Supervised Machine-Learning as a Decision Support Aid in Sea Lice Control for Norwegian Salmon Farmers

MSc Research Project
MSCDATOP

Julia Streckfuss
Student ID: 17158192

School of Computing
National College of Ireland

Supervisor: Dr Catherine Mulwa

National College of Ireland
Project Submission Sheet
School of Computing

Student Name:	Julia Streckfuss
Student ID:	17158192
Programme:	MSCDATOP
Year:	2023
Module:	MSc Research Project
Supervisor:	Dr Catherine Mulwa
Submission Due Date:	14/08/2023
Project Title:	Supervised Machine-Learning as a Decision Support Aid in Sea Lice Control for Norwegian Salmon Farmers
Word Count:	11598
Page Count:	32

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	J. Streckfuss
Date:	17th September 2023

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	☐
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	☐
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	☐

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Supervised Machine-Learning as a Decision Support Aid in Sea Lice Control for Norwegian Salmon Farmers

Julia Streckfuss
17158192

Abstract

In the country's goal of five-fold salmon production expansion until 2050, the biggest salmon producing nation, Norway, has put the strictest average female lice thresholds globally in place to adhere to standards of sustainable aquaculture. Previously conducted research has had the focus to simulate lice dispersal especially with the aim of estimating lice pressure for passing wild salmon fish, however no available academic research, to the knowledge of the researcher, exists which helps salmon farmers to comply with government-set lice thresholds. The presented study fills this gap by utilizing supervised classification models to identify the best point of warning to take preventative counter-actions for lice threshold exceedance, to classify whether farm localities are expected to exceed thresholds for future points in time, and to propose treatments proven most successful to them as well as to provide risk reduction estimates to avert the risk of exceeding if this is achievable. The focus of the study therefore lies at predicting the risk of exceedance of government-set lice thresholds at individual farm level and to provide treatment recommendations at a time when exceedance can still be prevented, which, to the knowledge of the researcher, has not been attempted by any other study published in the research domain yet. Furthermore, this study is the first study, to the knowledge of the researcher, which uses machine learning, in particular classification models, for the purpose of estimating lice counts at individual farm level in Norway and beyond. The strongest performing models have been found to be Random Forest, XGBoost, and AdaBoost with ROC-AUC scores between 0.997 and 0.985. Little to no degradation of models was found in comparing classifier performance from warning point 4 to 8 weeks prior to exceedance, and while metrics are similarly strong for both precision and recall, the preferred methods show slightly higher scores of precision than recall which demonstrates their ability to keep false positives low to mitigate the risk of unnecessary treatment costs encouraged.

1 Introduction

While previously considered a commonly occurring and generally harmless parasiticide on wild salmon fish returning from oceanic feeding sites by Norwegian fishermen (Misund 2019, p.403), salmon lice resulting from large-scale net pen farming and the associated risk of them spreading to wild fish have received increased political attention in the debate over sustainable aquaculture during the past two decades of salmon farming expansion in Norway (Misund 2019, p.405). Following scientific research attesting to the potential risk of farms transmitting lice to wild salmon, especially during migration periods, the country has gradually increased preventive measures to counter-act this risk and has implemented the strictest average female lice-thresholds across all salmon producing nations together with a so-called 'traffic-light system' which regulates aquaculture growth if elevated lice levels combined with other disadvantageous environmental factors could have the potential to harm wildlife (Vormedal & Larsen 2021, p.11-12).

Research in the domain has had the primary focus of quantifying the risk for salmon lice in net-pens to spread to surrounding wild fish populations (Johansen et al. 2011, Torrissen et al. 2013, Kristoffersen et al. 2018) by means of hydrodynamic and lice dispersal models combining knowledge on movement and growth of lice and lice particles (Sandvik et al. 2016, Myksvoll et al. 2018, Sandvik et al. 2020, Bøhn et al. 2022) under given environmental conditions such as temperature and salinity (Stige et al. 2021, Rittenhouse et al. 2016) [1] while another dominant research area covers models of lice transmission within and across farms (Aldrin et al. 2013, Harrington et al. 2023) [2] and of lice abundance at individual farm level (Aldrin et al. 2017, 2019). Most of the models discussed in the literature are based on scientific principles aiming at the explainability of lice abundance and dispersal through contributing physical and/or biological factors or at precisely simulating infestation pressure at different parts of Norway's coastline under given farming (also including treatment) and environmental conditions (Salama et al. 2013, Aldrin et al. 2017).[3]

Another alley of research has been focusing on the economic impact of salmon lice for salmon farmers (Liu & Bjelland 2014, Abolofia et al. 2017) regarding decreased growth rates or mortality (Walde et al. 2021, 2022) caused by infestation or treatment of infestation, and in relation to cost of treatment itself (Dean et al. 2021).

Lastly, advantages and disadvantages of different treatment options have been discussed covering environmental concerns over originally applied substance-based treatments and their potential to cause resistance (Revie et al. 2005, Aaen et al. 2015) [4] in salmon lice over time, as well as newer methods such as cleaner fish (Blanco Gonzalez & de Boer 2017) or mechanical removal.

While previous research in the domain has contributed largely to the body of knowledge on growth stages of lice, dispersal of lice under varying conditions, effect of treatment methods on the health or growth of fish, and has been helpful in identifying contributors to lice abundance and risk areas for spread to wild salmon fish, only a few studies estimate lice abundance at individual farm level, and there is no evidence in the literature, to the knowledge of the researcher, for the existence of a predictive systems for farmers that has the ability to warn them of the risk of exceeding the permitted lice threshold set by the Norwegian government for a given point in the future, and that is able to guide them in taking effective measures against lice infestations before they occur/get to an uncontrollable state.

The Research Question thereby is as follows:

RQ: "*How accurately can classification models (CART, Random Forest, AdaBoost, Gradient Boosting, KNN, XGBoost) solely utilizing publicly available data predict if sal-*

[1] Rittenhouse et al., 2016: Canadian study.

[2] Canadian studies.

[3] Salama et al., 2013: Scottish study.

[4] Revie et al., 2005: Scotland; Aaen et al., 2015: Resistance mechanisms across main salmon producing nations.

mon farm localities in Norway will exceed the permitted lice threshold set by the government for a given point in the future?"

Sub-Research-Questions are:

SUB-RQ 1: "*At what point in time before exceeding the permitted lice threshold can a warning based on that prediction be beneficial to salmon farmers in Norway in allowing them to take counteractions to prevent this outcome?*"

SUB-RQ 2: "*To what extent can an acceptable trade-off between precision and recall be achieved to minimize the level of false positives as a risk mitigation method for unnecessary treatment costs while identifying as many exceeders as possible?*"

SUB-RQ 3: "*To what extent does the predictive power of a model suffer with predicting for a later point in time in the future?*"

SUB-RQ 4: "*To what extent can beneficial treatment recommendations for farmers be obtained by suggesting treatments that have proven most efficient by other farming localities with similar attributes?*"

To answer the research questions outlined, the research aims to meet the following objectives outlined inTable 1 below:

Table 1: Research Objectives

ID	Description
1	Extract weekly fish health data from BarentsWatch API. Clean data, interpolate missing values, assess data quality. Merge weekly overview data per locality with detailed farm data on capacity, water temperature, treatments applied. Merge manually downloaded thresholds data with merged data and derive binary outcome variable of average female lice threshold exceedance.
2	Create features for classification models which incorporate information on farming cycle, treatment history, lice history, current lice count, and sea temperature of individual farms, and aggregated lice figures and sea temperature for farms in 10,20, and 50 km haversine distance.
3	By means of a CART predicting successful prevention of lice thresholds for up to two months, obtain warning signals farmers act upon that can successfully prevent exceeding of thresholds; derive summary statistics for when farmers exceed when missing to act upon found signals.
4	Implement classification models (CART, Random Forest, AdaBoost, Gradient Boosting, KNN, XGBoost) to predict lice threshold exceedance for 3 different points in the future, and compare trade-off between point of warning and model performance.
5	Evaluate models for balance between recall and precision to assure models perform strongly in identifying exceeders of lice thresholds while keeping false-positives as low as possible.
6	Integrate current week's recorded treatments in best performing classifier from Obj. 4 to demonstrate effect of different applied treatment options over no treatment applied and consequent risk reduction and prevention of exceeding of threshold.

Anticipated weaknesses of the research are that approximately 6 percent of data in the outcome variable are missing which is assumed to be be due to farms being exempt from reporting on lice in specific weeks during the farming cycle [5] (Vormedal & Larsen 2021, p.6), and that while average lice counts are provided by the farms on a weekly basis, biomass (total count of salmon fish) is updated only monthly by the farmers (Sandvik et al. 2020, p.747). Furthermore, due to strict regulations, the exceeding of the permitted lice count is a less common outcome than is the staying within limits of permitted lice levels. Thorough measures were taken to assure that missing values, especially in the outcome variable, were interpolated in a realistic manner, and that balancing of classes in the outcome variable was conducted prior to modelling.

[5] One example for this is them being exempt if water temperature falls below 4 degrees Celsius.

Lastly, data for cost of treatment options was not available to the researcher and therefore recommendations on possible treatments cannot include an elaboration on monetary risks of each treatment over others or over not taking the measure; this responsibility will remain with the farmers.

Assumptions are that during fallow periods or periods of low temperature, lice count is not reported by the farmers (Vormedal & Larsen 2021), and that during production cycles, farmers provide accurate lice counts.

The research presented here fills the gap of previously conducted research in the fields by means of producing a deterministic classification system that warns farmers ahead of a lice outbreak at a locality, and that provides them with treatment options to counteract the threat. Furthermore, it integrates research on the best point of warning for the beneficiaries of the farmers to assure that warnings and proposed countermeasures are provided at a time when actions can still be taken to prevent the outcome.

The achieved impact therefore is the provision of a decision aid tool for farmers to understand the risk of exceeding the allowed average female lice-threshold at a specific point in time in the future, to be shown the expected outcome of application of most efficient treatment options at the time of warning to avert the risk, and to bring their own knowledge and expertise to decide which measures prove appropriate and cost-effective.

The rest of the technical report starts with an in-depth analysis of current research in the problem domain in chapter 2, which differentiates and contrasts contributions of those studies to the knowledge of research with that of presented research. In chapter 3, the research methodology is outlined, and the rationale for choosing it is provided. Following that, the design specifications are discussed in chapter 4 before the implementation of the project solution is presented, and model results are evaluated. In chapter 6, main results of the study are summarized and evaluated against set research objectives, before conclusions and future work are discussed to outline where future research can continue in the given subject matter.

2 Related Work – Salmon Lice Dispersal, Treatment, and Abundance in Norway and Beyond in the Last 18 Years

Along with the expansion of salmon production in Norway and globally, multiple studies on sea lice dispersal, abundance, and treatment have been conducted in the last 18 years contributing to knowledge discovery in the domain in different ways.

While no time or geographic limit was set to compile sources of relevance for the conducted study, most of the research around salmon lice abundance, dispersal, and treatment can be found in the last 10-12 years; around the time (2012) when lice reporting by farmers in Norway was enforced weekly as opposed to monthly before (Aldrin et al. 2013, p.9). Most sources discussed in the following literature review relate to salmon lice management in Norway, however some of the studies presented here were conducted in Eastern Canada, Chile, or the Faroe Islands, and were included because they describe to solve similar problems to the ones addressed in this study.

