

Towards Concept Change Detection in Marine Ecosystems

Daniel Lukats^{*,¶}, Elmar Berghöfer^{*}, Frederic Stahl^{*,†}, Janina Schneider^{*,‡},
Daniela Pieck[‡], Mobin M. Idrees[†], Lars Nolle^{*,§}, Oliver Zielinski^{*,‡}

^{*}Marine Perception Research Department

German Research Center for Artificial Intelligence, Oldenburg, Germany

[†]Department of Computer Science

University of Reading, Reading, United Kingdom

[‡]Institute for Chemistry and Biology of the Marine Environment

University of Oldenburg, Oldenburg, Germany

[§]Department of Engineering Sciences

Jade University of Applied Sciences, Wilhelmshaven, Germany

[¶]Corresponding author: Daniel.Lukats@dfki.de

Abstract—The research presented in this paper aims to accelerate the natural science research process by partially automating the execution of experiments using AI-assisted Concept-Change Detection (C-CD), e.g., for monitoring systems and studying biodiversity and ecosystem functions. The purpose of C-CD is to detect concept changes, also known as concept drift, that may be relevant to the study or ecosystem state. For example, in intertidal marine ecosystems, the event of sudden flooding can lead to dramatic changes in biodiversity. It could also be of scientific interest to take sensor samples more frequently in the period leading up to such events. The paper proposes an architecture for C-CD to customize AI-based analysis of sensor data streams. Furthermore, the paper implements portions of the architecture and is applied on sensor data from the Spiekeroog Coastal Observatory (SCO) as a feasibility study. The study demonstrates C-CD’s ability to detect anomalies that are either of scientific or technical interest to the operation and exploration activities of SCO.

Index Terms—Anomaly Detection, Concept Change Detection, Marine Ecosystems, Marine Sensor Systems

I. INTRODUCTION

Oceans and coastal areas are critical for food supply, climate regulation, transportation, energy production, and quality of life. There is an urgent need for easier fact-based ocean management to protect the oceans and ensure sustainable use of ocean resources. This calls for reliable general monitoring of changes in marine environments and health characteristics, as well as monitoring for sustainable and cost-effective industrial exploitation of oceanic and coastal areas. Data from such systems will give valuable input to industry, governmental bodies, and researchers focusing on oceans and climate.

The overarching motivation for the United Nations’ *Decade of Ocean Science for Sustainable Development (2021-2030)* [1] is to “support efforts to reverse the cycle of decline in ocean health and create improved conditions for sustainable development of the ocean”. In the JPI Oceans Strategic Research and Innovation Agenda 2015-2020 [2], supported by the European Union, a key expectation is “strengthening observation and monitoring capacities through enabling technologies,

new platforms and sensors; addressing under-sampling, and ensuring that new environmental parameters can be rapidly and accurately measured”. Moreover, the OECD report [3] emphasizes “[...] the drive for miniaturization and automation, the growing demand for low-power, low-cost devices for the measurement and graphic display of the physical environment” and moves to endow the sensor itself with intelligence. The last three decades have seen substantial progress in sensors and (smart) sensor systems, enhancing our monitoring capacities. Still there are two challenges, firstly essential biological and geochemical parameters are not accessible (yet) with sensors [4] and secondly, sampling strategies need to be adaptive and intelligent [5].

To overcome these challenges, a first prototype of an automatic Concept-Change Detection (C-CD) technique in marine environments based on machine learning algorithms is presented. In machine learning, concept change is the drift of the relationship between input and output data over time. For example, sensor readings may change over time such that $P_{t_1}(X) \neq P_{t_2}(X)$ given $t_1 \neq t_2$, with X denoting a random variable representing the sensor reading. Whereas concept change denotes long-term changes, anomalies represent short-term deviations. Both concept changes and anomalies might indicate the presence of events pertaining to a natural scientist’s research interests. C-CD can be used for systems monitoring and applications, such as investigating biodiversity and ecosystem functioning. The system automatically detects concept changes that might be relevant to the investigation or the health of the ecosystem. For example, in intertidal marine ecosystems the event of sudden flooding may cause dramatic changes to biodiversity, and it may also be of scientific interest to take more frequent sensor samples during the time leading up to such events. The presented technique can be used to trigger more frequent data sampling.

C-CD extends to multimodal data streams, for example a water level sensor on its own may not be enough to predict sudden flooding, only to detect the flood as it is already

happening. However, a volatile water level combined with a change of the direction of wind, and reports of rain in certain locations, may indicate imminent flooding. Here the system needs to combine and analyze diverse kinds of data in real-time. C-CD draws on existing algorithms to detect changes in data streams such as Drift Detection Method, Early Drift Detection Method, Micro-Cluster Nearest Neighbor [6]–[8] concept change detection techniques and anomaly detection algorithms such as Isolation Forests or Local Outlier Factor [9], [10]. However, such techniques expect clean and accurate data and are often limited to certain types of change [8].

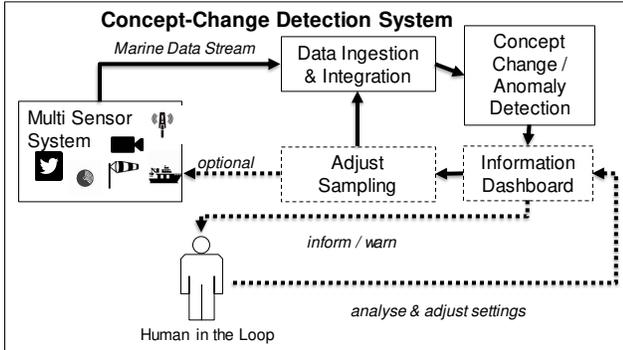


Fig. 1: Concept-Change Detection system overview

Fig. 1 depicts C-CD conceptually. Solid lines represent aspects of C-CD realized and presented in this paper, and dashed components are ongoing developments. C-CD is applicable in real-time and allows prompt responses to sensors and scientists by automating the data pipeline from data ingestion and integration, sensor feedback to real-time result production. The analyst also works as a human in the loop and may adjust parameter settings in real-time. A prototype of C-CD is tested using real data from Spiekeroog Coastal Observatory—as a proof of concept—and is presented in this paper. The results show C-CD’s capability to detect changes in multimodal marine sensor systems in real-time.

The paper is organized as follows: Section II discusses related work and Section III discusses the developed C-CD system requirements, architecture and algorithms. The following section, Section IV, provides an experimental evaluation of the C-CD prototype. Finally, ongoing and future work is discussed in Section V; followed by concluding remarks in Section VI.

II. RELATED WORK

A. Marine Environment Monitoring

Biodiversity is changing at an unprecedentedly high rate [11], reflecting the anthropogenic alteration of Earth’s ecosystems [12]. Consequently, research on biodiversity-ecosystem relationships including experimental setups and monitoring initiatives has become of major interest [13]–[15].

