

Data-driven Regression Models for Voyage Cost Optimisation Based on the Operating Conditions of the SA Agulhas II

by
Petrus Gerhardus Durandt

*Thesis presented in partial fulfilment of the requirements for the degree
of Master of Engineering (Mechatronic) in the Faculty of Engineering at
Stellenbosch University*



Supervisor: Prof. A. Bekker

December 2020

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Date: 2020/11/14

Copyright © 2020 Stellenbosch University
All rights reserved.

Abstract

Data-driven Regression Models for Voyage Cost Optimisation Based on the Operating Conditions of the SA Agulhas II

P.G. Durandt

*Department of Mechanical and Mechatronic Engineering,
University of Stellenbosch,
Private Bag X1, Matieland 7602, South Africa.*

Thesis: MEng (Mechatronic)

December 2020

The maritime industry is a cornerstone in the modern globalised economy. Efficient operation of ocean-going vessels is of great importance from both financial and environmental perspectives. Carbon emissions from maritime activities are projected to increase significantly in the coming decades. Short term strategies to address the carbon footprint issue calls for research around topics such as efficiency optimisation of ocean-going vessels.

Emerging digital twin platforms are allowing asset owners and operators to manage the vast information networks that monitor asset performance. Digital twins provide a way to plan, monitor and simulate various operating environments to find optimum configurations. Machine learning methods are harnessed to provide an innovative solution to modelling of data-driven problems which could be very useful in the prediction of asset responses for various operational scenarios. Speed and route optimisation with the use of data-driven models are prerequisites in the attempt to provide decision support capacity to gain tactical foresight for maritime operations.

The SA Agulhas II (SAAII) is a polar supply and research vessel owned and operated by the South African Department of Environment, Forestry and Fisheries (DEFF). This vessel is of particular importance due to the large quantity and variety of data, for both open water and ice navigation, that are recorded

during annual voyages to Antarctica, Marion and Gough Islands. Data is comprised of physical measurements from on-board sensors and diligent observations of ocean and ice conditions. Reconciliation and synchronisation of observed and machine data from the ship's central measurement unit (CMU) was successful and paved the way towards effective data-driven modelling. Two different machine learning models, support vector regression (SVR) and artificial neural networks (ANN), were trained to predict the powering performance of the SAAII for open water and ice navigation while subjected to various atmospheric and ocean conditions. Output power is directly relatable to fuel consumption and was successfully estimated from trained models. A non-linear relationship between power and speed is observed and provides an opportunity to optimise ship operations in terms of cost or time.

Speed optimisation illustrates the financial cost-benefit impact of operating at higher speeds and power levels. A pilot exercise is defined to assess the applicability of data-driven models in a route selection context. A dynamic optimisation technique is successfully implemented to account for the stochastic, time-series characteristics of weather conditions over a voyage path. Data-driven modelling and optimisation offer breakthrough opportunities to ensure the modernisation and sustainability of the SAAII in the context of a South African presence within Antarctic and Southern Ocean research.

Uittreksel

Datagedrewe Regressie Modelle vir Vaart Kosteoptimering Gebasseer op die Operasionele Omstandighede van die SA Agulhas II

*(“Data-driven Regression Models for Voyage Cost Optimisation Based on the
Operating Conditions of the SA Agulhas II”)*

P.G. Durandt

*Departement Meganiese en Megatroniese Ingenieurswese,
Universiteit van Stellenbosch,
Privaatsak X1, Matieland 7602, Suid Afrika.*

Tesis: MIng (Megatronies)

Desember 2020

Die seevaartbedryf is ’n hoeksteen van die moderne wêrelddeksonomie. Doeltreffende see-verwante bedrywighede uit beide ’n finansiële en omgewingsbewustheids perspektief is van belang. Die bydraes van grootskaalse koolstofvrylatinge uit die seevaartbedryf word na verwagting beraam om noemenswaardig toe te neem. Korttermyn doelwitte om die koolstofvoetspoor aan te spreek, vra van kundiges om onderwerpe soos doeltreffendheidsoptimering van skepe te ondersoek. Innoverende digitale platforms is besig om bate eienaars en operateurs te bemagtig met die vermoë om ’n geweldige hoeveelheid inligting van sensornetwerke bestuur kan word. Hierdie digitale platforms skep die geleentheid om beplanning, kontrolering en simulاسies vir verskeie operasionele omstandighede uit te oefen, sodat die optimale konfigurasie van veranderlikes geïdentifiseer kan word. Masjienleermetodes word gebruik om ’n oplossing vir die modellering van datagedrewe probleme te bied. Spoed- en roeteoptimering, met die gebruik van datagedrewe modelle, is voorvereistes in die poging om tegnologie te ontwikkel wat in verband met toekomsgerigte taktiese besluitneming ondersteuning te bied.

Die SA Agulhas II (SAAII) is ’n polêre verskaffing- en navorsingskip wat deur die Suid-Afrikaanse Departement van Omgewing, Bosbou en Vissery besit

word. Hierdie skip is van belang weens die beskikbaarheid van 'n groot hoeveelheid en verskeidenheid data uit jaarlikse ekspedisies na Antarktika, sowel as Marion- en Gough eiland. Hierdie data was tydens oop water en ysnavigasie omstandighede opgeneem. Die datastel bestaan uit metings deur sensors wat op die skip geïnstalleer is, asook ys- en golfobservasies wat deur vrywilligers aangeteken word. Die rekonsiliase en sinkronisasie van al die databronne was suksesvol en baan die weg na doeltreffende modellering van die skip se gedrag. Twee verskillende masjienleer modelle, naamlik ondersteuningsvektor regressie (*support vector regression*) en kunsmatige neurale netwerke (*artificial neural networks*), is ondersoek. Die modelle is geleer om die drywing van die SAII, met die effek van weers- en ysomstandighede in ag geneem, suksesvol te voorspel. 'n Nie-linêre verhouding tussen drywing en spoed is waargeneem. Tesame met die gevolgtrekking dat brandstofverbruik direk gekoppel is aan die skip se uitset drywing, skep dit geleentheid om die uitvoering van operationele planne in terme van koste of tyd te optimeer.

Spoedoptimering het die koste-voordeel impak van operasies met hoë spoed en drywing geïllustreer. 'n Loodsoefening is gedefinieer om die toepassingswaarde van 'n datagedrewe model, met betrekking tot 'n roete-keuse konteks, te wys. 'n Dinamiese optimeringstegniek wat voorsiening te maak vir veranderende en tydafhanklike weersomstandighede oor die lengte van 'n seeroete, is geïmplementeer. Datagedrewe modellering en optimering skep nuwe geleenthede om die modernisering en volhoubaarheid van die SAII, binne die konteks van 'n Suid-Afrikaanse teenwoordigheid in die Antarktiese navorsingsgemeenskap, te verseker.

Acknowledgements

This thesis took hard work and dedication complete. Foremost, I praise God for blessing me with the wisdom and strength to complete this piece of work. "I can do all things through Him who strengthens me," (Philippians 4 verse 13).

I would like to thank my supervisor, Prof. Annie Bekker, for her continuous support, guidance and knowledge throughout the course of this study. Without her enthusiasm and drive, I would not have received the once in a lifetime opportunity to be part of the SANAE 57 relief efforts to Antarctica. She has a passion for research, and walks the extra mile to see her students reach success in their own studies.

Lastly, I want to express my sincere gratitude to friends and family who has been part of this journey from start to finish. Their unconditional support did not go unnoticed. I would like to thank my loving parents, Pieter and Rina, my brother, André, and my dear friend, Francisca, for their guidance and motivation.

Contents

Declaration	i
Abstract	ii
Uittreksel	iv
Acknowledgements	vi
Contents	vii
List of Figures	ix
List of Tables	xiii
1 Introduction	1
1.1 Background	1
1.2 Motivation	3
1.3 Objectives	6
2 Literature review	9
2.1 Introduction	9
2.2 Modelling ship dynamics	11
2.3 The SA Agulhas II - a valuable asset for data-driven modelling and optimisation	14
2.4 Introduction to machine learning theory	16
2.5 Chapter summary	20
3 Data acquisition and processing	22
3.1 Data collection	22
3.2 Synchronisation problem of the CMU data	24
3.3 Observations from synchronised data	28
3.4 Chapter summary	35
4 Machine learning	37
4.1 Introduction	37

4.2	Model architecture	37
4.3	Model validation	48
4.4	Chapter summary	52
5	Voyage cost optimisation	55
5.1	Introduction	55
5.2	Theoretical overview of optimisation methods	55
5.3	Particle swarm optimisation in the open water environment	57
5.4	Results	60
5.5	Chapter summary	68
6	Conclusion	69
6.1	Introduction	69
6.2	Reflection on modelling success	73
6.3	Future work	74
	Appendices	77
A	Algorithms	78
A.1	Synchronisation of CMU data	78
A.2	Synchronisation ice and wave observations with CMU data	82
A.3	Support vector regression	90
A.4	Feed-forward neural network	95
A.5	Particle swarm optimisation	100
B	Observations from previous voyage data	104
C	Fuel cost calculation	107
C.1	Calculation of running cost	107
C.2	Calculation of fuel cost	108
	List of References	109

List of Figures

1.1	The SA Agulhas II at Neumayer Station (January 2018).	2
1.2	Route for 2017-2018 relief voyage from Cape Town (1) to Bouvet Island (2), Antarctica (3) and South Georgia (4).	3
1.3	Change in CO ₂ emission and intensity according to ship classes (Olmer <i>et al.</i> , 2017).	5
1.4	Flow diagram of project objectives.	8
2.1	Roadmap from data measurement to decision aiding (Bekker, 2017)	10
2.2	Perspectives gained from full-scale operational data.	10
2.3	Various operational and environmental factors that effect energy efficiency (Yoo and Kim, 2018).	12
2.4	Time series dependency of ship powering dynamics for steady state operation (Yoo and Kim, 2018).	13
2.5	Power versus speed curves for different Beaufort numbers (Yoo and Kim, 2018).	15
2.6	Diagram of dynamic optimisation.	16
2.7	Main differences between classical programming and machine learning (Chollet, 2018)	17
2.8	Example of a decision boundary for linearly separable problems. Adapter from Pedregosa <i>et al.</i> (2011).	19
2.9	Flow diagram of a general neural network architecture. Adapted from Chollet (2018).	20
3.1	Synchronisation problem between the machine control and navigation data sets. Data from 2017-2018 relief voyage.	25
3.2	Time domain synchronisation convergence plot of every 10 th navigation sample.	26
3.3	Histogram of the synchronisation error for the 2017-2018 relief voyage data with a temporal resolution of 3 minutes.	27
3.4	Synchronised data and corresponding route for the 2017-2018 Antarctic relief voyage.	28
	(a) Synchronised machine control and navigation data.	28
	(b) Route for the 2017-2018 Antarctic relief voyage.	28
3.5	Synchronised data and corresponding route for the 2019-2020 Antarctic relief voyage.	30

(a)	Synchronised machine control and navigation data.	30
(b)	Route for the 2019-2020 Antarctic relief voyage.	30
3.6	Scatter plots of power versus SOG showing stationary, ice and open water modes.	31
(a)	Scatter plot of 2017-2018 relief voyage.	31
(b)	Scatter plot of 2019-2020 relief voyage.	31
3.7	Histogram of noteworthy CMU parameters during open water navigation (2017-2018 relief voyage).	33
(a)	SOG	33
(b)	Starboard power	33
(c)	Starboard propeller pitch	33
(d)	Starboard shaft speed	33
(e)	Wind speed	33
(f)	Relative wind direction	33
3.8	Histogram of noteworthy CMU parameters during ice navigation (2017-2018 relief voyage).	34
(a)	SOG	34
(b)	Starboard power	34
(c)	Starboard propeller pitch	34
(d)	Starboard shaft speed	34
3.9	Pie charts of ratios between open water, ice and stationary data.	35
(a)	2017-2018 voyage.	35
(b)	2019-2020 voyage.	35
4.1	Predictive performance of the SVR open water model.	42
(a)	Power and SOG plot over time with estimated power from the open water SVR model (10 Dec. - 19 Dec. 2019).	42
(b)	Scatter plot of open water model predictions.	42
4.2	Predictive performance of the SVR ice model.	43
(a)	Power and SOG plot over time with estimated power from the ice navigation SVR model (29 Dec. - 30 Dec. 2019).	43
(b)	Scatter plot of ice model predictions.	43
4.3	Convergence of MAE on test data during training iterations of the open water and ice navigation neural networks.	45
(a)	MAE for the open water model	45
(b)	MAE for the ice model	45
4.4	Predictive performance of the FFNN open water model.	46
(a)	Power and SOG plot over time with estimated power from the open water FFNN model (29 Dec. - 30 Dec. 2019).	46
(b)	Scatter plot of open water model predictions.	46
4.5	Predictive performance of FFNN ice model.	47
(a)	Power and SOG plot over time with estimated power from the ice navigation FFNN model (01:00 to 05:00 on 30 Dec. 2019).	47

(b)	Scatter plot of ice model predictions.	47
4.6	Predictions from FFNN models on 2019-2020 test data.	49
(a)	Tolerance zone indicating the upper and lower limits of the MAE for open water navigation during the 2019-2020 Antarctic relief voyage. $MAE = 46,5 kW$	49
(b)	Tolerance zone indicating the upper and lower limits of the MAE for ice navigation during the 2019-2020 Antarctic relief voyage. $MAE = 117.43 kW$	49
4.7	Predictions from FFNN models on 2017-2018 test data.	50
(a)	Tolerance zone indicating the limits of the MAE for open water navigation during the 2017-2018 Antarctic relief voyage. $MAE = 153.13 kW$	50
(b)	Tolerance zone indicating the limits of the MAE for ice navigation during the 2017-2018 Antarctic relief voyage. $MAE = 308.67 kW$	50
4.8	Predictions from FFNN models on Weddel sea data (2018-2019).	52
(a)	Section of open water data from the Weddel sea expedition. Prediction from FFNN model trained on 2019-2020 relief voyage. $MAE = 412 kW$	52
(b)	Section of open water data from the Weddel sea expedition. Prediction from FFNN model trained on 2017-2018 and 2019-2020 relief voyage. $MAE = 191 kW$	52
4.9	Tolerance zone indicating the limits of the MAE for ice navigation during the Weddel sea expedition. $MAE = 738.62 kW$	53
5.1	Diagram of proposed cost structure.	59
5.2	Convergence of the PSO algorithm on the objective function $f(\mathbf{x})$	61
5.3	The effect of the Beaufort number on the power versus speed curve.	61
5.4	The effect of wind direction on the voyage costs.	62
5.5	Costs at sea in terms of power demand and voyage time.	63
(a)	Cost vs voyage time	63
(b)	Cost vs power output	63
5.6	Power and SOG plots for open water navigation from 2-8 Jan 2019.	64
5.7	Route options for the hypothetical voyage.	66
B.1	Histogram of noteworthy CMU parameters during open water navigation (2019-2020 relief voyage).	105
(a)	SOG	105
(b)	Starboard power	105
(c)	Starboard propeller pitch	105
(d)	Starboard shaft speed	105
(e)	Wind speed	105
(f)	Relative wind direction	105

B.2 Histogram of noteworthy CMU parameters during ice navigation (2019-2020 relief voyage).	106
(a) SOG	106
(b) Starboard power	106
(c) Starboard propeller pitch	106
(d) Starboard shaft speed	106

List of Tables

2.1	Specification of 4600 TEU class container ship and propulsion system (Yoo and Kim, 2018).	15
3.1	CMU Variables and units.	23
3.2	Parameters gauged from ice and wave observations	24
4.1	Computer specifications.	38
4.2	Pearson correlation matrix between most noteworthy CMU parameters.	39
4.3	Training variables for the machine learning model.	40
4.4	MAE scores for SVR open water and ice navigation models.	41
4.5	MAE scores for FFNN open water and ice navigation models based on the 2019-2020 test data.	45
5.1	Variables and constant selected for PSO.	60
5.2	Weather scenarios noted in Figure 5.5a.	63
5.3	Comparison between cost and time optimisation for a 3000 km voyage.	64
5.4	Ice navigation optimisation results for a 300 km ice route.	65
5.5	Breakdown of weather vectors for two possible routes.	67
5.6	Route optimisation results.	67
C.1	Breakdown of estimated hourly costs (ZAR).	107

Chapter 1

Introduction

The maritime industry plays an integral role of modern day life. Globalisation and availability of goods from around the world, which is commonplace in the modern era, would not be possible without international shipping. It connects countries from across the world to facilitate trade and international relations and can be considered as a cornerstone to the international economy (Cosci, 2018).

Maritime operations make it possible to conduct research activities in some of the most remote and isolated regions on the planet. The logistical solutions that ships offer make it easy to maintain permanently staffed research bases in areas such as Antarctica and the islands of the Southern Ocean. The amount of data gathered from voyages to remote environments contribute to how we understand the effect that climate change has on the oceans, atmosphere, plants and animals which are endemic to island habitats. The preferential location of South Africa allows access to some of the most oceanographically and biologically diverse routes to the southernmost continent in the world. (Ansorge, Skelton, Bekker, de Bruyn, Butterworth, Cilliers, Cooper, Cowan, Dorrington, Fawcett *et al.*, 2017).

The sustainability of the maritime industry is important for the wellbeing of the modern economy and research in geological, environmental and engineering sciences.

1.1 Background

The SA Agulhas II (SAAI) is a South African polar supply and research vessel owned and operated by the South African Department of Environment, Forestry and Fisheries (DEFF). The ship, as shown in Figure 1.1, was built to the PC-5 ice class specification, meaning that she can operate year round in medium first-year ice with some old ice inclusions (DNV-GL, 2017). She measures 121 m in length, 21.7 m wide and is powered by two 4500 kW elec-



Figure 1.1: The SA Agulhas II at Neumayer Station (January 2018).

tric motors connected to drive shafts that turn controllable pitch propellers (CPP). The propulsion system makes it possible for the ship to reach a reported maximum speed of 18 knots in open water and 5 knots in 1 m thick ice.

The SAAII is the ship used by the South African National Antarctic Programme (SANAP) to resupply the research stations in Antarctica, Marion and Gough Island. The voyages to these locations allow for oceanographic and engineering research activities while at sea. The Sound and Vibration Research Group (SVRG) from Stellenbosch University (SU) have been researching the dynamic behaviour of the SAAII since 2012. The rough sea conditions of the Southern Ocean, where the SAAII mostly operates, make it an ideal engineering laboratory to study the drivers of ship vibration, hull loads and operating performance in open water and in ice.

The SAAII undergoes an annual relief voyage to Antarctica to resupply the research station of the South African National Antarctic Expedition (SANAE IV) located in the Queen Maud Land area. The route of the 2017-2018 relief voyage is plotted in Figure 1.2. The ship departed on 8 December 2017 from Cape Town harbour (1) and sailed via Bouvet Island (2) towards Penguin Bukta in Antarctica (3) where provisions for SANAE-IV were offloaded. The ship spent more than a month at the Antarctic ice shelf, navigating through ice between the German Neumayer station and Penguin Bukta. When relief activities were completed, the ship departed for South Georgia (4)

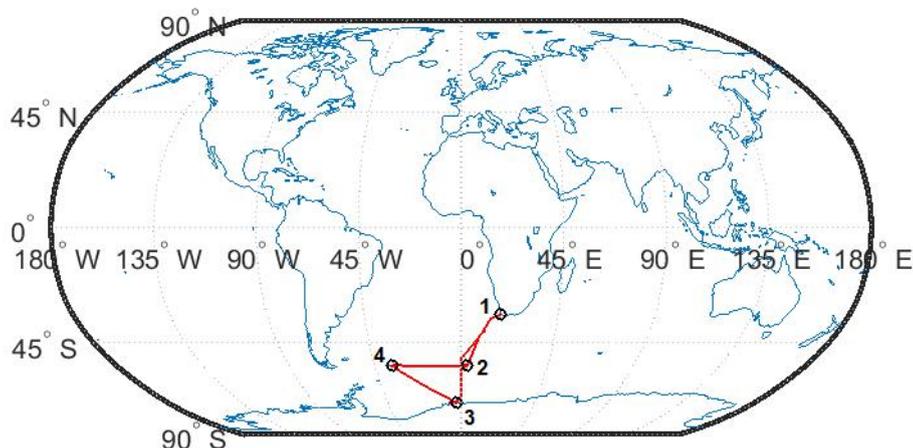


Figure 1.2: Route for 2017-2018 relief voyage from Cape Town (1) to Bouvet Island (2), Antarctica (3) and South Georgia (4).

and arrived on 31 January 2018. From South Georgia the ship sailed back towards Bouvet Island before returning to Cape Town, arriving in South Africa on 13 February 2018. This brief description is typical for an annual relief voyage to Antarctica during summer. Other voyages to the Marion and Gough islands rarely expose the ship to ice due to their locations north of the marginal ice zone. Relief voyages to Antarctica add an unique perspective into the performance of the SAAGI by exposing the ship to extreme conditions, yielding data that is rich in various open water and ice navigation scenarios.

1.2 Motivation

1.2.1 From a climate change and environmental perspective

Climate change is a global phenomenon that the scientific community is only beginning to grasp. It is a mainstream topic in international discussions to find environmentally sustainable policies. This is a key driving force behind innovation in industry, especially in sustainable and renewable energy, which need to be accounted for when realising massive investments for new polar vessels such as the *Polarstern II* (Germany) and the *RRS Sir David Attenborough* (United Kingdom).

During 2012, the maritime industry was responsible for close to 938 million tonnes of CO₂ emissions, accounting for roughly 2.6% of the global total. It is

projected that maritime related CO₂ emissions will increase significantly within the next few decades. In the period up to the year 2050, and depending on the future economic climate, maritime emissions could increase by 50% to 250%. Emission projections show that improvements in the energy efficiency of shipping is an important element in the effort to decrease the rate of CO₂ emission growth (IMO, 2015).

The strategy adopted by the International Maritime Organization (IMO) consists of short, medium and long term measures to reach the goal of reducing greenhouse gas (GHG) emissions by 40% by 2030 (Cosci, 2018). While the medium and long term measures rely heavily on a political drive from participating countries, the short term counter measures are more applicable to current engineering research. Cosci (2018) mentions that the strategy suggests a number of methods to improve shipping efficiency, which include: funding research into low carbon fuels; the development of more efficient ports; and lastly to research route, speed and power optimisation techniques to improve energy efficiency.

In contrast, according to a report from the International Council on Clean Transportation (ICCT), the fuel demand from ships has increased despite the efforts to improve their efficiency. Fuel consumption from international shipping have increased from 291 million tons in 2013 to 298 million tons in 2015 (Olmer, Comer, Roy, Mao and Rutherford, 2017). This report claims that should international shipping be treated as a country, it would have been the sixth largest emitter of energy-related CO₂ in 2015. The graph in Figure 1.3 shows the change in CO₂ intensity with the change in total CO₂ emissions of different ship classes. The yellow bars indicate the change in CO₂ intensity and the blue showing the change in total CO₂ emissions. For almost all of the classes the intensity of CO₂ emissions decreased, in some cases as much as 9%. This figure reinforces that, from either a design or operational perspective, ships are becoming more efficient in terms of energy usage. However, due to the increased demand for shipping during the period of the study, the efforts to improve efficiency have been countered by higher fuel usage. The ICCT report suggests that the mismatch between CO₂ intensity and emissions are unlikely to be substantially reduced by normal business-as-usual improvements (Olmer *et al.*, 2017).

Within this context of climate change, there is a global call for the shipping industry to reduce its environmental footprint. The cost of fuel has become one of the largest items associated with the operating costs (OPEX) of a vessel. Presently, fuel is accounting for almost 50% of the total voyage cost (Bialystocki and Konovessis, 2016). Keeping in mind that, depending on the size and purpose of a vessel, the amount of fuel that is consumed on a voyage can be in the order of a few tonnes per day. A 5% error in estimating the fuel consumption can translate into a substantial financial expense. Another

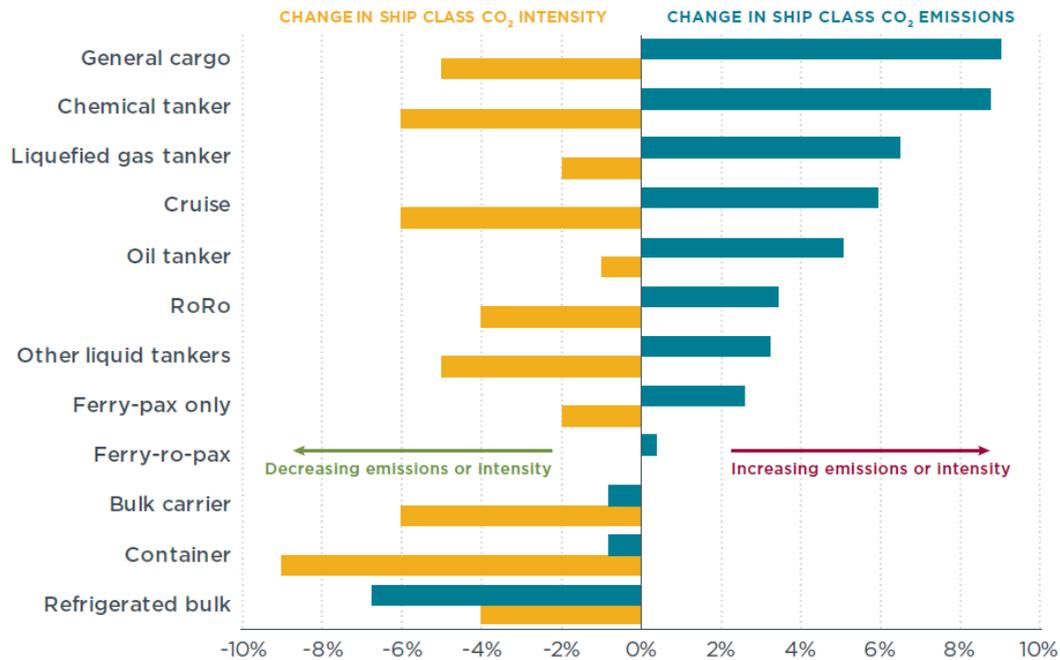


Figure 1.3: Change in CO₂ emission and intensity according to ship classes (Olmer *et al.*, 2017).

important point to mention is the fuel used in shipping is a non-renewable source of energy, which emphasises the fact that it must be consumed in an optimal and responsible manner. Researching methods to optimise fuel usage can find innovative ways to reduce operating costs and CO₂ emissions. These goals are inline with drives from the IMO to reduce the carbon footprint of the maritime industry.

1.2.2 From a technology and innovation perspective

Innovation has always been a driver to obtain a competitive edge in industry. It is fuelled by the prospect of securing new markets or refining and reducing costs in existing ones. Digital platforms allow industries to monitor and understand their processes. The insight from these platforms could identify inefficiencies and inspire ways to resolve them.

The flow of information is an integral part of modern industrial activities. Sensors installed across mines, processing plants or any other industrial assets provide valuable information on the productivity and condition of machinery. Digital solutions should form part of any plant's control and instrumentation infrastructure to manage the flow of information. Modern supervisory control and data acquisition (SCADA) systems are typical examples of this digital infrastructure, although not without its limitations. The challenges to man-

age these vast amounts of data is increasing rapidly. Cost-effective and more readily available sensors provide real or near real-time measurements, and are transforming plants into an industrial Internet of Things (IoT). About 20% of operational budgets can be attributed to poor information management (DNV-GL, 2016). Not only is the management of data important but also the interpretation thereof. A system limited to process monitoring is completely reactive to machinery failure, resulting in costly unplanned down time. On the other hand, a system that has some kind of foresight will enable operators to make corrective decisions in time before faults occur. Digital services should not just be a representation of physical systems but deliver value to the end user.

From various corners of industry, the notion of an *asset as a sensor* is becoming more apparent. Real-time and full scale measurements of assets could be beneficial by advising on the correct use and management thereof. This technology is a cornerstone when considering future endeavours such as the automation of assets. DNV-GL (2016) introduces a digital twin concept where a cloud-based virtual image is created to provide a platform for analysis, insight and diagnostics of an asset. This concept can be part of the solution to address the historical weakness of poor information management while still accommodating the increasing demand for real-time asset monitoring (DNV-GL, 2016). The digital twin, along with advanced analytics and data-driven techniques such as machine learning, can change the way how asset condition and performance is monitored (DNV-GL, 2016). It paves the way for decision aiding technologies with predictive capacity, which aims towards optimising efficient operations (Bekker, 2017) and to improve condition and load monitoring systems (Bekker, Lu, van Zijl, Matthee and Kujala, 2019). Industry is pushing for digital solutions that accomplishes this goal.

1.3 Objectives

In the light of the current economic, environmental and technological climate, it is of interest to find solutions that assist with the management and efficient use of assets within the maritime industry. It is proposed to harness the digital twin concept to investigate data-driven modelling and its contribution to decision support systems within the operational context of the SAII. Challenges include the stochastic and ever changing nature of weather conditions and the complexity of ice-ship interactions that influence the performance characteristics of the vessel. Data-driven modelling and cost optimisation could benefit the ship's operators by creating a tool for route planning which provides a sense of tactical foresight. Ice and weather conditions change daily and routes are often planned from satellite images that are sometimes delayed by a number of hours. It is envisioned that a ship such as the SAII have the technology available to assist with the planning and optimisation of routes, especially in

the Antarctic regions, which does not solely rely on the use of satellite imagery. Routes could be recommended in terms of the quickest voyage time between waypoints, or in terms of minimum cost by means of route selections that improve a ship's efficiency (Zhang, Zhang, Zhang and Mao, 2019). It is worth exploring the applications of this idea within the operational context of the SAII.

Cost optimisation in terms of time, energy efficiency and speed are the first step toward route optimisation. The objectives of this study are focussed on the development of a data-driven model that characterises the performance of the SAII, which is valid for a defined range of environmental and operating conditions. This data-driven model will be used to optimise the operating costs for a unit of distance travelled by the ship. It is not the purpose of this model to find the best route but rather to find the optimum speed to minimise costs. The results from this will be an input to a route optimisation problem. The four main objectives are listed as follows:

1. The first objective is to gather and process the operational data from the SAII's central measurement unit (CMU) and environmental observations which was obtained from previous voyages. Analysis of the data is required to show the distribution and correlations between variables. Lastly, the data has to be prepared for regression model training.
2. The second objective is to use suitable machine learning algorithms and train a data-driven regression model of the output power based on operational data from previous voyages. The validity of this model must be tested for both open water and ice navigation.
3. The third objective is to use the regression model in an optimisation problem to minimise operating costs by finding the optimum speed in simulated operating conditions. The cost function will be expressed as the sum of fuel and overhead costs.
4. The fourth objective is to illustrate the decision support value of data-driven modelling and cost optimisation in a pilot cost-benefit exercise for route recommendation and selection under simulated operating conditions. The models should predict the best route based on waypoints and artificial weather conditions.