While the problem of salmon lice has been explored through research in many ways, not all of them will be discussed in the following review. Topics that are consciously omitted revolve around image processing as a means of lice counting, and the financial impact of salmon lice to farmers by means of decreased growth rates and increased mortality through infestations and treatments of those, and by means of the cost of treatment itself. While these areas are relevant to research in the domain overall, they are not in scope for the research proposed in this study. Research on treatment efficiency at different times during the infestation process however is included to legitimize the need for an early warning system for the farmers.

The review starts with a collection of studies conducted to provide insights around the spread of salmon lice from net-pen farms to wild fish populations, because these studies relate directly to the government's initiatives to keep salmon aquaculture sustainable during the expansion of production, which appears to be the primary focus of studies around salmon lice. Next, the review discusses research around the dispersal of salmon lice within farm localities and between farms, and contrasts the aim of those studies to the one presented here. As a next step, studies around treatment, treatment efficiency and timing are discussed before lastly, lice abundance estimate studies for individual farm localities are in the focus of the review. The last section lies closest to the aim of the study presented here.

2.1 A Review of Studies on Salmon Lice Spread to Wild Salmon and Identified Gaps

While the aim of the conducted study lies at counter-acting the exceeding of government-set lice-thresholds and not at analysing whether instantiated regulations prove effective at minimizing the risk for lice occurring at salmon farms to spread to wild populations, it is still found beneficial to integrate research on that sub domain in the following review to discuss the value of its contributions to the research domain and to demonstrate how the findings of those studies serve a different beneficiary; namely the government in being able to take informed decisions on setting regulations around salmon lice. The research conducted here however aims at benefiting the farmers in complying with set regulations.

One of the most prominent studies that aims to simulate lice infestation pressure along Norway's coast to complement the manually conducted national salmon lice monitoring program in Norway was published by Sandvik et al in 2016 (Sandvik et al. 2016). The researchers feed outputs from a known oceanic model called NorKyst800 providing hourly data points on currents, salinity and temperature into a salmon lice dispersal model that simulates active lice swimming behaviour and passive drift through the currents to predict infestation levels at different farm locations in an area called Hardangerfjord on the west coast of Norway for the years 2012 to 2015 for the months of April to July. Sandvik et al. evaluate model performance by means of comparing modelled infestation pressure levels to lice counts of salmon fish that are captured in sentinel cages to demonstrate the effect infestation pressure on farms can have on wild salmon populations. For the years 2012-2015, the researchers predict 4 classes with a POD, what they refer to as 'probability of detection', of 78 percent, and for 2015, they predict two class-levels (elevated vs non-elevated infestation pressure) with a POD of 89 percent. While Sandvik et al. refer to their model objective as 'predicting' lice infestation pressure, the technical report does not comment on the point in time for which those predictions are being made; instead, the researchers describing the value of the model lying in the ability to simulate scenarios of optimizing production zones, strategies of control etc appears to confirm that it is a real-time simulation model meant to support government decisions as opposed to being a predictive forecasting model for salmon farmers (Sandvik et al. 2016).

A similar study estimating lice infestation pressure, however along the entire Norwegian coastline, was published by Myksvoll et al. in 2018 (Myksvoll et al. 2018). The research follows a mechanistic approach which combines a hydrodynamic model integrating further environmental factors such as circulation, wind, and regional differences with a behavioural model for salmon lice and validates the model outcomes against lice abundance in wild char and trout fish. The researchers measure model performance by means of Spearman rank correlation, and with a value above 0.7 for areas bigger than 13x13 state the model to be suitable for evaluating pressure of infestation including the pressure between stations of monitoring (Myksvoll et al. 2018).

A third study following both studies described above was published in 2020 by Sandvik et al. (Sandvik et al. 2020), It extends the binary prediction of elevated and non-elevated lice infestation pressure from the previous study to determine three levels of lice infestation pressure in the Hardangerfjord area for the years 2012-2017. The study utilises the previously described oceanic model, a lice behavioural model which includes additional parameters derived from lab experiments and includes predictors such as weekly lice counts and monthly biomass figures. Model performance is measured in comparison to lice counts in sentinel cages, and a hit rate of 0.70, a false alarm rate of 0.11, and a POD of 0.76 is achieved. The researchers state the model to be generalizable for the entire Norwegian coast (Sandvik et al. 2020).

A study by Bøhn et al. published in 2022 (Bøhn et al. 2022) applies Bayesian statistics to model the risk for wild sea trout to exceed acceptable levels of lice infestation due to infestation pressure in the environment, salinity in the ocean, and speed of currents and utilizes the previously described hydrodynamic dispersal model in Myksvoll et al. and Sandvik et al. (2020) to obtain those parameters. The researchers combine a Bernoulli and a Gamma mixed model in a so-called ZAG model and derive importance of examined parameters in determining infestation levels in wild trout. They reach the conclusion that lice infestation pressure in the environment has the strongest impact, followed by currents while salinity shows a non-significant impact when predicting lice in wild fish.

One of the most recent studies contributing to the knowledge on estimation of infestation pressure was published by Stige et al. in 2021 (Stige et al. 2021). In following a statistical mechanistic approach, the researchers were able to prove that the accuracy of a baseline model provided by Kristoffersen et al. in 2018 (Kristoffersen et al. 2018) can be improved by changing the formulation of the existing model by means of integrating infectivity dependent on temperature and salinity and by changing the approximation for time of development from one that is based on degree-days to one that is fractional (Stige et al. 2021).

While all the above studies contribute to assessing the risk of lice infestation pressure on wild fish and can thereby support government decisions on regulating lice by means of adjustment of thresholds, farm capacity and preventive measures, the beneficiary of those studies is the Norwegian government, and not the farmers, as in the conducted study described here. Furthermore, the studies appear to be simulation studies aiming to integrate science-based knowledge on hydrodynamics and lice behaviour/growth and the interaction of those with the aim to provide real-time estimates on lice infestation pressure along the Norwegian coast without the need to manually monitor lice counts by means of manual counting of lice on fish that are kept in sentinel cages. The aim of the conducted study however is to function as a warning system for farmers to prevent infestation pressure to be reaching high levels in the first place.

2.2 Critical Review of Research on Salmon Lice Dispersal Within and Across Farms and Identified Gaps

Another large area of research aims to find insights on how lice disperse within and across farms in Norwegian salmon farming.

One of the earlier studies on this topic was published by Aldrin et al. in 2013 (Aldrin et al. 2013). It describes a stochastic space-time model that is able to model monthly sea lice abundance at farm level in Norway based on lice counts in months prior, on lice levels at farms nearby and on other external unknown factors (e.g., lice from wild fish) that may increase lice infestation levels at a given farm while taking predictive variables such as water temperature, distance to surrounding farms and biomass at farm level into account. The researchers point out that estimates can be improved with weekly lice count data and that the inclusion of treatment may prove beneficial. They find that lice abundance is primarily due to infestation levels within farms, followed by external infestation pressure from surrounding farms and that other external sources for infestation show the smallest effect on the outcome. Aldrin et al. describe their contribution to the research area in providing a framework for tools able to simulate lice abundance and spread under set control measures (Aldrin et al. 2013).

Kristoffersen et al. published a similar study following a statistical modelling approach in 2014 (Kristoffersen et al. 2014) in which they calculate the farm-internal infection pressure and the times of development from egg-state to pre-adults and male adult lice stage and utilize the described model by Aldrin et al. as a function for seaway distance between farms. They do this for newly stocked farms to assure no previous lice exposure and find that external infestation pressure constitutes a main predictor for development of lice at individual farm level. They claim that their study can be used as a base for real-time models displaying local lice abundance (Kristoffersen et al. 2014).

A related study conducted for the Faroe Islands was published by Kragesteen et al. in 2021 (Kragesteen et al. 2021). The authors present a time series analysis on salmon lice data from 2011 to 2018 which estimates the external infestation pressure and the growth rate for salmon lice at individual farm level. The researchers state their work on parameter estimates to prove beneficial in the development of a dynamical lice population model to support management strategies (Kragesteen et al. 2021).

One of the most recent studies in the problem domain of inter-farm lice dispersal was published in British Columbia, Canada, which calculates the distribution of arrival time for lice dispersal from one farm to another by means of a mechanistic model integrating a hydrodynamic model combined with a model for particle tracking. The main findings of the simulation study are that spacing between neighbouring farms can impact infection across farms due to longer period of times for the lice to mature on the way to the neighbouring farm given currents are strong enough and that warmer temperatures can have an intensifying effect (Harrington et al. 2023).

Above studies are beneficial in differentiating the difference between within farm infestation pressure and risk of exposure to lice from surrounding farms and may be used in understanding risks originating from a farm locality itself vs the risk of exposure to lice from surrounding farms and knowledge could be transferred to estimations at individual farm level, however they do not utilize models to predict lice threshold exceeding for a future point in time or provide recommendations on counter-active measures to prevent lice spread between farms. Those studies may be beneficial in organizing coordinated lice prevention measures in farm systems, but this deviates from the objective of described study which goes to the core of the problem; namely preventing lice to exceed the recommended thresholds so that risk for cross-farm dispersal is minimised.

2.3 Critique of Academic Work on Salmon Lice Treatment and Identified Gaps

While Norway's treatment strategy moved from a reactive to a proactive approach due to the government's requirement of keeping lice levels under 0.5 for non-migration and 0.2 for migration periods instead of priorily enforcing treatment when lice levels of 0.5 [6] and higher were reported, there is, to the knowledge of the researcher, no other study published which models the effects of different lice treatments against a proactively set threshold. There are however studies on different aspects around treatment, timing of treatment, and treatment efficiency which are outlined in the following section of this review.

One of the first studies conducted to understand lice development patterns taking treatment measures of in-feed and bath treatments into account was published in 2011 by Gettinby et.al (Gettinby et al. 2011). The study mathematically models lice abundance for spring and autumn stocking farms and derives insights on differences in treatment patterns and efficiency for those two groups for the region of the Hardangefjord for 2004 till 2007. The contribution of this work lies in the comparison of observed patterns for spring and autumn stocked salmon while providing insights on how farmers' decisions are influenced by economic aspects such as preferring in-feed treatments in the first year of stocking due to the size of the fish being small enough to making it a cost-efficient treatment method which in addition can have a 'residual effect' of lasting for about 2-5 months (Gettinby et al. 2011, p.168)[7]. A disadvantageous aspect of this study is the fact that there are no performance measures provided which could allow for a comparison of modelled vs observed lice.

Murray et al. published a simulation study in 2016 for the area of Scotland which applies a simplified model to assess effects of synchronized fallowing and treating of salmon lice for management areas consisting of different individual farms. It finds that fallowing is the most beneficial way to keep lice levels at a controllable level, followed by synchronized treatment depending on treatment efficiency and hydrodynamic factors (Murray & Salama 2016, p.39). The authors state however that the model is basic and that local conditions are needed to get clearer estimates on synchronized treatment benefits (Murray & Salama 2016, p.46).

