadratic error term, and relied on additional constraints to regularize the P_{mus} waveform, which gives rise to a quadratic programming problem.

The herein proposed GP framework provides a very natural link between the batch estimation approach and recursive state estimation methods. The augmented state-space model (5.34) shares its general structure with other state-space models designed for recursive P_{mus} estimation [159, 182, 196]. These previously published approaches have also used a sub-block of the process matrix to model the dynamics of P_{mus} and Saatçi et al. [182] also used a random walk behavior for the parameters. In contrast to the GP model introduced in this chapter, previous P_{mus} models are much more simplistic and do not appear to regularize the estimate towards physiologically plausible solutions. In fact, the P_{mus} model proposed by Scheel et al. [196] can be seen as a special case of the quasi-periodic kernel when only the first harmonic oscillator is used. Moreover, the P_{mus} model by Saatçi et al. [182] represents a simplified version of the time-varying mean model $m_t(\phi(t))$, but in contrast to GPs, it did not explicitly account for deviations from the fixed template waveform. The close link to these previously proposed methods is precisely the reason why a Bayesian framework is so useful. As a unifying framework it offers the possibility of systematically analyzing the underlying assumptions of different algorithms.

5.3 Data analysis

The Gaussian process framework was evaluated on the study data introduced in chapter 3. The data preprocessing followed the same steps as in section 3.2.2. To briefly reiterate, the lung volume signal V was calculated by integrating the flow

\dot{V} and by correcting for leakages/drifts. Envelope signals were calculated for both surface EMG channels, and the channel with the higher SNR was selected to represent the electrical activity of the respiratory muscles. In this chapter, an estimate of the patient's breathing phase $\phi(t)$ is needed for the GP models. Thus, the sign of the flow was used to detect inspirations and expirations and then construct a piece-wise linear phase function, which linearly maps inspirations to the interval $[0, \pi]$ and expirations to the interval $[\pi, 2\pi]$.

In addition to the six P_{mus} models listed in table 5.1, two additional baseline models were tested. The first baseline model is given by an uninformative P_{mus} kernel, which means that its information matrix K^{-1} is zero, and thus, no prior knowledge about the statistics of P_{mus} is incorporated. Using the uninformative prior, the probabilistic model from section 5.2.1 simplifies to a Bayesian regression model with autocorrelated errors and the inference can be solved in a straightforward manner. The second baseline model is the OLS estimator that was already employed in chapter 4. OLS can also be derived as a special case of this chapter's probabilistic framework. It is given by the maximum of the posterior when the P_{mus} prior and parameter priors are uninformative, P_0 is constant and the errors ν_1 and ν_2 are white. In that sense, OLS represents the most basic version of the Bayesian model proposed in this chapter. The P_{mus} signal is calculated from the OLS parameters using the pneumatic reconstruction that was introduced in the previous chapter in section 4.2.2.

The main goal of this chapter is to investigate if the incorporation of a GP prior provides benefits for the P_{mus} estimation and what kind of prior knowledge (smoothness, periodicity, average waveform) is most beneficial. To reduce the influence of model mismatch, for this analysis, a subset of patients was selected that had no severe airflow limitation. Thus, 30 patients with sufficiently long (>0.35 s) occlusion maneuvers and with $i\text{PEEP} < 3$ cmH₂O were included. The data were split into a training and a test dataset at a 1:1 ratio, leading to *15 training* and *15 test* patients. All reported performance metrics were evaluated on the test data. The hyperparameters and waveforms of the probabilistic model were learned on the training dataset across *all available pressure support (PS) levels*. Then the learned model was used to estimate P_{mus} on the test patients during the *5 cmH₂O PS phase*. This PS level usually leads to a sufficiently high patient activity with reasonably good patient-ventilator interaction, and thus, provides the most useful data for comparing the properties of GP models. The other PS levels are less suited for evaluating GP performances: under CPAP, lung parameters are often not identifiable, while at higher PS levels, the sEMG amplitude becomes quite small and there is a higher degree of patient-ventilator asynchrony — while these effects are interesting from a clinical point of view, they complicate the model evaluation and were (for now) not the focus of this investigation. Thus, as opposed to chapter 4, the estimation performance of GPs is only reported in *one* PS level, namely the relatively simple case of 5 cmH₂O. The inference was solved by means of Kalman smoothing using the equivalent state-space model in equation (5.34)

on measurements sampled at 20 Hz. For the P_{mus} GP, a Matérn kernel with $\nu = 5/2$ was used, leading to a state-space model with 3 states. The squared exponential kernel was approximated using a state-space model with 7 states. The quasi-periodic kernel was approximated using 6 harmonic oscillators (12 states) and a decay model with 5 states, leading to a total of 60 states.

5.3.1 Model learning

The GP models introduced in section 5.2.1 were optimized to approximate the dynamics of invasively measured P_{mus} in the training dataset. Thus, all steps described in this section presume availability of esophageal pressure on the training data (but not necessarily on the test data). Training of models was done on *full recordings*, including *all available PS levels*. The P_{es} signals were processed as in section 3.2.2, i.e., P_{mus} signals were calculated by removing cardiac artifacts and correcting for the elastic recoil of the chest wall $E_{\text{cw}}V$. As a first step, $m_1(t)$ and $\sigma_1(t)$ were estimated in each PS level across all P_{mus} samples from all training datasets. To make this estimation robust, P_{mus} outliers were detected as points that are more than three standard deviations from the mean, and any samples within 1.5 s from these outliers were excluded from the estimation. Then, $m_1(t)$ and $\sigma_1(t)$ were calculated via the empirical mean and standard deviation across all remaining samples. The time-varying mean functions $m_2(t)$ and $\sigma_2(t)$ were estimated in a similar fashion: P_{mus} signals from the training data were first mapped to the breathing phase $\phi(t)$, and ensembles of all breathing waveforms over the interval $[0, 2\pi]$ were formed in each PS level. As the average P_{mus} waveform is highly susceptible to single outliers and artifacts, a highly sensitive outlier detector inspired by the *Hampel filter* [84] was used. Briefly, any breaths in which 15% of the waveform are more than K scaled median absolute deviations [98] from the ensemble median were discarded. For the cutoff K of the Hampel filter, Miller et al. [139] proposed values between 2 and 3, but empirical tests showed that more conservative cutoffs are also feasible on the analyzed data, thus, $K = 3.25$ was used. After removing outliers, the phase-dependent functions $m_2(t)$ and $\sigma_2(t)$ were determined by the ensemble average and ensemble standard deviation in each PS level.

The hyperparameters of kernels listed in table 5.1 were learned by minimizing the negative marginal log likelihood on the P_{mus} measurements. To reduce the computation time, learning relied on stochastic gradient descent (SGD) on mini-batches. For the Matérn and the squared exponential kernels, 10 s patches were randomly cropped from the training data, and in each batch one patch from each patient and PS level was included, which resulted in a batch size of 60. For the quasi-periodic kernels, longer patches of 15 s were used to better account for inter-breath correlation. All models were trained for 200 epochs using SGD with momentum, the learning rate was set to 3×10^{-5} and the exponential decay was set to 0.85.

To generate training data for the autocorrelated error term in ν_1 , the residual ϵ_1 of the pneumatic respiratory model was calculated by fitting the parameters in

$$P_{\text{aw}}(t_k) + P_{\text{mus}}(t_k) = E \cdot V_{\text{rs}}(t_k) + \epsilon_{1,k}, \quad (5.41)$$

to each patient in the training dataset, where V_{rs} was calculated as described in chapter 4 and P_{mus} is the reference signal derived from P_{es} . The Matérn kernel of the pneumatic error model was intended to only account for short-term autocorrelation, thus, the residual $\epsilon_{1,k}$ of equation (5.41) was high-pass filtered prior to GP training (using a 5th order Butterworth filter with 3 Hz cutoff). Notice that the slow error dynamics are already accounted for by the time-varying offset P_0 , therefore they will not need to be learned here and can be safely filtered out prior to training. The standard deviation of the Matérn kernel was calculated across all residual samples $\epsilon_{1,k}$ in the training data. Again, for this estimation, outliers were removed by discarding 1.5 s of data around points that are more than 3 standard deviations from zero. The length-scale of the Matérn kernel was learned by maximizing the negative log marginal likelihood on mini-batches of 12 s patches, and training was done using the same approach as above, i.e., one patch was taken from each PS level and patient. Again, the optimization was done via SGD with momentum. Empirically, it was observed that the training process was subject to more noise, thus the learning rate was chosen to be slightly lower than before at 2.5×10^{-5} and the number of epochs was set to 400 to ensure convergence.

Training data for the autocorrelated error term in ν_2 were generated in a similar fashion. The EMG was fitted to P_{es} -derived P_{mus} via

$$P_{\text{mus}}(t_k) = K_{\text{EMG}} \cdot \text{EMG}(t_k) + \epsilon_{2,k} \quad (5.42)$$

and the hyperparameters of the Matérn kernel were trained on the residual $\epsilon_{2,k}$. Again, as the focus is on modeling short-term errors, the residuals were high-pass filtered using a 5th order Butterworth filter with a 3 Hz cutoff. The standard deviation was then determined across all residuals $\epsilon_{2,k}$ in the training dataset and outliers were removed as described above. The length-scale was learned using SGD with momentum on mini-batches and was configured the same way as the training of ν_1 .

Parameter priors

Parameter priors were taken from the literature or tuned manually. Priors for R and E were chosen based on reference values for healthy and COPD lungs reported in [14]. The prior lung resistance was parameterized as $\mu_R = 14 \text{ cmH}_2\text{O s/L}$ and $\sigma_R = 4 \text{ cmH}_2\text{O s/L}$, the prior lung elastance was parameterized as $\mu_E = 25 \text{ cmH}_2\text{O/L}$ and $\sigma_E = 4 \text{ cmH}_2\text{O/L}$. The prior for the neuromechanical coupling was set using the values reported in chapter 3. The standard deviation $\sigma_{K_{\text{EMG}}}$ was set to $4 \text{ cmH}_2\text{O}/\mu\text{V}$

and the mean $\mu_{K_{\text{EMG}}}$ was set to $10 \text{ cmH}_2\text{O}/\mu\text{V}$, which introduces a slight bias towards larger neuromechanical couplings, and thus, better accounts for the outliers of the distribution that were observed in chapter 3. As a second variant, the prior information given by airway occlusion maneuver was also incorporated. This was done via a prior for K_{EMG} with $\mu_{K_{\text{EMG}}} = \alpha \hat{K}_{\text{EMG}}^{(\text{occl})}$ and $\sigma_{K_{\text{EMG}}}^2 = 10^{-5} \text{ cmH}_2\text{O}/\mu\text{V}$, refer to section 3.2.3 for further details on the occlusion technique. In the following, all GP models using the occlusion-based prior have the suffix “+occl”. The hyperparameters of the offset process P_0 were tuned to represent the inter-breath dynamics of iPEEP and volume drift. Thus, a Matérn kernel with $\nu = 5/2$, standard deviation of $3 \text{ cmH}_2\text{O}$ and length-scale $\ell = 7.5 \text{ s}$ was used. Finally, the noise standard deviations σ_{n1} and σ_{n2} were tuned manually and set to $2.5 \text{ cmH}_2\text{O}$ and $3 \text{ cmH}_2\text{O}$, respectively. At a first glance, both of these terms appear to be quite large, but keep in mind that they must account for any remaining model errors, such as model mismatch and nonlinearities. Similarly, Petersen [159] also reported that relatively wide priors are needed for the noise variables.

5.3.2 Performance metrics

The performance of the GP framework was evaluated on the 15 test patients using the P_{es} -derived reference for P_{mus} . As opposed to chapter 4, where estimators were evaluated in a *breath-wise* fashion, in this chapter a greater focus is put on the *sample-wise* correspondence between estimated and measured P_{mus} (because this is where the influence of temporal GP priors will be most noticeable). The following metrics were used to quantify the relation between the posterior $P_{\text{mus,post}}$ and the reference P_{mus} :

Correlation. The sample-wise correlation between the posterior mean $m_{P_{\text{mus,post}}}$ and the reference P_{mus} was quantified via Pearson’s correlation index r .

Deviation. The deviation between $m_{P_{\text{mus,post}}}$ and P_{mus} was quantified via the sample-wise mean absolute deviation (MAD).

Log likelihood. Finally, as the previous two metrics only incorporate the posterior mean, the uncertainty calibration of the posterior was evaluated by means of the *empirical cross entropy*, which is a widely used loss function in the machine learning literature. The metric was calculated as the average *negative log likelihood*

$$-\frac{1}{N} \sum_{k=1}^N \log \mathcal{N}(P_{\text{mus}}(t_k); m_{P_{\text{mus,post}}}(t_k), k_{P_{\text{mus,post}}}(t_k, t_k)) \quad (5.43)$$

across all included samples t_k with $k = 1, \dots, N$.

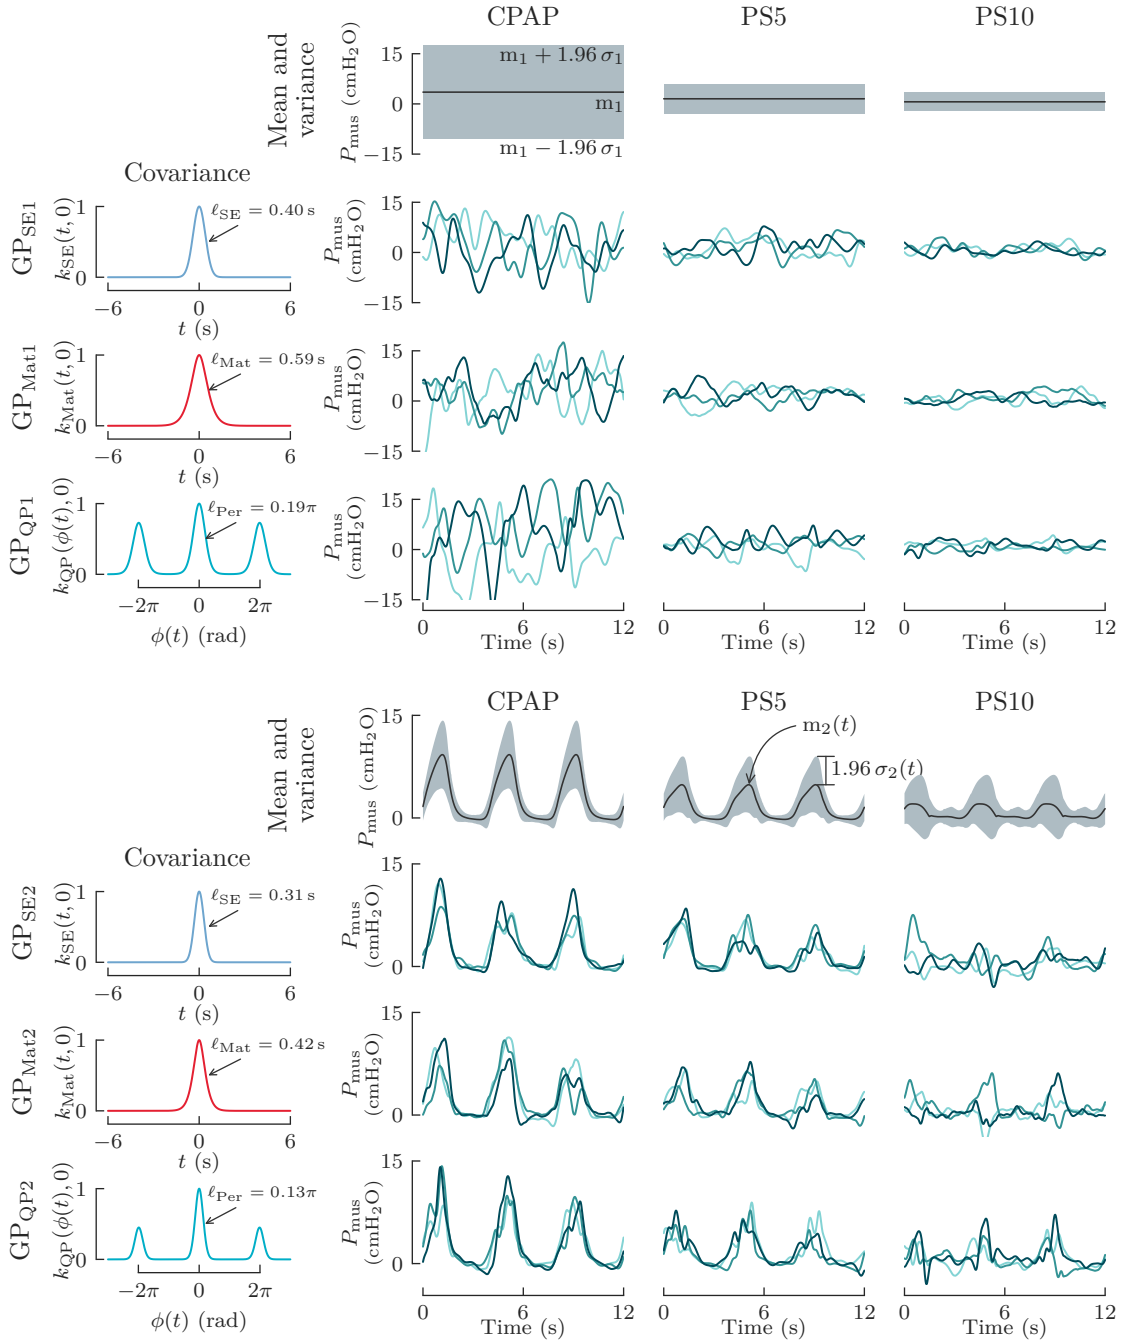


Fig. 5.4: An overview of P_{mus} GP priors after optimization of hyperparameters. Columns correspond to different pressure support (PS) levels and rows correspond to different kernels. The upper panel provides results for stationary means/variances, the bottom panel for non-stationary means/variances. Kernels are evaluated and plotted by their first input as $k(t, 0)$. Means and 95% confidence intervals for each PS level are plotted over three breathing cycles, assuming $T_{\text{insp}} = 1.5$ s and $T_{\text{exp}} = 2.5$ s. The inner part of both panels contains random functions sampled from the GPs.

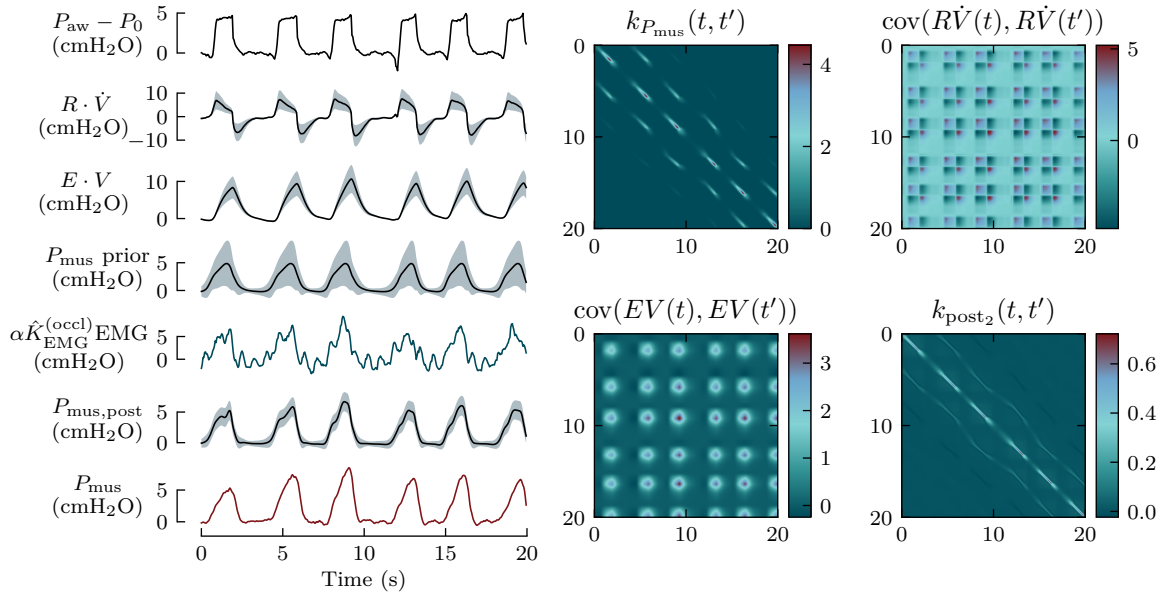


Fig. 5.5: An excerpt of patient data at pressure support of 5 cmH₂O. The left panel shows prior means and 95 % confidence intervals of the pneumatic data ($E \cdot V$ and $R \cdot \dot{V}$), of the quasi-periodic P_{mus} prior, and of the electrical data ($\alpha \hat{K}_{\text{EMG}}^{(\text{occl})}$ EMG). The EMG prior has very small confidence bounds due to the high certainty that was put on the occlusion-derived neuromechanical coupling index ($\hat{K}_{\text{EMG}}^{(\text{occl})}$). The P_{mus} posterior from equation (5.28) ($P_{\text{mus,post}}$) closely follows the esophageal pressure-derived signal (P_{mus}). The right panels depict prior and posterior covariance matrices ($k_{P_{\text{mus}}}(t, t')$ is the prior kernel from (5.1), $k_{\text{post}_2}(t, t')$ is the posterior kernel from (5.28)).

5.4 Results

The mean and standard deviation of P_{mus} GP priors across the 15 training patients are plotted in figure 5.4 for three support levels (CPAP, 5 cmH₂O and 10 cmH₂O). As expected, the amplitude and variance of the P_{mus} distribution decreased with increasing ventilatory support. The phase-dependent effort model (defined by m_2 and σ_2) showed a waveform that nicely reflects typical breathing behavior: on average, the activity during inspirations was higher than during expirations and also had a larger variance. The inspiration was initiated by a positive deflection in P_{mus} followed by a maximum in the middle of the inspiration and a decay behavior initiating the expiration. These observations coincide with what has been reported by Mecklenburgh and Mapleson [136]. The average shape of efforts was clearly moderated by the PS level, which indicates the necessity to learn a new shape in each level (this also means that during training P_{es} should be measured across multiple PS levels). Convergence of hyperparameter training was observed within the maximum number of epochs for

all GP models. Length-scale parameters of P_{mus} kernels are provided in figure 5.4 and the resulting GP dynamics are visualized by means of P_{mus} sample functions. The learned length-scales of stationary GPs (defined by m_1 and σ_1) were generally larger than those of non-stationary GPs (defined by m_2 and σ_2), because for the latter a large part of the slow dynamics are already covered by the mean function. The resulting sample functions exhibit a wide range of waveforms and thus nicely capture the variety of effort shapes that one would expect from mechanically ventilated patients. Very importantly, the samples are not restricted to a fixed parametric function as in the previous P_{mus} models introduced in section 5.1. The two quasi-periodic kernels in figure 5.4 reveal that there is a strong coupling between successive breaths, which is more pronounced in the model with constant mean (m_1). Again, this is to be expected, because the non-constant mean (m_2) already accounts for (strictly) periodic dynamics. The length-scale of the quasi-periodic decay was identified as $\ell_{\text{SE}} = 2.51\pi$ for the stationary and as $\ell_{\text{SE}} = 1.58\pi$ for the non-stationary model, which translates into a breath-to-breath correlation (between points with the same respiratory phase) of 0.73 and 0.45, respectively. The hyperparameters of the pneumatic noise model were identified as $\ell_{\text{Mat}} = 0.11$ s and $\sigma_{\text{Mat}} = 0.32$ cmH₂O. The hyperparameters of the EMG noise model were identified as $\ell_{\text{Mat}} = 0.13$ s and $\sigma_{\text{Mat}} = 0.57$ cmH₂O. These results provide important insights into the amplitude and temporal dynamic of the autocorrelated errors terms in equations (5.41) and (5.42).

