

Interactive comment on “Change-point detection in wind turbine SCADA data for robust condition monitoring with normal behaviour models” by Simon Letzgus

Simon Letzgus

simon.letzgus@tu-berlin.de

Received and published: 10 August 2020

We would like to sincerely thank the reviewer for taking the time to review our paper and for the constructive feedback. We are pleased that our work has been well received and the value of our contribution has been acknowledged. We found the review to be reasonable and there are no particular points we feel need to be rebutted. Indeed the reviewers questions and comments helped to provide an improved and more comprehensible presentation of our work. Please find the answers to the reviewer’s individual remarks in the supplement below.

[Printer-friendly version](#)

[Discussion paper](#)



Please also note the supplement to this comment:

<https://wes.copernicus.org/preprints/wes-2020-38/wes-2020-38-AC1-supplement.pdf>

Interactive comment on Wind Energ. Sci. Discuss., <https://doi.org/10.5194/wes-2020-38>, 2020.

WESD

Interactive
comment

Printer-friendly version

Discussion paper



Referee#1: Please find below the answers to the individual remarks ordered from general to specific.

Q1: A general point: although CPs are most likely due to some change in set points, maybe after maintenance, for example, many could be considered as genuine faults (a sensor drift or comms problem is a fault after all). Perhaps the authors could comment on how to differentiate between CPs which are related to faults and those which are not as they need to be treated differently in such as NBM.

A1: The distinction between CPs induced by faults and CPs caused by other effects is indeed difficult and without additional information, such as SCADA log-files or maintenance reports, speculative to some degree. In such a case, the only indication we see lays in the CP characteristics itself (as discussed in chapter 2.2). Firstly, changes in signal behaviour can be classified as being permanent or temporary. The latter ones consist of two CPs, where signal behaviour returns to its original pattern after a limited period (not longer than an interval of periodic inspections). For such changes, commonly, the first CP was caused by a malfunction, therefore being fault-related, which was consecutively corrected. The corrective action induced the second, non-fault related CP. A permanent change in signal behaviour, on the other hand, is not reverted. Permanent changes are more likely to be attributed to maintenance- and therefore non-fault-related. However, there still is the possibility of the change being induced by a fault which has not been discovered or was judged not to be severe enough to be fixed. Another distinction can be made between gradual and abrupt changes. Gradual changes can almost exclusively be attributed to be fault-related whereas abrupt changes could be either. Lastly, some physics of failure considerations might help to correctly identify the nature of an observed change. For temperature measurements, for instance, changes that manifest themselves in overall higher temperatures are more likely to be attributed to failures whereas changes leading to lower temperatures are more likely to be attributed to maintenance actions. For sensors like oil-pressure measurements, the exact opposite would be the case. Taking all these criteria together should enable the analyst to make an informed guess about the nature of the observed change, although some uncertainty remains. We think that the reviewer's question raises a point worth discussing and therefore incorporated this line of thought into the restructured section 2 (revised manuscript p. 4, line 9 ff. and p. 5, line 3 ff.).

With respect to the application of NBMs we would argue that the impact of fault and non-fault related CPs depends on the concrete question to be addressed. For the provision of clean training data, the main practical issue faced in real-world NBM application and therefore the focus of our contribution, both kinds of CPs have the same distorting effect on model training. Thus, any kind of CP violates the central assumption of NBMs and consequently has to be removed from the training data sets to ensure the method's feasibility.

Q2: Section 5.2: I am not sure I understand the analysis by CP – who exactly are the time series split? I assume there is one section of data with a CP within it, but where is the split made?

A2: Section 5 presents the performance evaluation of the algorithm. Performance is evaluated for the 600 selected signals, each covering two years of operation. Additionally, each of the signals is split exactly in the middle, resulting in 1200 signals, each covering one year of operation. With this split, a two-year signal, that contains only one CP indeed results in one signal with and one without a CP, each of length one year. In case multiple CPs are present, they might end up in either of the two shorter

signals, depending when they occurred (first half/second half). The evaluation for these 1200 shorter signals was conducted for three main reasons:

- i. To demonstrate the method's applicability to SCADA signals of different length. We think that this generalization property is an essential feature.
- ii. Previous work of the authors has shown that at least one year of SCADA data is required to train robust NBMs and which is in line with other publications explicitly recommending training data covering all four seasons (compare [1] and [2]). Having in mind the application of providing clean training data sets we think that demonstration on the one-year signals is closer to the application setting and therefore valuable.
- iii. The experiment showed that in many cases less dominant CPs could be successfully detected when a dominant CP was removed by the splitting procedure, which inspired the idea of an iterative CP removal (as discussed in section 5.3).

Thanks to the referee's remark we realised that the splitting procedure might not have been motivated adequately. This has been updated for the revised manuscript in section 2. There the split is initially discussed regarding its impact on the CP statistics. All three points mentioned above were explicitly incorporated (compare revised manuscript p. 6 lines 16 ff.).

Q3: Page 19, line 2: why would you want to remove such a trend? I assume this technique could be used to identify 'clean' sections of data which can then be analysed with fault detection algorithms. Taking out the trend would then be counter-productive

A3: In the given context the removal of overall signal trends is suggested only for signal pre-processing as part of the CP detection process. The data used for NBM training and application would still contain the trend but the training periods would be selected based on the outcome of the CP algorithm. The reasoning behind the trend-removal suggestion is that a steady trend which is present throughout the observed period does represent a shift in the signal's distribution but this steady shift itself is not changing and therefore should not be flagged as a CP.

This being said, the distinction between rising temperatures due to normal wear, which an NBM then would have to account for as 'normal' and an increased wear leading up to an early end of component life might be difficult. To our knowledge has not been addressed in literature so far and would be an interesting point for further research, since the presence of trends in the training data has been reported to be potentially indicative for slowly developing component problems (compare [3]).

Q4: Page 17, line 1: it is said that there are two CPs in 11a, but the figure shows only one shaded region.

A4: Shaded regions represent homogeneous periods with no change-points. True change-points are then indicated by the change in background colour. In figure 11a) the two true change points are in February (background colour changes from grey to red) and May (background colour changes from red to grey) of the second depicted year of operation. This way of visualizing the results was chosen to ensure both types of CPs, true and detected, are visible also in case of an exact detection where they overlap (compare the first CP in figure 11a). The figure captions have been updated to enhance clarity. '*Change in background colour indicates true CPs, dashed lines detection.*' was replaced by '*Each change in background colour indicates a true CP, each dashed line indicates a detected CP*'.

Q5: Page 14, line 14: should the statement ‘whereas splitting them might require detection of a less severe change in one half of the signals’ be something like ‘whereas splitting them might result in only the detection of a less severe change in one half of the signal time series.

A5: The sentence in question was changed to ensure comprehensibility: *‘Secondly, in case the two-year signal contains multiple CPs, detection of only the most significant one is enough for the signal to be evaluated as correctly classified (TP). When splitting this two-year signal into two one-year signals to analyse and evaluate them separately, detection of a less severe change in one of the signals might be required for both signals to be evaluated as correctly classified (both TP).’* (compare revised manuscript page 14, lines 19 ff.).

Q6: In general, ensure that all symbols in equations are properly explained.

i. **Q6.1:** Page 7, line 29: what is the bracketed (1) meant to represent?

A6.1: This was meant to be a reference to Equation (1) which was therefore corrected to *‘(compare Eq. (1))’*.

ii. **Q6.2:** Page 8, line4: what is P?

A6.2: P stands for the penalty term which acts as a regulariser for model complexity. For clarification, the missing reference was added inline as follows (bold): *‘Therefore, a regularisation term $P(\tau)$ was proposed for example by Lavielle (2005) which penalises for every additional CP and therefore reduces the complexity of the segmentation (compare Eq. (2))’*.

iii. **Q6.3:** Page 12, line 17: should the maximum std not be evaluated over k values rather than i?

A6.3: We agree with the reviewer and have updated the manuscript accordingly.

iv. **Q6.4:** In equation 10, what are D and T?

A6.4: T stands for the total number of time-steps the signal consists of, as defined in section 3.1. However, we agree that this should be stated again in proximity to Equation 10. The naming of D, which here stands for the actual number of segments, was named N in the earlier problem formulation in section 3.1 (compare Equation (1)). This inconsistency was corrected accordingly. Moreover, we noticed that the choice of D within this publication was not reported. The paragraph before equation 10 was therefore updated accordingly (compare revised manuscript page 13, lines 8 ff.).

Q7: Page 17, line 14: the reference to ‘signal drifts’ should be changed to ‘signal changes’. A drift suggests a problem with the sensor itself, whereas what is suggested is that the temperature change is genuine but just due to bearing wear.

A7: We agree with the reviewer and have updated the manuscript accordingly.

Q8: Some typos:

- i. Page 3, line 17: 'multiple'
- ii. Page 3, line 21: space between 'specifies' and 'the'
- iii. Page 4, line 19: better to be explicit in terms of 'oil pressure' (not just pressure which could be atmospheric pressure)
- iv. Page 7, line 25: 'let us' rather than 'let's' (avoid contractions in formal writing)
- v. Page 8, line 24: 'calculation'
- vi. Page 9, line 23: 'automated'
- vii. Page 10, line 8: 'occurrence'
- viii. Page 13, line 13: should be 'based on'
- ix. Page 13, line 22: 'reversely' should be 'conversely'
- x. Page 13, line 28: 'algorithm's'
- xi. Page 18, line 4: should be 'to ensure'
- xii. Page 18, line 8: should be 'able to'

A8: We agree with the reviewer and have updated the manuscript accordingly.

REFERENCES:

- [1] Bach-Andersen, M., Rømer-Odgaard, B., and Winther, O.: Flexible non-linear predictive models for large-scale wind turbine diagnostics, *Wind Energy*, 20, 753–764, 2017.
- [2] Letzgun S., Training data requirements for SCADA based condition monitoring using artificial neural networks, 2019, EAWE PhD Seminar, Nantes
- [3] Letzgun S, SCADA-based anomaly detection – challenges for automated application of artificial neural networks, 2018, EAWE PhD Seminar, Bruxelles

Interactive comment on “Change-point detection in wind turbine SCADA data for robust condition monitoring with normal behaviour models” by Simon Letzgus

Anonymous Referee #2

Received and published: 22 July 2020

The paper represents an valuable addition to the automatic data processing of SCADA data. By detecting the CPs reliably for further application of SCADA, such as monitoring and fault detections.

General question to the manuscript.

-In the manuscript it is mentioned that data 33 wind turbines from 3 different sites are analyzed in this paper. Do they represent the same turbine typology, i.e. geared versus direct drive, synchronous generator versus DFIG etc. What kind of site conditions they represent, complex terrain versus flat terrains. Age of the wind turbines? It is important in my opinion to discuss the representativeness of the data that are used in the analysis

[Printer-friendly version](#)

[Discussion paper](#)



as the Resulting method will be applied to different turbine types, ages, site conditions etc.

-the data analyzed here are temperature data according to the table B.1 What is the reason behind this choice. Are there vibration data from the wind turbine also available for the analysis. Will the algorithm change if other Type of sensors are analyzed, e.g. acceleration data.

-Change of operation modes. Do the algorithm consider changes in the operational state of the wind turbines? For example, down regulation of power due to grid demand, noise reduced operation due to noise regulation In the night with medium high wind speeds. These can look like CPs in the data possibly.