Another study aimed at demonstrating the benefit of controlled lice management, this time related to the area of Pacific Canada, was published by Peacock et al. in 2020 (Peacock et al. 2020). The researchers present a mechanistic model that, like studies described in section 2.1 of this literature review, estimates lice count in wild salmon populations depending on timing of migration based on temporal dynamics of lice dispersal in farmed salmon environments. This study however explores scenarios of timing of lice control and coordination of timing of neighbouring farms to support management of salmon farms in effectively controlling the parasiticide to protect migrating wild fish

[6] Vormedal and Larsen., 2021: for green licenses, thresholds are described to be even lower at 0.1 and 0.25 respectively. The green licence attribute obtained from the API however proved unrealiable and couldn't be utilized for the study here

[7] Due to inaccuracies found in farming cycles which doesn't always allow to derive year of stocking, this information could not be utilized for the study conducted here

populations and finds that proactive and controlled measures show the best protection of wild salmon, especially if migrations happen earlier than expected (p.6 Peacock et al. 2020, p.30).

An exploratory case study on the timing and effectiveness of lice control measures taken in a selected production area in Norway in 2013-2018 was published in 2019 by Jevne et al (Jevne & Reitan 2019). The study shows implications on different control practices over the years for the selected production area and includes an analysis of measure effectiveness given variations in environmental variables such as sea temperature. The main findings of the study are that even lice levels under the allowed thresholds should not be left untreated due to difficulties in controlling lice growth with high infestation pressure and that cleaner fish can delay lice growth for the start of a production cycle (Jevne & Reitan 2019, p.1583).

Another study by Jeong et al. published in 2021 supports the findings from the previously described research in attesting to the fact that it is found to be more beneficial to prevent infestation pressure as opposed to treating salmon that is lice infested (Jeong et al. 2021, p.6). While data on preventative measures is only partially available for the presented study, the findings described in Jeong et al.'s research promote the idea of an early warning system, as per the aim of the study conducted here.

While the last study included in this section of the review is not related to salmon lice, it is found beneficial to include it to show a similar approach in providing suggestions for corrective measures in aquaculture. Nayan et al. published a study in 2021 which analysis water quality in Bangladesh and effects it can have on fish health. The research applies machine learning to predict water quality and the risk for fish disease from water degradation and proposes countermeasures to avert this risk (Nayan et al. 2021, p.1). The study applies gradient boosting for that purpose and achieves a mean square error of 1.2 and an R squared of 75 percent. While the approach appears to be useful to farmers in identifying risks associated with water quality, the study does not train the algorithm against a ground truth of actual disease cases, however instead derives disease predictions from medical findings. (Nayan et al. 2021, p.6). Due to the prediction being one of a continuous variable (water quality) as opposed to classification of class in the study conducted here, the performance metrics used are not comparable to those in this study.

2.4 Review of Studies on Salmon Lice Abundance at Farm Localities and Identified Gaps

While there is no evidence for machine learning applied to estimate lice counts at individual farm level, other types of approaches have been applied to that purpose. One of those studies was published in 2017 by Aldrin et al. which follows a stage-structured Bayesian hierarchical approach (Aldrin et al. 2017). The researchers state that the benefit of their approach versus previously conducted studies on lice dynamics [8] lies in being fitted to actual lice counts at farm and cage level while assuring values for parameters were kept aligned with plausible biological ranges in the prior distributions whereas previous research being autoregressive was mostly dependent on data derived from experiments in labs (Aldrin et al. 2017, p.334). Unlike previously discussed research, Aldrin et al. show actual predictions for several weeks into the future for 3 different lice stages including

[8] Researchers categorize Aldrin et al. 2013 and Kristoffersen et al., 2014 from 2.2 as those autoregressive approaches.

adult female lice. They also integrate data on preventive measures such as cleaner fish and substance treatments, which is aligned with the approach followed in the study here. The approach is shown to reach a mean absolute percentage error of 74 for adult female lice in the first 30 days of prediction and of 104 for the remaining time of the prediction period and a 97 percent coverage of the 95 percent credible intervals for counts of adult female lice in the first 30 days of prediction and one of 87 percent for the later prediction periods of the production cycle [9]. The researchers conclude this study to serve as a 'mathematical laboratory' to see control efforts at individual farm level in favour of a coordinated control approach in an inter-connected farm system (Aldrin et al. 2017, p.347).

Another study on sea lice abundance at individual farm level was conducted by St.-Hilaire et al. for Chilean salmon production in 2018 (St-Hilaire et al. 2018). The researchers describe their model as a 'linear state-space model without stochastic component in the state equation' and predict lice abundance at farm level for isolated farms for a period of 8 upcoming weeks (St-Hilaire et al. 2018). Their model can predict adult sea lice numbers on average within 0.96 lice per salmon and for specific farms may predict within 1-2 lice deviating from actual lice counts. While the model is valuable in identifying lice abundance within individual farms, it is not taking inter-farm cross-infestation into account and is therefore only able to predict for those isolated cases.

A second study by Aldrin et al. from the year 2019 predicts lice counts at individual farm level for 1-8 weeks into the future in a model that is described as 'partly stage-structured' since it integrates a model on adult female lice and on other mobile stages of lice and does that based on functions including covariates of biological and physical nature (Aldrin et al. 2019, p.2) in an auto-regressive simulation study design. The researchers utilize publicly available and private farm data and estimate model parameters related to lag of lice counts, temperature dependency and others in relation to lice abundance of adult female lice and other mobile lice. Model performance against real lice abundance is only shown visually and no scientific measures are provided to understand reliability of the model.

Elghafghuf et al. published a comparative study on different sea lice estimation methods for the area of Bay of Fundy in New Brunswick Canada in 2020 by means of applying multivariate state-space models with the long-term goal to creating a tool which can predict lice abundance for individual farms similarly to the base goal of the study conducted here (Elghafghuf et al. 2020, p.2). The researchers compare 5 approaches for the timeframe May 2010 to October 2016 to predict pre-adult females and males combined with adult males (PAAM) and conclude that models which only include EIP (external infestation pressure) outperform models that only include IIP (internal infestation pressure) and that models which included both of these components performed the best. They compare model performance by means of Akaike's Information Criterion (AIC) and furthermore derive insights on how farm-connectivity requires local treatment strategies and on how female lice counts 3 weeks prior is a good indicator for PAAM at the time of prediction (Elghafghuf et al. 2020, p.11).

A further study by the same researchers published in 2021 applies the preferred method again using a multivariate autoregressive state-space model (MARSS) which attempts a rolling prediction method which is compared against a similar approach which requires re-fitting of the model. It is found that their model can reduce errors of pre-

[9]The authors do predict for chalimus and mobile stages also, but adult female lice is found most relevant in scope of study presented here.

diction while bringing advantages of less computation required (Elghafghuf et al. 2021, p.1). The researchers use absolute prediction error and mean absolute prediction error as measures to evaluate model performance.

Described research has contributed to the problem domain in the way that it changed the focus from estimation of infestation pressure on wild fish to dynamics at play amongst inter-connected farms and has brought insights into understanding the impact of internal vs external infestation pressure. Models described however fail to compare predictions against government-set thresholds and are therefore only partially beneficial for farmers. Also, while slowly moving towards data-driven fitting approaches they still predominantly work off biological and physical concepts partially derived from lab experiments and miss the opportunity of deriving the model from found vs assumed (attested) relationships in the data/of parameters through machine learning. One last disadvantage of those studies is that they lack in interpretability of model performance and often only showcase predictions for specific timeframes while machine learning applies clearer evaluation methods which make it easier to assess risk of false predictions vs risk of missing correct predictions. One positive thing to note however is the fact that the studies address data quality issues and how those were overcome, which has proven beneficial for data cleaning and feature engineering tasks required for the study presented here.

Overall, the review of related work on salmon lice has been critical in choosing and designing features for the models implemented in this study and has proven beneficial in shaping the applied research methodology outlined next in this report.

3 Scientific Research Methodology

The following chapter describes the applied methodology in the conducted research. It begins with an overview of the methodological framework chosen, continues with specifications on data collection, data exploration and cleaning, data transformation and feature engineering applied, and concludes with machine learning methods and subsequent evaluation techniques thereof selected to meet outlined research objectives. Furthermore, the tools used for the different steps in the process are described, and explanations on any calculations made while transforming the data and on underlying assumptions motivating those transformations are provided.

3.1 Lice Threshold Methodology Approach

The methodological framework chosen for conducted study as visualized in Figure 1 below is derived from the commonly applied CRISP-DM (Cross Industry Standard Process for Data Mining) framework.

Due to its integration of the understanding of the problem domain, it places an emphasis on the usability of a proposed solution to the beneficiary, which is the priority of the study conducted here. It assures that previous research in the domain is analysed prior to conducting the study to focus on identified gaps, but to also take learnings of other researchers into account for conducted research.

The methodology applied combines knowledge inclusion of previous research, knowledge discovery in terms of proposing a beneficial point of warning, and demonstration of effect of choice of different treatment options to guide salmon farmers in successfully complying with government standards, which in return will allow them to maintain or even grow their current farming capacity while protecting the health of their fish and of the fish in surrounding farm localities.

