Trend Detection in Crime-related Time Series with Change Point Detection Methods

Apostolos Konstantinou, Despoina Chatzakou, Ourania Theodosiadou, Theodora Tsikrika, Stefanos Vrochidis, and Ioannis Kompatsiaris

Information Technologies Institute, Centre for Research and Technology Hellas
{konstantinou,dchatzakou,raniatheo,theodora.tsikrika,stefanos,ikom}@iti.gr

Abstract. Time series analysis can be an asset in the hands of the authorities, as it can enable the understanding and monitoring of trends of criminal activities. In this work, a variety of methods is exploited to detect significant points of change in crime-related time series that may indicate the occurrence of events that require attention. In particular, change point analysis is applied in relevant time series, both offline (retrospective change detection when all data is available) and online (detection of changes as soon as they occur). The focus is on the Crimes in Boston and London Police Records datasets, examining how change point detection can benefit relevant authorities in understanding crime trends to better allocate and manage resources. The experimental results allow us to gain valuable insights, including the observation of seasonal patterns in some cases, with corresponding crimes peaking at specific times, the somewhat different change points identified in some cases by online and offline methods, and the observation that domain knowledge is desired for better method selection and parameters configuration.

Keywords: Trend detection · Change point detection · Crime-related time series.

1 Introduction

With the emergence of online platforms (such as social media, blogs, forums, etc.) and the Internet as a whole, new opportunities for delinquent and criminal activities have arisen, ranging from hacking and financial frauds [8] to even terrorism-related activities, including propaganda spreading, recruitment, and training, as well as hate spreading towards specific social groups [3]. Thus far, significant effort has been placed into developing a wide range of tools to tackle criminal activities from different perspectives, including real time detection of terrorism-related content [1], crime hotspots detection through spatio-temporal analysis [27], and linkage of online identities to criminal investigations [18].

Further to traditional data mining methods, time series have been effectively applied to a wide range of tasks, including the development of methods to detect and predict criminal activities (e.g., [6]). In this context, change point detection methods have been considered and employed for the detection of significant changes in time series, with the accurate and early detection of change points being a pivotal point for drawing valuable insights. Change point detection has many important applications in several areas, including, but not limited to, the financial sector [20], network traffic analysis [25], and climatology [26].

Regarding the fight against crime and terrorism, little effort has been placed thus far towards the detection of critical points of change, probably due to the difficulty in acquiring relevant (ground truth) data. One such example is the application of the Cumulative Sum change point detection method for the identification of statistically significant changes in Houston's daily crime totals during Hurricane Harvey [5]. In the same vein, a framework that builds on top of the E-Divisive change point detection algorithm has been proposed towards detecting statistically significant change points in terrorism-related time series [35]. In both cases, the focus has been on analyzing time series data in an *offline* manner, where change detection is applied retrospectively when all data is available.

Analyzing trends as well as identifying changes as soon as they occur in real-time (i.e., in an *online* manner) could be particularly valuable for the authorities, as it could enable a more effective response and allocation of resources in order to mitigate serious incidents. To this end, this work aims to investigate the effectiveness of both online and offline change point detection methods towards identifying critical changes in crime-related time series; to the best of our knowledge, there is no other work in the literature that performs such analyses on crime-related data in an online setting. Overall, to enable the effective evaluation of the most popular online and offline change point detection methods, first, a wide range of ground-truth datasets from different domains are examined, ultimately leading to the development of a framework that allows for the identification of trends and significant change points in an effective manner. The applicability of the proposed framework in the crime-related domain is demonstrated on two popular relevant datasets, namely the 'Crimes in Boston' and the 'London Police Records' datasets; these datasets do not though have an associated ground truth and thus an insightful qualitative analysis is performed.

2 Related Work

Change point detection methods are typically divided into online and offline [2]. Although offline methods are characterized by higher accuracy, one of their main disadvantages is that they need access to the entire time series, which makes them inapplicable in real time scenarios. Contrary, online methods process data in real time, thus being more suitable for crime detection in real world applications.

Offline Change Point Detection Methods (supervised and unsupervised). Supervised methods include Decision Trees [29], Bayesian Networks [16], Hidden Markov Models [17], and Gaussian Mixture Models [12]. A key drawback of such methods is their need for large amount of annotated data for training, while most real world data is sparsely annotated or not annotated at all. Training can be done on artificial data, but such models usually do not generalize well.