Marine environments are influenced by a wide diversity of anthropogenic and natural substances and organisms that may

have adverse effects on human health and ecosystems. Real-time measurements of biochemical parameters and marine pollutants across a range of spatial scales are required to adequately monitor ecosystem health and potential hazards [16]. Significant technological advancements have been made in recent years for the detection and analysis of the marine ecosystem status [17]. In Germany this is especially focused on observatories, like COSYNA [18], [19] or the Spiekeroog Coastal Observatory (SCO, operated by the Institute for Chemistry and Biology of the Marine Environment). In particular, multispectral sensors deployed on a variety of mobile and fixed-point observing platforms provide a valuable means to assess status, dynamics and hazards alike [20]. The authors of [21] classified sensors by their adaptability to various platforms, addressing large, intermediate, or small areal scales, identifying an urgent need for new sensors to detect changes in marine ecosystems at all scales in autonomous real-time mode.

Current progress in sensor technology is expected to depend on the development of small-scale smart sensor technologies with a high sensitivity and specificity towards target analytes or organisms [22]. However, deployable systems must comply with platform requirements as these connect the three areal scales. Future developments will include the integration of data stream mining and machine learning [23] into complex and operational sensing systems, enabling a comprehensive situational awareness and long-term monitoring. The last three decades have seen substantial progress in sensors and (smart) sensor systems, enhancing our monitoring capacities [24]. Still there are two challenges a) as essential biological and geochemical parameters are not accessible (yet) with sensors [4] and b) sampling strategies need to be adaptive and intelligent [5].

B. Concept Change and Anomaly Detection

Traditional machine learning for data mining builds its models on static batch training sets, enabling several iterations over the data. This is different in Data Stream Mining (DSM) as (new) models need to be built in linear or sub-linear time complexity [25]. Furthermore, DSM techniques need to enable dynamic adaptations to concept changes, which are changes in the patterns encoded in the stream [25]. Data mining models must reflect an accurate representation of the current pattern. In natural sciences, sensor systems represent an application for DSM techniques, where detecting such change is essential for applications.

Techniques to detect concept change exist, e.g., EDDM, DDM, MC-NN [7], [8] to name a few. However, they expect relatively clean and accurate data and are limited to certain types of change—gradual, sudden, recurring or incremental for example. These change detectors can identify events of interest to the natural scientist. However, research of obtaining tangible event identification is often tailored for specific applications, such as telecommunications [26] or chemical process industry [27]. There is very little work towards general purpose early

event detection, and thus C-CD will close this gap for natural science applications making use of sensor networks.

As mentioned earlier, change detection techniques expect relatively clean and accurate data, yet quite often this is not the case for Natural Science sensor networks. Here issues such as calibration changes, faults or extreme weather conditions influence the quality of the data. Hence, change detectors need to be accompanied by appropriate pre-processing techniques. Yet, off-the-shelf pre-processing techniques from standard machine learning toolboxes are often not fit for this purpose due to the variability of the data [8]: For example minimum and maximum values of a sensor may change over time, rendering previous normalization operations inaccurate. An attempt at integrating several pre-processing techniques into a single multi-purpose tool has been co-developed by one of the investigators [8].

Some fully real-time adaptive (data instance by data instance) algorithms do exist that mimic standard data mining techniques such as the adaptive predictive Concept Drift Very Fast Decision Tree (CVFDT) [28], Hoeffding Rules [29], Very Fast Decision Rules (VFDR) [30], or descriptive adaptive algorithms such as CluStream (Cluster Analysis) [31], Adaptive Generalised Rules [32], etc. However, please note that predictive DSM algorithms tend to be supervised and thus are not suitable for all natural science applications, as they need frequent and timely feedback about their predictive performance in order to adapt. Some predictive hybrids exist that have alternative feedback mechanisms in addition to predictive performance, such as [33], [34].

Compared with concept change, anomalies are short-term changes of the pattern encoded in the data stream, in extreme cases an anomaly may comprise only a single data instance, i.e., an outlier. Anomalies are difficult to learn from machine learning algorithms, since typically there are few examples to learn from. However, a common approach to detect anomalies is to learn a model that reflects normal behaviour and to flag up data that does not fit this model well. Common algorithms used in anomaly detection are for example Isolation Forests [9], Local Outlier Factor [10] and k-Nearest Neighbors [35]. A recently published paper describes an ensemble anomaly detection method for detecting computer network attacks. A comprehensive survey on anomaly detection methods is available in [36].

C. Data Stream Processing Frameworks and Libraries

A few relevant software frameworks and libraries have been developed in recent years. A workbench for benchmarking DSM algorithms and implementing DSM applications is the Massive Online Analysis (MOA) framework [23]. MOA implements a couple of standard predictive and descriptive DSM algorithms, but also offers other important DSM techniques, such as outlier detection, concept drift detection, etc. However, MOA does not offer workflow building mechanisms, lacks pre-processing techniques and only offers limited visualisation facilities. The Open Mobile Miner [37] is an early platform for situation aware DSM. However,

OMM is tailored for resource constraint mobile environments such as smartphones and implements versions of DSM algorithms which can be computationally tailored to available computational resources. However, this often comes with a trade-off to analytics performance. Some recent developments on data stream processing frameworks are Apache Flink (<https://flink.apache.org/>), Storm (<http://storm.apache.org/>) and SAMOA (<https://samoa.incubator.apache.org/>). Apache Flink offers an execution framework for building real-time data stream mining workflows. There, Apache Flink excels at scalability by parallelizing the workflow execution. Another parallel stream processing framework is Storm. However, both Apache Flink and Storm lack true instance by instance DSM algorithms as needed by natural scientists in the field. SAMOA on the other hand, incorporates instance by instance parallel data stream mining techniques. Furthermore, the parallelization capability comes as a considerable overhead, and thus Apache Flink, SAMOA & Storm are more suitable to real-time data processing in dedicated computer clusters rather than on mobile hardware systems suitable for field campaigns. Odysseus (<https://www.offis.de/offis/projekt/odysseus.html>) is a framework for constructing custom event stream management systems and provides a set of modules as foundation for processing event streams. Although Odysseus is relatively lightweight, a holistic implementation of C-CD for mobile field computing infrastructure is outside the scope of this paper and subject to future research.

III. C-CD: CONCEPT CHANGE DETECTOR SYSTEM

This section describes the C-CD System in detail. Hereby, Section III-A explains the requirements of C-CD and Section III-B discusses the C-CD architecture, while Section III-C gives details on the algorithms and workflow implemented in the C-CD prototype.