The flow diagram in Figure 1.4 provides a graphical representation of the four defined project objectives. The completion of all four stages presents an opportunity to attempt comprehensive route optimisation for both open water and ice, which falls within the overarching goal from the IMO to find operational strategies that improve efficiency (Cosci, 2018).

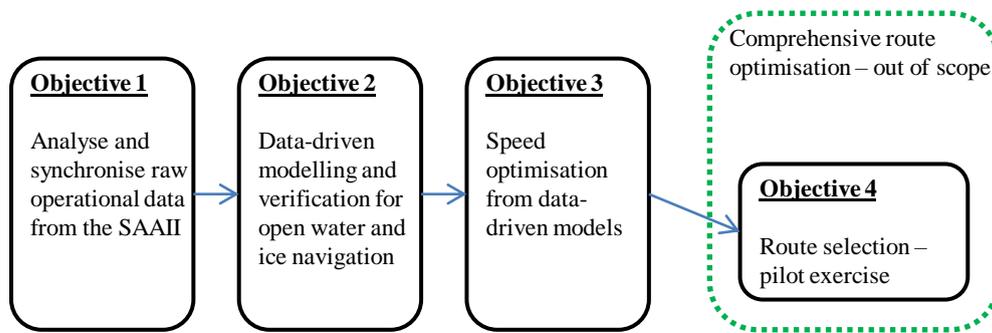


Figure 1.4: Flow diagram of project objectives.

Over the past few years, the SAAII has been fitted with many different types of sensors to measure structural vibration, hull loads, ship dynamics, machine settings and navigational parameters. Massive amounts of data is available from past voyages to Antarctica, Marion and Gough Islands. It is the ideal vessel to base this project on. The success thereof will benefit both the SAAII's crew and owners from an operations and financial perspective.

Chapter 2

Literature review

2.1 Introduction

The modern shipping industry is faced with demands to reduce costs and increase efficiency. Innovations must align with the directives set out by the International Maritime Organisation (IMO) to reduce the carbon footprint of the sector. Energy efficiency can be optimised from a design, operational or strategic point of view (Zhang *et al.*, 2019). It would be a slow process to wait for new and more energy efficient ships to replace the ones currently operating (Johnson and Andersson, 2016), which implies that design-based innovations are not a feasible option in the short and medium term. Instead, research efforts should focus on finding improved operational strategies such as speed optimisation, route selection and effective asset management. The modelling and optimisation of vessels are necessitated by this global drive. Nonetheless, ship operators should not sacrifice effective operational risk and safety management for gains in efficiency. Digital twin solutions aim to provide asset owners with valuable real-time information to make decisions that reduce operating costs and downtime arising from unplanned maintenance (DNV-GL, 2016).

The biggest contributing factors to the operating costs of the SAII is maintenance and fuel. With current provisions in the operating budget, the SAII will have significant budgetary shortfalls from 2020 to 2023. Due to these constraints, the ship cannot spend the desired 160 days per year out at sea (Devanunthan, 2019). This serves as motivation to use the extensive sensor networks on-board the SAII as a platform to explore the possibilities of digitisation and modelling of ship responses to obtain predictive capacity (Bekker *et al.*, 2019).

The progression of information from initial measurement to decision aiding ability is shown in Figure 2.1. The first two stages, measurement and analysis, have been documented in terms of the structural vibration (Soal, Bekker and

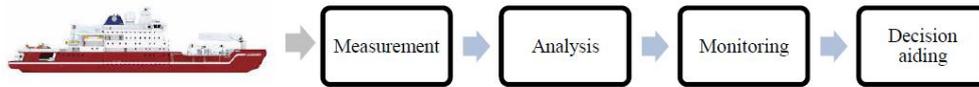


Figure 2.1: Roadmap from data measurement to decision aiding (Bekker, 2017)

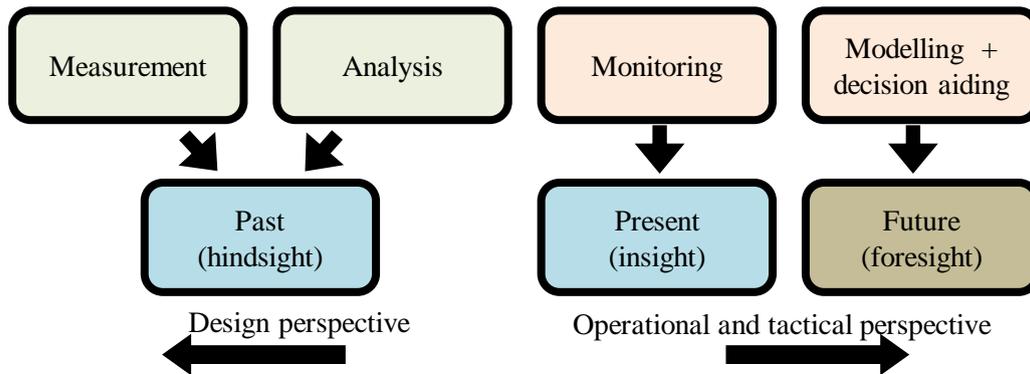


Figure 2.2: Perspectives gained from full-scale operational data.

Bienert, 2015), ice load estimation (Bekker *et al.*, 2019), detection of wave slamming sites (Omer and Bekker, 2016) and the human response thereof (Omer and Bekker, 2017). These studies have contributed to an extensive experience-driven operational and tactical knowledge-base for the SAAII. However, real-time monitoring and decision aiding capabilities are areas that still need attention.

Each step of data processing, Figure 2.1, describe the ship's responses from three distinct perspectives. The flow diagram in Figure 2.2 indicates that the data can be interpreted from a *past*, *present* and *future* orientated point of view. For example:

1. Measurement and analysis both report on what happened in previous voyages.
2. Monitoring systems show the real-time state of the ship.
3. Modelling with predictive capability estimates what the future responses of the ship would be, subject to various operational environments.

A hindsight perspective provides useful feedback for the iterative design process with medium and long term outputs looking into the development of improved components, parts and ship structures. This requires extensive analysis and

investment from stakeholders in the maritime industry. Comprehensive real-time monitoring requires extensive control and instrumentation infrastructure to implement successfully. These systems assist the crew with their operational, *in the moment*, decision making. Modelling and decision support, which aims to provide a sense of *foresight*, is necessary to produce a tactical tool that assists with the planning of shipping speeds and routes. A large potential for improvement in energy efficiency, with noteworthy economic gain, is yet to be exploited (Johnson and Andersson, 2016).

The digital twin is a virtual representation of an asset that allows single source access to information in all three time frames outlined in Figure 2.2. Historical information would include construction reports, quality acceptance tests and historical voyage data. Real-time processes could be monitored and compared to future estimates predicted from the digital model (DNV-GL, 2016). Literature indicates that the key to useful predictive analytics is the accurate modelling of ship dynamics from historical data (Bialystocki and Konovessis, 2016; Gkerekos, Lazakis and Theotokatos, 2019; Yoo and Kim, 2018).

2.2 Modelling ship dynamics

Ship dynamics describes the responses observed from propulsion, buoyancy and environmental forces that are exerted on an ocean-going vessel. These forces originate from the propulsion and steering systems within a highly variable operating environment. The powering performance is predominantly dependent on speed but environmental factors induce a considerable amount of variance in the power-speed relationship (Yoo and Kim, 2018). The SAARI predominantly operates in the Southern Ocean and around the coast of Antarctica. The load profiles of open water compared to ice navigation are very different. Characterising and understanding these significant differences are the key to developing a successful power performance model.

2.2.1 Performance indicators

Equation 2.2.1 describes the energy efficiency operational indicator (EEOI), which is a common metric used to quantify shipping performance in terms of energy efficiency. Guidelines for its use set out by the IMO (IMO, 2009).

$$EEOI = \frac{\sum_j FC_j \times C_{carbon}}{m_{cargo} \times d} \quad (2.2.1)$$

In Equation 2.2.1, j represents the fuel type; FC_j the total fuel consumption for a voyage; C_{carbon} is the carbon content of the fuel type j ; m_{cargo} is the mass of the cargo; and lastly d which is the total distance for a given voyage. The formulation shows that an improvement in energy efficiency would translate

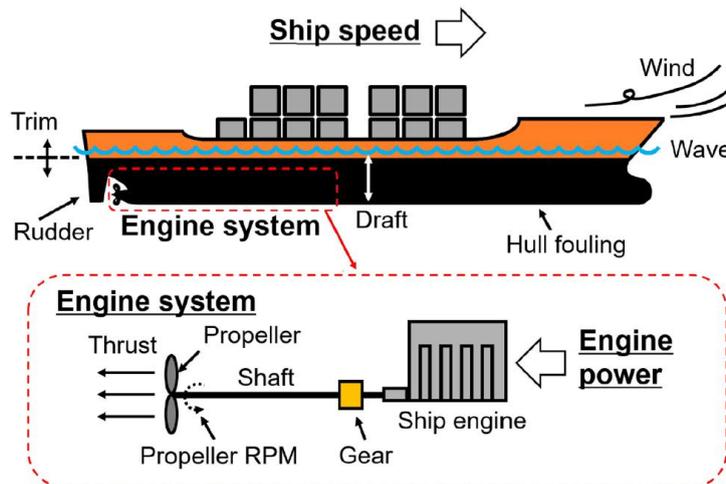


Figure 2.3: Various operational and environmental factors that effect energy efficiency (Yoo and Kim, 2018).

into a decrease of EEOI (Zhang *et al.*, 2019). The fuel type, vessel tonnage and distance are constants to a specific voyage and difficult to influence. The best strategy to improve the efficiency would be to decrease the fuel consumption of the vessel (Wang *et al.*, 2018; Zhang *et al.*, 2019). Wang *et al.* (2018) concludes that speed optimisation for fuel consumption reduction could improve profits significantly.

2.2.2 Factors influencing power demand

Power is required to push a ship through water, as is the case with any mechanical system that does work. Non-linear hydrodynamic forces between the hull and water induce drag that load the propulsion system. Fuel consumption is directly related to power output. To guide fuel consumption estimates, a power versus speed curve is calculated for new vessels during sea trails. However, a single curve is insufficient to describe the powering performance of a vessel for its whole life cycle (Bialystocki and Konovessis, 2016).

Some of the main factors that influence power demand are shown Figure 2.3. The power generated from the engine turns the drive shafts and propellers which create thrust and push the ship forwards. Unless otherwise specified, all references to output power should be considered as the mechanical shaft power driving the ship's propellers. Engine power increases in proportion to speed as a result of non-linear hydrodynamic drag between the hull and water. It is also affected by environmental factors and operational settings (Yoo and Kim, 2018). Bialystocki and Konovessis (2016) refers to three main factors that contribute additionally to the load:

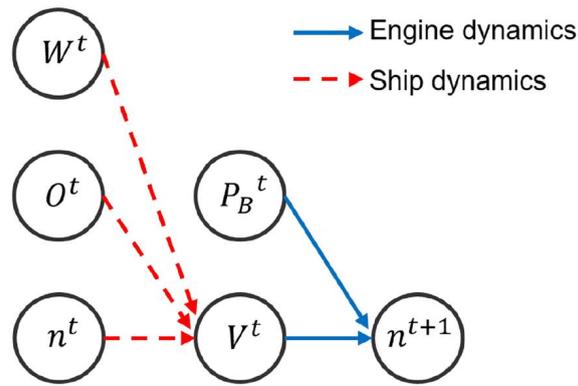


Figure 2.4: Time series dependency of ship powering dynamics for steady state operation (Yoo and Kim, 2018).

1. Increased draft and thus water displacement,
2. Adverse ocean and atmospheric weather conditions,
3. Wear and deterioration of the hull and propeller roughness.

Draft and displacement are the operational parameters that can be adjusted by the crew using hydrostatic and stability tables (Bialystocki and Konovessis, 2016). Weather conditions refer to oceanic (wave height, direction and length) and atmospheric factors (wind speed and direction). Lastly, wear and deterioration of hull and propeller roughness refer to the adverse affects of prolonged bio-fouling and cavitation on the efficiency of a vessel. Its effect can be mitigated by conducting routine maintenance (Bialystocki and Konovessis, 2016).

Apart from the listed elements, the time dependence between them should also be considered. Yoo and Kim (2018) represents the interconnected relationships for steady state conditions as a graphical probability model in Figure 2.4. Weather, W^t , operational settings, O^t , and engine rotational speed, n^t contribute to the ship dynamics at time t which has an effect on the speed, V^t through water. The speed and engine power, P_B^t , directly influence the rotational speed of the engines at the next time step $t + 1$. The continuous co-dependence of the variables, influenced by stochastic weather conditions, make the reliable modelling of ship dynamics very complex (Yoo and Kim, 2018). Analytical or empirical formulations between power and speed may be difficult to determine and could contain too many uncertainties (Zhang *et al.*, 2019). For ice-going ships this relation may be even more complex compared to open water shipping, with no official guidelines in place for the application of energy efficiency strategies for polar navigation (Zhang *et al.*, 2019).

2.2.3 Modelling ice interactions

Ice-ship interactions bring a lot of uncertainties into the modelling of ice navigation, mainly due to inaccurate ice data or the measurement thereof (Zhang *et al.*, 2019). Li *et al.* (2020) describe the *ice breaking* process in detail. Initially, ice will be crushed and start to shear along the edges when a ship enters a sheet of level sea ice. A bending moment is exerted on the sheet due to the vertical contact force between the hull and ice. This bending moment causes the ice to break and rotate parallel to the hull. Some of the ice pieces stay submersed under the hull where the ice heavily interacts with the hull and other pieces of ice (Li *et al.*, 2020).

Ship-ice and ice-water interaction occur on a localised scale which in turn contribute to a ship's performance on the global scale. The scope of this study which aims to predict powering performance in steady state conditions simplifies the problem with modelling ice interactions. On a global scale the randomness of ice thickness and strength would have a limited effect on the power demand estimate, should the mean ice conditions remain relatively constant over time (Li *et al.*, 2020).

In contrast to open water efficiency optimisation, which is usually the shortest and fastest route, is fuel consumption for ice navigation dependent on the selected route. The resistance from one route to the next may not be the same due to differing ice and environmental conditions, resulting in fuel consumption also being subject to route selection (Zhang *et al.*, 2019). This provides a *bigger picture* point of view towards the problem of efficient ice navigation. Apart from finding an optimal speed, route selection is of equal importance for effective and safe operation.

2.3 The SA Agulhas II - a valuable asset for data-driven modelling and optimisation

The SAII is equipped with a multi-sensor data acquisition network that measure vibration, hull strain, operational and environmental parameters (Bekker, 2017; Bekker *et al.*, 2019). Apart from Bialystocki and Konovessis (2016), who used noon reports for a statistical model, most other sources that had access to similar data sets to that of the SAII used various machine learning techniques to successfully model ship propulsion performance (Gkerekos *et al.*, 2019; Wang *et al.*, 2018; Yoo and Kim, 2018). Literature indicates that models could make predictions that account for changes in weather and sea state. Figure 2.5 shows the power versus speed curves for varying sea states estimated by a Gaussian process (GP) model (Yoo and Kim, 2018). A container ship larger than the SAII was involved with the research to produce the reported curves. The ship's specifications are listed in Table 2.1. The Beaufort number is an indi-

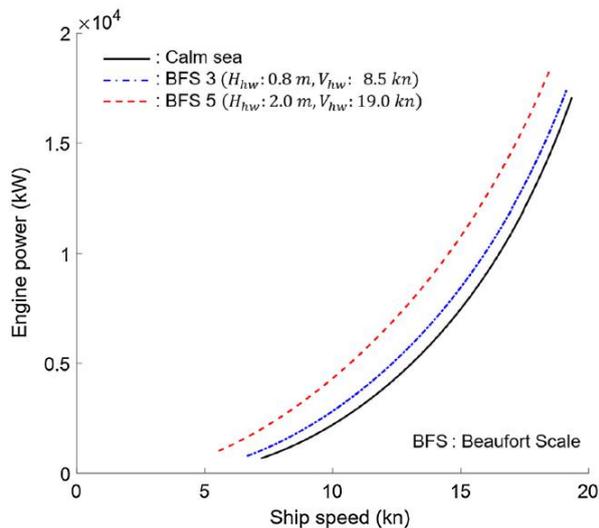


Figure 2.5: Power versus speed curves for different Beaufort numbers (Yoo and Kim, 2018).

Table 2.1: Specification of 4600 TEU class container ship and propulsion system (Yoo and Kim, 2018).

Ship feature	Specification
Overall length	254.7 m
Breadth	37.5 m
Design draft	12 m
Diesel engine rating	25 040 kW @ 95 rpm

cation of the sea state. A higher number represents harsher ocean conditions which in turn translate into a larger power demand to maintain a constant speed. The success of machine learning models, along with the availability of full-scale operational data, make a strong case for a similar approach towards the goal of a data-driven performance model for the SAII. Machine learning is an overarching term that could reference to numerous model architectures. Gkerekos *et al.* (2019) trained nine different regression models using various machine learning algorithms. All of the models were able to make accurate predictions, but the two that were among the best performers were support vector machines (SVM) and artificial neural networks (ANN), which achieved accuracy scores of 95% and higher (Gkerekos *et al.*, 2019). These two methods will be considered for training a powering performance model from the SAII's data. The model would accept input parameters such as speed, shaft and propeller settings and environmental conditions to make motor power demand predictions.

The next step towards a model with decision aiding capability would be opti-

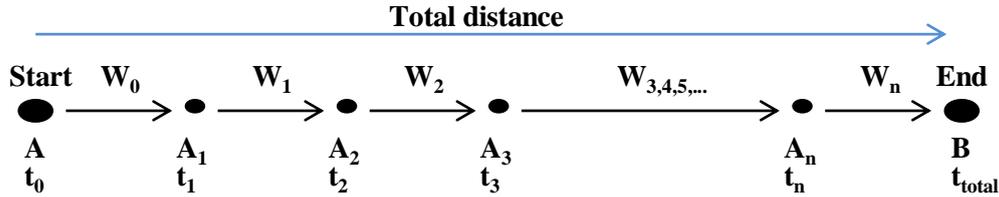


Figure 2.6: Diagram of dynamic optimisation.

optimisation in terms of energy efficiency, fuel consumption or cost. As discussed previously, for a single voyage the EEOI is most affected by fuel consumption (Wang *et al.*, 2018), which has a direct cost implication. Optimisation in terms of fuel consumption would automatically minimise the energy efficiency and voyage costs. The novelty of the model accounting for weather conditions is a possible drawback in the optimisation context. The uncertainty and time-series dependent nature of environmental factors could decrease the accuracy for weather forecasts over longer periods of time. It was shown that weather does affect the power requirement as illustrated in Figure 2.5. Static optimisation methods cannot ensure a reliable recommended speed if the weather changes significantly along the voyage route. To address this problem a dynamic optimisation method is proposed by Wang *et al.* (2018).

The dynamic optimisation method aims to compensate for the time-varying environmental factors along the voyage distance. Figure 2.6 illustrates the methodology behind the method. The voyage distance from point A to B is divided into segments labelled $A_1, A_2, A_3, \dots, A_n$ with the time steps indicated by $t_0, t_1, t_2, \dots, t_n$. For each segment a unique weather vector, W_j , is determined specific to the conditions a ship would see at the given location and time. The optimum sailing speed can then be determined for each segment to compensate for disturbances from changing environmental conditions (Wang *et al.*, 2018).

2.4 Introduction to machine learning theory

Machine learning is a subfield of computer science that is well adapted to process and analyse large, complex data sets (Géron, 2017). The philosophy behind it is very different from classical programming. The flow chart in Figure 2.7 illustrates the differences between the two approaches. According to this diagram, with classical programming the programmer will code the rules according to which the data should be analysed. Hence, when the data is fed into the program, it will result in answers according to the predefined rules (Chollet, 2018). The problem with this approach is that complex data may need complex rules, leading to a scenario where long lists of rules are required for proper analysis. This is not feasible to do by hand. Machine learning will

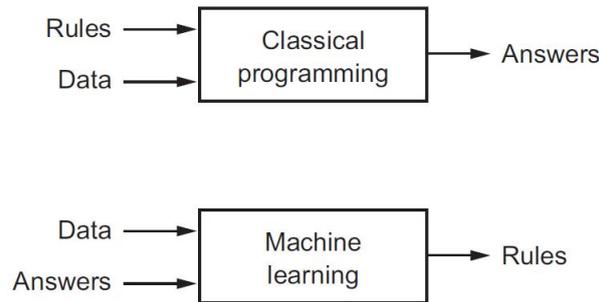


Figure 2.7: Main differences between classical programming and machine learning (Chollet, 2018)

most often simplify a program and give better results than conventional methods (Géron, 2017). With machine learning, the data and expected answers are given simultaneously as inputs to the program. The algorithm will then determine the rules that correlate the input data and expected solution i.e. it creates a data-driven model. New data can now be presented to the model to produce predicted results based on the lessons learnt from the original training data (Chollet, 2018).

There are many techniques that fall under the machine learning field and are mostly classified according to the amount of supervision necessary during training (Géron, 2017). The main types of learning are unsupervised, reinforcement and supervised learning. Unsupervised learning is useful in problems where training data is only available as inputs and the algorithm's goal is to highlight correlations observed from input data (Gkerekos *et al.*, 2019). Reinforcement learning requires the algorithm to make decisions and perform actions. The algorithm is penalised or rewarded based on the success of predicted outcomes as it learns the best strategy to solve a problem (Géron, 2017). A typical simple application is an algorithm which learns the best strategy to win a game. Supervised learning is more in line with the idea shown in Figure 2.7 where trainable input data is given to an algorithm along with the expected results. Supervised learning is the basis for classification (discrete number of outputs) and regression problems (continuous target variables). Therefore, the challenge of predicting power output over time is a regression problem due to the presence of continuous variables (weather, load and speed) that affect the overall resistance of the vessel (Gkerekos *et al.*, 2019).

Care must be taken to ensure that a machine learning model does not learn unwanted trends. The ability of an algorithm to make reliable, repeatable and accurate predictions from new data is the ultimate goal of developing a model in the first place. Poor quality data is one of the main reasons for inaccurate and unreliable models (Géron, 2017). Insufficient volumes and non-representative data are common pitfalls (Géron, 2017). Complex models that

are based on less data will often be outperformed by simple models that are exposed to vast quantities of data (Halevy, Norvig and Pereira, 2009). Gkerekos *et al.* (2019) concludes that the quality of a model is dependent on the quality of the training data. If the data represents only a portion of the sample space or contain irrelevant features, then the model would learn meaningless trends.

Overfitting is another hurdle that requires consideration. A model can learn the correct representations from the training data but may not generalise well to new examples. In such cases overfitting has occurred (Géron, 2017). Yoo and Kim (2018) note that ship performance models based on machine learning algorithms alone are especially vulnerable. It is suggested to include domain knowledge of the physical ship into the design of a regression model to reduce the likelihood of overfitting. Domain knowledge refers to physical ship dynamics such as the fact that speed cannot be increased or maintained at cruising levels without a corresponding supply of power from the engines. If these rules are violated then the model is invalid.

2.4.1 Support vector machines

A support vector machine (SVM) is a supervised learning technique used for classification, regression and outlier detection problems. A SVM does this by mathematically constructing a decision boundary, called a hyperplane, in a higher dimensional space to achieve good separation between different classes of training data. In general, the larger the distance between the decision boundary margins, referred to as support vectors, the lower the generalisation error will be (Pedregosa *et al.*, 2011). Figure 2.8 is a good example of a linearly separable problem. The orientation of the decision boundary is chosen by finding two parallel lines, called support vectors, that separate the red and blue data points with the largest distance, a , between them. SVM's are known as kernel methods, where the name refers to a kernel function which represent the hyperplane that define the decision boundary. Kernel functions are typically determined by hand while the hyperplane is learned from the training data (Chollet, 2018).

SVM's are very memory efficient. Unlike other machine learning methods, SVM's are very well understood and backed by theory and thorough mathematical analysis (Chollet, 2018; Pedregosa *et al.*, 2011). A major drawback of SVM's are that they do not scale very well to high dimensional problems. This could be an issue for its application on the SAAII's data set, which contains more than a dozen features. SVM's also require that internal hyper-parameters be selected appropriately according to the given problem. This gives rise to the issue that the algorithm may need to be optimised for a specific problem by tuning hyper-parameters until the best configuration is achieved. Nonetheless,

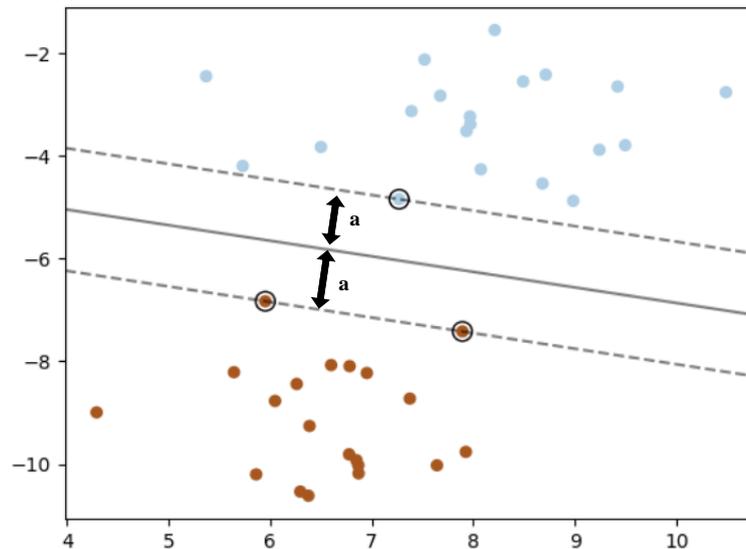


Figure 2.8: Example of a decision boundary for linearly separable problems. Adapter from Pedregosa *et al.* (2011).

because of its success in Gkerekos *et al.* (2019), the method might still have value as a baseline to compare with other techniques.

2.4.2 Artificial neural networks

An ANN is a technique of machine learning where the data is represented in a layered approach. The data is transformed from one layer to the next where, from the algorithm's perspective, it becomes increasingly informative of the final result (Chollet, 2018). The ANN uses these layers to learn connections between the input data and the desired outcomes. The parameterisation of how the input data is transformed in a layer is described by the layer's weights. For ANN, learning happens by tuning the values of the weights so that the network correctly maps the inputs to desired outputs (Chollet, 2018). A general flow diagram that illustrates the learning process is shown in Figure 2.9.

Initially, the values of the weights are random and result in a network with meaningless outputs. Feedback is required so that the network has a way to observe the error between a prediction and the desired value. A loss function is defined for this purpose and in turn calculates a loss score. The key to a ANN's success lies in the process where the weights are adjusted in the direction that minimizes the loss score. The central process behind an ANN is the *backpropagation algorithm* which facilitates this optimisation loop. Sufficient iterations of this training loop, typically tens of iterations over thousands of examples, will result in weights that are tuned to the point where the loss function is minimized. This yields a trained network with minimal loss between the predicted outputs and the target values (Chollet, 2018).

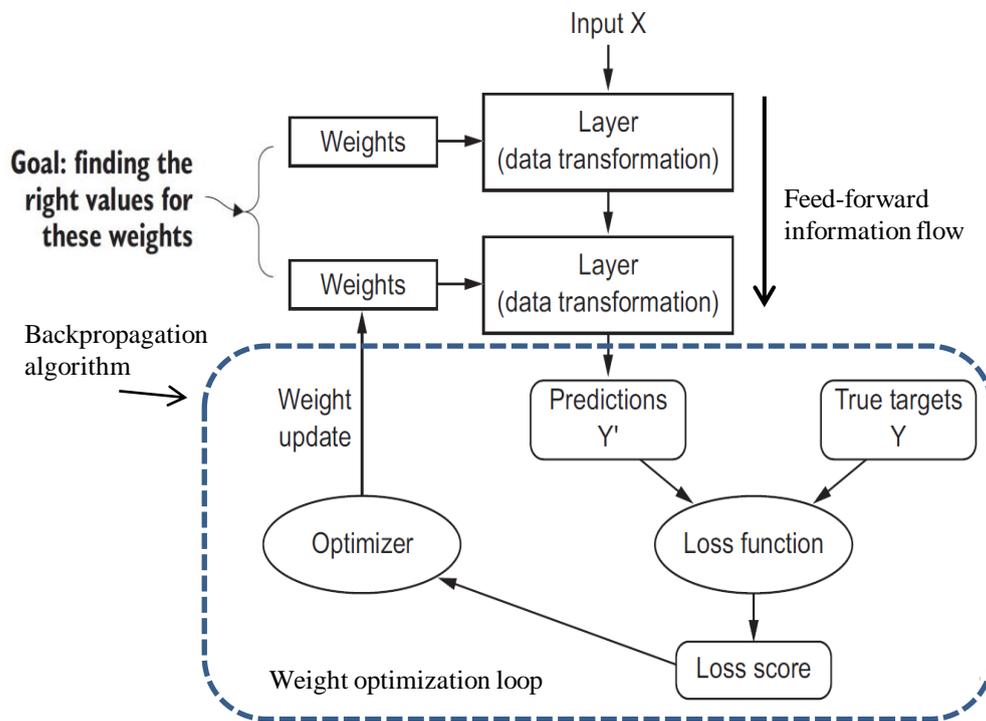


Figure 2.9: Flow diagram of a general neural network architecture. Adapted from Chollet (2018).

The topology of ANN allows information to be processed through the various layers from the input to output. This process can be interpreted as *feed-forward propagation* of information. ANNs display very broad approximation characteristics and can therefore be referred to as *universal approximators* (Bishop, 2006).

2.5 Chapter summary

A digital twin is the gateway for industrial assets to advance into the modern IoT environment. It is a step towards digitising expensive assets to better plan its construction, operation, maintenance and end-of-life phases. The goals defined for this study focuses on the efficient operation of the SAAIL. The building blocks of accurate modelling and optimisation were discussed to show the underlying technologies to access digital twin solutions.

Ships operate in extremely challenging environments that are difficult to forecast and model analytically or empirically. Machine learning provides a fresh perspective that is adapted for modelling of non-linear functions (Bishop, 2006). Combined with optimisation methods it could provide additional fore-

sight for efficient tactical or operational action. The next step towards modelling the powering performance of the SAAM is to inspect, clean and correct the data obtained from its extensive on-board sensor network. Defining the domain of a model is central to determining its application value. This can be observed by examining the quality and distribution of the data available from the 2017-2018 and 2019-2020 Antarctic relief voyages.