-minor comments and edits can be found in the attached PDF file

Please also note the supplement to this comment:

<https://wes.copernicus.org/preprints/wes-2020-38/wes-2020-38-RC2-supplement.pdf>

Interactive comment on Wind Energ. Sci. Discuss., <https://doi.org/10.5194/wes-2020-38>, 2020.

Printer-friendly version

Discussion paper



Referee#2: Please find below the answers to the individual remarks ordered from general to specific.

Q1: In the manuscript, it is mentioned that data 33 wind turbines from 3 different sites are analysed in this paper. Do they represent the same turbine typology, i.e. geared versus direct drive, synchronous generator versus DFIG etc.? What kind of site conditions they represent, complex terrain versus flat terrains? Age of the wind turbines? It is important in my opinion to discuss the representativeness of the data that are used in the analysis as the Resulting method will be applied to different turbine types, ages, site conditions etc... .

A1: The turbines are from different manufacturers and all of them are geared and equipped with DFIGs. All turbines were commissioned later than 2013 and the analysed periods fall within the first five years of operation. The sites can be characterised as moderately complex with mild elevation changes and occasional vegetation. This information was indeed missing and will be incorporated in section 2.1 of the revised manuscript accordingly (revised manuscript: p. 4, line 14 ff.).

Even though the turbines represent a rather homogeneous set we expect the method to perform equally well on temperature measurements along the drive train from turbines with different configurations where different sensors might be in place. This is due to the method's good performance over the wide range of different temperature signals as well as the different characteristics of the detected change-points. Neither do we expect the method's performance to decrease for older turbines or turbines in different site conditions. However, these characteristics might influence the presented cp statistics, with older turbines or turbines exposed to higher loads showing an increased amount of change-points due to increased wear and consecutive maintenance actions. This line of thought was added to section 2 as well: 'Even though these findings might vary across different turbine types, ages and site conditions the order of magnitude of CP presence highlights the necessity of a robust CPD methodology.' (compare revised manuscript p. 6, line 14 ff.).

Q2: The data analysed here are temperature data according to the table B.1 What is the reason behind this choice?

A2: In SCADA based monitoring of wind turbines using NBMs two approaches can be distinguished - performance and temperature monitoring. The former aims to detect abnormal deviations from the turbines usual power output, whereas the latter aims to detect deviations from the healthy thermal equilibrium conditions. Although both approaches have proven to be valuable (particularly in combination) temperature monitoring is better suited for detecting malfunctions in the components along the drive train, which account for the majority of turbine downtime (compare [1]). Moreover, the challenge of change-points in wind turbine SCADA data was mainly reported in the context of temperature monitoring in literature. Therefore, we decided to focus on temperature data. Nevertheless, the methods performance over a wide range of different temperature signals as well as over the different characteristics of the detected change-points suggests that the method can potentially be extended to other signals found in SCADA systems, a proposition that has been incorporated into the outlook section 6 of the revised manuscript (compare revised manuscript p.21, lines 11 ff.). Thanks to the referee's comment it also became clear, that neither the distinction between temperature and performance monitoring nor our motivation for variable selection were stated explicitly enough. Therefore, they incorporated into the introductory section (compare

updated manuscript p. 2, line 21) as well as the data set description of section 2 (compare revised manuscript p. 4, 18ff.).

Q3: Are there vibration data from the wind turbine also available for the analysis. Will the algorithm change if other types of sensors are analysed, e.g. acceleration data?

A3: Vibration/Acceleration data were not available for this study. Nevertheless, we assume that in principle the suggested kernel-based change-point detection algorithm should also be useful to analyse measurements from these kinds of sensors. [2] for example presents experimental results of kernel-based change-point detection being successfully applied to the segmentation of audio signals. In terms of structure and time resolution, audio signals are much closer to vibration data than the SCADA data analysed in this study. One particular challenge we see at this point is that the high data resolution could impose numerical challenges for computing the respective gram-matrix. In any case, the proposed data pre-processing method would need to be adjusted to the different types of data and could potentially help to overcome these problems. We think this is an interesting question that could be addressed in the future and therefore incorporated it into the outlook section 6 (compare revised manuscript p. 21, 10 ff.).

Q4: Change of operation modes. Does the algorithm consider changes in the operational state of the wind turbines? For example, downregulation of power due to grid demand, noise-reduced operation due to noise regulation in the night with medium/high windspeeds. These can look like CPs in the data possibly.

A4: The proposed framework considers changes in operational states of the wind turbine in two different ways. Firstly, the pre-processing procedure acts as a normalization which puts the measured temperature in relation to the operational state. Secondly, by averaging the signals over a full day, which was originally motivated by computational considerations, the impact of such presumable sub-day events is further reduced. This helps the algorithm to focus on the most significant and long-lasting changes and is part of the reason, why the pre-processing has such a crucial effect on the algorithm's performance (compare section 5.3).

Q5: Minor comments and edits can be found in the attached PDF file:

i. **Q5.1:** The pre-processing takes care of seasonal effects. What about diurnal effects?

A5.1: To reduce the numerical effort of computing the gram-matrix a daily averaging of the signals is part of the pre-processing procedure. This also removes all diurnal effects. Moreover, diurnal effects would be detected only with a penalty much lower than the proposed one, since the reduction in the cost function would need to compensate for as many change-points as days in the analysed period. Seasonality on the other hand induces 2 to 4 false CPs in each seasonal signal when not handled prior to the CP optimisation step and is therefore much more likely to be flagged by the algorithm.

ii. **Q5.2:** Page 4, line 4: 1Hz sampling is usually possible, the only problem is they are not being stored due to data storage reasons. The second reason you don't see them is because OEMs don't give access to wind farm operators.

A5.2: We agree with the reviewer and have updated the manuscript accordingly (compare revised manuscript p. 4, lines 3 ff.).

iii. **Q5.3:** Notes on spelling/grammar in the PDF-file.

A5.3: We agree with the reviewer and have updated the manuscript accordingly.

REFERENCES:

[1] Dao, C., Kazemtabrizi, B., Crabtree, C.: Wind turbine reliability data review and impacts on levelised cost of energy, *Wind Energy*, 22, 1848-1871, 2019

[2] Arlot, S., Celisse, A., and Harchaoui, Z.: Kernel change-point detection, arXiv preprint [arXiv:1202.3878](https://arxiv.org/abs/1202.3878), 2012.

List of relevant changes in the revised manuscript:

1. Abstract: Minor changes in formulations to keep the abstract more compact.

2. Incorporation of reviewer comments as described above in detail.

3. Section 2 was restructured from:

2. Structural breaks in SCADA data measurements

2.1 Data base and change-point annotation

2.2 Signal and change-point characterisation

to:

2. Change-points in wind turbine SCADA data

- update SCADA system specification (Referee2-Q5.2)

2.1 Data set and change-point annotation

- include background about turbines (Referee2-Q1)

- highlight motivation for signal selection (Referee2-Q2)

- paragraph about quantitative findings was moved to 2.3

2.2 Qualitative change-point evaluation

- incorporate discussion regarding fault-relation of cps (Referee1-Q1)

2.3 Quantitative change-point evaluation

- after including the above-mentioned updates, it seemed more coherent to separate the paragraph about quantitative findings of cp-presence, which initially was part of section 2.1, and append it as a new sub-section (quantitative CP discussion after the qualitative CP discussion).

4. Minor changes to ensure correct spelling, correct grammar, coherent equation numbering etc.

Change-point detection in wind turbine SCADA data for robust condition monitoring with normal behaviour models

Simon Letzgus

Technische Universität Berlin, Machine Learning Group; Straße des 17. Juni 135, 10623 Berlin, Germany

Correspondence: simon.letzgus@tu-berlin.de

Abstract. Analysis of data from wind turbine supervisory control and data acquisition (SCADA) systems has attracted considerable research interest in recent years. ~~The data is predominantly used to gain insights into~~ Its predominant application is to monitor turbine condition without the need for additional sensing equipment. Most ~~successful~~ approaches apply semi-supervised anomaly detection methods, also called normal ~~behaviour~~ behaviour models, that ~~use~~ require clean training data sets to establish healthy component baseline models. ~~However, one of the major challenges when working with wind turbine SCADA data in practice is~~ In practice, however, the presence of ~~systematic changes in signal behaviour~~ change-points induced by malfunctions or maintenance actions poses a major challenge. Even though this problem is well described in literature ~~it has not been systematically addressed so far. This~~ this contribution is the first to ~~comprehensively analyse the presence of change-points in wind turbine SCADA signals and introduce an algorithm for their automated detection~~ systematically evaluate and address the issue. 600 signals from 33 turbines are analysed over an operational period of more than two years. During this time ~~one third~~ one-third of the signals ~~are~~ were affected by change-points. ~~Kernel~~ which highlights the necessity of an automated detection method. Kernel-based change-point detection methods have shown promising results in similar settings ~~but their performance strongly depends on the choice of several hyperparameters. This contribution presents a comprehensive comparison between different kernels as well as kernel-bandwidth and regularisation-penalty selection heuristics. Moreover,~~ an appropriate. ~~We, therefore, introduce an appropriate SCADA data pre-processing procedure~~ is introduced to ensure their feasibility and conduct comprehensive comparisons across several hyperparameter choices. The results show that the combination of Laplace kernels with a newly introduced bandwidth and ~~penalty-regularisation-penalty~~ selection heuristic robustly outperforms existing methods. ~~In a signal validation setting more~~ More than 90% of the signals were classified correctly regarding the presence or absence of change-points, resulting in ~~a~~ an F1-score of 0.86. For ~~a~~ an automated change-point-free sequence selection, the most severe 60% of all CPs could be automatically removed with a precision of more than 0.96 and therefore without ~~a~~ any significant loss of training data. These results indicate that the algorithm can be a meaningful step towards automated SCADA data pre-processing which is key for ~~data-driven~~ data-driven methods to reach their full potential. The algorithm is open source and its implementation in Python publicly available.

1 Introduction

Wind energy plays a major role in the decarbonisation of energy systems around the world. It has developed into a mature technology over the past decades and its levelised cost of electricity (LCOE) has reached a ~~competitive~~ competitive level (IRENA (2019)). At the same time costs for operation and maintenance (O&M), which account for approximately ~~one-quarter~~ one-quarter of the LCOE, have seen only minor reductions (IRENA (2019)). An effective strategy to further reduce O&M costs is to switch from a scheduled maintenance scheme to condition-based maintenance. Under such a scheme maintenance decisions are based on information about the turbine's actual condition rather than on periodic inspections. The necessary information can be acquired through dedicated condition monitoring (CM) systems which can be for instance vibration-, oil- or acoustic emission-based (for a comprehensive review of state-of-the-art wind CM systems please refer to (Coronado and Fischer (2015))). On the other hand, each wind turbine is equipped with a variety of sensors in its supervisory control and data acquisition (SCADA) system. Utilisation of operational SCADA data for CM has attracted considerable research interest since it provides insights with no need for additional equipment. A wide range of methods have proven to be able to detect developing malfunctions at an early stage, often months before they resulted in costly component failures (see e.g. Zaher et al. (2009), Schlechtingen and Santos (2011), Bangalore et al. (2017), Bach-Andersen et al. (2017)). For a comprehensive review refer to (Tautz-Weinert and Watson (2016)). SCADA ~~data-based condition monitoring~~ therefore represents a cost efficient data-based condition monitoring, therefore, represents a cost-efficient and effective complement to ~~state-of-the-art~~ state-of-the-art CM-solutions. Its primary task is to classify the state of a turbine or one of its components as either healthy or faulty. However, the available SCADA data represents predominantly healthy operation with no or only comparatively few instances of faulty condition. In such a setting semi-supervised anomaly detection, often called normal behaviour ~~modeling~~ modelling, has proven to be useful (Chandola et al. (2009)). Normal behaviour models (NBMs) are trained on healthy ~~turbine~~ data to represent the class corresponding to the normal state. ~~Deviations~~ Subsequently, deviations between model output and the measured ~~SCADA~~ sensor values can be processed and evaluated to identify anomalies (compare Figure 1). For wind turbines, performance and temperature monitoring can be distinguished. The former aims to detect abnormal deviations from the turbine's usual power output, whereas the latter aims to detect deviations from the healthy thermal equilibrium conditions. We will focus on temperature monitoring which is better suited for detecting malfunctions in the components along the drive train, which account for the majority of turbine downtime (compare Dao et al. (2019)). Zaher et al. (2009) were among the first to apply the approach in the wind domain and prove its feasibility. Many publications with successful early detection of malfunctions followed (compare e.g. Butler et al. (2013), Kusiak and Verma (2012), Sun et al. (2016), Bangalore et al. (2017) and Bach-Andersen et al. (2017)).