On the contrary, unsupervised methods do not require any kind of annotations and include likelihood ratio methods, probabilistic, kernel-based, and graph-based approaches, as well as clustering methods [2]. The first attempts at unsupervised change point detection have been with the Cumulative Sum Control Chart (CUSUM), which allows for step detection in time series [28]. In the same direction, one of the most commonly used methods is the Binary Segmentation [31], which is characterized by low complexity and operates in a sequential manner. Pruned Exact Linear Time (PELT) [22] is another high-performance offline algorithm that can also be used on multivariate signals. PELT is both computationally efficient and versatile and in many cases outperforms Binary Segmentation making it one of the top change point detection algorithms. Finally, the Prophet forecasting and change point detection tool [34] implements several models and selects the most appropriate for the data at hand.

Online Change Point Detection Methods. Online (real time) methods run concurrently with the activity being monitored (e.g. crime rate), processing data, one point at a time, as it becomes available. Such a point could be the temperature at a location, the price of a stock, crime rate in a area [32], or the effect of an outside factor (e.g., changes on cannabis regulation laws [24]) on crime rates.

Many common algorithms for online change point detection are often variations of their offline counterparts. CUSUM is a typical example with several variations for online detection (e.g. [30, 37]). Another commonly used approach is the Bayesian Online Change Point Detection method [11], which allows for effective detection of long-term changes in online setups. Moreover, Change Finder [33] is an online learning framework based on a probabilistic model that enables the detection of outliers and change points in streaming time series data. Online methods often perform worse compared to their offline counterparts, since they require data from both before and after a data point to effectively determine if it is a change point, with different methods requiring different amounts of such data; their performance is thus a trade-off between the amount of data after the point being taken into account and the time criticality of the task at hand.

3 Methodology

This section briefly overviews the methods considered in this work to ultimately enable effective detection of changes in crime-related time series.

3.1 Offline Change Point Detection (CPD) Methods

Binary Segmentation [31] is characterized by low complexity and uses a recursive approach: first a change point in the complete input signal is detected, then the series is split into two parts around this change point, and the operation is repeated in each part. The process stops when a specified number of change points is detected; in case the number is unknown, a penalty parameter is given.

Pruned Exact Linear Time (PELT) [22] relies on a pruning rule and detects change points by minimizing a cost function over their possible numbers and locations. In particular, it combines optimal partitioning and pruning, and achieves efficient computational cost, while maintaining high accuracy, thanks to the pruning rule that discards many indexes under the assumption that they can never be minima in terms of the minimization performed at each iteration.

Cumulative Sum (CUSUM) [15] requires a set of parameters to be calculated first in order to condition the change detection, namely the mean, the standard deviation, the shift of interest (which is the smallest deviation we wish to detect), the allowance parameter K, and the decision parameter H that determines whether a change has occurred or not. Various values of H have been used; e.g., H was set to 100 in [15], while values between 1 and 40 were explored in [10].

Segment Neighborhood [4] searches the entire segmentation space by first defining a maximum number of change points, denoted as Q. By computing a cost function for all possible segments, then all segmentations with change point between 0 and Q are considered. Due to the exhaustive search performed, an important drawback of this method is the significant computational cost.

Prophet [34] first determines a large number of possible change points at which the rate is allowed to change. It then places a sparse priority on the magnitudes of rate changes (equivalent to L1 regularization) to limit the number of the possible change points to use. By default, 25 potential change points are specified that are uniformly placed in the first 80% of the time series.

3.2 Online Change Point Detection (CPD) Methods

Bayesian Online Change Point Detection (BOCPD) [11] focuses on generating an accurate distribution of the next datum in the series given only the previously observed data. Central point to the algorithm is the time since the last checkpoint, i.e. the run length. The algorithm assumes that the points in the observed time series can be partitioned into non-overlapping segments.

BOCPD with model selection (BOCPDMS) [23] extends BOCPD by introducing multiple models and a method for online model selection; it aggregates over all models and pruns run lengths, keeping the most probable ones per model.

SWAB. Two common approaches to CPD are top-down and bottom-up; top-down approaches (e.g., Binary Segmentation) start with the entire time series and recursively segment it until a halting condition is met, while in the bottom-up, the process starts with the maximum number of segments and merges them using a cost function, until a stopping threshold. SWAB [21] combines a sliding window with bottom-up segmentation, achieving good results in online setups.