Chapter 3

Data acquisition and processing

The data generated from the voyages of the SAAII is valuable as it can be used to develop data-driven models of the ship's responses. A model is only as good as the data used to train it. Therefore, a key concern is that the data is of the required quality and is representative of the whole operating range of the SAAII. Training a model on biased data will inevitably produce a biased model and, therefore, results in an inaccurate digital representation of the SAAII. Model training often requires very large data sets to learn the interconnected relationships between parameters. Thus, the purpose of this chapter is to discuss the methods used to process the raw data from the central measurement unit (CMU) into a large data set of acceptable quality that would be used to train machine learning models. Such a data set must be representative of the ship's operating range and contain enough volumes of data in order for a trained model to be generally applied to new scenarios.

3.1 Data collection

Measurements on the SAAII were recorded and stored on-board on the CMU. These variables relate to the operating parameters of the ship and the surrounding environmental conditions. It was not possible to automatically record and store ice or wave measurements on the CMU. Therefore, visual observations of ice and wave conditions were conducted to include this data of the surrounding sea and ice states. All observations and measurements were made according to UTC standard time.

3.1.1 CMU data

Operational data stored on the CMU contains the operating modes and navigation parameters of the ship. The CMU data is divided into a machine control and navigational set with the variables for both data sets listed in Table 3.1. The machine control data was recorded at 0.5 Hz and comprises of measure-

Table 3.1: CMU Variables and units.

Machine Control Data	Unit	Navigational Data	Unit
Motor current (port)	A	NavTime	N/A
Motor power (port)	kW	Latitude	Deg
Motor speed (port)	rpm	Longitude	Deg
Motor voltage (port)	V	Speed over ground (SOG)	kn
Motor current (starboard)	A	Coarse over ground (COG)	Deg
Motor power (starboard)	kW	Heading	Deg
Motor speed (starboard)	rpm	Relative wind direction	Deg
Motor voltage (starboard)	V	Wind speed	kn
Rudder order (port)	N/A	Water depth	m
Rudder order (starboard)	N/A		
Rudder position (port)	Deg		
Rudder position (starboard)	Deg		
Propeller pitch (port)	%		
Propeller pitch (starboard)	%		
Indicated shaft speed (port)	rpm		
Indicated shaft speed (starboard)	rpm		

ments from the control surfaces and machinery of the ship. The navigational data is sampled at 1 Hz and consists out of navigational and wind related data (Bekker *et al.*, 2019). Sensors which record these measurements are part of the ship's infrastructure and serves as a method to log the operations that were carried out during a voyage. All CMU variables are recorded in real-time and are not time averaged. Due to international construction standards for maritime vessels, on-board instrumentation can be assumed to adhere to industry specifications and good practice.

3.1.2 Ice and wave observations

The navigational data set from the CMU is limited in the sense that it does not have any information about sea states or ice conditions. Therefore, ice and wave observations were conducted to capture this data. Ice observations were only done during ice passage, with observers working rotating 3 hour shifts to ensure uninterrupted observed ice data. Ice conditions can change very rapidly, necessitating a high observation rate. The ice parameters, listed in Table 3.2, were observed every minute with an average calculated for every 10 minute interval. Ice and brash ice concentration were estimated as a fraction out of 10, where 0 referred to no ice and 10 to full ice cover. A zero value for both parameters is interpreted as open water navigation. During ice manoeuvres, the ice tends to rotate along the sides of the ship when it passes through, which allows a view of the thickness. This parameter was estimated by a 1.5 m long ruler protruding from to the side of the ship. The other parameters

Table 3.2: Parameters gauged from ice and wave observations

Ice observations	Wave observations
Snow cover	Beaufort number
Brash ice concentration	Wave direction
Ramming count	Average wave height
Vibration intensity	Max swell height
Ice concentration	Wave length
Ice thickness	Average wave period
Flow size	Average encounter frequency

were estimated based on the observer's judgement. Manual wave observations were done on an hourly bases during visible day time hours and recorded the variables listed in Table 3.2. The Beaufort number is a descriptive variable between 0 and 12 which describe the wind and sea conditions on an increasing scale of severity.

Variability of manual observations are expected. Fatigue, differences between the judgement of observers and the visibility during and observation interval could influence how observations are recorded. This inherent variability should be acknowledged and kept in mind when interpreting results based of visual observations. The most variability could be expected for the ice observation data as this task has the highest observation rate, require many fields to be filled in, and required observers to keep to a rolling 3 hour shift cycle to ensure uninterrupted observations. Variability is most likely be expected in the ice concentration, ice thickness, wave height and wave length observations.

3.2 Synchronisation problem of the CMU data

Data stored on the CMU is not sampled at same sampling rate. The dissimilar sample rates, 1 Hz for navigation and 0.5 Hz for machine control parameters, cause synchronisation problems and prohibits direct comparisons between the two data sets. The unsynchronised data from the 2017-2018 relief voyage is plotted in Figure 3.1, with the speed over ground (SOG) shown in red and the starboard motor power shown in blue. The SOG and power originates from the navigation and machine control data sets respectively. Point A show the last sample for the navigation set, while point B indicate the last sample for the machine control set. Both A and B represent the same point in time. It is illustrated that a sample number from the navigation set, that for example correspond to 13 December 2017, will not align with the same sample number in the machine control set, which represents a different day completely. It is not meaningful to try and find correlations or trends in the data while in this form.

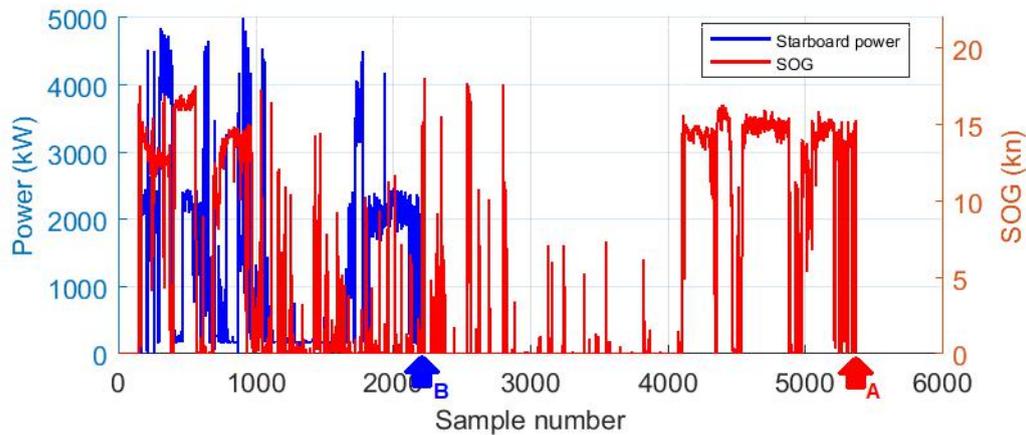


Figure 3.1: Synchronisation problem between the machine control and navigation data sets. Data from 2017-2018 relief voyage.

To preserve the integrity of the time-series data, interpolation cannot be used to create and increase the number of samples for the machine control data. Instead, data points from the navigation set has to be selected and aligned with the machine control data. A simple but somewhat naïve solutions would be to multiply the navigation set with a factor to scale it down to the size of the machine control set. This method was initially explored, however, it became apparent that the missing data points would cause this method to be invalid. This may occur due to a faulty sensor or the system that had to reboot. The CMU does not increment the sample number when no measurement is recorded and does not take into account the time that had passed since the previous successful sample. As a result, the data does not scale linearly using a simple scaling factor.

A more appropriate way to synchronise was to compare and match the time stamps of each data point to within a defined tolerance. By doing so, the data from the navigation set that does not have a match will be ignored. The ship operates mostly in steady state conditions during long distance voyages. It can be argued that the parameters recorded by the CMU is typically slow to respond to environmental changes. Therefore, a maximum synchronisation tolerance of 5 seconds is suggested based on these assumptions. Samples that could not converge within this tolerance were discarded. An algorithm was written in MATLAB for the synchronisation. The interested reader is referred to Appendix A. All data points had a corresponding time stamp that was converted into epoch time format. Epoch time is a real number format that represents a point in time as the total number of seconds that have passed since 1 January 1970 at 00:00:00 UTC. The task of comparing the time stamp information was much simpler in this format. Time stamp information of a machine control samples were compared to the time stamps from the naviga-

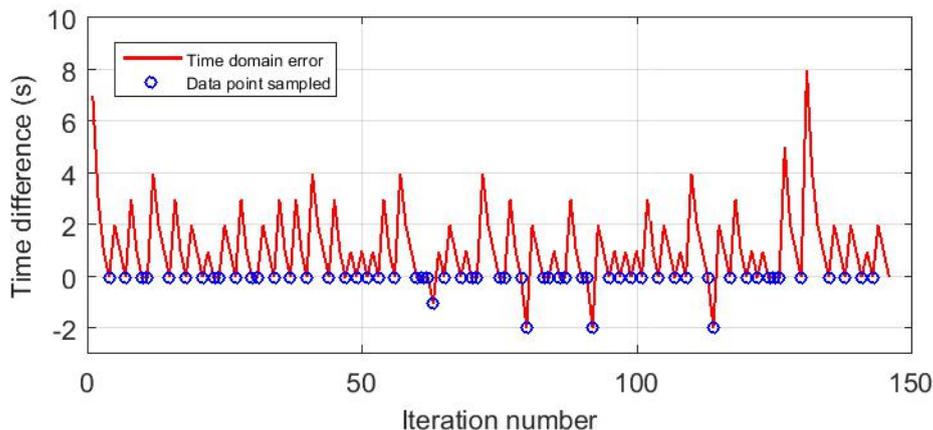


Figure 3.2: Time domain synchronisation convergence plot of every 10th navigation sample.

tion set. The algorithm incremented through the navigation set until it was within the time tolerance. A convergence plot of the first 60 data points is shown in Figure 3.2. The red lines show how the time difference between an individual machine control data point and a corresponding navigation data point converge to zero. The data points are saved and stored in a separate array when it is within the time tolerance, as illustrated in Figure 3.2 as the blue circles. Figure 3.2 shows the majority of data points were sampled with zero time difference between them while a small percentage of points were misaligned by 1 or 2 seconds.

Ice and wave observations were conducted by observers by populating an Excel spreadsheets with the parameters listed in Table 3.2. A very similar method was used to synchronise the observations to the CMU data. The Excel spreadsheet was loaded into MATLAB as an array. Time stamp information was then converted into epoch time format to make useful comparisons. The script incremented through each line of the CMU data, comparing date and time fields until the closest match was found. The CMU data within a 10 minute observation interval was allocated the corresponding discrete information from the ice and wave observations. Averaged data from a 10 minute ice observation interval is therefore spread over a whole 10 minute period of CMU data which corresponds to the same date and time when the observations were made. The temporal resolution of ice conditions are likely to be limited by the observation interval which serves as justification to use average values when synchronising observations with the CMU data. Similarly, data from wave observations were spread over a corresponding 1 hour period.

The 2017-2018 and 2019-2020 Antarctic relief voyage CMU datasets contained enough data points to validate data-driven techniques such as machine learning practices. The combined length of data from both voyages is in the order of

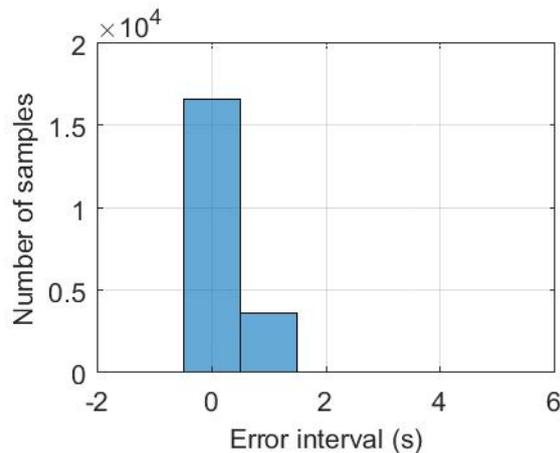
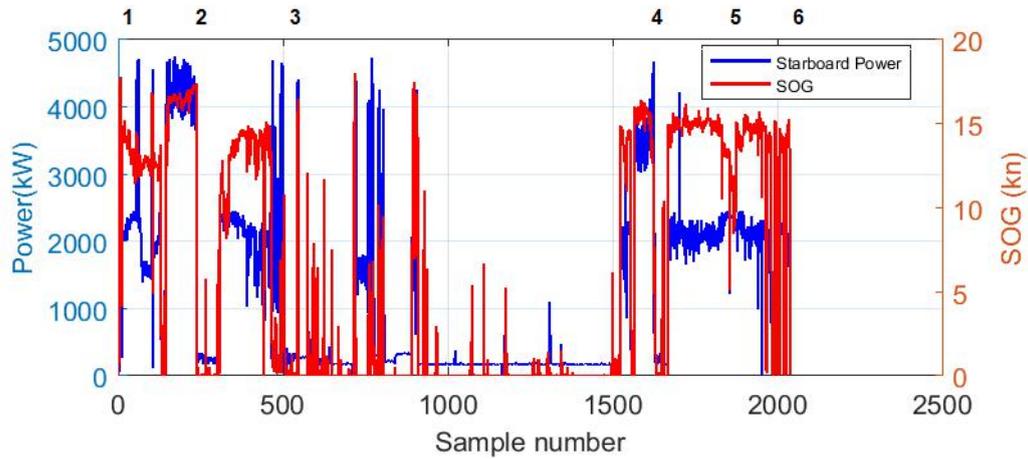


Figure 3.3: Histogram of the synchronisation error for the 2017-2018 relief voyage data with a temporal resolution of 3 minutes.

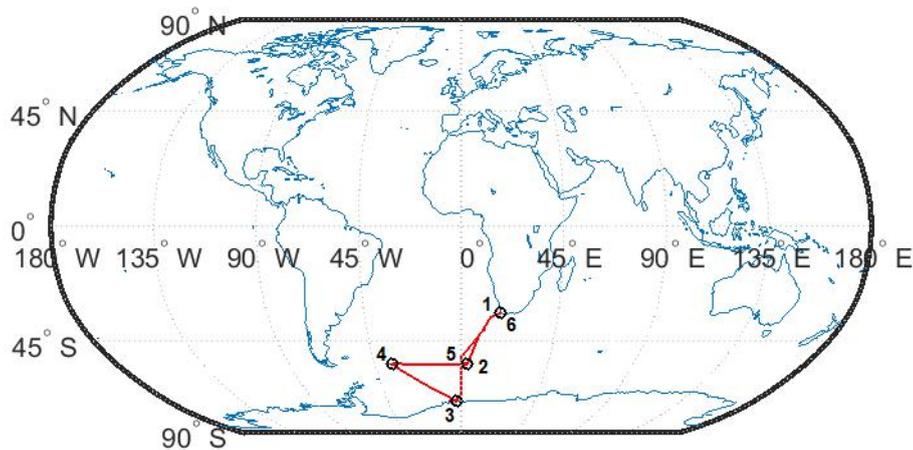
4.5 million and 13 million for the machine control and navigation sets respectively, with each sample containing various fields of measurement, Table 3.1. The synchronisation of all the samples is computationally expensive and time-consuming. The proposed synchronisation method allows for the opportunity to reduce the size of data sets. Every n^{th} sample from the navigation set can be selected and synchronised, thereby creating trade-off between temporal resolution and computational time.

Every 100th sample of the navigation data was used during the synchronisation for open water data. From the 2017-2018 relief voyage this method produced 21 952 synchronised samples of the whole voyage. This translates into a temporal resolution of roughly 3 minutes, which could be considered as acceptable for steady state open water passage. The resolution can be improved but requires more computational resources. Ice data for instance is highly erratic over a short time span and requires a higher temporal resolution. The dates for when ice passage took place were determined from the dates recorded from observations. The corresponding CMU data within these time periods were isolated and synchronise to a resolution of 15 seconds.

The histogram in Figure 3.3 give an indication of the amount of samples that are misaligned after processing with a resolution of 3 minutes. More than 80% of samples have a synchronisation error of zero seconds and 18% with an error of 1 second. Apart from a few outlier samples that are misaligned by more than 1 second, almost all of the data were synchronised successfully. Only 76 samples were not able to converge and was subsequently discarded. This accounts for less than 0.5% of the data set, proving that the technique was an appropriate data synchronisation tool. The same technique was then used to synchronise the 2019-2020 voyage data.



(a) Synchronised machine control and navigation data.



(b) Route for the 2017-2018 Antarctic relief voyage.

Figure 3.4: Synchronised data and corresponding route for the 2017-2018 Antarctic relief voyage.

3.3 Observations from synchronised data

3.3.1 2017-2018 relief voyage

A graph of the synchronised machine control and navigation data for 2017-2018 is presented in Figure 3.4a. The sample numbers on the x-axis do not represent the total number of samples available for model training; instead a reduced number of samples is used for representation purposes. SOG follows the power very well, especially at locations on the graph where the change in power is almost vertical. This indicates that the synchronisation was successful.

Comparisons between the two data sets can now be made with reasonable confidence.

The numbers 1-6 on top of Figure 3.4a cross-reference to the waypoints indicated on the map in Figure 3.4b. The first leg of the voyage, waypoints 1, 2 and 3, consisted of open water navigation for the most part with some ice passage when approaching the Antarctic ice shelf (3). Notice the erratic power output with low SOG at around 3 in Figure 3.4a. A number of weeks on the timeline between 3 and 4 passed as the ship conducted logistical activities, typically at low speed. At waypoint 4 the ship stopped at the island of South Georgia before departing towards Cape Town (6) via Bouvet Island (2, 5). In open water the ship typically operates at a continuous power level of 2500 kW with one exception between waypoints 1 and 2. Here the ship was operated at power levels higher than 4000 kW to obtain speeds of 15-17 kn. This may prove useful during model training and optimisation as this portion of data contain the ship's responses at higher than usual power levels.

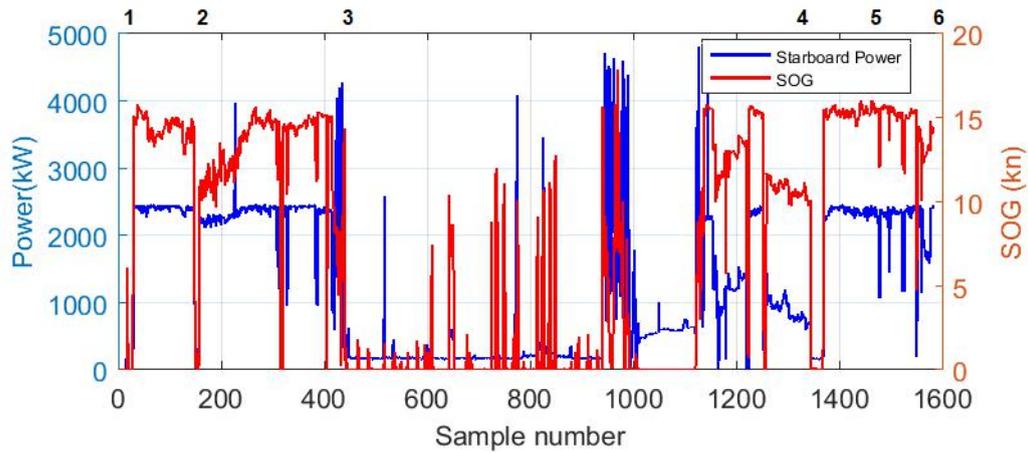
3.3.2 2019-2020 relief voyage

A successfully synchronised plot of the 2019-2020 CMU data is presented in Figure 3.5a with the corresponding voyage route in Figure 3.5b. Again, the sample numbers on the x-axis was reduced for representation and does not indicated the quantity of data available for model training.

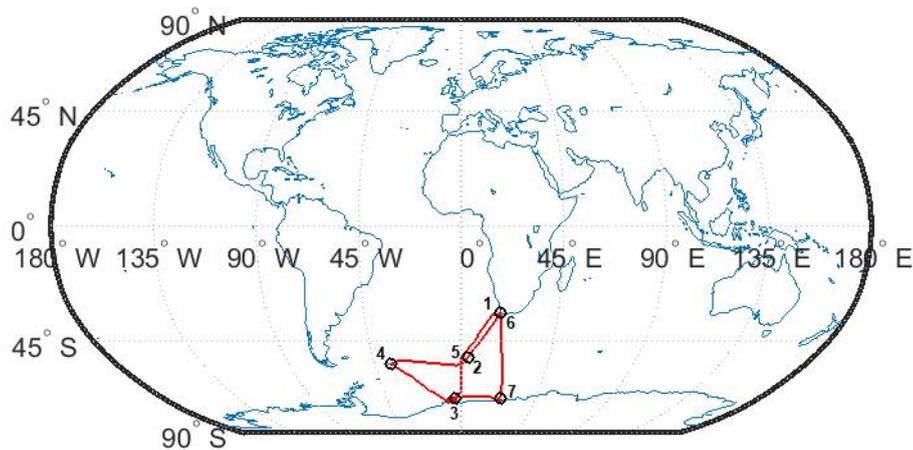
The waypoints indicated on the map, numbers 1-7, correspond to the numbers shown at the top of time Figure 3.5a. The ship left Cape Town (1) and sailed via waypoint 2 toward Penguin Bukta (3) in Antarctica. Power fluctuations are observed with peaks up to 4000 kW in the region of 3, see Figure 3.5a, as the ship passes through ice to reach the Antarctic coast. The ship spent a few weeks at the ice shelf to partake in relief operations before departing towards South Georgia (4). Afterwards the ship headed east towards waypoint 5 before setting a course Cape Town (6). Afterwards, the ship sailed back to Antarctica (7 and 3) a few days later to complete the relief mission. However, the machine control data from the CMU did not record any values from this point onwards. Therefore, from waypoints 1 to 6 is the only portion of the 2019-2020 relief voyage that is useful for model training.

3.3.3 Data correlations

The CMU data was divided into three groups according to the ship's operating modes: open water; ice navigation; and cases where the ship remained stationary or at very low speed and power levels. The open water and ice navigation cases were defined according to the ice concentration parameter from visual observations. A value equal to zero indicated open water movement, while an ice concentration value greater than zero indicated ice navigation. The



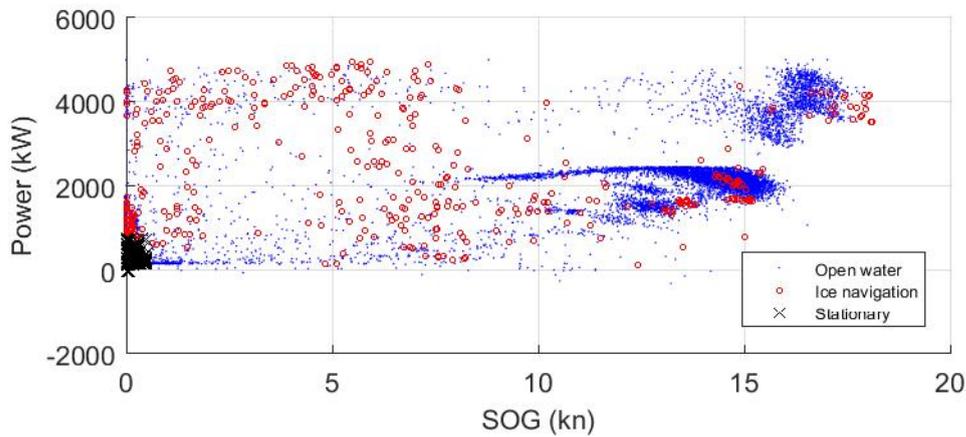
(a) Synchronised machine control and navigation data.



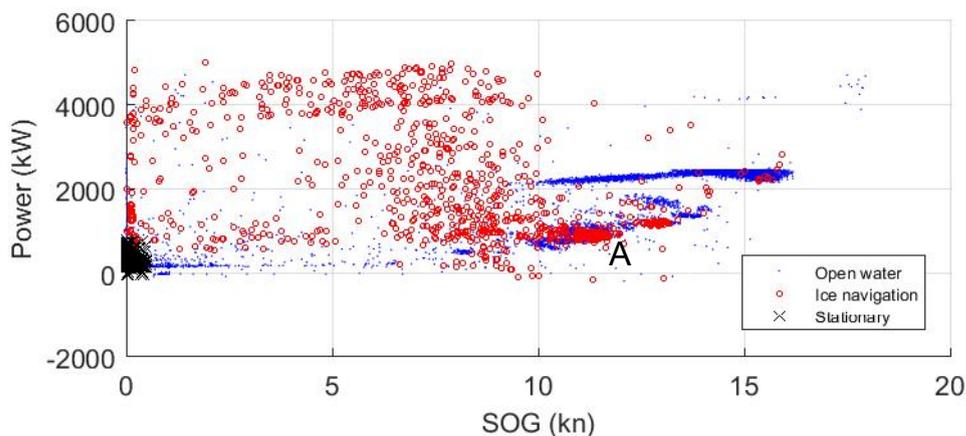
(b) Route for the 2019-2020 Antarctic relief voyage.

Figure 3.5: Synchronised data and corresponding route for the 2019-2020 Antarctic relief voyage.

stationary case was reserved for situations where the output power is simultaneously less than 800 kW with a SOG lower than 0.5 kn. The scatter plots in Figures 3.6a and 3.6b illustrate the relationship between the power and SOG for the two relief voyages. The data is plotted according to the three operating modes. There is a general non-linear relationship visible between power and SOG during open water navigation for both voyages, shown in blue. This is typical of open water operations where the load profile is caused by the hydrodynamic resistance between water and the ship's hull (Yoo and Kim, 2018). The data seems to be tightly grouped, especially around 2000 kW and 15 kn,



(a) Scatter plot of 2017-2018 relief voyage.



(b) Scatter plot of 2019-2020 relief voyage.

Figure 3.6: Scatter plots of power versus SOG showing stationary, ice and open water modes.

which could be promising for determining an appropriate model. The graph for the 2017-2018 voyage, Figure 3.6a, show a cluster of data at high SOG and power levels which is not present in the 2019-2020 voyage. This is a valuable grouping that could subject a data-driven model to data describing the ship's responses under high speed operation, which is considered financially expensive due to increased fuel consumption. Therefore, knowledge of such cases could help models distinguish between normal and high speed operations.

On the other hand, a scattered and seemingly disorganised pattern is present between the power and SOG for ice conditions as indicated by the red circles - a significantly different observation compared to the open water scenario. This chaotic pattern introduces added complexity, which implies increased difficulty for training a model for an ice environment. The power levels are very high, around 4000 kW, for SOG's ranging from 0-10 kn. However, for some instances

the same SOG is recorded at abnormally lower power levels of below 2000 kW, marked **A** on Figure 3.6b. This cluster could be a result of the ship moving through a stretch of calm water called leads or polynyas while ice observation were still conducted. Sea ice tends to dampen water motion, thereby resulting in very calm stretches of water located in the middle of ice fields. The calm water provides a lower resistance to the ship's motion and could explain why the ship can sail at speeds of 10-12 kn with lower than normal power levels.

The poor grouping and scattered correlation of ice navigation data is indicative of the extraneous effects of ice loads on the propulsion system. At low speeds, below 8 kn, the hydrodynamic resistance is very low which demands minimal power to maintain a constant speed. Ice interactions change the resistance to motion rapidly and in a very complex fashion (Li *et al.*, 2020), and as a result the power demand increases significantly up to the ship's maximum limit. Modelling such behaviour is possible with enough volumes of data that are directly relevant to the power-SOG relationship for an ice passage environment.

A data-driven model can only produce estimates for ranges of input values that were available during training. Various histograms of the most noteworthy parameters are presented in Figure 3.7 to illustrate regions where the most data originated during open water navigation. SOG, Figure 3.7a, is concentrated around 15 kn while most of the power samples, Figure 3.7b, cluster near 2500 kW. The similar grouping supports the connection between power and speed as illustrated by the scatter plots in Figure 3.6. Propeller pitch, Figure 3.7c, and shaft rotational speed, Figure 3.7d, were mostly maintained at constant levels during open water passage. When using these parameters as inputs to a data-driven model the configuration of both variables, for the port and starboard side, should be kept in the same general region as shown in the histograms to ensure the validity of a trained model's predictions. This constraint follows the notion that a data-driven model can only predict from data it was trained on. Lastly, Figures 3.7f and 3.7e show most common wind speeds and relative directions measured during the 2017-2018 relief voyage.

The data distribution during ice navigation, on the other hand, is much different. Figure 3.8 presents histograms for power, SOG, propeller pitch and shaft speed during ice operations. The SOG of the ship, Figure 3.8a, is concentrated in the slower end of the speed spectrum. Power demand is uniform over the available range with two peaks visible around 1300 kW and 4000 kW, Figure 3.8b. The relationship between power and SOG is not apparent from inspecting the two histograms, as power sample concentrations are not complemented by a corresponding grouping of SOG samples. This supports the idea that external ice interactions are the leading factor driving power demand. In some cases, the ship may get stuck in ice and must reverse to free itself. Therefore, the propeller pitch often changes position, Figure 3.8c. The shaft speed on the other hand is kept fairly constant near 140 rpm, Figure 3.8d.