The covariance of the different GPs are visualized in figure 5.5 and a comparison of all P_{mus} estimates on an illustrative excerpt of data is given in figure 5.6. Performance metrics across all 15 test datasets are provided in figure 5.7. The best performing estimators were the ones combining the non-stationary GP with the neuromechanical coupling prior derived from occlusions ($\text{GP}_{\text{SE2+occl}}$, $\text{GP}_{\text{Mat2+occl}}$ and $\text{GP}_{\text{QP2+occl}}$). For the Matérn kernel, an average correlation to P_{mus} of 0.93 ± 0.04 and an average MAD of 0.60 ± 0.25 was reached, and similar performances were also attained by the squared exponential and quasi-periodic kernel, refer to figure 5.7. The two baseline models (OLS and GP_{uninf}) clearly failed to obtain viable estimates of P_{mus} : both models did not succeed in producing a physiological waveform, which was reflected in a low correlation, and they also failed to correctly estimate the amplitude of P_{mus} , which resulted in a large MAD. Thus, these results demonstrate the benefits of using the proposed GP framework and of incorporating prior knowledge about P_{mus} . Indeed, all investigated GP priors provided some benefit to the estimation. The simplest GP models (GP_{SE1} , GP_{Mat1} , GP_{QP1}), using constant mean m_1 and standard deviation σ_1 , attained a higher correlation to P_{mus} than the uninformative model, which can be mainly attributed to the fact that they smooth out artifacts in the posterior P_{mus} curve. The non-stationary models (GP_{SE2} , GP_{Mat2} , GP_{QP2}), using the non-constant mean m_2 and standard deviation σ_2 , performed very well in terms of both correlation and MAD. The analysis did not show a significant difference between the performances of the three covariance functions, i.e., the Matérn, squared exponential and quasi-

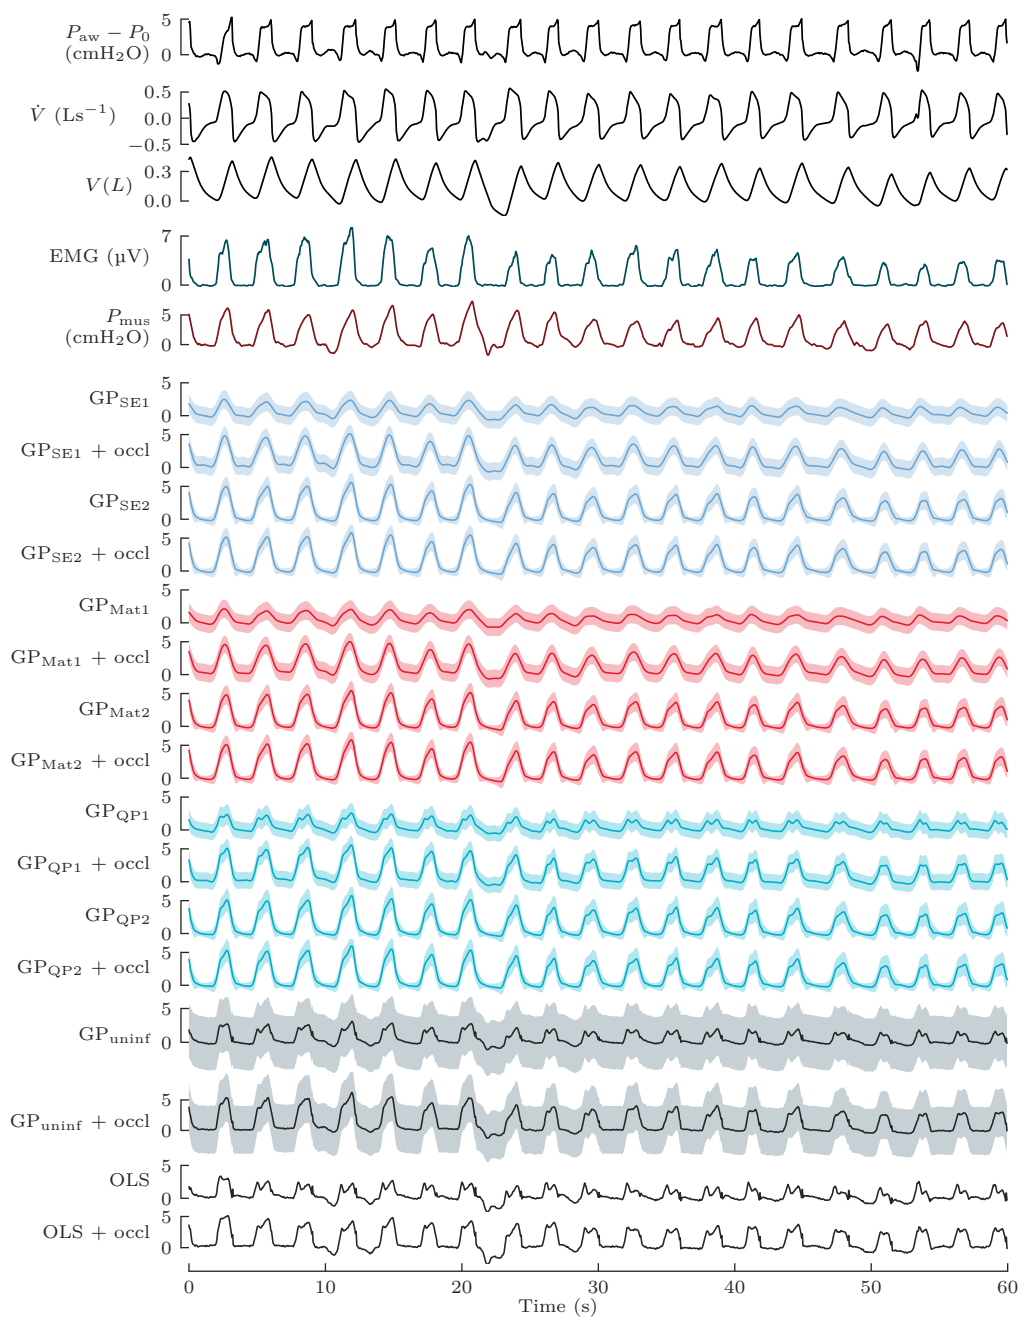


Fig. 5.6: A 60s excerpt of patient data and GP posteriors at pressure support of 5 cmH₂O. The P_{mus} reference signal was derived from esophageal pressure. The mean and 95% confidence intervals of GP posteriors were calculated using the state-space formulation from section 5.2.4. The unit for all P_{mus} estimates is cmH₂O.

periodic kernels achieved very similar results. Across all models (baseline and GP models), the incorporation of prior knowledge derived from occlusions was found to be highly beneficial and consistently improved the performance. The results in figure 5.7 also give insights into the capability of different models to estimate the posterior uncertainty, which was quantified by means of the negative log likelihood. Interestingly, the uninformative GP model had a relatively small negative log likelihood, which it achieved by correctly predicting very wide confidence bounds, refer to figure 5.6 for an illustrative example. The most prominent uncertainty miscalibration was observed in the three stationary GP models (GP_{SE1} , GP_{Mat1} , GP_{QP1}), which all tend to be overly confident, indicated by high negative log likelihood values.

5.5 Discussion

This chapter has introduced a novel GP framework for modeling the respiratory muscle pressure P_{mus} in mechanically ventilated patients. GPs provide a highly flexible class of models that bridge the gap between existing approaches to P_{mus} estimation, including regularization techniques and Kalman filtering. Several key characteristics of P_{mus} signals such as smoothness, periodicity and the shape of efforts were implemented via different kernel functions, and incorporated into a Bayesian sensor fusion model. The main result of this chapter is the following: The incorporation of prior knowledge in the form of a GP offers substantial benefits over simpler (uninformative) models for estimating P_{mus} .

Important contributions towards a physiological, stochastic P_{mus} model have been made in this chapter. Previous models were rather simplistic, relying on fixed parametric functions [104, 182] or on simple state-space models such as harmonic oscillators [159, 196]. The herein proposed GP framework provides a comprehensive family of stochastic models that cover many relevant effects, including short-time dynamics, periodicity, and phase-dependent waveforms. As opposed to most previous studies, all model parameters were identified directly on esophageal pressure recordings in mechanically ventilated patients. The resulting GPs provide, for the first time, insight into the dynamics of P_{mus} signals and their covariance both within and across breaths. This adds to the findings of previous authors, such as Bellemare et al. [30] who have studied the temporal characteristics of isolated diaphragm twitches and Petersen [159] who identified simple state-space models to represent P_{mus} signals. Moreover, to the best of the author's knowledge, no previous study has investigated the average P_{es} -derived muscle pressure waveform across a similarly large cohort of patients. The work closest to this one would be that by Mecklenburgh and Mapleson [136], who determined an average inspiratory effort waveform. However, they did not directly measure P_{mus} from esophageal pressure, but used a calculated P_{mus} signal. Another contribution of this chapter is the investigation of the dynamics of measurement errors in the pneumatic

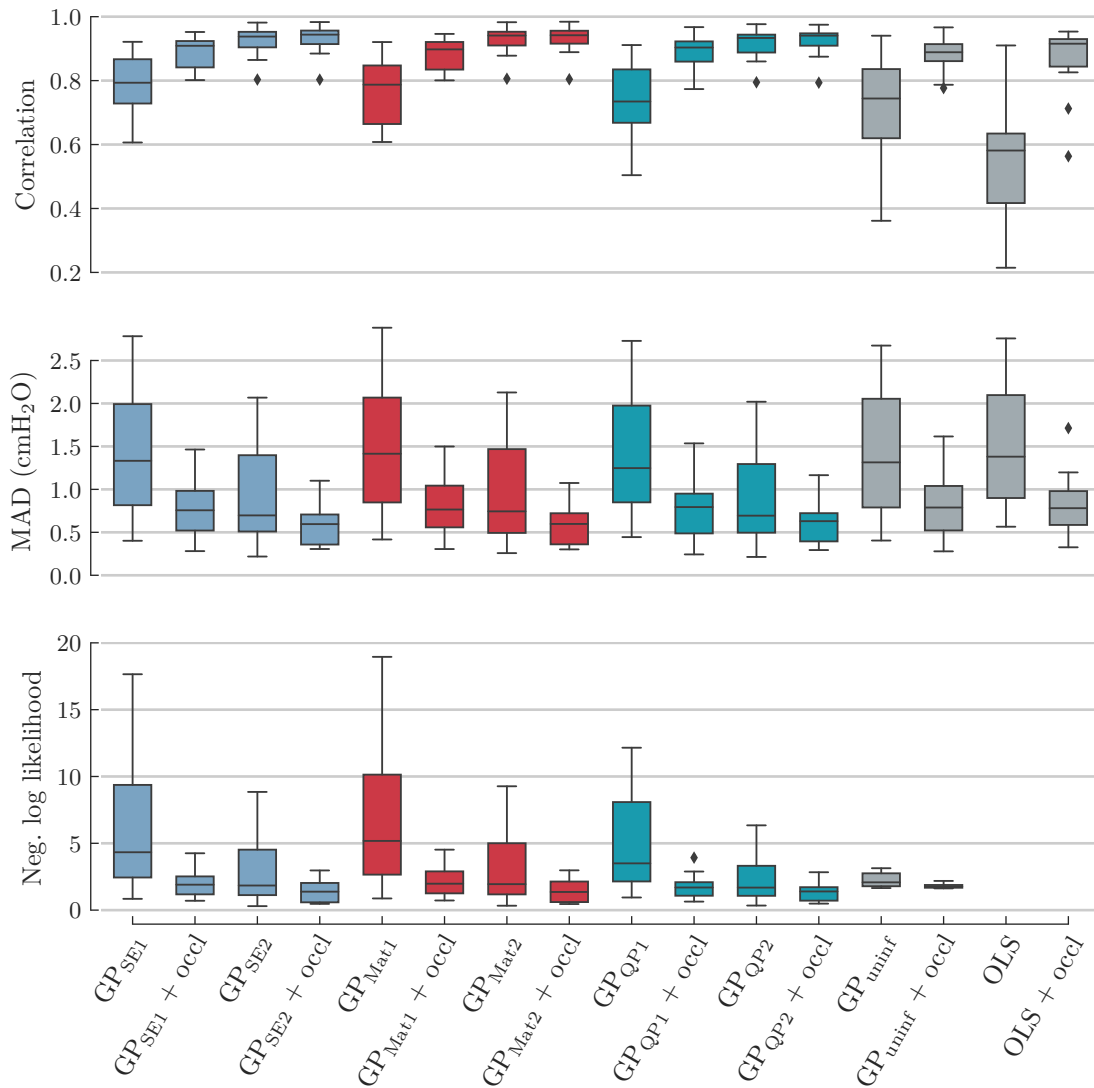


Fig. 5.7: Performance metrics on the test data ($n = 15$) for all considered GP models. The figure provides box plots of Pearson correlation, mean absolute deviation (MAD) and average negative log likelihood, all calculated in a sample-wise fashion.

and the EMG model. For both equations, the training data had considerable autocorrelation on a scale of roughly 0.1 s. One important takeaway is that sampling times significantly below 0.1 s will not provide much benefit for solving the problem posed in section 5.2.3, because the samples would become strongly autocorrelated. In other words, a sampling time of roughly 0.1 s would be sufficient to solve the inference. In

this evaluation, a slightly smaller sampling time of 0.05 s has been used to guarantee that all relevant information are captured.

One key goal of this chapter was to investigate what kind of prior knowledge needs to be incorporated to estimate P_{mus} non-invasively from ventilator and EMG data. The results show that the proposed GP models clearly outperform uninformative priors. Moreover, all informative priors (using squared exponential, Matérn, or quasi-periodic covariance functions) were found to have very similar performance profiles. Most surprisingly, the periodic prior did not outperform the smooth GP models, despite periodicity being one of the most distinctive features of respiratory signals. A strong correlation of subsequent points with the same respiratory phase was indeed learned by the quasi-periodic GP model on the training data, but it appears that the incorporation of this property is not needed to solve the inference problem. One possible explanation is that most of the measured ventilatory data are themselves already periodic, and thus implicitly regularize the result toward periodic solutions. Still, periodic P_{mus} models could be beneficial in other domains, for example, when the data are very noisy or in the context of generative models. On the problem at hand, the Matérn kernel provides the best trade-off between model complexity (indicated by its state-space dimension) and estimation performance. Another key result of the analysis is that the phase-dependent models for the mean and standard deviation (m_2 and σ_2) are very effective in improving performance, and would allow to estimate P_{mus} on a single PS level without any disruptive maneuvers. The other decisive prior proved to be the neuromechanical coupling derived from occlusion maneuvers, which improved performance across all models and also provided viable estimates in combination with the uninformative GP.

The GP framework presented in this chapter complements the techniques that were investigated in chapter 4. While the focus in chapter 4 was on robustness, this chapter has made the case for incorporating prior knowledge about P_{mus} dynamics. While chapter 4 was focused on breath-wise estimation, in this chapter a sample-wise evaluation was pursued (because this is where GP models have the greatest influence). The results show that both directions improve P_{mus} estimation — due to the different types of metrics and different test datasets, a direct comparison of the two approaches has not yet been conducted, but is planned in the future. Both chapters, ultimately, have addressed the shortcomings of simpler regression techniques such as OLS, which tend to generate unphysiological P_{mus} curves due to their strong sensitivity to artifacts, phase delays and model mismatch. From a Bayesian viewpoint, chapter 4 has focused on the choice of the likelihood function, in particular regarding fat-tailed distributions and the weight of samples, and chapter 5 has focused on the choice of the prior. Thus, the two perspectives nicely complement each other and could also be combined to derive novel estimators, for example by using GP priors together with non-Gaussian likelihood functions. Such models would generally give rise to intractable inference problems, but a wide range of techniques described in the literature offer approximate

solutions, e.g., by means of expectation propagation or variational inference [149]. An important advantage of the herein proposed GP-based framework is that it connects many of the previously proposed methods for estimating P_{mus} . GP inference can be posed as a regularized optimization problem and thus provides direct links to the P_{mus} estimation techniques studied in [176, 228]. It was also shown that equivalent state-space models can be derived, allowing to solve the inference via Kalman filtering and smoothing techniques, which in turn provides links to the publications by Saatçi and Akan [182], Petersen [159] and Scheel et al. [196]. For this reason, the GP framework seems to be an extremely valuable tool for unifying the various algorithms and for studying their structure as well as their underlying assumptions.

Several extensions of the methodology seem worth to be pursued further. Firstly, a systematic investigation of more complex respiratory models beyond the linear equation of motion would be valuable. For instance, several authors have proposed nonlinear differential equations for the respiratory mechanics [16], which could also be easily combined with the GP framework. Again, closed-form solutions are intractable but approximations in the form of extended or unscented Kalman filters/smoothers are available. An increase in model complexity offers the potential to reduce model mismatch, but at the same time complicates parameter identification. In this context, the use of GP priors could be very helpful to regularize the parameters towards physiological solutions. Another interesting direction is the extension of the framework to include more measurements such as esophageal pressure, gastric pressure or EAdi. In the spirit of the sensor fusion paradigm, combining multiple independent data could further improve the estimation. The proposed GP framework can be flexibly adapted to novel applications and data by adding measurement equations and by assessing what kind of prior information needs to be included.

Chapter 6

Scalable Gaussian Process Regression for Spatio-Temporal Biomedical Data

Gaussian processes are a versatile tool to model the spatio-temporal dynamics of biomedical data. However, the application of GPs to large datasets remains challenging due to heavy memory and computational requirements that come with the classical GP regression formulation. In this chapter, these practical limitations will be addressed by investigating techniques for *scalable GP regression*. In the previous chapter, one possible strategy has already been introduced, namely Kalman filtering with equivalent state-space models—this approach, however, may be prohibitive on *spatio-temporal* datasets when the state-space dimension becomes too large. This chapter discusses another orthogonal approach to scalability which is based on exploiting existing structure in the kernel matrix. In particular, this chapter investigates the separation of physiological processes with a non-stationary phase, which is a common problem in biomedical signal processing (e.g., in respiratory or cardiac signals). The herein proposed method employs sets of non-equidistant inducing points to retrieve *Kronecker and Toeplitz* structure in the kernel matrix, allowing for efficient matrix operations. The approach is demonstrated on two biomedical applications: (1) separation of maternal and fetal ECG signals and (2) reconstruction of periodic processes in electrical impedance tomography (EIT).

The chapter begins with an introduction to structure exploiting methods for scalable GP regression (section 6.1), focusing in particular on state-space and Toeplitz/Kronecker methods. The following section 6.2 discusses the proposed scalable technique for GP-based separation of biomedical data. Sections 6.3 and 6.4 then proceed to present first results in two biomedical case studies, and section 6.5 concludes the chapter with a discussion. Parts of this chapter were published in a conference paper that was presented at the ICML in 2021 [71] with a preprint available on arXiv [70] and in an abstract submitted to the AUTOMED in 2020 [79]. Both of these works were first-authored by the author of the thesis. He developed the methodology, conducted the numerical experiments and wrote most of the manuscript of both papers. The case study conducted in section 6.4 has not yet been published elsewhere.

6.1 Introduction and previous work

As discussed in section 2.1.3, the main bottleneck for GP inference and learning on large datasets comes from solving $(K + \sigma_n^2 I)^{-1}y$, and computing the log determinant $\log |K + \sigma_n^2 I|$, which is usually done via Cholesky decomposition of the kernel matrix $K + \sigma_n^2 I$. This requires $\mathcal{O}(N^3)$ operations and $\mathcal{O}(N^2)$ memory, where N is the number of data points, which proves to be intractable for all but the smallest datasets. Typical spatio-temporal datasets — with high temporal or spatial sampling rates — easily exceed the tolerable size of data points, which renders standard GPs inapplicable to many relevant biomedical domains, including for instance electrical impedance tomography. This section discusses scalable GP inference and learning, focusing on techniques that exploit structure in the kernel matrix. It will touch upon (1) state-space representation methods, and (2) methods exploiting Toeplitz and Kronecker matrix structure.

6.1.1 Equivalent state-space representation

The state-space approach described by Särkkä et al. [85, 191, 212] is a first instance of structure exploitation in the sense that it uses a factorization of the GP’s spectral density when it has rational function form, refer to section 2.1.4 for details. The conversion of temporal kernels into state-space form enables highly scalable inference and marginal likelihood evaluation by means of Kalman filtering and smoothing, which have complexity $\mathcal{O}(TM^3)$, where T is the number of time-steps and M the number of states. Equivalent state-space models can also be derived for spatio-temporal kernels, leading to stochastic partial differential equations, which in turn can be solved using infinite-dimensional Kalman filter and smoother techniques proposed by Särkkä et al. [192]. The state-space approach allows for easy integration of physical models and non-stationarities as was demonstrated in the previous chapter 5 and in application papers [86, 193].

Despite in principle having linear complexity, in practice, spatio-temporal GPs often require high-dimensional state-space models and thus there is a large cost involved at every time step, e.g., when the spatial domain is large or when the temporal spectral density has a high-order polynomial factorization. In regard of the challenges involved in large-scale Kalman filtering/smoothing problems, in this chapter, an alternative direction will be explored, namely scalable operations on the batch kernel matrix.

6.1.2 Toeplitz and Kronecker methods

In recent years, the exploitation of structure in the kernel matrix has received much attention and has spawned a wide range of scalable GP inference/learning frameworks [63, 127, 237, 240]. In the following, a brief introduction to the most prevalent techniques will be provided.

When a *stationary* kernel ($k(t, t') = k(t + \tau, t' + \tau)$) is evaluated on a set of equispaced points t_1, \dots, t_N , the corresponding kernel matrix $k(t_{1:N}, t_{1:N})$ has Toeplitz structure, that is, it has constant diagonals: $k(t_i, t_j) = k(t_{i+1}, t_{j+1})$. As described in detail by Cunningham et al. [43], an $N \times N$ Toeplitz matrix can be embedded into a larger $(2N - 1) \times (2N - 1)$ circulant matrix, which enables fast *matrix-vector multiplications* (MVMs) via convolution of a zero padded argument vector with one row of the larger circulant matrix. Thus, MVMs can be calculated in $\mathcal{O}(N \log N)$ using fast Fourier transforms, and one only ever needs to store the first row of the circulant matrix, reducing memory requirements to $\mathcal{O}(N)$. The expression $(K + \sigma_n^2 I)^{-1} y$ can then be solved efficiently in an iterative fashion using *conjugate gradients* (CG), which only requires fast MVMs with the full kernel matrix $K + \sigma_n^2 I$ and usually converges to machine precision within a small number of steps $J \ll N$. The naive batch calculation for the log determinant of Toeplitz matrices is also not efficient, but some scalable approximations have been proposed, refer to [51].

An orthogonal approach is to exploit Kronecker structure, which arises when input points are on a multi-dimensional rectilinear lattice and the kernel is *product separable* along its input dimensions $d = 1, \dots, D$, that is, it admits the form

$$k(x, z) = \prod_d k_d(x^{(d)}, z^{(d)}). \quad (6.1)$$

Here, contrary to FFT-based methods, inputs need not necessarily be equispaced. The corresponding kernel matrix can be written as a Kronecker product $K = K_1 \otimes \dots \otimes K_D = \bigotimes_{d=1}^D K_d$ in terms of smaller dimension-wise matrices K_d , which substantially reduces memory usage to $\mathcal{O}(DN^{\frac{2}{D}})$ and enables a number of efficient matrix operations that can be used for scalable GP inference and learning. Saatçi [183] provides a comprehensive list of relevant properties involving the Kronecker product, some of which will be repeated in the following. In particular, the *eigendecomposition* $K = QVQ^\top$ of the full kernel matrix can be calculated by means of the Kronecker products $Q = \bigotimes_{d=1}^D Q_d$ and $V = \bigotimes_{d=1}^D V_d$, where Q_d and V_d are given by eigendecompositions $K_d = Q_d V_d Q_d^\top$ of the smaller matrices. The expression $(K + \sigma_n^2 I)^{-1} y$ can then be rewritten using

$$(K + \sigma_n^2 I)^{-1} y = Q(V + \sigma_n^2 I)^{-1} Q^\top y, \quad (6.2)$$

where inversion is trivial, because $V + \sigma_n^2 I$ is a diagonal matrix [183]. Kronecker structure allows for MVMs in $\mathcal{O}(DN^{\frac{D+1}{D}})$ time [240], making the evaluation of (6.2) highly efficient. The eigendecomposition can also be used to efficiently calculate the log determinant in closed form [240] using the identity

$$\log |K + \sigma_n^2 I| = \sum_k \log(V^{(k,k)} + \sigma_n^2). \quad (6.3)$$

Toeplitz and Kronecker structures are complementary in the sense that a stationary, separable kernel evaluated over an equispaced lattice gives rise to a Kronecker product

of Toeplitz matrices, which can be exploited for highly efficient MVMs. A combination of both techniques makes most sense when one of the dimensions becomes prohibitive, i.e., when the eigendecomposition of one of the Kronecker factors is too expensive. One can then resort to solving the system of linear equations via CG by exploiting fast MVMs. No simple closed-form solution is known for the log determinant of a Toeplitz matrix, but iterative methods can be considered for computing the eigenvalues [181].

Inducing points and structured kernel interpolation

Both Toeplitz and Kronecker methods are restricted to a few specialized problems due to the requirements on the kernel, which must be stationary or separable, and on the input points, which must be equispaced or form a rectilinear lattice. The restrictions on the placement of input points were lifted through the *structured kernel interpolation* (SKI) framework by Wilson et al. [237, 240]. It employs a dedicated set of input points u_1, \dots, u_M — the so called *inducing points* — to recover structure.