This section presented various online and offline CPD methods, with each one approaching the CPD problem in a different way; e.g., Binary segmentation follows a recursive approach, while PELT builds on a pruning rule. While some methods tend to perform better overall (e.g. [36]), variations in their performance can be observed depending on the data and the choice of initial parameters (if any). This is also evident in our experiments (Section 5) indicating that domain knowledge is also needed for choosing the best method in each case.

4 Datasets

To evaluate the different change point detection methods described in Section 3, two types of datasets are considered: (i) a collection of 42 datasets for which the ground truth labels are known and thus a direct comparison among different methods is feasible; and (ii) two crime-related datasets for which there are no ground-truth labels and thus a more qualitative analysis takes place.

4.1 TCPD Benchmark and Dataset Collection

A benchmarking framework consisting of a collection of methods and time series data (also referred to as TCPD) has been proposed that allows testing and comparing CPD algorithms [9]. Overall, it consists of 37 annotated datasets and 5 artificial control datasets, including e.g. the daily closure price of Apple Inc. stock, the price of bitcoin, and the GDP of several countries. All these datasets have been annotated by one or more field experts in time series analysis.

4.2 Crimes in Boston Dataset

The 'Crimes in Boston' dataset [7] contains crimes reported to the Boston Police Department from June 15, 2015 to September 29, 2019. For each reported crime, additional information is available, e.g. its incident number, the date and time of the crime, and the location. In total, there are 576 different crime types. To allow for a more coarse-grained analysis, we manually grouped the different types in eight general categories, e.g. crimes related to 'Theft', 'Robbery', and 'Burglary' are grouped into the *Theft* category. The crime categories along with their frequency are: Person-related: 125, 747 (18.79%), Assault: 124, 750 (18.64%), Theft: 119, 788 (17.90%), Fraud: 108, 983 (16.29%), Traffic: 58, 844 (8.79%), Narcotics: 23, 928 (3.58%), Misc 22, 450 (3.35%), and Other: 84, 702 (12.66%).

Figure 1 depicts the time series for each category, with each point corresponding to the number of crimes committed in each of the 52 months in the dataset. Through visual inspection one can identify underlying patterns that may be of interest and assess potential critical points of change. For instance, we observe that for 'Theft' a periodicity appears that could be useful for police authorities to make a better allocation of resources with the aim of dealing with this type of crime as best as possible. Moreover, for 'Assault', there appear to be four points (at around the 10, 15, 35 and 50 points on the x-axis) of change that may require further study to draw useful conclusions. On the other hand, for the 'Person' and 'Fraud' crimes there do not seem to be any obvious points of interest, which suggests that there is a relatively stable pattern regarding these crimes.

To determine if there is any connection and co-occurrence between crimes, we also examined the correlation matrix of each crime type against every other, using the Pearson's correlation coefficient [14]. Table 1 indicates that there is for instance high correlation between 'Traffic' and 'Fraud' (0.856), and 'Fraud' and 'Assault' (0.690), but also between 'Person-related' crimes and 'Fraud' (0.822). The correlations that emerge from such an analysis can be useful, as they can

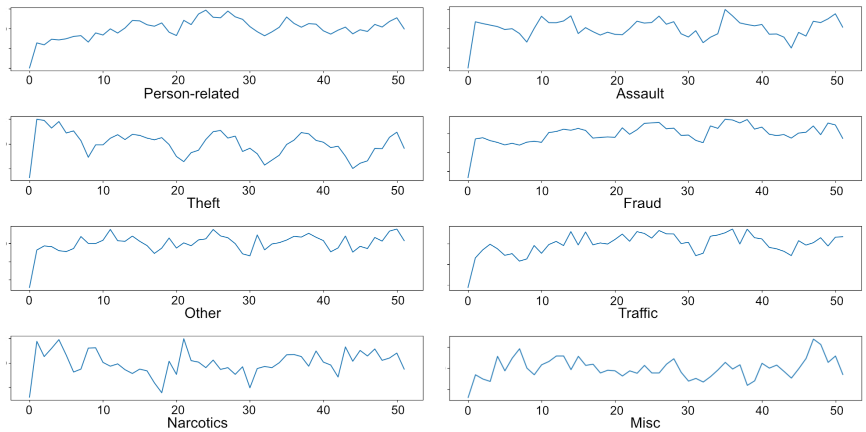

Fig. 1: Boston crime series: count of crimes per category in each month.

Table 1: Crimes in Boston dataset: correlation matrix.