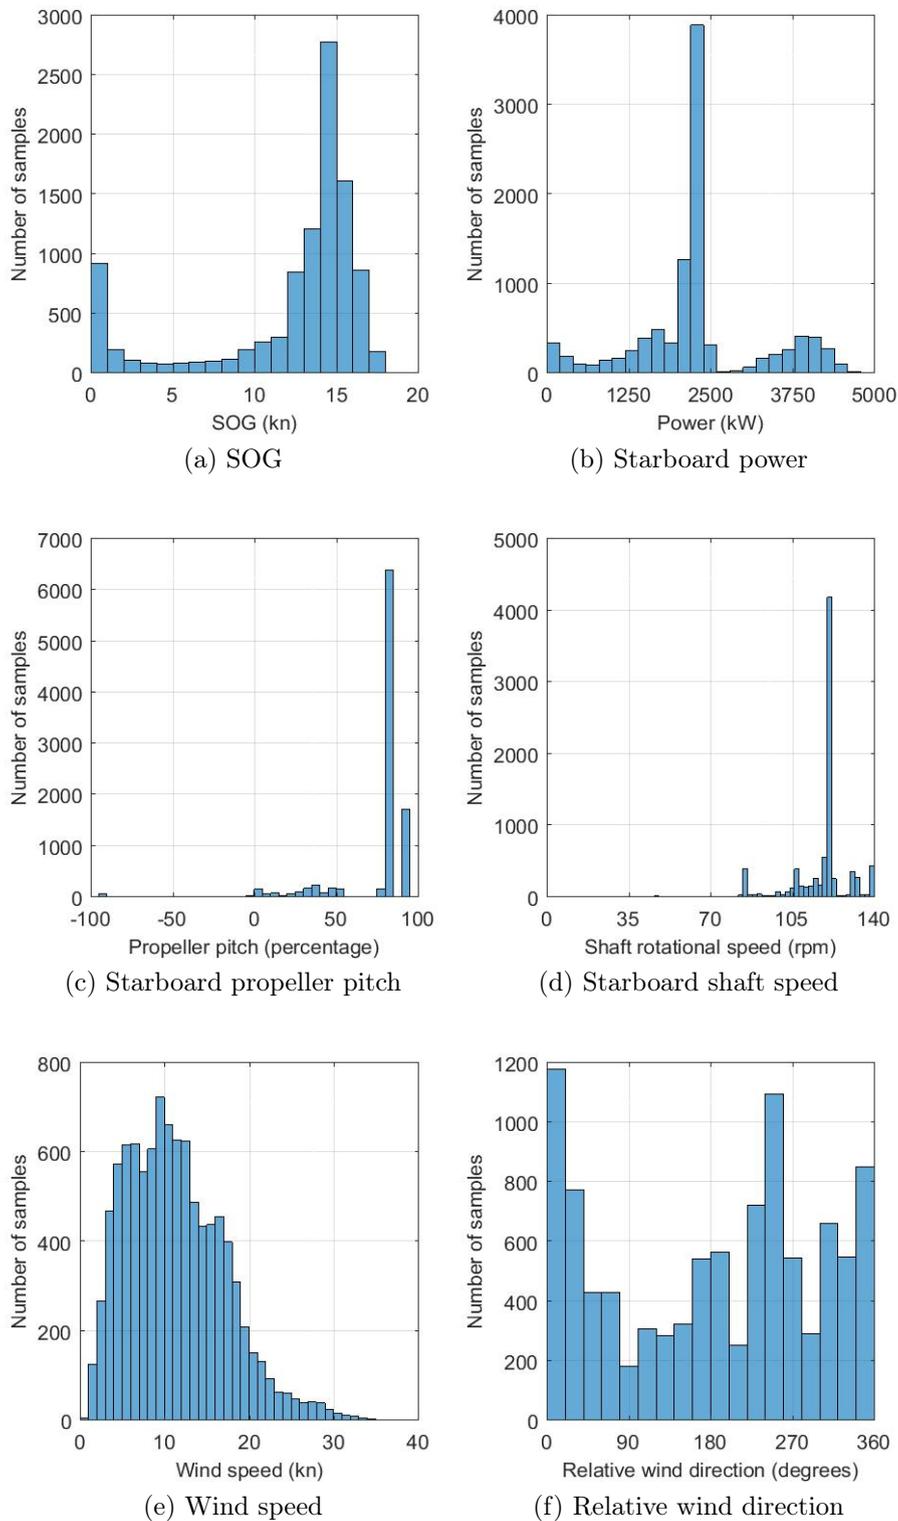


Figure 3.7: Histogram of noteworthy CMU parameters during open water navigation (2017-2018 relief voyage).

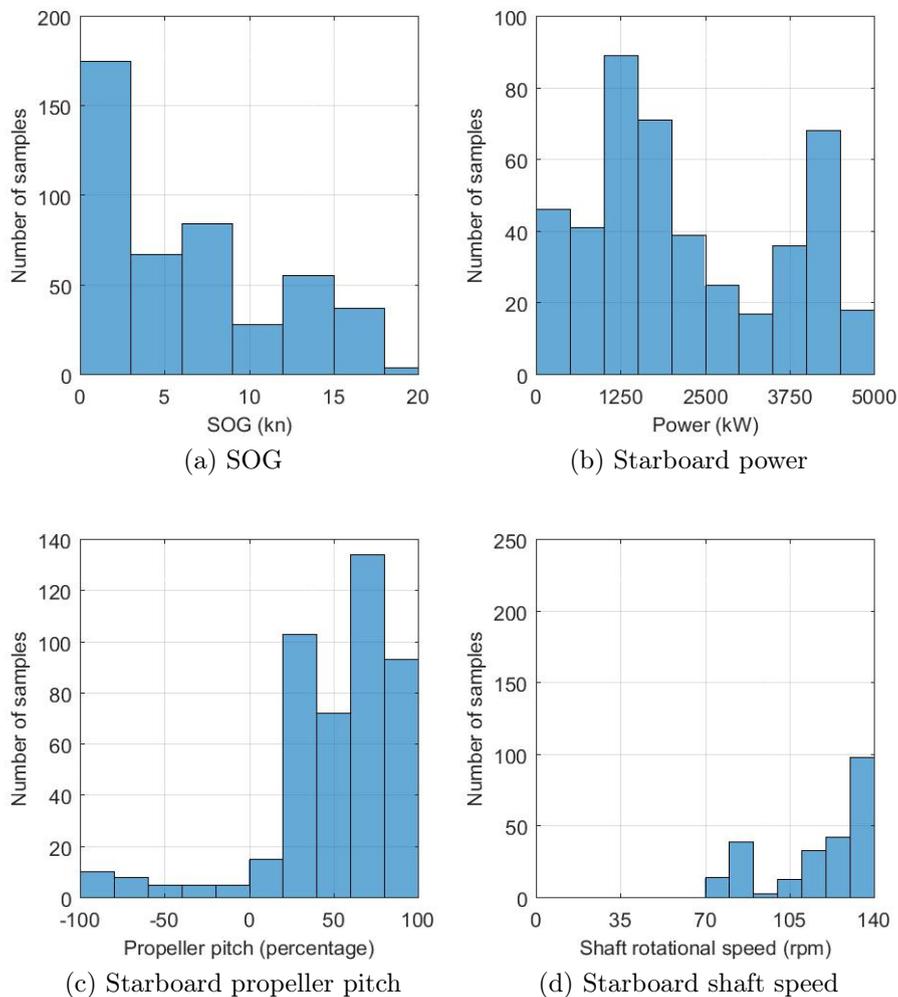


Figure 3.8: Histogram of noteworthy CMU parameters during ice navigation (2017-2018 relief voyage).

The 2019-2020 relief voyage experienced very similar conditions and operation settings with some minor exceptions such as instances of high power demand during high speed open water passage. Refer to Appendix B for similar histograms detailing the distribution and concentration of data for open water and ice navigation during this voyage.

3.3.4 Limitations of CMU data

The quantity of ice data that is available from the CMU is much lower in relation to the open water and stationary data. This might be an issue when training a single model for universal use. The pie charts in Figure 3.9 illustrates ratio of ice, open water and stationary data available from each voyage. In both voyages the ship spent roughly half of the time in open water and a very

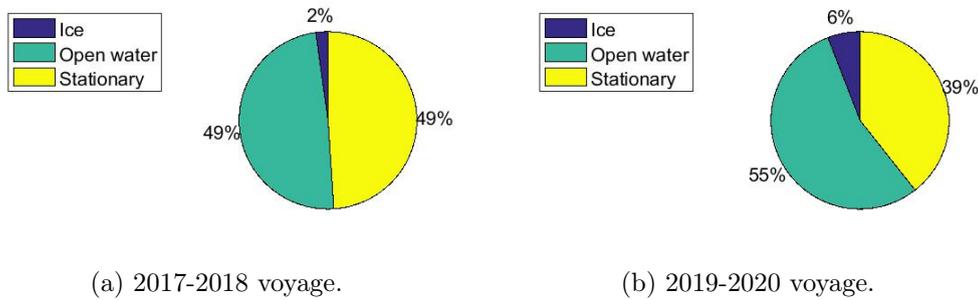


Figure 3.9: Pie charts of ratios between open water, ice and stationary data.

small percentage in active ice navigation. Even when disregarding the amount of time that the ship remains stationary, the proportion of ice to open water data is disproportionately small. This highlights a significant data limitation, as the ship spends most of the time on the voyage's departure and return legs. Ice navigation is limited to essential movement near the Antarctic coast. When the ship reaches the ice shelf, it idles and remains stationary during the offloading of cargo. It would be computationally inexpensive to rather create separate models for the two scenarios of open water and ice navigation. This way, accurate models can be developed separately, whereas an universal model will require significantly more data and computational power to fully define the complex non-linear relationships between the data parameters. Due to hardware limitations, an universal model is not feasible for the scope of the project. Further, a model trained on a disproportionately uneven dataset of ice to open water navigation will produce skewed predictions for ice scenarios, where ice interactions would most likely be misinterpreted for special cases of open water navigation.

Therefore, two separate models are developed using data-driven techniques. The first should be trained on the open water and stationary data. An open water model can then be used to predict the power requirement exclusively for open water navigation. The second model should be trained on the ice data. It might be helpful to improve the temporal resolution so that more ice data is available for training. It should be noted that the 10 minute resolution obtained from ice observations could still be too slow for reliable predictive modelling. With this approach, the two operating modes of the ship can be isolated during model training and would limit the overwhelming bias that the open water data would have on the predictive ability of the model for ice cases.

3.4 Chapter summary

The CMU data from the 2017-2018 and 2019-2020 Antarctic relief voyages were synchronised successfully with minimal error. The data was separated

into three groups based on the operating modes of the SAAIL: stationary, open water, and ice navigation. Scatter plots of all three groups, Figure 3.6, showed that the relationship between power and SOG is very different for open water passage compared to ice navigation. The pie charts in Figure 3.9 indicated that the amount of open water data is significantly more than the ice data for both relief voyages. Using this data to train an accurate model for universal open water and ice application would be unfeasible due to hardware limitations and would require significantly more data. Therefore, two models are recommended rather than one universal model. Histograms of the data showed the spread of the data where ship operations took place most of the time. The distribution of these parameters, which will be the inputs to a data-driven model, defines the domain over which the model would be accurate and robust. The quantity and quality of the data was found to be acceptable for open water navigation. It was noted that a lack of sufficient ice condition data might become a constraint in developing a robust model for this mode of operation.

Chapter 4

Machine learning

4.1 Introduction

In the previous chapter, the cleaning of the data to an acceptable quality standard was discussed. The next step forward is to use this data to train a machine learning model. As previously discussed, machine learning is a term describing a set of computing techniques used to train data-driven models for problems with high orders of complexity (Chollet, 2018; Géron, 2017). Similar to a mathematical model, the data-driven model would make predictions of desired outcomes based on the input data it receives. This chapter provides a walk through of how two sets of models, SVR and ANN, were developed for the SAAII to describe the behaviour of the propulsion system in ice and open water. Models were trained on data from past voyages and were evaluated retrospectively, meaning that the accuracy of predictions could be verified by the actual measurement taken at the time of the voyage. A reliable model would be one that consistently makes accurate predictions of the output power over the whole sample period. Such a model could be a helpful tool that has the predictive ability to estimate power usage over a whole voyage.

4.2 Model architecture

As motivated by literature two techniques were used to develop models for the SAAII: SVR and ANN (Gkerekos *et al.*, 2019). Both methods were implemented with the Python programming language and Jupyter Notebook as the user interface. Scikit-Learn and Keras are open-source libraries that contain functions and algorithms to perform machine learning operations. SVR was implemented using Scikit-Learn while Keras was used to create ANN. Features from the CMU data were selected and normalised. The data was separated according to operating conditions to train models for open water and ice passage independently.

Table 4.1: Computer specifications.

Hardware	Specification
Device model	Dell Inspiron 1545
Operating system	Windows 10
System type	64-bit
Processor	Intel Core Duo CPU - 2.20 GHz
Installed RAM	4.00 GB

4.2.1 Computing hardware

Computing hardware is an important limiting factor that dictates the complexity and size of machine learning models that can be trained. Processor chip design has improved significantly in the past 30 years allowing small ANN to run on the average modern computer with a central processing unit (CPU) (Chollet, 2018). However, more advanced graphic processing units (GPU's) are required for applications that require very high computational power such as computer vision or speech recognition (Chollet, 2018). The hardware specifications that were available for model training and data processing are listed in Table 4.1. Although the use of a GPU is recommended, the data is processed up to a temporal resolution to allow for training on a CPU which would complete within a reasonable amount of time.

4.2.2 Parameter selection, preparation and training methodology

The selection of input parameters from the synchronised CMU data are of utmost importance to ensure meaningful model predictions. Géron (2017) cleverly puts it as: *garbage in, garbage out*. Eliminating irrelevant features could reduce the risk of a model learning skewed associations. Pearson correlation analysis is a commonly used method to find underlying linear connections between parameters. A summarised version of the most noteworthy correlation coefficients are given in Table 4.2. The reported correlation numbers range between -1 and 1, where numbers close to -1 indicate a strong negative correlation and numbers close to 1 a strong positive correlation. A number near zero indicates no linear correlation. Zhang *et al.* (2019) conducted a Pearson analysis on data from a PC6 class cargo ship operating in the Arctic and reported similar correlation numbers between power-SOG, power-ice and SOG-ice variables.

The coefficients in Table 4.2 were calculated from data of the 2019-2020 relief voyage and includes all of the ship's operating modes. In cases where no ice or wave observation were recorded, a Not-a-Number (NaN) variable was introduced to indicate that no data was available for that specific parameter. Intuitively, shaft speed and propeller pitch parameters are directly responsible for the propulsion motor's output power, seeing that these are the mechanisms

Table 4.2: Pearson correlation matrix between most noteworthy CMU parameters.

	1.	2.	3.	4.	5.	6.
1. Stbd power	1					
2. Stbd shaft speed	0.804	1				
3. Stbd propeller pitch	0.797	0.598	1			
4. SOG	0.831	0.693	0.902	1		
5. Beaufort number	-0.021	-0.023	-0.022	-0.273	1	
6. Ice concentration	0.425	0.446	-0.384	-0.769	0.599	1

by which the torque and rotational speed are controlled. Therefore, the correlation coefficients indicate their relationship to output power is strong and positive. SOG contributes to the hydrodynamic load experienced by the ship, particularly during open water navigation. This is seen through the strong positive correlations between SOG, power and propeller pitch. Ice concentration has a positive correlation to power with a coefficient of 0.435, which is a weaker connection as reported by Zhang *et al.* (2019). This is odd due to the ship reaching its peak power level usually when in ice mode. In addition to the low temporal resolution of observed ice data, a very low percentage of recorded data was captured during ice navigation, see Figure 3.9, which could contribute to the lower than expected correlation coefficient. As per expectation, SOG and ice concentration has a relatively strong negative relationship.

Apart from choosing the correct input data, the scaling is also an important consideration. For example: the SOG varies between 0 and 18 kn while the wind direction has a range of 0-360 degrees. Typically, the performance of machine learning algorithms are severely hindered when input data has attributes with different scales (Géron, 2017). To correct this, the data is normalised to a range from 0 to 1. Scikit-Learn has a function called *MinMaxScaler* that automatically performs the normalisation by subtracting the minimum value from the current data point and then dividing it by the difference between the maximum and minimum values as shown in Equation 4.2.1. The machine learning algorithm can reliably compare normalised input data to each other and make useful connections.

$$x_i^* = \frac{x_i - x_{min}}{x_{max} - x_{min}} \quad (4.2.1)$$

The correlation matrix, Table 4.2, indicates that individual variables, which was obtained from various sampling methods and sources, could have an influence each other and on the performance of the ship. A systematic approach was adopted by training a baseline model on the minimum number of features which showed the highest correlation to power and then incrementally including other data fields in subsequent models. The effect that new parameters

Table 4.3: Training variables for the machine learning model.

Training variables			
Target	Machine data	Environmental data	
	CMU	CMU	Observations
Power	Port motor speed Starboard motor speed Port propeller pitch Starboard propeller pitch SOG	Wind direction Wind speed	Brash ice Ice concentration Ice thickness Beaufort number Wave direction Wave length

have on the model could then be evaluated against the baseline. Table 4.3 lists the features that were used to train the model and are separated into groups that indicate their origin. At first, the machine data from the CMU was used for training as per Table 4.3. Environmental data, both from the CMU and observations was later introduced incrementally. Steady state operating conditions are of concern and thus it is assumed that both motors operate at the same power levels. Therefore, the target training parameter is the average power between the port and starboard motors.

Mean absolute error (MAE), Equation 4.2.2, is used as a metric to determine the accuracy of the trained models. Here, x_i and y_i represent the predicted and target values of each sample of the n respectively.

$$MAE = \frac{\sum_{i=1}^n |x_i - y_i|}{n} \quad (4.2.2)$$

This metric calculates the absolute error between the predictions and targets over the whole sample set. It can be used to create a tolerance zone around the target values where the predictions are likely to be located. The unit of the MAE score is in kilowatt and indicates the average variance between the predicted and target power for the whole time period. A smaller MAE translates into a tighter tolerance zone which suggests improved predictive accuracy achieved from the model. All models for both open water and ice conditions are evaluated according to the MAE.

4.2.3 SVR model

A SVR model was trained with a radial basis function (RBF) kernel to model non-linearity. Discussed previous chapters, the problem with the SVR method is that the hyper-parameters need to be tuned to suite the desired application. The accuracy of untuned SVR models were significantly worse compared to the optimised models. Scikit-Learn's *GridSearch* function was used to find the best configuration of hyper-parameters. This process was computationally

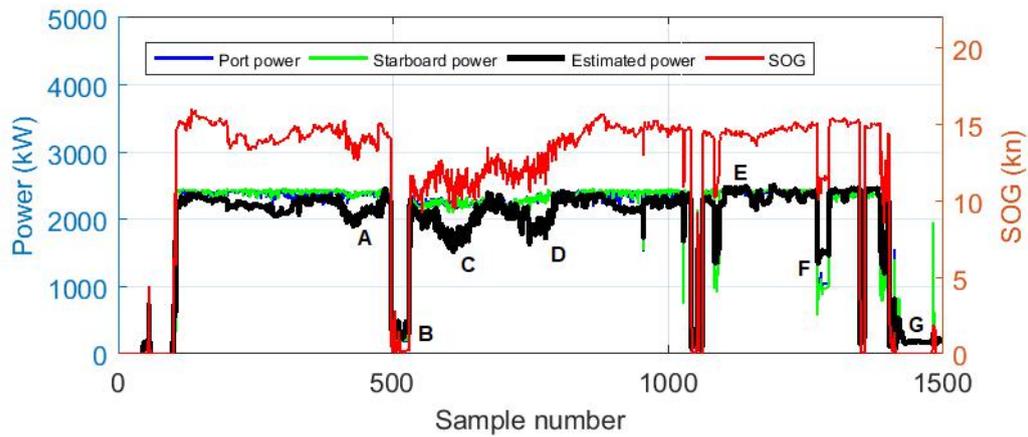
Table 4.4: MAE scores for SVR open water and ice navigation models.

Model	Tuned with GridSearch (Yes/No)	MAE (kW)
Open water model	No	581,08
	Yes	139,58
Ice model	No	388.63
	Yes	191.16

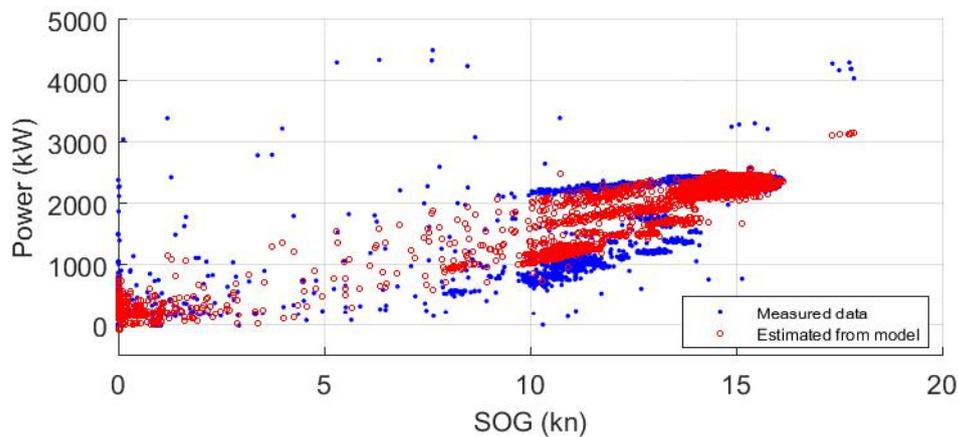
very expensive, especially when introducing large amounts of data. Typical training sessions lasted between 45-60 minutes and increased as more features were added to the input vector.

Table 4.4 lists the MAE scores for the best performing SVR models out of 54 experiments. The MAE for the given models are calculated from a separate test set that was unknown to the model during training. Notice the reduced MAE after tuning the hyper-parameters. The open water data spanned over most of the voyage and has a temporal frequency of roughly 3 minutes. On paper the SVR open water model performed well, achieving a MAE of 140 kW. The ice data is concentrated to areas where the ship operated near the Antarctic coast and, as hypothesised, proved to be much more volatile. A finer temporal resolution is preferable to show more detail of ice loads on the propulsion system. Therefore, the sampling rate of the ice data was increased by re-synchronising the data to obtain a temporal frequency of 90 seconds. Computational limits meant that this was the best achievable resolution to train SVR models. The SVR ice model did not achieve the same MAE score as the SVR open water model. However, this is expected due to volatile and low quantities of ice related data. The best MAE score that could be achieved for ice navigation was 191 kW.

The MAE score does not give any indication of the output's accuracy over time. Nor does it show if the output has overfitted to any of the parameters. Figure 4.1a shows a plot of SOG and power for the time period from 10 to 19 December 2019 during open water navigation. This model was trained on all open water related parameters listed in Table 4.3. The data originates from the test set which was not used during training. Thereby, the model's is forced to make new predictions from previously unseen data. The measured power outputs from both motors, in blue and green, are plotted alongside the predicted power output from the SVR model, shown in black. The predicted power follows the measured power, especially in cases where SOG is relatively constant (see points E and G). However, at points A, C and D the prediction seems to be very closely connected to SOG, suggesting that the model might have overfitted to this parameter. If SOG changes slightly then the prediction changes with it. Furthermore, the model performs adequately where the output is either at 2500 kW (E) or idling (B and G), but struggles in the intermediate power range (F). This is not acceptable and challenges the precision



(a) Power and SOG plot over time with estimated power from the open water SVR model (10 Dec. - 19 Dec. 2019).



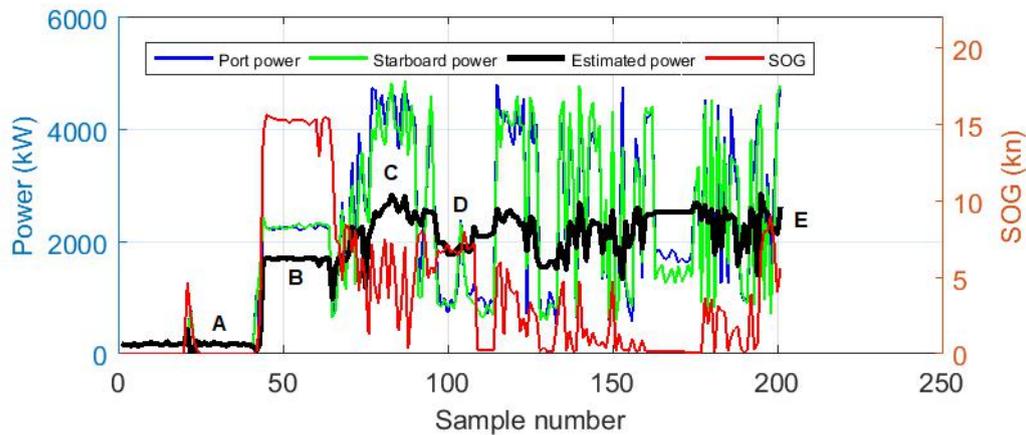
(b) Scatter plot of open water model predictions.

Figure 4.1: Predictive performance of the SVR open water model.

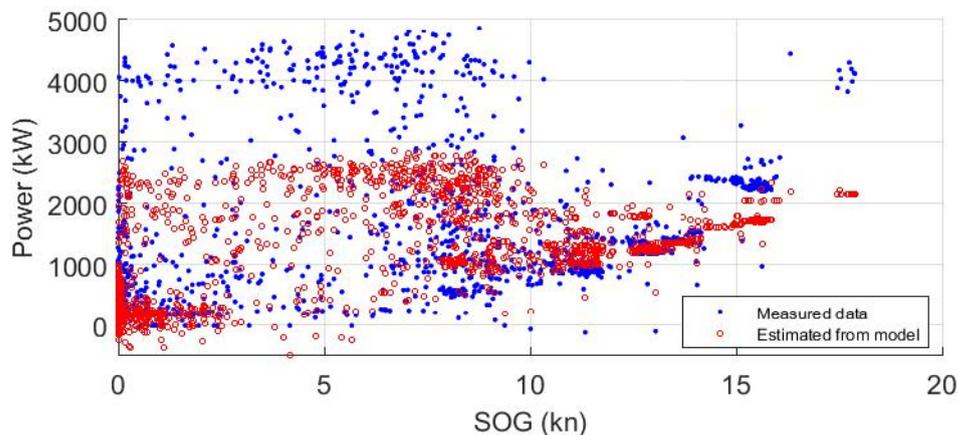
and robustness of the model.

A scatter plot, Figure 4.1b, paints a picture of how well the predicted power matches the actual measured data. The model's predictions, shown in red, follows the expected general non-linear power versus speed trend but struggles to describe the width of the actual power output, which is most probably caused by variations in environmental factors. The most accurate estimates lie in the regions of 2500 kW at 15 kn and when the ship is idling. This provides insight into why the estimated power in Figure 4.1a does not perform satisfactory.

The second challenge is to characterise the dynamic behaviour typical of ice navigation. A brief 15 hour period of ice manoeuvres are plotted in Figure 4.2a showing the SOG, actual power outputs of both motors and the predicted power from the SVR model. This model was trained on all ice related parame-



(a) Power and SOG plot over time with estimated power from the ice navigation SVR model (29 Dec. - 30 Dec. 2019).



(b) Scatter plot of ice model predictions.

Figure 4.2: Predictive performance of the SVR ice model.

ters listed in Table 4.3. It appears that the model's predicted power does follow its measured test equivalent over the course of time. However, the model is unable to describe the high (C) and low (D) peak values. It is as if the model's response is limited to an average power output of around 2500 kW. An offset of roughly 500 kW between the predicted and actual power is visible at point B where the graph is fairly constant. The error is amplified in the region around point C to almost 2000 kW. In the scatter plot, Figure 4.2b, the model is unable to make any predictions above 2500 kW. The cumulative effect of adding up the error over time could translate into a significant mismatch between projected and actual power demand. For this reason, the SVR model is not acceptable for characterising the power requirement for ice navigation.

The SVR model performed well as a first step into understanding how measurements and observations from the SAAII can be used in a predictive capacity.

The open water model was able to make adequate estimates but was limited within two operating regions, either at cruising speeds or when stationary. The ice model was crude and lacked the ability to characterise the dynamic nature of ice loadings on the propulsion system. The SVR learning process was also computationally expensive. As more variables were introduced the training time increase up to a point where it was intractable. In some instances training took a few hours to complete.

4.2.4 Neural network model

The alternative approach to the problem was using ANN. Again, separate models for open water and ice modes were trained. A feed-forward neural network (FFNN) was created consisting of one input layer and two hidden layers. The *rectified linear unit* (relu) function, Equation 4.2.3, was used as the activation function in the first two layers. No activation function was applied to the output layer, leaving the model free to predict values in any range. The optimiser from the *backpropogation* algorithm was instructed to minimise the mean square error (MSE) between the predicted values and the actual data in order to train the network.

$$f_{relu} = \begin{cases} x & x \geq 0 \\ 0 & x < 0 \end{cases} \quad (4.2.3)$$

Both models were trained on the same 2019-2020 data set as the SVR models with the same variables as listed in Table 4.3. The training process for neural networks completed much faster than the SVR method, typically within a 10 minute time frame. Figure 4.3 shows the convergence plots for the open water and ice navigation models. Both plots show the MAE as they progress through optimisation iterations. Each iteration is known as an epoch. The model for ice navigation, Figure 4.3b, took longer to converge compared to the open water model, Figure 4.3a. This is most likely due to the unpredictable nature of the ice data. Both models were trained for 80 epochs where, for the first few iterations, the MAE score improves dramatically and flattens out later on. MAE does not significantly improve beyond 80 iterations, which indicates that 80 epochs would be the around optimum number of training iterations to achieve the best fit for the present model proposal.

The MAE scores are calculated based on the test data. In Table 4.5 the MAE scores of both ice and open water models are listed. MAE for the open water and ice models are 46 kW and 117 kW respectively. This is a significant improvement in comparison with SVR's performance. When the range of the ship's power output of around 4500 kW, the relative error is small at about 2% for both open water and ice scenarios. This realisation enables the FFNN model to make predictions with reasonable confidence and reliability.

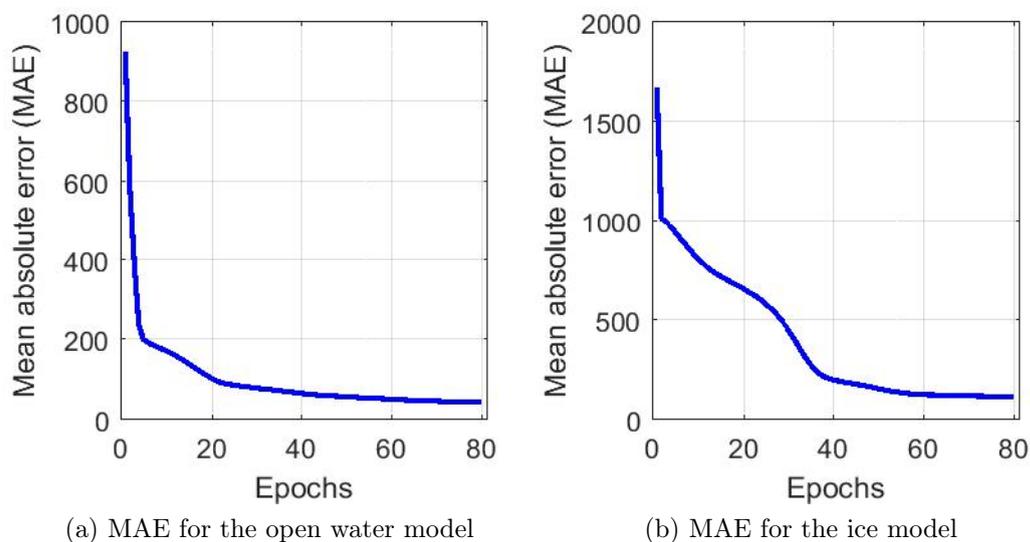
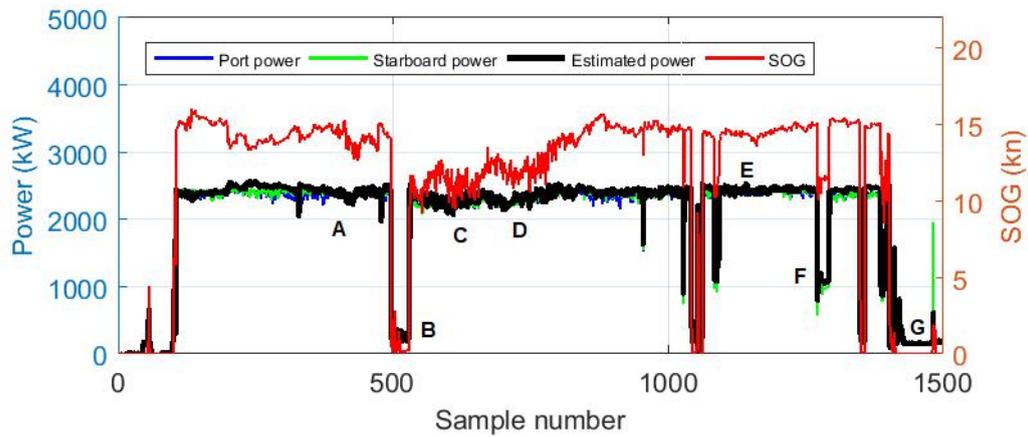


Figure 4.3: Convergence of MAE on test data during training iterations of the open water and ice navigation neural networks.