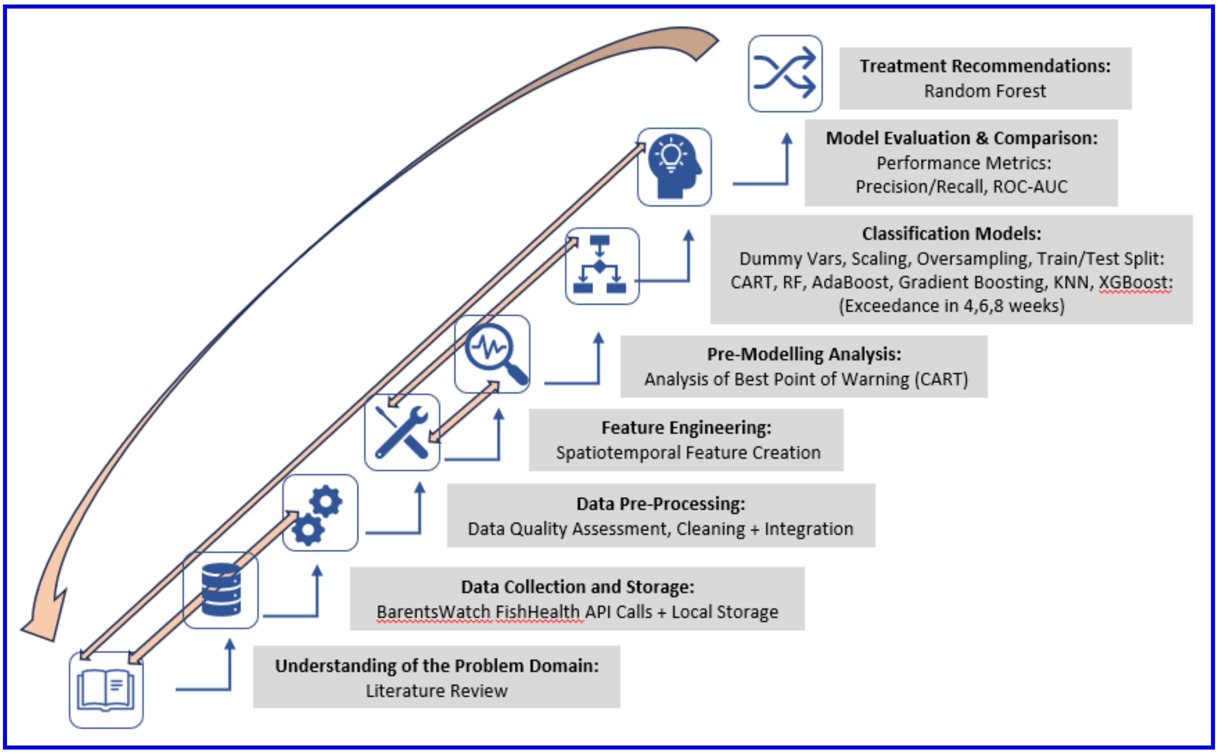

Figure 1: Methodological Research Framework

3.2 Data Collection

After thorough analysis of previous research in the problem domain (as discussed in chapter 2), data was collected from the publicly available Fishhealth API accessible through the BarentsWatch website (BarentsWatch 2023*a*)(as per Obj. 1 of this research). BarentsWatch is the official information system for monitoring of marine and coastal activity in Norway, and as a governmental entity under the authority of the Ministry of Trade, Industry, and Fisheries, collaborates with several ministries, administrative agencies and research institutes in the country (BarentsWatch 2023*b*). The source of the data, used by numerous researchers over the years, can thereby be assessed as reliable and accredited and was therefore found to be suitable for the purpose of the study presented. After free-of-cost sign-up to the API service, data was obtained by means of request calls in Python using bearer authentication tokens which need to be manually retrieved through curl commands in the prompt. Anaconda Prompt and Jupyter Notebooks were used for this activity. The API call for weekly data collected records for the period 2012 (week 1) until 2023 (week 15), and with 1014038 records and 2581 unique farm localities combines data for all farm localities in Norway per week from the beginning of weekly fish health collection in Norway as opposed to monthly collection prior to 2012. Lice medication events and escapes had to be retrieved by farm and year, and detailed farm by farm and week. Due to long processing times for data retrieval especially for the detailed farm data (2000 farms *52 weeks * 10 years = 1 million separate API calls) and losing of validity for the tokens while running the iteratively sequenced API call loops, only data for relevant farms were selected for each year for those additional attributes (400k API calls). Relevant were considered farms that have salmonoids, and which are not fallow, which means that they are in a production cycle. After joining part of the treatment data and detailed farm data with weekly data, the records were reduced to active farm data. In addition to selecting only 'non-fallow' records that are marked as having salmonoids, another attribute which identifies farm localities to being on land was used to exclude those. 327238 active farm records remained after filtering, and 1147 unique farm localities. The number of records per farm ranged from a minimum value of 1, to 290 records at the 25th percentile, to 385 records at the median number of weekly records, to 434 records per farm at the 75th percentile, and a max value of 566 unique active weekly records per farm. The average number of records per farm lied at 354.4 unique active farming weeks (0.1 percent of all records per farm). Information contained in the data outside previously mentioned attributes ranged from lice count attributes (female, mobile, stationary lice) to sea temperature, lat/lon coordinates of farm localities, to municipality name and identifier, to production area name and identifier, to information on species farmed, production types, farming capacity, occurrence of Pancreatic Disease or Infectious Salmon Anaemia (viral diseases in salmon), to commonly conducted treatment types of lice such as mechanical removal, substance treatments (in-feed or bath), and cleaner fish deployed.

The data was stored locally in csv format. API calls for some farm localities failed with unknown cause for the detailed farm data which contains information on stationary and mobile lice, on sea temperature, and on capacity. Data on thresholds was not obtainable through the API, and thus had to be downloaded from the BarentsWatch website manually. [10] There can be two different thresholds during the year, and depending on the location of a farm locality along the Norwegian coast, thresholds are lowered for different farms at different times. Information on lice (female, stationary, and mobile), on sea temperature, and on location (latitude, longitude) were obtained manually together with the threshold data to accommodate for missing values from the API calls for detailed weekly farm data.

3.3 Data Quality Assessment and Cleaning

For data quality assessment and data processing (Obj.1), Python was used in the PyCharm Community IDE, and GitHub Desktop was used for version control. Anaconda was used as the underlying interpreter. Data was first assessed for duplicates and missing values. If lice data and sea temperature data was complete from weekly and detailed data combined for a farm (data obtained through API calls), this data was used, otherwise the data from the threshold data download was used instead together with the features of latitude and longitude. The reasoning for this was to assure that reported sea temperatures and coordinates are aligned with lice counts from exactly those locations.

As a next step, farm cycles were identified by means of comparing fallow statuses of a given week to fallow statuses of week prior and week after. Because the data was recorded from the beginning of 2012 and several farms were already in the middle of a current cycle, only data was considered which lies after the point when the last farm who was in a cycle at the beginning of 2012 had completed their cycle. This was done, because week within cycle was aimed to be used as a predictor itself since previous research suggests that lice

[10] (BarentsWatch 2023*c*)`https://www.barentswatch.no/nedlasting/fishhealth/lice`

counts at the beginning of a cycle lie at 0, because fish who have just been placed in seawater have not been exposed to sea lice before that point. Farm cycle information was then utilized to interpolate missing values for fish lice and sea temperature in the manner that each cycle was interpolated separately linearly after setting lice counts at zero at the beginning of cycles if they were missing. Data on capacity, species, and other farm level information was interpolated with ffill and consequently applied bfill considering not only active weeks for a farm but also inactive weeks to assure that the last known, registered capacity figure for a farm could be obtained.

A finding was made that particularly at the end of farming cycles, lice counts were not as actively reported as at the beginning or in the middle of a cycle. Furthermore, there are unusually short cycles, sometimes of 1 week, for a small number of farms where a farm is fallow in the week prior and the week after but not fallow in the current week which is not explainable based on descriptions of salmon farming whereby fish can be stocked in the spring or autumn and remain in the locality until slaughtering after 1.5 years [11]. For consequent modelling purposes, data from those 1 week cycles were not included. There are also a small number of farms which have no fallow week in several years, which indicates that restocking of salmon smolts must have happened after slaughtering, but that fallow periods for those farms are either not required or that the rule to fallow is not abided by. While this occurrence doesn't seem to be in line with expected farming cycle length, it is found that due to the absence of fallow periods when farm localities can normally recover from high lice levels, no adjustment must be made in terms of splitting long cycles into several shorter cycles. Data on escapes could not be used because it is unclear how to interpret the data. Conflicting features make it impossible to know which numbers (features) to trust.

The overall number of records after pre-processing lied at 256325 with 1072 unique farm localities and an average number of 283.4 records (0.11 percent of all records) per farm. The minimum number of records per farm lied at 2, the 25th percentile in number of records per farm lied at 238 unique week records, which was followed by 307 median records per farm, 344 weekly records at the 75th percentile and a maximum number of records of 444 per farm. The time frame ranged from week 32 in 2014 to week 15 in 2023. The loss of active farming records after pre-processing lied at 70913, which is 21.7 percent of all active farming records collected from the API.

3.4 Feature Engineering

As per Obj. 2 of this study, features for the consequently fitted classification models were designed to best represent attributes assumed to impact female lice levels at individual farming localities for a future point in time:

1. Lice threshold for average female lice in current week, lice threshold in target week, and number of weeks till reduction of threshold within next 8 weeks at farm level.
2. Cycle week, capacity (tonnes of salmon), and derived number of fish at farm level [12].
3. Lice levels (average female, average mobile, average stationary) in previous week, and mean, max aggregates in previous month (including current week) at farm level.

[11] (Gettinby et al. 2011)

[12] A fish is assumed to weigh 3kgs on average with freshly stocked fish known to weigh between 100 and 1000 grams and grown fish to weigh 5-6kgs. (Dempsey et al. 2023)

4. Treatments (mechanical removal, cleaner fish, and feed/bath treatments of different substances both partial and complete locality) in previous week and aggregates of treatment categories applied in previous month (excluding current week's treatments) at farm level.

5. Total, mean, and max capacity for farms within 10, 20, and 50km radius.

6. Mean, max sea temperature, and mean, max lice levels (average female, average mobile, average stationary) for current week for farms within 10, 20, and 50km radius.

Lice levels and treatment at farm level for previous week/month were generated by means of utilizing an anchor date (current week's date: Tuesday of every week) and a relative time delta from that anchor date combined with a loc function to obtain aggregate figures. For lice figures in previous month, this week's lice counts were included to accommodate situations where there are less than 4 completed weeks in the current farming cycle for farm localities. For treatment history, the same procedure was not duplicated since the current week's treatment is meant to be left out and added as a treatment recommendation for farms to achieve the best possible outcome regarding avoiding exceeding the threshold in the target week. It is also assumed that treatments can occur and will thereby be recorded only later in a given farming cycle week; thus, current week's treatment data may not be available until the week has been completed and would not be able to be incorporated if this predictive system was used for real-time predictions.

To obtain aggregate lice and capacity figures for surrounding farms in a specified radius, haversine distance was applied using pairwise distance of longitude/latitude locations converted into radians. Once farms within a specific distance were obtained, the indexes for those farms were utilized to generate the aggregate figures across those neighbouring farms for individual farms. A Python package called 'searoute' was tested at first because the assumption was that sea route, in terms of shipping route, would give a better representation of farms within a given distance due to Norway's fjord system, which doesn't always allow for direct access of two points by means of direct seaway connection; since most farm localities however are in the open water, haversine distance proved more accurate when testing individual farm distances.

4 Design Specification

The design of the conducted study follows a three-layered approach as shown in Figure 2 below which integrates data extraction, pre-processing, feature/response variable generation, and data model preparation in the data layer, pre-modelling analysis (deriving knowledge on best point of warning from a classification tree), classification models predicting class of exceedance/non-exceedance of lice threshold for 3 points in the future, and a classification model integrating choice of treatment for current week in the domain knowledge and modelling layer, and an evaluation layer which compares knowledge gain of the conducted study against that of previously conducted research, which discusses possible future studies, and which foremost assesses usability of fitted models against needs of the beneficiaries, the Norwegian salmon farmers. All layers are interwoven and arranged in a circular architecture to accommodate for inclusion of knowledge gain for future steps of depicted research and for future iterations of research in the problem domain.