	Assault	**Fraud**	**Misc**	**Narcotics**	**Other**	**Person-related**	**Theft**	**Traffic**
Assault	1	0.690	0.454	0.335	0.729	0.584	0.688	0.645
Fraud	0.690	1	0.322	0.279	0.754	0.822	0.317	0.856
Misc	0.454	0.322	1	0.180	0.401	0.310	0.278	0.195
Narcotics	0.335	0.279	0.180	1	0.383	0.038	0.166	0.135
Other	0.729	0.754	0.401	0.383	1	0.649	0.400	0.603
Person-related	0.584	0.822	0.310	0.038	0.649	1	0.202	0.820
Theft	0.688	0.317	0.278	0.166	0.400	0.202	1	0.264
Traffic	0.645	0.856	0.195	0.135	0.603	0.820	0.264	1

be important piece of information on how and whether it makes sense to deal with not just one crime at a time, but a set of crimes in a more effective way.

4.3 London Police Records Dataset

The 'London Police Records' [13] dataset consists of a list of crimes committed in the area of London from June 2014 to May 2017. Overall, it consists of 14 crime types: 1. Vehicle crime: $262,309$ (8.90%), 2. Violence and sexual offences: $596,107$ (20.23%), 3. Antisocial Behavior: $708,264$ (24.04%), 4. Bicycle theft: $54,649$ (1.85%), 5. Other theft: $333,817$ (11.33%), 6. Theft from the person: $109,168$ (3.71%), 7. Other crime: $29,208$ (0.99%), 8. Drugs: $106,836$ (3.63%), 9. Burglary: $213,125$ (7.23%), 10. Public order: $130,653$ (4.43%), 11. Shoplifting: $135,780$ (4.61%), 12. Criminal damage and arson: $184,772$ (6.27%), 13. Robbery: $68,920$ (2.34%), and 14. Possession of weapons: $12,871$ (0.44%).

From Figure 2 we observe that 'Burglary' and 'Bicycle theft' crimes seem to have quite distinguishable change points, while periodicity is also observed in both cases; e.g. for 'Burglary', a seasonality is observed, with crime rates

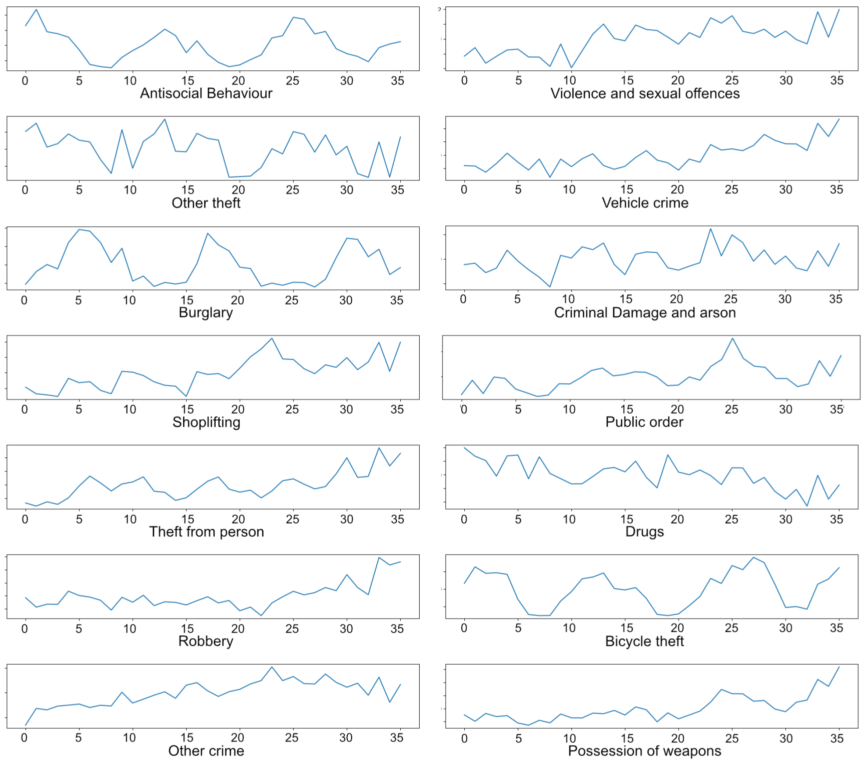

Fig. 2: London crime series: count of crimes per category in each month.

peaking each year during the months of November (5, 17, 29), December (6, 18, 30), and January (7, 19, 31). 'Bicycle theft' is also observed in the summer period during June (0, 12, 24), July (1, 13, 25), and August (2, 14, 26), while 'Antisocial behavior' seems to peak every July (1, 13, 25). Change points that may be of interest and could receive more attention can also be seen in the crimes of 'Violence and sexual offences', 'Criminal damage and arson', and 'Shoplifting'. For the rest of the crimes, no particularly obvious change points seem to appear.