Inducing point methods have been used for a long time to deal with the computational requirements of GP regression, refer to [167] for comprehensive overview. Classic inducing point approaches alleviate the computational and memory burden by using a set of M inducing points and approximating the kernel by means of the $M \times M$ -dimensional matrix $K_{UU} := k(u_{1:M}, u_{1:M})$. For instance, the subset of regressors (SoR) method uses the approximation $k(x, x') \approx k_{\text{SoR}}(x, x') = k(x, u_{1:M}) K_{UU}^{-1} k(u_{1:M}, x')$, with a kernel matrix of rank (at most) M . It reduces computational cost to $\mathcal{O}(M^2N + M^3)$ and memory cost to $\mathcal{O}(MN + M^2)$ [167]. Classic inducing point methods do not exploit any structure within the kernel matrix and rely simply on using a small number $M \ll N$ of inducing points. These methods work best when data are densely sampled compared to the characteristic length-scale of the kernel, however, they are less suited to spatio-temporal measurements because the “extending domain” has to be filled with a sufficient amount of inducing points to guarantee a good approximation of the kernel [211].

Building upon the classical inducing point methods, the SKI framework by Wilson et al. [237, 240] aimed to improve scalability with respect to the number of inducing points by exploiting structure in the matrix K_{UU} . To this end, they proposed an approximation to the cross-covariance matrix $K_{XU} := k(x_{1:N}, u_{1:M})$ of the form

$$K_{XU} \approx \tilde{K}_{XU} = WK_{UU}, \tag{6.4}$$

where W is an $N \times M$ interpolation matrix that maps from the inducing points $u_{1:M}$ to the original inputs $x_{1:N}$. The matrix W can be constructed to be very sparse. For example, when the inducing points $u_{1:M}$ are on an equispaced lattice, one can use cubic convolution interpolation [109], which means that W only has 4 non-zero entries per

row, requiring only $\mathcal{O}(N)$ storage. Using the SoR scheme and substituting \tilde{K}_{XU} , the full kernel matrix $K_{XX} := k(x_{1:N}, x_{1:N})$ can be approximated via

$$K_{XX} \stackrel{\text{SoR}}{\approx} K_{XU} K_{UU}^{-1} K_{XU}^\top \stackrel{\text{eq. (6.4)}}{\approx} W K_{UU} W^\top =: K_{\text{SKI}}. \quad (6.5)$$

The SKI approximation $W K_{UU} W^\top$ enables highly efficient MVMs. Multiplication with the sparse interpolation matrix W can be done in $\mathcal{O}(N)$, and, by placing inducing points on an equispaced lattice, multiplication with K_{UU} allows for exploitation of Kronecker/Toeplitz structure, leading to a worst-case runtime of $\mathcal{O}(M \log M)$ for Toeplitz structure [238]. The total runtime for MVMs is thus $\mathcal{O}(N + M \log M)$. Exploitation of both Toeplitz and Kronecker structure leads to an MVM runtime of $\mathcal{O}(N + g(M))$ with $g(M) < M \log M$, and in many cases quasi-linear complexity $g(M) \approx M$ can be reached [238]. Contrary to previous inducing point methods, SKI places almost no restrictions on the number M of inducing points, and even $M > N$ is feasible. Using a large number of inducing points empowers SKI to reach very close approximations to the original kernel matrix. Since its introduction, the viability of SKI has been demonstrated in a wide range of large-scale regression and classification problems, refer to [239].

6.2 Scalable separation of GPs with a non-stationary phase

This chapter addresses the application of GPs to large biomedical datasets with temporal/spatio-temporal structure. Particular attention is given to solving the separation problem for *additive mixtures* of physiological processes that have a *non-stationary phase* $\phi(t)$. Such problems often arise in biomedical data, for instance, many patient datasets contain a superposition of cardiogenic and respiratory activity, which are inherently non-stationary due to natural fluctuations in the heart or respiratory rate, and have to be separated prior to further processing, see for instance [19, 162, 164]. Thus, in the following section, structure exploitation for non-stationary kernels will be investigated. More specifically, an extension of SKI to mixtures of kernels with a non-stationary phase is discussed, which enables to solve highly relevant biomedical source separation and regression problems.

6.2.1 Model definition

Generally, this chapter is concerned with additive mixtures of the form

$$\begin{aligned} f_{\text{m}}(x) &= \sum_i f_{\phi,i}(x) \\ y &= \mathcal{H} f_{\text{m}}(x) + \epsilon, \end{aligned} \quad (6.6)$$

where $x \in \mathbb{R}^D$ is the input variable, and measurements $y \in \mathbb{R}^N$ are obtained via the deterministic linear functional \mathcal{H} with additive Gaussian white noise $\epsilon \sim \mathcal{N}(0, \sigma_{\text{n}}^2 I)$.

Linear functionals of GPs have been discussed in detail by Särkkä et al. [186] and may for example represent derivatives/integrals of the GP, convolution kernels corresponding to the effect of a linear system, or partial differential equations (PDE) modeling spatial physical processes. In the simplest case, the functional \mathcal{H} corresponds to evaluating a fixed set of points within the input domain, which will be used in section 6.3. In this chapter, the components of the additive mixture were modeled via $f_{\phi_i}(x) \sim \mathcal{GP}(0, k_i(\phi_i(x), \phi_i(x')))$, where k_i are dimension-wise product separable and stationary kernels, and $\phi_i : \mathcal{D}_{\text{in}} \rightarrow \mathcal{D}_i$ are nonlinear transformations of the inputs with $\mathcal{D}_{\text{in}} \subseteq \mathbb{R}^D$, $\mathcal{D}_i \subseteq \mathbb{R}^D$, which are used to account for the non-stationarity of physiological data. The objective is to separate the non-stationary GPs, that is, to estimate the sources $f_{\phi_i}(x)$. The full kernel can be stated as

$$k_{\text{m}}(x, x') = \sum_i k_{\phi_i}(x, x') = \sum_i k_i(\phi_i(x), \phi_i(x')). \quad (6.7)$$

Notice that the kernels $k_{\phi_i}(x, x')$ are in general not product separable with respect to the input dimensions due to the nonlinear warping of the inputs through ϕ_i .

Contrary to the prominent deep kernel learning (DKL) framework [239] (which also employs nonlinear transformations of the input variable), in this chapter the assumption was made that the functions ϕ_i are ‘well-behaved’, which here means that they are invertible and do not have singularities in the input domain \mathcal{D}_{in} . The kernel property would be preserved even if ϕ_i are not invertible, but focusing on invertible functions will prove to be useful later on and does not pose strong limitations with regard to solving biomedical problems.

Source separation

The functional \mathcal{H} can be thought of as evaluating the result of a linear operator at certain points, e.g., evaluating the solution to PDEs or derivatives/integrals. The solution to $\mathcal{H}f_{\text{m}}(x)$ can in practice be computed using appropriate numerical methods which will not be discussed in detail, but the interested reader is referred to [189] for an overview. A few methods are worth to be mentioned: for instance, solutions to PDEs may be obtained by means of point collocation methods, finite difference approaches or basis function expansion methods (such as FEMs), for which a practical example will be given in section 6.4. For the remainder of the chapter, it will be assumed that a numerical solution to $\mathcal{H}f_{\text{m}}(x)$ is available and can be represented (or at least approximated) by evaluating $f_{\text{m}}(x)$ at a finite set of input points x_1, \dots, x_L , which means that the result can be stated as $Hf_{\text{m}}(x_{1:L})$, where $f_{\text{m}}(x_{1:L})$ is a vector of function values and $H \in \mathbb{R}^{N \times L}$ is a matrix. The inference problem to be solved is the following: given the model stated in equation (6.6) and a vector of measurements y , calculate the mean of the posterior $f_{\phi_j}(x)|y$ corresponding to the j th term in the

mixture. All functionals are linear and all priors have Gaussian distributions, thus, the posterior mean can be calculated in closed-form via

$$m_{\text{post},j}(x) = k_{\phi,i}(x, x_{1:L}) H^\top \left(H K_{m,XX} H^\top + \sigma_n^2 I \right)^{-1} y \quad (6.8)$$

with $K_{m,XX} = k_m(x_{1:L}, x_{1:L})$, which is a generalization of the standard GP regression equation. As in standard GP regression, $k_{\phi,i}(x, x_{1:L}) H^\top$ is a $1 \times N$ vector, and $H K_{m,XX} H^\top$ is an $N \times N$ covariance matrix. The main challenge of using equation (6.8) comes from $\mathcal{O}(N^3)$ computational cost required for solving the system of linear equations. Classic GP regression would be recovered by defining the functional as $\mathcal{H}f(x) = [f(x_1), \dots, f(x_N)]^\top$, leading to

$$m_{\text{post},j}(x) = k_{\phi,i}(x, x_{1:N}) \left(K_{m,XX} + \sigma_n^2 I \right)^{-1} y \quad (6.9)$$

with $K_{m,XX} = k_m(x_{1:N}, x_{1:N})$.

6.2.2 Structure exploitation

The goal here is to alleviate the computational and memory burden of the source separation problem in equations (6.8) and (6.9) by exploiting existing structure within the kernel matrix. The numerical computation relies on solving $(H K_{m,XX} H^\top + \sigma_n^2 I)^{-1} y$ and $(K_{m,XX} + \sigma_n^2 I)^{-1} y$, respectively. Assuming that multiplication with H is efficient, scalable matrix operations for the mixture kernel matrix $K_{m,XX}$ will be explored below. No restrictions are to be imposed on the placement of input points x_1, \dots, x_L , that is, no special structure (grid or equispaced) should be required. At first, the exploitation of structure is not straightforward because classical requirements to the kernel $k_m(x, x')$ such as stationarity and product separability are not met. In particular, the warping functions ϕ_i lead to the loss of Toeplitz structure even for equispaced inputs and the summation of multiple kernels prevents Kronecker structure of the full kernel matrix even when each of the summands can be written as a Kronecker product. In the following, a variant of structured kernel interpolation is proposed that employs multiple sets of non-equidistant inducing points to recover structure and enable scalable solutions to the inference problem.

Structured kernel interpolation for $K_{m,XX}$

This section begins by focusing on one of the terms $k_\phi(x, x') = k(\phi(x), \phi(x'))$ in the mixture. The proposed approach is based on the observation, first made in the context of DKL [239], that SKI can be applied directly to the transformed inputs $z = \phi(x)$ when k is stationary and product separable. This gives rise to the approximation

$$K_{\phi,XX} \approx W_Z K_{UU} W_Z^\top, \quad (6.10)$$

where $K_{\phi,XX} = k_{\phi}(x_{1:L}, x_{1:L})$, and W_Z maps from the inducing points U to the transformed points $Z = \Phi(X)$ with Φ denoting the operation that transforms points according to $z_j = \phi(x_j)|_{j=1:L}$. Thus, Kronecker and Toeplitz structure in K_{UU} can be imposed through placement of the inducing points and there are no restrictions to the structure of inputs. It is also worth noting that in equation (6.10), the non-stationarity of the model is fully embedded in the sparse interpolation matrix W_Z .

In the following, a reformulation of the above idea is proposed: instead of applying SKI to the nonlinearly transformed inputs Z , a set of inducing points \hat{U} will be used to directly recover structure when evaluating the non-stationary kernel $k(\phi(x), \phi(x'))$. To this end, inducing points are constructed via $\hat{U} = \Phi^{-1}(U)$ with U having an equispaced, rectilinear grid structure. This gives rise to the approximation

$$K_{\phi,XX} \approx W_X K_{\phi,\hat{U}\hat{U}} W_X^\top = W_X K_{UU} W_X^\top =: K_{\phi\text{SKI}}, \quad (6.11)$$

where the interpolation matrix W_X maps from \hat{U} to the input points X , which uses the fact that $K_{\phi,\hat{U}\hat{U}} = K_{UU}$. The main difference to standard SKI is that this reformulation uses an inducing point set \hat{U} with a warped, non-equidistant structure. The reformulations of SKI in equations (6.10) and (6.11) are equivalent (when ϕ is invertible) and one can apply both depending on which of the two is more straightforward. In some cases it might be easier to account for non-stationarity by transforming the inputs with ϕ , in other cases it might be more intuitive to directly construct the set of non-stationary inducing points \hat{U} . With regard to biomedical data, \hat{U} could for example be constructed by placing points in rhythm with the cardiac or respiratory cycle. Interestingly, this allows to model non-stationarity just by defining the point sets U and \hat{U} without ever explicitly formulating the warping function ϕ . Thus, the above reformulation in (6.11) provides useful features beyond standard SKI, allowing to build highly flexible and expressive kernels. All involved matrices are visualized in figure 6.1.

The next step involves specifying the approximation to the full mixture kernel matrix

$$K_{m,XX} \approx \sum_i K_{\phi\text{SKI},i} = \sum_i W_{X,i} K_{U_i U_i} W_{X,i}^\top, \quad (6.12)$$

where now multiple matrices $W_{X,i}$ are used to map from the non-equidistant inducing point sets \hat{U}_i to the inputs X . Thus, the full kernel matrix can be written in terms of sparse interpolation matrices and Kronecker/Toeplitz matrices, substantially relaxing storage requirements and enabling fast MVMs. Another advantage of equation (6.12) is that the size of inducing point sets can be adjusted individually, depending on what kind of accuracy is required for different kernel matrices. As in standard SKI, the computational cost of MVMs is $\mathcal{O}(L + g(\hat{M}))$ with $g(\hat{M}) \leq \hat{M} \log \hat{M}$ where L is the number of inputs X and \hat{M} is an upper bound on the number of inducing points across all \hat{U}_i .

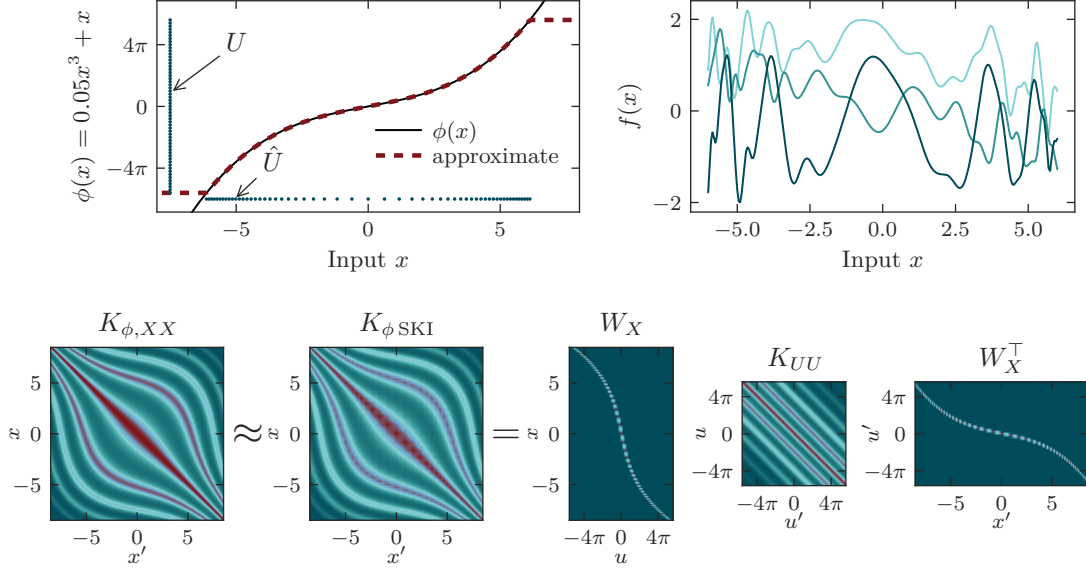


Fig. 6.1: An illustration of the SKI approximation for a quasi-periodic kernel $k_{\text{QP}}(\phi(x), \phi(x'))$ with a non-stationary phase. The upper left panel shows the phase warping function $\phi(x) = 0.05x^3 + x$ and its approximation obtained via inducing points. Two sets of inducing points are shown: U contains equispaced points, $\hat{U} = \Phi^{-1}(U)$ contains the points that are transformed by the inverse of ϕ . The top right panel depicts random functions sampled from the non-stationary, quasi-periodic GP. The bottom panel shows the involved matrices: the exact kernel matrix $K_{\phi,XX}$ is approximated by $K_{\phi\text{SKI}}$ in terms of the sparse interpolation matrix W_X and the Toeplitz matrix K_{UU} .

Fast source separation

The source separation problem stated in equation (6.8) can now be computed efficiently. To this end, the solution α of the linear system is first approximated by

$$\tilde{\alpha} = \left(H \sum_i K_{\phi\text{SKI},i} H^\top + \sigma_n^2 I \right)^{-1} y \quad (6.13)$$

using conjugate gradients and exploiting fast MVMs with the full kernel matrix (and here it will also be assumed that multiplication with H is efficient). An approximate solution $\tilde{\alpha}$ can usually be obtained in a small number of iterations. Given $\tilde{\alpha}$, the posterior mean of the j th source term can be computed efficiently using

$$m_{\text{post},j}(x) \approx w_{x,j} K_{U_j U_j} W_{X,j}^\top H^\top \tilde{\alpha}, \quad (6.14)$$

where $w_{x,j}$ denotes a sparse interpolation vector that maps from U to x .

Fast model learning

The hyperparameters θ of the kernels $k_{\phi,i}$ can be learned by jointly maximizing the marginal log likelihood $\mathcal{L}(\theta|y)$ of the full model. Unfortunately, due to the additive non-stationary structure, the log determinant $\log|K_{m,XX}|$ can not be calculated in closed-form, and previously proposed direct solutions such as the scaled eigenvalue approximation [238] or the Kronecker-based decomposition [183] are not applicable. Therefore, iterative eigenvalue solvers can be employed, i.e., methods that produce a sequence of increasingly accurate estimates of the largest eigenvalues and only access the matrix by evaluating MVMs, such as Lanczos methods, refer to [181] for an overview.

For the optimization of hyperparameters, however, the main requirement is a scalable solution to the log likelihood *and* its derivative, for which Dong et al. [51] recently presented a framework based on *stochastic trace estimation*. To give a brief summary, this technique uses $\log|K| = \text{trace}(\log(K)) = \mathbb{E}[z^\top \log(K)z]$, where z is a random vector with each element being sampled from a Rademacher distribution. The expectation can be estimated by the empirical mean over a small number of probe vectors z . What remains to be solved is the product $\log(K)z$, which is not straightforward, but approximate solutions based on MVMs with the kernel matrix K have been introduced by Dong et al. [51]. The authors also showed how derivatives of the log determinant can be calculated at very low additional cost. The stochastic trace estimation framework only relies on fast MVMs for scalability and thus is directly compatible with the herein proposed form for the kernel matrix $K_{m,XX}$, enabling highly efficient kernel learning.

Finally, other techniques may also be considered for large-scale biomedical problems. A particularly useful technique for learning short-term dynamics models is to split the data into multiple patches (e.g., along the temporal axis), and assume that patches are statistically independent. The optimization can then be solved efficiently on mini-batches using stochastic gradient descent—such techniques have already been employed in the previous chapter 5. However, this approach should be used with caution, because it does not allow to capture any effects with length-scales exceeding the length of patches.

6.2.3 Spatio-temporal kernels

In many relevant biomedical applications, refer to [39, 71, 165, 193], authors have modeled the nonlinear functions ϕ_i to be elementwise functions, that is, they assumed that the functions can be written as a vector of one-dimensional functions

$$\phi_i(x) = \left[\phi_i^{(1)}(x^{(1)}), \dots, \phi_i^{(D)}(x^{(D)}) \right]^\top, \quad (6.15)$$

allowing to formulate the full kernel as

$$k_m(x, z) = \sum_i \prod_d k_i^{(d)} \left(\phi_i^{(d)}(x^{(d)}), \phi_i^{(d)}(z^{(d)}) \right), \quad (6.16)$$

where product separability of the summands is recovered. An important special case of equation (6.16) and the main focus of this chapter are spatio-temporal mixture kernels

$$k_m(s, s', t, t') = \sum_i k_{s,i}(s, s') k_{t,i}(\phi_i(t), \phi_i(t')), \quad (6.17)$$

where s denotes spatial inputs, t denotes time, $k_{s,i}$ models spatial covariance, $k_{t,i}$ models temporal covariance, and non-stationarity is only assumed for the temporal kernel. In such models, the nonlinear functions $\phi_i(t)$ have been termed *time warping* functions [146]. They are typically modeled to be monotonic and account for changes in velocity of the underlying dynamic process. When the temporal kernel is periodic/quasi-periodic, $\phi(t)$ provides the phase of the process by mapping time points to multiples of 2π (this is referred to as the *unwrapped phase*). The instantaneous frequency of the periodic functions is given by the derivative of the phase: $f(t) = \frac{1}{2\pi} \frac{d\phi(t)}{dt}$.

Despite its simplicity, the model defined by equation (6.17) can be applied to a wide range of biomedical problems that are governed by non-stationary periodic patterns. Variants of equation (6.17) have been demonstrated by the author of the thesis and others for removing physiological noise in functional magnetic resonance images [193], extracting ultradian sleep rhythms in newborn infants [39], predicting body surface potential maps [71], and separating cardiac/respiratory components in esophageal high-density manometry [165] or electrical impedance tomography images [71]. In most of these problems, the complexity of using large GP models necessitates efficient solutions to inference and model learning. Sections 6.3 and 6.4 of this chapter will treat two temporal/spatio-temporal applications that benefit greatly from the structure exploitation techniques.

The methods for scalable inference and learning introduced in the previous section 6.2.2 are directly applicable to the spatio-temporal kernel defined by equation (6.17). When the same spatial measurements are available at every time instant, Kronecker structure between the spatial and temporal kernel can be exploited directly (without using any inducing points), and Toeplitz structure can be recovered for the temporal kernel using non-equispaced inducing points as proposed above. When the spatial domain is high-dimensional and spatial points unstructured, inducing points with lattice structure (as in standard SKI) can also be used to impose structure on the spatial kernel matrix. For a given application, one has to assess individually which structure is most useful to be exploited—usually Toeplitz structure is particularly important in temporal data due to the potentially long time axis, which may cause Kronecker methods to become prohibitive with respect to the temporal domain.

6.2.4 Relation to other methods

The herein presented ideas are closely related to the deep kernel learning (DKL) framework by Wilson et al. [239]. DKL uses a non-stationary kernel $k(g(x, \theta), g(x', \theta))$,

where g is a nonlinear function defined by a deep neural network with weights θ and k is a spectral mixture kernel. The key idea of DKL is to jointly optimize the neural network weights and the kernel hyperparameters by maximizing the marginal model likelihood. DKL also employs SKI on the warped inputs $z = g(x, \theta)$ to reach scalability. The model in this work is conceptually similar, but focuses on *mixtures* of non-stationary processes. Thus, in contrast to DKL, multiple sets of non-equispaced warped inducing points are employed, using the reformulation in equation (6.11). On biomedical data it proves beneficial to construct the phase functions from available physiological features (e.g., respiratory, cardiac, gaiting or circadian rhythms), which in turn allows to extract highly interpretable sources from the data. On biomedical data it is often very difficult to also learn the phase warping functions ϕ , as it is done in DKL, due to the large number of local minima.

Another closely related work is the one by Liutkus et al. [124], who investigated the separation of convolutive GPs. The authors presented two different approaches to scalability. Firstly, they considered a *patch-wise* solution to inference and learning (referred to as *framing* in their paper), which however can not account for long temporal correlations. Secondly, they proposed to employ circulant approximations for stationary kernels, thus, enabling efficient matrix decomposition by means of Fourier transforms. This chapter contributes to this line of work by exploiting *existing* structure when kernels in the mixture have non-stationary phases. Very importantly, the herein presented approach also provides good approximations for processes with large length-scales.

6.3 Case study: fetal ECG extraction

This first case study aims to demonstrate the benefits of structure exploitation on a large *temporal* source separation problem. To this end, this section treats the extraction of the *fetal ECG* signal measured non-invasively via electrodes placed on the maternal abdomen. The fetal ECG has been shown to provide valuable information about fetal well-being, including indicators of fetal distress, which is why authors have advocated to monitor fetal ECG morphology [42]. However, in practice, the analysis of non-invasive fetal ECG is hindered by its low signal-to-noise ratio and a strong contamination by *maternal ECG*. The separation of the fetal and maternal ECG components is a longstanding problem and poses a major challenge due to substantial overlap between the two spectra [42]. Several methods have been proposed for solving the multi-channel separation problem, including blind source separation, independent component analysis, or adaptive filtering [184]. For single-channel data, Kalman filtering methods have been investigated [151]. It has also already been demonstrated by Niknazar et al. [150] that the data can be modeled using non-stationary periodic Gaussian processes. However, the authors did not employ any means of structure

exploitation for solving the inference which renders their approach inapplicable to larger datasets.