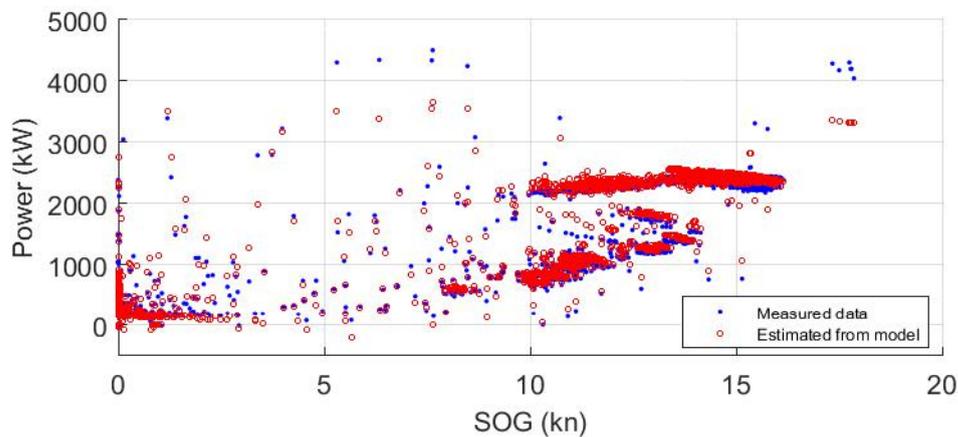
Table 4.5: MAE scores for FFNN open water and ice navigation models based on the 2019-2020 test data.

Model	MAE (kW)
Open water model	46.47
Ice model	117.42

A plot of power and SOG over time for the FFNN open water response is given in Figure 4.4a for the same time period as with the SVR model. The labels, A to G, mark the same points of interest noted for the SVR response, Figure 4.1a. Comparing the two figures, it is illustrated that the response of FFNN follows the actual power much better than the SVR model. Recall that at points A, C and D, the SVR model deviated significantly from the targeted response and seemed to have overfitted to SOG. This is not the case for FFNN as the model does not seem to be overly affected by the variations in SOG at A, C and D. The FFNN is better at predicting power levels in the intermediate range (F). SVR could predict the overall trend of the power output, but at intermediate power levels displayed an offset error of 500 kW or more. The FFNN does not struggle with this problem and gives accurate predictions for low (B and G), intermediate (F) and high power levels (A, C, D and E). The scatter plot of power versus SOG also proves this point, Figure 4.4b. The estimated power, shown as red circles, overlaps the measured power much better compared to the SVR model, Figure 4.1b. The predicted power from FFNN almost completely covers the measured test data, which indicates that the model was able to



(a) Power and SOG plot over time with estimated power from the open water FFNN model (29 Dec. - 30 Dec. 2019).

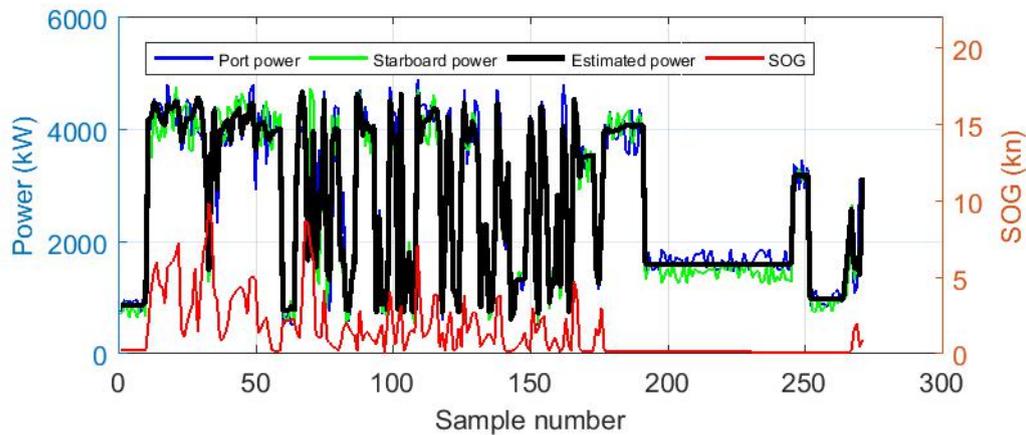


(b) Scatter plot of open water model predictions.

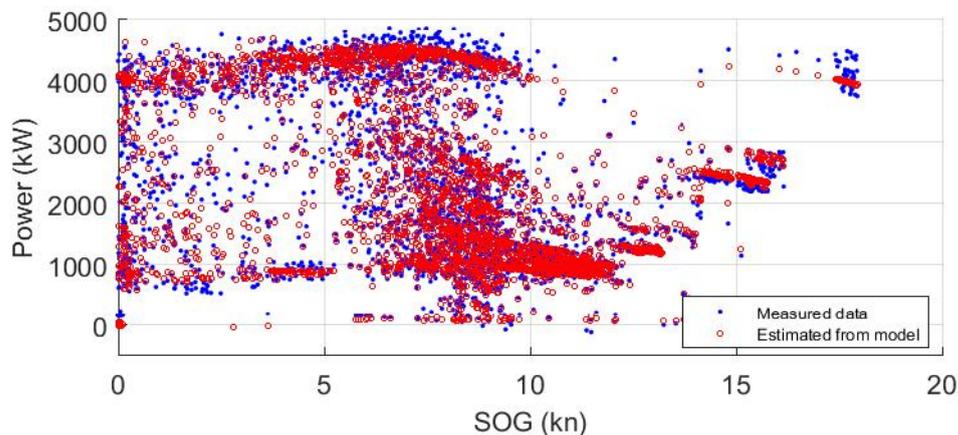
Figure 4.4: Predictive performance of the FFNN open water model.

introduce variance caused by extraneous ocean and atmospheric conditions.

The time to train a neural network model was significantly faster in comparison with the SVR model. Depending on the volume of data, training time was typically 4-8 minutes for the FFNN compared to the 40 to 60 minutes of SVR models. For this reason, it was still economical to train the FFNN on larger volumes of data which would have been impractical for SVR. The training data for ice mode cases were thus sampled at a much higher rate, resulting in temporal resolution of 15 seconds. A huge improvement to the 90 second resolution available to SVR. Figure 4.5a shows the power and SOG response over time during ice navigation. The power output oscillates heavily and is inversely related to SOG. Peaks from the power variable correspond to SOG valleys and vice versa. This is an indication of the ship interacting with heavy ice, driving the power output to high levels with little increase in speed.



(a) Power and SOG plot over time with estimated power from the ice navigation FFNN model (01:00 to 05:00 on 30 Dec. 2019).



(b) Scatter plot of ice model predictions.

Figure 4.5: Predictive performance of FFNN ice model.

Describing this behaviour is particularly difficult as it does not follow a typical polynomial model as for open water. Observations of ice thickness, ice concentration and ice loading are very limited and subjective, making the use of an analytical or numerical model unrealistic. However, the FFNN does characterise the power output very well by following the fluctuating power parameter and correctly estimating the peak values. The scatter plot for the power versus SOG, Figure 4.5b, show how well the FFNN was able to make predictions. Estimates from the model, shown as red circles, cover the whole spectrum of power levels from ice data, especially at mid-range speeds where higher power levels are not driven by hydrodynamic forces between the hull and water, but from ice interactions instead. By comparison to the SVR model, Figure 4.2b, the FFNN is superior for this application.

Both ice and open water neural network models showed improved performance

in comparison with the SVR models. On inspection it does not seem that the models overfit to any of the input parameters. The MAE was much lower for FFNN compared to SVR with the FFNN models displaying accurate and reliable predictive ability for both open water and ice conditions.

4.3 Model validation

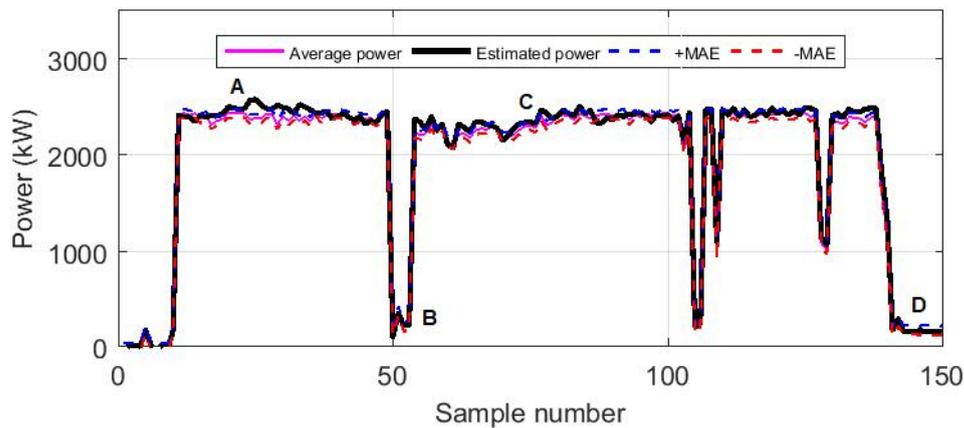
The FFNN models achieved much better MAE scores and displayed superior predictive ability compared to SVR on the test data sets. Therefore, only FFNN will be considered from for model validation. Although the models achieved good performance, it must still be verified on data that is unique and not used during training or testing. The idea is to test the model on completely new scenarios to find its limitations.

4.3.1 Antarctic relief voyage (2019-2020)

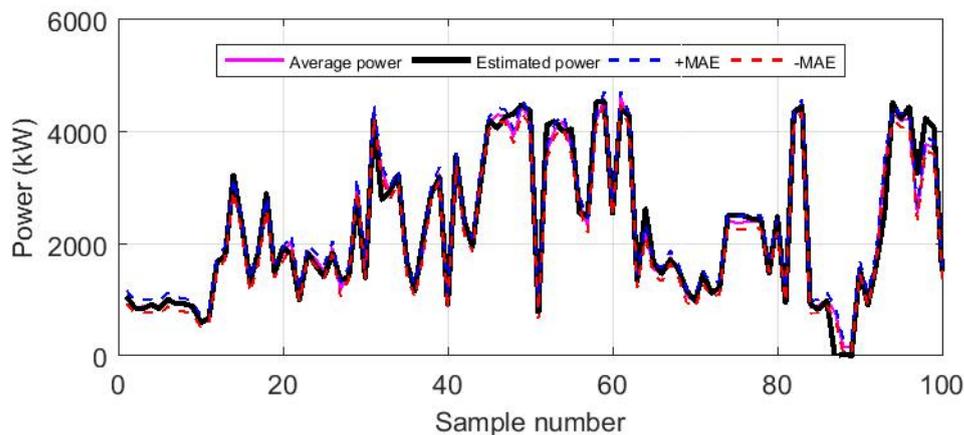
A small portion of the 2019-2020 test data is plotted for open water passage, Figure 4.6a, and ice navigation, Figure 4.6b, to compare the actual average power to the model's projections. The upper and lower limits of the MAE are indicated by dotted lines around the actual average power to show how closely the average predicted power follows the target value. As mentioned, MAE only shows the mean error and thus it is expected for the predicted value to travel in and out of these limits. A glimpse into the limitations of the model is provided when this happens and indicates where the model can be improved.

Points of interest are indicated as A, B, C and D on Figure 4.6a. At point A the predicted value moves out of the MAE band but returns quickly. At points B, C and D the predicted value is almost exactly the same as the target value. Taking into consideration that almost half of the data is from stationary operations may suggest that the MAE score is biased towards low speed and low power scenarios. Overexposing the model to stationary points may not yield a true reflection of the error during open water navigation. Power output during stationary operations can be predicted with relative ease and can create a false sense of accuracy for open water estimates. The ice model would not have this problem as the training data did not contain any stationary data points. Nonetheless, with fuel consumption estimations we are not concerned with local inaccuracies as identified at A. Instead, it is required that the model be consistently accurate over the whole time period, which is evident in the FFNN predictions for open water.

The power predictions for the FFNN ice model is presented in Figure 4.6b. The volatile nature of the power output is estimated very well. However, upon critical reflection it may not be evident that this model is robust as yet. The test data for 2019-2020 is very similar to the training set used to develop the



(a) Tolerance zone indicating the upper and lower limits of the MAE for open water navigation during the 2019-2020 Antarctic relief voyage. $MAE = 46,5 \text{ kW}$



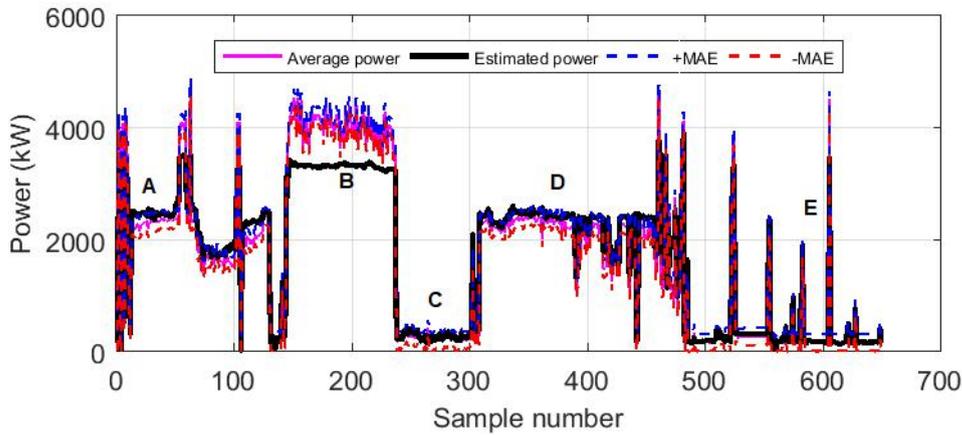
(b) Tolerance zone indicating the upper and lower limits of the MAE for ice navigation during the 2019-2020 Antarctic relief voyage. $MAE = 117.43 \text{ kW}$

Figure 4.6: Predictions from FFNN models on 2019-2020 test data.

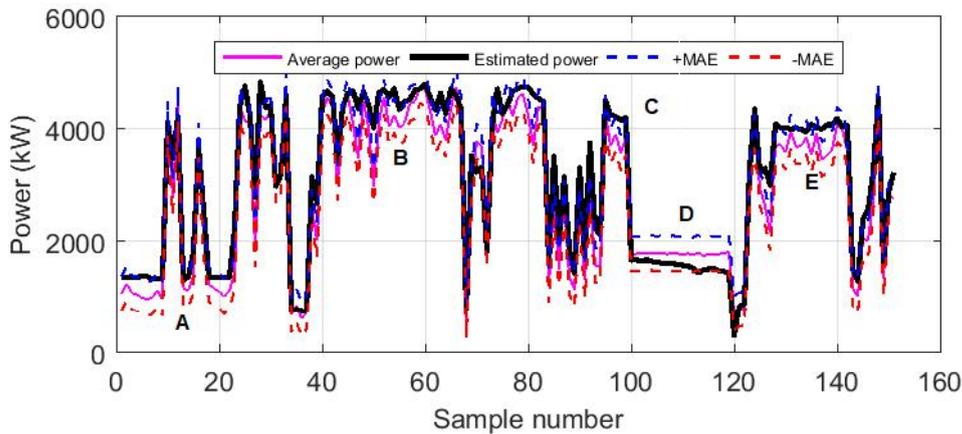
model. The extraordinary fit seen in Figure 4.6b could be attributed to the model *memorising* the expected outcomes and would therefore show excellent predictive ability. Further validation is required on data sets outside of the 2019-2020 relief voyage data for both ice and open water models to conclude whether training was successful.

4.3.2 Antarctic relief voyage (2017-2018)

Data from the Antarctic relief voyage of 2017-2018 was used to create predictions for the FFNN model, which was trained on the 2019-2020 relief voyage data. A small section of the open water predictions is plotted in Figure 4.7a with a new MAE tolerance band specific to this data set shown as dotted lines. The predicted output does follow the actual average power very well



(a) Tolerance zone indicating the limits of the MAE for open water navigation during the 2017-2018 Antarctic relief voyage. $MAE = 153.13 \text{ kW}$



(b) Tolerance zone indicating the limits of the MAE for ice navigation during the 2017-2018 Antarctic relief voyage. $MAE = 308.67 \text{ kW}$

Figure 4.7: Predictions from FFNN models on 2017-2018 test data.

with slight overestimation at points A, C, D and E, but otherwise mostly staying within the MAE band. The ship was operated slightly different for a short period of time during the 2017-2018 voyage. To achieve higher speeds the ship was operated at power levels of 4000 kW in open water, which is a configuration usually reserved for ice passage. This mode of operation was not present in the training data and explains why the model has poor predicting ability at point B. Without any examples of this behaviour in the training set, the model cannot learn from these extreme trends. The only way to improve the predicting ability at point B is to include this in the training process. Other than that, the model's predictive ability is acceptable around A, C, D and E, which does not show any wild trends that might suggest that the model has overfitted. Another aspect to take into account is that wave observations were not conducted during the 2017-2018 voyage which results in a lack of available

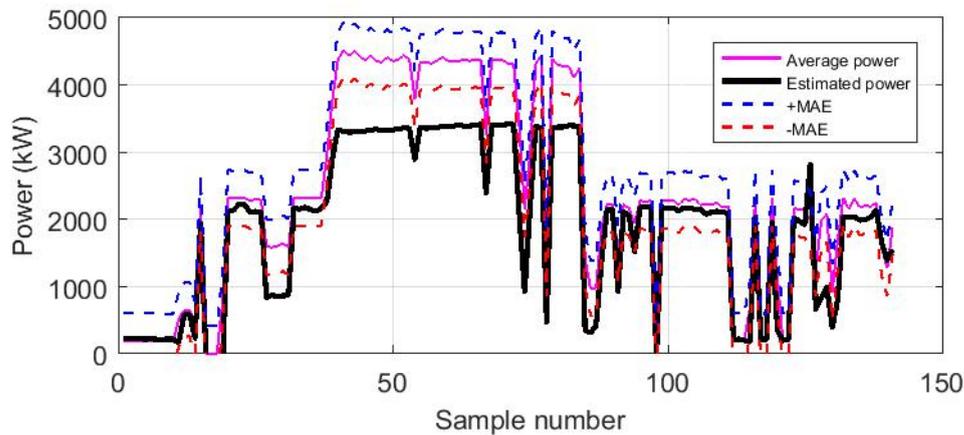
input variety. It is not apparent that the model is unaffected by this, which would suggest that knowledge of the ocean conditions are currently not the primary driver for the model's estimates.

Figure 4.7b shows a sample of the power over time plot for the ice data from the 2017-2018 relief voyage. The estimated and actual average power is plotted along with the MAE tolerance bands calculated specifically for this data set. At first glance the model performs very well, showing the ability to accurately predict oscillations in the power output (points A, B, D and E). The predictions also mostly stay within the MAE band, although it is a looser fit compared to the 2019-2020 test data. At some occasions, the model does get estimations wrong, such as at the peaks outside the MAE band at point C. However, as previously stated, local inaccuracies will not adversely affect the models ability to estimate energy use over time. It is more important to consistently follow the actual power, which is true for the model predictions in Figure 4.7b. The limited quantity is not yet apparent as the 2017-2018 data was very similar to the 2019-2020 test set. The model could still rely on *memorisation* to make predictions and therefore the success of the ice model is still inconclusive.

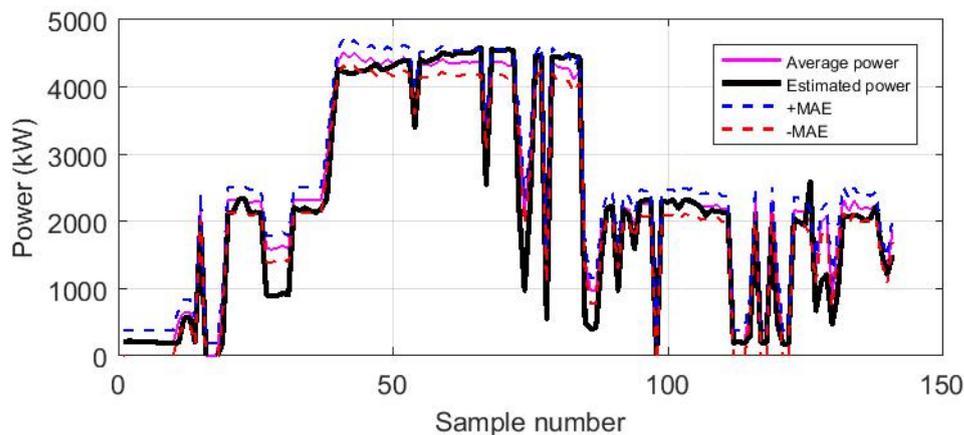
4.3.3 Weddel sea expedition (2018-2019)

A third data set from the Weddel sea expedition (2018-2019) was available for model validation. The effect of training a model on representative data is illustrated in the two plots in Figures 4.8a and 4.8b. At the beginning of the Weddel sea expedition the ship was operated at 4000 kW to achieve a higher SOG. The FFNN model trained on the 2019-2020 data was not exposed to high-speed manoeuvres which constrained the domain of the model. This effect is visible in Figure 4.8a where the maximum power estimate of roughly 3200 kW fall significantly short of the actual measured power, which was in the region of 4200 kW. In contrast, Figure 4.8b plots the estimated power for the same data but with an FFNN model trained on both 2017-2018 and 2019-2020 data sets. As previously mentioned, the 2017-2018 relief voyage contained a portion of high power open water operations similar the that experienced on the Weddel sea expedition. By widening the range of the training data has a positive effect on the FFNN model which can now predict the powering performance of high-power and high-speed open water operations. The improved model achieves a better MAE score of 191 kW and predicts consistently within the MAE bands indicated in Figure 4.8b. This verification instils confidence that the FFNN model trained of both relief voyages will be able to reliably predict the output power for a wider range of operating conditions.

The limitations of insufficient data become apparent when examining the power predictions for ice navigation data from the Weddel sea expedition. Figure 4.9 shows a portion of the data during ice navigation. The ice observations from the Weddel sea expedition were recorded in a slightly different



(a) Section of open water data from the Weddel sea expedition. Prediction from FFNN model trained on 2019-2020 relief voyage. $MAE = 412 \text{ kW}$



(b) Section of open water data from the Weddel sea expedition. Prediction from FFNN model trained on 2017-2018 and 2019-2020 relief voyage. $MAE = 191 \text{ kW}$

Figure 4.8: Predictions from FFNN models on Weddel sea data (2018-2019).

manner with a temporal resolution of 5 minutes. This presents completely new data to the FFNN model which meant it could not rely on memorised patterns to make predictions. Very poor predictive performance is observed in Figure 4.9. Apart from the model's poor fit to the measured power, the MAE score is the worst thus far at 739 kW. This is an unacceptable result that questions the robustness of the FFNN ice model when applied to new data.

4.4 Chapter summary

Two machine learning algorithms were trained on the CMU data from the 2019-2020 Antarctic relief voyage. A SVR and FFNN were proposed for experimentation. The correlations between the various input parameters were

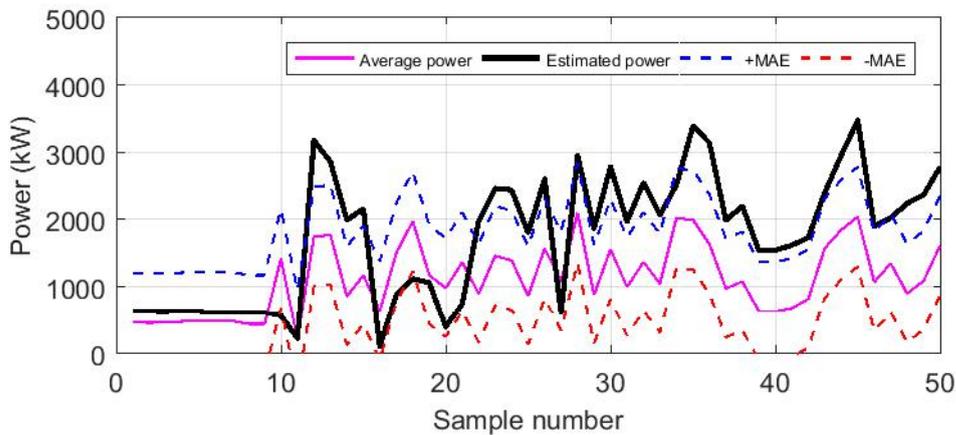


Figure 4.9: Tolerance zone indicating the limits of the MAE for ice navigation during the Weddel sea expedition. $MAE = 738.62 \text{ kW}$

calculated and discussed to find the strongest connections to the output power. This led to the identification of 13 variables suitable for training either the ice or open water models, Table 4.3. A training methodology was discussed to explain the importance of scaling the input data. MAE was defined as a metric to measure the performance of the SVR and FFNN models.

A SVR model was trained and evaluated on a test portion of the 2019-2020 CMU data. Optimising the SVR hyper-parameters was computationally expensive with the training process taking a significant amount of time to complete. The open water SVR model overfitted to the SOG parameter, showing an inability to make accurate predictions over the voyage period. In addition, SVR was unable to characterise the volatility of ice data, Figure 4.1. The ice model could only predict in the general area where it would be expected to find mean power estimates. Considering both flaws, the models were found to be unacceptable and would be ineffective as a predictive tool to estimate energy use over time.

FFNN proved to perform much better than the SVR. The training process was much simpler and completed in a shorter time frame. Once again a test portion of the 2019-2020 CMU data was used to evaluate the model's initial performance. Both ice and open water models were able to predict the output power within reasonable accuracy. At first, the ice model particularly impressed by predicting the volatile ice data accurately in terms of both frequency and amplitude. The FFNN models were validated on CMU data from the 2017-2018 Antarctic relief voyage and proved to be effective on data previously unknown to the models but lacked the ability to estimate power demand for high-speed operation. The model was trained again on both 2017-2018 and 2019-2020 relief voyages to increase the range of the available training data. The improved model was successfully validated on an open water portion of data from the

Weddel sea expedition. The insufficient quantity of ice-related data proved to be troublesome. The validation process dismissed the validity of the ice model which is not sufficiently trained for general applications.

Chapter 5

Voyage cost optimisation

5.1 Introduction

Data-based models describing the dynamic behaviour of the SAII in a range of weather and operating conditions presents an opportunity to examine the historical response of the vessel. This *hindsight* perspective that a model brings is useful in a design review or maintenance planning context and may provide insight into how the ship responds to specific operating environments. On the other hand, an optimisation problem challenges the robustness and accuracy of the data-based models. It requires the model to have *foresight* and know the effect that some input parameters have on others.

A proof of concept optimisation exercise is presented to demonstrate the value of a data-driven model within a route planning and financial context. A comparison between classical optimisation and evolutionary algorithms are discussed. Evolutionary methods are recommended and the costs for a fictitious voyage optimised using particle swarm optimisation (PSO).

5.2 Theoretical overview of optimisation methods

Mathematical optimisation is the process of formulating and finding the optimum solution to a problem subject to constraints. The origin of a problem could stem from a practical or theoretical source. General classical constrained optimisation can be mathematically expressed as (Snyman, 2005):

$$\begin{aligned}
 & \underset{\text{w.r.t. } \mathbf{x}}{\text{minimize}} && f(\mathbf{x}), \mathbf{x} = [x_1, x_2, \dots, x_n]^T \in R^n \\
 & \text{subject to} && g_i(\mathbf{x}) \leq 0, i = 1, 2, 3 \dots, m. \\
 & && h_j(\mathbf{x}) = 0, j = 1, 2, 3 \dots, r.
 \end{aligned} \tag{5.2.1}$$

Here $g_i(\mathbf{x})$ refers to inequality constraint functions and $h_j(\mathbf{x})$ to equality constraint functions. In the context of this study, the objective function $f(\mathbf{x})$ represents the voyage costs of the SAII with \mathbf{x} a vector of input parameters used for model training, Table 4.3.

The objective function $f(\mathbf{x})$ utilises an FFNN model as part of the formula to calculate costs of travelling over a set distance. Estimating curve characteristics, such as rates of change or partial derivatives, are not as simple for data-driven models compared to a conventional mathematical function. The model only has knowledge of the ship dynamics specific to the available training data. The possible implications of this range from multiple minima points to significant numerical noise. The subject literature does not always agree on the correct approach to such problems. Snyman (2005) argues that, despite the presence of inhibiting factors, gradient-based methods are appropriate and more suitable compared to computationally expensive evolutionary approaches. The judicious use of gradient-based methods can solve problems with numerical noise and multiple minima much faster (Snyman, 2005). On the other hand, Lindfield and Penny (2017) emphasize that evolutionary methods are applicable to non-linear problems with particular reference to problems with multiple minima points. Lindfield and Penny (2017) note that PSO would find the global minimum where gradient-based methods could fixate on local minima. PSO was successfully applied in other studies where data-driven modelling was used (Wang *et al.*, 2018). This previous work provides a strong argument for its use within the current context.