Finally, similarly to before, we also estimated the correlation matrix as presented in Table 2. Overall, a high correlation is observed between Antisocial behavior and Bicycle theft (0.889), Sexual offences and Vehicle crimes (0.703), Violence and Public order (0.825), as well as Violence and Shoplifting $(0, 705)$.

5 Experimental Results

This section first presents the results obtained on the TCPD dataset collection. Then, the best performing (offline and online) methods are employed to conduct a qualitative analysis on the two crime-related datasets (i.e. 'Crimes in Boston' and 'London Police Records'), for which no ground truth labels are available.

Table 2: London Police Records dataset: correlation matrix.

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
1	1	0.70	0.15	0.37	0.09	0.73	0.58	-0.46	0.04	0.59	0.71	0.49	0.84	0.84
2	0.70	1	0.23	0.31	0.25	0.46	0.79	-0.22	-0.17	0.82	0.71	0.68	0.46	0.71
3	0.15	0.23	1	0.89	0.64	-0.29	0.01	0.29	-0.57	0.53	-0.07	0.50	0.08	0.29
4	0.37	0.31	0.89	1	0.55	-0.13	0.17	0.11	-0.60	0.62	0.06	0.54	0.24	0.45
5	0.09	0.25	0.64	0.55	1	-0.09	-0.06	0.34	-0.05	0.37	-0.13	0.59	0.10	0.01
6	0.73	0.46	-0.29	-0.13	-0.09	1	0.35	-0.61	0.36	0.32	0.61	0.25	0.83	0.59
7	0.58	0.79	0.01	0.17	-0.06	0.35	1	-0.33	-0.20	0.68	0.74	0.51	0.22	0.54
8	-0.46	-0.22	0.29	0.11	0.34	-0.61	-0.33	1	-0.05	-0.18	-0.42	-0.04	-0.40	-0.35
9	0.04	-0.17	-0.57	-0.60	-0.05	0.36	-0.20	-0.05	1	-0.43	-0.06	-0.24	0.23	-0.25
10	0.59	0.82	0.53	0.62	0.37	0.32	0.68	-0.18	-0.43	1	0.54	0.75	0.37	0.70
11	0.71	0.71	-0.07	0.06	-0.13	0.61	0.74	-0.42	-0.06	0.54	1	0.56	0.47	0.66
12	0.49	0.68	0.50	0.54	0.59	0.25	0.51	-0.04	-0.24	0.75	0.56	1	0.28	0.43
13	0.84	0.46	0.08	0.24	0.10	0.83	0.22	-0.40	0.23	0.37	0.47	0.28	1	0.74
14	0.84	0.71	0.29	0.45	0.01	0.59	0.54	-0.35	-0.25	0.70	0.66	0.43	0.74	1

5.1 Experimental Results on the TCPD Dataset Collection

In the TCPD framework (see Section 4.1) various CPD methods have been implemented [9]. To gain a good understanding of the methods most commonly considered in the literature (Section 3), we tested and compared them against each other (using the settings defined in the aforementioned framework) based on two metrics: (i) the cover metric, which is based on the Jaccard Index (also known as Intersection over union); and (ii) the F_1 metric that treats change points detection as a classification problem [9].

Table 3 presents the corresponding results, through averaging the scores obtained across the 42 datasets presented in Section 4. Overall, TCPD consists of both univariate and multivariate datasets[1], but as in our case of interest (i.e. crime rate monitoring per observation period to gain insights) the focus is on univariate data analysis, we proceed only with the univariate ones. The results indicate that the offline CPD methods perform better compared to the online ones, which is expected as offline methods have all the data available to perform the relevant analysis. In particular, Binary Segmentation achieves 0.672 cover and 0.698 F1 scores, while PELT achieves 0.652 and 0.674, respectively.