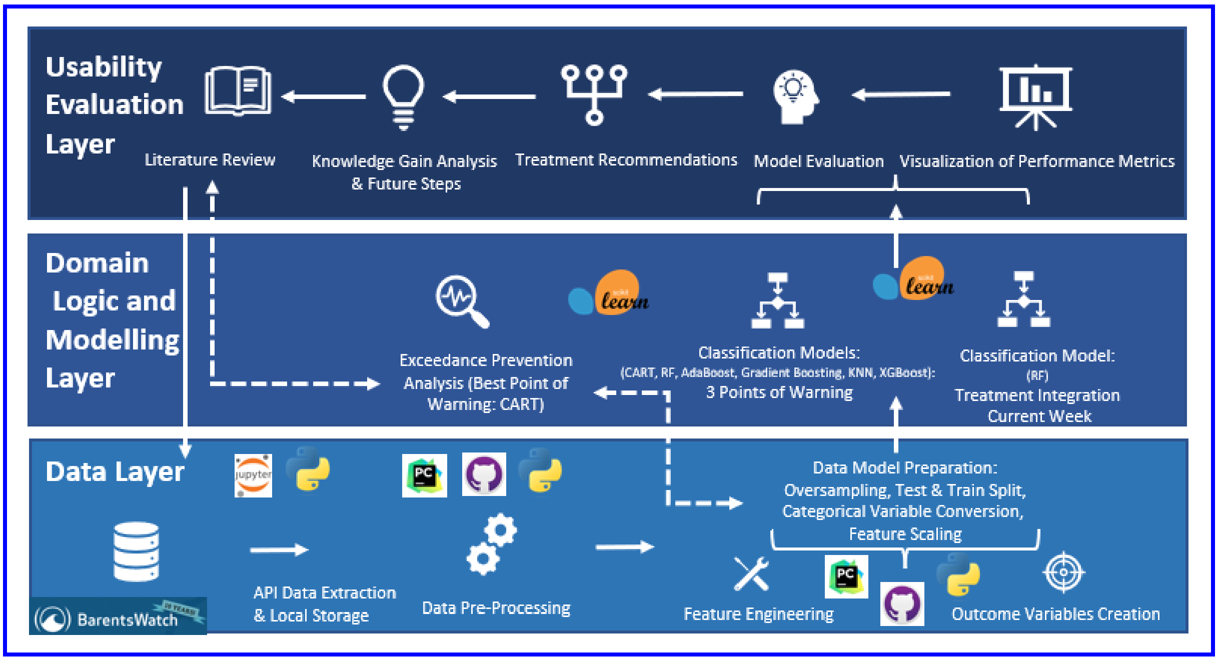

Figure 2: Design Specification

5 Implementation, Evaluation, and Results of Sea Lice Classification Models and Treatment Recommendations

The following chapter presents the implementation and evaluation of the individual algorithms trained for the purpose of the pre-modelling analysis, the classification of threshold exceedance, and for the classification of exceedance including treatment recommendations.

The chapter starts with a legitimization of models selected for aforementioned modelling purposes and explains why those modelling techniques are suitable for the data utilized for this research and continues with explaining pre-processing and evaluation techniques applied before it presents implementation, evaluation, and results for specific algorithms used.

For the classification of threshold exceedance, which represents the main part of this research, all models were trained for three different exceedance points in the future. Since 4 weeks before exceedance was found to be the most suitable point of warning for salmon farmers in Norway, and since performance indicators show very similar results for the classifiers of different warning points, model performance is only visualized for 4 weeks after current week while performance metrics for the other two warning points will be provided in a tabular overview for comparison purposes.

5.1 Supervised Classification Techniques

Due to the auto-correlated nature of the data this research utilizes, tree models were chosen for the classification problems in the research as they are best suited in dealing with auto-correlated data sets. Logistic regression as an alternative modelling technique, for example, relies on the assumption that data points are independent, and was thereby discarded as a possible classification technique, because data points in presented research represent lice and treatment data of farms of individual weeks while all the other active farming weeks of those farms are also contained in the data set and thereby one observation cannot be considered independent of other observations in the data. For the pre-modelling analysis, interpretability of the decision nodes of the model were crucial to obtain information on when farmers who can successfully prevent exceedance of lice thresholds for up to 8 weeks start treating proactively and what cues (e.g., in lice status, sea temperature etc) they recognize (consciously or subconsciously) as relevant to start treating. Therefore, a classical CART was utilized for that purpose while for the main part of the research, the focus was placed on the use of ensemble tree models such as Random Forest, AdaBoost, Gradient Boosting, and XGBoost, because they showed to outperform unpruned and pruned CART algorithms and additionally benefit from advantages such as prevention of over-fitting by means of bootstrapping in Random Forest, and sequential error reducing learning in AdaBoost, Gradient Boosting and XGBoost, with the latter outperforming traditional models such as SVM in speed and performance orientation.

KNN was tested as a model based on Euclidean distance to assess whether this modelling technique could prove beneficial for the research; it however showed less convincing results compared to the ensemble tree models applied. SVM was attempted as an additional alternative method, however the dataset proved too large for the algorithm to be trained in a timely manner. Therefore, this algorithm was discarded.

5.2 Modelling Data Preparation

5.2.1 Handling Categorical Variables

Due to controversy in the data science discourse over whether tree models require categorical variables to be changed to dummy variables, especially when using sci-kit learn, which was used for conducted research, it was found the preferred solution to do so.

5.2.2 Scaling of Numerical Variables

KNN and XGBoost require numerical variables to be scaled to overcome differences in scale in the numerical features utilized in the classification models and to not give features of large scale higher importance than those of small scale. Standard Scaler was applied for that purpose.

5.2.3 Handling Class Imbalance

SMOTE (Synthetic Minority Oversampling Technique) was applied for all three classification purposes. In the pre-modelling analysis, the minority class represented the farm localities that did exceed the threshold while treating proactively in up to 8 weeks after the current week while for the classification of threshold exceedance excluding and including the current week's treatment variable it represented the target class of farms that exceed the given lice threshold in the target week. The method was applied to assure that both classes would be equally represented in the data before model fitting.

5.3 Performance Metrics and Model Evaluation Techniques

Performance metrics used to assess the predictive power of individual classification models were accuracy, precision, recall, f1-score, and ROC-AUC score. While for the simple CARTs 10 k-fold cross-validation was applied in combination with grid search to obtain most robust model results and prevent the models from over-fitting, for the more complex ensemble models, it was decided to apply only hold-out cross-validation due to long processing times for fitting the models. For that purpose, the data was split into training and test set with an 80/20 ratio and stratified splitting which assures that the same percentage of both classes is contained in both training and test set.

5.4 Pre-Modelling Analysis. Implementation, Evaluation, and Results of CART algorithm

5.4.1 Pre-Modelling Analysis

As described above, the pre-modelling analysis was conducted for the purpose of identifying warning signs for farmers (consciously or subconsciously recognized) which makes them treat successfully in a proactive manner so that they don't exceed set government thresholds for up to 8 weeks after the current week. This was achieved by means of following the decision nodes of a fitted decision tree classifying whether a farm that is treating preventative is able to successfully prevent lice exceedance for up to 8 weeks and subsequently by deriving summary statistics of farms who do not treat following these cues. This step in the research is important to assure that the model is fitted for the right point of warning that allows farmers to successfully prevent the exceedance of a future week lice threshold. It assures the model is useful to the beneficiaries and can help them comply with government regulations.

5.4.2 Implementation, Evaluation, and Results

Implementation:

For the pre-modelling analysis, only those observations were considered where farms apply a treatment despite the fact that they are not currently exceeding the set thresholds. Variables included were treatment of previous weeks, of current week, lice counts in current week and in previous weeks, as well as sea temperature in current week. Once the categorical variables were turned into dummy variables, data was then balanced in regards to the number of observations in the target class (successful preventers of exceedance for up to 8 weeks) and the non-target class (unsuccessful preventative treaters). Subsequently, the data was split into 80 percent training observations and 20 percent test observations and a CART was fitted with grid search using ten-fold cross-validation. The model output was then examined for decision nodes of successful preventers.

Evaluation and Results:

As shown in Figure 3, farms that treat successfully treat at a max female lice aggregate of larger than 0.16 in the previous month (including this current week's female lice count), at a minimum stationary lice count aggregate for the previous month of larger than 0 and at an average sea temperature larger than 7.521 as per farms reported in a 50 km radius.

Those obtained cues were then applied as filtering criteria to farms that do not treat when criteria apply to them and that do exceed set lice threshold within the upcoming

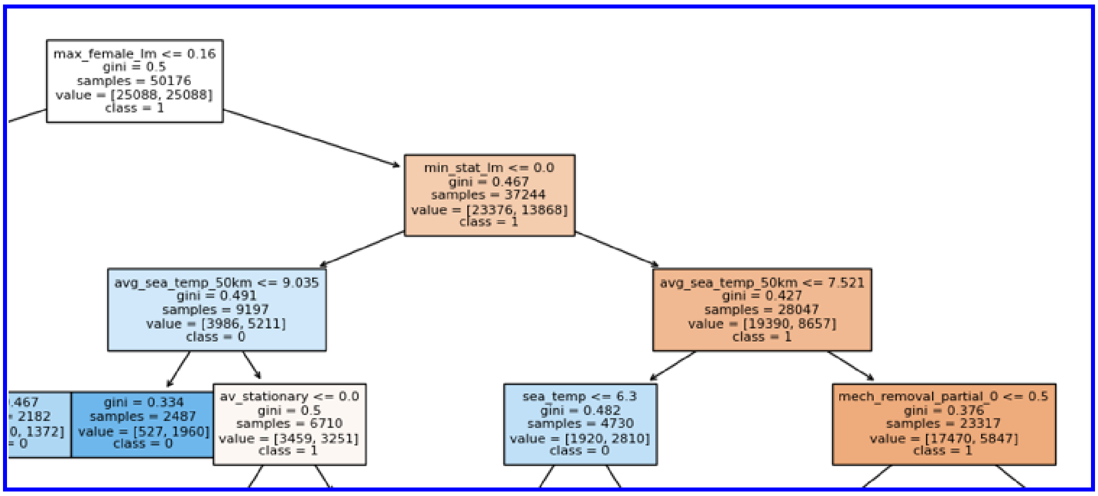

Figure 3: Pre-Modelling CART: Successful Prevention of Exceedance

8 weeks. The mean value of number of weeks until farms exceed was identified to be 3.6 and derived as a summary statistics from the farms currently under threshold missing to act on those cues as displayed in above Figure 4. This step was taken to legitimize the best point of warning for salmon farmers to take preventative measures as subsequently utilized in designing the set-up of the exceedance classification presented in 5.5.

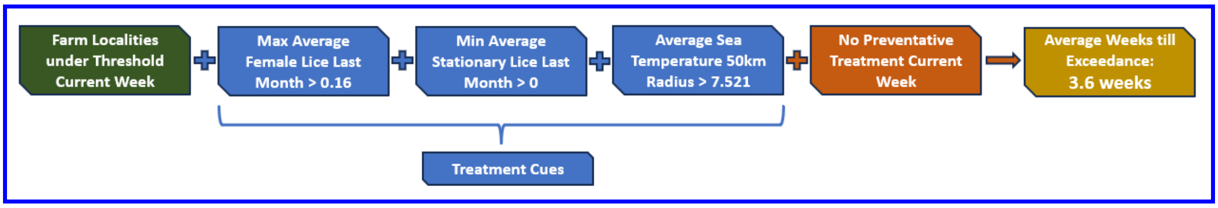

Figure 4: Missed Treatment Cues and Average Weeks Till Exceedance

The model predicted with an accuracy of 0.75, precision of 0.75, recall of 0.74, and and ROC-AUC score of 0.85. While these results could be improved by means of increasing the tree depth or by means of applying ensemble methods as applied in the threshold exceedance classifications, the accuracy was accepted in favour of easy explainability of the model to obtain treatment cues.