A. Requirements

For the C-CD System under development, the following requirements have been identified:

a) *Anomaly detection*: The definition of an anomaly depends on the application domain, in this work we assume that anomalies are outliers. Due to a potentially large quantity of data in marine data streams labelling, anomalies may be prohibitively expensive. Therefore, unsupervised anomaly detection algorithms must be employed, which commonly detect anomalies by identifying outliers [38].

b) *Concept change detection*: Concept change may occur suddenly, incrementally, gradually or recurring as illustrated in Fig. 2. We argue that at least three types of concept changes exist in marine ecosystems, incrementally changing, sudden changes and gradual changes. The tidal cycle, the day and night cycle as well as seasonal effects may cause incremental changes, e.g., to the wind direction, water temperature or air humidity [39]. Sudden changes may occur during abrupt weather changes such as storms or floods. Lastly, sensor degradation may be caused by constant exposure to saline water and in turn lead to gradual changes of the pattern

encoded in the sensor data stream. Hence, there is a need for C-CD to detect sudden, incremental and gradual concept changes over time.

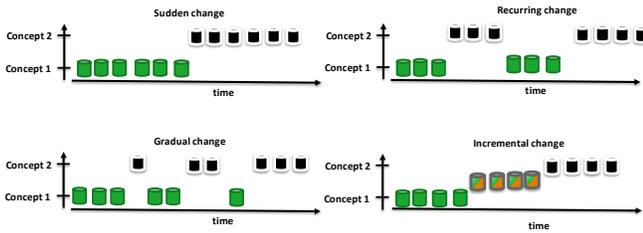


Fig. 2: Types of concept change

c) *Missing value imputation*: A variety of reasons can cause missing values. Individual sensor outages as well as power outages can cause loss of individual sensor readings or the loss of an entire data frame respectively. The same holds true for maintenance, during which parts of the system need to be shut down. Thus, strategies for processing or imputing missing values must be incorporated.

d) *Self-contained mobile platform for fieldwork*: A system such as C-CD needs to be applicable in the field within different contexts, for example on board of a research vessel. Therefore, the system needs to be integrated into a self contained mobile mission control system, which is easy to configure for different applications of C-CD.

B. System Architecture

The architecture of C-CD and its context in natural sciences are given in Fig. 3. The dotted line represents the system’s boundaries.

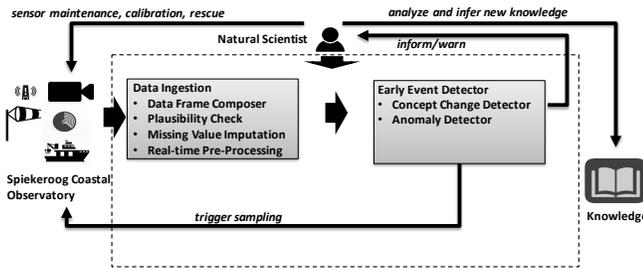


Fig. 3: C-CD system architecture

Outside the boundaries are the users (natural scientists), the sensor ecosystem (data hub)—in our case the *Spikeeroog Coastal Observatory* (SCO)—and knowledge derived by the natural scientists. User-defined requirements determine the use-cases and create application-specific workflows to be executed in the system. In the figure, the sensor data are generated by the sensors (SCO data hub) and flows from left to right. Essentially, a natural scientist defines meta-data which describes the investigated scenario and data or system configurations at a high level, i.e., which sensors are available and what data is expected from each sensor. The system then ingests the data and processes it using implemented machine

learning techniques to detect imminent concept changes of the sensor system and their causality.

Based on the type of identified imminent change, the system may issue a warning or inform the natural scientist to act upon. For example, the scientist may make use of the detected change to infer new knowledge in their chosen field, or they may adjust the sensor system through calibration, maintenance or rescue missions. However, the C-CD system may also communicate with the sensor system directly: an imminent flooding event may be interesting to the research question (as specified in the meta-data) and thus C-CD triggers an automated sampling system to produce more fine-grained research data for the scientist. The system comprises 2 modules, (a) data ingestion and (b) early event detector:

a) *Data Ingestion*: This module provides a data frame composer, plausibility checker and real-time missing value imputation (see left-hand side of Fig. 3). In multimodal data streams, sensor readings may be taken at different points in time. It is the data frame composer’s responsibility to resolve these issues and create a sensible data frame. In the SCO case study, this component was not required, and its development is still ongoing.

The plausibility checker optionally checks the consistency of the data received as compared with the required meta-data format, data resolution specification, provenance etc. to ensure that the data frames received satisfy the requirements and to avoid processing any miss-matched data types or obviously corrupted data. For example, detecting a negative wind speed is not possible and should be flagged by the plausibility checker. The plausibility checker is the only but a very much necessary component, which is not based on AI.

Sensors may not always deliver readings, and thus missing values are a very likely occurrence. The missing value imputation module implements strategies which are either defined by the natural scientist or by the system itself, depending on the context and circumstances. Such strategies range from simple averaging, deletion of data records with missing values, to the more complex prediction of missing values.

b) *Early Event Detector*: The early event detector receives the cleaned data frames from the data ingestion component (see right-hand side of Fig. 3). The sub-components of this module are based on artificial intelligence techniques. Intelligent adaptive pre-processing prepares the data for the later components, it comprises techniques such as real-time normalization or outlier detection. Normalization techniques are provided by the data stream mining framework River [40].

The change detector component comprises various techniques to detect potential medium to longer term changes of the pattern encoded in the data stream. As ground truth concept changes are rarely—if ever—available, concept change detectors typically are supervised in the sense that they require labelled data to observe supervised algorithms. Although no labels are available in our case study, we use supervised concept change detectors due to a lack of readily available unsupervised concept change detectors. In the future, the use of unsupervised concept change detectors is desired.

Whereas concept changes are longer lasting pattern changes, the anomaly detector can flag unusual data readings that potentially constitute short term pattern changes in the data. These short-term changes might indicate to the natural scientist that an event such as a severe weather change or damage to the system is happening [41].