Despite the promising NBM examples reported in literature scaling the method to large fleets of wind turbines comes with practical challenges. Leahy et al. (2019) analysed 12 studies that apply the concept of NBM to wind turbine SCADA data and found that all but one reported significant manual efforts in data pre-processing due to data quality and data ~~access-related~~ access-related issues. That is why researchers have developed different filtering methods ~~with the aim~~ to ensure healthy training data without traces of malfunctions. They can be divided into domain-knowledge-based-, alarm-based-, work-order-based-,




Figure 1. Scheme of normal behavior model-based anomaly detection with offline model preparation (left) and online application (right).

or statistical-approaches (Leahy et al. (2019)). Manual selection of representative operational patterns from the SCADA data sets would be an example of domain-knowledge-based filtering and can be found for instance in Zaher et al. (2009). Another common procedure is to filter NBM data against a certain threshold of active power production ~~in order~~ to exclude transitions between operational and non-operational states as well as corrupted sensor measurements during standstill (compare e.g. Sun et al. (2016), Bangalore et al. (2017), Tautz-Weinert (2018)). Schlechtingen and Santos (2011) were among the first to describe a more systematic semi-automated data pre-processing procedure. It consists of a domain-knowledge-based parameter range check, data scaling, handling of missing values and lag removal. These measures have been extended by multivariate statistical ~~filter-filtering~~ methods to automatically remove outliers (compare e.g. Bangalore et al. (2017)). However, a much more severe problem than missing, invalid or poorly processed data is caused by structural changes in sensor measurements (~~Tautz-Weinert and Watson (2017)~~). ~~Abrupt changes in the underlying data generating regime at a specific point in time (change-point) which~~ have been reported in different publications (e.g. Schlechtingen and Santos (2011) or Tautz-Weinert and Watson (2017)). They can be caused by sensor or component malfunctions as well as by maintenance actions. In an ideal setting, all potential causes would be ~~detected, corrected, fully documented and quickly detected and corrected with the corresponding information being~~ available to the respective data analyst. Unfortunately, this is rarely the case in practice (Tautz-Weinert and Watson (2017) and Leahy et al. (2019)) which has severe implications for NBMs. ~~NBM training represents statistical parameter estimation of an underlying process which can only be successful if training data is stochastically homogeneous. NBMs trained~~ ~~Trained~~ on data containing ~~change-points (CPs) abrupt changes in the underlying data-generating regime at a specific point in time (change-point), NBMs~~ are fit to multiple, potentially even ~~abnormal faulty~~, states of operation causing them to fail their intended task. Since ~~CPs change-points (CPs)~~ can make the NBM-approach infeasible in practice, this has been identified as the most serious issue for their application (Tautz-Weinert and Watson (2017)).

Based on the ~~above findings~~ ~~findings described above~~ this study aims to be the first ~~one~~ to conduct a systematic analysis regarding the presence of CPs in SCADA signals. Moreover, an approach for robust detection of structural changes in SCADA measurements will be suggested. Non-parametric kernel-based change-point detection (CPD) methods will be adapted to the problem at hand. This includes recommendations for the choice of respective hyperparameters and useful signal pre-processing steps based on evaluation across a large range of SCADA signals from ~~multiple multiple~~ wind farms. The result represents a step towards scalability of SCADA based NBM which is essential for the promising method to reach its full potential. The remainder

of this paper is organised as follows: Section 2 ~~gives an overview of CPs presence in the SCADA database with summary statistics and characteristic examples~~ presents the SCADA data used in this study and evaluates the presence and characteristics of CPs. Section 3 presents the method utilised in this study by formalising the CPD problem, introducing kernel-based CPD algorithms and their respective evaluation metrics. Section 4 ~~specifies the~~ specifies the CPD algorithm with its pre-processing steps and the selection of hyperparameters. Section 5 presents the performance over a range of hyperparameter configurations with respect to concerning different evaluation objectives followed by a discussion of results. Section 6 concludes with a summary and outlook.

2 Structural breaks Change-points in wind turbine SCADA data measurements


Wind turbine SCADA systems record measurements from sensors placed all over the turbine. Available signals usually include temperature measurements, electrical measures, pressure values, speed counters, timers, status parameters and environmental conditions. Modern SCADA systems often record more than 100 different signals ~~at sampling rates of 1 Hz. However, the~~ typical temporal resolution is 10-minute average values, although systems with a resolution as high as 1 Hz exist but are far less common. Additionally, some manufacturers store the signals' standard deviation as well as minimum and maximum values during the averaging period. Structural breaks in these measurements manifest themselves as an abrupt change in sensor behaviour at a specific time instant τ called a CP. The various potential causes can be classified into being available for analysis is 10-minute average values due to data storage limitations and access restrictions. Change-points in wind turbine SCADA signals can be induced by various causes. Generally, they can be sensor, component or maintenance related maintenance-related. Sensor related structural breaks can be are often caused by sensor drifts, sensor failures or malfunctions in the communication system. Component related breakpoints CPs can originate from particularly strong wear changes in component physics or component failure. Lastly, maintenance induced changes can be attributed to specific actions like the exchange of operating materials, replacement of components or sub-components and control adjustment. Although, the presence of CPs has been described in multiple publication as a challenge when working with wind turbine SCADA data (compare Tautz-Weinert and Watson (2017) and Leahy et al. (2019)), this study is the first to systematically evaluate While sensor and component related CPs can be considered as genuine faults, specific maintenance activities, such as changes in set-points, are another common cause. The following sections first describe the SCADA data used in this study, the signal selection and the CP annotation process. Subsequently, qualitative CP characteristics, their relation to potential causes as well as their implications for detection are discussed. Finally, the presence of CPs in wind turbine SCADA signals, to the author's knowledge the data sets is evaluated quantitatively.

2.1 Data base set and change-point annotation

For the current study SCADA data from 33 multi-MW turbines from different manufacturers were used. All turbines are equipped with gearboxes and double-fed induction generators and were commissioned later than 2013. They are located at 3 different sites was used of moderate complexity. For each turbine SCADA data representing more than 2 two full years

of continuous operation within the first 5 years after commissioning was present. Each turbine's SCADA system records between 30 and 100 signals in the typical 10-minute resolution. From the almost 2000 time series, 600 were selected for CPD based on the signal's potential for NBM temperature monitoring using NBMs. Therefore, all power train related temperature and oil pressure values were selected. Additionally, temperatures from the pitch system, the electrical system, and ambient conditions were chosen. The left pie-chart in Figure 2 shows the allocation of the 600 analysed signals to the respective components. ~~Next to the self-explanatory component classes the category 'Others' contains signals such as shaft bearing, nacelle and brake temperatures.~~ Generator and gearbox-related signals represent half of the overall selection. These ~~are also the components~~ components are also typically targeted by SCADA based NBMs for temperature monitoring (compare Tautz-Weinert and Watson (2016)). The high number of pitch-related signals is due to the availability of ~~temperature measurements from multiple sub-components on multiple sensors in~~ each blade's pitch ~~systems~~system. A full list of the analysed signals and their mapping to the respective components can be found in Appendix B1. ~~In addition~~ Next to the sensor data time series, SCADA-log files ~~,~~and information about major maintenance activities was present. ~~This information was~~ They were combined with a visual inspection of all analysed ~~time series signals~~ signals to manually annotate CPs. The raw signals, their de-trended and normalised transformations (compare section 4.1), ~~as well as their summary statistics (minimum, mean, median and maximum)~~ were compared with, were compared using different temporal resolutions. ~~Additionally, the time series were contextualised by comparing~~ The comparison of all signals related to the same component ~~. Often changes in one component were reflected in multiple component related signals at the same point in time. Such coherent findings helped to increase confidence during the annotation process. Additionally, each signal was compared to its~~ often lead to coherent findings in case of CP presence, which further increased confidence in the annotation. Moreover, signals were compared to their equivalent from at least 5 neighbouring turbines in the farm. This ~~so called~~so-called trending approach is well known in SCADA analysis for monitoring wind turbines (compare Tautz-Weinert and Watson (2016)) and helped to highlight the difference between normal signal behaviour and abrupt changes. ~~Moreover, after annotating each signal for the full two years, smaller segments were analysed to judge the signal's behaviour in a different temporal context. The result~~ The results of this tedious task ~~was~~ were reviewed by fellow researchers in order to secure utmost objectivity ~~. Nevertheless, not every annotated CP could be confirmed with the maintenance information or the SCADA logs which is attributed to the incomplete available information, a typical problem for this field of research (compare e.g. Tautz-Weinert and Watson (2017) and Leahy et al. (2019)). However, the described signal inspection procedure is assumed to~~ and reduce the number of ~~mis-annotations~~ false annotations to a minimum.

Figure 2 shows the results of the signal annotation process. The central chart represents a histogram over the number of CPs per signal. Exactly one third of the analysed signals were affected by changes over the approximately 2.5 year period. However, generally only a few CPs were found per signal. Actually, less than 5% of the affected signals exhibit 3 or more CPs. The right-hand diagram of Figure 2 compares the share of signals corrupted by changes for each component category. Gearbox related signals are most affected with more than half the signals containing CPs. For pitch related and ambient condition signals around one 30% of the time series were found to be affected. The high number of pitch related CPs were caused by systematic disturbances in the pitch motor temperature sensors for one of the wind farms. In case of ambient conditions a range of temperature sensors was found to be affected by severe drifts. The extent of CP presence highlights the necessity



Number of signals per component (left), number of CPs per signal (center), and share of signals with CPs per component (right) for the full 2-years time horizon.

Figure 2. Number of signals per component (left), number of CPs per signal (center), and share of signals with CPs per component (right) for the full two-year time horizon.

of a robust CPD methodology. The presented figures reflect the CP summary statistics across the selected signals for the full available period. In addition the 600 signals will be split into 1200 signals (covering approximately 1-year each) to analyse the algorithms ability to generalise to different signal lengths. This obviously changes the summary statistics which can be found in Appendix A1.