The aim of this case study is to replicate the experiment by Niknazar et al. [150] and to demonstrate scalability to large datasets using the techniques proposed in section 6.2. The study focuses on one of the recordings (channel 4 in subject R01) from the Physionet fetal ECG database [102]. Similar to Niknazar et al. [151], the additive mixture model

$$s_m(t_k) = s_{\text{maternal}}(t_k) + s_{\text{fetal}}(t_k) + \epsilon_k \quad (6.18)$$

was employed, where $s_{\text{maternal}}(t)$ and $s_{\text{fetal}}(t)$ are zero-mean GPs with covariances $k_{\text{maternal}}(\phi_1(t), \phi_1(t'))$ and $k_{\text{fetal}}(\phi_2(t), \phi_2(t'))$, respectively, and measurements are taken at time-steps $k = 1, \dots, N$. Both k_{maternal} and k_{fetal} are quasi-periodic kernels with a period length of 2π , the time warping functions ϕ_1 and ϕ_2 provide the cardiac phases, and ϵ_k is Gaussian white noise with variance σ_n^2 . The non-stationarity of this model corresponds to variability of the fetal and maternal heart rates. Notice that this formulation did not include a physical measurement model, therefore, the functional \mathcal{H} from equation (6.6) was just defined as $\mathcal{H}s_m(t) = s_m(t_{1:N})$. The model learning and separation problem can be solved for very long recordings (with high sampling rates) using the techniques proposed in this chapter.

As discussed by Niknazar et al. [150], the hyperparameter optimization should be guided by hyperpriors to regularize the solution towards physiologically reasonable results, and it may also prove beneficial to fix some of the hyperparameters. Thus, in equation (6.18), fixed values for the length-scales of the quasi-periodic kernels were used—these parameters can usually be adjusted quite easily based on prior physiological knowledge. For instance, the local length-scale ℓ_{Per} of the quasi-periodic kernel can be adjusted to be slightly smaller than the width of R peaks and the global length-scale ℓ_{SE} should roughly reflect how strongly correlated neighbouring cardiac waveforms are. Thus, in the following experiments $\ell_{\text{Per},\text{fetal}} = 0.1$ rad and $\ell_{\text{Per},\text{maternal}} = 0.25$ rad with $\ell_{\text{SE}} = 18$ rad for both kernels were used. The remaining hyperparameters, namely the standard deviations σ_{maternal} and σ_{fetal} of the two processes, were then optimized using the maximum likelihood approach. In the herein considered dataset, the nonlinear functions ϕ_1 and ϕ_2 could be derived directly by detecting R peaks in the available reference signals and assuming a linear phase between R peaks. It is also possible to learn the cardiac phase using marginal likelihood maximization, which has been demonstrated for ECG data in [152], but shall not be the focus of this work.

To demonstrate the effectiveness of the proposed method, the approximate structure exploiting approach was first compared to the classical GP regression on a short excerpt of $N = 5000$ data points, corresponding to 10 s of data sampled at 500 Hz. In order to demonstrate scalability, a large-scale stresstest was then performed on $N = 10^5$ data points, corresponding to 100 s of data sampled at 1000 Hz. Separation success was quantified by calculating the SNR improvement of the fetal ECG—this metric

has been described in the context of ECG denoising by Petersen et al. [162]. Two sets of non-equispaced inducing points were used to recover Toeplitz structure, which enables efficient inference and learning by means of fast Fourier transforms. The hyperparameter optimization was executed using stochastic trace estimation with 20 probe vectors. The tolerance of CG was set to 10^{-1} for parameter learning and $5 \cdot 10^{-3}$ for source separation. Figure 6.2 shows an excerpt of the large-scale source separation and a comparison of exact/approximate marginal likelihood evaluations. Performance metrics are provided in table 6.1.

Table 6.1: Comparison of different experiments conducted on the fetal ECG dataset.

	batch GP	SKI	SKI
Input points	5000	5000	10^5
Inducing points	–	3400 (m) + 4800 (f)	14300 (m) + 21600 (f)
Time for Inference	2.60 s	0.27 s	4.14 s
Time for learning	28.5 s	47.4 s	462.3 s
σ_{fetal}	5.08	5.95	5.11
σ_{maternal}	21.48	21.42	32.11
SNR improvement	18.6 dB	18.1 dB	18.2 dB

SNR: signal-to-noise ratio; batch GP: the standard solution to Gaussian process regression using Cholesky decomposition; SKI: the variant of structured kernel interpolation proposed in this chapter;




Fig. 6.2: An illustration of the fetal ECG learning and separation problem. The left panel shows an excerpt of data from the $N = 10^5$ stresstest, including both raw data and separated ECG signals. The right panel shows the negative log likelihood for different values of σ_{maternal} on the smaller $N = 5000$ dataset. The black line corresponds to the exact solution, the red line corresponds to the approximation done by means of stochastic trace estimation with 20 probe vectors.

6.4 Case study: reconstruction of periodic processes in electrical impedance tomography (EIT)

After having demonstrated the benefits of temporal structure exploitation in a single-channel source separation problem, this second case study is focused on *spatio-temporal* source separation. Thus, scalable GP regression techniques are applied to *electrical impedance tomography* (EIT) data. This imaging modality relies on injecting small currents through pairs of electrodes attached to the skin surface and measuring the resulting voltages. As opposed to other imaging techniques, EIT is radiation free and supports high temporal sampling rates. The main focus of this study is the estimation of dynamic changes in lung tissue impedance, which can be measured with a belt of electrodes placed around the chest, and is increasingly being used for bedside monitoring of intensive care patients [60]. Impedance changes measured by thoracic EIT are mainly caused by variations of the pulmonary air content during respiration. It allows to derive a number of relevant clinical parameters such as the spatial heterogeneity of ventilation distribution or regional lung compliances [45]. To a lesser extent, thoracic EIT is influenced by changes in the blood content of the lung tissue, which are induced by cardiac activity and superposed on the ventilation. Thus, it has also been suggested that EIT enables monitoring of pulmonary perfusion [60]. The amplitude of cardiac related impedance changes is quite small compared to the more prominent lung ventilation, making this evaluation more challenging, but there is emerging evidence for its diagnostic value, refer to [60]. In mechanically ventilated patients, the separation of these two effects—lung ventilation and perfusion—is an important problem, that to this date remains largely unsolved. In particular, point-wise temporal Fourier filtering is not applicable because the higher harmonics of the ventilation usually overlap with the frequencies of the perfusion signal. Several algorithms have been proposed based on gating [58], principal component analysis (PCA) [164], or empirical mode decomposition [110]. Yet, none of the methods manage to solve the problem to a satisfactory degree: they either require initial training phases leading to lack of adaptivity, or rely entirely on modeling temporal dynamics while omitting relevant spatial effects. It is also worth noting that most previously proposed EIT separation methods solve the problem in two separate steps, that is, they first solve the inverse problem in each time step individually and then apply filtering to the pixels of the reconstructed image series.

The following section proposes a novel approach to EIT perfusion/ventilation separation based on a spatio-temporal mixture of non-stationary Gaussian processes as in equation (6.6). This method enables to incorporate prior knowledge about the statistics of both signal components. Very importantly, the reconstruction and separation is treated in an *end-to-end* fashion, that is, a probabilistic forward model is formulated and then the inverse problem is solved with respect to the perfusion and

ventilation images in closed form. The method exploits both Kronecker and Toeplitz structure for fast MVMs, rendering the inference scalable. This work builds upon the spatio-temporal EIT model proposed by Graßhoff and Rostalski [79], who applied GPs to reconstructed EIT images, but did not yet incorporate a solution to the inverse problem, i.e., images were reconstructed prior to the GP filtering. The following section will first provide necessary background on EIT reconstruction and then proceed by presenting the novel structure exploiting model.

6.4.1 Temporal EIT solver

This section largely adopts the problem formulation and notation from the classical EIT reconstruction literature, refer for example to [6, 7]. The goal in EIT is to solve the inverse problem with respect to the spatio-temporal resistivity distribution $\rho(s, t)$ of the lung tissue (measured in Ω m for 3D or Ω for 2D domains). In practice, the tissue conductivity $\sigma(s, t) = 1/\rho(s, t)$ (measured in S m⁻¹ in 3D or S in 2D domains) is often more convenient and employed by most authors. In the following, it is assumed that ρ and σ are real-valued quantities. The EIT forward problem can be formulated in terms of a nonlinear operator \mathcal{F} that maps from the conductivity distribution $\sigma(s, t_k)$ to a Q -dimensional vector of electrical potentials $v(t_k)$ measured on the boundary of the domain for a set of Q electrodes at timesteps t_k :

$$v(t_k) = \mathcal{F}(\sigma(s, t_k)) + \epsilon_k, \quad (6.19)$$

where $\epsilon_k \sim \mathcal{N}(0, \sigma_n^2 I)$ is Gaussian measurement noise and the forward operator \mathcal{F} is established mathematically in terms of Maxwell's equations with a certain set of boundary conditions, which are provided in appendix A.3. The forward operator also depends on the stimulation pattern, the type of electrodes used and the amplitude of the injected current. Here, the main focus is the *temporal* EIT reconstruction problem, that is, $\sigma(s, t)$ is to be estimated from noisy measurements $v(t_k)$ at timesteps $k = 1, \dots, T$.

It was assumed that the conductivity is given by an additive mixture

$$\sigma(s, t_k) = \sigma_{\text{vent}}(s, t_k) + \sigma_{\text{perf}}(s, t_k) + \sigma_0(s) \quad (6.20)$$

of a spatio-temporal perfusion and a ventilation process, denoted by $\sigma_{\text{vent}}(s, t)$ and $\sigma_{\text{perf}}(s, t)$, respectively, and $\sigma_0(s)$ denotes a constant homogeneous background conductivity. The goal is to reconstruct and separate the components of model (6.20) using the voltage measurements given by equation (6.19). Note that equations (6.19) and (6.20) are a nonlinear variant of the general mixture model that has been introduced in equation (6.6), i.e., the nonlinear functional \mathcal{F} takes the role of \mathcal{H} . Based on the considerations in section 6.2, a scalable temporal EIT solver can be formulated that exploits prior knowledge of the perfusion and ventilation process. As a model for the

temporal evolution of the data, periodic Gaussian process priors were used on both the perfusion and ventilation component:

$$\text{cov}(\sigma_{\text{vent}}(s, t), \sigma_{\text{vent}}(s, t')) = k_{\phi, \text{vent}}(t, t') = k_{\text{vent}}(\phi_{\text{vent}}(t), \phi_{\text{vent}}(t')) \quad (6.21)$$

$$\text{cov}(\sigma_{\text{perf}}(s, t), \sigma_{\text{perf}}(s, t')) = k_{\phi, \text{perf}}(t, t') = k_{\text{perf}}(\phi_{\text{perf}}(t), \phi_{\text{perf}}(t')), \quad (6.22)$$

where k_{vent} and k_{perf} are quasi-periodic kernels with period length 2π , and ϕ_{vent} and ϕ_{perf} denote the nonlinear phases of the respiratory and cardiac cycle, respectively. To numerically solve the forward problem, the spatial domain was discretized using FEM. This allows to represent the conductivity function in timestep t_k via a vector $\sigma_k \in \mathbb{R}^L$ corresponding to the L tetrahedral or triangular elements of an FEM mesh. The forward problem could thus be written as

$$v_k = F(\sigma_k) + \epsilon_k = F(\sigma_{\text{vent},k} + \sigma_{\text{perf},k} + \sigma_0) + \epsilon_k, \quad (6.23)$$

where $F(\cdot)$ denotes the nonlinear mapping calculated by means of an FEM solver, and $\sigma_{\text{vent},k}$, $\sigma_{\text{perf},k}$ and σ_0 are vectors representing the FEM elements of the ventilation, perfusion and constant signal component, respectively. A common practice in EIT is to linearize the forward problem for small variations around the reference conductivity vector σ_0 , which is often referred to as *difference EIT* [7]. The Jacobian (also referred to as sensitivity matrix) of the forward FEM solution is defined as $J^{(i,j)} = \left. \frac{\partial F^{(i)}(\sigma)}{\partial \sigma^{(j)}} \right|_{\sigma=\sigma_0}$, allowing to approximate the relationship with

$$y_k = Jx_k + \epsilon_k, \quad (6.24)$$

where $x_k = \sigma_k - \sigma_0$ and $y_k = v_k - v_0$. Here v_0 denotes a reference potential vector, which can in practice be obtained by averaging measured voltages across end-expiratory time points (i.e., when the conductivity is relatively stable).

For the spatial domain, the matrix $P \in \mathbb{R}^{Q \times Q}$ is defined to model prior covariance between the discrete elements of the FEM—several prior matrices have been proposed for EIT, a common choice is to use the inverse of the *discrete Laplacian filter* matrix evaluated over the FEM grid, refer to [6]. This prior provides a simple model for smoothness across the spatial domain. It does not correspond to a proper GP distribution, but can be easily calculated in FEM meshes and proves to perform very well in EIT reconstruction problems. Assuming separability of the spatial and temporal process, the prior covariance for the vector $\tilde{x} = [x_1^\top, \dots, x_T^\top]^\top$, concatenating FEM elements of all time points, can be given in terms of the Kronecker product

$$K_m = (K_{\phi, \text{perf}} + K_{\phi, \text{vent}}) \otimes P, \quad (6.25)$$

where $K_{\phi, \text{perf}} = k_{\phi, \text{perf}}(t_{1:T}, t_{1:T})$ and $K_{\phi, \text{vent}} = k_{\phi, \text{vent}}(t_{1:T}, t_{1:T})$. Using the techniques proposed in this chapter, the source separation problem with respect to σ_{vent} and σ_{perf}

can be solved efficiently. For the convenience of the reader, the complete solution is provided here: given the vector $\tilde{y} = [y_1^\top, \dots, y_T^\top]^\top$ consisting of all voltage measurements, the posterior mean for the perfusion/ventilation at a test point t_* is given by

$$\sigma_{\text{vent,post}}(t_*) = (k_{\phi,\text{vent}}(t_*, t_{1:T}) \otimes P) \tilde{J}^\top \tilde{\alpha} = \left(k_{\phi,\text{vent}}(t_*, t_{1:T}) \otimes (PJ^\top) \right) \tilde{\alpha} \quad (6.26)$$

$$\sigma_{\text{perf,post}}(t_*) = (k_{\phi,\text{perf}}(t_*, t_{1:T}) \otimes P) \tilde{J}^\top \tilde{\alpha} = \left(k_{\phi,\text{perf}}(t_*, t_{1:T}) \otimes (PJ^\top) \right) \tilde{\alpha} \quad (6.27)$$

with

$$\tilde{\alpha} = \left(\tilde{J}K_m\tilde{J}^\top + \sigma_n^2 I \right)^{-1} \tilde{y} \quad (6.28)$$

$$= \left((K_{\phi,\text{perf}} + K_{\phi,\text{vent}}) \otimes (JPJ^\top) + \sigma_n^2 I \right)^{-1} \tilde{y}, \quad (6.29)$$

where $\tilde{J} = I \otimes J$. Similar equations have been used in the spatio-temporal reconstruction framework by Adler et al. [7]. One can immediately exploit Kronecker structure between the spatial and temporal domain for fast MVMs. Toeplitz structure can be induced via the techniques described in section 6.2.2. Thus, the temporal kernels were approximated via $K_{\phi,\text{vent}} \approx W_1 K_{\phi,\text{vent},\hat{U}_1\hat{U}_1} W_1^\top$ and $K_{\phi,\text{perf}} \approx W_2 K_{\phi,\text{perf},\hat{U}_2\hat{U}_2} W_2^\top$, where \hat{U}_1 and \hat{U}_2 are the inducing points, respectively. Then the solution to the system of linear equations can be calculated efficiently using CG as described in section 6.2.2, which ensures scalability to very long time axes with $T \gg 10^4$ samples. As opposed to previous methods, the solution in equations (6.26)–(6.29) allows to directly estimate the perfusion/ventilation images from the measured voltages in an end-to-end fashion.

Hyperparameter learning

What remains to be solved is the tuning of kernel hyperparameters. As discussed before, the length-scales of periodic processes can often be adjusted manually based on prior knowledge of the data, for instance, the local length-scale is chosen to roughly reflect short-term smoothness of the data while the global length-scale is tuned to reflect the correlation between adjacent respiratory/cardiac cycles. Here, fixed values of $\ell_{\text{SE,vent}} = 5$ rad, $\ell_{\text{Per,vent}} = 0.75$ rad were used for the ventilation prior, and $\ell_{\text{SE,perf}} = 30$ rad, $\ell_{\text{Per,perf}} = 0.4$ rad for the perfusion prior. The proposed GP solution appeared to be not too sensitive regarding the precise choice of length-scales, thus, no lengthy tuning was necessary. Alternatively to manual adjustment, one may also consider patch-based optimization techniques for learning the temporal length-scales (as in chapter 5).

For learning the standard deviation of the noise and kernel, Graßhoff and Rostalski [79] used marginal likelihood maximization on the reconstructed images, which however tends to converge very slowly due to the noise involved in the stochastic trace estimation. Therefore, in this work, a novel, highly scalable method is proposed, inspired by

Adler and Guardo [3], who used a noise amplification metric for adjusting regularization parameters. The herein proposed method is based on the observation that —once the length-scale parameters of $k_{\phi,\text{perf}}$ and $k_{\phi,\text{vent}}$ are fixed— the posterior means in equations (6.26) and (6.27) only depend on the *ratio* between the noise variance σ_n^2 and the variances of the perfusion/ventilation priors, denoted by σ_{perf}^2 and σ_{vent}^2 , respectively. Thus, one can simplify the hyperparameter tuning problem by only adjusting two ratios, defined as $\beta^2 = \frac{\sigma_{\text{vent}}^2}{\sigma_{\text{perf}}^2}$ and $\lambda^2 = \frac{\sigma_n^2}{\sigma_{\text{perf}}^2 + \sigma_{\text{vent}}^2}$, where β balances the influence of the two source processes and λ controls the trade-off between data fit and prior. Both β and λ take the role of classical regularization parameters. Note that the posterior means in equations (6.26) and (6.27) are fully defined by β and λ (which can be easily seen when factoring out σ_n^2 from the inverse term). Interestingly, after fixing β and λ , this leaves one degree of freedom for choosing *absolute* values for σ_n^2 , σ_{perf}^2 and σ_{vent}^2 without influencing the result of equations (6.26) and (6.27) — this degree of freedom can be understood quite naturally as scaling the noise variance σ_n^2 in concert with the prior variance $\sigma_{\text{perf}}^2 + \sigma_{\text{vent}}^2$. Notice that the above proposed simplification (of only tuning the ratios β and λ) does not suffice when the posterior uncertainty is also of interest, because for a well calibrated uncertainty estimate all standard deviations (σ_n^2 , σ_{perf}^2 and σ_{vent}^2) would have to be optimized for (e.g., using the marginal likelihood approach). But, as in this case study the main concern was on reconstructing a good posterior mean, a full marginal likelihood optimization was not necessary. Empirical tests showed that $\beta = 2$ seemed to be a robust choice across most scenarios, implementing the prior belief that amplitudes of the ventilation process are roughly twice as high as those of the perfusion process. Then, the only remaining parameter was λ , which was optimized via the approach by Adler et al. [3], who proposed to use the reconstruction’s noise amplification as a performance metric. To this end, the *noise figure* $\text{NF} = \text{SNR}_i / \text{SNR}_o$ was calculated as the ratio of SNRs of the input voltage and the reconstructed output images, denoted by SNR_i and SNR_o , respectively. Following recommendations in [4], λ was optimized such that $\text{NF} = 0.5$ was reached.

As in section 6.3, fixed respiratory/cardiac phase functions ϕ_{perf} and ϕ_{vent} were used, which could be derived from available reference signals (such as ECG or ventilator data). Alternatively, it may also sometimes be possible to detect breaths and heart beats directly in one of the measured EIT voltages without reconstructing a full EIT image.

Dual reconstruction model

This work employs a 2D-3D dual model for solving the image reconstruction [5]. This modeling approach is based on the idea of solving the FEM forward problem on a detailed 3D model and then projecting the solution vertically to a coarse 2D mesh. The image reconstruction was thus performed in 2D, making it computationally efficient, while also preserving relevant properties of the higher-dimensional model [5]. This




Fig. 6.3: The dual (2.5D) reconstruction model. The mesh on the left represents the 3D thorax model on which the FEM forward problem was solved. The solution was then projected vertically to the 2D mesh on the right, which was used for solving the inverse problem.

technique, referred to as 2.5D modeling, outperforms standard 2D FEM solutions [37] and is provided as one of the preset models in the EIDORS library [8]. EIDORS was used to calculate the Jacobian J_{3D} on the 3D phantom model depicted in figure 6.3 (generated with NETGEN [198]), consisting of $\sim 2.5 \times 10^4$ tetrahedral elements. The corresponding 2D model depicted in figure 6.3 can be thought of as a cross-section through the 3D phantom at the height of the electrodes, and it was constructed as a partial 64×64 square grid, consisting of $\sim 6.5 \times 10^3$ triangles. Using EIDORS a projection matrix C was calculated (called “coarse-to-fine mapping” in EIDORS), where the entry $C^{(i,j)}$ is the fraction of the i th 3D element contained in the j th 2D element along the vertical axis [69]. The Jacobian J of the coarse model is thus given by $J = J_{3D}C$, which allows for efficient MVMs and is therefore used in the GP-based source separation.

6.4.2 Results on simulation and clinical data

To demonstrate feasibility of the herein proposed perfusion/ventilation separation technique, results on both simulation and experimental data are presented in the following.

Simulation case study

The 3D phantom model depicted in figure 6.4 was used to simulate pulsatile sources, corresponding to cardiac and respiratory activity. To this end, an FEM mesh was generated with EIDORS [8] and NETGEN [198] with two embedded lung geometries and a spherical volume for the heart, comprising a total of $\sim 1.4 \times 10^5$ elements. Around the model, 16 electrodes were positioned equidistantly at half of its height, and an adjacent stimulation pattern was used to simulate output voltages. All geometries

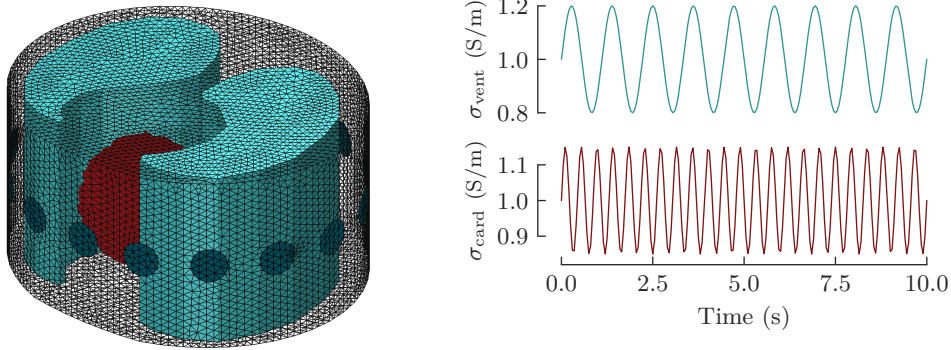


Fig. 6.4: The model that was used to generate artificial EIT data. The mesh on the left represents the 3D model which contains embedded geometries for the lung and heart. Respiratory and cardiac activity was simulated by changing the conductivity of the mesh elements according to the functions σ_{vent} and σ_{card} plotted on the right.

were modeled to have homogeneous, time-varying conductivities and all remaining elements were given a constant background conductivity of $\sigma_0 = 1 \text{ S m}^{-1}$. Similar to previous publications [216], the heart and lung cycles were simulated via sinusoidal functions, which are plotted in figure 6.4. The resulting voltages were calculated using the EIDORS solver. The reconstruction of the ventilation and perfusion signals was solved via the GP based separation techniques introduced above, the CG tolerance was set to 5×10^{-4} , results are plotted in figure 6.5a. This case study demonstrates the feasibility of separating pulsatile components in an end-to-end fashion: notably, in the region between the heart and lung (pixel 2 in figure 6.5a) the two sources are separated successfully despite their spatial proximity.