For online methods, the best performance is obtained with BOCPD, a quite popular approach, followed by BOCPDMS, which trains multiple models choosing the appropriate model each time. Although one would expect BOCPDMS to perform better compared to BOCPD, it was observed that in some datasets (e.g. with small time series length) the performance was particularly poor, consequently leading to an overall reduced performance. Moreover, the BOCPDMS is a variant of the Bayesian method, targeting mainly multivariate data; when

[1] Multivariate data analysis involves more than two dependent variables to result in an outcome, compared to univariate where only one variable at a time is considered.

Table 3: Experimental results on TCPD dataset collection.

Method	Cover score	F1 score
Offline CPD methods		
Binary Segmentation	**0.672**	**0.698**
CUSUM	0.526	0.572
Segment Neighborhoods	0.642	0.635
PELT	0.652	0.674
Prophet	0.522	0.47
Online CPD methods		
Bayesian online change point detection (BOCPD)	**0.594**	**0.662**
BOCPD with model selection (BOCPDMS)	0.590	0.495
SWAB	0.543	0.487

evaluating BOCPDMS on TCPD's multivariate datasets, better performance was observed compared to BOCPD (0.496 vs. 0.455 in the cover score).

In the following sections, a qualitative analysis on the crime-related datasets is conducted. Due to space limits for each case (online and offline), we proceed with one method per case. In particular, focusing on the offline methods, although Binary Segmentation achieved better performance, we proceed with PELT (the second best performing method), as according to the literature PELT leads in most cases to a more accurate detection of change points [36]. In relation to the online methods, BOCPD is used for the remaining analyses.

5.2 Experimental Results on the Crime-related Datasets

Crimes in Boston dataset. Focusing illustratively (due to space limits) on the 'Assault' crime, Figure 3 depicts the corresponding time series in addition to the identified change points (depicted by dashed lines) as detected with PELT (offline) and BOCPD (online). Overall, we sampled the time series with various sample rates (i.e. grouping of crime rates by hour, day, week, and month), and here we indicatively present the results at a monthly rate. Based on the illustrated results, different change points are identified with the offline vs. online methods. This could be attributed to the fact that online methods have limited access to posterior information and consequently sometimes lead to a different interpretation of the trend evident in the time series, which also affects the change points detection process. Moreover, although PELT is known for its efficacy, it relies heavily on the choice of the initial parameters, meaning that for instance a small *penalty* (hyperparameter) would make the method too sensitive and thus leading to too many points being predicted. Overall, the depicted results indicate that there is no method that fits all, and therefore targeted configuration (based on domain knowledge) is required to best describe the data at hand.

London Police Records. Similarly to above, here we indicatively focus on the 'Antisocial behavior' (Figure 4a), and 'Burglary' (Figure 4b) crimes. As the figures show, there are in some cases seasonal patterns (e.g. crimes related to

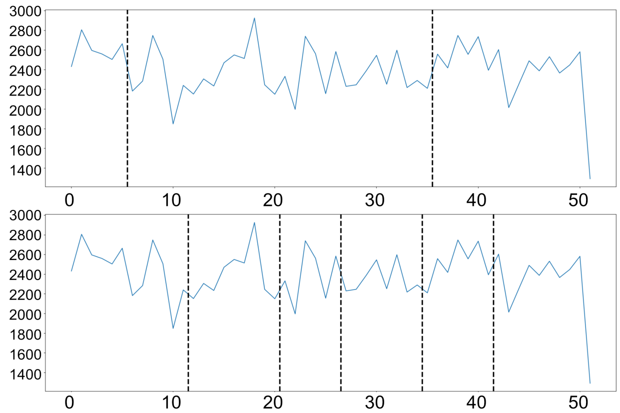

Fig. 3: Crimes in Boston (assault crimes): PELT (top) and BOCD (bottom).

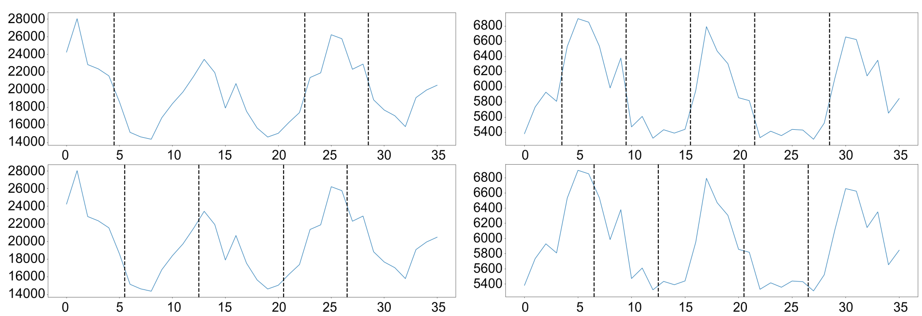

(a) Antisocial Behavior crime (b) Burglary crime

Fig. 4: London Police Records: PELT (top) and BOCD (bottom).

antisocial behavior tend to be more intense in the winter months) indicating that the respective crimes peak at certain times during the year.