5 2.2 Signal and change-point characterisation

2.2 Qualitative change-point evaluation

Changes manifest themselves in a wide range of different signal behaviours, due to the multitude of potential reasons for structural changes, as well as the unique statistical natures of each signal. This is why a unifying framework to detect changes in SCADA measurements has to account for the diversity of signals and changes. Generally CPs can be classified into being abrupt or gradual as well as permanent or temporary. Structural changes in SCADA signals manifest themselves in a wide range of different signal behaviours. This is due to the multitude of potential causes in combination with the unique statistical nature of each signal. Often the cause of a change-point is closely related to how it manifests itself in the signal. Changes in signal behaviour can, for instance, be classified as permanent or temporary. Temporary changes consist of two CPs, where signal behaviour returns to its original pattern after a limited period (usually not longer than an interval of periodic inspections). For such changes, it is very common that the first CP was caused by a malfunction or fault which was consecutively fixed by a corrective maintenance action. A permanent change in signal behaviour, on the other hand, is not reverted and more likely to be attributed to a preventive maintenance action or control changes. However, there still is the possibility of a permanent change being induced by a fault which has either not been discovered or has not been considered to be severe enough to fix. Another distinction can be made between gradual and abrupt changes. Gradual changes can almost exclusively be attributed to be fault-related whereas abrupt changes could be either. Furthermore, some physics of failure considerations might explain the nature of an observed change. For temperature measurements, for instance, it is rather unlikely that a component failure




Figure 3. Exemplary SCADA signals exposing different structural changes. Each change in background color indicates a CP.

manifests itself in overall lower temperatures. For sensors like oil-pressure measurements, the exact opposite would be the case. Figure 3 shows three inherently different SCADA signals that exemplify different exemplary types of structural changes. In order to highlight the changes non-operational data was were excluded and the signals were normalised with their respective median to facilitate a comparison. Figure 3 a) shows a gearbox bearing temperature. Since the gearbox is an actively cooled component the signal has a comparably small, static range and a well-defined maximum. Therefore, the CP in February of the depicted year is easy to recognise. It occurred after a scheduled maintenance during which a cooling fluid was exchanged and the bearing consequently operates at clearly elevated temperatures. Such abrupt changes are mostly induced by maintenance actions or spontaneous sensor and component failures. Their detection and especially their exact localisation is usually easy in comparison to gradual changes due to their abruptness. Figure 3 b) displays a turbine's hydraulic oil temperature over a period of for two years. The signal is characterised by a comparatively large range and significant seasonality. A hydraulic fault in October of the second depicted year of operation causes the temperature to steadily rise compared to pre-CP conditions. The overlapping effects of seasonality and the gradual nature of the underlying change make the determination of the CP as well as its exact point in time challenging. This holds true even if additional information is available. A reported finding in a maintenance protocol, for instance, confirms the presence of a structural change but represents only an upper bound on the actual time of the event. Moreover, normal wear of components can also induce a steady gradual change and it can be difficult to tell normal wear apart from additional gradual effects. Here, the comparison of signals from neighboring turbines has proven to be helpful. Lastly, Figure 3 c) shows a gear oil pressure signal over a period of one year. The signal with high variance and a relatively well defined maximum shows a temporary decline of a turbine's gear oil inlet pressure and its return to the initial level. This was caused by an issue with the lubrication oil filter which was fixed during a scheduled maintenance activity. Most temporary changes are caused by corrective maintenance action related to an initial change. A unifying framework to detect changes in SCADA measurements has to account for this diversity of signals and changes.

When formalising the CP detection problem in the next chapter it becomes will become clear that the individual signal characteristics and their change in behaviour CP characteristics translate into how statistically distinct and therefore, how easy

to detect γ -a CP is. ~~This depends on the~~ Next to the decisive ratio between the magnitude of change and the individual signal variance or noise level (Garreau (2017)), ~~but also to the duration of how long a certain change is present. Intuitively, temporary changes or CPs at the very beginning or end of a time series are therefore more difficult to detect.~~ qualitative characteristics play a major role too. Permanent changes are for example easier to detect than short temporary ones. Also, abrupt changes are
5 generally easier to detect than slowly developing gradual changes, especially when it comes to exact temporal localisation.
~~Exemplary SCADA signals exposing different structural changes. CPs are indicated by a change in background color.~~

2.3 Quantitative CP evaluation

Figure 2 shows the results of the quantitative CP evaluation. The central chart represents a histogram over the number of CPs per signal. Exactly one-third of the analysed signals were affected by changes over the approximately 2.5 year period. Generally,
10 only a few CPs were found per signal. Less than 5 % of the affected signals exhibit 3 or more CPs. The right-hand diagram of Figure 2 compares the share of signals corrupted by changes for each component category. Gearbox related signals are most affected with more than half the signals containing CPs. For pitch-related and ambient condition signals around one 30% of the time series were found to be affected. The high number of pitch related CPs were caused by systematic disturbances in the pitch motor temperature sensors for one of the wind farms. In the case of ambient conditions, a range of temperature sensors
15 was found to be affected by severe drifts. Even though these findings might vary across different turbine types, ages and site conditions the order of magnitude of CP presence highlights the necessity of a robust CPD methodology. The presented figures reflect the CP summary statistics across the selected signals for the full period where data was available. Additionally, each of the 600 signals was exactly in the middle resulting in 1200 sub-signals each covering approximately one year of operation. A two-year signal that contains only one CP, therefore, results in one signal with and one signal without a CP. The respective
20 summary statistics of the one-year signals can be found in Appendix A1. The algorithm will be evaluated on both, the two-year as well as the one-year signals to ensure its generalisation abilities over different signal lengths. Moreover, a period of one year seems to be closer to the current practice of NBM training data selection which represents the algorithm's target application (compare Letzgus (2019)).

3 **Method for change-point detection**

25 The detection of CPs in time series is a ~~well-studied~~ well-studied problem in statistics, signal processing and machine learning. The goal is to detect time instants at which the underlying data generation process and therefore the marginal distribution of the observations changes abruptly. In other words, the time series is to be split into statistically homogeneous segments (Brodsky and Darkhovsky (1993)). First works date back to the 1950s (e.g. Page (1955)) but the topic has stayed the subject of active research until today, with methods being further refined and applied to many different domains, such as remote sensing (Touati
30 et al. (2019)), audio signal processing (Rybach et al. (2009)), or medical condition monitoring (Malladi et al. (2013)). Refer to Aminikhanghahi and Cook (2017) for an overview ~~on~~ of time series CPD methods. The following section will describe, classify and formalise the CP problem at hand based on Brodsky and Darkhovsky (1993).

3.1 Problem formulation

Conceptually, the CPD problem can be divided into online and offline detection. The former, sometimes also referred to as sequential CP detection, aims to identify changes in real-time settings as early and confidently as possible. In contrast, the latter, also known as signal segmentation, aims to determine the CP a posteriori with the data acquisition process being completed at the time that the homogeneity hypothesis is checked. Offline CP-problems can be further classified with respect to the a priori ~~knowledge~~ knowledge of the respective task. Complexity is significantly lower if the number of true CPs is known which reduces the task to the precise estimation of their location. In most real-world applications, however, the number of CPs itself has to be estimated. The same applies for a priori information about the statistical characteristics of the respective signals. Prior knowledge allows for assumptions regarding the family of underlying distributions. Therefore, CPs can be detected by identifying a change in the parameters describing the distribution. Non-parametric methods, on the other hand, require no such prior information which makes them more flexible and therefore often better suited for ~~real-world~~ real-world problems. The present task of ensuring CP-free training data sets represents an offline CPD problem, where the number of true CPs is unknown. Even though it is expected that many SCADA signals are not affected by structural changes ~~it is possible that~~ more than one statistically homogeneous segment ~~exist per signal~~ per signal may exist. Lastly, the SCADA data set consists of various statistically different signals which do not allow for unifying assumptions regarding their family of distributions. Therefore, non-parametric methods will be applied.

Let ~~s~~ s formalise the given problem under the prevailing conditions. We assume $X = \{X_1, X_2, \dots, X_T\}$ to be a piece-wise stationary time series signal in \mathbb{R}^d consisting of T observations. Piece-wise stationarity implies that X can be divided into N ($N \geq 1$) segments where each segment is well described by some distribution which might differ for consecutive segments. The segments therefore represent homogeneous sets s which are characterised by $N - 1$ CPs at some unknown instants in time $\tau_1^* < \tau_2^* < \dots < \tau_{N-1}^*$ (compare Eq. (1)). Now, CP detection can be formulated as a model selection problem where the CPs τ are the model parameters to be estimated. This can be achieved by defining a cost function $C(\tau)$ that quantifies intra-segment dissimilarity with respect to the ~~true CPs τ^*~~ chosen CPs τ (compare Eq. (2)). A naive minimisation of this cost function would result in a segmentation into N segments of unit size. Therefore, a regularisation term $\mathcal{P}(\tau)$ was proposed for example by Lavielle (2005) which penalises for every additional CP and therefore reduces complexity of the segmentation (compare Eq. (2)).

$$s = \{s_1, s_2, \dots, s_N\} = \{\{X_0, \dots, X_{\tau_1}\}, \{X_{\tau_1+1}, \dots, X_{\tau_2}\}, \dots, \{X_{\tau_{N-1}}, \dots, X_T\}\} \quad (1)$$

$$\hat{\tau} \in \underset{\tau}{\operatorname{argmin}} C(\tau) + \mathcal{P}(\tau) \quad \text{where} \quad C(\tau) = \sum_{n=1}^N C(s_n) \quad (2)$$

Since the complexity of the optimisation problem grows quadratic with the number of data points a naive approach for minimising the cost function $C(\tau)$ can be computationally expensive. Several approximate search methods like a sliding

window or binary segmentation were developed (compare Truong et al. (2020)). They come with benefits regarding computing time but naturally compromise on precision. The optimal solution can still be obtained efficiently by applying an algorithm based on dynamic programming. It was originally introduced in 1958 (Bellman (1958)) for solving a shortest-path problem for traffic networks. Since then the algorithm has been developed further (see e.g. Guédon (2013)) and was successfully applied in the context of CPD. The method utilises the additive structure of the cost objective to recursively compute optimal CPs for multiple sub-signals among which the global minimum is then selected. An implementation of the algorithm is publicly available as part of the CP detection library `ruptures` in Python (Truong et al. (2020)) and was utilised within this study.