C. Algorithms

Two types of algorithms are currently implemented in C-CD, anomaly detection algorithms and concept change detection algorithms:

a) *Anomaly detectors*: In order to detect anomalies, a heterogeneous ensemble of outlier detectors is used. With few exceptions, most detectors used in this work detect outliers either through distance-based, spatial detection or through the construction of probability distributions. The ensemble consists of the following algorithms:

- Clustering Based Local Outlier Factor (CBLOF) determines the anomaly score by clustering input data and evaluating the size of the cluster the data was assigned to, as well as the distance of the data to large clusters [42].
- Copula Based Outlier Detector (COPOD) identifies outliers by computing copulas on the input data. Anomaly scores are calculated by predicting probabilities for each data point [43].
- Histogram-based Outlier Score (HBOS) predicts outliers by constructing independent histograms for each feature of the input data [44].
- Isolation Forest (iForest) performs anomaly detection by isolating anomalies rather than learning models of inliers [9].
- k-Nearest Neighbors (KNN) outlier detection determines anomaly scores as a measure of distance to each data point's nearest neighbours [35].
- Lightweight On-line Detector of Anomalies (LODA) employs an ensemble comprised of multiple weak anomaly detectors. For the weak detectors, one-dimensional histograms are constructed from the input data [45].
- Local Outlier Factor (LOF) identifies outliers by observing the density of the neighbourhood of each data point [10].
- Principal Component Analysis (PCA) outlier detection uses singular value decomposition and compares data points based on their eigenvectors and eigenvalues [46].

b) *Concept change detectors*: The most commonly used concept change detectors observe their learner's prediction error rate. In general, they signal a concept change, if the prediction error rate rises above a threshold determined by the respective detector. We use the following concept change detectors in our evaluation:

- Adaptive Windowing (ADWIN) detects concept changes by comparing two windows of recent prediction errors [47]. The size of the windows is automatically adapted by the algorithm, and concept change is determined by a statistical test of the average errors of the two windows.



Fig. 4: Island of Spiekeroog located in the German Bight in the North Sea

- Drift Detection Method (DDM) assumes a consistently decreasing error rate. A concept change is detected, if the learner's error rate increases and surpasses a threshold determined from the error rate and its standard deviation [6]
- Early Drift Detection Method (EDDM) is a modification of DDM that tracks not only the error rate but also the time between errors [7].
- $HDDM_A$ and $HDDM_W$ are two variants of a concept change detection method based on probability inequalities. $HDDM_A$ considers a moving average of the learner's error rate and derives a confidence bound for concept change detection from Hoeffding's inequality. $HDDM_W$ follows the same approach, however it observes a weighted moving average and derives its confidence bound from McDiarmid's inequality [48].

IV. EXPERIMENTAL EVALUATION ON SPIEKEROOG COASTAL OBSERVATORY CASE STUDY

This section describes the application of C-CD in a case study on data from the Spiekeroog Coastal Observatory (SCO). Section IV-A and IV-B give an overview of the case study area and the available data. In Section IV-C the workflow for the anomaly detection and concept change detection is explained, whereas the evaluation of the results for both is presented in Section IV-D.

A. Case Study Description

The data used in this paper originate from a large scale and long-term coastal observation onto and around the German island of Spiekeroog in the southern North Sea (see Fig. 4). As part of the Wadden Sea, it belongs to the UNESCO world natural heritage since 2009. Spiekeroog accommodates the Spiekeroog Coastal Observatory (SCO), which consists of different elements for marine and terrestrial ecosystem research and started in 2002 with a time series station to measure oceanographic, meteorological, and biogeochemical

data. Since then, it is under continuous development with different research sites distributed across the whole area and a dedicated research center added. As a part of the SCO, artificial islands and a weather measurement station were used for an environmental conditions and biodiversity study of the back-barrier salt-marsh of the island [39]. The weather station is located approximately 500 m north of the southern shoreline (53°45'57.10" N, 007°43'34.11" E) and installed at 10 m height. It was the primary data source for the evaluations, because of familiarity with the environment and easy access to raw data as well as domain expert feedback. Although there are no labeled anomalies available, we expect to find anomalies during rare weather events and consult Deutscher Wetterdienst's (German Meteorological Services) weather reports for the North Sea region [49] for validation.

B. Data Set

The data used in this paper originate from the previously described weather station and covers a period of about five years, from November 2014 until September 2019. The sampling interval is 1 min with an averaging time of 10 s. From the 18 variables recorded at the weather station, 8 were chosen for the application: wind speed, wind direction, air temperature, relative humidity, pressure, maximum brightness, brightness direction and precipitation. Among the 10 unused attributes there are 4 further attributes related to brightness, 3 more attributes related to precipitation, 2 attributes for solar elevation and azimuth angle and the date of the data frame. These were dropped due to redundancies or low correlation with other values.

Due to malfunction of the pressure sensor (July 2015 to March 2016) and maintenance and re-calibration by the manufacturer (April 2016 to September 2016 and July 2017 to February 2018), gaps in the data set exist accordingly. This results in a data availability of 64 % and therefore an overall amount of about 1.6 million data points. The raw data are provided by [39] upon request; processed data was published by the authors using PANGAEA data publisher services [50]–[52]. In their processing step, the authors erased defective recordings and removed outliers. As a basis for our analysis, the unprocessed raw data was used, provided by the authors of the original study. Some post-processing that was necessary to correct wind direction values (because the sensor could not be properly aligned to North) were reapplied.

C. Workflow

The requirements (a) anomaly detection, (b) concept change detection and (c) missing value imputation have been partially implemented. Requirement (d) self-contained mobile platform is subject to future work. The components have been implemented in Python.

Pre-processing operations such as missing value imputation and data normalization are executed first. Afterwards, anomaly detection and concept change detection can be performed simultaneously.

a) *Anomaly detection*: Anomaly detection is performed using a heterogeneous ensemble of unsupervised anomaly detectors, as outlined in Section III-C. The detectors are provided by the PyOD library [53]. We normalize the data using the min-max normalization:

$$x' \doteq \frac{x - \min(x)}{\max(x) - \min(x)}$$

where x' is the normalized data and x represents the training data. The same minimum and maximum values are used to normalize the test data.

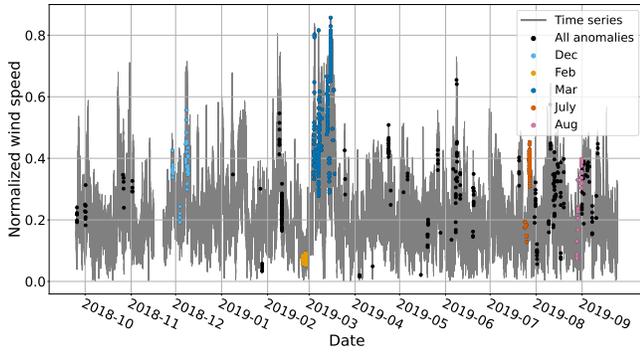
The anomaly detectors are trained on all data from 2014-11-19 through 2018-09-24, which leaves exactly one year of data for testing. This results in a train-test split of 68% training data and 32% test data. Then, the threshold for anomaly detection is determined by ordering anomaly scores on the training data and choosing the top λ , $0 \leq \lambda \leq 1$ anomaly score as a cut-off. Any anomaly score higher than this cut-off will be regarded as an anomaly in the test set. Assuming a similar distribution of anomaly scores in the training and test data, this should result in roughly the same proportion of the test data being labelled as anomalies. This value can be customized by the natural scientist, e.g., to account for sampling capacities of auto samplers. In this work, we determine the value of λ empirically and choose $\lambda = 0.002$.

b) *Concept change detection*: Concept change detection is based upon a modified pipeline in River [40]. Data are normalized using an online min-max normalization provided by the River framework, which maintains a running minimum and maximum.