5.5 Implementation, Evaluation, and Results of Sea Lice Exceedance Classification Models

5.5.1 Exceedance Classifier for 4,6, and 8 weeks after current week

Based on the results from the pre-modelling analysis, response variables were created to determine exceedance of set lice thresholds for three different future points in time.

For the subsequently trained classification models as per Obj.4 of this research, 4 weeks after the current week was found to be the most beneficial point of warning for salmon farmers in Norway to react to the risk of exceeding lice thresholds based on the results of the pre-modelling analysis presented in 5.4.

However, as an additional part of Obj.4 lies in answering SUB-RQ 3 whether model performance decreases for earlier points of warning, models were also created for exceedance 6 and 8 weeks after the current week. Earlier points of warning than 8 weeks before exceedance were not considered, to not encourage over-treatment (when treatment is not required) and to keep model results tangible for the farmers.

For this part of the research, only treatments of previous weeks were selected and not this current week's applied treatment for the purpose of generating model results that represent commonly found outcomes of sea lice levels irrespective of what type of treatment is applied in current week.

As outlined in 5.1, with the exception of KNN, tree models, in particular ensemble learners, were chosen for the exceedance classification models presented due to the autocorrelated nature of the data which violates the assumption that observations in the data are independent, which needs to be met in order to utilize models such as logistic regression. Also, tree classifiers are suited for the research purpose at hand because they can handle a combination of both numerical and categorical variables, as in the data present in this study.

5.5.2 Implementation, Evaluation, and Results of CART Model: Unpruned and Pruned

Implementation:

As for all trained algorithms trained as part of this research, categorical variables were turned into dummy variables, and oversampling was applied before model fitting. Consequently, data was split into 80 percent training observations and 20 percent test observations. The model was then trained on the training set and model performance indicators were obtained by comparing model predictions (classifications) to real outcomes of the training observations. At first, an unpruned CART was trained. Since unpruned CARTs can have the tendency to overfit the training data and therefore lack in performance on unseen observations, the same approach as presented in 5.4 was followed in consequence by means of applying grid search over commonly used model parameters in combination with ten-fold cross-validation while training the model. Both CARTs were trained for the outcome of exceedance in 4, 6, and 8 weeks, and performance parameters for the models were obtained.

Table 2: CART Performance

Classifier	**Week**	**Accuracy**	**Precision**	**Recall**	**f1-score**	**ROC-AUC**
CART Unpruned	**4**	0.945	0.94	0.95	0.95	0.94511
CART Unpruned	6	0.947	0.94	0.95	0.95	0.94669
CART Unpruned	8	0.946	0.94	0.95	0.95	0.94568
CART Grid Search	**4**	0.789	0.76	0.84	0.8	0.87359
CART Grid Search	6	0.776	0.75	0.84	0.79	0.85239
CART Grid Search	8	0.766	0.8	0.71	0.76	0.85154

Evaluation and Results:

As shown in Table 3 above, the unpruned CART performed with an accuracy of 94-95 percent on the test set with the best accuracy and ROC-AUC score achieved for classifying exceedance for 6 weeks after current week. Precision was consistently at 94 percent while recall was at 95 percent. This means that the unpruned classifier performs better in regards to identifying true exceeders than it is at falsely classifying non-exceeders as exceeders; this is not ideal since it may encourage 6 percent of farmers to take preventative actions when in reality no action is required to be taken.

The pruned CART performs lower overall with an accuracy of 79-77 percent and an ROC-AUC score of 87-85 percent (with 1 percent accuracy level drop with every two weeks further into the future). The same pattern in precision and recall can be observed as for the unpruned CART except for classifying exceedance for 8 weeks ahead, where slightly higher precision than recall figures could be detected. Overall however, the pruned CART's performance is too weak to be considered a strong performing model.

Since the unpruned CART does not stop fitting before the split at the nodes into classes is 100 percent completed, the prediction probabilities are always either 0 or 1. This is not very helpful in regards to understanding the risk reduction which can be achieved after applying specific treatments. The CART therefore does not prove an efficient method for classification of lice threshold exceedance.

5.5.3 Implementation, Evaluation, and Results of Random Forest Model

Implementation: The Random Forest was implemented in the same manner as the CART in regards to preparation of categorical variables, oversampling and train and test split. The number of estimators was set to 300 trees.

The method of Random Forest benefits from the advantage of training the individual estimators on different subsets of the data, so called bootstrap samples. This means that this model type tends to perform well on unseen observations. As an ensemble method, it furthermore benefits from majority voting across all trained classifiers and thus provides robust classification results.

Table 3: Random Forest Performance

Classifier	**Week**	**Accuracy**	**Precision**	**Recall**	**f1-score**	**ROC-AUC**
Random Forest	**4**	0.98	0.98	0.98	0.98	0.99786
Random Forest	6	0.98	0.98	0.98	0.98	0.99765
Random Forest	8	0.98	0.99	0.97	0.98	0.99742

Evaluation and Results:

Random Forest shows the strongest performance across all trained classifiers with an accuracy of 98 percent for all three predicted exceedance points, with 98 of precision and recall for exceedance classifications 4 weeks and 6 weeks ahead and 99 precision and 97 recall for 8 weeks ahead. The ROC-AUC consistely lies between 0.997 and 0.998 with slightly decreasing results with every two weeks later in point of prediction. That means that the model only has about 2 percent of false positive predictions and 2 percent false negative predictions. Most exceeding farms are detected with only few misclassifications of non-exceeders as exceeders.

5.5.4 Implementation, Evaluation, and Results of AdaBoost and Gradient Boosting Model

Implementations: Adaboost and Gradient Boosting are combined in this technical report because both boosting methods were applied without scaling as opposed to XGBoost which requires the numerical variables to be scaled.

After creation of dummy variables, oversampling was applied the data and a train and test split of 0.8 vs 0.2 was applied. To understand which of the two boosting methods

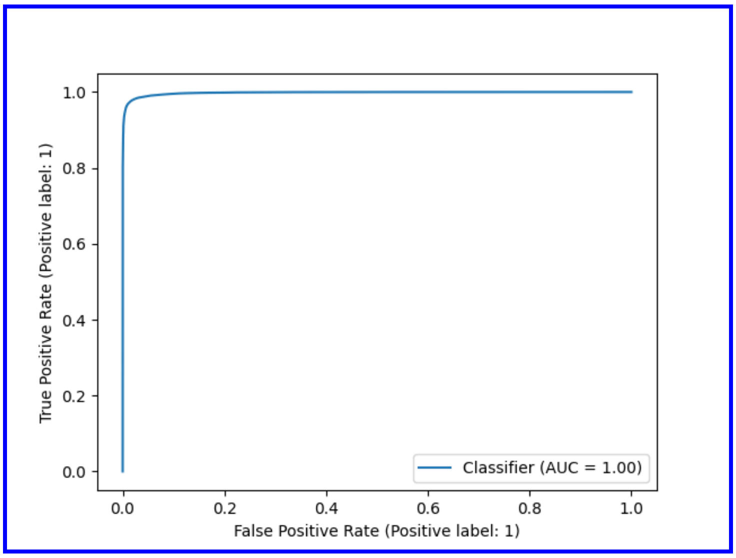

Figure 5: ROC-AUC Random Forest: 4 Weeks Out

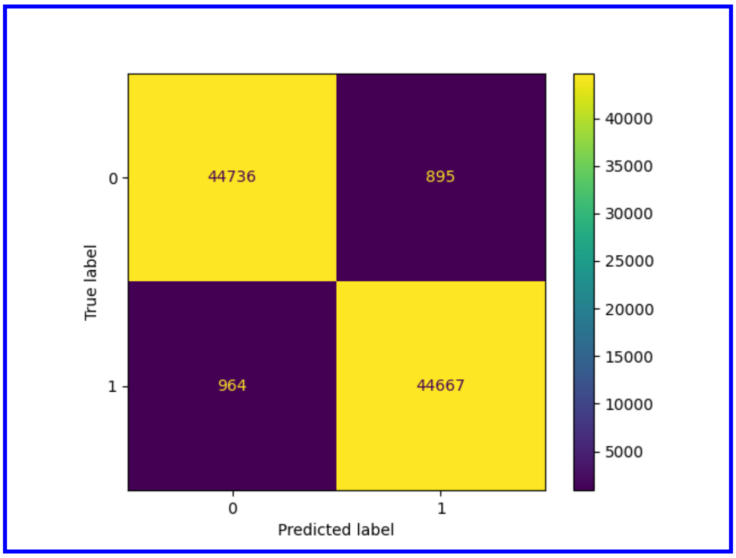

Figure 6: Confusion Matrix Random Forest: 4 Weeks Out

performs more strongly than the other, 300 was set as estimators for both modeling methods. The same concept was applied for setting the max depth of the weak learning as the base learner of the algorithm to 1 for both algorithms. AdaBoost, also called Adaptive Boosting, benefits from sequential correction of prediction errors by means of placing more weight on wrongly conducted predictions, while Gradient Boosting doesn't adjust weights but instead utilizes the residual errors of predecessor's estimators as labels for the prediction of the next estimator fitted.

Table 4: AdaBoost and Gradient Boost Performance

Classifier	**Week**	**Accuracy**	**Precision**	**Recall**	**f1-score**	**ROC-AUC**
AdaBoost	**4**	0.948	0.96	0.94	0.95	0.98771
AdaBoost	6	0.946	0.96	0.93	0.94	0.98515
AdaBoost	8	0.945	0.96	0.93	0.94	0.98501
Gradient Boosting	**4**	0.891	0.88	0.91	0.89	0.96286
Gradient Boosting	6	0.887	0.88	0.89	0.89	0.95921
Gradient Boosting	8	0.883	0.88	0.89	0.88	0.95557

Evaluation and Results: While both classifiers show strong performance across assessed performance indicators, AdaBoost is the preferred method since it outperforms GradientBoosting in regards to overall accuracy, and ROC-AUC score; but most importantly because the method shows slightly higher precision than recall which reduces the risk of false classification of exceeders while Gradient Boosting performs better at detecting exceeders over its ability to correctly classify only true exceeders as exceeders.

For both methods, recall decreases slightly for later points of predictions. This means that both models over time are missing to detect true exceeders and inaccurately classify them as non-exceeders. This however is only a minimal increase; the models still perform strongly over time.

5.5.5 Implementation, Evaluation, and Results of KNN Model

Implementation: The two remaning models, KNN and XGBoost were trained on scaled numerical variables in addition to pre-processing steps including dummy variable creation and oversampling. For KNN, this is required to assure that all features are given the same importance in the determination of Euclidean distance which is the basic concept of the algorithm. Number of neighbours was set to the square root of training observations.

Table 5: KNN Performance

Classifier	**Week**	**Accuracy**	**Precision**	**Recall**	**f1-score**	**ROC-AUC**
KNN	**4**	0.788	0.74	0.89	0.81	0.87937
KNN	6	0.774	0.72	0.89	0.8	0.86892
KNN	8	0.758	0.71	0.89	0.79	0.85195

Evaluation and Results: KNN performed similarly to the pruned CART in regards to performance metrics observed. Just as the pruned CART, that KNN shows consistently higher recall scores than it does precision scores. That means that is better able to detect true exceeders than it is to avoid misclassifications of non-exceeders as exceeders. Therefore, this method is also not a preferred method.