3.2 Kernel based change-point detection

Equation 2(2) represents a general cost-function for solving the signal segmentation task at hand but the result heavily depends on an appropriate measure for the intra-segment similarity. Harchaoui and Cappé (2007) proposed a kernel-based approach which does not rely on parametric assumptions but ~~is able to~~ can detect changes in the high order moments of the signal distribution. Kernel methods use mapping functions $\Phi : \mathbb{R}^d \rightarrow \mathcal{H}$ to implicitly project a signal into a potentially much higher dimensional Reproducing Kernel Hilbert Space (Scholkopf and Smola (2002)). With the ~~well known~~ ~~well known~~ kernel-trick, the distance or similarity of two data points in the high dimensional feature space can be calculated by directly applying the kernel function (compare Eq. (3)). Harchaoui and Cappé (2007) used this property to evaluate the adequacy of τ . They define a kernel least-squares criterion that measures the intra-segment scatter (see Eq. 4(4)). Intuitively, the second term of Equation 4(4) increases if the chosen segments are more similar to each other and in return maximises dissimilarity between segments due to the negative sign. Note that the intra-segment scatter requires the ~~calculaten~~ ~~calculation~~ of the kernel-gram matrix $G_{i,j} = K(X_i, X_j)$, which implies a quadratic computational complexity and therefore restricts the method regarding the size of the data sets. By minimising the criterion the best segmentation for a known number of CPs can be obtained. Conceptually, any positive semi-definite kernel can be applied in this framework. Popular candidates are the Linear (*i*), Laplacian (*ii*) or Gaussian (*iii*) kernel (compare Eq. 5(5)). Note, that Laplacian and Gaussian kernels need the selection of an appropriate bandwidth parameter h . Arlot et al. (2016) expanded the method to an unknown number of CPs by applying the concept of penalising for additional CPs (compare Eq. 2(2)). Since then the ~~kernel based~~ ~~kernel based~~ algorithm has been successfully applied to multiple real-world time series CPD problems (compare Arlot et al. (2016)).

$$k(x, x') = \langle \Phi(x), \Phi(x') \rangle_{\mathcal{H}} \quad \text{and} \quad k(x, x) = \|\Phi(x)\|_{\mathcal{H}}^2 \quad (3)$$

$$C(\tau) = \frac{1}{T} \sum_{t=1}^T k(X_t, X_t) - \frac{1}{T} \sum_{n=1}^N \left[\frac{1}{\tau_n - \tau_{n-1}} \sum_{i=\tau_{n-1}+1}^{\tau_n} \sum_{j=\tau_{n-1}+1}^{\tau_n} k(X_i, X_j) \right], \tau_0 = 0 \quad (4)$$

$$i) k_{lin}(x, y) = \langle x, y \rangle \quad ii) k_{lp}(x, y) = \exp\left(\frac{-\|x - y\|}{h}\right) \quad iii) k_{rbf}(x, y) = \exp\left(\frac{-\|x - y\|^2}{h}\right) \quad (5)$$

3.3 Performance evaluation

The performance of ~~the CPD algorithm~~ CPD algorithms can be evaluated using the classic notation of true positives (tpTP), false positives (fpFP), true negatives (tnTN) and false negatives (fnFN). ~~In order to~~ FN. To appropriately interpret the evaluation results the implications of false classifications have to be considered. In case of NBMs ~~a fn~~, a FN translates into a risk for model quality, a fpFP into loss of potentially valuable training data. However, the individual impact depends on the severity of the change, meaning the more distinct degree of distinction from normal signal behaviour as well as the longer its presence, the more severe its consequences duration of its presence. This goes well with the concept of the presented CPD algorithm since the notion of severity directly translates into a cost reduction by segmentation. For the concrete evaluation of CPD results two different evaluation objectives are distinguished:

1. Automatic training data validation: detect presence of CP in a given processed signal
2. Automatic training sequence selection: detect the number and exact locations of CPs in the processed signal

Automatic training data validation answers whether a CP is present in a given SCADA signal or not. Therefore, the CP-detection result is evaluated once for each signal. In case the algorithm indicates one or more CPs for a signal containing at least one CP the result is evaluated as a true positive. No CP detection in a CP-free signal represents a true negative, etc. In practice, this means that a CP detection result is evaluated as a true positive, even if the number and location of indicated CPs does not necessarily represent the ground truth. Nevertheless, this can be ~~a~~ useful information for validating signals against the presence of CPs. Especially, since the alternative is a full manual inspection of all signals. Automatic training data validation ~~therefore~~, therefore, pre-selects the signals for visual inspection, in which the actual locations of the CPs are subsequently determined. The next step towards a fully automated NBM approach is ~~auotmated~~ automated training sequence selection. This requires a more precise evaluation for each CP and each CP-indication individually. Therefore, an acceptable margin is selected around each true CP in which a detection is evaluated as a tpTP. While for automatic training data validation this margin was practically set to infinity a fixed number of days has to be chosen for automatic training sequence selection. A CP which is present in the signal but not indicated by the algorithm is evaluated as a fnFN. Detection outside the margin boundaries represents a fpFP. ~~A tn~~ FP. A TN represents a CP-free signal with no detection. The concept is visualised in Figure 4 where a tp, fp and a fn TP, FP and a FN are depicted. Note, that a CP indicated just outside the margin already leads to a "double" punishment by evaluating the indication as a fpFP and the true CP as not detected (fnFN). Intuitively, the overall detection result depends on the selected acceptable margin. Moreover, the margin corresponds to the amount of data around the detected CP to be automatically cut off in automated training data selection and therefore to a trade-off between data loss and accuracy. In this paper, the acceptable margin was selected to be +/- 60 days around the true CP. The choice is motivated by the fact that missing a dominant true CP can be much more critical for many applications, such as NBM for example, than a ~~the~~ reduction of training data by the given period. Moreover, this attributes to potential inaccuracy during the manual annotation process due to uncertainties about the actual occurrence in case of uncertainty about the exact temporal localisation




Figure 4. Exemplary evaluation of a CP-detection result with one ~~fn~~FN (left), ~~fp~~FP (center) and ~~tp~~TP (right).

of the change (compare section 2.1). With the classification of each detected CP into true or false positive or ~~negative, the well-known~~ ~~negative, the well-known~~ evaluation metrics Accuracy, Precision, Recall and the F1-Score can be calculated:

$$accuracy = \frac{\#tp + \#tn}{N} \frac{\#TP + \#TN}{\#TP + \#TN + \#FP + \#FN} \quad precision = \frac{\#tp}{\#tp + \#fp} \frac{\#TP}{\#TP + \#FP} \quad recall = \frac{\#tp}{\#tp + \#fn} \quad (6)$$

$$f1 - score = 2 * \frac{precision * recall}{precision + recall} \quad (7)$$

5 4 Algorithm for change-point detection

This section describes the detailed steps of signal processing applied to detect the CPs in this study. Signal-pre-processing, as well as the choice of ~~hyper-parameters~~ ~~hyperparameters~~, are discussed.

4.1 Data pre-processing

Wind turbine operation is highly volatile due to intermittent ambient conditions. This is reflected in the high variance of raw SCADA measurements and complicates CP detection because the change in signal behaviour might be small in relation to regular signal behaviour. Signal pre-processing methods can help to reduce the signal to its most valuable components for CP detection and therefore facilitate the process. Tautz-Weinert and Watson (2017) suggest the comparison of monthly maximums and percentiles to detect structural changes. This was found to be too granular to attribute for temporary changes of less than one month. Additionally, such an approach does not reduce seasonality which was found to be an important factor for successful kernel-CPD. Instead, the following pre-processing steps were taken in this study:




Figure 5. Exemplary SCADA signals after pre-processing (compare Figure 3). True CPs are indicated by a Each change in background color indicates a CP.

- Removal of non-operational periods
- Normalisation with operational state and ambient conditions
- Re-sampling with reduced temporal resolution

The removal of non-operational periods is a routine pre-processing step in SCADA data analysis (compare e.g. Sun et al. (2016), Bangalore et al. (2017), Tautz-Weinert (2018)) and was motivated by the reasoning that changes in operational conditions will become most apparent when the turbine is in operation. Also, it was observed that sensor values during non-operational periods, e.g. during maintenance, sometimes take pre-defined standard values. In order to To exclude such distorting effects all data points where the turbine is operating on less than 10% of its rated power were excluded. Furthermore, the signal measurements were normalised with respect to based on the prevailing operating conditions. Active power production and rotor rotational speed were found to be the most dominant to characterise the turbine’s operational state. Ambient temperature was identified to be a good regressor to exclude seasonality from the sensor measurements. Therefore, each signal was normalised using these three input variables in a linear regression (compare Eq. 8(8)). The model was found to adequately subtract the influences of external conditions, is computationally cheap and due to its simplicity does not allow overfitting the sensor signals. In a last step, the normalised signal was averaged over each day, if at least three hours of operational data were available. This allows to extract the normalised signal characteristics and additionally reduces the amount of data points which facilitates the computation of the kernel gram-matrix. An exemplary result of the pre-processing procedure is shown in Figure 5. It displays the three signals shown earlier (compare Figure 3), after pre-processing. Note in particular how the method facilitates a clear identification of the CP in the hydraulic oil temperature compared to the raw signal.

$$signal^* = f(X_P, X_{rpm}, X_{T_{amb}}) = w_1 \cdot X_P + w_2 \cdot X_{rpm} + w_3 \cdot X_{T_{amb}} + c \quad (8)$$

20 A minor disadvantage of this approach is that the regressing signals, meaning active power, rotor speed and ambient temperature cannot be pre-processed the same way. In fact, abrupt change in one of the regressors can induce a CP in highly correlated

signals. In this study, such a case occurred a few times when the ambient temperature sensor of a turbine was corrupted. However, by regressing the inputs themselves with signals from neighbouring turbines first and then running the algorithm on them to check for CPs can exclude those cases. Alternatively, a simple rule checking for simultaneous CP detections in all signals particularly correlated to the same regressor can do the trick as well.

5 4.2 Choice of hyperparameters

The CPD method described in section 3 requires adequate selection of several hyperparameters. Namely, the type of kernel, its respective bandwidth and the penalty term for additional CPs. All three were found to have a profound impact on the CPD performance. Choosing an appropriate kernel is a ~~well-studied~~ well-studied problem by itself in many applications. In the context of CPD the widely used Linear, Gaussian and Laplacian kernels have been used (Garreau (2017)). Therefore, all three will be compared within this study. The choice of an adequate bandwidth h is another problem often encountered when working with kernel methods. Looking at the definition of the Gaussian and the Laplacian kernels (see Eq. (5)) it becomes clear that a bandwidth chosen too large or too small will make the entries of the gram matrix go towards zero or one respectively and therefore valuable information will be lost. A common approach is therefore to choose the bandwidth in the range of the calculated distances. Gretton et al. (2012) for example suggest a median heuristic in the context of a kernel two-sample test (compare ~~Equation 9~~ Eq. (9)). This heuristic is heavily used in ML literature (Garreau (2017)) and is also applied in the CPD settings (Truong et al. (2020)). Arlot et al. (2016) on the other hand suggest to use the empirical standard deviation of the signal itself as the bandwidth. Both choices of bandwidths are tested and compared in this paper. Furthermore, it is argued here that estimation of an appropriate bandwidth based on a signal with abruptly changing properties might lead to a non-optimal choice. Therefore, a third approach is being introduced and tested where the signal is divided into k different segments $S_{bw} = \{\{X_0 \dots X_t\}, \dots, \{X_{(k-1)*t+1} \dots X_{k*t}\}\}$ of equal length t and the empirical standard deviation is calculated for each segment. The bandwidth is consequently chosen as the maximum of the ~~$i+k$~~ $i+k$ standard deviations (compare ~~Equation 9~~ Eq. (9)). In this study, k is selected to be 20. Consequently, each segment consists of roughly ~~2 months~~ 3 to 5 weeks of operational data. For the remainder of the paper ~~approach~~ approach c is ~~referred~~ referred to as batch-std bandwidth.

$$a) \quad h = \text{median}(\|X_i - X_j\|^n) \quad b) \quad h = \text{std}(X) \quad c) \quad h = \max(\text{std}(S_{bw})) \quad (9)$$

Another crucial hyperparameter choice is the selection of an appropriate penalty term (compare ~~Equation 2~~ Eq. (2)) which controls the number of CPs to be detected by the algorithm. If the penalty is selected too low, too many CPs will be detected and vice versa. A ~~data driven~~ data-driven approach for choosing the penalty in the context of minimisation of a penalised ~~ereriterion~~ criterion is the so-called slope-heuristic (Birgé and Massart (2007)). It was shown that the optimal penalty to avoid overfitting is approximately proportional to a minimal penalty which can be obtained based on a regression between the penalised quantity and the associated cost function without penalisation. In the context of CP detection this was firstly ~~applied by Lebarbier (2002)~~ and described by Lebarbier (2002), further refined by Baudry et al. (2012) and applied e.g. by Arlot et al. (2016). They suggest a minimal penalty based on two constants s_1 and s_2 which are obtained by ~~a regression between~~ regressing the cost function