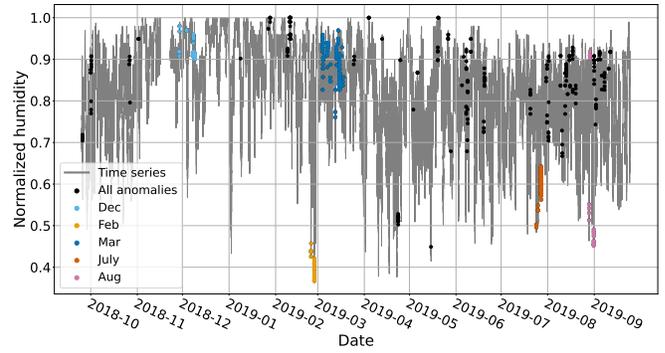
Classifiers are trained on data as they are observed in an interleaved test-train scheme in online machine learning. This means that classifiers are prompted to predict the next label first based on the provided features. Only after predicting the label do they receive the true label, which they may then use for further training. Whilst the data stream is being processed by the classifiers, the concept change detectors observe the error rate of their respective classifier. Once they detect a concept change, the corresponding classifier is re-trained starting with recent data.

Since our data stream contains no labels assigned by a domain expert, we generate a pseudo-label by binning temperature values. Temperature values were chosen as the pseudo-label, because analysis has shown that they are the most strongly correlated with other features. Finally, we create 10 bins of equal size resulting in roughly 160000 items per class. Naturally, the temperature readings are not included in the features anymore.

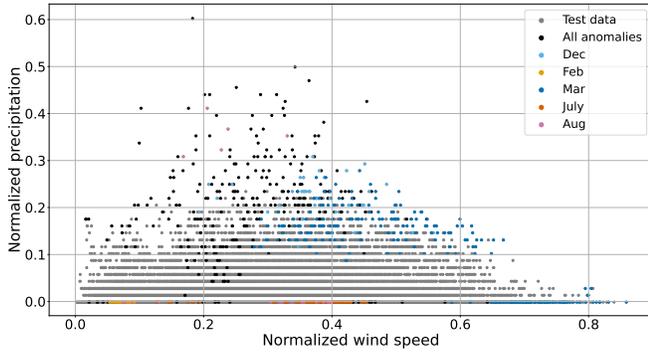
c) *Missing value imputation*: Missing value imputation is currently facilitated by replacing the missing value with the last known valid value. However, future versions of C-CD will implement more sophisticated missing value imputation strategies and develop missing value imputation strategies that take concept change into consideration. A possible approach that merits further testing is shown in Section IV-D.



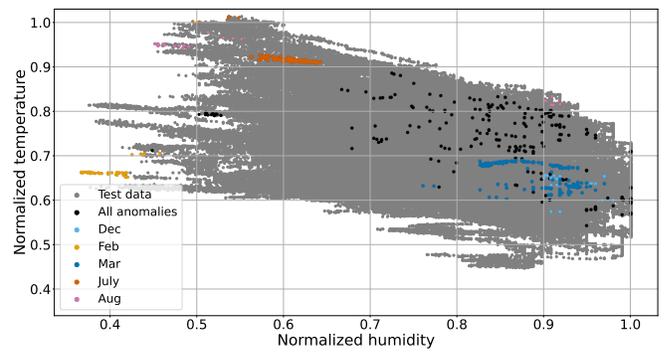
(a) Wind speed time series



(b) Humidity time series



(c) Wind speed & precipitation



(d) Humidity & temperature

Fig. 5: Anomaly detection results with 5 groups of anomalies highlighted for easier identification. In text, these groups are highlighted using the same color.

D. Evaluation

a) *Anomaly detection*: Currently C-CD implements the anomaly and concept change detection components. In order to evaluate detected anomalies and concept change, a domain expert about SCO was consulted to confirm anomalies and concept change detected by C-CD. Also, multiple anomalies are aligned with noteworthy weather events. After showing the impact of concept change detection in a classification task, the benefit on regression is demonstrated. In future development, regression combined with concept change detection could be used for missing value imputation [54].

In Fig. 5a and 5b, the detected anomalies are displayed on a time series plot of the wind speed and humidity, respectively. Additionally, scatter plots in Fig. 5c and 5d highlight anomalies as arguments of wind speed and precipitation as well as humidity and temperature. On the one hand, the plot in Fig. 5c was chosen, because it shows the prevalence of anomalies featuring high wind speed and precipitation; these anomalies might indicate the presence of a storm. On the other hand, Fig. 5d alerted us to the presence of heat waves and near-tropical weather conditions in the North Sea. In the following, certain months are highlighted in the same color as their counterparts in the aforementioned figures.

We confirmed that detected anomalies align with rare weather by visual inspection. Further comparison of detected

anomalies with weather reports for the North Sea region by Deutscher Wetterdienst (German Meteorological Services) show that several anomalies align with storms [49]:

- Anomalies in early **December 2018** align with the *MARIELOU* storm front, which featured above average wind speeds, precipitation and temperature.
- Anomalies in the first half of **March 2019** align with a series of storm fronts, which brought high wind speed and above average precipitation.

Other anomalies, such as the anomalies detected in late **February 2019**, late **July 2019** and late **August 2019**, feature low humidity paired with high temperatures. In fact, Deutscher Wetterdienst noted the same irregularities. They even declared the event in late July a heat wave and stressed the unusual temperatures in August [49].

However, not all anomalies can be attributed to such events: For example, low humidity anomalies in May 2019 do not align with unusual temperature events. Additionally, only a few anomalies in the first half of August 2019 can be attributed to storm fronts; other anomalies in that time span do not coincide with notable events from weather reports of Deutscher Wetterdienst.

These preliminary results indicate the capability of the proposed system to detect notable anomalies, which might merit further research by natural scientists for example through

analysis of water samples. On the other hand, they also show the need to further refine and improve the anomaly detection component of C-CD. We outline further steps to accomplish this goal in Section V.

b) Concept change detection: We evaluate the concept change detection component of C-CD by examining the performances of classifiers enhanced with concept change detectors and comparing them to base classifiers trained without concept change detectors. Again, since the SCO data sets are unlabeled, pseudo-labels are generated by creating 10 bins of equal size from temperature values. We assume that the presence of concept changes will have a notable impact on the performance of classifiers. Therefore, enhancing these classifiers with concept change detectors should result in a notable improvement of the respective classifier’s performance.