Classical optimisation with a gradient-based approach was derived from mathematics. The steepest descent method is well known and is typically used for unconstrained optimisation on its own. When subject to constraints, steepest descent can still be used but within the augmented Lagrangian method, which converts the problem into an unconstrained problem with Lagrangian multipliers (Snyman, 2005). Gradient-based methods were attempted on the cost optimisation problem, discussed in the subsequent subheading, but required knowledge of the first derivative of the objective function. This was impossible to calculate analytically. Numerical differentiation of the FFNN model with the use of finite difference approximations was possible. However, it remained difficult to discern between numerical noise and the approximate derivatives. Preliminary results from this approach were impractical, returning minimum costs as a large negative value. If this were true then the ship would generate profit by just being out at sea. Along with a tedious implementation and highly variable results, it was decided that classical methods would not be used in further experimentation.

Evolutionary optimisation algorithms draw inspiration from natural phenomena. PSO and genetic algorithms (GA's) are well established techniques in the evolutionary optimisation toolkit. These methods rely on randomised guesses

that use the collective knowledge of the whole population to migrate to either a local or global minimum. PSO draws inspiration from birds flocking and is well suited to a non-linear problem with multiple local minima where the global optimum is sought (Lindfield and Penny, 2017). PSO was selected as the method for experimentation due to its simple implementation and the fact that it does not require knowledge of the slope of the objective function.

5.3 Particle swarm optimisation in the open water environment

A basic optimisation exercise could illustrate the application value of data-driven technique as an integral part of the digital twin concept. The purpose of cost optimisation is to find the optimum configuration of the ship's control surfaces to minimize the costs associated with travelling a predefined distance. An adjustable input vector is defined with variables that describe the configuration of the ship's control surfaces: shaft rotational speed and propeller pitch. The PSO is therefore conducted over a two dimensional search space. The influence of ocean and atmospheric conditions are introduced as external, manually adjustable parameters which describe the contours of the search area.

5.3.1 Defining the objective function $f(\mathbf{x})$

There are many factors which contribute to the costs related to keeping a vessel such as the SAAIL in operation. Not only does the efficiency of a vessel play a significant role but also the overhead costs such as maintenance, crew wages and other items not associated with the direct power output from the engines. A simple cost structure is proposed to describe the operating expenses to a destination over change in time

The overhead costs can be broken down into two components: recurring costs for each day spent out at sea and initial fixed costs incurred before departure. The sum of the two equates to the total overhead costs which can be expressed as a linear function of time.

$$A(\Delta t) = m\Delta t + b \quad (5.3.1)$$

In Equation 5.3.1 the total fixed cost is expressed as a linear function of the time difference, Δt , to the destination. $A(\Delta t)$ is the total overhead costs, with m the daily expenses at sea and b the initial expenditure before departure. Daily expenses include crew wages, energy for heating, electricity generation and food provisions. Initial costs entail tasks required before departure - such

as harbour fees, maintenance costs, fumigation of the ship and land-based logistics.

Fuel consumption is linked to power demand of two propulsion motors over defined time period, expressed in the unit kWh. The cost for an amount of fuel is given by:

$$B(\Delta t) = 2kP(\mathbf{x}, v(\mathbf{x}), \mathbf{W})\Delta t \quad (5.3.2)$$

In Equation 5.3.2 the total fuel costs, $B(\Delta t)$, is proportional to the power output over a time difference, Δt . The constant of proportionality, k , scales the product of power output and time difference to obtain the correct relation between fuel costs and energy use. $P(\mathbf{x}, v(\mathbf{x}), \mathbf{W})$ denotes the power output at a specific shaft and propeller configuration, SOG and weather environment. This value is estimated by the FFNN model which was trained on data from both 2017-2018 and 2019-2020 relief voyages. Initial PSO results highlighted a flaw in the model's power predictions when the control surfaces and SOG could be adjusted independently from each other. The FFNN model could not describe the physical connection between SOG, propeller configurations and power demand. PSO found an optimum configuration for this model but required the ship to sail at full speed and low power levels by keeping shaft rotational speed and propeller pitch to a minimum - an impossible configuration to sustain. As a consequence, a secondary model was trained from the two relief voyage data sets, but with SOG as the main output target. The result of this SOG model was then processed as an input to the power model. The dependency of both SOG and power on the control settings were now enforced. \mathbf{x} represents the shaft and propeller configurations for the port and starboard side. $v(\mathbf{x})$ is the result obtained from the SOG FFNN model with \mathbf{x} as the input. The vector that contains relevant weather information is represented by \mathbf{W} .

Combining and factorising Equations 5.3.1 and 5.3.2 results in the total costs as a function of the time difference between departure and arrival.

$$C(\Delta t) = b + (m + 2kP(\mathbf{x}, v(\mathbf{x}), \mathbf{W}))\Delta t \quad (5.3.3)$$

Realizing that $\Delta t = \Delta d/v(x)$, and substituting into Equation 5.3.3 yields the objective function $f(\mathbf{x})$:

$$f(\mathbf{x}) = b + \frac{\Delta d(m + 2kP(\mathbf{x}, v(\mathbf{x}), \mathbf{W}))}{v(\mathbf{x})} \quad (5.3.4)$$

An exaggerated graphical representation of the cumulative costs, as described by the objective function, is presented in Figure 5.1. Voyage time, Δt , is represented on the x-axis with the costs on the y-axis. The maximum obtainable

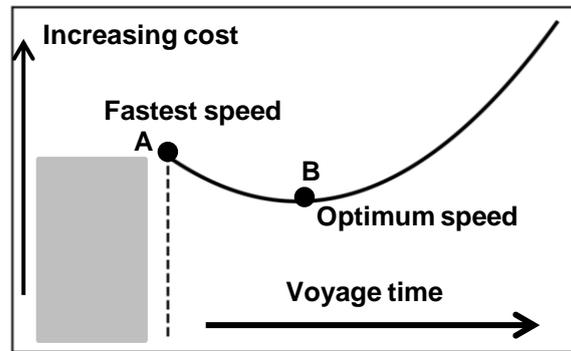


Figure 5.1: Diagram of proposed cost structure.

SOG dictates how quick the ship can reach the destination. Point A indicates the quickest attainable voyage time that can be achieved by operating the ship at full speed. The shaded area represents a region that falls outside of the ship's performance capabilities. The relationship between power and SOG for open water passage is non-linear, which leads to the possibility of locating an optimum point at B. Between A and B fuel consumption drives rising costs, whereas recurring overhead expenses become the dominant contributor as voyage times increase beyond point B. In addition, the implications from uncertainties caused by ocean and atmospheric conditions could result in a shifting optimum point which highlights the need scenario-specific cost optimisation.

5.3.2 Defining the search domain

The search space must be defined in accordance to the parameters that a ship's crew can control. The input variables are also limited by the domain over which the model was trained. Apart from route planning strategies and protocols, weather conditions are beyond the control of any ship operator. Four main assumptions have to be made in order to simplify the overarching optimisation problem:

1. Adjustable parameters are limited to the control surfaces of the ship: shaft rotational speed and propeller pitch. This forces PSO to exclusively adjust parameters that the crew have control over.
2. Weather conditions are stochastic and time series dependent, meaning that the current weather is influenced by the previous conditions. To simplify for the purpose of the optimisation exercise, it was assumed that weather would be constant over the whole distance in the optimisation interval, Δd .

Table 5.1: Variables and constant selected for PSO.

	Description	Domain	Symbol
Input variables	Shaft rotational speed	90 - 140 (rpm)	x_1
	Propeller pitch	33 - 100 (deg)	x_2
Constants	Voyage distance	3000 km	d
	Hourly cost	R 10 000	m
	Initial cost	R 100 000	b
	Cost per kWh	R 2.80	k

3. The pitch of both the starboard and port propellers are the same and similar to how it would be configured in practice under steady state conditions.
4. The rotational speed of both the starboard and port side drive shafts are identical and similar to how it would be configured in practice under steady state conditions.

By means of the above-mentioned assumptions, the problem is simplified and focusses only on parameters that the crew can control. Some of these assumptions may skew the accuracy of the calculated costs but should still allow the FFNN model to converge to an optimum recommended SOG. The domain of the input vector is also limited to regions where the model is best defined. These regions fall into the areas where the most data was available in the training set. The acceptable range of the shaft rotational speed and propeller pitch are given in Table 5.1.

The assumed values of constants in the objective function $f(\mathbf{x})$ is listed in Table 5.1. The constant relating to fuel consumption cost, k , was estimated based on correspondence with ship officers (Ligthelm, 2020) and a diesel price of R13.00 per litre. The reader is referred to Appendix C for the calculation of costs listed in Table 5.1. These costs are for illustrative purposes only and do not reflect actual voyage expenses. The distance of 3000 km was arbitrarily selected.

5.4 Results

The objective function $f(\mathbf{x})$ was optimised with a *global best position* PSO search with respect to the vector \mathbf{x} . Speed optimisation was done on the open water and ice models. The PSO algorithm typically converged within 50 iterations is illustrated by the convergence plot, Figure 5.2, where the x-axis shows the number of iteration and the y-axis represents costs in South African Rands (ZAR). As an additional pilot exercise, PSO was used to assist with a hypothetical route selection scenario.

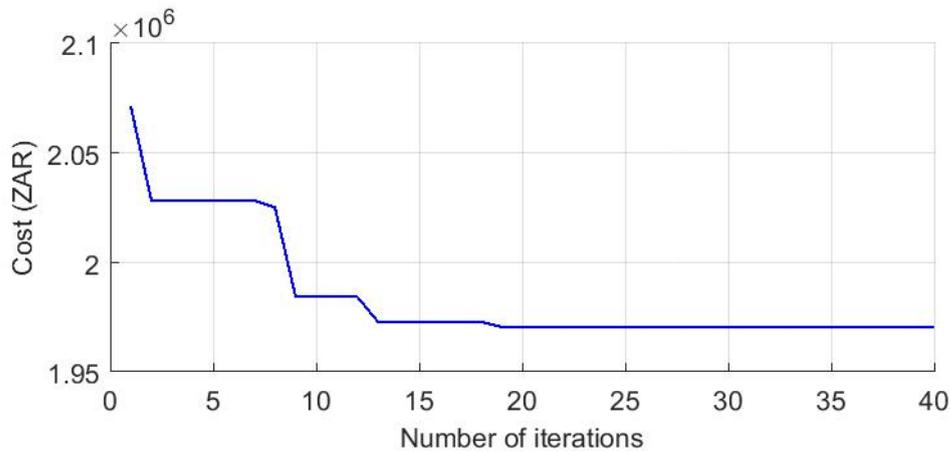


Figure 5.2: Convergence of the PSO algorithm on the objective function $f(\mathbf{x})$.

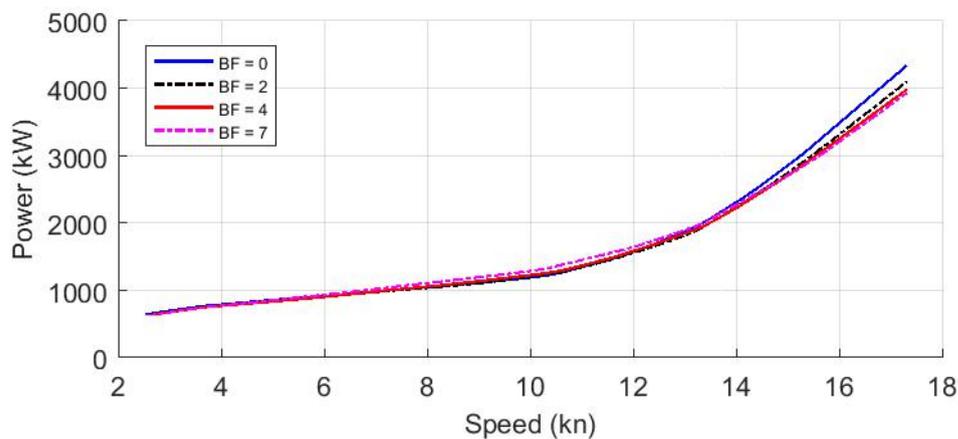


Figure 5.3: The effect of the Beaufort number on the power versus speed curve.

5.4.1 Open water model - speed optimisation

A good starting point for speed optimisation is to observe the power-speed relationship predicted by the model. Yoo and Kim (2018) demonstrated the usefulness of their data-driven model by plotting estimated power versus speed curves under various wave and wind conditions associated with the Beaufort scale, see Chapter 2. Similarly, the power versus speed curves for various Beaufort numbers, determined from the open water model, are presented in Figure 5.3. In-line with what is expected, a non-linear relationship is visible between the two variables with the slope increasing rapidly at speeds exceeding 12 kn. Compared to scatter plots of the data, Figure 3.6, the predicted curves have eliminated noise originating from other extraneous factors. This provides a sense of confidence that predictions from the data-driven models are in line with the ship's responses. Unfortunately, the curves in Figure 5.3 do not show significant variance between the four Beaufort numbers which could be a

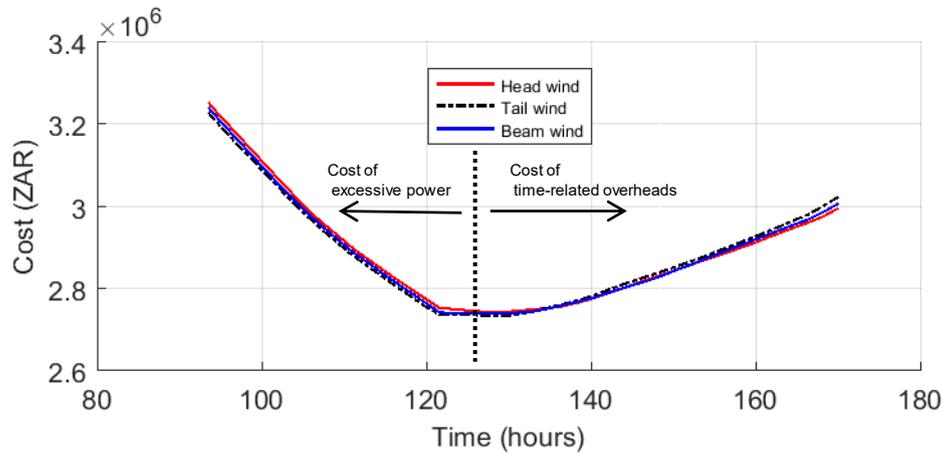


Figure 5.4: The effect of wind direction on the voyage costs.

result of poor distribution of sea state related data in the training set. Beaufort numbers were only recorded during daylight hours at typical cruising speeds during the 2019-2020 relief voyage. As a consequence, some portions of the operating domain is paired with recorded sea state data while others are not. It does not necessarily mean that the model cannot account for weather, but rather indicates that the Beaufort number on its own does not have a great influence on the prediction based on the current training data. Nonetheless, the Beaufort number only represents a general description of environmental conditions. As discussed previously, Chapter 4, wind and wave related data should be able to introduce variance which allows for distinctions between levels of weather severity.

A similar observation can be made regarding the effect of wind direction on the powering performance. Figure 5.4 shows the costs over voyage time plots for three distinct wind directions: head, beam (side), and tail winds. The voyage costs were calculated according to Equation 5.3.4 with the parameters set out in Table 5.1. A head wind intuitively produces the most resistance to motion and accordingly is the most expensive out of the three discrete directions. However, from Figure 5.4 the variance between the three directions does not seem to be significant. At the optimum speed of 12.5 kn, or roughly 130 hours of voyage time, the cost difference between a head and tail wind is R8 840. An insignificant cost compared to the total cost of R2.75 million. This may provide a glimpse into how the model weighs the effect of individual parameters on its predictions, as it is not always sensitive to individual environmental parameter changes. This is an unlikely event to be encountered in practice as, for example, a decrease in wind speed could affect the wave height and length. Instead, the model might rely on the simultaneous and collective changes to the environmental conditions as a whole.

Voyage cost efficiency is a balancing act between two cost drivers as indicated in

Table 5.2: Weather scenarios noted in Figure 5.5a.

Scenario	Beaufort number [n/a]	Wind speed [kn]	Wind direction [Degrees]	Wave direction [Degrees]
1	1.4	6.6	0	0
2	2.8	16.7	0	0
3	5.6	23.3	0	0
4	7	30.0	0	0

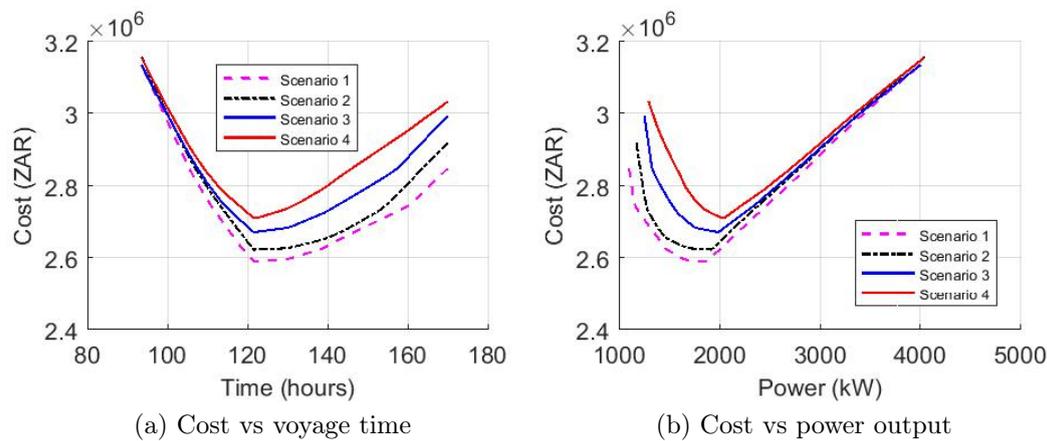


Figure 5.5: Costs at sea in terms of power demand and voyage time.

Figure 5.4. To reach a quicker voyage time translates into a much higher power demand, which increases fuel consumption. On the other hand, time-related overhead costs become significant when travelling slower than the optimum speed. In Figure 5.5 two plots of voyage cost are presented as functions of time and power for various weather scenarios with increasing severity from numbers 1 to 4. These plots illustrate the trade-off relationship between running costs, voyage time and power output to reach a destination. The weather vector, \mathbf{W} , for each of the four scenarios is detailed in Table 5.2. Each scenario was manually configured by adjusting parameters in conjunction with each other to describe the weather scenarios. As expected, the estimated minimum cost for ideal weather (Scenario 1) is much lower compared to harsher weather (Scenario 4). The increased resistance from more intense atmospheric and ocean conditions drive the cost difference between Scenarios 1 and 4 to roughly R118 000 at the optimum point.

Figure 5.5b highlights the balance between cost and average power output. By operating at higher power levels of 4000 kW in open water may improve the ship's arrival time, but can increase the cost by roughly 20%. A financial decision is to be made: incur significant extra fuel cost to minimise the time spent at sea or save on fuel by operating at optimum speeds. As discussed, routes

Table 5.3: Comparison between cost and time optimisation for a 3000 km voyage.

Optimisation goal	Power output [kW]	SOG est. [kn]	Cost est. [ZAR]	Time to destination [Hours]
Fastest time	4008	17.2	R 3 135 300	94
Lowest cost	1880	13.4	R 2 593 000	122

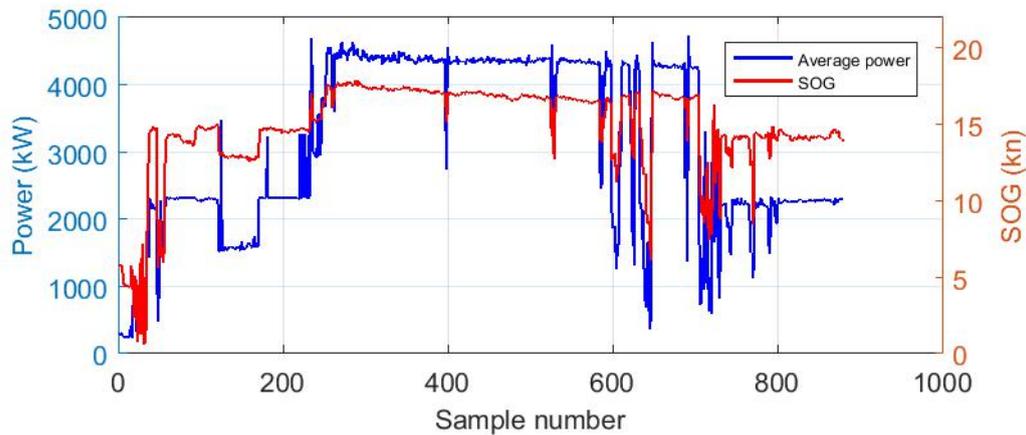


Figure 5.6: Power and SOG plots for open water navigation from 2-8 Jan 2019.

can be optimised in terms of voyage duration or cost. To achieve minimal voyage time is trivial, as maximum SOG will translate into minimum voyage time. The two contrasting optimisation goals are compared in Table 5.3. Considering Scenario 1 for the voyage of 3000 km, the destination is reached in 94 hours when sailing at full speed compared to 122 hours at optimal speed. However, it should be taken into consideration that this 28 hour gain comes at a significant extra cost estimate of roughly R542 300.

The SAII is not always operated at its optimum speed. Data from the Weddel sea expedition shows the ship travelling at speeds of 18 kn with the motors running at 4500 kW for long periods of time, Figure 5.6. The power output doubles for a 20% increase in SOG. The motivation for operating the ship at high power levels could override the need for financial efficiency. Such motivation could be ice avoidance or science and logistical activities that take priority over cost in the fixed time frame of the bigger expeditionary plans. Nonetheless, the eligibility of such decisions need to be weighed against the significant cost implications of excessive power usage.

Table 5.4: Ice navigation optimisation results for a 300 km ice route.

Ice conditions	Power output [kW]	SOG est. [kn]	Cost est. [ZAR]	Time to destination [Hours]
No ice	4276	6.6	R 829 400	25
Light	1316	4.7	R 598 800	34
Medium	1140	5.5	R 486 250	29
Hard	1157	8.7	R 304 700	19

5.4.2 Ice model - speed optimisation

Speed optimisation was conducted with the model trained on ice navigation data. Similar to how PSO was implemented for the open water model with varying weather, PSO has been used on the ice model for a voyage of 300 km with varying ice conditions. The results from this optimisation are listed in Table 5.4. The total costs were calculated using Equation 5.3.4 but with the initial expenses, b , set to R 0. Regrettably, there are some inconsistencies that make these results implausible. For instance, the highest power output that the model predicted was with no ice conditions. In addition, the fastest estimated time to the destination was 19 hours when the ship was intended to be in the toughest simulated ice conditions. The power was also not accurately predicted as noted in Table 5.4. Looking back at the actual data, the power fluctuated rapidly between 1000 kW and 4500 kW in ice, Figure 4.5a. However, the reported power estimates are averaged around the 1200 kW region and decrease despite worsening levels of simulated ice conditions.

Such contradictions have a negative impact on the credibility of the model. It may have accurately predicted the past voyages on the test data but struggles when brand new examples are presented. This problem most probably arose due to two main reasons: an insufficient quantity of ice data and poor modelling of the ice environment. The quality of models are dependent on the quality and quantity of the training data (Géron, 2017; Gkerekos *et al.*, 2019). A very small percentage of the 2017-2018 and 2019-2020 CMU data was measured during ice navigation, see Chapter 3. It is therefore expected that the ice model would not perform as well as the open water model. Also, due to the complex nature of how ice interacts with the ship, efficiency in ice is highly sensitive to the specific route. To further improve the reliability, a secondary navigation model of the environment is proposed for future research to simulate the routing options based on data from satellite imagery (Zhang *et al.*, 2019) and ship-mounted equipment such as radar or cameras with computer vision technology (Sandru, 2018).

5.4.3 Route recommendation for open water navigation - a proof of concept exercise

By optimising in terms of speed and cost paves the way to methods that provide recommendations for the best route choice. A simple hypothetical voyage is proposed in Figure 5.7 to illustrate the decision aiding potential of a data-driven model. The shortest distance that a ship can travel between point A and B is a straight line. However, in this hypothetical voyage, adverse weather is located midway between points A and B as indicated by the concentric circles. Route 1 is the shortest route but goes straight through the weather system. The other option, Route 2, deviates from the shortest route with an angle θ to bypass the storm. To facilitate a dynamic optimisation technique, as proposed by Wang *et al.* (2018), both routes are equally divided into six segments. A unique weather vector is assigned to each segment which aims to simulate changes in weather conditions over the length of the two routes. PSO is used to calculate the optimum ship speed and cost for each segment. The total cost for each route can be determined by summing the costs for all segments.

Table 5.5 provides a breakdown of the weather vectors for each of the six segments in both route options. The conditions were manually estimated to simulate how the ship would experience changing weather over the length of the two routes. The shortest option, Route 1, passes through a storm. Therefore, the wind conditions and Beaufort number intensifies from W_1 to W_3 as the ship approaches and enters the weather system. Higher Beaufort numbers of between 5 and 7 can be expected here. These conditions decrease again from W_4 to W_6 when the ship leaves the area. As an alternative option, Route 2, the ship sails around the storm, see Figure 5.7. This route is longer but the operating environment is less harsh. The weather conditions do increase in intensity as the ship approaches the edge of the storm, but not as much as for Route 1. As the ship turns towards point B, where the two routes converge again, the weather becomes increasingly similar between the two options. Apart from the wind direction, the vectors W_6 and W_6^* are exactly the same.

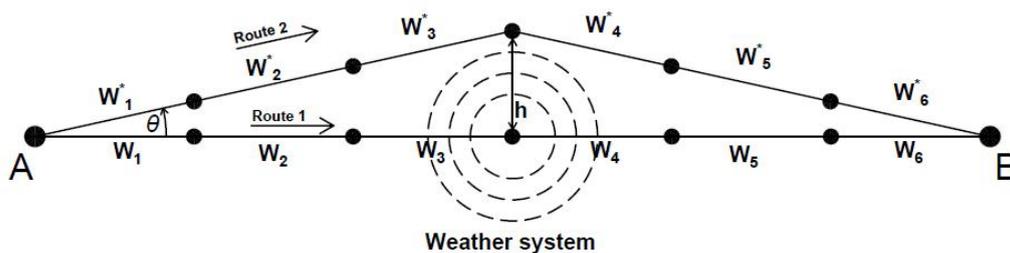


Figure 5.7: Route options for the hypothetical voyage.

Table 5.5: Breakdown of weather vectors for two possible routes.

	Vector	Wind direction [Deg]	Wind speed [kn]	Beaufort number [n/a]	Wave direction [Deg]
Route 1	W_1	323	6.5	1.4	310
	W_2	270	20.0	4.2	293
	W_3	270	31.5	6.7	310
	W_4	40	30.0	6.3	17
	W_5	18	20.0	4.6	17
	W_6	0	10.0	2.8	0
Route 2	W_1^*	341	6.7	1.4	17.3
	W_2^*	323	13.3	2.1	0
	W_3^*	323	16.7	3.5	17.3
	W_4^*	0	16.7	3.2	328
	W_5^*	341	14.0	2.8	338
	W_6^*	341	10.0	2.8	0

Table 5.6: Route optimisation results.

	θ [Deg]	h [km]	Distance [km]	SOG [kn]	Voyage time [hours]	Cost [ZAR]
Route 1	0	0	1000	12.7	42.6	R 849 850
Route 2	11.3	100	1020	13.3	41.5	R 853 100
	16.7	150	1043	13.2	42.6	R 873 460
	21.8	200	1077	13.3	43.7	R 900 915

The results obtained from dynamic PSO for both routes are listed in Table 5.6 and show the optimum speed with the expected cost. For Route 2, the PSO was done with three deviation angles, θ , to illustrate the trade-off between the costs of alternative routes compared to operational difficulty. Specific for the straight line routes in Figure 5.7, a bigger deviation angle translates into a larger perpendicular distance, h , from the weather system's centre. This also means that the alternative route's total distance from A to B increases as θ is increased. In addition, it is not assumed that the size of the weather system is fixed. In reality, the weather conditions should change as the distance to a storm's centre increases due to its fixed size. Instead, for this pilot exercise, the weather vectors remain independent of θ , meaning that they remain unchanged despite of the perpendicular distance to the storm, h , varying as a function of θ . This inherently implies that the size of the weather system fluctuates as θ is changed. However, it should not be considered as a problem for this pilot exercise. On the contrary, it highlights further complexity of automated route optimisation with decision-support potential. Weather modelling and the automatic import of ERA5 satellite wave data could provide improved

resolution of W_n along all directions and through time, which, in comparison to the two elementary routes currently available, allows for access to more and better routing options that can be evaluated.

The model that produced the results in Table 5.6 used the same cost structure given in Table 5.1, apart from the large initial cost, b , which was set to R 0. By comparison, the alternative scenario presented by Route 2 is very competitive in relation to Route 1. The three options for Route 2 are slightly more expensive due to longer distances that has to be travelled. However, estimated voyage times to the destination is around 42 hours on average. A very similar result compared to the more direct Route 1. Other factors such as comfort, ability to work on-board and safety factors enter the decision-making process. The human responses to rough weather conditions or long sustained periods at sea may negatively influence morale and crew productivity. With better weather resolution and factoring in comfort and safety, the cost-benefit ratio for possible routes can be better quantified. This could provide a tactical advantage for the ship officers in their decision making, as cost efficiency alone may not always be the best strategy.

5.5 Chapter summary

The argument was made that evolutionary optimisation techniques are better suited for the non-linear nature of working with ANN models. An objective function was defined that describes the fixed cost and fuel-related expenses expected from a voyage in an open water environment. PSO was used to define conditions for minimum voyage cost over the search domain. The data-driven model of the SAII's power output was able to reach convergence, which predicts an optimum speed for a given set of atmospheric and ocean conditions in open water navigation. A similar approach was used to optimise speed for ice navigation. However, due to insufficient data, the model was not able to make reliable predictions from new data. The need for environmental modelling was highlighted as a method to introduce ice floe interaction which produces the fluctuating power demand typical of ice navigation.

A basic route selection experiment was conducted with a dynamic optimisation approach. Similar to the definition of mathematical integration, the voyage distance would be broken down into smaller segments and optimised according to each segment's localised load conditions. The cumulative cost of each segment would then add up to the total cost for the route. From a cost efficiency point of view, alternative routes that pass through calmer oceans are very competitive compared to the most direct route for a hypothetical storm. The need to define comfort and safety was highlighted to provide supplementary guidance for route planning, as cost efficiency on its own, may not always be the best operational strategy.

