Poster: Extracting Insights from Censorship Measurement Data Using Statistical Techniques

Anna Crowder
University of Florida
annacrowder@ufl.edu

Kevin Butler
University of Florida
butler@ufl.edu

***Abstract*—As censorship methods implemented by governments worldwide have grown in sophistication, the open-source platforms used to measure censorship have increased in scale and number. These platforms are difficult to explore because of their complexity and require a strong understanding of the data collection process and the proper use of established data analysis methods to accurately extract information from them. We combine two statistical methods, control charts and the Mann-Kendall trend detection test, and implement them on data from three open-source censorship detection platforms: Censored Planet, GFWatch, and the Open Observatory of Network Interference (OONI). With these two methods, we analyze data from Ukraine, Russia, and China from March 2020 through December 2022. Our analysis correctly identifies when censorship activity occurs, which we correlate from ground truth reporting of these events. Our approach additionally uncovers changes in censorship data that previously went unreported without requiring the collection of in-country networking data from providers. We also discuss some of the challenges that come with investigating open-source censorship measurement data. Our methodology represents a starting point for deeper automated analysis of censorship data and a tool for monitoring censorship measurement data as it is published.**

I. Introduction

Controlling access to the Internet through censorship is a powerful tool for governments and authorities that has become increasingly common. While governments continue to increase the sophistication of their censorship apparatus, researchers are also improving censorship detection and measurement techniques. In this paper, we examine the landscape of publicly available censorship measurement data sets. We focus on data collected by Censored Planet [7], OONI [2], and GFWatch [3] from March 2020 to December 2022. We use data from Ukraine, Russia, and China, three countries that have experienced serious fluctuation in censorship over the past two years. Because censorship measurement data is complex and challenging, it is often included as just a supporting tool in censorship research. Our work is the first step towards using solely open-source censorship measurement data to understand censorship events.

II. Techniques for Time Series Analysis

After each data set is processed and each measurement is labeled as an anomaly (potentially censored) or not an anomaly, we perform preliminary exploratory data analysis and determine the best statistical methods for censorship measurement data. We use a non-overlapping averages control chart for OONI and GFWatch [6] and an Individuals control chart [6] for Censored Planet's Satellite data set, each combined with one nonparametric technique, the Mann-Kendall trend detection test [5]. Control charts are an established statistical tool for monitoring an ongoing process. They can indicate shifts in average performance or changes in variability. Data over a relatively stable period of time are used as a baseline to establish a process average and upper and lower control limits based on distributional assumptions. Data outside the control limits are no longer behaving in an expected way according to the established baseline, and thus are a signal. The Mann-Kendall trend detection test is used to indicate an overall upward or downward trend in the data. It can confirm that an ongoing shift occurred at the time the control chart signaled and further eliminate false positives. The statistical significance of the trend is determined based on a null hypothesis that there is no trend with α-level = 0.05, and a signal occurs when the trend is significant enough to reject the null hypothesis. The two separate methods in combination with each other improve the ability of the system to detect signals of potential changes to censorship levels. The decision process for choosing when the control chart and Mann-Kendall signals should be investigated further is shown in Table I.

III. Time Series Evaluation

A summary of the control chart and Mann-Kendall signals for OONI, Satellite, and GFWatch for Ukraine, Russia, and China are shown in Figure 1. Our focus is on points in time where the signals from the two methods overlap.

Figure 1a shows that the OONI data from Ukraine stayed fairly consistent until July 2021 when it went above the upper control limit (UCL) and stayed above. Around the same time in August 2021, the Mann-Kendall data trended upward for eleven consecutive windows. In December 2021 the Satellite data has windows of increasing Mann-Kendall trend and points that go above the control chart UCL. Both the OONI and Satellite Ukraine data have statistically significant changes in the Fall of 2021 when the control chart and Mann-Kendall signal at the same time. While little is known about the state of censorship in Ukraine before the conflict, our method indicates a disruption occurred well before Russia invaded in February 2022.

Tool	Signal	Decision
Control Chart	1 point outside the control limits	Monitor
	2 out 3 points outside the control limits	Investigate further because metric has shifted
Mann-Kendall Trend Detection Test	1 point with increasing or decreasing trend	Monitor
	2 out of 3 points with increasing or decreasing trend	Investigate in detail because metric is trending
Both	At least one point outside the control limits that also has an increasing or decreasing trend	Primary focus of further investigation

TABLE I: Different scenarios where each test indicates a censorship signal, and the corresponding action to take.