5.5.6 Implementation, Evaluation, and Results of XGBoost Model

Implementation: XGBoost, the same way as KNN, requires scaling of the numeric variables. Therefore pre-processing of data was conducted in the same manner as it was for the fitting of the KNN algorithm. XGBoost, which stands for Extreme Gradient Boosting is an advanced version of Gradient Boosting with improved generalization. It is known to outperform Gradient Boosting in training times and therefore is very suited for large datasets as the one utilized for the study presented here.

Table 6: XGBoost Performance

Classifier	Week	Accuracy	Precision	Recall	f1-score	ROC-AUC
XGBoost	**4**	0.971	0.98	0.96	0.97	0.99429
XGBoost	6	0.973	0.98	0.96	0.97	0.99376
XGBoost	8	0.973	0.99	0.96	0.97	0.99331

Evaluation and Results: XGBoost performed the second strongest after the Random Forest algorithm with consistent accuracy of 97 percent and ROC-AUC scores of 0.994 for all three predicted points in time. Recall consistently stayed at 0.96 for all three predicted time frames while precision is at 0.98 for 4, and 6 weeks out and improves to 0.99 for 8 weeks out. This means that just as with Random Forest and AdaBoost (the other two out of the three strongest performers), precision is slightly higher than recall which is to be preferred to rule out risk of mis-classification and risk of excessive unnecessary treatments.

5.6 Exceedance Classifier Including Treatment Recommendations

5.6.1 Implementation, Evaluation, and Results of Random Forest Classifier including Current Week's Treatment

Implementations: As per Obj. 6, an additional classification model (Random Forest as best performing model from 5.5) was then trained, this time including the current week's treatments applied, and farm localities predicted to exceed in 4 weeks as per the strongest performing classifier, Random Forest presented in 5.5.3, were then fed into the model to obtain outcomes for application of all possible (all observed) treatment options; including 'no treatment applied'. The model was trained in the same manner as the RF model from 5.5.3 except for the fact that this current week's treatment, which can be an individual treatment or a combination of treatments, was included as a predictor. This was done to be able to predict the same observation with all possible treatment options including the baseline treatment of no treatment applied.

Evaluation and Results: The model predicted with 0.979 accuracy, 0.98 precision and recall and with an ROC-AUC score of 0.9978. It is thereby able to provide strong predictions while neither misclassifying too many mo-exceeders as exceeders nor failing to classify true exceeders as exceeders. The probabilties to exceed for each observations for every possible treatment type or combination of treatments were then compared against the baseline proability to exceed when no treatment is applied to give farmers who are at risk of exceeding to decrease the risk of exceeding. Therefore only those recommendations are included that can change the classification result from exceed to not exceed and it orders the possible treatment options in order of biggest risk reduction from baseline risk of no treatment applied. For the shown treatment options, the cost of one treatment over the other is not integrated, because such data was not available to the researcher. Neither are specific enough constraints as to not being able to apply individual treatments based on weight of fish (Gettinby et al. 2011, p.168) or other constraining factors.

Figure 7 below shows an example of such a recommendation with the current risk of the general model from 5.5.3 to exceed lice thresholds in 4 weeks and a display of the predicted risk reduction to be achieved for the three most efficient treatment options applied compared to the treatment option of 'no treatment' applied. The firstly displayed risk can be interpreted as a general risk observed in farms with similar attributes not taking current week's applied treatments into account.

```
Predicted risk for lice threshold exceedance in 4 weeks is at 0.91.
 See predicted risk reduction over baseline risk of no treatment by applying following proposed treatments:

                                   proposed_treatment  baseline_risk  prob_to_exceed_with_treatment  risk_reduction
                    ['Diflubenzuron-4-InFeed-True']       0.926667                       0.460000        0.466667
                   ['Hydrogenperoksid-6-Bath-True']       0.926667                       0.463333        0.463333
['cleaner_fish', 'Emamectin benzoat-5-InFeed-True']       0.926667                       0.470000        0.456667
```

Figure 7: Recommendation Example 1

The baseline risk outlined refers to the risk predicted from the model that includes current week's treatments applied and that refers to the risk of not applying any treatment in current week. The probability to exceed with proposed treatment/treatment combination lists the three options that show the lowest probability to exceed for the predicted point in the future which is 4 weeks after the current week.

```
Predicted risk for lice threshold exceedance in 4 weeks is at 0.93.
 See predicted risk reduction over baseline risk of no treatment by applying following proposed treatments:

                  proposed_treatment  baseline_risk  prob_to_exceed_with_treatment  risk_reduction
             ['mech_removal_entire']       0.943333                       0.496667        0.446667
['Emamectin benzoat-5-InFeed-False']       0.943333                       0.496667        0.446667
```

Figure 8: Recommendation Example 2

The risk reduction demonstrates by how much applications of these treatments/treatment combinations can reduce the risk to exceed lice levels for the predicted point in the future. 'True' refers to treatment applied in entire locality whereas 'False' refers to treatment applied in partial locality. These identifiers are used for substance-based Bath or Feed treatments. The results are only generated for those treatments that can change the predicted outcome from 'exceed' to 'not exceed' by means of application of efficient treatments.

6 Summary of Study Results and Evaluation against Set Research Objectives

The following chapter summarizes the results and evaluates them in relation to achievement of set research objectives as well as in relation to knowledge or value gain in comparison to previously conducted research in the problem domain.

6.1 Pre-Modelling Analysis

The pre-modelling analysis was conducted to legitimize the best point of warning for the subsequently trained classifier of lice threshold exceedance. The chosen method of an individual CART whose model parameters were obtained through grid search of best accuracy validated with tenfold cross-validation to assure the model was not overfitting the training data was able to visualize successful prevention of exceedance for up to 8 weeks of farms treating proactively. Cues could be obtained which could then serve as filtering criteria for farms who miss out on treating when treatment has proven necessary, and aggregate statistics could be obtained that legitimize which point of warning may prove as beneficial for most farmers. SUB-RQ 1 could therefore be answered with 4 weeks prior to exceeding and Obj. 3 serving the purpose of answering the question was therefore successfully achieved.

6.2 Classifier Performance

As per model results shown in chapter 5, the strongest classifiers to predict exceedance of lice thresholds in 4, 6, and 8 weeks are Random Forest, XGBoost and AdaBoost with ROCs between 0.997 and 0.985. Only minimal performance decrease can be observed from models predicting lice exceedance in 4 weeks versus lice exceedance 6 weeks or 8 weeks later.

Precision scores are consistently slightly higher than recall scores which indicates that the model misclassifies less farm localities as exceeders when they in fact won't exceed (false positives) than wrongly predicting actual exceeders as non-exceeders (false negatives).

The divides between precision and recall become slightly bigger with growing time frames till exceedance. While the models miss out on a few more farms they should detect as exceeders, they misclassifies only few localities as exceeders when they in fact won't exceed. This can be in favour of the designed warning system that encourages farmers to apply costly treatments if they are predicted to exceed thresholds for a future point in time. If there are few false positives, the risk of financial damage for farmers applying treatment when no treatments are necessary can be kept low. Random Forest is the strongest performer with 98 percent Precision and Recall for exceedance both 4 weeks and 6 weeks after current week. Only for exceedance 8 weeks after the current week, the classifier shows a divide of 2 percent however increases precision to 99 percent while recall is at 97 percent. Lowest performers are CART with grid search and KNN with an ROC of only 0.85 to 0.87. SUB-RQ2 and SUB-RQ3 could thereby be answered and supporting objectives 4 and 5 were thus successfully achieved.

6.3 Treatment Recommendations

As per Obj. 6, treatment recommendations were obtained by means of sorting probabilities to exceed in ascending order for different possible treatments/treatment combinations. While the model results obtained through the classifiers trained on features excluding current week's treatments show the general probability to exceed, the model results from the model including the current week's treatment can compare the probability to exceed without any treatment applied with that of those treatments/ treatment recommendations that show the largest risk reduction from the baseline of no treatment applied.

Since information on treatment cost or treatment constraints was not included in this research, the recommended treatments do not take those factors into consideration. The results can be improved by future iterations of the study integrating constraining factors such as fish weight (some treatments prove too costly for large fish) and by aiming at minimizing cost, especially if two treatment options that differ in price show the same risk reduction of exceedance. Furthermore, it should be evaluated if treatment recommendations enforce treatment patterns that harm the environment, that potentially encourage excessive treatment or that harm the balance of coordinated treatments in certain coastal areas (if those coordinated treatment efforts exist or should be encouraged).

Therefore, SUB-RQ 4 can be answered in the way that the model is accurate in predicting risks of exceedance including possible treatments for current week, but that further optimization must be conducted to assure treatment recommendations are cost-efficient and in line with coordinated treatment efforts in the farming system of Norway. Objective 6 was successfully accomplished to answer this SUB-RQ.

6.4 Discussion

The conducted study was able to demonstrate how supervised machine-learning, especially ensemble tree classifiers, can accurately predict exceedance of government-set thresholds for different points in the future. It was furthermore able to provide legitimization for proposed points of warning to assure that the warning system proves beneficial for the beneficiaries of Norway's salmon farmers. Lastly, the research provided a framework to integrate data-driven recommendations on efficient treatment strategies for farms with different treatment backgrounds and lice figures in changing conditions of the environment including environmental factors such as sea temperature and lice pressure originating from surrounding farms. While the research cannot simulate movement of lice particles through current drift or to explain physical phenomena of lice transmission under specific environment conditions, as in previously conducted research in the domain, it can provide actionable and easily understandable insights to farmers to help them comply with government regulations. With an accuracy of 98 percent, it can predict whether a farm locality will exceed government-set lice levels for up to 2 months in the future. It can especially assist farmers who run at risk of exceeding to be equipped with recommendations on most-efficient treatment options while assuring the only a small percentage of farmers will be recommended to treat when no treatment is necessary, which minimizes the risk of unnecessary spending of the farmers while detecting most farm localities that will indeed exceed.

The biggest advantage of a warning-system based on data-driven results is that no assumptions must be integrated on how the natural phenomenon of lice abundance comes to be at individual farm localities. Previously conducted research in the domain has focused on this aspect and has served the purpose of providing valuable insights on lice transmission between different farm localities which has been used in this research by means of including found knowledge into designed features utilized by the algorithm. Furthermore, the study has shown that the prediction can be obtained by solely using publicly available data, without a need to integrate further farm specific information not available to the public. Since this data is already publicly available, the demonstrated solution can be turned into a live tool for farmers which could potentially be free of charge to them. Data quality and missing data has proven to be time-consuming to improve and address, and therefore automation would require further attention if this solution were to be commercialized.

7 Conclusion and Future Work

The conducted study achieved to demonstrate a working solution to support farmers in taking preventative actions in their fight against salmon lice. It has proven that publicly available can be used to achieve a 98 percent accuracy in predicting whether a farm locality will exceed the government set lice threshold for 3 different points in the future. Tree classifiers, in particular ensemble learners with boosting have proven most effective in predicting exceedance of lice thresholds based on historical data on treatments, lice figures, lice counts of current week at farm level and in surrounding farms. Sea temperature proved sufficient in integrating environmental features into the prediction.