Clinical case study

Next, a clinical case of a spontaneously breathing neonate in intensive care is presented. The data were previously recorded, and published in [88]. The neonate was in prone position and 16-electrode EIT data were recorded for 220 frames at 13 Hz, leading to a total of 45 760 independent measurements. This dataset has previously been investigated with GPs, refer to [79], but there the solution has not yet been calculated in end-to-end fashion. On this dataset, the phase functions ϕ_{perf} and ϕ_{vent} could be extracted directly from the data. To this end, the impedance time-course of two spatial points were calculated using standard spatial reconstruction available in EIDORS and evaluated. This reconstruction step was done prior to the full spatio-temporal reconstruction proposed in this chapter, but only requires to evaluate two individual pixels, thus, posing very small additional computational costs. The respiratory phase was estimated from a pixel located in the left lung by detecting the peaks of the respiratory waveform and assuming a piece-wise linear phase in between peaks. The

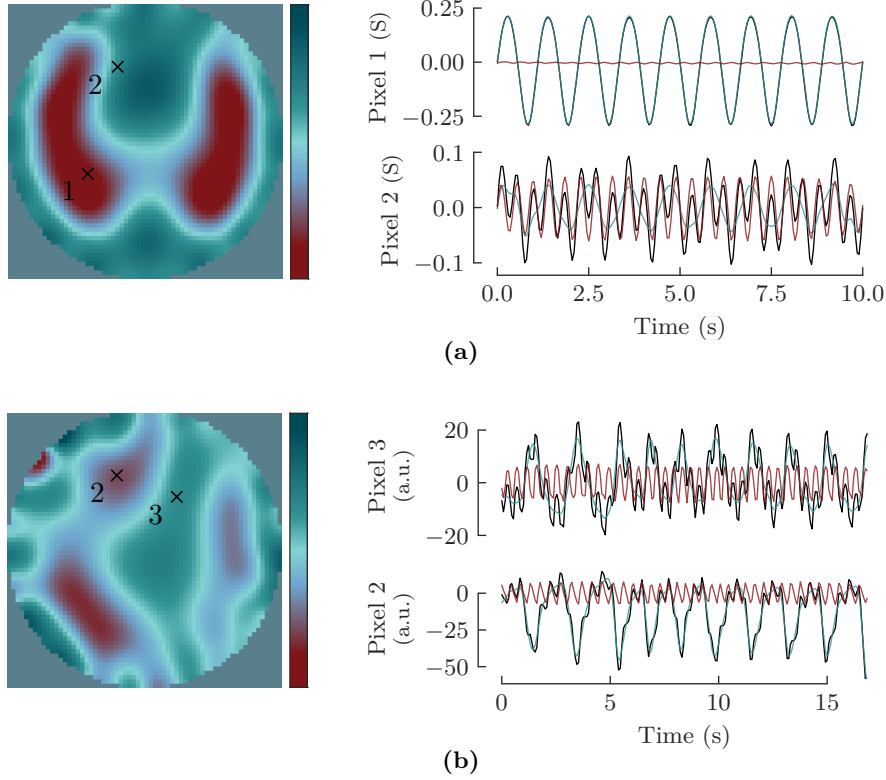


Fig. 6.5: An illustration of reconstructed EIT images and their temporal evolution at selected pixels. The red and light green line show posterior means of σ_{perf} and σ_{vent} , respectively, the black line shows $\sigma_{\text{perf}} + \sigma_{\text{vent}}$. (a) Result of the simulation case study. At both pixels the separated signals recover the sinusoidal patterns from figure 6.4. (b) Result of the clinical case study. Both in the right lung (pixel 2) and in the mediastinum (pixel 3) a superposition of cardiac and respiratory activity can be observed, which is successfully separated by means of the proposed GP model.

cardiac phase was extracted in the same way by detecting peaks of the cardiac waveform from a pixel close to the heart (where the perfusion signal was strongest). Then, the perfusion and the ventilation images were solved for with the structure exploiting techniques proposed above. As in the simulation experiment, the CG tolerance was set to 5×10^{-4} . An excerpt of the full kernel matrix $K_{\phi, \text{vent}}$ corresponding to two spatial locations is shown in figure 6.6. The separated signals are plotted in figure 6.5b. This case study demonstrates the applicability of the proposed EIT framework to clinical datasets with non-stationary respiratory and cardiac phases.




Fig. 6.6: An excerpt of the prior kernel matrix $K_{\phi,vent}$. The matrix was evaluated at two spatial points (s_1 and s_2) on the clinical dataset from figure 6.5b. The non-stationarity of the data stems from fluctuation in respiratory rate and is reflected by the warped near-Toeplitz structure of the temporal submatrices.

6.5 Discussion

This chapter has investigated additive mixtures of non-stationary GPs for modeling large biomedical datasets. The approach exploits Toeplitz/Kronecker structure of the kernel matrix, enabling highly scalable inference and model learning. To account for non-stationarity, the technique uses SKI with multiple sets of non-equispaced/warped inducing points. The validity of this approach was demonstrated in two relevant temporal/spatio-temporal source separation problems in sections 6.3 and 6.4. The findings of the chapter indicate that the proposed techniques have a high potential of uncovering new and relevant structure in biomedical datasets, where GPs could not be applied so far, in particular in long spatio-temporal recordings of respiratory/cardiac activity. This work is part of a larger push in the recent GP literature that aims to alleviate computational requirements of classic GP regression by accessing the kernel matrix only through fast matrix vector multiplications [51, 63, 237, 240]. As in standard SKI, the herein proposed method reduces computational cost for inference and model learning to a worst case of $\mathcal{O}(N + M \log M)$ when exploiting only Toeplitz structure. Thus, the presented method complements existing works on GP source separation, e.g., by Liutkus et al. [124], who investigated convolutive GP mixtures, but relied on framing and circulant approximations for scalability. The approach by Liutkus et al. did not account for potential non-stationarities of the mixture terms, which limits applicability to biomedical problems. Being closely related to deep kernel learning (DKL), the main difference of the herein proposed technique is that

it employs multiple mapping functions ϕ to model non-stationarities of (biomedical) data. On temporal/spatio-temporal data it would be much more difficult to infer the nonlinear mapping functions ϕ_i and would require strong hyperpriors regularizing the optimization, which has been discussed for ECG data by Noorzadeh et al. [152]. Generally for biomedical applications, simply using fixed phase functions—derived from available features—is a viable option and allows to extract very meaningful structure from the data.

The case studies in section 6.3 and 6.4 aim to demonstrate the potential of the proposed source separation techniques on large biomedical datasets. To this end, the first case study replicated experiments by Niknazar et al. [150] on extracting fetal/maternal ECG signals from single-channel abdominal recordings. Scalability to 10^5 data points was achieved for both inference and learning by inducing Toeplitz structure. The performance metrics in table 6.1 reveal a strong improvement of the fetal ECG’s SNR, which is in the same order of magnitude as the classical batch GP solution. The second case study was concerned with the separation of perfusion and ventilation processes in electrical impedance tomography. Building upon earlier work by Graßhoff and Rostalski [79], the thesis has developed the first *end-to-end* solution to this problem, which is based on using a spatio-temporal prior and then solving the inverse problem with respect to the mixture components. Scalability was reached by exploiting existing Kronecker structure and inducing Toeplitz structure on the temporal domain via inducing points. The proposed solution was demonstrated on simulation data and a first clinical case. The presented test cases already prove feasibility of the approach, but further experimental data will be needed for a more comprehensive validation.

Non-stationary phase functions can also be implemented via equivalent state-space models using the Kalman filtering techniques described in section 2.1.4, which theoretically enables linear complexity in the number of timesteps. However in practice, these state-space models often require a high state-dimension to reach a good approximation, making the solution computationally expensive. It has also been recognized that the state-space approach tends to become slow when gaps between data points Δt_i are uneven or when kernel inputs depend on a nonlinear function $\phi(t)$, because much of the computation time is spent on the recalculation of the discrete-time transition matrices [149]. Nickisch et al. [149] proposed a remedy for this problem by precomputing transition matrices on a fixed set of Δt_i values. When choosing between the *batch* approach investigated in this chapter and the *recursive* state-space approach one has to trade-off the different factors influencing the runtime for the specific problem at hand. One of the main advantages of the state-space approach is that, by design, it estimates uncertainties of the prediction in each timestep. In contrast, the CG based methods presented in this chapter provide scalable estimates for the *posterior mean*, but are less suited for estimating the *posterior uncertainty* (because it would require to solve N linear systems). As an extension to techniques discussed in this chapter,

scalable sampling-based methods could be considered to also estimate uncertainties via matrix-vector multiplications, refer to [127] for more details.

Chapter 7

Conclusion

Completely non-invasive techniques for estimating respiratory activity may facilitate the assessment of clinically relevant quantities in patients under mechanical ventilation. This thesis has explored surface electromyography and Gaussian process regression, two techniques enabling to quantify the respiratory effort of patients. The main focus of this work was on the development of model-based estimators that allow to infer latent parameters and states of the respiratory system by integrating pneumatic and electrical measurements. The following sections provide a summary of the thesis' key findings and an outlook to future research directions.

7.1 Summary and key insights

The current clinical gold standard for assessing respiratory effort relies on measuring the esophageal pressure with an invasive, nasogastric catheter. This technique is rarely used in clinical practice as it is considered to be error-prone and time-consuming. Surface electromyography (sEMG) represents a promising alternative: Its measurement only requires placement of a few electrodes on the patient's thorax and gives direct insight into the electrical activity of respiratory muscles. This thesis has made important contributions to respiratory effort estimation by combining ventilator and sEMG data in a single model. To this end, it has explored a range of statistical models, including estimators based on robust regression and Gaussian process regression. These considerations have also spurred the development of new methods for scalable Gaussian process regression on non-stationary biomedical data.

The first part of the investigation (chapter 3) was concerned with characterizing the relationship between sEMG and respiratory muscle pressure P_{mus} . The analysis proved, for the first time, that an appropriately preprocessed sEMG signal is well correlated to P_{mus} , and it also showed that the relation between these two quantities is approximately linear. The chapter has proposed a method for leveraging the availability of multiple channels: a simple channel selection technique based on the signal-to-noise ratio significantly improved correlation to P_{mus} . Thus, it provided strong evidence for the benefits of selecting the channel with the most favorable measurement conditions, whereas earlier studies [27] tended to favor channel averaging. The key insight from

the chapter is that, despite the substantial level of noise, sEMG represents a viable surrogate for the total pressure exerted by the respiratory muscles.

The following parts of the thesis (chapter 4 and 5) were concerned with the development of novel inference algorithms that incorporate sEMG and available pneumatic data to estimate P_{mus} completely non-invasively. All methods developed in these chapters were based on a sensor fusion paradigm, which means that noise in the estimated quantity (P_{mus}) is reduced by combining all available measurements. Two different directions have been pursued: the first one was focused on robustness against artifacts, the other one on incorporating prior knowledge about the unknown respiratory effort signal. Both directions are intended to address the shortcomings of simpler least squares estimators, which often fail to produce viable estimates because they are overly sensitive to outliers, signal delay and model mismatch [111, 147, 180, 215].

Chapter 4 introduced several new techniques for increasing the robustness of respiratory parameter estimation. This includes a technique for determining the respiratory time constant τ as well as a new integrated variant of the equation of motion, and using robust regression methods instead of classical least squares regression. The results prove that the performance of P_{mus} estimation could be significantly improved through these approaches— notably, a very high breath-wise correlation of 0.95 with invasively measured P_{mus} was attained by combining all the new techniques mentioned above. One of the key takeaways of this chapter is that model-based methods offer a remedy for many of the shortcomings of respiratory sEMG: The level of noise is reduced substantially and the scaling of the sEMG signal can be identified without using any disruptive maneuvers. There are still limitations regarding the composition of the input data, for instance, parameters can only be identified if the data contain sufficient variation of patient-ventilator workload shares.

The family of models presented in chapter 5 sought to exploit prior knowledge about the statistics of P_{mus} signals. Earlier approaches to modeling P_{mus} were either based on highly nonlinear neurological models [142–144, 209] or rather simplistic models, such as template functions [61, 136, 182] or (damped) harmonic oscillators [159, 196]. In contrast, chapter 5 employed Gaussian processes (GPs), a Bayesian nonparametric model, to capture key characteristics of the data by means of dedicated kernel functions. The estimation of respiratory effort could then be posed as a Bayesian inference problem, to which two different solutions have been presented. The results prove that the incorporation of prior knowledge is highly beneficial and might allow to estimate P_{mus} on a single pressure support level without using any disruptive maneuvers. However, in both chapter 4 and 5, the highest performances were reached by estimating the neuromechanical coupling index during short airway occlusion maneuvers, which is an important takeaway regarding the practical use of sEMG.

Many authors have employed Gaussian processes, because they generalize well on small datasets and are inherently suited for quantifying prediction uncertainties [12, 170]. The results obtained in this thesis support the use of Gaussian processes as an

effective tool for modeling non-stationary biomedical signals with temporal or spatio-temporal structure. In particular, they provide straightforward means to account for complex physiological processes that are governed by mixtures of quasi-periodic patterns, which is very common in biomedical applications. Moreover, Gaussian processes are closely related to other widely used methods, such as Tikhonov regularization and Kalman filtering/smoothing, which makes them very useful as a unifying framework and for analyzing the relation between different algorithms.

A recurring theme of this thesis was the reduction of computational complexity: the runtime of the classical GP regression formulation scales cubically in the number of data points, which is intractable for many (spatio-temporal) biomedical datasets. Two different directions have been investigated. Chapter 5 has used a reformulation in terms of equivalent state-space models, allowing to solve the inference via Kalman smoothing with quasi-linear complexity. This approach is highly useful, but still breaks down on many spatio-temporal datasets, because the state-dimension often becomes excessively large. Thus, chapter 6 has investigated the exploitation of existing Toeplitz/Kronecker structure within the kernel matrix. To this end, it has been demonstrated that structured kernel interpolation [240] with multiple sets of non-equispaced inducing points can be used to account for non-stationarity of the data. This technique enables highly scalable inference and hyperparameter learning. It has led to the first end-to-end solution for reconstructing pulsatile sources (perfusion and ventilation) in electrical impedance tomography (EIT) data.

7.2 Outlook

The findings made in this thesis confirm the great potential of sEMG as a monitoring tool in intensive care. Still, several limitations exist that warrant further clinical studies and methodological research. The following section points out some key research gaps and paths for further development. It also provides an outlook on potential therapeutic applications of the technologies presented in the thesis.

A crucial step towards a practical application of sEMG-based respiratory effort monitoring will be its validation on severely ill patients undergoing long-term mechanical ventilation. The data used in chapters 3–5 did not involve any patients with acute or chronic respiratory failure and previous studies on ICU patients have relied on rather small cohorts [27, 159]. It seems evident that the simple equation of motion is not applicable to represent severe ARDS and COPD. Therefore, the models used throughout this thesis should be extended to better account for these pathologies and validated on further study data. Refer to [90, 91] for recent work in that direction.

Another important practical hurdle are the limited real-time capabilities of the herein presented algorithms. All estimators used in this thesis have been implemented as offline methods and were solved on pre-recorded data. Note that real-time require-

ments depend on the specific use-case: for purely diagnostic questions, offline analyses may suffice, but for therapeutic purposes (e.g., adjusting ventilator settings at the bedside), online processing with small delays would be required. The development of an online solution is challenging. Firstly, a real-time capable algorithm for detecting and removing ECG artifacts is needed — to this end, gating seems to be a good choice, because it allows for relatively small delays and could be easily combined with wavelet filter banks. Secondly, for online P_{mus} estimation, sequential algorithms such as recursive least squares or Kalman filters can be used, which was already demonstrated by Petersen [159], but should be validated on a larger cohort. A pragmatic solution to this problem would be to perform parameter estimation and P_{mus} calculation separately: during an initial training phase, parameters could be identified with an offline algorithm and then be treated as constant values during online P_{mus} reconstruction. A similar approach has been taken for example in [114, 200, 241] and in proportional pressure support ventilation modes.

This raises another important research question: when and how often should maneuvers be used? Each maneuver entails a disruption to the natural breathing rhythm of a patient, but at the same time improves identifiability of respiratory parameters. A practical solution to this trade-off is complex and requires further research. Ideally, maneuvers would be triggered automatically when a certain level of parametric uncertainty is exceeded. To make matters more complicated, the identifiability of parameters is also influenced differently depending on the maneuver and the spontaneous breathing pattern of patients. Thus, maneuvers should be selected individually for each patient in a way that is least disruptive to them. This thesis has investigated mainly airway occlusion maneuvers and pressure support steps, but other options could also prove useful, such as a variation of trigger and cycling criteria, refer to a recent patent application [56], which was co-written by the author of this thesis.

There are several paths for further methodological development. A combination of the methods used in chapter 4 and chapter 5 seems to be a natural next step towards improving the estimation performance. To this end, a Gaussian process prior for P_{mus} could be combined with a robust non-Gaussian likelihood function. Inference in such a model is usually intractable and necessitates the use of approximate methods, such as expectation propagation, variational inference, or variants thereof. An interesting line of methodological research that is currently very active is the development of scalable sequential solutions to this class of inference problems, refer to [149, 236] for recent developments. Yet another, completely different direction is to investigate fully data-driven approaches, that is, to use machine learning methods such as *deep neural networks*. A first proof-of-concept of this idea has been published in [35], a paper co-written by the author of this thesis. In these methods, no or very little prior knowledge would be exploited, that is, the relationship between measurements and latent variables has to be learned entirely from the available data. While state-of-the-art machine learning makes it possible to represent complex nonlinear relationships,

its performance relies on the availability of large training datasets that reflect clinical variability of patients and treatments. To this date, there is no single dataset of mechanically ventilated patients that could fulfill this requirement, because collecting data from critically ill patients is costly, requires ethical approval and often involves time-consuming procedures. Research in this direction could be advanced by fusing existing datasets and augmenting simulated or synthetic data.

Finally, it should be pointed out that potential therapeutic applications may be derived from the herein proposed techniques. This work was concerned mainly with the extraction of diagnostic information from the sEMG signal, in particular for the estimation of respiratory effort. A potential future clinical use could rely on “closing the loop”, that is, using the sEMG signal directly to control ventilatory parameters. This idea has an enormous potential for improving patient care and relieving clinical staff, but of course it entails major technological and regulatory hurdles that must be overcome. A first approach in this direction is to trigger the ventilator on the sEMG signal (or on the estimated P_{mus} signal) to reduce patient-ventilator asynchronies. Another highly promising approach is an sEMG-based proportional pressure ventilation mode, where ventilator assistance is given in proportion to the estimated P_{mus} signal. This would give patients full control over the breathing pattern, thus, enabling a much more physiological mode of ventilation. Notably, proportional pressure support ventilation has been hypothesized to offer both lung- and diaphragm-protection because it retains the patient’s physiological control mechanisms against lung injury and prevents diaphragm atrophy [106].

Appendix

A.1 Study data

Table A.1: Clinical characteristics of included patients ($n = 43$) from [76].

Characteristic	Result
Age, mean \pm SD year	64 \pm 11
Men, n (%)	34 (79)
Weight, mean \pm SD kg	79.0 \pm 8.9
BMI, mean \pm SD kg m ⁻²	26 \pm 6
TLC, mean \pm SD l	6.7 \pm 2.4
VC, mean \pm SD l	3.6 \pm 1.4
FEV1 % predicted, mean \pm SD %	82 \pm 29
Tiffeneau index, mean \pm SD %	69 \pm 19
RV, mean \pm SD l	3.4 \pm 2.4
RV/TLC, mean \pm SD %	45 \pm 19
iPEEP, mean \pm SD cmH ₂ O	2.1 \pm 1.4
Diagnosis, n (%)	
OSAS	4 (9)
COPD	16 (37)
GOLD I	2 (5)
GOLD II-III	14 (33)
ACOS	3 (7)
Bronchial asthma	5 (19)
ILD	7 (17)
Lung cancer	20 (47)
Infectious or rheumatic diseases	11 (26)

BMI: Body-mass index; TLC: total lung capacity; VC: vital capacity; FEV1: forced expiratory volume in 1 s; RV: residual volume; iPEEP: intrinsic positive end-expiratory pressure; OSAS: obstructive sleep apnea syndrome; COPD: chronic obstructive pulmonary disease; ACOS: asthma–COPD overlap syndrome; ILD: interstitial lung disease;

A.2 Kalman filtering and smoothing algorithm

The following section provides algorithms for executing the linear Kalman filter [108] and Rauch-Tung-Striebel smoother [172] using the state-space model

$$\begin{aligned}x_k &= A_{k-1}x_{k-1} + q_{k-1} \\y_k &= H_kx_k + r_k,\end{aligned}\tag{A.1}$$

with system matrix A_{k-1} , measurement matrix H_k , additive process noise $q_{k-1} \sim \mathcal{N}(0, Q_{k-1})$ and additive measurement noise $r_k \sim \mathcal{N}(0, R_k)$. The negative log marginal likelihood of the model can be calculated alongside a forward Kalman filter pass [187, section 12.3.2].

Algorithm 1: Linear Kalman filter and log marginal likelihood evaluation

input : Initial mean m_0 ; initial covariance P_0 ; discrete-time state-space model matrices A_{k-1} , Q_{k-1} , H_k , R_k depending on hyperparameters θ ; measurements y_k for $k = 1, \dots, N$

output: Mean m_k and covariance P_k of filtered state estimates $p(x_k|y_{1:k})$ for $k = 1, \dots, N$; negative log marginal likelihood ϕ_N of hyperparameters θ

$\phi_0 \leftarrow 0$

for $k \leftarrow 1$ **to** N **do**

 // prediction step

$m_k^- \leftarrow A_{k-1}m_{k-1}$

$P_k^- \leftarrow A_{k-1}P_{k-1}A_{k-1}^\top + Q_{k-1}$

 // update step

$v_k \leftarrow y_k - H_k m_k^-$

$S_k \leftarrow H_k P_k^- H_k^\top + R_k$

$K_k \leftarrow P_k^- H_k^\top S_k^{-1}$

$m_k \leftarrow m_k^- + K_k v_k$

$P_k \leftarrow P_k^- - K_k S_k K_k^\top$

 // negative log marginal likelihood

$\phi_k \leftarrow \phi_{k-1} + \frac{1}{2} \log |2\pi S_k| + \frac{1}{2} v_k^\top S_k^{-1} v_k$

end

Algorithm 2: Linear Rauch-Tung-Striebel smoother

input : Kalman filter estimates m_k and P_k , and discrete-time state-space model matrices A_{k-1} , Q_{k-1} for $k = 1, \dots, N$

output: Mean m_k^s and covariance P_k^s of smoother state estimates $p(x_k|y_{1:N})$ for $k = 1, \dots, N$

$m_N^s \leftarrow m_N$
 $P_N^s \leftarrow P_N$
for $k \leftarrow N - 1$ **to** 1 **do**

$m_{k+1}^- \leftarrow A_k m_k$
$P_{k+1}^- \leftarrow A_k P_k A_k^\top + Q_k$
$G_k \leftarrow P_k A_k^\top (P_{k+1}^-)^{-1}$
$m_k^s \leftarrow m_k + G_k (m_{k+1}^s - m_{k+1}^-)$
$P_k^s \leftarrow P_k + G_k (P_{k+1}^s - P_{k+1}^-) G_k^\top$

end

A.3 EIT forward equations

The following equations are reproduced from [213]. Let σ be the conductivity distribution within a domain Ω with a piece-wise smooth boundary $\partial\Omega$. The forward electrical impedance tomography (EIT) model uses a static variant of Maxwell's equations, which is given by

$$\nabla \cdot (\sigma \nabla u) = 0 \quad \text{in } \Omega \tag{A.2}$$

with u being the scalar electrical potential. The domain boundary comprises L electrodes with boundaries e_l and contact impedances z_l ($l = 1, \dots, L$). The complete electrode model (CEM) is defined by the following set of boundary conditions:

$$u + z_l \sigma \frac{\partial u}{\partial v} = U_l \quad \text{on } e_l \text{ for } l = 1, \dots, L \tag{A.3}$$

$$\int_{e_l} \sigma \frac{\partial u}{\partial v} = I_l \quad \text{for } l = 1, \dots, L \tag{A.4}$$

$$\sigma \frac{\partial u}{\partial v} = 0 \quad \text{on } \partial\Omega \setminus \bigcup_{l=1}^L e_l, \tag{A.5}$$

where v is the outward normal vector on $\partial\Omega$, U_l is the potential at the l th electrode, and I_l is the current injected at the l th electrode.

References

- [1] Abbas, H., Graßhoff, J., Rostalski, P., and Brinkmann, R. “On the estimation of optoacoustic waves in retinal laser therapy using Gaussian processes”. In: *Proceedings on Automation in Medical Engineering*. Vol. 1. 1. 2020.
- [2] Abunurah, H., Russell, D., and Lowman, J. “The validity of surface EMG of extra-diaphragmatic muscles in assessing respiratory responses during mechanical ventilation: A systematic review”. *Pulmonology* 26.6 (2020), pp. 378–385.
- [3] Adler, A. and Guardo, R. “Electrical impedance tomography: regularized imaging and contrast detection”. *IEEE Transactions on Medical Imaging* 15.2 (1996), pp. 170–179.
- [4] Adler, A., Arnold, J. H., Bayford, R., Borsic, A., Brown, B., Dixon, P., Faes, T. J. C., Frerichs, I., Gagnon, H., Gärber, Y., Grychtol, B., Hahn, G., Lionheart, W. R. B., Malik, A., Patterson, R. P., Stocks, J., Tizzard, A., Weiler, N., and Wolf, G. K. “GREIT: a unified approach to 2D linear EIT reconstruction of lung images”. *Physiological Measurement* 30.6 (2009), S35–S55.
- [5] Adler, A., Borsic, A., Polydorides, N., and Lionheart, W. R. “Simple FEMs aren’t as good as we thought: experiences developing EIDORS v3.3”. In: *9th EIT conference 2008*. 2008.
- [6] Adler, A. and Boyle, A. “Electrical Impedance Tomography: Tissue Properties to Image Measures”. *IEEE Transactions on Biomedical Engineering* 64.11 (2017), pp. 2494–2504.
- [7] Adler, A., Dai, T., and Lionheart, W. R. B. “Temporal image reconstruction in electrical impedance tomography”. *Physiological Measurement* 28.7 (2007), S1–S11.
- [8] Adler, A. and Lionheart, W. R. B. “Uses and abuses of EIDORS: an extensible software base for EIT”. *Physiological Measurement* 27.5 (2006), S25–S42.
- [9] Agostoni, E., Sant’Ambrogio, G., and Portillo Carrasco, H. “Electromyography of the diaphragm in man and transdiaphragmatic pressure”. *Journal of Applied Physiology* 15.6 (1960), pp. 1093–1097.