Chapter 6

Conclusion

6.1 Introduction

Maritime operations contribute significantly to the global economy but also have a large impact on climate change (IMO, 2015; Olmer *et al.*, 2017). In the short term, methods that further operational efficiency should be the focus as advances in ship design technologies are likely on long term horizons (Johnson and Andersson, 2016). This study aimed to use operational data from the SAAI's Antarctic relief voyages to produce data-driven power performance models of the vessel. Apart from the goal to seek a sense of tactical foresight, this modelling effort pioneers for cost and route optimisation of this vessel, which could enable future systems to have decision support capability. Four objectives were defined aimed at achieving this, namely:

1. Filter, clean and synchronise the SAAI's data stored on the CMU and consolidate with ice and wave observations.
2. Train machine learning models to describe the power performance of the vessel. The best performing models were used in objectives 3 and 4, which was trained on the data from both 2017-2018 and 2019-2020 data. Model performance was validated with data recorded during the Weddel sea Expedition (2018-2019).
3. Use suitable optimisation techniques to find the optimum speed with the lowest fuel consumption for various operational conditions.
4. Implement the data-driven models to predict the best and most efficient route based on waypoints and environmental conditions.

All of the four objectives were implemented - each providing a unique perspective on the modelling process and possible future developments.

6.1.1 Data processing and quality assessment (objective 1)

The raw data from the CMU and observations were unsynchronised and stored in various formats which could not be used to train data-driven models - a common problem encountered in practice. The data required synchronisation into one complete set that contained all features measured by the CMU and observed during the two relief voyages. The biggest problem with the CMU, in addition to the unsynchronised data, was that in some cases the navigation parameters were recorded while measurements for the machine control data, such as propulsion and rudder settings, failed. This might have been caused by sensor errors or software issues on the ship's computer network. A placeholder for the missing measurements was unfortunately not recorded which lead to gaps in the time-series sequence for some measurement periods. To rectify this problem, a synchronisation algorithm was written in MATLAB to align the timestamps of navigational and machine control data. A similar technique was adopted to integrate ice and wave observations with the CMU data. **This algorithm is now a permanent asset to benefit future vessel data management.**

Data quality and veracity are important factors that determine the reliability and accuracy of data-driven models (Chollet, 2018). After comparing parameters such as power and speed, it was determined that multiple models would be required to predict the performance of the SAII in ice and open water independently. These contrasting operating conditions produce widely different power demands on the ship's propulsion system. The ratio of ice data to open water was disproportionately low. This meant that a single universal model would most likely classify ice data as outliers cases for open water passage and would be discarded. Therefore, to avoid diluting ice data, a secondary model trained specifically for ice navigation was proposed.

Large amounts of good quality open water data were available. A typical non-linear relationship was observed between power and speed, similar to data reported by Yoo and Kim (2018). Regions of high data concentrations were identified near the ship's cruising speed which foreshadowed the range where trained models would be best defined. High volumes of open water data with complete wave observations instilled confidence that training a reliable model would be possible.

Ice data was difficult to extract and relied mostly on ice observations to help distinguish between open water and ice cases. Parameters from the CMU was erratic, typically caused by sudden load changes due to the crushing and breaking forces exerted on the ship and from backing and ramming manoeuvres when passing through of fields of ice. Ice observations provided the only form of environmental data available of which the temporal resolution was particularly troublesome. A significant constraining factor would be the 10 minute

observation intervals which could at best provide average ice condition information on a macro scale. Local interactions between the hull and ice, which causes significant fluctuating loads, could not be measured visually.

6.1.2 Machine learning models (objective 2)

Two distinct machine learning approaches were trained on the ice and open water data sets. SVR and ANN were selected based on results reported in literature (Gkerekos *et al.*, 2019; Zhang *et al.*, 2019). SVR was implemented with the RBF kernel and for the ANN method, a FFNN architecture was used. Both ice and open water data sets were randomly divided into two groups for model development and validation. 70% of the data was allocated for training, while the other 30% was reserved for testing the accuracy of the model's predictions. From comparisons between the actual recorded power levels and estimates by both SVR and FFNN, it was determined that FFNN was superior. **The FFNN model could better describe the complex relationship between the output power, speed and environmental conditions and was, therefore, the preferred model of choice for further experimentation.**

The open water FFNN model was first trained on the 2019-2020 relief voyage data, which contained a tight clustering of power versus speed data around the rated cruising speed of the ship. Wave observations were conducted during this voyage and could therefore further contribute to how weather influenced powering performance. The model showed promising results on the 2019-2020 test data with a MAE deviation of 46 kW, which translates roughly into a 2% error. The model was further validated on the 2017-2018 data and illustrated the issue of insufficient and non-representative data. During the voyage of 2017-2018, the ship was operated at full power to achieve speeds of close the 18 kn at power levels higher than 4000 kW. Since this information was not present during training, the model could not accurately estimate power levels at higher than normal cruising speeds. This would be an issue when attempting optimisation. Therefore, a final FFNN model for open water powering performance was trained with data from both relief voyages to increase the quantity of training data and to widen the range for power predictions. Data from the Weddel sea expedition served as validation for the improved model. **The results were verified which instilled confidence that accurate and reliable power predictions for open water navigation was possible.**

A FFNN model for ice navigation was also trained and showed promise upon initial inspection. Despite the low quantity of usable ice data, the model was able to consistently predict the fluctuating power parameter when compared to similar data from both relief voyages. The typical MAE was 117 kW which corresponds to a 2.5% error measured from a maximum power output

of 4500 kW. However, the results were not convincing enough to rule out the likelihood that the model predicted from *memory*, meaning that estimates are based on similarities between the training and test data. The Weddel sea expedition produced ice data with different ice observation rates than what was available in the training set. This served as an external source of previously unseen data that could challenge the robustness of power estimates. **The ice model did not perform up to expectation and highlighted the serious limitation of insufficient training data.**

6.1.3 Speed and cost optimisation (Objective 3)

The non-linear relationship between power and speed creates an opportunity to optimise operational settings to achieve minimal running costs. The open water FFNN model, trained on data from both relief voyages, was able to predict these non-linear characteristics. Evolutionary optimisation techniques were found to be most versatile when implemented in conjunction with data-driven models. PSO was selected as an appropriate optimisation method. An objective function was intuitively defined and optimised for a hypothetical voyage of 3000 km. The PSO was successful and illustrated the effect of operating the SAAIL at higher than recommended speeds in open water. **A 20% increase in fuel expenditure could be expected if the ship was operated at its top speed of 18 kn instead of the ideal 13.5 kn.**

PSO was used with the ice model to find optimum speeds for ice navigation. Again, PSO was able to find optimum points but exposed limitations in the reliability and accuracy of the ice model. The predicted power was incoherent, with the model estimating very high power levels in zero or light ice conditions. Meanwhile, for a tough ice environment, the model predicted low power levels with high speeds. Possible reasons for this contradiction include insufficient quantities of training data, as validated by Weddel sea expedition data, but also the lack of ice environment modelling.

The use of data-driven models in optimisation problems showed some success but also highlighted the limitations caused by poor data quality and limited data quantities. The optimisation problem had to be constrained in such a way to simulate practical operational settings as close as possible, otherwise the PSO would find global minimum points that are not possible in reality. Such illogical predictions originate from data combinations that are never physically encountered. The search domain has to be constrained to ranges where the data-driven model is valid. The domain is limited to the distribution of training data as illustrated in Chapter 3. **This has shed light on the need to generate new data as much as possible to increase a model's exposure to various operating conditions.**

6.1.4 Route recommendation (Objective 4)

A pilot exercise for route recommendation and selection was conducted in an attempt to show how data-driven models could provide a sense of tactical foresight for ship operators. The exercise consisted of two route options in a hypothetical voyage between two points. The direct option passed through a weather system with harsh weather conditions, while the alternative route progressed around the adverse weather in search for a longer but easier route. A coarse dynamic optimisation method was manually introduced that could account for weather and sea state changes (Wang *et al.*, 2018).

The exercise showed that an alternative route can be very competitive with respect to cost and voyage time. There exists an optimum point between the additional loading from rough weather and the increased distance to the destination. The success of the exercise has highlighted the need for more complete weather data with improved resolution. The dynamic optimisation could also be automated to improve resolution and to evaluate other better-suited route options.

6.2 Reflection on modelling success

Modelling and decision support technologies address the grey areas around efficient voyage planning and its tactical execution. It is envisioned that officers in control of the SAII have access to an innovative tool that can support their decision making with regards to route choices and optimum speeds that reduce operating costs. This study attempted to research the fundamentals for training data-driven models and their application on the SAII. The model for open water powering performance was successfully trained, tested and used in optimisation exercises. The training data originated from the ship's installed infrastructure and provides an excellent opportunity to continuously re-train, maintain and update the model to improve its reliability and accuracy. The ice model initially did perform well on the test data but was unable to make good predictions during model validation and speed optimisation exercises. As mentioned, the most prominent reason for this is insufficient data. The model could be improved by further training using ice related CMU data from future voyages.

The origin of ice data remains an issue. Satellite imagery and on-board radar systems should to be investigated as an alternative data sources for ice conditions on a macro scale. Visual ice observations have limitations and are subject to the observer's judgement. Transitioning to machine vision technologies (Sandru, 2018) could improve the temporal resolution of ice observations and presents the opportunity to identify individual interactions between the ship and ice, and link them to fluctuations observed in the power over time plots. Both models could benefit from including other relatable features in the train-

ing data such as draft, displacement and trim. These parameters could extend the scope of both models, which could allow accurate predictions for a wider range of operating conditions.

Optimisation problems with the data-driven models were structured in such a way evaluate the feasibility of the decision support concept. The dynamic optimisation technique can be improved up to a point that is still computationally viable. However, it did show its value in terms of highlighting how operational changes could improve cost and emissions efficiency; thereby addressing one of the key benefits of digital twin technology (DNV-GL, 2016). The models could possibly help improve the efficiencies of the SAII in terms of cost, speed and route planning. The philosophy behind finding alternative routes to avoid adverse weather can be adapted to ice navigation problems where suitable ice routes can be planned in advance. Even so, route optimisation is not exclusively a cost-benefit problem but also has to address factors such as passenger comfort and safety (Zhang *et al.*, 2019). Energy management and cost reduction is only viable under acceptable levels of risk.

Modelling ship performance is a complex problem with reported analytical and empirical methods found to be lacking (Zhang *et al.*, 2019). Machine learning methods proved to be very useful and could be applied to almost every vessel type (Gkerekos *et al.*, 2019). The observed results showed that cost and speed optimisation is possible with the use of data-driven models. In addition, greater insight of a vessel's performance capabilities are available compared to a typical power versus speed curve which was estimated during sea trials (Bialystocki and Konovessis, 2016). Data-driven models could make massive advancements in how performance degradation, shipping operations and engine emissions are tracked over time (Gkerekos *et al.*, 2019).

6.3 Future work

The process from raw data to data-driven modelling were outlined and achieved reasonably successful results. Nonetheless, further research and development are still required to hone model performance and to reap the benefits of data-driven optimisation. The following future work is proposed in order to make advances in digital twin technology for the SAII:

1. The poor performance of the ice model during optimisation illustrated an opportunity for further investigation and refinement. Research can be directed towards the complex interactions between the ship and ice, and how these interactions effect the power demand. Alternative sources of environmental data such as satellite imagery, on-board radar systems, camera footage and drone surveillance could compliment the manual ice observations that were available for this study.

2. Develop a working route optimisation and decision support system for open water and ice navigation based on the data-driven models of the ship. Weather and ice data originating from external sources could provide a continuous, real-time source of weather and ice data that span across vast geographical areas. Literature has reported optimisation approaches such as Ant Colony Algorithms (ACA) that were successfully implemented for ice route optimisation in the Arctic (Zhang *et al.*, 2019). A comprehensive study is proposed to research these innovations to the benefit of the maritime industry based on the extensive full-scale data of the SAII in Antarctica.
3. Risk will always be part of maritime activities and should be reduced to acceptable levels to ensure the safety of crew and passengers. By incorporating risk matrices into route optimisation could be crucial for advancing such technologies into the maritime navigation toolbox. Future work is proposed to define the risks involved and assess whether cost reduction strategies are done under acceptable levels.
4. A financial perspective could provide an objective judgement on the economic success of data-driven speed and route optimisation efforts. Such a study would entail comprehensive analysis of maritime economics and how the finances of the SAII compares to the general industry. Voyages should be conducted, with preference given to the selected routes suggested by the data-driven models, and compared to the expenses from previous voyages with conventionally selected routes. Without a financial point of view it would be difficult to definitively conclude if route optimisation provides any realisable cost benefit to the ship owners and its operators.

Machine learning and artificial intelligence (AI) is a very powerful computational tool that is still in its infancy (Chollet, 2018). Advancements in data-driven modelling which predict asset responses for universal environmental conditions are key to the development digital twin platforms. The complete life cycles of assets could be modelled, simulated and optimised before start of any construction. The maritime industry may have formed the basis for data-driven modelling in this study, but the versatility of machine learning in an information intensive world make it possible for a wide range of applications in other industrial sectors such as mining, manufacturing, logistics and day-to-day business operations. In addition, driving efficiency is the key toward achieving the goal set out by the IMO to reduce carbon emissions from maritime activities by 40% by 2030 (Cosci, 2018). Green operations should be a focus point of innovations to ensure sustainability in the maritime industry. The data-driven performance models of the SAII will hopefully contribute to

the vessel's effective management and efficient operations, which would secure safe and sustainable scientific activities in the southern ocean and Antarctica.

Appendices

Appendix A

Algorithms

A.1 Synchronisation of CMU data

This MATLAB script automatically synchronises the machine control data and the navigational data from the CMU

```
clear all
% load data
load('CMU_Data_2017_2018.mat');
disp('Data upload complete')

% ----- DATA SYNC -----

% clear important variables
clear A

start_index = 157000; % define sample number to start sync
stop_index = length(machctrl.time); % define sample number to stop sync;

scaling = 1.5; % scaling factor for nav dataset
skip_factor = 100; % scaling factor to thin out data
dynamic_scaling = 1;

i = start_index; % main counter
j = round(scaling*start_index); % counter
n = 1; % counter for output array
m = 1; % counter in while loop
timediff = 1; % load timediff variable to enter while loop
overshoot_no = 0;
overshoot_flag = 0;
```

```
clear B
clear A

% flags
flag_while_entry = 0;

progressbar('Sync data...');

while i < stop_index

% Get machine control time instance
mach_str = datestr(machctrl.time(i,1));
mach_time = datetime(mach_str);

while (timediff > 0)

j = j + 1*dynamic_scaling; % increment

nav_str = datestr(nav.time(j,1));
nav_time = datetime(nav_str);

timediff_prev = timediff;

% convert datetime variable into epoch time and
% calculate the difference. The minimum difference
% will be the nearest time value.
timediff = posixtime(mach_time) - posixtime(nav_time);

B(m,1) = j;
B(m,2) = timediff;
B(m,3) = dynamic_scaling;

m = m + 1;

% set scaling factor
if timediff > 2
dynamic_scaling = round(timediff/2);
elseif timediff < -2 % overshoot correction
overshoot_no = overshoot_no + 1;
j = j - 1*dynamic_scaling;
dynamic_scaling = round(-(timediff/2));
else
```

```
dynamic_scaling = 1;
end

flag_while_entry = 1; % set flag

A(n,1) = j; % save index number with sychronised nav data points
A(n,2) = i; % corosponding machctrl data point
A(n,3) = timediff; % save time difference whith sychronised data points.
A(n,4) = m; % save point where final result is recorded

if A(n,3) < -5
A(n,3) = nan; % save time difference whith sychronised data points.
overshoot_flag = overshoot_flag + 1;
break
end

end

progressbar((i - start_index)/(stop_index - start_index));

% reset timediff variable to enter while loop
timediff = 1;

i = i + 1*skip_factor; % increment
n = n + 1; % increment
flag_while_entry = 0; % reset flag

if skip_factor > 1
j = j + round(skip_factor*scaling); % adjust j if skip factor applied
end
end

progressbar(1); % close progress bar
disp('Synchronization complete')

% Create array for synchronised features
clear data
clear D

% convert timestamp to epoch time
D(:,1) = posixtime(datetime(datestr(machctrl.time(A(:,2),1))));
```

```

% machine control data
data(:,1) = 1:length(A); % index number
data(:,2) = machctrl.time(A(:,2),1); % time and date number
data(:,3) = machctrl.data(A(:,2),1); % PortPropMotorCurrent
data(:,4) = machctrl.data(A(:,2),2); % PortPropMotorPower
data(:,5) = machctrl.data(A(:,2),3); % PortPropMotorSpeed
data(:,6) = machctrl.data(A(:,2),4); % PortPropMotorVoltage
data(:,7) = machctrl.data(A(:,2),5); % StbdPropMotorCurrent
data(:,8) = machctrl.data(A(:,2),6); % StbdPropMotorPower
data(:,9) = machctrl.data(A(:,2),7); % StbdPropMotorSpeed
data(:,10) = machctrl.data(A(:,2),8); % StbdPropMotorVoltage
data(:,11) = machctrl.data(A(:,2),9); % RudderOrderPort
data(:,12) = machctrl.data(A(:,2),10); % RudderOrderStbd
data(:,13) = machctrl.data(A(:,2),11); % RudderPositionPort
data(:,14) = machctrl.data(A(:,2),12); % RudderPositionStbd
data(:,15) = machctrl.data(A(:,2),13); % PropellerPitchPort
data(:,16) = machctrl.data(A(:,2),14); % PropellerPitchStbd
data(:,17) = machctrl.data(A(:,2),15); % ShaftRPMIndicationPort
data(:,18) = machctrl.data(A(:,2),16); % ShaftRPMIndicationStbd

disp('machctrl data array complete')

nav data

data(:,19) = nav.data(A(:,1),1); % NavTime
data(:,20) = nav.data(A(:,1),2); % Latitude
data(:,21) = nav.data(A(:,1),3); % Longitude
data(:,22) = 2*nav.data(A(:,1),4); % SOG
data(:,23) = nav.data(A(:,1),5); % COG
data(:,24) = nav.data(A(:,1),6); % HDT
data(:,25) = nav.data(A(:,1),7); % WindDirRel
data(:,26) = nav.data(A(:,1),8); % WindSpeed
data(:,27) = nav.data(A(:,1),9); % Depth
data(:,28) = D(:,1); % time and date epoch format

disp('nav data array complete')

% Save synchronised result in Excel file
xlswrite('My synchronised CMU data.xlsx', data);
disp('write file complete')

```

A.2 Synchronisation ice and wave observations with CMU data

This MATLAB script synchronises the ice and wave observations to the CMU data.

```
% Load data
%ice_obs_data = xlsread('Relief 2017_Ice Obs_trimmed sheet', 'A5:AX1866');
ice_obs_data = xlsread('Relief 2019_Ice Obs_trimmed sheet', 'A125:AP2465');
sea_obs_data = xlsread('SVRG Wave Observations Antarctica 2019-2020_trimmed',
, 'A4:AD915');
cmu_data = xlsread('My synchronised CMU data');

disp('Observation sheet upload complete')

ice_obs_data(:,2:7) = floor(ice_obs_data(:,2:7));
dates_ice = datetime(datestr(ice_obs_data(:,2:7)));

ice_obs_data(:,1) = posixtime(dates_ice);

% do the same for sea state data
sea_obs_data(:,2:7) = floor(sea_obs_data(:,2:7));
dates_sea = datetime(datestr(sea_obs_data(:,2:7)));

sea_obs_data(:,1) = posixtime(dates_sea);

% ----- Ice obserbation sync -----

progressbar('Sync data...'); % initialize progressbar function

start_cmu_index = 1; % index starting point for CMU data
stop_cmu_index = length(cmu_data); % index stopping point for CMU data

start_obs_index = 19; % index starting point for ice obs data
stop_obs_index = length(ice_obs_data); % index stopping point for
ice obs data

time_diff = 1; % initialize variable
time_diff_prev = 1;
```

```
i = start_cmu_index;
n = start_obs_index;

overshoot = 0;

override_flag = 0; % override flag to convert NaN numbers to 0
load_nan_flag = 0; % flag to load array with NaN numbers

if override_flag == 1 % keep NaN if 1
x_brash = 0;
x_ram = 0;
x_ice_conc = 0;
x_ice_thickness = 0;
x_flow_size = 0;
end

if load_nan_flag == 1
cmu_data(:,29:33) = nan;
end

while i < stop_cmu_index

% calculate time difference
time_diff_prev = time_diff;
time_diff = ice_obs_data(n+1,1) - cmu_data(i,28);

% timestamp is found
if time_diff < 0

% check if not a number
if override_flag == 0;

x_brash = isnan(ice_obs_data(n,12));
x_ram = isnan(ice_obs_data(n,13));
x_ice_conc = isnan(ice_obs_data(n,27));
x_ice_thickness = isnan(ice_obs_data(n,42));
x_flow_size = 1;%isnan(ice_obs_data(n,50));

end

if abs(time_diff) > abs(time_diff_prev) % overshoot case

overshoot = overshoot + 1;
```

```
i = i - 1;

% set first run
if n == start_obs_index
    index_prev = i;
end

if x_brash == 1
    cmu_data(index_prev:i,29) = 0;
else
    cmu_data(index_prev:i,29) = ice_obs_data(n,12); % brash ice count
end

if x_ram == 1
    cmu_data(index_prev:i,30) = 0;
else
    cmu_data(index_prev:i,30) = ice_obs_data(n,13); % Ramming count
end

if x_ice_conc == 1
    cmu_data(index_prev:i,31) = 0;
else
    cmu_data(index_prev:i,31) = ice_obs_data(n,27); % Ice concentration average
end

if x_ice_thickness == 1
    cmu_data(index_prev:i,32) = 0;
else
    cmu_data(index_prev:i,32) = ice_obs_data(n,42); % Ice thickness average
end

if x_flow_size == 1
    cmu_data(index_prev:i,33) = 0;
else
    cmu_data(index_prev:i,33) = ice_obs_data(n,50); % Flow size average
end

else % normal case

% set first run
if n == start_obs_index
    index_prev = i;
```

```
end

if x_brash == 1
    cmu_data(index_prev:i,29) = 0;
else
    cmu_data(index_prev:i,29) = ice_obs_data(n,12); % brash ice count
end

if x_ram == 1
    cmu_data(index_prev:i,30) = 0;
else
    cmu_data(index_prev:i,30) = ice_obs_data(n,13); % Ramming count
end

if x_ice_conc == 1
    cmu_data(index_prev:i,31) = 0;
else
    cmu_data(index_prev:i,31) = ice_obs_data(n,27); % Ice concentration average
end

if x_ice_thickness == 1
    cmu_data(index_prev:i,32) = 0;
else
    cmu_data(index_prev:i,32) = ice_obs_data(n,42); % Ice thickness average
end

end

index_prev = i; % save the previous index number
progressbar((n - start_obs_index)/(stop_obs_index - start_obs_index));
n = n + 1;

end

i = i + 1;

if n == stop_obs_index
    break;
end

end
```

```
progressbar(1) % close progressbar
disp('end')

% ----- Wave observation sync -----

progressbar('Sync data...'); % initialize progressbar function

start_cmu_index = 2; % index starting point for CMU data
stop_cmu_index = length(cmu_data); % index stopping point for CMU data

start_obs_index = 214; % index starting point for wave data
stop_obs_index = length(sea_obs_data); % index stopping point for
wave obs data

time_diff = 1; % initialize variable
time_diff_prev = 1;
%time_diff_prev = 1.5e+9;

i = start_cmu_index;
n = start_obs_index;

overshoot = 0;

override_flag = 0; % override flag to convert NaN numbers to 0
load_nan_flag = 0; % flag to load array with NaN numbers

if override_flag == 1 % keep NaN if 1
x = 0;
x_wavedir = 0;
x_waveheight = 0;
x_ice_thickness = 0;
x_flow_size = 1;
end

if load_nan_flag == 1
cmu_data(:,34:40) = nan;
end

while i < stop_cmu_index

% calculate time difference
```

```
time_diff_prev = time_diff;
time_diff = sea_obs_data(n+1,1) - cmu_data(i,28);

% timestamp is found
if time_diff < 0

A(n) = i;
%disp('diff success')
% check if not a number
if override_flag == 0;

x_bf = isnan(sea_obs_data(n,14));
x_wd = isnan(sea_obs_data(n,19));
x_wh = isnan(sea_obs_data(n,20));
x_sh = isnan(sea_obs_data(n,21));
x_wl = isnan(sea_obs_data(n,23));
x_wp = isnan(sea_obs_data(n,26));
x_ef = isnan(sea_obs_data(n,30));

end

if abs(time_diff) > abs(time_diff_prev) % overshoot case

overshoot = overshoot + 1;
i = i - 1;

% set first run
if n == start_obs_index
index_prev = i;
end

if x_bf == 1
cmu_data(index_prev:i,34) = 0;
else
cmu_data(index_prev:i,34) = sea_obs_data(n,14); % Beaufort number
end

if x_wd == 1
cmu_data(index_prev:i,35) = 0;
else
cmu_data(index_prev:i,35) = sea_obs_data(n,19); % Wave direction
rel to ship
end
```

```
if x_wh == 1
cmu_data(index_prev:i,36) = 0;
else
cmu_data(index_prev:i,36) = sea_obs_data(n,20); % Ave wave height
end

if x_sh == 1
cmu_data(index_prev:i,37) = 0;
else
cmu_data(index_prev:i,37) = sea_obs_data(n,21); % Max swell heighy
end

if x_wl == 1
cmu_data(index_prev:i,38) = 0;
else
cmu_data(index_prev:i,38) = sea_obs_data(n,23); % Wave length lambda
end

if x_wp == 1
cmu_data(index_prev:i,39) = 0;
else
cmu_data(index_prev:i,39) = sea_obs_data(n,26); % Average Wave period
end

if x_ef == 1
cmu_data(index_prev:i,40) = 0;
else
cmu_data(index_prev:i,40) = sea_obs_data(n,30); % Average encounter
frequency
end

else % normal case

% set first run
if n == start_obs_index
index_prev = i;
end

if x_bf == 1
cmu_data(index_prev:i,34) = 0;
else
```

```
cmu_data(index_prev:i,34) = sea_obs_data(n,14); % Beaufort number
end

if x_wd == 1
cmu_data(index_prev:i,35) = 0;
else
cmu_data(index_prev:i,35) = sea_obs_data(n,19); % Wave direction
rel to ship
end

if x_wh == 1
cmu_data(index_prev:i,36) = 0;
else
cmu_data(index_prev:i,36) = sea_obs_data(n,20); % Ave wave height
end

if x_sh == 1
cmu_data(index_prev:i,37) = 0;
else
cmu_data(index_prev:i,37) = sea_obs_data(n,21); % Max swell heighy
end

if x_wl == 1
cmu_data(index_prev:i,38) = 0;
else
cmu_data(index_prev:i,38) = sea_obs_data(n,23); % Wave length lamda
end

if x_wp == 1
cmu_data(index_prev:i,39) = 0;
else
cmu_data(index_prev:i,39) = sea_obs_data(n,26); % Average Wave period
end

if x_ef == 1
cmu_data(index_prev:i,40) = 0;
else
cmu_data(index_prev:i,40) = sea_obs_data(n,30); % Average encounter
frequency
end
end

index_prev = i; % save the previous index number
progressbar((n - start_obs_index)/(stop_obs_index - start_obs_index));
```

```

n = n + 1;
end

%progressbar((i - start_cmu_index)/(stop_cmu_index - start_cmu_index));
i = i + 1;

if n == stop_obs_index
break;
end

end

progressbar(1) % close progressbar
disp('end')

% Write to Excel File

xlswrite('My synchronised data set.xlsx', cmu_data);
disp('write file complete')

```

A.3 Support vector regression

```

In [ ]: # Support vector regression training

# Load required python packages
import pandas as pd
import numpy as np

# Load packages from Scikit-Learn
from sklearn.svm import SVR
from sklearn.model_selection import train_test_split
from sklearn.preprocessing import MinMaxScaler
from sklearn.model_selection import GridSearchCV

# Import pickle library to save models
import pickle

# Import data from Excel
# Open water data from 2019-2020 relief voyage
df1 = pd.read_excel('CMU_2019_2020_open water.xlsx')