While treatment recommendations are derived from a 98 percent accurate Random Forest algorithm, they would benefit from optimization including constraints that affect costs or synchronized treatment patterns of farmers. This element could thereby be improved in future iterations of the study. Furthermore, additional automation in data quality control could prove beneficial especially when commercializing the proposed solution.

Overall, the study was aimed at assisting farmers to take corrective, pro-active measures to counter-act growing lice levels, especially amidst the country's efforts to fiftyfold their salmon farming capacity. This goal was achieved. The fact that publicly available data was used for the study could furthermore aid the beneficiary in benefiting from a tool that might be available to be used free of charge to them with low maintenance. Model-drift should be observed and a change in the landscape of treatment methods.

8 Acknowledgement

I would like to thank Dr Catherine Mulwa for her expertise and guidance while conducting my research and for the support of my family, partner, friends, work colleagues, manager, and company in the course of pursuing my dream of creating a career in data analytics. I would like to thank NCI and Springboard for providing me with the opportunity to gain invaluable industry-acknowledged skills, confidence in independent research, and for allowing me to make great friends in the course of this journey. I am extremely thankful for that.

References

Aaen, S. M., Helgesen, K. O., Bakke, M. J., Kaur, K. & Horsberg, T. E. (2015), 'Drug resistance in sea lice: a threat to salmonid aquaculture', *Trends in Parasitology* **31**(2), 72–81.

Abolofia, J., Asche, F. & Wilen, J. (2017), 'The cost of lice: Quantifying the impacts of parasitic sea lice on farmed salmon', *32* (3), 329–349.

Aldrin, M., Huseby, R., Stien, A., Grøntvedt, R., Viljugrein, H. & Jansen, P. (2017), 'A stage-structured bayesian hierarchical model for salmon lice populations at individual salmon farms – estimated from multiple farm data sets', *Ecological Modelling* **359**, 333–348.

Aldrin, M., Jansen, P. & Stryhn, H. (2019), 'A partly stage-structured model for the abundance of salmon lice in salmonid farms', *Epidemics* **26**, 9–22.

Aldrin, M., Storvik, B., Kristoffersen, A. B. & Jansen, P. (2013), 'Space-time modelling of the spread of salmon lice between and within norwegian marine salmon farms', *PLoS One* **8**(5).

BarentsWatch (2023*a*), `https://www.barentswatch.no/bwapi/openapi/index.html?urls.primaryName=Fishhealth%20API`. Accessed: 2023-04-07.

BarentsWatch (2023*b*), `https://www.barentswatch.no/en/about/`. Accessed: 2023-04-15.

BarentsWatch (2023*c*), `https://www.barentswatch.no/nedlasting/fishhealth/lice`. Accessed: 2023-04-15.

Blanco Gonzalez, E. & de Boer, F. (2017), 'The development of the norwegian wrasse fishery and the use of wrasses as cleaner fish in the salmon aquaculture industry', *Fisheries Science* **83**(5), 661 – 670.

Bøhn, T., Nilsen, R., Gjelland, K. , Biuw, M., Sandvik, A. D., Primicerio, R., Karlsen, & Serra-Llinares, R. M. (2022), 'Salmon louse infestation levels on sea trout can be predicted from a hydrodynamic lice dispersal model', *Journal of Applied Ecology* **59**(3), 704 – 714.

Dean, K. R., Aldrin, M., Qviller, L., Helgesen, K. O., Jansen, P. A. & Bang Jensen, B. (2021), 'Simulated effects of increasing salmonid production on sea lice populations in norway', *Epidemics* **37**, 1–8.

Dempsey, D. P., Reid, G. K., Lewis-McCrea, L., Balch, T., Cusack, R., Dumas, A. & Rensel, J. (2023), 'Estimating stocking weights for atlantic salmon to grow to market size at novel aquaculture sites with extreme temperatures', *Frontiers in Marine Science* **10**.

Elghafghuf, A., Vanderstichel, R., Hammell, L. & Stryhn, H. (2020), 'Estimating sea lice infestation pressure on salmon farms: Comparing different methods using multivariate state-space models', *Epidemics* **31**.

Elghafghuf, A., Vanderstichel, R., Hammell, L. & Stryhn, H. (2021), 'State-space modeling for inter-site spread of sea lice with short-term population predictions', *Ecological Modelling* **452**.

Gettinby, G., Robbins, C., Lees, F., Heuch, P. A., Finstad, B., Malkenes, R. & Revie, C. W. (2011), 'Use of a mathematical model to describe the epidemiology of lepeophtheirus salmonis on farmed atlantic salmon salmo salar in the hardangerfjord, norway', *Aquaculture* **320**(3-4), 164 – 170.

Harrington, P. D., Cantrell, D. L., Foreman, M. G. G., Guo, M. & Lewis, M. A. (2023), 'Timing and probability of arrival for sea lice dispersing between salmon farms', *Royal Society Open Science* **10**(2).

Jeong, J., Stormoen, M., McEwan, G. F., Thakur, K. K. & Revie, C. W. (2021), 'Salmon lice should be managed before they attach to salmon: Exploring epidemiological factors affecting lepeophtheirus salmonis abundance on salmon farms', *Aquaculture* **541**.

Jevne, L. & Reitan, K. (2019), 'How are the salmon lice (lepeophtheirus salmonis krøyer, 1837) in atlantic salmon farming affected by different control efforts: A case study of an intensive production area with coordinated production cycles and changing delousing practices in 2013–2018', *Journal of Fish Diseases* **42**(11), 1573–1586.

Johansen, L.-H., Jensen, I., Mikkelsen, H., Bjørn, P.-A., Jansen, P. & Bergh, O. (2011), 'Disease interaction and pathogens exchange between wild and farmed fish populations with special reference to norway', *Aquaculture* **315**(3-4), 167 – 186.

Kragesteen, T. J., Simonsen, K., Visser, A. W. & Andersen, K. H. (2021), 'Estimation of external infection pressure and salmon-louse population growth rate in faroese salmon farms', *Aquaculture Environment Interactions* **13**, 21 – 32.

Kristoffersen, A. B., Jimenez, D., Viljugrein, H., Grøntvedt, R., Stien, A. & Jansen, P. A. (2014), 'Large scale modelling of salmon lice (lepeophtheirus salmonis) infection pressure based on lice monitoring data from norwegian salmonid farms', *Epidemics* **9**, 31–39.

Kristoffersen, A. B., Qviller, L., Helgesen, K. O., Vollset, K. W., Viljugrein, H. & Jansen, P. A. (2018), 'Quantitative risk assessment of salmon louse-induced mortality of seaward-migrating post-smolt atlantic salmon', *Epidemics* **23**, 19–33.

Liu, Y. & Bjelland, H. v. (2014), 'Estimating costs of sea lice control strategy in norway', *Preventive Veterinary Medicine* **117**(3-4), 469 – 477.

Misund, A. (2019), 'From a natural occurring parasitic organism to a management object: Historical perceptions and discourses related to salmon lice in norway', *Marine Policy* **99**, 400–406.

Murray, A. G. & Salama, N. K. (2016), 'A simple model of the role of area management in the control of sea lice', *Ecological Modelling* **337**, 39 – 47.

Myksvoll, M. S., Sandvik, A. D., Albretsen, J., Asplin, L., Johnsen, I. A., Karlsen, , Kristensen, N. M., Melsom, A., Skardhamar, J. & Ådlandsvik, B. (2018), 'Evaluation of a national operational salmon lice monitoring system—from physics to fish', *PLoS ONE* **13**(7).

Nayan, A.-A., Saha, J., Mozumder, A. N., Mahmud, K. R., Al Azad, A. K. & Kibria, M. G. (2021), 'A machine learning approach for early detection of fish diseases by analyzing water quality', *Trends in Sciences* **18**(21).

Peacock, S. J., Krkošek, M., Bateman, A. W. & Lewis, M. A. (2020), 'Estimation of spatiotemporal transmission dynamics and analysis of management scenarios for sea lice of farmed and wild salmon', *Canadian Journal of Fisheries and Aquatic Sciences* **77**(1), 55 – 68.

Revie, C., Robbins, C., Gettinby, G., Kelly, L. & Treasurer, J. (2005), 'A mathematical model of the growth of sea lice, lepeophtheirus salmonis, populations on farmed atlantic salmon, salmo salar l., in scotland and its use in the assessment of treatment strategies.', *Journal of fish diseases* **28**(10), 603 – 613.

Rittenhouse, M. A., Revie, C. W. & Hurford, A. (2016), 'A model for sea lice (lepeophtheirus salmonis) dynamics in a seasonally changing environment', *Epidemics* **16**, 8–16.

Salama, N., Collins, C., Fraser, J., Dunn, J., Pert, C., Murray, A. & Rabe, B. (2013), 'Development and assessment of a biophysical dispersal model for sea lice', *Journal of Fish Diseases* **36**(3), 323 – 337.

Sandvik, A., Bjørn, P., Ådlandsvik, B., Asplin, L., Skardhamar, J., Johnsen, I., Myksvoll, M. & Skogen, M. (2016), 'Toward a model-based prediction system for salmon lice infestation pressure', *Aquaculture Environment Interactions* **8**, 527 – 542.

Sandvik, A. D., Johnsen, I. A., Myksvoll, M. S., Sævik, P. N. & Skogen, M. D. (2020), 'Prediction of the salmon lice infestation pressure in a norwegian fjord', *ICES Journal of Marine Science* **77**(2), 746 – 756.

St-Hilaire, S., Patanasatienkul, T., Yu, J., Kristoffersen, A., Stryhn, H., Revie, C., Ibarra, R., Tello, A. & McEwan, G. (2018), 'Tool for predicting caligus rogercresseyi abundance on salt water salmon farms in chile', *Preventive Veterinary Medicine* **158**, 122 – 128.

Stige, L. C., Helgesen, K. O., Viljugrein, H. & Qviller, L. (2021), 'A statistical mechanistic approach including temperature and salinity effects to improve salmon lice modelling of infestation pressure', *Aquaculture Environment Interactions* **13**, 339 – 361.

Torrissen, O., Jones, S., Asche, F., Guttormsen, A., Skilbrei, O., Nilsen, F. & Horsberg, T.E.and Jackson, D. (2013), 'Salmon lice — impact on wild salmonids and salmon aquaculture', *Journal of Fish Diseases* **36**(3), 171–194.

Vormedal, B. I. & Larsen, M. L. (2021), 'Regulatory processes for setting sensitive-period sea-lice thresholds in major salmon producer jurisdictions: An evaluation', `https://asc-aqua.org/wp-content/uploads/2022/03/Regulatory-Processes-for-Setting-Sensitive-Period-and-Sea-lice-Thresholds-in-Major-pdf`. Accessed: 2023-04-15.

Walde, C. S., Jensen, B. B., Pettersen, J. M. & Stormoen, M. (2021), 'Estimating cage-level mortality distributions following different delousing treatments of atlantic salmon (salmo salar) in norway', *Journal of Fish Diseases* **44**(7), 899–912.

Walde, C. S., Stormoen, M., Pettersen, J. M., Persson, D., Røsæg, M. V. & Bang Jensen, B. (2022), 'How delousing affects the short-term growth of atlantic salmon (salmo salar)', *Aquaculture* **561**, 1–11.