- [10] Akoumianaki, E., Maggiore, S. M., Valenza, F., Bellani, G., Jubran, A., Loring, S. H., Pelosi, P., Talmor, D., Grasso, S., Chiumello, D., Guérin, C., Patroniti, N., Ranieri, V. M., Gattinoni, L., Nava, S., Terragni, P.-P., Pesenti, A., Tobin, M., Mancebo, J., and Brochard, L. “The Application of Esophageal Pressure Measurement in Patients with Respiratory Failure”. *American Journal of Respiratory and Critical Care Medicine* 189.5 (2014), pp. 520–531.
- [11] Aliverti, A., Cala, S. J., Duranti, R., Ferrigno, G., Kenyon, C. M., Pedotti, A., Scano, G., Sliwinski, P., Macklem, P. T., and Yan, S. “Human respiratory muscle actions and control during exercise”. *Journal of Applied Physiology* 83.4 (1997), pp. 1256–1269.
- [12] Álvarez, M., Luengo, D., and Lawrence, N. D. “Latent Force Models”. In: *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 5. 2009, pp. 9–16.
- [13] Anders, C., Brose, G., Hofmann, G. O., and Scholle, H.-C. “Evaluation of the EMG–force relationship of trunk muscles during whole body tilt”. *Journal of Biomechanics* 41.2 (2008), pp. 333–339.
- [14] Arnal, J.-M., Garnero, A., Saoli, M., and Chatburn, R. L. “Parameters for Simulation of Adult Subjects During Mechanical Ventilation”. *Respiratory Care* 63.2 (2018), pp. 158–168.
- [15] Aron, E. “Der intrapleurale Druck beim lebenden, gesunden Menschen”. *Archiv für Pathologische Anatomie und Physiologie und für Klinische Medicin* 160.1 (1900), pp. 226–234.
- [16] Athanasiades, A., Ghorbel, F., Clark, J. W., Niranjana, S. C., Olansen, J., Zwischenberger, J. B., and Bidani, A. “Energy Analysis of a Nonlinear Model of the Normal Human Lung”. *Journal of Biological Systems* 8.2 (2000), pp. 115–139.
- [17] Avanzolini, G., Barbini, P., Cappello, A., Cevenini, G., and Chiari, L. “A new approach for tracking respiratory mechanical parameters in real-time”. *Annals of Biomedical Engineering* 25.1 (1997), pp. 154–163.
- [18] Avilés-Rojas, N. and Hurtado, D. E. “Whole-lung finite-element models for mechanical ventilation and respiratory research applications”. *Frontiers in Physiology* 13 (2022).
- [19] Bartolo, A., Roberts, C., Dzwonczyk, R. R., and Goldman, E. “Analysis of diaphragm EMG signals: comparison of gating vs. subtraction for removal of ECG contamination”. *Journal of Applied Physiology* 80.6 (1996), pp. 1898–1902.
- [20] Bates, J. H. T. *Lung Mechanics: An Inverse Modeling Approach*. Cambridge University Press, 2009.

-
- [21] Baydur, A., Behrakis, P., Zin, W. A., Jaeger, M. J., and Milic-Emili, J. “A simple method for assessing the validity of the esophageal balloon technique”. *The American Review of Respiratory Disease* 126.5 (1982), pp. 788–91.
- [22] Beaton, A. E. and Tukey, J. W. “The Fitting of Power Series, Meaning Polynomials, Illustrated on Band-Spectroscopic Data”. *Technometrics* 16.2 (1974), pp. 147–185.
- [23] Becher, T. H. “Bestimmung der Atemmechanik bei unterstützter Spontanatmung durch minimale Variation des Unterstützungsdrucks”. PhD thesis. Christian-Albrechts-Universität zu Kiel, 2013.
- [24] Beck, J., Gottfried, S. B., Navalesi, P., Skrobik, Y., Comtois, N., Rossini, M., and Sinderby, C. “Electrical Activity of the Diaphragm during Pressure Support Ventilation in Acute Respiratory Failure”. *American Journal of Respiratory and Critical Care Medicine* 164.3 (2001), pp. 419–424.
- [25] Beck, J., Sinderby, C., Lindström, L., and Grassino, A. “Effects of lung volume on diaphragm EMG signal strength during voluntary contractions”. *Journal of Applied Physiology* 85 (1998), pp. 1123–34.
- [26] Behrakis, P. K. “Respiratory Mechanics during Anesthesia in Humans”. PhD thesis. McGill University, 1986.
- [27] Bellani, G., Bronco, A., Arrigoni Marocco, S., Pozzi, M., Sala, V., Eronia, N., Villa, G., Foti, G., Tagliabue, G., Eger, M., and Pesenti, A. “Measurement of Diaphragmatic Electrical Activity by Surface Electromyography in Intubated Subjects and Its Relationship With Inspiratory Effort”. *Respiratory Care* 63.11 (2018), pp. 1341–1349.
- [28] Bellani, G., Laffey, J. G., Pham, T., Fan, E., Brochard, L., Esteban, A., Gattinoni, L., Haren, F., Larsson, A., McAuley, D. F., Ranieri, M., Rubenfeld, G., Thompson, B. T., Wrigge, H., Slutsky, A. S., and Pesenti, A. “Epidemiology, Patterns of Care, and Mortality for Patients With Acute Respiratory Distress Syndrome in Intensive Care Units in 50 Countries”. *JAMA* 315.8 (2016), p. 788.
- [29] Bellani, G., Mauri, T., Coppadoro, A., Grasselli, G., Patroniti, N., Spadaro, S., Sala, V., Foti, G., and Pesenti, A. “Estimation of Patient’s Inspiratory Effort From the Electrical Activity of the Diaphragm”. *Critical Care Medicine* 41.6 (2013), pp. 1483–1491.
- [30] Bellemare, F., Bigland-Ritchie, B., and Woods, J. J. “Contractile properties of the human diaphragm in vivo”. *Journal of Applied Physiology* 61.3 (1986), pp. 1153–1161.
- [31] Bertoni, M., Spadaro, S., and Goligher, E. C. “Monitoring Patient Respiratory Effort During Mechanical Ventilation: Lung and Diaphragm-Protective Ventilation”. *Critical Care* 24.1 (2020).

- [32] Blanch, L., Sales, B., Montanya, J., Lucangelo, U., Garcia-Esquirol, O., Villagra, A., Chacón, E., Estruga, A., Borelli, M., Burgueño, M. J., Oliva, J., Fernandez, R., Villar, J., Kacmarek, R., and Murias, G. “Validation of the Better Care® system to detect ineffective efforts during expiration in mechanically ventilated patients: a pilot study”. *Intensive Care Medicine* 38.5 (2012), pp. 772–780.
- [33] Bland, J. and Altman, D. G. “Statistical methods for assessing agreement between two methods of clinical measurement”. *The Lancet* 327.8476 (1986), pp. 307–310.
- [34] Bland, J. M. and Altman, D. G. “Agreement Between Methods of Measurement with Multiple Observations Per Individual”. *Journal of Biopharmaceutical Statistics* 17.4 (2007), pp. 571–582.
- [35] Bockelmann, N., Graßhoff, J., Hansen, L., Bellani, G., Heinrich, M. P., and Rostalski, P. “Deep Learning for Prediction of Diaphragm Activity from the Surface Electromyogram”. *Current Directions in Biomedical Engineering* 5.1 (2019), pp. 17–20.
- [36] Boudnik, L., Graßhoff, J., Vollmer, F., and Rostalski, P. “Muscle Artifact Removal in Single-Channel Electrocardiograms using Temporal Convolutional Networks”. In: *Proceedings of the 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2022.
- [37] Boyle, A. and Adler, A. “Modelling with 2.5D Approximations”. In: *17th Conference on Electrical Impedance Tomography*. 2016.
- [38] Braun, N. M., Arora, N. S., and Rochester, D. F. “Force-length relationship of the normal human diaphragm”. *Journal of Applied Physiology* 53.2 (1982), pp. 405–412.
- [39] Brügge, N. S., Graßhoff, J., Weigenand, A., and Rostalski, P. “Multi-Task Gaussian Process Regression for the Detection of Sleep Cycles in Premature Infants”. In: *ICASSP 2022 – 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022.
- [40] Carreaux, G., Lyazidi, A., Cordoba-Izquierdo, A., Vignaux, L., Joliet, P., Thille, A. W., Richard, J.-C. M., and Brochard, L. “Patient-Ventilator Asynchrony During Noninvasive Ventilation”. *Chest* 142.2 (2012), pp. 367–376.
- [41] Clancy, E., Morin, E., and Merletti, R. “Sampling, noise-reduction and amplitude estimation issues in surface electromyography”. *Journal of Electromyography and Kinesiology* 12.1 (2002), pp. 1–16.
- [42] Clifford, G. D., Silva, I., Behar, J., and Moody, G. B. “Non-invasive fetal ECG analysis”. *Physiological Measurement* 35.8 (2014), pp. 1521–1536.

-
- [43] Cunningham, J. P., Shenoy, K. V., and Sahani, M. “Fast Gaussian Process Methods for Point Process Intensity Estimation”. In: *Proceedings of the 25th International Conference on Machine Learning (ICML)*. ACM, 2008, pp. 192–199.
- [44] Daoud, E. G., Katigbak, R., and Ottochian, M. “Accuracy of the Ventilator Automated Displayed Respiratory Mechanics in Passive and Active Breathing Conditions: A Bench Study”. *Respiratory Care* 64.12 (2019), pp. 1555–1560.
- [45] Dargaville, P. A., Rimensberger, P. C., and Frerichs, I. “Regional tidal ventilation and compliance during a stepwise vital capacity manoeuvre”. *Intensive Care Medicine* 36.11 (2010), pp. 1953–1961.
- [46] De Luca, C. J. “The Use of Surface Electromyography in Biomechanics”. *Journal of Applied Biomechanics* 13.2 (1997), pp. 135–163.
- [47] De Troyer, A. and Boriek, A. M. “Mechanics of the Respiratory Muscles”. *Comprehensive Physiology* (2011), pp. 1273–1300.
- [48] Di Mussi, R., Spadaro, S., Volta, C. A., Bartolomeo, N., Trerotoli, P., Staffieri, F., Pisani, L., Iannuzziello, R., Dalfino, L., Murgolo, F., and Grasso, S. “Continuous assessment of neuro-ventilatory drive during 12 h of pressure support ventilation in critically ill patients”. *Critical Care* 24.1 (2020).
- [49] Dionne, A., Parkes, A., Engler, B., Watson, B. V., and Nicolle, M. W. “Determination of the best electrode position for recording of the diaphragm compound muscle action potential”. *Muscle & Nerve* 40.1 (2009), pp. 37–41.
- [50] Donders, F. C. “Beitrage zum Mechanismus der Respiration und Circulationim gesunden und kranken Zustande”. *Zeitschrift für Rationelle Medicin* 3 (1953), pp. 287–319.
- [51] Dong, K., Eriksson, D., Nickisch, H., Bindel, D., and Wilson, A. “Scalable Log Determinants for Gaussian Process Kernel Learning”. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS)*. 2017, pp. 6327–6337.
- [52] Downey, R. “Anatomy of the Normal Diaphragm”. *Thoracic Surgery Clinics* 21.2 (2011), pp. 273–279.
- [53] Dres, M., Dubé, B.-P., Mayaux, J., Delemazure, J., Reuter, D., Brochard, L., Similowski, T., and Demoule, A. “Coexistence and Impact of Limb Muscle and Diaphragm Weakness at Time of Liberation from Mechanical Ventilation in Medical Intensive Care Unit Patients”. *American Journal of Respiratory and Critical Care Medicine* 195.1 (2017), pp. 57–66.

- [54] Duiverman, M. L., Eykern, L. A., Vennik, P. W., Koëter, G. H., Maarsingh, E. J. W., and Wijkstra, P. J. “Reproducibility and responsiveness of a noninvasive EMG technique of the respiratory muscles in COPD patients and in healthy subjects”. *Journal of Applied Physiology* 96.5 (2004), pp. 1723–1729.
- [55] Eger, M., Petersen, E., Graßhoff, J., and Rostalski, P. *Process and signal processing unit for determining a pneumatic parameter with the use of a lung-mechanical model and of a gradient model*. DE102020133460A1 / US11890416B2 / CN113143250B, 2020.
- [56] Eger, M., Rostalski, P., Petersen, E., and Graßhoff, J. *Method and signal processing unit for determining the respiratory activity of a patient*. DE102019007717B3 / CN114599277A / US20220379057A1 / WO2021089215A1, 2019.
- [57] Esteban, A., Anzueto, A., Alía, I., Gordo, F., Apezteguía, C., Pálizas, F., Cide, D., Goldwasser, R., Soto, L., Bugedo, G., Rodrigo, C., Pimentel, J., Raimondi, G., and Tobin, M. J. “How Is Mechanical Ventilation Employed in the Intensive Care Unit? An International Utilization Review”. *American Journal of Respiratory and Critical Care Medicine* 161.5 (2000), pp. 1450–1458.
- [58] Eyuboglu, B. M. and Brown, B. H. “Methods of cardiac gating applied potential tomography”. *Clinical Physics and Physiological Measurement* 9.4A (1988), pp. 43–48.
- [59] Farina, D., Merletti, R., and Stegeman, D. F. “Biophysics of the Generation of EMG Signals”. In: *Electromyography. Physiology, Engineering, and Noninvasive Applications*. John Wiley & Sons, Ltd., 2004. Chap. 2.
- [60] Frerichs, I., Amato, M. B. P., Kaam, A. H., Tingay, D. G., Zhao, Z., Grychtol, B., Bodenstein, M., Gagnon, H., Böhm, S. H., Teschner, E., Stenqvist, O., Mauri, T., Torsani, V., Camporota, L., Schibler, A., Wolf, G. K., Gommers, D., Leonhardt, S., and Adler, A. “Chest electrical impedance tomography examination, data analysis, terminology, clinical use and recommendations: consensus statement of the TRanslational EIT developmeNt stuDy group”. *Thorax* 72.1 (2016), pp. 83–93.
- [61] Fresnel, E., Muir, J.-F., and Letellier, C. “Realistic human muscle pressure for driving a mechanical lung”. *EPJ Nonlinear Biomedical Physics* 2.1 (2014).
- [62] Friedrich, S., Groll, A., Ickstadt, K., Kneib, T., Pauly, M., Rahnenführer, J., and Friede, T. “Regularization approaches in clinical biostatistics: A review of methods and their applications”. *Statistical Methods in Medical Research* 32.2 (2022), pp. 425–440.

-
- [63] Gardner, J. R., Pleiss, G., Bindel, D., Weinberger, K. Q., and Wilson, A. G. “GPYtorch: Blackbox Matrix-matrix Gaussian Process Inference with GPU Acceleration”. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS)*. 2018, pp. 7587–7597.
- [64] Gattinoni, L. and Quintel, M. “Fifty Years of Research in ARDS – Why Is Acute Respiratory Distress Syndrome So Important for Critical Care?” *American Journal of Respiratory and Critical Care Medicine* 194.9 (2016), pp. 1051–1052.
- [65] Goligher, E. C., Brochard, L. J., Reid, W. D., Fan, E., Saarela, O., Slutsky, A. S., Kavanagh, B. P., Rubenfeld, G. D., and Ferguson, N. D. “Diaphragmatic myotrauma: a mediator of prolonged ventilation and poor patient outcomes in acute respiratory failure”. *The Lancet Respiratory Medicine* 7.1 (2019), pp. 90–98.
- [66] Goligher, E. C., Fan, E., Herridge, M. S., Murray, A., Vorona, S., Brace, D., Rittayamai, N., Lanys, A., Tomlinson, G., Singh, J. M., Bolz, S.-S., Rubenfeld, G. D., Kavanagh, B. P., Brochard, L. J., and Ferguson, N. D. “Evolution of Diaphragm Thickness during Mechanical Ventilation. Impact of Inspiratory Effort”. *American Journal of Respiratory and Critical Care Medicine* 192.9 (2015), pp. 1080–1088.
- [67] Goligher, E. C., Jonkman, A. H., Dianti, J., Vaporidi, K., Beitler, J. R., Patel, B. K., Yoshida, T., Jaber, S., Dres, M., Mauri, T., Bellani, G., Demoule, A., Brochard, L., and Heunks, L. “Clinical strategies for implementing lung and diaphragm-protective ventilation: avoiding insufficient and excessive effort”. *Intensive Care Medicine* 46.12 (2020), pp. 2314–2326.
- [68] Goligher, E. C. et al. “Lung- and Diaphragm-Protective Ventilation”. *American Journal of Respiratory and Critical Care Medicine* 202.7 (2020), pp. 950–961.
- [69] Gong, B., Schullcke, B., Krueger-Ziolek, S., Vauhkonen, M., Wolf, G., Mueller-Lisse, U., and Moeller, K. “EIT Imaging Regularization Based on Spectral Graph Wavelets”. *IEEE Transactions on Medical Imaging* 36.9 (2017), pp. 1832–1844.
- [70] Graßhoff, J., Jankowski, A., and Rostalski, P. “Scalable Gaussian Process Regression for Kernels with a Non-Stationary Phase”. *arXiv preprint arXiv:1912.11713* (2019).
- [71] Graßhoff, J., Jankowski, A., and Rostalski, P. “Scalable Gaussian Process Separation for Kernels with a Non-Stationary Phase”. In: *Proceedings of the 37th International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 3722–3731.

- [72] Graßhoff, J., Männel, G., Abbas, H. S., and Rostalski, P. “Model Predictive Control using Efficient Gaussian Processes for Unknown Disturbance Inputs”. In: *Proceedings of the 58th IEEE Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 2708–2713.
- [73] Graßhoff, J., Männel, G., and Rostalski, P. “Nonparametric modeling of quasi-periodic signals – application to esophageal pressure filtering”. In: *Annual Conference of the German Society for Biomedical Engineering (BMT)*. DGBMT, 2018.
- [74] Graßhoff, J., Petersen, E., Becher, T., and Rostalski, P. “Automatic Estimation of Respiratory Effort using Esophageal Pressure”. In: *Proceedings of the 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2019, pp. 4646–4649.
- [75] Graßhoff, J., Petersen, E., Eger, M., Bellani, G., and Rostalski, P. “A Template Subtraction Method for the Removal of Cardiogenic Oscillations on Esophageal Pressure Signals”. In: *Proceedings of the 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2017, pp. 2235–2238.
- [76] Graßhoff, J., Petersen, E., Rostalski, P., and Walterpacher, S. “Surface EMG-based Quantification of Inspiratory Effort: A Quantitative Comparison with Pes”. *Critical Care* 25.441 (2021).
- [77] Graßhoff, J., Petersen, E., Rostalski, P., and Walterspacher, S. “A Comparison of Pes- and sEMG-derived Measures for the Quantification of Inspiratory Effort”. In: *Online Meeting of the Pleural Pressure Working Group (PLUG)*. 2021.
- [78] Graßhoff, J., Petersen, E., Walterspacher, S., and Rostalski, P. “Model-Based Estimation of Inspiratory Effort Using Surface EMG”. *IEEE Transactions on Biomedical Engineering* 70.1 (2023), pp. 247–258.
- [79] Graßhoff, J. and Rostalski, P. “Spatio-Temporal Gaussian Processes for Separation of Ventilation and Perfusion Related Signals in EIT Data”. In: *Proceedings on Automation in Medical Engineering*. Vol. 1. 1. 2020.
- [80] Grassino, A., Goldman, M. D., Mead, J., and Sears, T. A. “Mechanics of the human diaphragm during voluntary contraction: statics”. *Journal of Applied Physiology* 44.6 (1978), pp. 829–839.
- [81] Grinnan, D. C. and Truwit, J. “Clinical review: Respiratory mechanics in spontaneous and assisted ventilation”. *Critical Care* 9.5 (2005), pp. 472–484.

-
- [82] Gruner, J., Schmid, N., Männel, G., Graßhoff, J., Abbas, H. S., and Rostalski, P. “Recursively Feasible Model Predictive Control using Latent Force Models Applied to Disturbed Quadcopters”. In: *Proceedings of the 61st IEEE Conference on Decision and Control (CDC)*. IEEE, 2022.
- [83] Hamnegard, C., Wragg, S., Mills, G., Kyroussis, D., Road, J., Daskos, G., Bake, B., Moxham, J., and Green, M. “The effect of lung volume on transdiaphragmatic pressure”. *European Respiratory Journal* 8.9 (1995), pp. 1532–1536.
- [84] Hampel, F. R. “The Influence Curve and its Role in Robust Estimation”. *Journal of the American Statistical Association* 69.346 (1974), pp. 383–393.
- [85] Hartikainen, J. and Särkkä, S. “Kalman filtering and smoothing solutions to temporal Gaussian process regression models”. In: *2010 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*. 2010, pp. 379–384.
- [86] Hartikainen, J., Seppänen, M., and Särkkä, S. “State-Space Inference for Non-Linear Latent Force Models with Application to Satellite Orbit Prediction”. In: *Proceedings of the 29th International Conference on Machine Learning (ICML)*. 2012.
- [87] Heckman, C. and Enoka, R. M. “Motor Unit”. *Comprehensive Physiology* (2012), pp. 2629–2682.
- [88] Heinrich, S., Schiffmann, H., Frerichs, A., Klockgether-Radke, A., and Frerichs, I. “Body and head position effects on regional lung ventilation in infants: an electrical impedance tomography study”. *Intensive Care Medicine* 32.9 (2006). Data from https://eidors3d.sourceforge.net/data_contrib/if-neonate-spontaneous/, p. 1392.
- [89] Hennig, P., Osborne, M. A., and Kersting, H. P. *Probabilistic Numerics: Computation as Machine Learning*. Cambridge University Press, 2022.
- [90] Hennigs, C., Bilda, F., Graßhoff, J., Waltersbacher, S., and Rostalski, P. “A switching lung mechanics model for detection of expiratory flow limitation”. *at – Automatisierungstechnik* (in press, 2024).
- [91] Hennigs, C., Scholleman, F., Graßhoff, J., Hardel, T., and Rostalski, P. “A simple mathematical lung model of Chronic Obstructive Pulmonary Disease (COPD)”. In: *Proceedings on Automation in Medical Engineering*. Vol. 2. 1. 2023.
- [92] Herrero, J. L., Khuvis, S., Yeagle, E., Cerf, M., and Mehta, A. D. “Breathing above the brain stem: volitional control and attentional modulation in humans”. *Journal of Neurophysiology* 119.1 (2018), pp. 145–159.

- [93] Heunks, L. and Ottenheim, C. “Diaphragm-Protective Mechanical Ventilation to Improve Outcomes in ICU Patients?” *American Journal of Respiratory and Critical Care Medicine* 197.2 (2018), pp. 150–152.
- [94] Heyer, L., Baconnier, P. F., Eberhard, A., Biot, L., Viale, J.-P., Perdrix, J.-P., and Carry, P.-Y. “Non-invasive detection of respiratory muscles activity during assisted ventilation”. *Comptes Rendus Biologies* 325.4 (2002), pp. 383–391.
- [95] Hickling, K. G. “The Pressure–Volume Curve Is Greatly Modified by Recruitment: A Mathematical Model of ARDS Lungs”. *American Journal of Respiratory and Critical Care Medicine* 158.1 (1998), pp. 194–202.
- [96] Hof, A. L. “The relationship between electromyogram and muscle force”. *Sportverletzung - Sportschaden* 11.03 (1997), pp. 79–86.
- [97] Huber, P. J. “Robust Estimation of a Location Parameter”. *The Annals of Mathematical Statistics* 35.1 (1964), pp. 73–101.
- [98] Huber, P. J. *Robust statistics*. Wiley Series in Probability and Statistics, John Wiley & Sons, 1981.
- [99] IngMar Medical Ltd., *ASL 5000 User’s Manual*. SW 3.6, Rev. 1. 2016.
- [100] Iotti, G. A., Braschi, A., Brunner, J. X., Smits, T., Olivei, M., Palo, A., and Veronesi, R. “Respiratory mechanics by least squares fitting in mechanically ventilated patients: Applications during paralysis and during pressure support ventilation”. *Intensive Care Medicine* 21.5 (1995), pp. 406–413.
- [101] Jansen, D., Jonkman, A. H., Roesthuis, L., Gadgil, S., Hoeven, J. G., Scheffer, G.-J. J., Girbes, A., Doorduyn, J., Sinderby, C. S., and Heunks, L. M. A. “Estimation of the diaphragm neuromuscular efficiency index in mechanically ventilated critically ill patients”. *Critical Care* 22.1 (2018).
- [102] Jezewski, J., Matonia, A., Kupka, T., Roj, D., and Czabanski, R. “Determination of fetal heart rate from abdominal signals: evaluation of beat-to-beat accuracy in relation to the direct fetal electrocardiogram”. *Biomedical Engineering / Biomedizinische Technik* 57.5 (2012). Data from <https://physionet.org/content/adfecgdb/1.0.0/>.
- [103] Jiang, T.-X., Reid, W. D., Belcastro, A., and Road, J. D. “Load Dependence of Secondary Diaphragm Inflammation and Injury after Acute Inspiratory Loading”. *American Journal of Respiratory and Critical Care Medicine* 157.1 (1998), pp. 230–236.
- [104] Jodat, R. W., Horgan, J. D., and Lange, R. L. “Simulation of Respiratory Mechanics”. *Biophysical Journal* 6.6 (1966), pp. 773–785.