$C(\tau)$ and $\log\left(\frac{T-1}{D_\tau-1}\right) \cdot C(\tau_n)$ against $\log\left(\frac{T-1}{n-1}\right)$ as well as $\frac{D}{T}$ for $D \in [0.6 \cdot D_{max}, D_{max}]$ (compare 10). $\frac{n}{T}$ corresponds to the total number of data points and $n \in [0.6 \cdot N_{max}, N_{max}]$, where N_{max} is an estimation of the maximal number of segments N (compare Eq. (1)). Based on our findings from 2, we chose $N_{max} = 6$. Finally, the ~~minimal~~ minimum penalty is multiplied with the factor α to obtain the final optimal penalty (compare Eq.(10)). Even though the optimal choice of α is problem specific

5 $\alpha = 2$ was reported as a suitable choice by Arlot et al. (2016).

$$pen_{opt-slope}(\tau_n) = \alpha_{slope} \cdot \frac{1}{T} \cdot (-s_1 \cdot \log\left(\frac{T-1}{n-1}\right) - s_2 \frac{D_\tau n}{T}) \quad (10)$$

In this study, the slope heuristic is compared to a simpler approach chosen based on the following consideration: signals which are inherently similar to themselves are by default characterised by a relatively low initial cost value and vice versa. This means that each CP by default leads to a larger cost reduction for more dissimilar signals. Therefore, the penalty term is chosen based ~~to the on~~ the sum of costs without any CP. Figure 6 supports this reasoning. Here, a CP was enforced on all signals without changes. The resulting reduction in the cost function is shown over the initial average cost (left). An approximately quadratic relation between cost reduction and initial average cost can be observed. The right side of Figure 6 shows the relative reduction normalised with $C(\tau = 0)^2$. Consequently, the normalised cost reductions are ~~distributed~~ much more uniformly ~~distributed~~ and facilitate the selection of a single penalty value over all signals. Moreover, the penalty term can now be easily calculated ~~form from~~ the signal characteristic itself. This is considered an advantage over the more complex methods found in literature. The findings indicate that a reasonable choice of the penalty factor α_{cost} would be in the range between 75 and 150 for a Laplace kernel with a bandwidth selected according to the batch-std heuristic. A penalty factor larger than 200 can be considered a conservative choice with only a few false positives. ~~Reversely the reduction~~ ~~Conversely, the reduction of the cost function~~ induced by a CP can be interpreted as a confidence measure. Note that these values depend on the kernel configuration.

$$20 \quad pen_{opt-cost}(\tau) = \alpha_{cost} \cdot C(\tau = 0)^2 \quad (11)$$

5 Results for change-point detection

The result section presents the algorithm's performance on automatic training data validation and selection. For both evaluation objectives, different hyperparameter configurations in terms of kernel-, bandwidth- and penalty-selection are compared. Additionally, the effect of signal length is investigated ~~in order~~ to ensure the algorithms generalisation abilities. Results are analysed on a cumulative as well as on a component level. Finally, the results are discussed and implications for the algorithm's practical application derived.




Figure 6. Absolute (left) and normalised (right) cost reduction of healthy signals by imposing one CP versus average cost.

5.1 Results-automatic-Evaluation of automated training data validation

In this section, the algorithm's ability to distinguish between signals with and without CPs is evaluated. Figure 7 shows the results achieved by different configurations on the full two-year signal length (left) and the half one-year signal length (right); representing approximately 2 and 1 years of operation. Both, F1- and accuracy-scores are compared for different kernels, bandwidth choices and penalty selection schemes. A clear ranking can be identified when comparing kernels. For configurations with cost-based penalties, Laplacian kernels perform best followed by Gaussian kernels. Linear kernels perform much worse. For penalties chosen according to the slope-heuristic, the contrary is the case. Linear kernels perform best, closely followed by Laplace and Gaussian configurations. In terms of bandwidth selection, the intra-kernel ranking differs, but the leading Laplacian configurations use the batch-std heuristic. It clearly which outperforms established standard deviation or median heuristics. When comparing penalty selection schemes the cost-based penalty estimation suggested in this paper clearly outperforms performs better than the slope-heuristic. All discussed qualitative observations hold for both time horizons. However, a clear performance loss in terms of F1-score for the shorter signals can be observed. This is attributed to the design of the evaluation scheme. Firstly, there is a shift in the distribution between affected and not affected signals (compare Figures 2 and A1). Secondly, in case of multiple CPs it is enough to flag the two-year signal contains multiple CPs, detection of only the most significant one to be classified as a true positive in the two-year signal, whereas splitting the might require is enough for the signal to be evaluated as correctly classified (TP). When splitting this two-year signal into two one-year signals to analyse and evaluate them separately, detection of a less severe change in one half of the signals might be required for both signals to be evaluated as correctly classified (both TP).

In absolute terms the overall best performing configurations are able to, the overall best-performing configurations can classify more than 90% of the signals correctly regarding the presence or absence of CPs. The wrongly classified signals are approximately one-quarter one-quarter false positives and three-quarters three-quarters false negatives. This translates into F1-scores of 0.87 and 0.76 for the different signal lengths. This performance is reached for a penalty factor of $\alpha_{cost_2} = 145$ and $\alpha_{cost_1} = 80$ respectively. The best result using the slope-heuristic for penalty selection was achieved for penalty factors of $\alpha_{slope} = 12.5$ which is much higher than the $\alpha_{slope} = 2$ suggested in literature. At the same time, F1-scores range around 20%




Figure 7. Validation F1-score and accuracy for different hyperparameter configurations on 2-year two-year and 1-year one-year signals

behind the leading cost-penalty based configuration. Table 1 displays the overall results as well as the results per component for the best performing CPD configuration on the 2-year two-year signals in detail. The algorithm reaches high performance across components. Ambient condition signals, as well as gearbox and pitch related signals, were classified with particularly high accuracy. Correctly classified CPs are often characterised by sharp transitions between the states and a significant difference in relation to the regular signal noise. Examples for of successful detections can be found in Figure 8. The top chart shows the correctly classified change in gear oil pressure after maintenance. The second chart shows the correctly identified shift and its return to normal behaviour in a nacelle temperature signal which was induced by problems in the generator cooling system and its consecutive fix. At the same time, there is a relatively high number of false positives for gearbox related signals. This is mostly caused by bearing temperatures that are gradually rising due to normal wear (compare Figure 9 a)). The algorithm detects the drift in the signal's distribution which, under the given evaluation framework, represents a false positive. In the broader context of NBM, this information is still valuable since it highlights the need for periodical model re-training. False negatives were mostly caused by short temporary changes which were not pronounced enough to compensate for the penalty of two CPs, which would be required to flag them correctly. An example is shown in Figure 9 b) which depicts a generator bearing temperature with temporary high temperatures. The detailed results for all configuration by penalty and time horizon can be found in Appendix C.

5.2 Evaluation of automated training sequence selection

In this section, the algorithm's ability to automatically select periods without CPs for each signal is analysed. Therefore, performance is evaluated for each CP individually rather than for each signal (compare section 3.3). Figure 10 shows the CPD results analogously to the results for automated training validation in the previous section. Qualitatively, the findings with respect to concerning kernel selection and configuration are equivalent. Laplace kernels with batch-std bandwidths perform

Table 1. Automated signal-training data validation results per component for best configuration (Laplace / batch-std / $\alpha_{cost} = 145$) on 2-year-signalstwo-year signals

Component	<u>tn-TN</u>	<u>tp-TP</u>	<u>fn-FN</u>	<u>fp-FP</u>	accuracy	precision	recall	f1-score
Gearbox	40	57	6	8	0.874	0.877	0.905	0.891
Generator	162	21	11	1	0.938	0.955	0.656	0.778
Pitch system	98	38	13	1	0.907	0.974	0.745	0.844
Electrical system	24	4	2	0	0.933	1.0	0.667	0.8
Ambient conditions	30	16	0	0	1.0	1.0	1.0	1.0
Others	50	11	3	2	0.924	0.846	0.786	0.815
Total	404	147	35	12	0.91	0.925	0.808	0.862




Figure 8. Examples of successful CP detection in gear oil pressure (a) and nacelle temperature (b). Change-Each change in background colour-color indicates real-a true CP, each dashed lines-detectionsline indicates a detected CP.

best. In comparison with the results from the previous section, the more difficult evaluation objective manifests itself in overall lower performance scores. Even though accuracies reach well above 80%, F1-scores drop to 0.73 and 0.71 for the two time horizonstime-horizons. However, the performance between the two analysed time horizons are very similar which attributes for the algorithm's ability to generalise across different signal length. The optimal penalty factors remain time horizon specific but stable across evaluation metric with $\alpha_{cost_2} = 150$ and $\alpha_{cost_1} = 80$. The general advantage of cost-basedcost-based penalties is preserved with F1-scores approximately 15% above the best slope-heuristics results which are achieved at $\alpha_{slope_2} = 11.5$ and $\alpha_{slope_1} = 4$.

Table 2 displays the overall results as well as the results per component for the best performing CPD configuration on the 2-yeartwo-year signals in detail. To explain the drop in F1-score the different false classification of each component were analysed. Gearboxes show both, a relatively high number of fns-FNs as well as fps. Approximately 50% of the fns-FNs




Figure 9. Examples of misclassification in gearbox (a) and generator (b) bearing temperature. Change-Each change in background color indicates a true CPsCP, each dashed lines-detection line indicates a detected CP.

can be attributed to the coexistence of large and comparatively small changes in the same signal. An example is shown in Figure 11 a) where an oil temperature signal undergoes two significant changes with the second one not being detected. The initial dissimilarity, based on which the penalty is calculated, is dominated by the first change and therefore detection of the second change cannot compensate for the high penalty value. Another 20% of the gearbox-related fns-FNs are caused by short temporary changes and further 20% by detections outside the 60 days margin (compare e.g. the gearbox bearing temperature in Figure 11 b)). These represent at the same time approximately 25% of the fps. However, the majority of gearbox related fps is caused by the described signal drifts due to normal wear in gearbox bearings (compare Figure 9 a)). At the same time around two-thirds-two-thirds of CPs are correctly detected in gearbox related signals, often representing major changes such as a drop in gear-oil pressure after a maintenance (compare Figure 8 a)). For the generator related signals the main cause of fns-FNs are relatively short temporary changes, such as the temporary high temperatures in a generator bearing displayed in 9 b). The same reason causes the majority of fns-FNs in pitch related signals. -In fact, (35 out of the 40 false negatives were of the kind shown in Figure 11 c)). These shifts in pitch motor winding temperature signals were caused by systematic communication problems. At the same time, many shifts were distinct enough to be detected which explains the high number of tps-TPs in pitch related signals. It can be summarised that signal drifts-changes due to normal wear, temporary changes and the coexistence of CPs with different significance levels represent challenges which have to be addressed in the future. Nevertheless, the algorithm gives reasonable results and was able to identify the majority of CPs present in the signals. The detailed results per component of all configurations by penalty and time horizon can be found in Appendix C.

5.3 Discussion of pre-processing, results, and application

From the presented results it can be concluded that Laplacian kernels in combination with bandwidths chosen based on the batch-std heuristic are best suited for the problem at hand. This configuration in combination with cost-based penalties clearly outperformed all other configurations. Analysis showed that correctly classified CPs are often characterized by a permanent




Figure 10. Selection F1-score and accuracy for different hyperparameter configurations on 2-year two-year and 1-year signals

Table 2. Automated signal-training data selection results per component for best configuration (Laplace / std-max / $\alpha_{cost} = 150$) on 2-year two-year signals.