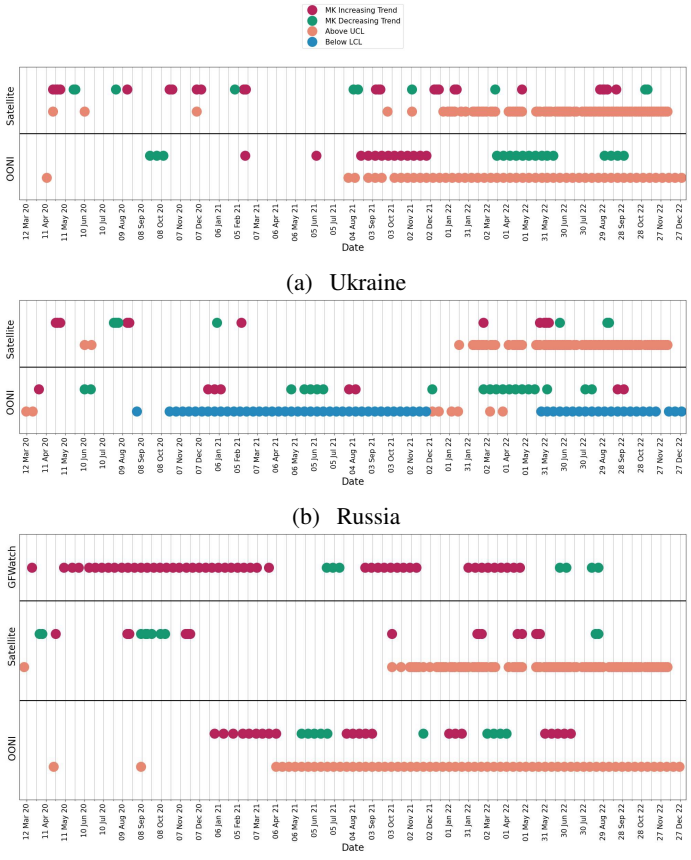

(a) Ukraine

(b) Russia

(c) China

Fig. 1: Summary of control chart and Mann-Kendall trend detection signals.

In Russia both methods indicate dramatic shifts in the OONI data starting in December 2021 and returning in May 2022. In contrast, the Satellite data does not start shifting significantly until February 2022 when both methods signal an increase. The details of these signals can be seen in Figure 1b. The differences between the OONI data and the Satellite data highlight the importance of using multiple open-source measurement data sets when monitoring for potential censorship events.

When we perform the Mann-Kendall trend test on the GFWatch data, it has an increasing trend from the very beginning which continues throughout almost the entire two years of data. At the start of January 2021, the OONI data starts increasing according to the Mann-Kendall trend, and around April 2021 the data goes above the UCL. The data stays above the UCL but regularly alternates between trending upward and downward. In contrast, Satellite does not signal until October 2021, when the data goes above the UCL and experiences an increasing Mann-Kendall trend followed by small windows of upward trends in February and April 2022. The differences between OONI and Satellite point to the strengths and weaknesses of each data set. The OONI data responds quickly to small shifts in anomalies caused by China's ever-changing censorship policies which often react to current political events [1]. GFWatch and Satellite data provide a consistent view of the well-documented [4], ever-increasing level of censorship in China.

IV. Conclusion

Our work demonstrates the use of statistical methods to pull key information from open-source censorship measurement data. We use control charts and the Mann-Kendall trend detection test on data from Censored Planet, OONI, and GFWatch in Ukraine, Russia, and China over two years. By combining these two methods we pinpoint signals in censorship measurement data. We see an increase in censorship in Ukraine and Russia a few months before Russia invaded Ukraine. We see that censorship in China has been steadily increasing over the past two years with small fluctuations in 2021 and 2022. The signals identified provide a starting point for deeper investigation into changes to censorship data.

References

[1] Kerry Allen. How China's Covid protests are being silenced. *BBC News*, 2022. URL: https://www.bbc.com/news/world-asia-china-63788477.

[2] Arturo Filastò and Jacob Appelbaum. OONI: Open observatory of network interference. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2012.

[3] Nguyen Phong Hoang, Arian Akhavan Niaki, Jakub Dalek, Jeffrey Knockel, Pellaeon Lin, Bill Marczak, Masashi Crete-Nishihata, Phillipa Gill, and Michalis Polychronakis. How great is the Great Firewall? Measuring China's DNS censorship. In *Proceedings of the USENIX Security Symposium*, 2021.

[4] Freedom House. China: Freedom on the Net 2022 Country Report. URL: https://freedomhouse.org/country/china/freedom-net/2022.

[5] Pacific Northwest National Laboratory. Mann-kendall test for monotonic trend. URL: https://vsp.pnnl.gov/help/vsample/design_trend_mann_kendall.htm.

[6] Douglas C. Montgomery. *Introduction to Statistical Quality Control*, pages 261, 268. Wiley, 2013.

[7] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. Censored Planet: An Internet-wide, longitudinal censorship observatory. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2020.