In contrast to the Crimes in Boston dataset, the PELT and BOCPD methods here make very similar predictions. This could be attributed to the difference in the length of the datasets (Boston covers 52 months, while London 36 months). Based on how PELT works, keeping the *penalty* hyperparameter constant, the shorter the length, the more the number of change points detected. For a more effective detection of change points, the value of *penalty* should be set taking also into account the size of the time series; there are works focusing on the appropriate selection of this value (e.g. [19]). Finally, in the case of the 'Antisocial behavior' and 'Burglary' crimes (London) stronger fluctuations are observed in the time series compared to the 'Assault' crimes (Boston), e.g. 'Antisocial behavior' has 3493.87 standard deviation, while 'Assault' has 268.34. As BOCPD is based on detecting changes in variance, this could explain why it performs better on the London dataset and predicts more similar change points to PELT.

6 Conclusions and Future Work

This work examined offline and online CPD methods to enable effective understanding and detection of trends and change points in time series. The focus was on crime-related time series, having first performed a fairly extensive analysis of data coming from other domains, but characterized by ground truth labels. The analysis conducted indicates that CPD methods can be a valuable tool for police authorities as they will be able to better understand the trend on topics of interest so they can then proceed with better management of resources. Through time series analysis, patterns can be identified (such as seasonality), while at the same time the detected change points can be pivotal points for decision-making. In the future, we intend to conduct a similar study to additional crime-related datasets, while also deep neural network-based approaches will be examined to allow for an even better detection of changes in crime-related time series data.

Acknowledgements This project has received funding from the European Union's H2020 research and innovation programme as part of the STARLIGHT (GA No 101021797) project.

References

1. Abrar, M.F., Arefin, M.S., Hossain, M.S.: A framework for analyzing real-time tweets to detect terrorist activities. In: ECCE. pp. 1–6. IEEE (2019)
2. Aminikhanghahi, S., Cook, D.J.: A survey of methods for time series change point detection. Knowledge and information systems **51**(2), 339–367 (2017)
3. Asongu, S.A., Orim, S.M.I., Nting, R.T.: Terrorism and social media: Global evidence. J. Glob. Inf. Technol. **22**(3), 208–228 (2019)
4. Auger, I.E., Lawrence, C.E.: Algorithms for the optimal identification of segment neighborhoods. Bulletin of mathematical biology **51**(1), 39–54 (1989)
5. Augusto, D.: Change-point analysis of houston crime during hurricane harvey. Journal of applied security research **16**(1), 1–18 (2021)
6. Borowik, G., Wawrzyniak, Z.M., Cichosz, P.: Time series analysis for crime forecasting. In: ICSEng. pp. 1–10. IEEE (2018)
7. Boston, A.: Crimes in boston. https://kaggle.com/datasets/AnalyzeBoston/crimes-in-boston (2023)
8. Burden, K., Palmer, C.: Internet crime: Cyber crime—a new breed of criminal? Computer Law & Security Review **19**(3), 222–227 (2003)
9. Van den Burg, G.J., Williams, C.K.: An evaluation of change point detection algorithms. arXiv preprint arXiv:2003.06222 (2020)
10. Burkatovskaya, Y., Kabanova, T., Khaustov, P.: Choice of the parameters of the cusum algorithms for parameter estimation in the markov modulated poisson process. In: ITSMSSM. pp. 456–462. Atlantis Press (2016)
11. Chowdhury, M.F.R., Selouani, S.A., O'Shaughnessy, D.: Bayesian on-line spectral change point detection: a soft computing approach for on-line asr. International Journal of Speech Technology **15**, 5–23 (2012)
12. Cleland, I., Han, M., Nugent, C., Lee, H., McClean, S., Zhang, S., Lee, S.: Evaluation of prompted annotation of activity data recorded from a smart phone. Sensors **14**(9), 15861–15879 (2014)