Component	<u>tn-TN</u>	<u>tp-TP</u>	<u>fn-FN</u>	<u>fp-FP</u>	accuracy	precision	recall	f1-score
Gearbox	40	60	35	18	0.65	0.769	0.632	0.694
Generator	162	22	24	7	0.86	0.759	0.478	0.587
Pitch system	98	82	39	8	0.79	0.911	0.678	0.777
Electrical system	24	3	6	1	0.79	0.75	0.333	0.462
Ambient conditions	30	16	0	0	1	1.0	1.0	1.0
Others	50	12	6	3	0.87	0.8	0.667	0.727
Total	404	195	110	37	0.8	0.841	0.639	0.726

nature, sharp transitions between states and a significant difference in relation to the regular signal noise. The latter two qualities are particularly amplified by the pre-processing procedure (compare section 4.1). In-order-to-To demonstrate its importance the algorithm was run on the database with only a minimum of signal pre-processing, namely a daily averaging of the measurements in-order-to ensure computational feasibility. Results show a drop in F1-scores from 0.83 to 0.6 for validation of the 2-year two-year signals and an even more dramatic decline from 0.73 to 0.27 for the selection task. This highlights the pre-processing procedure as an essential part of the approach. The detailed results of the run without pre-processing can be found in Appendix D1.

However, differentiated considerations are required to adequately interpret the presented results. While the algorithm is able to judge the signals with an accuracy of at least 80% across all evaluation objectives there is a significant difference in F1-scores between automated signal validation and training data selection. The clear maximum of F1-scores for the validation of 2-year




Figure 11. Examples of partially correct classified oil temperature (a), gearbox bearing temperature (b) and pitch motor winding temperature (c). Change Each change in background color indicates a true CPsCP, each dashed lines-detectionline indicates a detected CP.

two-year signals suggests that this is the application the algorithm is suited best for, but not limited to. Analysis has shown that the reduction in performance is predominantly caused by a few challenges common across signals. One of them being fps due to drifts induced by normal wear. A trend-removal step in the pre-processing procedure is suggested to mitigate the effect of regular wear. The challenge of multiple CPs with different significance levels can be tackled by an iterative application of the algorithm to the automatically selected training sub-sequenced. In fact, the The results from the two different time horizons have shown that by dividing the changes of different significant levels into two sub-signals each can be detected successfully. Lastly, the impact of temporary changes on NBM training depends on the significance of the change as well as on the duration of its presence. Short and significant temporary changes can be removed with existing statistical filtering approaches (compare e.g. Bangalore et al. (2017)). A combined application with the presented CPD algorithm is recommended. These measures will help to improve the performance of the algorithm in an application scenario beyond the presented results.

A more conservative approach would be to aim for maximal precision instead of maximal F1-scores. This corresponds to minimal training data loss while still identifying the most significant CPs. As an example, Table 3 shows the algorithm's results for conservative cost-based penalty factors 50 points above the optimal F1-Scores for the Laplace kernel configuration across the different time horizons and objectives. The remaining few fps can be exclusively attributed to normal wear phenomena like shown in Figure 9 a) which can be a useful indicator by itself in a-an NBM setting, as discussed before. This means that without significant loss in training data the algorithm is able to identify and correctly flag the 62%/50% most severe cases among the affected signals. When automatically selecting training data with these conservative penalty-values the 44%/41%

Table 3. Conservative penalty choices and their performance for Laplace kernels and batch-std.

evaluation objective	penalty-values	time horizon	tn -TN	tp -TP	fn -FN	fp -FP	accuracy	precision	recall	f1-score
validation	$\alpha_{cost} = 200$	2 years	412	113	69	4	0.88	0.966	0.621	0.756
validation	$\alpha_{cost} = 130$	1 year	972	107	109	10	0.9	0.914	0.495	0.643
selection	$\alpha_{cost} = 200$	2 years	412	142	162	13	0.76	0.916	0.467	0.619
selection	$\alpha_{cost} = 130$	1 year	972	125	182	17	0.85	0.88	0.407	0.557

most severe CPs are automatically excluded. For illustration, the CPs depicted in Figure 3 as well as the successful detections depicted in Figures 8 and 11 were all correctly identified with the conservative penalty factors. Therefore, the method shows a clear advantage over classical pre-processing procedures.

An alternative and potentially even more effective way to apply the algorithm in the context of NBM is to run it directly on the training error once a model is considered well trained. Conceptually it is clear that CPs in the model input or target induce CPs in the model error. ~~In fact~~ Actually, any CP in the model training error represents a change in conditions the model was not able to adapt to and is therefore worth investigating. The presented pre-processing procedure itself exposes similarities with early approaches of NBM when simple linear models with basic SCADA inputs were used (compare e.g. Schlechtingen and Santos (2011)). This suggests that an application to the training error should be effective and the hyperparameter suggestions from this study applicable. However, these assumptions need to be confirmed with further experiments. A disadvantage of the training ~~error-based~~ error-based approach is that it requires computational expensive model training before validation of the training period. In fact, a combination of both approaches might be the best practice.

6 Summary and Outlook

Literature points out systematic changes in sensor behaviour as one of the most severe challenges when analysing wind turbine SCADA data for early failure detection. This is ~~due to the fact that~~ because most approaches require a clean baseline data set to fit their respective models. This study therefore systematically analysed and, for the first time, quantified the presence of CPs in wind turbine SCADA data. 600 signals from 33 Turbines were analysed for an operational period of more than ~~2~~ two years. During this time ~~one-third~~ one-third of the signals showed one or more significant changes in behaviour induced by sensor and component malfunctions or maintenance actions. This finding highlights the need ~~of~~ for an automated CP detection method. A kernel-based offline CP detection algorithm was introduced which consists of a normalising pre-processing procedure and recommendations on how to choose ~~a number of~~ several crucial hyperparameters. Performance of the algorithm was evaluated across Linear, Gaussian and Laplace kernel configurations, different kernel-bandwidths, and penalty selection schemes. Laplace kernels in combination with newly introduced heuristics for bandwidth and penalty selection performed best and clearly outperformed existing alternative approaches. Signals containing a CP were labelled as such with ~~a~~ an F1-score of up to 0.86 which translates into approximately 50 misclassifications among the 600 analysed signals. Evaluation on a per-CP

basis resulted in a maximum F1-score of 0.73. Despite the reduction in performance, the algorithm was able to automatically exclude the most significant 40% to 60% of all true CPs without significant loss of training data. Therefore, the presented algorithm represents a valuable tool for SCADA data pre-processing and will help ~~data-driven~~ data-driven methods to become more robust despite widely spread data quality issues. ~~Future research has to confirm the presented results for different SCADA~~
5 ~~data sets.~~ Moreover, an extension of signal pre-processing, an iterative application of the algorithm and the combination with existing statistical filtering methods hold the potential for further improve ~~performance.~~ the presented algorithm's performance.
Future research has to confirm the presented results for different SCADA data sets and could aim to extend the method not only beyond the SCADA signals selected in this study but to data from other sensing equipment for condition monitoring in wind turbines as well. In Arlot et al. (2012), for example, kernel CPD has been successfully applied to the segmentation of audio
10 signals which in terms of structure and time resolution are much closer to vibration or acceleration data than the SCADA data analysed in this study. Further development is encouraged by making the code available under the GNU general public license.

Code and data availability. The code of the kernel-based CPD-algorithm publicly available (see Letzger (2020)). The SCADA data-set used during this study is proprietary but several exemplary pre-processed signal samples are published along with the code.

Appendix A: CP summary statistics for 1-year one-year signals




Figure A1. Number of signals per component (left), number of CPs per signal (center), and share of signals with CPs per component (right) for the 1-year one-year time horizon.

Appendix B: List of analysed signals

Table B1. Full list of analysed signals

Component	Signal	Number of signals
Gearbox	Gear Bearing Temperature	48
Gearbox	Gearbox Temperature	18
Gearbox	Gear Oil Temperature	30
Gearbox	Gear Oil Pressure	15
Generator	Generator Bearing Temperature	66
Generator	Generator Winding Temperature	81
Generator	Cooling Temperature	48
Pitch system	Pitch Converter Temperature	90
Pitch system	Pitch Motor Temperature	45
Pitch system	Hydraulic Oil Temperature	15
Electrical system	Transformer Temperatures	15
Electrical system	Box Temperatures	15
Ambient conditions	Ambient temperature	33
Ambient conditions	Tower Temperature	15
Others	Shaft Bearing Temperature	18
Others	Nacelle Temperature	33
Others	Rotor Break Temperature	18
Total		600

Appendix C: Detailed results per component

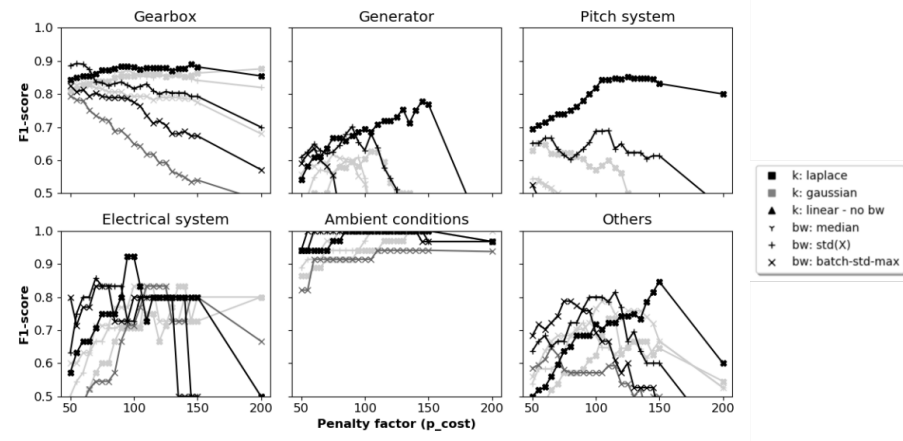


Figure C1. ~~2-year~~ Two-year signal validation: F1-Scores per component for different hyperparameter configurations and penalty values

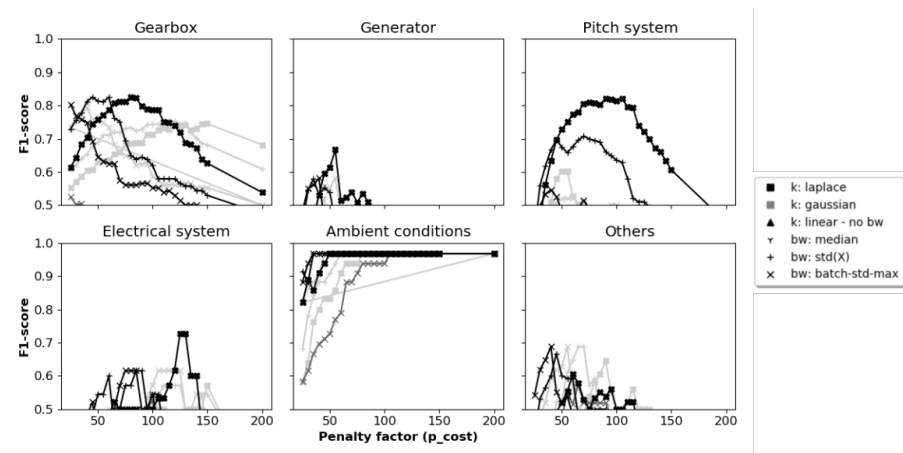


Figure C2. ~~1-year~~ One-year signal validation: F1-Scores per component for different hyperparameter configurations and penalty values.