To this end, an F_1 score is computed over the entire data stream and a global performance metric across all 10 classes is obtained by performing micro-averaging [55]. The F_1 score is the harmonic mean of two different metrics, the precision and the recall. According to [55], precision denotes the “proportion of instances classified as positive that are really positive”, whereas recall denotes the proportion of positive instances classified as such. In a multi-class setting like in this case study, precision and recall are computed by regarding a single class as positive and all others as negative. We do not state the precision and recall in our evaluation, since the micro-averaging results in the global precision and recall having the same value as the global F_1 score.

The concept change detectors given in Section III-C are tested with two base classifiers each, a Hoeffding Tree (HT) and a Naive Bayes (NB) classifier [55], as the detectors depend on the classifier’s prediction error rate.

Table I contains the F_1 scores calculated over the entire data stream. The base classifiers—HT and NB—show notably worse performances than their enhanced counterparts, with NB classifier performances being slightly lower than that of HT in general. As the use of concept change detectors leads to an improvement of the F_1 score of at least 0.67 in the case of the NB classifier and an improvement of at least 0.37 in the case of the HT, we argue that concept changes are present in the given data.

To further support our claim and highlight a possible immediate use of concept change detection in the context of marine ecosystems, we perform a regression on temperature values. As before, we compare a base regressor with regression methods supported by concept change detectors and measure their performance by tracking the mean squared error of each regressor. The mean squared error is given by $\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$, with y denoting the true value and \hat{y} denoting the predicted value.

The mean squared errors of a HT regressor and a k-nearest neighbor (KNN) regressor with as well as without an ADWIN concept change detector are displayed in Table II. Resetting the models after detecting a concept change leads to a clear improvement of the performance in the case of the HT; the average mean squared error over the entire data stream is

TABLE I: Comparison of base classifiers and classifiers enhanced with concept change detectors on binned temperatures

Classifier	Concept change detector	F_1 score
Hoeffding Tree	—	0.542
Hoeffding Tree	ADWIN	0.941
Hoeffding Tree	DDM	0.912
Hoeffding Tree	EDDM	0.950
Hoeffding Tree	HDDM _A	0.960
Hoeffding Tree	HDDM _W	0.952
Naive Bayes	—	0.214
Naive Bayes	ADWIN	0.934
Naive Bayes	DDM	0.885
Naive Bayes	EDDM	0.945
Naive Bayes	HDDM _A	0.963
Naive Bayes	HDDM _W	0.951

TABLE II: Comparison of mean squared errors of temperature value regression for base regressors and those supported by ADWIN

Regressor	Concept change detector	Mean squared error
Hoeffding Tree	—	2.352
Hoeffding Tree	ADWIN	0.601
KNN	—	0.044
KNN	ADWIN	0.042

~ 1.7 lower. In the case of the KNN the performance is improved very little with an average mean squared error of 0.044 for the base regressor and one of 0.042 for the enhanced one. Nevertheless, this result supports the claim that concept changes are present in the SCO data set and that various algorithms may benefit from concept change detection—albeit to differing degrees.

V. ONGOING AND FUTURE WORK

Since this first feasibility study of C-CD on data from Spiekeroog Coastal Observatory shows promising results, further developments of the system are ongoing or planned to improve support for natural scientists.

In order to allow the use of further sensor systems in the future, a data frame composer (see Fig. 3) is currently being developed to allow the use of data sources with different measurement frequencies. This component will implement methods to align data from different sources/sensors with potentially different times of measurement and different temporal resolution.

Due to sensor faults and maintenance, sensor readings may be corrupted or outright missing. To maintain reliable anomaly detection and concept change detection in the presence of these issues, reliable missing value imputation methods are currently being developed.

As highlighted in Fig. 1, an information dashboard is subject to future work to provide users an overview of current data stream properties, anomalies and concept changes. Based on

the information presented in this dashboard, users will then be able to adjust the sampling rates of samplers attached to the system, e.g., of automated water samplers. Furthermore, the system shall be integrated into a mobile mission control system hardware to enable the use of C-CD in field experiments.

Finally, the prototype developed for this paper uses supervised concept change detectors that must observe a learner's error rate. In turn, this means that a supervised classification algorithm is required to allow the use of these detectors. Because there are no labels assigned to the data set used in the case study, we used the temperature as a pseudo-label. For a more widely applicable solution, unsupervised concept change detectors that operate on the data stream with no immediately available ground truth are desired, investigation in this direction is ongoing.

VI. CONCLUSION

In this paper we outlined the motivation and aims of a Concept-Change Detection (C-CD) system that detects short-term as well as long-term changes in data patterns to support the research of natural scientists. The architecture of C-CD contains pre-processing techniques as well as the aforementioned change detection algorithms, namely anomaly detection and concept change detection algorithms. On top of developed components, opportunities for further development and improvement of the system were highlighted. First and foremost, more sophisticated missing value imputation techniques and a data frame composer enabling the use of multimodal data sources will be implemented.

The anomaly detection and concept change detection components included in C-CD were evaluated on a case study with data from the Spiekeroog Coastal Observatory. In this case study, several anomalies could be detected and matched to noteworthy and unusual weather events. Furthermore, the presence of concept changes was empirically demonstrated as well as the benefits employing concept change detection can have on regression for missing value imputation.

Based on the results of this study we conclude that the desired concept change detection system is feasible. We will further expand the used data sources and evaluate C-CD's capabilities in cooperation with natural scientists in a real-time setting in taking meaningful water samples in the Spiekeroog Coastal Observatory.

ACKNOWLEDGMENT

The research in this paper is supported by the Ministry for Science and Culture, Lower Saxony, Germany, through funds from the Niedersächsische Vorab (ZN3480, ZN3683, and ZN3727). This paper has received partial funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101000825 (NAUIT-LOS). We thank Kai Schwalfenberg and Claudia Thölen for their support with the Spiekeroog Coastal Observatory. Thanks to the Institute for Chemistry and Biology of the Marine Environment's workshop team and all persons involved in fieldwork on Spiekeroog.