13. Dane, S.: London police records. https://kaggle.com/datasets/sohier/london-police-records (2023)
14. Derrick, T.R., Bates, B.T., Dufek, J.S.: Evaluation of time-series data sets using the pearson product-moment correlation coefficient. MSSE **26**(7), 919–928 (1994)
15. Granjon, P.: The cusum algorithm-a small review (2013)
16. Grzegorczyk, M., Husmeier, D.: Non-homogeneous dynamic bayesian networks for continuous data. Machine Learning **83**, 355–419 (2011)
17. Han, M., Vinh, L.T., Lee, Y.K., Lee, S.: Comprehensive context recognizer based on multimodal sensors in a smartphone. Sensors **12**(9), 12588–12605 (2012)
18. Han, X., Wang, L., Cui, C., Ma, J., Zhang, S.: Linking multiple online identities in criminal investigations: A spectral co-clustering framework. IEEE Transactions on information forensics and security **12**(9), 2242–2255 (2017)
19. Haynes, K., Eckley, I.A., Fearnhead, P.: Computationally efficient changepoint detection for a range of penalties. J Comput Graph Stat **26**(1), 134–143 (2017)
20. Jeon, S.Y., Ryou, H.S., Kim, Y., Oh, K.J.: Using change-point detection to identify structural changes in stock market: application to russell 2000. Quantitative Bio-Science **39**(1), 61–69 (2020)
21. Keogh, E., Chu, S., Hart, D., Pazzani, M.: An online algorithm for segmenting time series. In: International conference on data mining. pp. 289–296. IEEE (2001)
22. Killick, R., Fearnhead, P., Eckley, I.A.: Optimal detection of changepoints with a linear computational cost. JASA **107**(500), 1590–1598 (2012)
23. Knoblauch, J., Damoulas, T.: Spatio-temporal bayesian on-line changepoint detection with model selection. In: International Conference on Machine Learning. pp. 2718–2727. PMLR (2018)
24. Lu, R., Willits, D., Stohr, M.K., Makin, D., Snyder, J., Lovrich, N., Meize, M., Stanton, D., Wu, G., Hemmens, C.: The cannabis effect on crime: Time-series analysis of crime in colorado and washington state. Justice Quarterly **38**(4), 565–595 (2021)
25. Lung-Yut-Fong, A., Lévy-Leduc, C., Cappé, O.: Distributed detection/localization of change-points in high-dimensional network traffic data. Statistics and Computing **22**(2), 485–496 (2012)
26. Maidstone, R.: Efficient analysis of complex changepoint problems. Lancaster University (United Kingdom) (2016)
27. Malleson, N., Andresen, M.A.: Spatio-temporal crime hotspots and the ambient population. Crime science **4**, 1–8 (2015)
28. Page, E.S.: Continuous inspection schemes. Biometrika **41**(1/2), 100–115 (1954)
29. Reddy, S., Mun, M., Burke, J., Estrin, D., Hansen, M., Srivastava, M.: Using mobile phones to determine transportation modes. TOSN **6**(2), 1–27 (2010)
30. Romano, G., Eckley, I., Fearnhead, P., Rigaill, G.: Fast online changepoint detection via functional pruning cusum statistics. arXiv preprint:2110.08205 (2021)
31. Scott, A.J., Knott, M.: A cluster analysis method for grouping means in the analysis of variance. Biometrics pp. 507–512 (1974)
32. Shamsuddin, N.H.M., Ali, N.A., Alwee, R.: An overview on crime prediction methods. In: ICT-ISPC. pp. 1–5. IEEE (2017)
33. Takeuchi, J.i., Yamanishi, K.: A unifying framework for detecting outliers and change points from time series. TKDE **18**(4), 482–492 (2006)
34. Taylor, S.J., Letham, B.: Forecasting at scale. The American Statistician **72**(1), 37–45 (2018)
35. Theodosiadou, O., Pantelidou, K., Bastas, N., Chatzakou, D., Tsikrika, T., Vrochidis, S., Kompatsiaris, I.: Change point detection in terrorism-related online content using deep learning derived indicators. Information **12**(7), 274 (2021)

36. Wambui, G.D., Waititu, G.A., Wanjoya, A.: The power of the pruned exact linear time (pelt) test in multiple changepoint detection. AJTAS **4**(6), 581 (2015)
37. Wei, S., Xie, Y.: Online kernel cusum for change-point detection. arXiv preprint arXiv:2211.15070 (2022)