```

```
# Stationary data from the 2019-2020 relief voyage
df2 = pd.read_excel('CMU_2019_2020_stationary.xlsx')

# Data for ice navigation is imported in the same manner

# Combine into one data frame
df = df1.append(df2, sort=False)
=====

# Split data set into a training and test set
training, test = train_test_split(df, test_size = 0.33, random_state = 42)

=====
# Data has to be normalised and presented in the correct shape and format
# to train a model successfully
=====

# Normalisation of training data

# Transform par 1
scaler = MinMaxScaler()
scaler.fit(df['PortPropMotorSpeed'].values.reshape(-1,1))
train_par_1 = scaler.transform(training['PortPropMotorSpeed']
                                .values.reshape(-1,1))

# Transform par 2
scaler.fit(df['StbdPropMotorSpeed'].values.reshape(-1,1))
train_par_2 = scaler.transform(training['StbdPropMotorSpeed']
                                .values.reshape(-1,1))

# Transform par 3
scaler.fit(df['PropellerPitchPort'].values.reshape(-1,1))
train_par_3 = scaler.transform(training['PropellerPitchPort']
                                .values.reshape(-1,1))

# Transform par 4
scaler.fit(df['PropellerPitchStbd'].values.reshape(-1,1))
train_par_4 = scaler.transform(training['PropellerPitchStbd']
                                .values.reshape(-1,1))

# Transform par 5
scaler.fit(df['SOG'].values.reshape(-1,1))
train_par_5 = scaler.transform(training['SOG']
                                .values.reshape(-1,1))
```

```
# Transform par 6
scaler.fit(df['WindDirRel'].values.reshape(-1,1))
train_par_6 = scaler.transform(training['WindDirRel']
                               .values.reshape(-1,1))

# Transform par 7
scaler.fit(df['WindSpeed'].values.reshape(-1,1))
train_par_7 = scaler.transform(training['WindSpeed']
                               .values.reshape(-1,1))

# Transform par 8
scaler.fit(df['Beaufort number'].values.reshape(-1,1))
train_par_8 = scaler.transform(training['Beaufort number']
                               .values.reshape(-1,1))

# Transform par 9
scaler.fit(df['Wave direction'].values.reshape(-1,1))
train_par_9 = scaler.transform(training['Wave direction']
                               .values.reshape(-1,1))

# Transform par 10
scaler.fit(df['Wave length'].values.reshape(-1,1))
train_par_10 = scaler.transform(training['Wave length']
                               .values.reshape(-1,1))

#=====

# Normalisation of ice data

# Transform par 1
scaler = MinMaxScaler()
scaler.fit(df['PortPropMotorSpeed'].values.reshape(-1,1))
test_par_1 = scaler.transform(test['PortPropMotorSpeed']
                              .values.reshape(-1,1))

# Transform par 2
scaler.fit(df['StbdPropMotorSpeed'].values.reshape(-1,1))
test_par_2 = scaler.transform(test['StbdPropMotorSpeed']
                              .values.reshape(-1,1))

# Transform par 3
scaler.fit(df['PropellerPitchPort'].values.reshape(-1,1))
test_par_3 = scaler.transform(test['PropellerPitchPort']
```

```
.values.reshape(-1,1))

# Transform par 4
scaler.fit(df['PropellerPitchStbd'].values.reshape(-1,1))
test_par_4 = scaler.transform(test['PropellerPitchStbd']
                              .values.reshape(-1,1))

# Transform par 5
scaler.fit(df['SOG'].values.reshape(-1,1))
test_par_5 = scaler.transform(test['SOG']
                              .values.reshape(-1,1))

# Transform par 6
scaler.fit(df['WindDirRel'].values.reshape(-1,1))
test_par_6 = scaler.transform(test['WindDirRel']
                              .values.reshape(-1,1))

# Transform par 7
scaler.fit(df['WindSpeed'].values.reshape(-1,1))
test_par_7 = scaler.transform(test['WindSpeed']
                              .values.reshape(-1,1))

# Transform par 8
scaler.fit(df['Beaufort number'].values.reshape(-1,1))
test_par_8 = scaler.transform(test['Beaufort number']
                              .values.reshape(-1,1))

# Transform par 9
scaler.fit(df['Wave direction'].values.reshape(-1,1))
test_par_9 = scaler.transform(test['Wave direction']
                              .values.reshape(-1,1))

# Transform par 10
scaler.fit(df['Wave length'].values.reshape(-1,1))
test_par_10 = scaler.transform(test['Wave length']
                              .values.reshape(-1,1))

#=====

# Transform training data into the correct shape

X_1 = np.reshape(train_par_1,-1)
X_2 = np.reshape(train_par_2,-1)
X_3 = np.reshape(train_par_3,-1)
```

```
X_4 = np.reshape(train_par_4,-1)
X_5 = np.reshape(train_par_5,-1)
X_6 = np.reshape(train_par_6,-1)
X_7 = np.reshape(train_par_7,-1)
X_8 = np.reshape(train_par_8,-1)
X_9 = np.reshape(train_par_9,-1)
X_10 = np.reshape(train_par_10,-1)

X_train = (X_1, X_2, X_3, X_4, X_5, X_6, X_7, X_8, X_9, X_10)
train_x = np.transpose(X_train)

# Transform test data into the correct shape

X1 = np.reshape(test_par_1,-1)
X2 = np.reshape(test_par_2,-1)
X3 = np.reshape(test_par_3,-1)
X4 = np.reshape(test_par_4,-1)
X5 = np.reshape(test_par_5,-1)
X6 = np.reshape(test_par_6,-1)
X7 = np.reshape(test_par_7,-1)
X8 = np.reshape(test_par_8,-1)
X9 = np.reshape(test_par_9,-1)
X10 = np.reshape(test_par_10,-1)

X_test = (X1, X2, X3, X4, X5, X6, X7, X8, X9, X10)
test_x = np.transpose(X_test)

# Data is now prepared and ready for training
#=====

# SVR training !!!!!!!!!!!!!!!!

# Train SVR model - hyperparameters not optimised

clf = SVR(C=1.0, epsilon=0.2)
clf.fit(train_x, train_y)

# Check accuracy score
SVR_score_untuned = clf.score(test_x, test_y)
print(SVR_score_untuned)

#=====
# Optimise SVR hyper-paramters
```

```

parameters = {'kernel': ('linear', 'rbf', 'poly'), 'C':(1,4,8,16,32)}
svr = SVR()

# Train model
clf2 = GridSearchCV(svr, parameters, cv=5)
clf2.fit(train_x, train_y)

# Check accuracy score
SVR_score_tuned = clf.score(test_x, test_y)
print(SVR_score_tuned)

# Save optimised SVR model which can be called when needed
with open('My SVR model.pkl', 'wb') as f:
    pickle.dump(clf2,f)

# End of SVR training
#=====

```

A.4 Feed-forward neural network

```

In [ ]: # Feed-forward neural network training

# Import required python packages
import pandas as pd
import numpy as np

# Packages from Scikit-Learn
from sklearn.svm import SVR
from sklearn.model_selection import train_test_split
from sklearn.preprocessing import MinMaxScaler

# Packages from Keras and Tensorflow
import tensorflow as tf
from tensorflow import keras
from keras import models
from keras import layers

# Import data from Excel
# Open water data from 2019-2020 relief voyage
df1 = pd.read_excel('CMU_2019_2020_open water.xlsx')

# Stationary data from the 2019-2020 relief voyage
df2 = pd.read_excel('CMU_2019_2020_stationary.xlsx')

```

```
# Open water data from 2017-2020 relief voyage
df3 = pd.read_excel('CMU_2017_2018_open water.xlsx')

# Data for ice navigation is imported in the same manner

# combine two data sets
df_1 = df1.append(df2, sort=False)
df = df_1.append(df3, sort=False)
=====

# Split data set into a training and test set
training, test = train_test_split(df, test_size = 0.33, random_state = 42)

=====
# Data has to be normalised and presented in the correct shape and format
# to train a model successfully
=====

# Normalisation of training data

# Transform par 1
scaler = MinMaxScaler()
scaler.fit(df['PortPropMotorSpeed'].values.reshape(-1,1))
train_par_1 = scaler.transform(training['PortPropMotorSpeed']
                                .values.reshape(-1,1))

# Transform par 2
scaler.fit(df['StbdPropMotorSpeed'].values.reshape(-1,1))
train_par_2 = scaler.transform(training['StbdPropMotorSpeed']
                                .values.reshape(-1,1))

# Transform par 3
scaler.fit(df['PropellerPitchPort'].values.reshape(-1,1))
train_par_3 = scaler.transform(training['PropellerPitchPort']
                                .values.reshape(-1,1))

# Transform par 4
scaler.fit(df['PropellerPitchStbd'].values.reshape(-1,1))
train_par_4 = scaler.transform(training['PropellerPitchStbd']
                                .values.reshape(-1,1))

# Transform par 5
scaler.fit(df['SOG'].values.reshape(-1,1))
```

```
train_par_5 = scaler.transform(training['SOG']
                                .values.reshape(-1,1))

# Transform par 6
scaler.fit(df['WindDirRel'].values.reshape(-1,1))
train_par_6 = scaler.transform(training['WindDirRel']
                                .values.reshape(-1,1))

# Transform par 7
scaler.fit(df['WindSpeed'].values.reshape(-1,1))
train_par_7 = scaler.transform(training['WindSpeed']
                                .values.reshape(-1,1))

# Transform par 8
scaler.fit(df['Beaufort number'].values.reshape(-1,1))
train_par_8 = scaler.transform(training['Beaufort number']
                                .values.reshape(-1,1))

# Transform par 9
scaler.fit(df['Wave direction'].values.reshape(-1,1))
train_par_9 = scaler.transform(training['Wave direction']
                                .values.reshape(-1,1))

# Transform par 10
scaler.fit(df['Wave length'].values.reshape(-1,1))
train_par_10 = scaler.transform(training['Wave length']
                                .values.reshape(-1,1))

#=====

# Normalisation of ice data

# Transform par 1
scaler = MinMaxScaler()
scaler.fit(df['PortPropMotorSpeed'].values.reshape(-1,1))
test_par_1 = scaler.transform(test['PortPropMotorSpeed']
                                .values.reshape(-1,1))

# Transform par 2
scaler.fit(df['StbdPropMotorSpeed'].values.reshape(-1,1))
test_par_2 = scaler.transform(test['StbdPropMotorSpeed']
                                .values.reshape(-1,1))

# Transform par 3
```

```

scaler.fit(df['PropellerPitchPort'].values.reshape(-1,1))
test_par_3 = scaler.transform(test['PropellerPitchPort']
                              .values.reshape(-1,1))

# Transform par 4
scaler.fit(df['PropellerPitchStbd'].values.reshape(-1,1))
test_par_4 = scaler.transform(test['PropellerPitchStbd']
                              .values.reshape(-1,1))

# Transform par 5
scaler.fit(df['SOG'].values.reshape(-1,1))
test_par_5 = scaler.transform(test['SOG']
                              .values.reshape(-1,1))

# Transform par 6
scaler.fit(df['WindDirRel'].values.reshape(-1,1))
test_par_6 = scaler.transform(test['WindDirRel']
                              .values.reshape(-1,1))

# Transform par 7
scaler.fit(df['WindSpeed'].values.reshape(-1,1))
test_par_7 = scaler.transform(test['WindSpeed']
                              .values.reshape(-1,1))

# Transform par 8
scaler.fit(df['Beaufort number'].values.reshape(-1,1))
test_par_8 = scaler.transform(test['Beaufort number']
                              .values.reshape(-1,1))

# Transform par 9
scaler.fit(df['Wave direction'].values.reshape(-1,1))
test_par_9 = scaler.transform(test['Wave direction']
                              .values.reshape(-1,1))

# Transform par 10
scaler.fit(df['Wave length'].values.reshape(-1,1))
test_par_10 = scaler.transform(test['Wave length']
                              .values.reshape(-1,1))

#=====

# Transform training data into the correct shape
X_1 = np.reshape(train_par_1,-1)

```

```

X_2 = np.reshape(train_par_2,-1)
X_3 = np.reshape(train_par_3,-1)
X_4 = np.reshape(train_par_4,-1)
X_5 = np.reshape(train_par_5,-1)
X_6 = np.reshape(train_par_6,-1)
X_7 = np.reshape(train_par_7,-1)
X_8 = np.reshape(train_par_8,-1)
X_9 = np.reshape(train_par_9,-1)
X_10 = np.reshape(train_par_10,-1)

X_train = (X_1, X_2, X_3, X_4, X_5, X_6, X_7, X_8, X_9, X_10)
train_x = np.transpose(X_train)

# Transform test data into the correct shape

X1 = np.reshape(test_par_1,-1)
X2 = np.reshape(test_par_2,-1)
X3 = np.reshape(test_par_3,-1)
X4 = np.reshape(test_par_4,-1)
X5 = np.reshape(test_par_5,-1)
X6 = np.reshape(test_par_6,-1)
X7 = np.reshape(test_par_7,-1)
X8 = np.reshape(test_par_8,-1)
X9 = np.reshape(test_par_9,-1)
X10 = np.reshape(test_par_10,-1)

X_test = (X1, X2, X3, X4, X5, X6, X7, X8, X9, X10)
test_x = np.transpose(X_test)

# Data is now prepared and ready for training
#=====

# Define neural network architecture

def build_model():
    model = models.Sequential()
    model.add(layers.Dense(64, activation='relu', input_shape=(10,)))
    model.add(layers.Dense(64, activation='relu'))
    model.add(layers.Dense(1))
    model.compile(optimizer='rmsprop', loss='mse', metrics=['mae'])
    return model

model = build_model()

```

```

# train model for 100 epochs on training set
history = model.fit(train_x, train_y, epochs = 100, batch_size=16, verbose=0)

# Return MSE and MAE scores
test_mse_score, test_mae_score = model.evaluate(test_x, test_y)

# save model to file
model.save("My neural network model.h5")
print("Saved model to disk")

# End of training procedure
#=====

```

A.5 Particle swarm optimisation

```
In [ ]: # Partical swarm optimisation
```

```

# import required python packages
import tensorflow as tf
from tensorflow import keras
from keras import models
from keras import layers
import numpy as np
import pandas as pd
import math
import pyswarms as ps
from pyswarms.single.global_best import GlobalBestPSO

# import neural network model trained on the 2019-2020 and 2017-2018 data
ship_model = models.load_model('Run 67_NN.h5')
SOG_model = models.load_model('Run 66_NN_SOG estimate.h5')

# define objective function
#=====
# function - the cost function to be evaluated
def function(x):

# define voyage constants - all of these are estimates
    d = 3000 # distance
    m = 10000 # R10 000 per hour - estimate
    c = 100000 # Minimum cost of R 100 000 before starting voyage
    b = 2.8 # Rough calculation - R 2,8/kWh

```

```

# weather variables
windd = 0.25 # wind_dir - head wind
ws = 1 # wind speed
BF = 0.5 # beaufort number
wave_d = 0 # wave direction
wl = 0.035 # wave length

l = 0
p = 0

x2 = np.empty([10,10])
x_SOG_est = np.empty([10,4])

x1 = np.reshape(x, (10,2))

x_SOG_est[:,0] = x1[:,0]
x_SOG_est[:,1] = x1[:,1]
x_SOG_est[:,2] = x1[:,0]
x_SOG_est[:,3] = x1[:,1]

# SOG estimate from SOG FFNN model
SOG_est = SOG_model.predict(x_SOG_est) # output in knots

constants = np.reshape(np.array([windd, ws, BF, wave_d, wl]), (1,5))

# populate array for
for l in range(10):
    for p in range(10):
        if p < 4:
            x2[l,p] = x_SOG_est[l,p]
        elif p == 4:
            # convert to scale between 0 and 1
            x2[l,p] = SOG_est[l,0]/17.94
        else:
            x2[l,p] = constants[0,p-5]

x_shaped = np.reshape(x2, (-1,10))

# neural network
nn = 2*ship_model.predict(x_shaped)

# convert SOG value to km/h
SOG = 1.852*(np.reshape(SOG_est[:,], (10,1)))

```

```

    # cost function - total costs = fixed costs + voyage costs
    fx = m*(d/SOG) + c + b*nn*(d/SOG)
    fx = np.reshape(fx, (10,))

    return fx
# end of function definition
#=====

# Define PSO variables

# create bounds of input parameters
max_bound= 1*np.ones(2)
min_bound = 0.65*np.ones(2)
bounds = (min_bound, max_bound)

# initialize swarm
options = {'c1': 0.5, 'c2': 0.5, 'w':0.9}

# initiate optimizer
optimizer = GlobalBestPSO(n_particles=10, dimensions=2,
                           options=options, bounds=bounds)

# now run the optimization for 1000 iterations
cost, pos = optimizer.optimize(function, 1000)

x_input = np.array([pos, pos])
x_input = np.reshape(x_input, (-1,4))

# Get SOG estimate from optimum input conditions
SOG_estimate = SOG_model.predict(x_input)

# Print outputs

# Print SOG
print("Optimum SOG:")
print(SOG_estimate)

# Print best propeller and shaft speed configuration
print("Optimum machine control settings:")
print(pos)

# Print minimum cost
print("Minimum cost:")

```

```
print(cost)

# End of PSO algorithm
# =====
```

Appendix B

Observations from previous voyage data

Various histograms of noteworthy parameters from the 2019-2020 Antarctic relief voyage are presented in Figure B.1 to illustrate regions where the most data originated during for open water navigation. SOG, Figure B.1a, is concentrated around 15 kn while most of the power samples, Figure B.1b, cluster near 2500 kW. The similar grouping supports the relationship between power and SOG. Propeller pitch, Figure B.1c, and shaft rotational speed, Figure B.1d, were mostly maintained at constant maximum levels during open water passage. When using these parameters as inputs to a data-driven model, the configuration of both variables, for the port and starboard side, should be kept in the same general region as shown in the histograms to ensure the validity of a trained model's predictions. Lastly, Figures B.1f and B.1e show most common wind speeds and relative directions measured during the 2019-2020 relief voyage. The ship rarely experience wind speeds higher than 20 kn. Head winds were also the most common wind direction encountered.

The data distribution during ice navigation is different to open water passage. Figure B.2 presents histograms for power, SOG, propeller pitch and shaft speed during ice operations. The SOG of the ship, Figure B.2a, is concentrated around speeds between 5 to 10 kn. The most samples for power demand are located at lower power levels of 1250 kW. However, a number of samples are also located at peak power levels of 4000 kW, Figure B.2b. The relationship between power and SOG are not apparent from inspecting the two histograms, as power sample concentrations are not complemented by a corresponding grouping of SOG samples. This supports the idea that external ice interactions are the leading factor driving power demand. In some cases, the ship may get stuck in ice and must reverse to free itself. Therefore, the propeller pitch often changes position, Figure B.2c. The shaft speed on the other hand is kept fairly constant near 140 rpm, Figure B.2d.

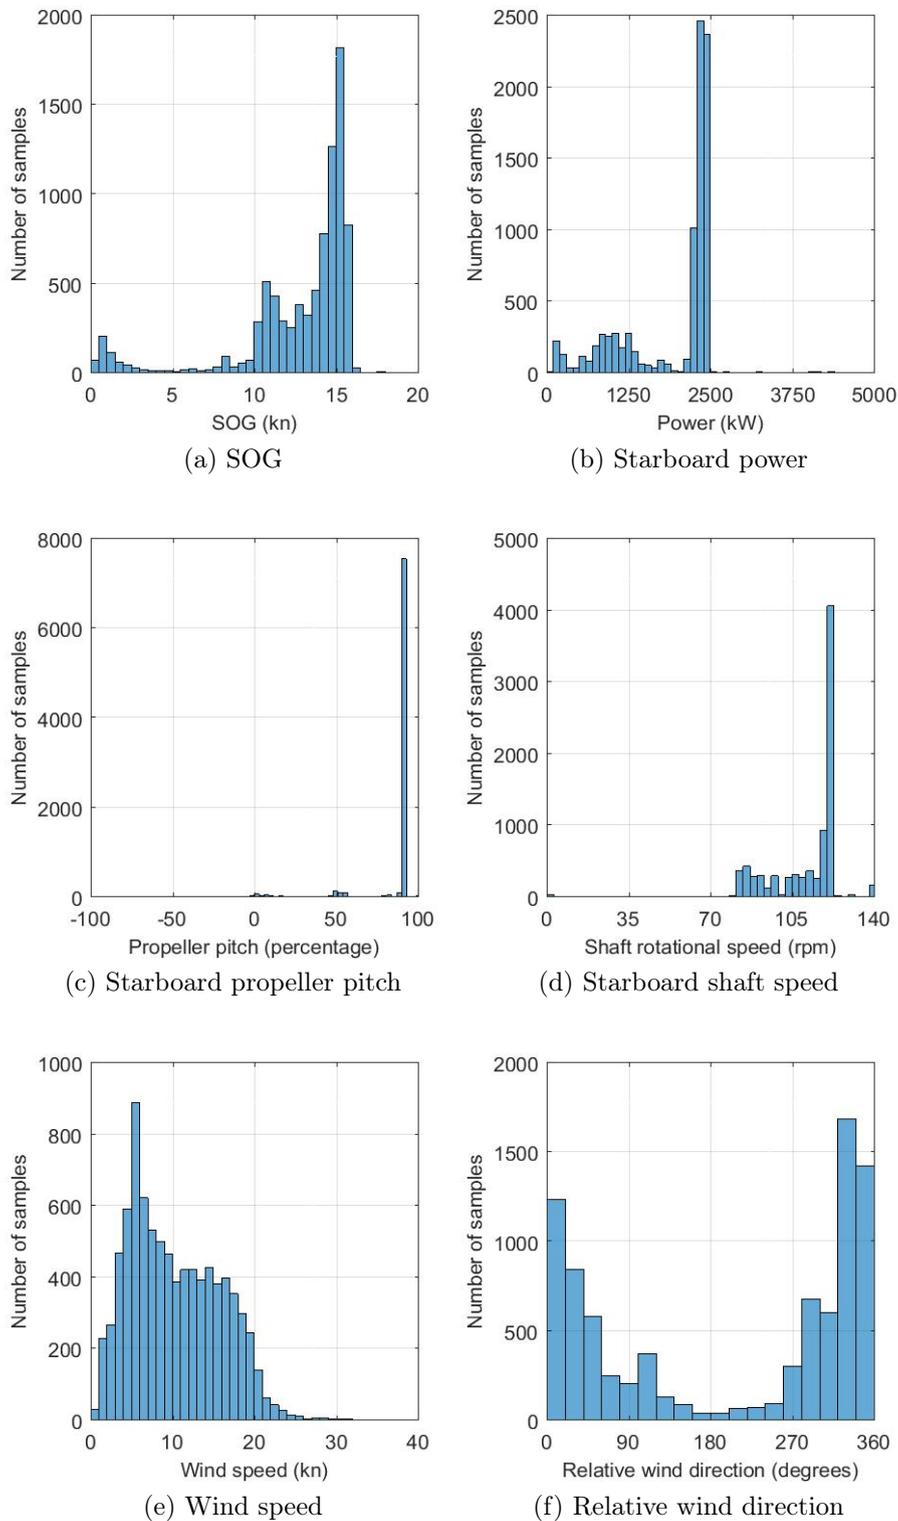


Figure B.1: Histogram of noteworthy CMU parameters during open water navigation (2019-2020 relief voyage).

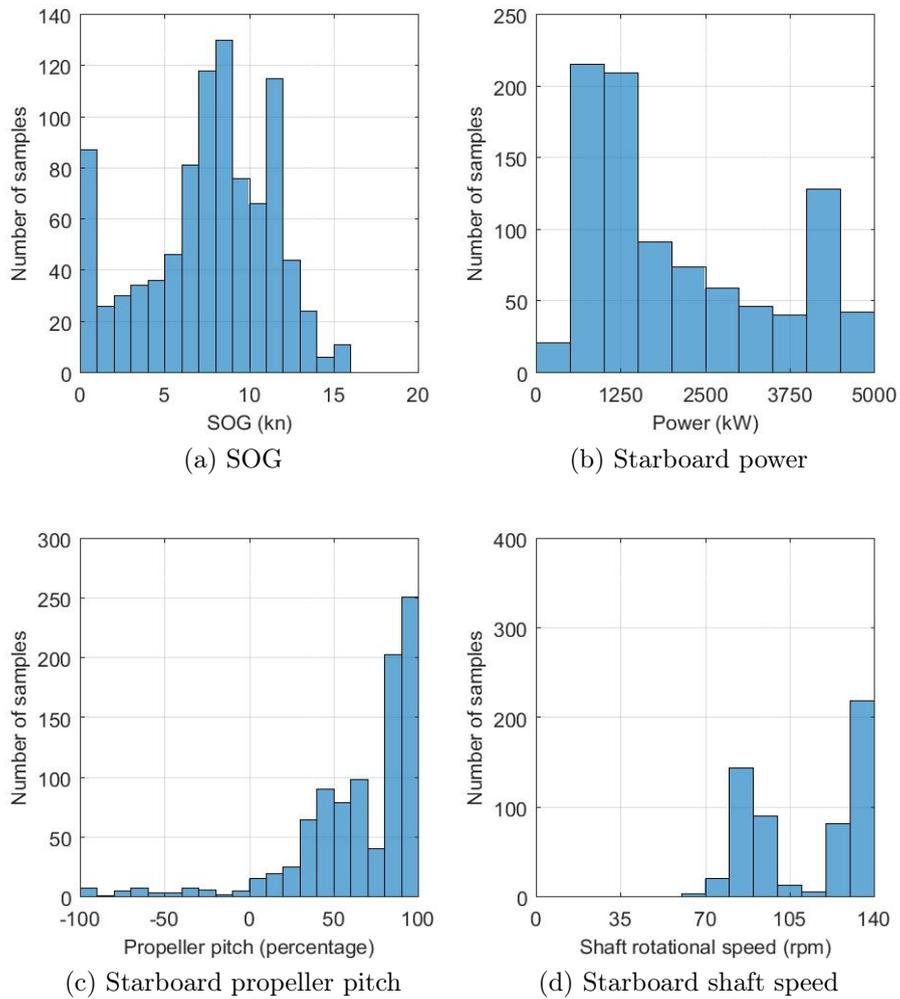


Figure B.2: Histogram of noteworthy CMU parameters during ice navigation (2019-2020 relief voyage).

Appendix C

Fuel cost calculation

C.1 Calculation of running cost

An hourly overhead cost was introduced into the speed optimisation problem as discussed in Chapter 5. A breakdown of the estimated hourly cost calculation is presented in Table C.1. These budgeted expenses were estimated to facilitate the optimisation problem in Chapter 5. All amounts are listed in South African Rand (ZAR). The estimated expenses added up to R9 840 and was subsequently rounded up to R10 000.

Table C.1: Breakdown of estimated hourly costs (ZAR).

Position	Number of crew	Estimated annual earnings	Estimated hourly earnings	Total
Captain	1	R1 200 000	R520	R520
Chief Engineer	1	R1 200 000	R520	R520
Officer	12	R500 000	R200	R2 400
Deckhand	30	R200 000	R80	R2 400
Total per hour for crew wages				R5 840
Miscellaneous hourly expenses				R4 000
Grand total per hour				R9 840

C.2 Calculation of fuel cost

The estimated fuel consumption for the SAAII of 179 g/kWh was communicated by the ship's officers (Lighthelm, 2020). Assuming the price for diesel at R13 per litre with a fuel density of $\rho = 0.8323$ kg/litre, the price per kilogram can be calculated as

$$\text{Price per kilogram} = R13 \times \rho \approx R15.62/\text{kg}. \quad (\text{C.2.1})$$

The price per kWh can be approximated as

$$\text{Price per kWh} = R15.62/\text{kg} \times 179\text{g/kWh} \approx R2.80/\text{kWh}. \quad (\text{C.2.2})$$

List of References

- Ansorge, I.J., Skelton, P., Bekker, A., de Bruyn, N.P., Butterworth, D., Cilliers, P., Cooper, J., Cowan, D.A., Dorrington, R., Fawcett, S. *et al.* (2017). Exploring South Africa's southern frontier: A 20-year vision for polar research through the South African National Antarctic Programme. *South African Journal of Science*, vol. 113, no. 5-6, pp. 1-7.
- Bekker, A. (2017). From (big) data to insight—a roadmap for the SA Agulhas II. *High Performance Marine Vehicles*, pp. 212-223.
- Bekker, A., Lu, L., van Zijl, C., Matthee, J. and Kujala, P. (2019). Correlation between bow ice loads and operational responses during ice navigation in the weddell sea. In: *Port and Ocean Engineering under Artic Conditions*. Delft, Netherlands.
- Bialystocki, N. and Konovessis, D. (2016). On the estimation of ship's fuel consumption and speed curve: A statistical approach. *Journal of Ocean Engineering and Science*.
- Bishop, C. (2006). *Pattern recognition and machine learning*. Springer. ISBN 978-0387-31073-2.
- Chollet, F. (2018). *Deep Learning with Python*. Manning Publications Co., 20 Baldwin Road, Shelter Island, NY 11964. ISBN 9781617294433.
- Cosci, G. (2018). The IMO's GHG strategy: A step toward meeting the 2°target? Tech. Rep., Climate Institute, 1201 New York Avenue, NW, Suite 400 Washington DC. Available: <http://climate.org> [2019, Jan 31].
- Devanunthan, N. (2019 November). Value, function, challenges and costs of operating the SA Agulhas II [presentation]. Department of Environmental Affairs. SAAII Mini Conference. Stellenbosch, South Africa.
- DNV-GL (2016). Making your asset smarter with the digital twin. Article on DNV GL website. Available: <https://www.dnvgl.com/article/> [2020, 1 May].
- DNV-GL (2017 January). Rules for classification - ships. Tech. Rep., Det Norske Veritas. Available: <https://rules.dnvgl.com> [2020, Apr 14].
- Gkerekos, C., Lazakis, I. and Theotokatos, G. (2019). Machine learning models for predicting ship main engine fuel oil consumption: A comparative study. *Ocean Engineering*, vol. 188, p. 106282.

- Géron, A. (2017). *Hands On Machine Learning with Scikit-Learn and TensorFlow*. O'Reilly Media Inc.
- Halevy, A., Norvig, P. and Pereira, F. (2009). The unreasonable effectiveness of data. *IEEE Intelligent Systems*, vol. 24, no. 2, pp. 8–12.
- IMO (2009). Guideline for voluntary use of the ship energy efficiency operational indicator (eoi). Tech. Rep., International Maritime Organization (IMO), 4 Albert Embarkment, London. Available: <http://www.imo.org/> [2020, Aug 10].
- IMO (2015). Third IMO greenhouse gas study 2014. Tech. Rep., International Maritime Organization (IMO), 4 Albert Embarkment, London.
- Johnson, H. and Andersson, K. (2016). Barriers to energy efficiency in shipping. *WMU Journal of Maritime Affairs*, vol. 15, no. 1, pp. 79–96.
- Li, F., Goerlandt, F. and Kujala, P. (2020). Numerical simulation of ship performance in level ice: A framework and a model. *Applied Ocean Research*, vol. 102, p. 102288.
- Ligthelm, F. (2020 January). Main engine fuel consumption. E-mail correspondence with Bekker, A.
- Lindfield, G.R. and Penny, J. (2017). *Introduction to Nature-Inspired Optimization*. San Diego: Elsevier Science and Technology. ISBN 9780128036662.
- Olmer, N., Comer, B., Roy, B., Mao, X. and Rutherford, D. (2017). Greenhouse gas emissions from global shipping, 2013-2015. Tech. Rep., The International Council on Clean Transportation (ICCT). Available: <https://www.theicct.org/> [2019, Jan 31].
- Omer, H. and Bekker, A. (2016). Detection of wave slamming sites from ship deflections. *Research and Development Journal of South Africa*, vol. 32, pp. 50–57.
- Omer, H. and Bekker, A. (2017). Human responses to wave slamming vibration on a polar supply and research vessel. *Applied Ergonomics*.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830. Available: <https://scikit-learn.org/> [2020, May 18].
- Sandru, A. (2018). *Sea Ice Field Analysis Using Machine Vision*. Master's thesis, School of Electrical Engineering, Aalto University, Espoo, Finland.
- Snyman, J. (2005). *Practical Mathematical Optimization. An Introduction to Basic Optimization Theory and Classical and New Gradient-Based Algorithms*. 1st edn. Springer US.

- Soal, K., Bekker, A. and Bienert, J. (2015). Structural vibration analysis on the polar supply and research vessel the SA Agulhas II in Antarctica. In: *Proceedings of the International Conference on Port and Ocean Engineering Under Arctic Conditions*.
- Wang, K., Yan, X., Yuan, Y., Jiang, X., Lin, X. and Negenborn, R. (2018). Dynamic optimization of ship energy efficiency considering time-varying environmental factors. *Transportation Research*.
- Yoo, B. and Kim, J. (2018). Probabilistic modeling of ship powering performance using full-scale operational data. *Applied Ocean Research*.
- Zhang, C., Zhang, D., Zhang, M. and Mao, W. (2019). Data-driven ship energy efficiency analysis and optimization model for route planning in ice-covered Arctic waters. *Ocean Engineering*.