REFERENCES

- [1] United Nations, "The Science we Need for the Ocean We Want: The United Nations Decade of Ocean Science for Sustainable Development (2021-2030)," Paris, 2020. [Online]. Available: https://www.oceanandecade.org/assets/The_Science_We_Need_For_The_Ocean_We_Want.pdf
- [2] JPI Oceans, "Strategic Research and Innovation Agenda 2015–2020," 2015.
- [3] OECD, *The Ocean Economy in 2030*, 2016. [Online]. Available: <https://www.oecd-ilibrary.org/content/publication/9789264251724-en>
- [4] Konsortium Deutsche Meeresforschung, "Altonaer Erklärung," 2015. [Online]. Available: <https://www.deutsche-meeresforschung.de/de/altona>
- [5] MBARI, "Technology roadmap." [Online]. Available: <https://www.mbari.org/2014-technology-roadmap>
- [6] J. Gama, P. Medas, G. Castillo, and P. Rodrigues, "Learning with Drift Detection," in *Advances in Artificial Intelligence – SBIA 2004*, ser. Lecture Notes in Computer Science, A. L. C. Bazzan and S. Labidi, Eds. Berlin, Heidelberg: Springer, 2004, pp. 286–295.
- [7] M. Baena-Garcia, J. del Campo-Ávila, R. Fidalgo, A. Bifet, R. Gavalda, and R. Morales-Bueno, "Early drift detection method," in *Fourth international workshop on knowledge discovery from data streams*, vol. 6, 2006, pp. 77–86.
- [8] M. S. Hammoodi, F. Stahl, and A. Badii, "Real-time feature selection technique with concept drift detection using adaptive micro-clusters for data stream mining," *Knowledge-Based Systems*, vol. 161, pp. 205–239, Dec. 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0950705118304039>
- [9] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in *2008 Eighth IEEE International Conference on Data Mining*, Dec. 2008, pp. 413–422.
- [10] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: identifying density-based local outliers," *ACM SIGMOD Record*, vol. 29, no. 2, pp. 93–104, May 2000. [Online]. Available: <https://doi.org/10.1145/335191.335388>
- [11] J. Rockström, W. Steffen, K. Noone, Å. Persson, F. S. Chapin, E. F. Lambin, T. M. Lenton, M. Scheffer, C. Folke, H. J. Schellnhuber *et al.*, "A safe operating space for humanity," *nature*, vol. 461, no. 7263, pp. 472–475, 2009.
- [12] H. Hillebrand and B. Matthiessen, "Biodiversity in a complex world: consolidation and progress in functional biodiversity research," *Ecology letters*, vol. 12, no. 12, pp. 1405–1419, 2009.
- [13] T. Balke, K. Löhmus, H. Hillebrand, O. Zielinski, K. Haynert, D. Meier, D. Hodapp, V. Minden, and M. Kleyer, "Experimental salt marsh islands: A model system for novel metacommunity experiments," *Estuarine, Coastal and Shelf Science*, vol. 198, pp. 288–298, 2017.
- [14] G. Canonico, P. L. Buttigieg, E. Montes, F. E. Muller-Karger, C. Stepien, D. Wright, A. Benson, B. Helmuth, M. Costello, I. Sousa-Pinto *et al.*, "Global observational needs and resources for marine biodiversity," *Frontiers in Marine Science*, vol. 6, p. 367, 2019.
- [15] V. P. Mega, "The future of climate will be decided in cities and oceans," in *Eco-Responsible Cities and the Global Ocean*. Springer, 2019, pp. 85–119.
- [16] P. Farcy, D. Durand, G. Charria, S. J. Painting, T. Tamminen, K. Collingridge, A. J. Grémare, L. Delauney, and I. Puillat, "Toward a european coastal observing network to provide better answers to science and to societal challenges; the jerico research infrastructure," *Frontiers in Marine Science*, vol. 6, p. 529, 2019.
- [17] B. Baschek, F. Schroeder, H. Brix, R. Riethmüller, T. H. Badewien, G. Breitbach, B. Brügge, F. Colijn, R. Doerffer, C. Eschenbach *et al.*, "The coastal observing system for northern and arctic seas (cosyna)," *Ocean science*, vol. 13, no. 3, pp. 379–410, 2017.
- [18] H. M. Benway, L. Lorenzoni, A. E. White, B. Fiedler, N. M. Levine, D. P. Nicholson, M. D. DeGrandpre, H. M. Sosik, M. J. Church, T. D. O'brien *et al.*, "Ocean time series observations of changing marine ecosystems: an era of integration, synthesis, and societal applications," *Frontiers in Marine Science*, vol. 6, p. 393, 2019.
- [19] C. Moore, A. Barnard, P. Fietzek, M. R. Lewis, H. M. Sosik, S. White, and O. Zielinski, "Optical tools for ocean monitoring and research," *Ocean Science*, vol. 5, no. 4, pp. 661–684, 2009.
- [20] O. Zielinski, J. A. Busch, A. D. Cembella, K. L. Daly, J. Engelbrektsen, A. K. Hannides, and H. Schmidt, "Detecting marine hazardous substances and organisms: sensors for pollutants, toxins, and pathogens," *Ocean Science*, vol. 5, no. 3, pp. 329–349, 2009.