-
- [105] Jonkman, A. H., Warnaar, R. S. P., Baccinelli, W., Carbon, N. M., D’Cruz, R. F., Doorduyn, J., Doorn, J. L. M., Elshof, J., Estrada-Petrocelli, L., Graßhoff, J., Heunks, L. M. A., Koopman, A. A., Langer, D., Moore, C. M., Nunez Silveira, J. M., Petersen, E., Poddighe, D., Ramsay, M., Rodrigues, A., Roesthuis, L. H., Rossel, A., Torres, A., Duiverman, M. L., and Oppersma, E. “Analysis and applications of respiratory surface EMG: report of a round table meeting”. *Critical Care* 28.1 (2024).
- [106] Jonkman, A. H., Vries, H. J., and Heunks, L. M. A. “Physiology of the Respiratory Drive in ICU Patients: Implications for Diagnosis and Treatment”. *Critical Care* 24.1 (2020).
- [107] Kahl, L., Petersen, E., Graßhoff, J., and Rostalski, P. *Process and device for determining a respiratory and/or cardiogenic signal*. US20220330837A1 / DE112020000232A5 / WO2021063601A1, 2019.
- [108] Kalman, R. E. “A New Approach to Linear Filtering and Prediction Problems”. *Journal of Basic Engineering* 82.1 (1960), pp. 35–45.
- [109] Keys, R. “Cubic convolution interpolation for digital image processing”. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 29.6 (1981), pp. 1153–1160.
- [110] Ko, Y.-F., Cheng, K.-S., and Su, P.-L. “Separation of Heart and Lung-related Signals in Electrical Impedance Tomography Using Empirical Mode Decomposition”. *Current Medical Imaging* 18.13 (2022), pp. 1396–1415.
- [111] Kondili, E., Alexopoulou, C., Xirouchaki, N., Vaporidi, K., and Georgopoulos, D. “Estimation of inspiratory muscle pressure in critically ill patients”. *Intensive Care Medicine* 36.4 (2010), pp. 648–655.
- [112] Koo-Poeggel, P., Neuwerk, S., Petersen, E., Grahoff, J., Mölle, M., Martinetz, T., and Marshall, L. “Closed-loop acoustic stimulation during an afternoon nap to modulate subsequent encoding”. *Journal of Sleep Research* 31.6 (2022).
- [113] Koopman, A. A., Blokpoel, R. G. T., Eykern, L. A., Jongh, F. H. C., Burgerhof, J. G. M., and Kneyber, M. C. J. “Transcutaneous electromyographic respiratory muscle recordings to quantify patient–ventilator interaction in mechanically ventilated children”. *Annals of Intensive Care* 8.1 (2018).
- [114] Kreit, J. W., Capper, M. W., and Eschenbacher, W. L. “Patient work of breathing during pressure support and volume-cycled mechanical ventilation.” *American Journal of Respiratory and Critical Care Medicine* 149.5 (1994), pp. 1085–1091.
- [115] Kusche, R., Graßhoff, J., Oltmann, A., Boudnik, L., and Rostalski, P. “A Robust Multi-Channel EMG System for Lower Back and Abdominal Muscles Training”. *Current Directions in Biomedical Engineering* 7.2 (2021), pp. 159–162.

- [116] Kusche, R., Graßhoff, J., Oltmann, A., and Rostalski, P. “A Multichannel EMG System for Spatial Measurement of Diaphragm Activities”. *IEEE Sensors Journal* 22.23 (2022), pp. 23393–23402.
- [117] Kusche, R., Oltmann, A., and Graßhoff, J. *Method for vectorial detection and processing of electromyographic signals*. DE102021126194A1 / EP4162878A1, 2021.
- [118] Kusche, R., Oltmann, A., Graßhoff, J., and Rostalski, P. “Comfortable Body Surface Potential Mapping by Means of a Dry Electrode Belt”. In: *Proceedings of the 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2022.
- [119] Langdon, R., Docherty, P. D., Chiew, Y.-S., Möller, K., and Chase, J. G. “Use of basis functions within a non-linear autoregressive model of pulmonary mechanics”. *Biomedical Signal Processing and Control* 27 (2016), pp. 44–50.
- [120] Lansing, R. and Savelle, J. “Chest surface recording of diaphragm potentials in man”. *Electroencephalography and Clinical Neurophysiology* 72.1 (1989), pp. 59–68.
- [121] Lauzon, A. M. and Bates, J. H. “Estimation of time-varying respiratory mechanical parameters by recursive least squares”. *Journal of Applied Physiology* 71.3 (1991), pp. 1159–1165.
- [122] Lee, C.-T., Chien, J.-Y., Hsu, M.-J., Wu, H.-D., and Wang, L.-Y. “Inspiratory muscle activation during inspiratory muscle training in patients with COPD”. *Respiratory Medicine* 190 (2021), p. 106676.
- [123] Leuteren, R. W., Bekhuis, R. E., Waal, C. G., Jongh, F. H., Kaam, A. H., and Hutten, G. J. “Diaphragmatic electromyography in preterm infants: The influence of electrode positioning”. *Pediatric Pulmonology* 55.2 (2019), pp. 354–359.
- [124] Liutkus, A., Badeau, R., and Richard, G. “Gaussian Processes for Underdetermined Source Separation”. *IEEE Transactions on Signal Processing* 59.7 (2011), pp. 3155–3167.
- [125] Loeliger, H.-A., Dauwels, J., Hu, J., Korl, S., Ping, L., and Kschischang, F. R. “The Factor Graph Approach to Model-Based Signal Processing”. *Proceedings of the IEEE* 95.6 (2007), pp. 1295–1322.
- [126] Lucangelo, U., Bernabé, F., and Blanch, L. “Respiratory mechanics derived from signals in the ventilator circuit”. *Respiratory Care* 50.1 (2005), pp. 55–67.
- [127] Luttinen, J. and Ilin, A. “Efficient Gaussian Process Inference for Short-Scale Spatio-Temporal Modeling”. In: *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 22. Proceedings of Machine Learning Research. PMLR, 2012, pp. 741–750.

-
- [128] Maarsingh, E. J. W., Eykern, L. A., Sprikkelman, A. B., Hoekstra, M. O., and Aalderen, W. M. C. “Respiratory muscle activity measured with a noninvasive EMG technique: technical aspects and reproducibility”. *Journal of Applied Physiology: Respiratory, Environmental and Exercise Physiology* 88.6 (2000), pp. 1955–1961.
- [129] MacKay, D. J. C. “Introduction to Gaussian Processes”. In: *Neural Networks and Machine Learning*. Springer Verlag, 1998.
- [130] Männel, G., Graßhoff, J., Rostalski, P., and Abbas, H. S. “Iterative Gaussian Process Model Predictive Control with Application to Physiological Control Systems”. In: *Proceedings of the 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021.
- [131] Matérn, B. “Stochastic models and their application to some problems in forest surveys and other sampling investigations”. *Spatial Variation* 49.5 (1960).
- [132] Matthay, M. A., Zemans, R. L., Zimmerman, G. A., Arabi, Y. M., Beitler, J. R., Mercat, A., Herridge, M., Randolph, A. G., and Calfee, C. S. “Acute respiratory distress syndrome”. *Nature Reviews Disease Primers* 5.1 (2019).
- [133] Mauri, T., Bellani, G., Grasselli, G., Confalonieri, A., Rona, R., Patroniti, N., and Pesenti, A. “Patient–ventilator interaction in ARDS patients with extremely low compliance undergoing ECMO: a novel approach based on diaphragm electrical activity”. *Intensive Care Medicine* 39.2 (2012), pp. 282–291.
- [134] Mauri, T., Yoshida, T., Bellani, G., Goligher, E. C., Carreaux, G., Rittayamai, N., Mojoli, F., Chiumello, D., Piquilloud, L., Grasso, S., Jubran, A., Laghi, F., Magder, S., Pesenti, A., Loring, S., Gattinoni, L., Talmor, D., Blanch, L., Amato, M., Chen, L., Brochard, L., Mancebo, J., and PLeUral pressure working Group (PLUG—Acute Respiratory Failure section of the European Society of Intensive Care Medicine), “Esophageal and transpulmonary pressure in the clinical setting: meaning, usefulness and perspectives”. *Intensive Care Medicine* 42.9 (2016), pp. 1360–1373.
- [135] McCully, K. K. and Faulkner, J. A. “Length-tension relationship of mammalian diaphragm muscles”. *Journal of Applied Physiology* 54.6 (1983), pp. 1681–1686.
- [136] Mecklenburgh, J. S. and Mapleson, W. W. “Ventilatory assistance and respiratory muscle activity. 2: Simulation with an adaptive active (“aa” or “a-squared”) model lung”. *British Journal of Anaesthesia* 80.4 (1998), pp. 434–439.
- [137] Menezes, D., Prata, D., Secchi, A., and Pinto, J. “A review on robust M-estimators for regression analysis”. *Computers & Chemical Engineering* 147 (2021), p. 107254.

- [138] Milic-Emili, J. and Zin, W. A. “Relationship Between Neuromuscular Respiratory Drive and Ventilatory Output”. *Comprehensive Physiology* (1986), pp. 631–646.
- [139] Miller, J. “Short Report: Reaction Time Analysis with Outlier Exclusion: Bias Varies with Sample Size”. *The Quarterly Journal of Experimental Psychology Section A* 43.4 (1991), pp. 907–912.
- [140] Moerer, O., Barwing, J., and Quintel, M. “„Neurally adjusted ventilatory assist“ (NAVA): Ein neuartiges assistiertes Beatmungsverfahren”. *Der Anaesthetist* 57.10 (2008), pp. 998–1005.
- [141] Mojoli, F., Iotti, G. A., Torriglia, F., Pozzi, M., Volta, C. A., Bianzina, S., Braschi, A., and Brochard, L. “In vivo calibration of esophageal pressure in the mechanically ventilated patient makes measurements reliable”. *Critical Care* 20.1 (2016).
- [142] Molkov, Y. I., Abdala, A. P. L., Bacak, B. J., Smith, J. C., Paton, J. F. R., and Rybak, I. A. “Late-Expiratory Activity: Emergence and Interactions With the Respiratory CPG”. *Journal of Neurophysiology* 104.5 (2010), pp. 2713–2729.
- [143] Molkov, Y. I., Bacak, B. J., Dick, T. E., and Rybak, I. A. “Control of breathing by interacting pontine and pulmonary feedback loops”. *Frontiers in Neural Circuits* 7 (2013).
- [144] Molkov, Y. I., Shevtsova, N. A., Park, C., Ben-Tal, A., Smith, J. C., Rubin, J. E., and Rybak, I. A. “A Closed-Loop Model of the Respiratory System: Focus on Hypercapnia and Active Expiration”. *PLoS ONE* 9.10 (2014), e109894.
- [145] Morton, S. E., Knopp, J. L., Chase, J. G., Docherty, P., Howe, S. L., Möller, K., Shaw, G. M., and Tawhai, M. “Optimising mechanical ventilation through model-based methods and automation”. *Annual Reviews in Control* 48 (2019), pp. 369–382.
- [146] Müller, M. “Dynamic time warping”. *Information Retrieval for Music and Motion* 2 (2007), pp. 69–84.
- [147] Natalini, G., Buizza, B., Granato, A., Anibaldi, E., Pisani, L., Ciabatti, G., Lippolis, V., Rosano, A., Latronico, N., Grasso, S., Antonelli, M., and Bernardini, A. “Non-invasive assessment of respiratory muscle activity during pressure support ventilation: accuracy of end-inspiration occlusion and least square fitting methods”. *Journal of Clinical Monitoring and Computing* 35.4 (2020), pp. 913–921.
- [148] Nawaz, A., Herzog né Hoffmann, C., Graßhoff, J., Pfeiffer, S., Lichtenberg, G., and Rostalski, P. “Probabilistic model-based fault diagnosis for the cavities of the European XFEL”. *at – Automatisierungstechnik* 69.6 (2021), pp. 538–549.

-
- [149] Nickisch, H., Solin, A., and Grigorevskiy, A. “State Space Gaussian Processes with Non-Gaussian Likelihood”. In: *Proceedings of the 35th International Conference on Machine Learning (ICML)*. 2018.
- [150] Niknazar, M., Rivet, B., and Jutten, C. “Fetal ECG extraction from a single sensor by a non-parametric modeling”. In: *2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*. 2012, pp. 949–953.
- [151] Niknazar, M., Rivet, B., and Jutten, C. “Fetal ECG Extraction by Extended State Kalman Filtering Based on Single-Channel Recordings”. *IEEE Transactions on Biomedical Engineering* 60.5 (2013), pp. 1345–1352.
- [152] Noorzadeh, S., Niknazar, M., Rivet, B., Fontecave-Jallon, J., Gumery, P.-Y., and Jutten, C. “Modeling quasi-periodic signals by a non-parametric model: Application on fetal ECG extraction”. In: *Proceedings of the 36th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2014.
- [153] Otis, A. B., Fenn, W. O., and Rahn, H. “Mechanics of Breathing in Man”. *Journal of Applied Physiology* 2.11 (1950), pp. 592–607.
- [154] Papazian, L., Forel, J.-M., Gacouin, A., Penot-Ragon, C., Perrin, G., Loundou, A., Jaber, S., Arnal, J.-M., Perez, D., Seghboyan, J.-M., Constantin, J.-M., Courant, P., Lefrant, J.-Y., Guérin, C., Prat, G., Morange, S., and Roch, A. “Neuromuscular Blockers in Early Acute Respiratory Distress Syndrome”. *New England Journal of Medicine* 363.12 (2010), pp. 1107–1116.
- [155] Parker, P. A. and Merletti, R. *Electromyography. Physiology, engineering, and noninvasive applications*. IEEE Press series in Biomedical Engineering. IEEE Press, 2004. 1494 pp.
- [156] Pelosi, P., Rocco, P. R., and Abreu, M. G. “Use of computed tomography scanning to guide lung recruitment and adjust positive-end expiratory pressure”. *Current Opinion in Critical Care* 17.3 (2011), pp. 268–274.
- [157] Pepin, J.-L., Le-Dong, N.-N., Cuthbert, V., Coumans, N., Tamisier, R., Malhotra, A., and Martinot, J.-B. “Mandibular Movements are a Reliable Noninvasive Alternative to Esophageal Pressure for Measuring Respiratory Effort in Patients with Sleep Apnea Syndrome”. *Nature and Science of Sleep* 14 (2022), pp. 635–644.
- [158] Peslin, R., Silva, J. F., Chabot, F., and Duvivier, C. “Respiratory mechanics studied by multiple linear regression in unsedated ventilated patients”. *European respiratory journal* 5.7 (1992), pp. 871–878.
- [159] Petersen, E. “Model-based Probabilistic Inference for Monitoring Respiratory Effort using the Surface Electromyogram”. PhD thesis. University of Lübeck, 2021.

- [160] Petersen, E., Buchner, H., Eger, M., and Rostalski, P. “Convolutive blind source separation of surface EMG measurements of the respiratory muscles”. *Biomedical Engineering / Biomedizinische Technik* 62.2 (2017), pp. 171–181.
- [161] Petersen, E., Graßhoff, J., Eger, M., and Rostalski, P. “Surface EMG-based Estimation of Breathing Effort for Neurally Adjusted Ventilation Control”. In: *Proceedings of the 20th IFAC World Congress*. Vol. 53. 2. Elsevier BV, 2020, pp. 16323–16328.
- [162] Petersen, E., Sauer, J., Graßhoff, J., and Rostalski, P. “Removing Cardiac Artifacts From Single-Channel Respiratory Electromyograms”. *IEEE Access* 8 (2020), pp. 30905–30917.
- [163] Petit, J. M., Milic-Emili, G., and Delhez, L. “Role of the diaphragm in breathing in conscious normal man: an electromyographic study”. *Journal of Applied Physiology* 15.6 (1960), pp. 1101–1106.
- [164] Pikkemaat, R. and Leonhardt, S. “Separation of ventilation and perfusion related signals within EIT-data streams”. *Journal of Physics: Conference Series* 224 (2010), p. 012028.
- [165] Prüßmann, J., Graßhoff, J., and Rostalski, P. “Exploitation of Kronecker Structure in Gaussian Process Regression for Efficient Biomedical Signal Processing”. *Current Directions in Biomedical Engineering* 7.2 (2021), pp. 287–290.
- [166] Putensen, C., Zech, S., Wrigge, H., Zinserling, J., Stüber, F., Spiegel, T., and Mutz, N. “Long-Term Effects of Spontaneous Breathing During Ventilatory Support in Patients with Acute Lung Injury”. *American Journal of Respiratory and Critical Care Medicine* 164.1 (2001), pp. 43–49.
- [167] Quinero Candela, J. and Rasmussen, C. E. “A Unifying View of Sparse Approximate Gaussian Process Regression”. *Journal of Machine Learning Research (JMLR)* 6 (2005), pp. 1935–1959.
- [168] Ramsay, M., Mandal, S., Suh, E.-S., Steier, J., Douiri, A., Murphy, P. B., Polkey, M., Simonds, A., and Hart, N. “Parasternal electromyography to determine the relationship between patient-ventilator asynchrony and nocturnal gas exchange during home mechanical ventilation set-up”. *Thorax* 70.10 (2015), pp. 946–952.
- [169] Ramsook, A. H., Koo, R., Molgat-Seon, Y., Dominelli, P. B., Syed, N., Ryerson, C. J., Sheel, A. W., and Guenette, J. A. “Diaphragm Recruitment Increases during a Bout of Targeted Inspiratory Muscle Training”. *Medicine & Science in Sports & Exercise* 48.6 (2016), pp. 1179–1186.
- [170] Rasmussen, C. E. and Williams, C. K. I. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2008. 248 pp.

-
- [171] Rassier, D. E., MacIntosh, B. R., and Herzog, W. “Length dependence of active force production in skeletal muscle”. *Journal of Applied Physiology* 86.5 (1999), pp. 1445–1457.
- [172] Rauch, H. E., Tung, F., and Striebel, C. T. “Maximum likelihood estimates of linear dynamic systems”. *AIAA Journal* 3.8 (1965), pp. 1445–1450.
- [173] Al-Rawas, N., Banner, M., Euliano, N., Tams, C., Brown, J., Martin, A., and Gabrielli, A. “Expiratory time constant for determinations of plateau pressure, respiratory system compliance, and total resistance”. *Critical Care* 17 (2013), R23.
- [174] Redmond, D. P., Docherty, P. D., Chiew, Y. S., and Chase, J. G. “A polynomial model of patient-specific breathing effort during controlled mechanical ventilation”. In: *Proceedings of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 2015, pp. 4532–4535.
- [175] Reece, S., Ghosh, S., Rogers, A., Roberts, S., and Jennings, N. R. “Efficient State-Space Inference of Periodic Latent Force Models”. *Journal of Machine Learning Research* 15.1 (2014), pp. 2337–2397.
- [176] Reinders, J., Hunnekens, B., Wouw, N., and Oomen, T. “Noninvasive Breathing Effort Estimation of Mechanically Ventilated Patients Using Sparse Optimization”. *IEEE Open Journal of Control Systems* 1 (2022), pp. 57–68.
- [177] Roberts, T. J. and Gabaldon, A. M. “Interpreting muscle function from EMG: lessons learned from direct measurements of muscle force”. *Integrative and Comparative Biology* 48.2 (2008), pp. 312–320.
- [178] Rohrer, F. “Der Strömungswiderstand in den menschlichen Atemwegen und der Einfluss der unregelmässigen Verzweigung des Bronchialsystems auf den Atmungsverlauf in verschiedenen Lungenbezirken”. *Pflüger’s Archiv für die Gesamte Physiologie des Menschen und der Thiere* 162.5–6 (1915), pp. 225–299.
- [179] Rohrer, F. “Der Zusammenhang der Atemkräfte und ihre Abhängigkeit vom Dehnungszustand der Atmungsorgane”. *Pflüger’s Archiv für die Gesamte Physiologie des Menschen und der Tiere* 165.8–10 (1916), pp. 419–444.
- [180] Ruiz Ferrón, F. and Serrano Simón, J. “La monitorización convencional no es suficiente para valorar el esfuerzo respiratorio durante la ventilación asistida”. *Medicina Intensiva* 43.4 (2019), pp. 197–206.
- [181] Saad, Y. *Numerical methods for large eigenvalue problems*. Revised Edition. Society for Industrial and Applied Mathematics (SIAM), 2011.

- [182] Saatçi, E. and Akan, A. “Lung Model Parameter Estimation by Unscented Kalman Filter”. In: *Proceedings of the 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2007, pp. 2556–2559.
- [183] Saatçi, Y. “Scalable Inference for Structured Gaussian Process Models”. PhD thesis. University of Cambridge, 2011.
- [184] Sameni, “A Review of Fetal ECG Signal Processing Issues and Promising Directions”. *The Open Pacing, Electrophysiology & Therapy Journal* (2010).
- [185] Särkkä, S. “Recursive Bayesian Inference on Stochastic Differential Equations”. PhD thesis. Helsinki University of Technology, 2006.
- [186] Särkkä, S. “Linear Operators and Stochastic Partial Differential Equations in Gaussian Process Regression”. In: *Artificial Neural Networks and Machine Learning – ICANN 2011*. Springer Berlin Heidelberg, 2011, pp. 151–158.
- [187] Särkkä, S. *Bayesian filtering and smoothing*. Cambridge University Press, 2013.
- [188] Särkkä, S., Alvarez, M. A., and Lawrence, N. D. “Gaussian Process Latent Force Models for Learning and Stochastic Control of Physical Systems”. *IEEE Transactions on Automatic Control* 64.7 (2019), pp. 2953–2960.
- [189] Särkkä, S. and Hartikainen, J. “Infinite-Dimensional Kalman Filtering Approach to Spatio-Temporal Gaussian Process Regression”. In: *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, 2012, pp. 993–1001.
- [190] Särkkä, S. and Piche, R. “On convergence and accuracy of state-space approximations of squared exponential covariance functions”. In: *2014 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2014.
- [191] Särkkä, S. and Solin, A. *Applied Stochastic Differential Equations*. Cambridge University Press, 2019.
- [192] Särkkä, S., Solin, A., and Hartikainen, J. “Spatiotemporal Learning via Infinite-Dimensional Bayesian Filtering and Smoothing: A Look at Gaussian Process Regression Through Kalman Filtering”. *IEEE Signal Processing Magazine* 30.4 (2013), pp. 51–61.
- [193] Särkkä, S., Solin, A., Nummenmaa, A., Vehtari, A., Auranen, T., Vanni, S., and Lin, F.-H. “Dynamic retrospective filtering of physiological noise in BOLD fMRI: DRIFTER”. *NeuroImage* 60.2 (2012), pp. 1517–1527.
- [194] Sauer, J., Graßhoff, J., Carbon, N. M., Koch, W. M., Weber-Carstens, S., and Rostalski, P. “Automated characterization of patient–ventilator interaction using surface electromyography”. *Annals of Intensive Care* 14.1 (2024).