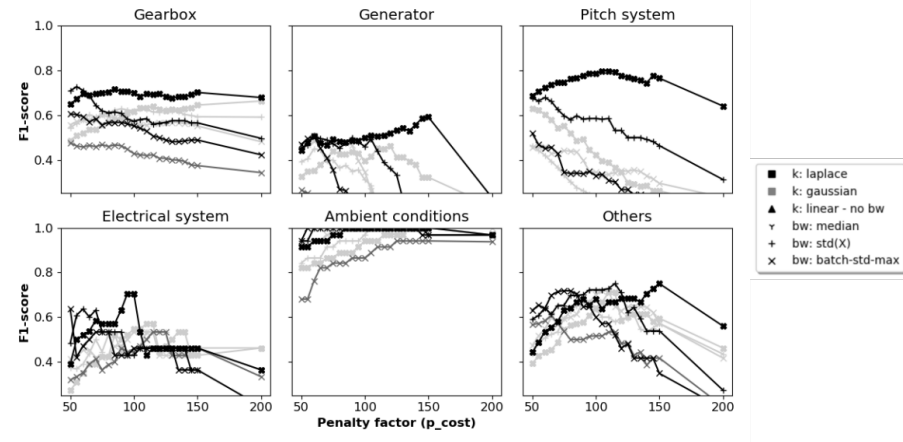


Figure C3. ~~2-year~~ Two-year signal selection: F1-Scores per component for different hyperparameter configurations and penalty values.

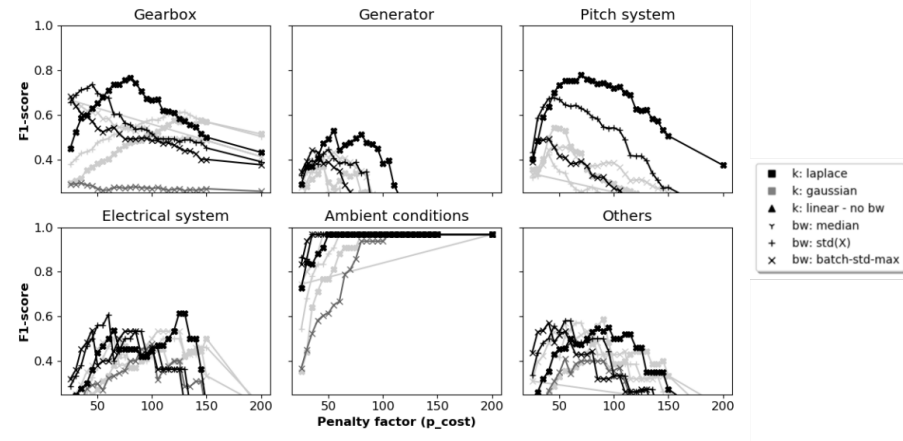


Figure C4. ~~1-year~~ One-year signal selection: F1-Scores per component for different hyperparameter configurations and penalty values.

Table C1. Results validation per component for best configuration (Laplace / std-max / $\alpha_{cost} = 80$) on ±one year signals.

Component	tn -TN	tp -TP	fn -FN	fp -FP	accuracy	precision	recall	f1-score
Gearbox	128	66	8	20	0.874	0.767	0.892	0.825
Generator	349	15	19	7	0.933	0.682	0.441	0.536
Pitch system	220	54	15	11	0.913	0.831	0.783	0.806
Electrical system	45	5	2	8	0.833	0.385	0.714	0.5
Ambient conditions	78	15	1	0	0.989	1.0	0.938	0.986
Others	110	8	8	6	0.894	0.571	0.5	0.533
Total	930	163	53	52	0.91	0.758	0.755	0.756

Table C2. Results selection per component for best configuration (Laplace / std-max / $\alpha_{cost} = 80$) on ±one year signals.

Component	tn -TN	tp -TP	fn -FN	fp -FP	accuracy	precision	recall	f1-score
Gearbox	128	77	18	29	0.81	0.726	0.811	0.766
Generator	349	21	27	13	0.9	0.618	0.438	0.512
Pitch system	220	87	33	24	0.84	0.784	0.725	0.753
Electrical system	45	5	4	8	0.81	0.385	0.556	0.455
Ambient conditions	78	15	1	0	0.99	1.0	0.938	0.968
Others	110	9	10	6	0.88	0.6	0.474	0.529
Total	930	214	93	80	0.87	0.728	0.697	0.712

Appendix D: Results of algorithm without pre-processing

Table D1. Performance of the algorithm without pre-processing on ~~2-year~~two-year signals.

evaluation objective	penalty	time horizon	tn -TN	tp -TP	fn -FN	fp -FP	accuracy	precision	recall	f1-score
validation	$\alpha_{cost} = 7$	2 years	264	155	38	172	0.67	0.47	0.8	0.6
selection	$\alpha_{cost} = 30$	2 years	376	101	212	357	0.46	0.22	0.32	0.26

Competing interests. The author declares no competing interests.

Acknowledgements. The author greatly acknowledges support by the Berlin International Graduate School in Model and Simulation based Research (BIMoS) and the Open Access Publication Fund of TU Berlin. Moreover, the author thanks Greenbyte AB, in particular Pramod Bangalore, for the cooperation and the fruitful discussions.

References

- Aminikhanghahi, S. and Cook, D. J.: A survey of methods for time series change point detection, *Knowledge and information systems*, 51, 339–367, 2017.
- Arlot, S., Celisse, A., and Harchaoui, Z.: Kernel change-point detection, arXiv preprint arXiv:1202.3878, 6, 2012.
- 5 Arlot, S., Celisse, A., and Harchaoui, Z.: A kernel multiple change-point algorithm via model selection, arXiv preprint arXiv:1202.3878, 2016.
- Bach-Andersen, M., Rømer-Odgaard, B., and Winther, O.: Flexible non-linear predictive models for large-scale wind turbine diagnostics, *Wind Energy*, 20, 753–764, 2017.
- Bangalore, P., Letzgs, S., Karlsson, D., and Patriksson, M.: An artificial neural network-based condition monitoring method for wind
10 turbines, with application to the monitoring of the gearbox, *Wind Energy*, 20, 1421–1438, 2017.
- Baudry, J.-P., Maugis, C., and Michel, B.: Slope heuristics: overview and implementation, *Statistics and Computing*, 22, 455–470, 2012.
- Bellman, R.: On a routing problem, *Quarterly of applied mathematics*, 16, 87–90, 1958.
- Birgé, L. and Massart, P.: Minimal penalties for Gaussian model selection, *Probability theory and related fields*, 138, 33–73, 2007.
- Brodsky, E. and Darkhovsky, B. S.: *Nonparametric methods in change point problems*, vol. 243, Springer Science & Business Media, 1993.
- 15 Butler, S., Ringwood, J., and O’Connor, F.: Exploiting SCADA system data for wind turbine performance monitoring, in: 2013 Conference on Control and Fault-Tolerant Systems (SysTol), pp. 389–394, IEEE, 2013.
- Chandola, V., Banerjee, A., and Kumar, V.: Anomaly detection: A survey, *ACM computing surveys (CSUR)*, 41, 15, 2009.
- Coronado, D. and Fischer, K.: *Condition Monitoring of Wind Turbines: State of the Art, User Experience and Recommendations*, 2015.
- Dao, C., Kazemtabrizi, B., and Crabtree, C.: Wind turbine reliability data review and impacts on levelised cost of energy, *Wind Energy*, 22,
20 1848–1871, <https://doi.org/10.1002/we.2404>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2404>, 2019.
- Garreau, D.: *Change-point detection and kernel methods*, Ph.D. thesis, PSL Research University, 2017.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A.: A kernel two-sample test, *Journal of Machine Learning Research*, 13, 723–773, 2012.
- Guédon, Y.: Exploring the latent segmentation space for the assessment of multiple change-point models, *Computational Statistics*, 28,
25 2641–2678, 2013.
- Harchaoui, Z. and Cappé, O.: Retrospective multiple change-point estimation with kernels, in: 2007 IEEE/SP 14th Workshop on Statistical Signal Processing, pp. 768–772, IEEE, 2007.
- IRENA: *Renewable power generation costs in 2018*, 2019.
- Kusiak, A. and Verma, A.: Analyzing bearing faults in wind turbines: A data-mining approach, *Renewable Energy*, 48, 110–116, 2012.
- 30 Lavielle, M.: Using penalized contrasts for the change-point problem, *Signal processing*, 85, 1501–1510, 2005.
- Leahy, K., Gallagher, C., O’Donovan, P., and O’Sullivan, D. T.: Issues with Data Quality for Wind Turbine Condition Monitoring and Reliability Analyses, *Energies*, 12, 201, 2019.
- Lebarbier, É.: *Quelques approches pour la détection de ruptures à horizon fini*, Ph.D. thesis, Paris 11, 2002.
- Letzgs, S.: *Training data requirements for SCADA based condition monitoring using artificial neural networks*, eawe PhD-seminar 2019,
35 2019.
- Letzgs, S.: *slzgs/KernelCPD_WindSCADA: public review WES*, <https://doi.org/10.5281/zenodo.3728023>, 2020.

- Malladi, R., Kalamangalam, G. P., and Aazhang, B.: Online Bayesian change point detection algorithms for segmentation of epileptic activity, in: 2013 Asilomar Conference on Signals, Systems and Computers, pp. 1833–1837, IEEE, 2013.
- Page, E.: A test for a change in a parameter occurring at an unknown point, *Biometrika*, 42, 523–527, 1955.
- Rybach, D., Gollan, C., Schluter, R., and Ney, H.: Audio segmentation for speech recognition using segment features, in: 2009 IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 4197–4200, IEEE, 2009.
- 5 Schlechtingen, M. and Santos, I. F.: Comparative analysis of neural network and regression based condition monitoring approaches for wind turbine fault detection, *Mechanical systems and signal processing*, 25, 1849–1875, 2011.
- Scholkopf, B. and Smola, A.: Support Vector Machines and Kernel Algorithms, *The Handbook of Brain Theory and Neural Networks*, 2002.
- Sun, P., Li, J., Wang, C., and Lei, X.: A generalized model for wind turbine anomaly identification based on SCADA data, *Applied Energy*, 10 168, 550–567, 2016.
- Tautz-Weinert, J.: Improved wind turbine monitoring using operational data, Ph.D. thesis, Loughborough University, 2018.
- Tautz-Weinert, J. and Watson, S. J.: Using SCADA data for wind turbine condition monitoring—a review, *IET Renewable Power Generation*, 11, 382–394, 2016.
- Tautz-Weinert, J. and Watson, S. J.: Challenges in using operational data for reliable wind turbine condition monitoring, 2017.
- 15 Touati, R., Mignotte, M., and Dahmane, M.: Multimodal Change Detection in Remote Sensing Images Using an Unsupervised Pixel Pairwise-Based Markov Random Field Model, *IEEE Transactions on Image Processing*, 29, 757–767, 2019.
- Truong, C., Oudre, L., and Vayatis, N.: Selective review of offline change point detection methods, *Signal Processing*, 167, 2020.
- Zaher, A., McArthur, S., Infield, D., and Patel, Y.: Online wind turbine fault detection through automated SCADA data analysis, *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 12, 574–593, 2009.