- [21] P. Fischer, "Intelligent sensor technology: A 'must-have' for next-century marine science," in *AI Technology for Underwater Robots*. Springer, 2020, pp. 19–36.
- [22] K. Malde, N. O. Handegard, L. Eikvil, and A.-B. Salberg, "Machine intelligence and the data-driven future of marine science," *ICES Journal of Marine Science*, vol. 77, no. 4, pp. 1274–1285, 2020.
- [23] A. Bifet, G. Holmes, B. Pfahringer, J. Read, P. Kranen, H. Kremer, T. Jansen, and T. Seidl, "Moa: a real-time analytics open source framework," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2011, pp. 617–620.
- [24] F. E. Muller-Karger, P. Miloslavich, N. J. Bax, S. Simmons, M. J. Costello, I. Sousa Pinto, G. Canonico, W. Turner, M. Gill, E. Montes *et al.*, "Advancing marine biological observations and data requirements of the complementary essential ocean variables (eovs) and essential biodiversity variables (ebvs) frameworks," *Frontiers in Marine Science*, vol. 5, p. 211, 2018.
- [25] M. M. Gaber, A. Zaslavsky, and S. Krishnaswamy, "A survey of classification methods in data streams," *Data streams*, pp. 39–59, 2007.
- [26] C. Wrench, F. Stahl, G. D. Fatta, V. Karthikeyan, and D. Nauck, "A rule induction approach to forecasting critical alarms in a telecommunication network," in *2019 International Conference on Data Mining Workshops (ICDMW)*, 2019, pp. 480–489.
- [27] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0098135409000076>
- [28] G. Hulthen, L. Spencer, and P. Domingos, "Mining time-changing data streams," in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 2001, pp. 97–106.
- [29] T. Le, F. Stahl, M. M. Gaber, J. B. Gomes, and G. Di Fatta, "On expressiveness and uncertainty awareness in rule-based classification for data streams," *Neurocomputing*, vol. 265, pp. 127–141, 2017.
- [30] J. Gama and P. Kosina, "Learning decision rules from data streams," in *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- [31] C. C. Aggarwal, S. Y. Philip, J. Han, and J. Wang, "A framework for clustering evolving data streams," in *Proceedings 2003 VLDB conference*. Elsevier, 2003, pp. 81–92.
- [32] F. Stahl, T. Le, A. Badii, and M. M. Gaber, "A frequent pattern conjunction heuristic for rule generation in data streams," *Information*, vol. 12, no. 1, 2021. [Online]. Available: <https://www.mdpi.com/2078-2489/12/1/24>
- [33] V. M. Souza, D. F. Silva, J. Gama, and G. E. Batista, "Data stream classification guided by clustering on nonstationary environments and extreme verification latency," in *Proceedings of the 2015 SIAM International Conference on Data Mining*. SIAM, 2015, pp. 873–881.
- [34] M. Kontaki, A. Gounaris, A. N. Papadopoulos, K. Tsihlias, and Y. Manolopoulos, "Efficient and flexible algorithms for monitoring distance-based outliers over data streams," *Information systems*, vol. 55, pp. 37–53, 2016.
- [35] F. Angiulli and C. Pizzuti, "Fast Outlier Detection in High Dimensional Spaces," in *Principles of Data Mining and Knowledge Discovery*, G. Goos, J. Hartmanis, J. van Leeuwen, J. G. Carbonell, J. Siekmann, T. Elomaa, H. Mannila, and H. Toivonen, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, vol. 2431, pp. 15–27, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/3-540-45681-3_2
- [36] S. Thudumu, P. Branch, J. Jin, and J. J. Singh, "A comprehensive survey of anomaly detection techniques for high dimensional big data," *Journal of Big Data*, vol. 7, no. 1, pp. 1–30, 2020.
- [37] P. D. Haghghi, S. Krishnaswamy, A. Zaslavsky, M. M. Gaber, A. Sinha, and B. Gillick, "Open mobile miner: A toolkit for building situation-aware data mining applications," *Journal of Organizational Computing and Electronic Commerce*, vol. 23, no. 3, pp. 224–248, 2013.
- [38] T. Zemichal and T. G. Dieterich, "Anomaly detection in the presence of missing values for weather data quality control," in *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies*, ser. COMPASS '19. New York, NY, USA: Association for Computing Machinery, Jul. 2019, pp. 65–73. [Online]. Available: <https://doi.org/10.1145/3314344.3332490>
- [39] O. Zielinski, D. Meier, K. Löhms, T. Balke, M. Kleyer, and H. Hillebrand, "Environmental conditions of a salt-marsh biodiversity experiment on the island of Spiekeroog (Germany)," *Earth System Science Data*, vol. 10, pp. 1843–1858, 2018.
- [40] J. Montiel, M. Halford, S. M. Mastelini, G. Bolmier, R. Sourty, R. Vaysse, A. Zouitine, H. M. Gomes, J. Read, T. Abdesslem, and A. Bifet, "River: machine learning for streaming data in python," 2020.
- [41] S. Oehmcke, O. Zielinski, and O. Kramer, "Event Detection in Marine Time Series Data," in *KI 2015: Advances in Artificial Intelligence*, S. Hölldobler, R. Peñaloza, and S. Rudolph, Eds. Cham: Springer International Publishing, 2015, vol. 9324, pp. 279–286, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-319-24489-1_24
- [42] Z. He, X. Xu, and S. Deng, "Discovering cluster-based local outliers," *Pattern Recognition Letters*, vol. 24, no. 9, pp. 1641–1650, Jun. 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167865503000035>
- [43] Z. Li, Y. Zhao, N. Botta, C. Ionescu, and X. Hu, "COPOD: Copula-Based Outlier Detection," in *2020 IEEE International Conference on Data Mining (ICDM)*, Nov. 2020, pp. 1118–1123, iSSN: 2374-8486.
- [44] M. Goldstein and A. Dengel, *Histogram-based Outlier Score (HBOS): A fast Unsupervised Anomaly Detection Algorithm*, Sep. 2012.
- [45] T. Pevný, "Loda: Lightweight on-line detector of anomalies," *Machine Learning*, vol. 102, no. 2, pp. 275–304, Feb. 2016. [Online]. Available: <https://doi.org/10.1007/s10994-015-5521-0>
- [46] C. C. Aggarwal, *Outlier Analysis*, 2nd ed. Springer Publishing Company, Incorporated, 2016.
- [47] A. Bifet and R. Gavaldà, "Learning from Time-Changing Data with Adaptive Windowing," in *Proceedings of the 2007 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics, Apr. 2007, pp. 443–448. [Online]. Available: <https://epubs.siam.org/doi/10.1137/1.9781611972771.42>
- [48] I. Frías-Blanco, J. d. Campo-Ávila, G. Ramos-Jiménez, R. Morales-Bueno, A. Ortiz-Díaz, and Y. Caballero-Mota, "Online and Non-Parametric Drift Detection Methods Based on Hoeffding's Bounds," *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 3, pp. 810–823, Mar. 2015, conference Name: IEEE Transactions on Knowledge and Data Engineering.
- [49] Deutscher Wetterdienst, "Die Witterung in den deutschen Küstengebieten," n.d. [Online]. Available: <https://www.dwd.de/DE/leistungen/witterungskueste/witterungskueste.html?nn=480164>
- [50] O. Zielinski, D. Meier, and H. Hillebrand, "Continuous meteorological observations at BEFmate automatic weather station, Spiekeroog, Germany, 2014-11 to 2017-04," Institute for Chemistry and Biology of the Marine Environment, Carl-von-Ossietzky University of Oldenburg, Germany, 2017. [Online]. Available: <https://doi.org/10.1594/PANGAEA.870988>
- [51] O. Zielinski, D. Meier, M. Kleyer, and H. Hillebrand, "Continuous meteorological observations at BEFmate automatic weather station, Spiekeroog, Germany, 2017-04 to 2018-12," Institute for Chemistry and Biology of the Marine Environment, Carl-von-Ossietzky University of Oldenburg, Germany, 2020. [Online]. Available: <https://doi.org/10.1594/PANGAEA.910715>
- [52] O. Zielinski, D. Meier, M. Kleyer, K. Löhms, and H. Hillebrand, "Continuous meteorological observations at DynaCom automatic weather station (2019), Spiekeroog, Germany, 2019-01 to 2019-09," Institute for Chemistry and Biology of the Marine Environment, Carl-von-Ossietzky University of Oldenburg, Germany, 2020. [Online]. Available: <https://doi.org/10.1594/PANGAEA.910752>
- [53] Y. Zhao, Z. Nasrullah, and Z. Li, "Pyod: A python toolbox for scalable outlier detection," *Journal of Machine Learning Research*, vol. 20, no. 96, pp. 1–7, 2019. [Online]. Available: <http://jmlr.org/papers/v20/19-011.html>
- [54] S. Oehmcke, O. Zielinski, and O. Kramer, "kNN ensembles with penalized DTW for multivariate time series imputation," in *2016 International Joint Conference on Neural Networks (IJCNN)*, Jul. 2016, pp. 2774–2781, iSSN: 2161-4407.
- [55] M. Bramer, *Principles of Data Mining*, 3rd ed., ser. Undergraduate Topics in Computer Science. London: Springer, 2013.