-
- [195] Sauer, J., Streppel, M., Carbon, N. M., Petersen, E., and Rostalski, P. “Blind source separation of inspiration and expiration in respiratory sEMG signals”. *Physiological Measurement* 43.7 (2022), p. 075007.
- [196] Scheel, M., Berndt, A., and Simanski, O. “Application of Kalman filter for breathing effort reconstruction for OSAS patients in breathing therapy”. *at – Automatisierungstechnik* 66.12 (2018), pp. 1064–1071.
- [197] Schepens, T., Dres, M., Heunks, L., and Goligher, E. C. “Diaphragm-protective mechanical ventilation”. *Current Opinion in Critical Care* 25.1 (2019), pp. 77–85.
- [198] Schöberl, J. “NETGEN An advancing front 2D/3D-mesh generator based on abstract rules”. *Computing and Visualization in Science* 1.1 (1997), pp. 41–52.
- [199] Schölkopf, B. *Learning with Kernels. Support vector machines, regularization, optimization, and beyond*. Adaptive computation and machine learning. MIT Press, 2002. 1626 pp.
- [200] Schott, R. H. A., Blom, J. A., and Korsten, H. H. M. “Real-time Computation of a Patient’s Respiratory Effort During Ventilation”. *Journal of Clinical Monitoring and Computing* 20.3 (2006), pp. 193–200.
- [201] Schuessler, T. F., Gottfried, S. B., Goldberg, P., Kearney, R. E., and Bates, J. H. T. “An Adaptive Filter to Reduce Cardiogenic Oscillations on Esophageal Pressure Signals”. *Annals of Biomedical Engineering* 26.2 (1998), pp. 260–267.
- [202] Schwartzstein, R. M. *Respiratory physiology. A clinical approach*. The integrated physiology series. Lippincott Williams & Wilkins, 2006. 232 pp.
- [203] Siebert, M., Graßhoff, J., and Rostalski, P. “Uncertainty Analysis of Deep Kernel Learning Methods on Diabetic Retinopathy Grading”. *IEEE Access* 11 (2023), pp. 146173–146184.
- [204] Sieck, G. C., Ferreira, L. F., Reid, M. B., and Mantilla, C. B. “Mechanical Properties of Respiratory Muscles”. *Comprehensive Physiology* (2013), pp. 1533–1567.
- [205] Sinderby, C., Navalesi, P., Beck, J., Skrobik, Y., Comtois, N., Friberg, S., Gottfried, S. B., and Lindström, L. “Neural control of mechanical ventilation in respiratory failure”. *Nature Medicine* 5.12 (1999), pp. 1433–1436.
- [206] Sinderby, C., Liu, S., Colombo, D., Camarotta, G., Slutsky, A. S., Navalesi, P., and Beck, J. “An automated and standardized neural index to quantify patient-ventilator interaction”. *Critical Care* 17.5 (2013), R239.
- [207] Sinderby, C. A., Beck, J. C., Lindström, L. H., and Grassino, A. E. “Enhancement of signal quality in esophageal recordings of diaphragm EMG”. *Journal of Applied Physiology: Respiratory, Environmental and Exercise Physiology* 82.4 (1997), pp. 1370–1377.

- [208] Smith, J. and Bellemare, F. “Effect of lung volume on in vivo contraction characteristics of human diaphragm”. *Journal of Applied Physiology* 62.5 (1987), pp. 1893–1900.
- [209] Smith, J. C., Abdala, A. P. L., Koizumi, H., Rybak, I. A., and Paton, J. F. R. “Spatial and Functional Architecture of the Mammalian Brain Stem Respiratory Network: A Hierarchy of Three Oscillatory Mechanisms”. *Journal of Neurophysiology* 98.6 (2007), pp. 3370–3387.
- [210] Solin, A. “Stochastic Differential Equation Methods for Spatio-Temporal Gaussian Process Regression”. PhD thesis. Aalto University School of Science, 2016.
- [211] Solin, A., Hensman, J., and Turner, R. E. “Infinite-Horizon Gaussian Processes”. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS)*. 2018. arXiv: 1811.06588.
- [212] Solin, A. and Särkkä, S. “Explicit Link Between Periodic Covariance Functions and State Space Models”. In: *17th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2014.
- [213] Somersalo, E., Cheney, M., and Isaacson, D. “Existence and Uniqueness for Electrode Models for Electric Current Computed Tomography”. *SIAM Journal on Applied Mathematics* 52.4 (1992), pp. 1023–1040.
- [214] Spieth, P. M., Güldner, A., Huhle, R., Beda, A., Bluth, T., Schreiter, D., Ragaller, M., Gottschlich, B., Kiss, T., Jaber, S., Pelosi, P., Koch, T., and Gama de Abreu, M. “Short-term effects of noisy pressure support ventilation in patients with acute hypoxemic respiratory failure”. *Critical Care* 17.5 (2013), R261.
- [215] Stegmaier, P. A., Zollinger, A., Brunner, J. X., and Pasch, T. “Assessment of Pulmonary Mechanics in Mechanical Ventilation: effects of Imprecise Breath Detection, Phase Shift and Noise”. *The Journal of Clinical Monitoring and Computing* 14.2 (1998), pp. 127–134.
- [216] Stein, E., Chen, R., Battistel, A., and Moeller, K. “Separating Respiration and Perfusion in EIT: Harmonic Analysis on 2D-Thorax Simulation”. *Current Directions in Biomedical Engineering* 8.2 (2022), pp. 785–788.
- [217] Stein, M. L. *Interpolation of Spatial Data – Some Theory for Kriging*. Springer Series in Statistics. Springer Verlag, 1999.
- [218] Stender, M., Graßhoff, J., Braun, T., Möller, R., and Rostalski, P. “A Hybrid Factor Graph Model for Biomedical Activity Detection”. In: *2021 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)*. IEEE, 2021.
- [219] Stigler, S. M. “Gauss and the Invention of Least Squares”. *The Annals of Statistics* 9.3 (1981), pp. 465–474.

-
- [220] Straus, C., Arnulf, I., Similowsky, T., and Derenne, J.-P. “Control of breathing: neural drive”. In: *Applied Physiology in Respiratory Mechanics*. Springer Milan, 1998, pp. 1–19.
- [221] Terson de Paleville, D. G. L., McKay, W. B., Folz, R. J., and Ovechkin, A. V. “Respiratory Motor Control Disrupted by Spinal Cord Injury: Mechanisms, Evaluation, and Restoration”. *Translational Stroke Research* 2.4 (2011), pp. 463–473.
- [222] Tokizane, T., Kawamata, K., and Tokizane, H. “Electromyographic Studies on the Human Respiratory Muscles”. *The Japanese Journal of Physiology* 2 (1951), pp. 232–247.
- [223] Tukey, J. W. “A survey of sampling from contaminated distributions”. In: *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*. 1960, pp. 448–485.
- [224] Uhl, R. R. and Lewis, F. “Digital computer calculation of human pulmonary mechanics using a least squares fit technique”. *Computers and Biomedical Research* 7.5 (1974), pp. 489–495.
- [225] Uhlenbeck, G. E. and Ornstein, L. S. “On the Theory of the Brownian Motion”. *Physical Review* 36.5 (1930), pp. 823–841.
- [226] van Leuteren, R., Hutten, G., de Waal, C., Dixon, P., van Kaam, A., and de Jongh, F. “Processing transcutaneous electromyography measurements of respiratory muscles, a review of analysis techniques”. *Journal of Electromyography and Kinesiology* 48 (2019), pp. 176–186.
- [227] Varin, S. and Panagiotakos, D. B. “A review of robust regression in biomedical science research”. *Archives of Medical Science* 16.5 (2020), pp. 1267–1269.
- [228] Vicario, F., Albanese, A., Karamolegkos, N., Wang, D., Seiver, A., and Chbat, N. W. “Noninvasive Estimation of Respiratory Mechanics in Spontaneously Breathing Ventilated Patients: A Constrained Optimization Approach”. *IEEE Transactions on Biomedical Engineering* 63.4 (2016), pp. 775–787.
- [229] Vicario, F., Buizza, R., Truschel, W. A., and Chbat, N. W. “Noninvasive estimation of alveolar pressure”. In: *Proceedings of the 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 2016, pp. 2721–2724.
- [230] Vincent, J.-L., Boulanger, C., Mol, M. M. C., Hawryluck, L., and Azoulay, E. “Ten areas for ICU clinicians to be aware of to help retain nurses in the ICU”. *Critical Care* 26.1 (2022).
- [231] Vollmer, F., Graßhoff, J., and Rostalski, P. “Probabilistic Ultra-Wideband TDoA Localization with Bias Correction”. In: *2022 Proceedings of the 30th European Signal Processing Conference (EUSIPCO)*. IEEE, 2022.

- [232] Volta, C., Marangoni, E., Alvisi, V., Capuzzo, M., Ragazzi, R., Pavanelli, L., and Alvisi, R. “Respiratory mechanics by least squares fitting in mechanically ventilated patients: application on flow-limited COPD patients”. *Intensive Care Medicine* 28 (2001), pp. 48–52.
- [233] Wald, A., Jason, D., Murphy, T. W., and Mazzia, V. D. “A computers system for respiratory parameters”. *Computers and Biomedical Research* 2.5 (1969), pp. 411–429.
- [234] Walterspacher, S., Isaak, L., Guttman, J., Kabitz, H.-J., and Schumann, S. “Assessing Respiratory Function Depends on Mechanical Characteristics of Balloon Catheters”. *Respiratory Care* 59.9 (2014), pp. 1345–1352.
- [235] West, J. B. *Respiratory physiology. The essentials*. 8th edition. Lippincott Williams & Wilkins, 2009. 186 pp.
- [236] Wilkinson, W. J., Chang, P. E., Andersen, M. R., and Solin, A. “State Space Expectation Propagation: Efficient Inference Schemes for Temporal Gaussian Processes”. In: *Proceedings of the 37th International Conference on Machine-Learning (ICML)*. 2020.
- [237] Wilson, A. G., Dann, C., and Nickisch, H. “Thoughts on Massively Scalable Gaussian Processes”. *arXiv preprint arXiv:1511.01870* (2015).
- [238] Wilson, A. G., Gilboa, E., Nehorai, A., and Cunningham, J. P. “Fast Kernel Learning for Multidimensional Pattern Extrapolation”. In: *Advances in Neural Information Processing Systems 27 (NIPS)*. 2014, pp. 3626–3634.
- [239] Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. “Deep kernel learning”. In: *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2016, pp. 370–378.
- [240] Wilson, A. G. and Nickisch, H. “Kernel Interpolation for Scalable Structured Gaussian Processes (KISS-GP)”. In: *Proceedings of the 32nd International Conference on Machine Learning (ICML)*. 2015, pp. 1775–1784.
- [241] Yamada, Y., Shigeta, M., Suwa, K., and Hanaoka, K. “Respiratory muscle pressure analysis in pressure-support ventilation”. *Journal of Applied Physiology* 77.5 (1994), pp. 2237–2243.
- [242] Yamada, Y. and Du, H.-L. “Analysis of the mechanisms of expiratory asynchrony in pressure support ventilation: a mathematical approach”. *Journal of Applied Physiology* 88.6 (2000), pp. 2143–2150.
- [243] Yoshida, T., Amato, M. B. P., Grieco, D. L., Chen, L., Lima, C. A. S., Roldan, R., Morais, C. C. A., Gomes, S., Costa, E. L. V., Cardoso, P. F. G., Charbonney, E., Richard, J.-C. M., Brochard, L., and Kavanagh, B. P. “Esophageal Manometry and Regional Transpulmonary Pressure in Lung Injury”. *American Journal of Respiratory and Critical Care Medicine* 197.8 (2018), pp. 1018–1026.

-
- [244] Yoshida, T., Fujino, Y., Amato, M. B. P., and Kavanagh, B. P. “Fifty Years of Research in ARDS. Spontaneous Breathing during Mechanical Ventilation. Risks, Mechanisms, and Management”. *American Journal of Respiratory and Critical Care Medicine* 195.8 (2017), pp. 985–992.
- [245] Yoshida, T., Uchiyama, A., Matsuura, N., Mashimo, T., and Fujino, Y. “Spontaneous breathing during lung-protective ventilation in an experimental acute lung injury model: High transpulmonary pressure associated with strong spontaneous breathing effort may worsen lung injury”. *Critical Care Medicine* 40.5 (2012), pp. 1578–1585.
- [246] Younes, M. and Riddle, W. “A model for the relation between respiratory neural and mechanical outputs. I. Theory”. *Journal of Applied Physiology* 51.4 (1981), pp. 963–978.
- [247] Younes, M., Brochard, L., Grasso, S., Kun, J., Mancebo, J., Ranieri, M., Richard, J.-C., and Younes, H. “A method for monitoring and improving patient: ventilator interaction”. *Intensive Care Medicine* 33.8 (2007), pp. 1337–1346.
- [248] Younes, M., Kun, J., Masiowski, B., Webster, K., and Roberts, D. “A Method for Noninvasive Determination of Inspiratory Resistance during Proportional Assist Ventilation”. *American Journal of Respiratory and Critical Care Medicine* 163.4 (2001), pp. 829–839.
- [249] Younes, M., Webster, K., Kun, J., Roberts, D., and Masiowski, B. “A Method for Measuring Passive Elastance during Proportional Assist Ventilation”. *American Journal of Respiratory and Critical Care Medicine* 164.1 (2001), pp. 50–60.
- [250] Zhang, D., Lu, G., Zhu, X., Zhang, L., Gao, J., Shi, L., Gu, J.-H., and Liu, J.-n. “Neural Respiratory Drive Measured Using Surface Electromyography of Diaphragm as a Physiological Biomarker to Predict Hospitalization of Acute Exacerbation of Chronic Obstructive Pulmonary Disease Patients”. *Chinese Medical Journal* 131.23 (2018), pp. 2800–2807.
- [251] Zhang, Z. “Parameter Estimation Techniques: A Tutorial with Application to Conic Fitting”. *Image and Vision Computing* 15.1 (1997), pp. 59–76.
- [252] Zin, W. A., Pengelly, L. D., and Milic-Emili, J. “Single-breath method for measurement of respiratory mechanics in anesthetized animals”. *Journal of Applied Physiology* 52.5 (1982), pp. 1266–1271.

List of Figures

1.1	Graphical overview of the thesis' content and methodological chapters.	4
2.1	An illustration of likelihood and loss functions	11
2.2	An illustrative example of Gaussian process regression	14
2.3	Marginal log likelihood surface on a toy dataset	15
2.4	An illustration of commonly used kernel functions	17
2.5	A comparison of GP batch regression, Kalman filtering and smoothing	22
2.6	An illustration of the human chest including the lungs, rib cage and main respiratory muscles	24
2.7	The single-compartment model of lung mechanics	28
2.8	An illustrative simulation of pneumatic data and the Campbell diagram during assisted breathing	30
3.1	An overview of esophageal/gastric pressure and sEMG signal processing pipelines	41
3.2	An illustration of different metrics used for quantifying inspiratory effort	43
3.3	Airway pressure and sEMG curves during an airway occlusion maneuver	45
3.4	An illustrative excerpt of time-series data recorded in the clinical study	46
3.5	Distribution of pressure–time products, EMG–time products, minute ventilation and intrinsic PEEP	47
3.6	Distribution of neuromechanical coupling indices K_{EMG} and biases P_{bias}	48
3.7	The sEMG- P_{mus} relation for three selected patients	49
3.8	Scatter plot of neuromechanical coupling indices determined invasively and non-invasively via occlusions	50
4.1	Overview of different methods for estimating respiratory effort	59
4.2	Key steps involved in the robust P_{mus} estimation procedure	62
4.3	A simulation of V_{rs} during assisted and spontaneous breathing	63
4.4	Measurements and model-based P_{mus} estimates from one patient across all PS levels	68
4.5	Relation between estimated $PTP_{mus,EMG}$ and invasively measured PTP_{mus} values	70
4.6	Relation between estimated $PTP_{mus,EMG}/min$ and invasively measured PTP_{mus}/min values	71

List of Figures

4.7	Scatter plot of neuromechanical coupling indices determined invasively and non-invasively via the model-based estimator	73
4.8	Excerpts of measurements and P_{mus} estimates from two patients	75
5.1	Examples of parametric P_{mus} functions	81
5.2	A factor graph of the joint distribution of the probabilistic sensor fusion model	87
5.3	A comparison of exact kernel functions and their state-space approximations	89
5.4	An overview of P_{mus} GP priors after optimization of hyperparameters	98
5.5	Excerpt of GP priors/posteriors and covariance matrices from one patient	99
5.6	Excerpt of measurements and P_{mus} posteriors from one patient	101
5.7	Distribution of performance metrics for all considered GP models	103
6.1	An illustration of the SKI approximation for a quasi-periodic kernel with a non-stationary phase	115
6.2	An illustration of the fetal ECG learning and separation problem	120
6.3	The dual EIT reconstruction model	126
6.4	The FEM model used for generating artificial EIT data	127
6.5	An illustration of reconstructed EIT images and their temporal evolution at selected pixels	128
6.6	An excerpt of the prior kernel matrix $K_{\phi, \text{vent}}$	129

List of Tables

3.1	Pearson correlation between effort metrics	48
4.1	Performance metrics for baseline ordinary least squares estimators and different combinations of robust regression methods	69
4.2	Performance metrics for occlusion- and model-based methods	72
4.3	Parameter estimates of model-based methods	73
5.1	Overview of mean and kernel functions used to represent P_{mus}	85
6.1	Comparison of different experiments conducted on fetal ECG data	120
A.1	Clinical characteristics of included patients	139

List of Publications

Journal articles as first author

- Graßhoff, J., Petersen, E., Farquharson, F., Kustermann, M., Kabitz, H.-J., Rostalski, P., and Walterpacher, S. “Surface EMG-based Quantification of Inspiratory Effort: A Quantitative Comparison with Pes”. *Critical Care* 25.441 (2021).
- Graßhoff, J., Petersen, E., Walterspacher, S., and Rostalski, P. “Model-based Estimation of Inspiratory Effort using Surface EMG”. *IEEE Transactions on Biomedical Engineering* 70.1 (2023), pp. 247–258.

Conference papers as first author

- Graßhoff, J., Petersen, E., Eger, M., Bellani, G., and Rostalski, P. “A Template Subtraction Method for the Removal of Cardiogenic Oscillations on Esophageal Pressure Signals”. In: *Proceedings of the 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2017, pp. 2235–2238.
- Graßhoff, J., Männel, G., and Rostalski, P. “Nonparametric modeling of quasi-periodic signals – application to esophageal pressure filtering”. In: *Annual Conference of the German Society for Biomedical Engineering (BMT)*. DGBMT, 2018.
- Graßhoff, J., Männel, G., Abbas, H. S., and Rostalski, P. “Model Predictive Control using Efficient Gaussian Processes for Unknown Disturbance Inputs”. In: *Proceedings of the 58th IEEE Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 2708–2713.
- Graßhoff, J., Petersen, E., Becher, T., and Rostalski, P. “Automatic Estimation of Respiratory Effort using Esophageal Pressure”. In: *Proceedings of the 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2019, pp. 4646–4649.
- Graßhoff, J., Jankowski, A., and Rostalski, P. “Scalable Gaussian Process Separation for Kernels with a Non-Stationary Phase”. In: *Proceedings of the 37th International Conference on Machine Learning (ICML)*. Vol. 119. PMLR, 2020, pp. 3722–3731.

-
- Graßhoff, J. and Rostalski, P. “Spatio-Temporal Gaussian Processes for Separation of Ventilation and Perfusion Related Signals in EIT Data”. In: *Proceedings on Automation in Medical Engineering*. Vol. 1. 1. 2020.

Journal articles as co-author

- Petersen, E., Sauer, J., Graßhoff, J., and Rostalski, P. “Removing Cardiac Artifacts From Single-Channel Respiratory Electromyograms”. *IEEE Access* 8 (2020), pp. 30905–30917.
- Nawaz, A., Herzog né Hoffmann, C., Graßhoff, J., Pfeiffer, S., Lichtenberg, G., and Rostalski, P. “Probabilistic model-based fault diagnosis for the cavities of the European XFEL”. *at - Automatisierungstechnik* 69.6 (2021), pp. 538–549.
- Koo-Poeggel, P., Neuwerk, S., Petersen, E., Graßhoff, J., Mölle, M., Martinetz, T., and Marshall, L. “Closed-loop acoustic stimulation during an afternoon nap to modulate subsequent encoding”. *Journal of Sleep Research* 31.6 (2022).
- Kusche, R., Graßhoff, J., Oltmann, A., and Rostalski, P. “A Multichannel EMG System for Spatial Measurement of Diaphragm Activities”. *IEEE Sensors Journal* 22.23 (2022), pp. 23393–23402.
- Siebert, M., Graßhoff, J., and Rostalski, P. “Uncertainty Analysis of Deep Kernel Learning Methods on Diabetic Retinopathy Grading”. *IEEE Access* 11 (2023), pp. 146173–146184.
- Jonkman, A. H., Warnaar, R. S. P., Baccinelli, W., Carbon, N. M., D’Cruz, R. F., Doorduyn, J., Doorn, J. L. M., Elshof, J., Estrada-Petrocelli, L., Graßhoff, J., Heunks, L. M. A., Koopman, A. A., Langer, D., Moore, C. M., Nunez Silveira, J. M., Petersen, E., Poddighe, D., Ramsay, M., Rodrigues, A., Roesthuis, L. H., Rossel, A., Torres, A., Duiverman, M. L., and Oppersma, E. “Analysis and applications of respiratory surface EMG: report of a round table meeting”. *Critical Care* 28.1 (2024).
- Sauer, J., Graßhoff, J., Carbon, N. M., Koch, W. M., Weber-Carstens, S., and Rostalski, P. “Automated characterization of patient–ventilator interaction using surface electromyography”. *Annals of Intensive Care* 14.1 (2024).
- Hennigs, C., Bilda, F., Graßhoff, J., Waltersbacher, S., and Rostalski, P. “A switching lung mechanics model for detection of expiratory flow limitation”. *at - Automatisierungstechnik* (in press, 2024).

Conference papers as co-author

- Bockelmann, N., Graßhoff, J., Hansen, L., Bellani, G., Heinrich, M. P., and Rostalski, P. “Deep Learning for Prediction of Diaphragm Activity from the Surface Electromyogram”. *Current Directions in Biomedical Engineering* 5.1 (2019), pp. 17–20.
- Abbas, H., Graßhoff, J., Rostalski, P., and Brinkmann, R. “On the estimation of optoacoustic waves in retinal laser therapy using Gaussian processes”. In: *Proceedings on Automation in Medical Engineering*. Vol. 1. 1. 2020.
- Petersen, E., Graßhoff, J., Eger, M., and Rostalski, P. “Surface EMG-based Estimation of Breathing Effort for Neurally Adjusted Ventilation Control”. In: *Proceedings of the 20th IFAC World Congress*. Vol. 53. 2. Elsevier BV, 2020, pp. 16323–16328.
- Kusche, R., Graßhoff, J., Oltmann, A., Boudnik, L., and Rostalski, P. “A Robust Multi-Channel EMG System for Lower Back and Abdominal Muscles Training”. *Current Directions in Biomedical Engineering* 7.2 (2021), pp. 159–162.
- Männel, G., Graßhoff, J., Rostalski, P., and Abbas, H. S. “Iterative Gaussian Process Model Predictive Control with Application to Physiological Control Systems”. In: *Proceedings of the 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021.
- Prüßmann, J., Graßhoff, J., and Rostalski, P. “Exploitation of Kronecker Structure in Gaussian Process Regression for Efficient Biomedical Signal Processing”. *Current Directions in Biomedical Engineering* 7.2 (2021), pp. 287–290.
- Stender, M., Graßhoff, J., Braun, T., Möller, R., and Rostalski, P. “A Hybrid Factor Graph Model for Biomedical Activity Detection”. In: *2021 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)*. IEEE, 2021.
- Boudnik, L., Graßhoff, J., Vollmer, F., and Rostalski, P. “Muscle Artifact Removal in Single-Channel Electrocardiograms using Temporal Convolutional Networks”. In: *44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2022.
- Brügge, N. S., Graßhoff, J., Weigenand, A., and Rostalski, P. “Multi-Task Gaussian Process Regression for the Detection of Sleep Cycles in Premature Infants”. In: *ICASSP 2022 – 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022.
- Gruner, J., Schmid, N., Männel, G., Graßhoff, J., Abbas, H. S., and Rostalski, P. “Recursively Feasible Model Predictive Control using Latent Force Models Applied to Disturbed Quadcopters”. In: *Proceedings of the IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022.

-
- Kusche, R., Oltmann, A., Graßhoff, J., and Rostalski, P. “Comfortable Body Surface Potential Mapping by Means of a Dry Electrode Belt”. In: *Proceedings of the 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2022.
 - Vollmer, F., Graßhoff, J., and Rostalski, P. “Probabilistic Ultra-Wideband TDoA Localization with Bias Correction”. In: *2022 Proceedings of the 30th European Signal Processing Conference (EUSIPCO)*. IEEE, 2022.
 - Hennigs, C., Scholleman, F., Graßhoff, J., Hardel, T., and Rostalski, P. “A simple mathematical lung model of Chronic Obstructive Pulmonary Disease (COPD)”. In: *Proceedings on Automation in Medical Engineering*. Vol. 2. 1. 2023.